



UNIVERSITY OF CAMBRIDGE

**Inventing Intelligence:  
On the History of Complex Information Processing  
and Artificial Intelligence in the  
United States in the Mid-Twentieth Century**

This thesis is submitted for the degree of  
Doctor of Philosophy

In

History and Philosophy of Science

By

Jonathan Nigel Ross Penn ('Jonnie Penn')  
Pembroke College, Cambridge

Examiners:

Professor Simon Schaffer, University of Cambridge  
Professor Jon Agar, University College London

Supervisor: Dr. Richard Staley, University of Cambridge  
Advisor: Dr. Helen Anne Curry, University of Cambridge

Word Count: 78,033  
Date: 14 December 2020

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my thesis has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. It does not exceed the prescribed word limit for the Degree Committee of the Department of History and Philosophy of Science at the University of Cambridge.

## Abstract

### *Inventing Intelligence: On the History of Complex Information Processing and Artificial Intelligence in the United States in the Mid-Twentieth Century*

In the mid-1950s, researchers in the United States melded formal theories of problem solving and intelligence with another powerful new tool for control: the electronic digital computer. Several branches of western mathematical science emerged from this nexus, including computer science (1960s–), data science (1990s–) and artificial intelligence (AI). This thesis offers an account of the origins and politics of AI in the mid-twentieth century United States, which focuses on its imbrications in systems of societal control. In an effort to denaturalize the power relations upon which the field came into being, I situate AI's canonical origin story in relation to the structural and intellectual priorities of the U.S. military and American industry during the Cold War, circa 1952 to 1961.

This thesis offers a detailed and comparative account of the early careers, research interests, and key outputs of four researchers often credited with laying the foundations for AI and machine learning—Herbert A. Simon, Frank Rosenblatt, John McCarthy and Marvin Minsky. It chronicles the distinct ways in which each sought to formalise and simulate human mental behaviour using digital electronic computers. Rather than assess their contributions as discontinuous with what came before, as in mythologies of AI's genesis, I establish continuities with, and borrowings from, management science and operations research (Simon), Hayekian economics and instrumentalist statistics (Rosenblatt), automatic coding techniques and pedagogy (McCarthy), and cybernetics (Minsky), along with the broadscale mobilization of Cold War-era civilian-led military science generally.

I assess how Minsky's 1961 paper 'Steps Toward Artificial Intelligence' simultaneously consolidated and obscured these entanglements as it set in motion an initial research agenda for AI in the following two decades. I argue that mind-computer metaphors, and research in complex information processing generally, played an important role in normalizing the small- and large-scale structuring of social behaviour using mathematics in the United States from the second half of the twentieth century onward.

By: Jonathan Penn

# Table of Contents

<b>Title</b>	<b>1</b>
<b>Abstract</b>	<b>3</b>
<b>Table of Contents</b>	<b>4</b>
<b>Acknowledgements</b>	<b>6</b>
<b>Archive Collections</b>	<b>9</b>
<b>List of Figures</b>	<b>10</b>
<b><u>Chapter One – Introduction</u></b>	<b>11</b>
Prehistories: Psychology, Cybernetics and Computing	18
Cybernetics: The Field that Was Not a Field	22
Computer Programming: Materiality, Mind and Machine	26
Historiography from Within: Popular Accounts of AI History by Practitioners and their Kin	30
Scholarly Histories: Classical and Alternative Histories of Computing	36
Summary	42
<b><u>Chapter Two – On the Logic Behind the Logic Theory Machine</u></b>	<b>44</b>
1930-40s: Simon’s Scientific Account of Administration	46
1952-55: Simon’s Account of Administration is Translated into Code at RAND	51
1955-56: Building the Logic Theory Machine	61
1956-58: Framing the Logic Theory Machine	69
Conclusion	80
<b><u>Chapter Three – On Frank Rosenblatt’s Perceptron Theory, 1945-61</u></b>	<b>82</b>
1945-56: Rosenblatt and Civilian-Led Military Science in Post-War America	83
1955-58: From the Statistics of Personality Research to the Statistics of Brain Modelling	94
1958: Rosenblatt’s Statistical Approach to Brain Modelling	101
1958-71: On the ‘Second Direction’ of Brain Modelling and its Military Suitors	108
Conclusion	114
<b><u>Chapter Four – Applied Epistemology: On John McCarthy’s Account of Programmed Language as Intelligence, 1952-59</u></b>	<b>116</b>
1944-55: The Road to Automata	118
1955-56: From Automata to Automatic Programming and AI	125
1956: The Dartmouth Summer Research Project for McCarthy	135

1956-59: Programming as Applied Epistemology	146
Conclusion	157
<b><u>Chapter Five – Contextualising Marvin Minsky’s Agenda for Artificial Intelligence, 1950-61</u></b>	<b>159</b>
1944-54: Neural Networks as a ‘Plausible Analogy’ for Mental Behaviour	161
1954-56: A Change of Mind – Minsky’s Turn to Heuristic Programming	169
1956-61: Minsky Outlines a Research Agenda for AI	180
1961-1970s: AI and Military Infrastructure	191
Conclusion	197
<b><u>Chapter Six – Conclusion</u></b>	<b>197</b>
<u>An Emerging Discipline</u>	200
<u>Fitting AI into Histories of American Social Science Research in the Mid-Twentieth Century</u>	213
<b><u>Bibliography</u></b>	<b>221</b>

## Acknowledgements

‘There is no power out of the church that could sustain slavery an hour, if it were not sustained in it.’

- Albert Barnes<sup>1</sup>

I have reflected at various stages of this project on why it is not customary in the tradition of the history and philosophy of science that I have been trained in to explore, in a formal capacity as in an Acknowledgements section, the distinct context and positionality that informs one’s historical research. While these factors are often (but not always) considered in the body of a work, or pointed to in footnotes, at conferences, in presentations or over drinks, there is no conventional expectation to include, say, a Reflexivity Statement to consider one’s own contextual orientation upfront. Learning how to relate to the cultural, political and social forces that shape both me and my research—as well as the priorities of Western academic research generally—has been an important and ongoing part of this professionalizing experience so I will include a view brief remarks here alongside a set of well-deserved Acknowledgements.

One immediate consideration to name is that the University of Cambridge is currently closed due to the Covid-19 pandemic. To my knowledge, the last closure of this type was due to the Great Plague in the mid seventeenth century. The full impact of the virus remains to be felt. That two and a half billion people worldwide are currently in lockdown begs reflection on the extent to which academia challenges or is complicit in government systems that consistently fail to protect—and often to recognise—the vulnerable, even as they introduce such widescale compensatory measures. I deliver this PhD from a government-imposed quarantine in Western Canada, where I am with family to provide caring duties. For posterity, I will note that I am a dual Canadian-British, Caucasian, able-bodied, cis-male born as a settler in a settler colonial context. This positionality informs both the strengths and weaknesses of my current research project in ways that I continue to unpack.

During the course of this dissertation, from 2016-20, I have called a number of locations home. The first and most important is the Department of the History and Philosophy of Science at the University of Cambridge. I would like to thank the Rausing, Williamson and Lipton trust for their financial support, without which these words would not have been written. A sincere thank you also goes to Richard Staley, my supervisor. Richard is a dedicated scholar and generous friend. I will cherish our long, rolling conversations together in his office on Free School Lane—where a hand drawn Einstein portrait hangs on the wall. I hope that I can live up to the quality of excellence he has helped me to appreciate and to enjoy pursuing. I am equally indebted to Helen Anne Curry, my advisor. Helen’s skilful critiques have renewed my confidence in the authority of the written word. I feel I must learn them all again now. Her sharp eye and patient ear have enriched this project considerably. Any remaining faults in what follows are irredeemably my own.

Additional thanks are due to friends and colleagues at the University of Cambridge. I am grateful to Simon Schaffer, Anna Alexandrova, Joe Martin, Marina Velickovic, David Runciman, Andrew Buskell and Josh Nall for thoughtful guidance on various topics. Tamara Hug and her talented team made my archival work possible; I am grateful for these relationships too. A special thank you to Stephen Cave, Sarah Dillon, Susan K. Gowans and the

---

<sup>1</sup> As cited in: Douglass, ‘The Meaning of July Fourth for the Negro’.

Leverhulme Centre for the Future of Intelligence, where I have had the opportunity to foster an international community of scholars oriented toward historical understandings of AI. I also want to acknowledge the University's undergraduate community for insisting—often via civil disruption—that academic and government leaders take action to address mounting democratic and environmental crises. The University disentangles its debts to slavery and colonization at a snail's pace. Unsustainable investments remain intact. I am grateful to the hearts and minds who illuminate the paths to justice in these areas—I hope that academia will honour their unappreciated bravery (and its costs) as the freighter changes course.

These critiques and acknowledgements carry equally to Harvard University and the Massachusetts Institute of Technology, two institutions that I have also called home during various placements and fellowships over the past few years. I am moved by my colleagues at the Berkman Klein Center for Internet & Society, whose commitment to engage with political realities has revealed to me shapes of what academia *could* be. I want to thank Jonathan Zittrain and Rebecca Tabasky specifically for creating and sustaining a community in which such intellectual pursuits can proceed. I reflect often on conversations had with Mary Gray, Nathan Freitas, Momin M. Malik, Elizabeth Renieris, Elif Sert, Sabelo Mhlambi, Luke Stark, Boaz Sender, Ben Green, Salome Viljoen, Kathy Pham, Kendra Albert, Elettra Bietti, Primavera De Filippi, Jeannette Estruth, Dragana Kaurin, Jasmine McNealy, Amy X. Zhang, Jenny Korn, James Mickens, Sarah Newman, Keith Porcaro, Dean Jansen, Beau Sievers, Josh Simons, Sean McDonald and numerous others. I am grateful to Jennifer Light at MIT for welcoming me into HASTS as a visiting student and for coaching me on research strategies during our fleeting time together. I thank David Kaiser for showing me kindness at various turns along the way. At the MIT Media Lab, where I served as a BKMLA Fellow, I am grateful to Natalie Saltiel, Dhaval Adjodah and to Sandy Pentland for far reaching discussions about data governance, data cooperatives, Davos, transdisciplinarity, technology cults and artificial intelligence.

An unexpected joy over the course of this project has been the opportunity to collaborate and learn from many whom I look up to. I thank Stephanie Dick, Matthew Jones, Mustafa Ali and Aaron Plasek in this respect especially. I am also grateful to Margaret Boden, Ronald Kline, Joy Rankin, George Nagy, Sarah T. Hamid, William Lazonick, Daniel Volmar, Katya Babinsteva, Bo An, Theodora Dryer, Yaqub Chaudhary, George Dyson, Meredith Whittaker, Andrew Pickering, Rudi Seising, Colin Garvey, Rodrigo Ochigame, Michael and Caroline Running Wolf and others whose research and insights guide my way. I thank my friends for their jest and unreasonable patience along the way. Michael and Isabel Thornton, Jack Clark, Gesa Sophia Borgeest and Marcel Elias, Ed Saperia and the community at Newspeak House, Rowland Bagnall and Becky Chesshyre, Alex Fisher, Oliver Marler, Josh McKenty and the graduate student body at the Department of the History and Philosophy of Science, along with others I've already named and those I have unintentionally overlooked. I am also thankful to my family and friends in Canada, who are still not entirely sure what I've been up to for the past decade.

As a matter of disclosure, I have contributed to various external projects during the course of my research. I thank John Linsey et. al. for insights into European policy while I served as a Google Technology Policy Fellow in partnership with the European Youth Forum. At UNI Global Union, my deepest thanks go to Christina Colclough, a true lion, whom I am ever grateful to have found as a friend. I thank Chartwell Speakers Bureau and Maria Axente for connecting me to contemporary debates in policy and business related to AI for Social Good and the United Nation's 2030 Sustainable Development Goals. On this topic, I also thank Sean McGregor, Margaux Luck, Yoshua Bengio and the rest of the team(s) with whom I have

had the pleasure to collaborate with on the AI for Social Good Workshop Series at NeurIPS, ICML and ICLR. It has been a pleasure to connect with so many AI and machine learning researchers from around the globe on questions about critical practices.

One final and important home to mention was that of Marvin Minsky, whose family welcomed me into their lives to examine the extensive personal archives that Marvin left behind after his passing in 2016. I value this connection, which blossomed beyond our immediate shared interest in Marvin's research. I have tried not to let this familiarity alter the subjects I have chosen to write about, but I disclose the tie all the same. Thank you to Gloria Rudisch for teaching me the grace of youth, and for having me at your jazz nights. It has been a pleasure to reconstruct timelines with Margaret Minsky and Cynthia Solomon, and to break bread with extended family, including Charlotte Minsky. I regret that I could not integrate the archives of the AI Laboratory at MIT into my current research while in Boston. I bring new archival material to light on each of the four figures I study most closely, but the project would have been richer with this institutional perspective included as well. The same can be said of the broader racial, gender and geopolitical dimensions underlying my account, as well as questions of ability—within the scope of this dissertation, these areas could be stronger.

A chance experience from my initial archival research stays with me. In 2017, I fell into conversation with an archivist at the Library of Congress. It emerged that many graduate students in the history of science credited their decision to study certain scientists and/or subjects directly to the size of the grant apportioned to them for archival research by their academic funders. The archivist explained that the Library of Congress and its holdings were popular because D.C. was accessible on that budget. Together we wondered how this allotment had shaped understandings of science, biasing researchers toward the lionization of certain scientists and subjects at the expense of others.

It is experiences like this one that lead me to consider the contingencies shaping historical craftsmanship and, in turn, the need to consider a project's path dependencies and limitations. I develop these ideas more thoroughly in this dissertation as they pertain to my broader inquiry. I study knowledge production and bureaucracy in elite universities in mid-century America and cannot help but see Western academia as a form of locomotion. It has been interesting to witness the University of Cambridge persist as a set of logics amidst a deadly pandemic, just as MIT and Harvard did following the Jeffrey Epstein controversy. Collective etiquette and its harms are a force to behold. If this pandemic is a glimpse at the climate crises ahead, then I pray that scholarly norms and accompanying organizational logics will evolve. At times I fear that my culture will learn itself to death.

From the M-Room under a pink moon.

Jonnie Penn  
2020



## Archive Collections

Carnegie Mellon University Archives, Allen Newell Collection. Pittsburgh, PA.

Carnegie Mellon University Archives, Herbert A. Simon Collection. Pittsburgh, PA.

Cornell University Library, Division of Rare and Manuscript Collections, Dale R Corson Papers. Ithaca, NY. With permission from the President's Office.

Cornell University Library, Division of Rare and Manuscript Collections, Frank Rosenblatt Papers. Ithaca, NY.

IBM Corporate Archives, IBM Corporation. Poughkeepsie, NY

Library of Congress, Manuscript Division, Papers of John von Neumann. Washington, DC.

Library of Congress, Manuscript Division, Papers of Claude Shannon. Washington, DC.

Library of Congress, Manuscript Division, Papers of Maurice Rosenblatt. Washington, DC.

Library of Congress, Manuscript Division, Papers of Nathaniel Rochester. Washington, DC.

Massachusetts Institute of Technology, Institute Archives and Special Collections. Cambridge, MA.

National Archives and Research Administration, Archives II. College Park, MD.

University of Minnesota, Charles Babbage Institute, Archives and Special Collections. Twin Cities, MN.

Rockefeller Archive Center, Rockefeller Foundation Collection, Sleepy Hollow, NY. (Courtesy of Ronald Kline.)

Stanford University Libraries, Department of Special Collections and University Archives, John McCarthy Papers. Stanford, CA.

University of Cambridge, King's College Archive Centre, Papers of Alan Mathison Turing. Cambridge, UK.

University of Cambridge, St John's College Library, Papers of Maurice Wilkes. Cambridge, UK.

Personal Archives of George Dyson.

Personal Archives of Bruce W. Knight.

Personal Archives of George Nagy. Troy, NA.

Personal Archives of Gloria Rudisch. Brookline, MA.

Personal Archives of Grace Solomonoff.

## List of Figures

### Chapter Three

Figure 1.0 – General organization of a perceptron.

Figure 2.0 – A more advanced perceptron schematic from 1957.

### Chapter Four

Figure 3.0 – An IBM 704 Electronic Data-Processing Machine. Courtesy of International Business Machines Corporation, © 1954 International Business Machines Corporation.

Figure 4.0 – ‘The Artificial Intelligencers.’ Attendees at Dartmouth. Courtesy of Gloria Rudisch.

Figure 5.0 – A cartoon depicts the function of automatic programming. From the Historical and Interpretive Collections of The Franklin Institute, Philadelphia, PA.

### Chapter Five

Figure 6.0 – Overlap by category between bibliographical items in ‘Steps Toward Artificial Intelligence’ and ‘A Selected Descriptor-Indexed Bibliography to the Literature on Artificial Intelligence.’ Eighty eight of ninety-nine items in ‘Steps’ included. Seven did not repeat.

Figure 7.0 – Items in ‘A Selected Descriptor-Indexed Bibliography to the Literature on Artificial Intelligence’ labelled as relevant to the highest number of sub-categories. Ex: Article 323 was labelled as relevant to ‘Human learning behavior,’ ‘Conditional probability,’ and ten other sub-categories.

Figure 8.0 – 14 October 1962 – Artificial intelligence appears under ‘Programmers/Analysts.’

Figure 9.0 – 4 November 1962 – Artificial intelligence appears under ‘Programmers/Analysts.’

## Chapter One – Introduction

'Everything that humiliates labor also humiliates the intelligence, and vice versa'  
- Albert Camus<sup>2</sup>

In Western culture, idealised conceptions of mental behaviour such as 'intelligence' have long been used to naturalize techniques and characteristics that have also served as methods of political domination and control.<sup>3</sup> In antiquity, Aristotle advocated for the subjugation of women, slaves, animals and plant life on the assumption that the ability to reason was naturally endowed only to men of good birth. Alison Adam shows how analogous assumptions carried through Cartesian ideals of the rational man into the foundations of logicism, which inspired the epistemological concerns of mid-century artificial intelligence (AI) research and the machines it inspired.<sup>4</sup> Activists, journalists and academics deconstruct how these patterns repeat through the contemporary deployment of AI-powered decision making systems.<sup>5</sup> Questions of power, control and inequity linked to conceptions of mental behaviour recur across historical periods, even as the epistemologies and technologies through which they manifest have changed.

A thorough assessment of AI's political imbrications requires first an appreciation of the field's origins in elite research and manufacturing centers in the United States in the 1950s and 1960s. In this dissertation, I argue that AI has been a science of industry since its beginnings, in contrast to historical treatments that depict a purely academic venture concerned with the demonstration of intelligence by machines. I bring to light how administrative and economic theory on decision systems like corporations and decentralized markets informed key researchers' conceptions of how to model the human mind. These entanglements lionized conformity by treating rule following, be it for profits or efficiency, as tantamount to thought itself. This conflation also obscured the complexity of neural behaviour, disingenuously equating closed systems with open ones. Despite criticisms over a

---

<sup>2</sup> Camus and O'Brien, *Resistance, Rebellion, and Death*, 96.

<sup>3</sup> Cave, 'Intelligence: A History'; Cave, 'The Problem with Intelligence'.

<sup>4</sup> Adam, *Artificial Knowing*.

<sup>5</sup> As example, see: O'Neil, *Weapons of Math Destruction*; Eubanks, *Automating Inequality*; Broussard, *Artificial Unintelligence*; Keyes, 'The Misgendering Machines'; Buolamwini and Gebu, 'Gender Shades'; Benjamin, *Race After Technology*; Couldry and Mejias, *The Costs of Connection*.

lack of intellectual rigour, AI rapidly transitioned from a provisional to professional science between 1955-63. This was in part, I argue, due to the shifting epistemic status of automatic programming techniques, which reduced the labour required to develop operational brain models.

The phrase ‘artificial intelligence’ was coined by John McCarthy, an American mathematician, in 1955. It has travelled with a noticeably amorphous definition since. I introduce the term ‘brain model’ research to situate early AI efforts in relation to a set of conceptually adjacent modes of analysis that practitioners and commentators retrospectively annexed into ‘artificial intelligence’ after the late 1950s. These include complex information processing, heuristic programming and machine learning. I venture that AI should perhaps be understood, from a historical perspective, as a branch of political science parsed through the toolset of computer science. Historians must probe and challenge the extent to which certain social orders have tended to be reified under the guise of ‘AI’; ossifying, as it does, particular views of the nature of the mind.

AI’s viability as a discipline was not assured in the 1950s and early 1960s, the period I focus on herein. Luminaries in American mathematics like Claude Shannon and John von Neumann bristled at a new generation’s hubristic use of mind-as-computer metaphors, such as a suggestion by McCarthy in 1954 that mathematical automata could be ‘intelligent.’ Less than a decade later, however, the Artificial Intelligence Laboratory at the Massachusetts Institute of Technology operated with an annual budget purported to be in the millions—a transformative step into professionalization that has yet to be contextualized.

Canonical origin stories of AI overlook structural contingencies when they take this ascent as a given. The reorientation I propose requires a more precise articulation of the strategic choices made by those responsible for advancing brain model and AI theory as well as the strategic choices of the institutions that enabled them. The next four chapters are devoted to this interplay. I provide a detailed account of the early career research foci, institutional interactions and published works of Herbert A. Simon, Frank Rosenblatt, John McCarthy and Marvin Minsky—four men credited with having made foundational contributions to AI, machine learning or both.<sup>6</sup> Using previously unexamined archival

---

<sup>6</sup> This group’s explicit research programmes offer a reasonably comprehensive understanding of the central techniques and claims of AI in this period, diminishing the value of extending

material, I connect their assorted access to special hardware, fellowships, military research grants and other contingencies in the 1950s to significant U.S. military investments (and, subsequently, industrial investments) in search and storage techniques related to Cold War geopolitical anxieties about quantifiable risks and technological supremacy. The search for a theory of intelligence was expensive, but alignment with these institutions, and their motivating logics, helped to settle the bill.

At the heart of my story is the observation that each of these four figures sought a distinct abstract *language* with which to expressively communicate a computer's operations to a human programmer and vice versa. In their eyes, language was the key to synthesizing what amounted to a totalizing epistemology; a knowledge system limited only by the sophistication of its manufacturing. 'Once one system of epistemology is programmed and works no other will be taken seriously unless it also leads to intelligent programs,' McCarthy wrote in his diaries, 'The artificial intelligence problem will settle the main problems of epistemology in a scientific way.'<sup>7</sup> This positivist vision of scientific progress was a common lodestar for these men. Even as they disagreed over the correct measure of fidelity to biological phenomena in their modelling or the most appropriate rhetoric for their results, they positioned brain modelling as ripe for its Aristotle, Newton or Descartes; a singular genius who would synthesize a grand theory of cognition that would reduce epistemology to code, a possibility they all presumed existed.

In retracing this story, I call attention to one critical respect in which these men were mistaken—the abstract languages they sought were realized by *many* hands, not just a few. The rapid proliferation of automatic programming techniques between 1952-57 reduced the expertise and labour needed to program a digital computer. These automated refinements availed computing to a wider number of professionals, circumstances that conformed the mental and manual labour of new parties to the affordances of digital techniques. This, in turn, lent an aura of possibility to the notion that a machine could think. McCarthy himself, who coined 'AI,' saw the two domains, automatic programming and AI, as broadly

---

my detailed analysis to additional colleagues in computing, cybernetics and other cognate fields. Although my focus on influential individuals runs the risk of reinforcing 'great men' accounts of history, my reframing points to the importance of heterogenous programming communities, including industrial workers and corporations like IBM, in understandings of AI.

<sup>7</sup> McCarthy, 'Methodology of Work on the Artificial Intelligence Problem', 81.

synonymous. However, automatic programming was largely an *industrial* project removed from the rarefied study of neural behaviour. Manufacturers like the International Business Machines Corporation (IBM) invested heavily to standardize techniques and devices in order to enrol new clients, not to find the next Newton. The men I examine were active agents in this enrolment, proselytizing to other academics about the inevitable need to revisit one's expertise through the lens of computer programming. These activities and structural contingencies have not figured in AI's history until now.

The vocabulary Simon, Rosenblatt, McCarthy and Minsky chose to describe new techniques in major newspapers and scholarly journals informed Americans' still plastic understandings of what was possible, and indeed desirable, in the emerging information age. In the early 1950s, professional etiquette compelled Minsky to use scare quotes around his early mathematical notions of machine 'learning.' By the late 1950s, these quotes had disappeared. This scholarly transition, from metaphorical to literal descriptions of human capacities in machines, was not frictionless. During the mid to late 1950s, these men turned to clannishness, self-aggrandizement, speculative rhetoric, fluid definitions of key terms and poor citation practices to shore up legitimacy for their controversial new techniques—actions that drew attention toward questions of *how* to accomplish such aims and away from *whether* they were well founded.

Internalist accounts of brain model theory have tended to forgive these practices as customary. Recurring emphasis on narratives of discontinuity, such as on AI funding 'winters' or the fabled Dartmouth workshop where AI gained its name, have obscured broader methodological and structural continuities. One recurring yet distracting tendency in existing histories is a focus on the swing of fashions between two apparent poles: symbolic reasoning (e.g. the simulation of problem solving) and neural network research (e.g. the simulation of learning). I eschew that dichotomy herein to revisit what these pursuits ultimately had in common, such as the foundational beliefs that complex mental processes were susceptible to mathematical formalization, that digital electronic computers were the appropriate tool for that job, and that the product of this enterprise would be an abstract 'language.' As I note, imbrications with military and industrial partners were another commonality, as was the belief that these imbrications had no meaningful influence on the direction of new theory.

Artificial intelligence was certainly more than indiscrete posturing and the suggestive rhetoric with which key figures advanced their work. Still, the field's influence on culturally

specific notions of intelligence and its rapid rise in Western science compel us to take careful account of both the promise offered and threat hidden in what were at times disingenuous practices. This observation indicates the fundamentally political grounds of their work, engaging as it did the interrelated yet unequal scientific authority of social science and biological fact, as well as the entangled capacity to re-order social systems *using* AI.

I begin this introductory chapter with brief synopses of subsequent chapters followed by a survey of relevant historical precedents for brain model research in psychology, cybernetics and computer programming to give a picture of the intellectual milieu within which these men worked. I juxtapose popular understandings of AI's genesis with the historiographical frameworks that historians of computing have used to structure understandings of both the field and of computing more generally. I join with commentators who position the history of information and statistical technologies in close proximity to the histories of public administration. My comparative biographies call attention to particular choices that we should take account of throughout the development of AI, here examined until the early 1960s.

In Chapter Two, I isolate the specific competencies used by Herbert A. Simon, along with his collaborators Allen Newell and J. Clifford Shaw, to justify comparisons between the abilities of the Logic Theory Machine, the fabled 'first prototype of AI,' and the problem solving abilities of a human being. I illustrate how Simon's training as a political scientist at the University of Chicago in the 1930-40s shaped his positivist view of the social sciences and, in turn, the design specifications of the Logic Theory Machine. I demonstrate how the trio transmuted conceptual tools from Simon's *Administrative Behavior* and from Bertrand Russell and Alfred North Whitehead's *Principia Mathematica* into digital code between 1955-56.

The Logic Theory Machine was a model of a certain type of behaviour, namely human problem solving, as seen through the prism of early twentieth century symbolic logic and post-war American administrative logic. These logics were celebrated at the RAND Corporation, where Simon worked, amidst the rise of Cold War era civilian science, which provided the basis for his employment there. I show that this trio saw their work as a contribution to operations research, not artificial intelligence. Revisiting AI's roots in what they called 'complex information processing' and 'heuristic programming' clarifies the field's conceptual debts to systems designed to understand and order *social* dynamics—in this case, American administrative logics—rather than neural dynamics.

In Chapter Three, I revisit the early career trajectory of Frank Rosenblatt, a pioneer of neural network research about whom little serious historical scholarship has been undertaken. I draw on unexamined archival evidence to situate Rosenblatt's 1956 PhD on the explanatory limits of statistics and digital computation in personality science in relation to post-war optimism about the capacity of those two domains to bring about transformative scientific and social progress in the United States. The U.S military fed from and into this optimism to coordinate, fund and retain domestic civilian scientists during peacetime through new organizations like the Office of Naval Research. Military and scientific hopes entangled in Rosenblatt's post-PhD research on statistical theories of brain models, which began with his involvement in perception research for the U.S. Air Force at the Cornell Aeronautical Lab, where he spent his career.

I argue that while Rosenblatt did not support the view that the brain's complexity could be meaningfully synthesized using deterministic logics, as Simon, Minsky and McCarthy did, he fervently believed that *psychology* could be reduced to statistics in a far more comprehensive fashion than had yet been attempted—and that electronic computers would validate proposed techniques, as had been the premise of his PhD. Rosenblatt grounded this faith in Friedrich Hayek's writings on decentralized market behaviour. He used Hayek's economic theory as a basis from which to diverge from his own professed commitments to biological fidelity in mathematical brain modelling, arguing that brains, like economies, were so complex that it was unreasonable to think they could be centrally managed. Rosenblatt defined his work in contrast to a 'loyal opposition' in AI, a framing that subsequent historical accounts have followed. In contrast, I group Rosenblatt *within* this community, not outside it. I show that he indulged his own abstract contrivances in relation to cognition, and advanced his ideas, as did his 'opposition,' through military patronage and digital computing.

In Chapter Four, I consider the early career of John McCarthy, who is credited with coining the term 'artificial intelligence.' I survey his early formal training and publication record in academic mathematics to situate his definition of intelligence alongside the development of logic, partial differential equations and automata studies in the 1940-50s. McCarthy believed, by 1955, that the key to understanding the dynamics of intelligence lay in the study of language. I provide a detailed chronology of his attempts to qualify this assumption as co-editor of *Automata Studies* (1956), co-convenor of the Dartmouth Summer



Research Project on Artificial Intelligence (1956), and co-founder of both the MIT AI Group (1958) and the MIT AI Laboratory (1959).

By the late 1950s, McCarthy claimed that AI would reduce epistemology to a branch of applied mathematics. I situate his views in relation to the rapid development of automatic coding techniques between 1952-57. Albeit colloquially, McCarthy and others used the terms ‘AI’ and ‘automatic coding’ interchangeably in this period. The latter automated computational instruction sets to reduce the labour and expertise needed to program a computer. This set of contextual concerns has not yet been accounted for historically. By shifting from histories of scientific personality and research fashions to histories of language propagation, commercial interests and pedagogy, these structural contingences come one step closer into view.

In Chapter Five, using new archival evidence, I chart the development of two influential 1961 papers by Marvin Minsky, each republished for a wide audience in the 1963 volume *Computers and Thought*, a bestselling textbook in early AI. These papers helped to consolidate aforementioned ideas from Simon, McCarthy and others into a research agenda for AI in the 1960s and 70s. I begin with Minsky’s early career developing SNARC, for Stochastic Neural-Analogue Reinforcement Computer, the first mechanical neural network, as a doctoral candidate in mathematics at Princeton University in 1950. The mathematician read heavily in cybernetic theory and mathematical biophysics and was initially tentative about neural metaphors. This began to change by 1954 as he pursued ‘universal’ decision procedures. Buoyed by results in heuristic programming from Simon and others, he spent five years developing ‘Steps Toward Artificial Intelligence,’ which laid out a research agenda for the field based in mathematical understandings of efficiency.

Minsky’s intention to model optimal search and storage techniques aligned him with leaders in industry, government and the military, each with their own purposes for analogous functionality. He equated AI directly to these projects, as if the mind were an information-service like a research library. I argue that his two foundational contributions—a research agenda and a bibliography—can be understood through this archive-as-intelligence lens. Seeing AI’s roots in information services rather than neurophysiology alone helps us to recognise how esoteric academic ventures fit into larger geopolitical projects like the bureaucratization of academia and the military. With that groundwork laid, I discuss how the U.S. Air Force had come to see AI by the 1960s: as a profound new tool for bureaucratic

control. I demonstrate this in reference to the first AI-related job advertisements, posted in *The New York Times* in 1962. I argue that the expression ‘Good Old Fashioned Artificial Intelligence,’ which is used to characterise this early AI research, misrepresents and obscures this military legacy as apolitical or even twee.

In Chapter Six, I summarize key findings. I use my bibliographic vignettes to surface continuities, both theoretical and structural, between Simon’s development of complex information processing and heuristic programming, McCarthy and Minsky’s development of artificial intelligence, Rosenblatt’s development of machine learning, and postwar American social science more broadly. I argue that the procedures these men crafted to simulate aspects of cognition in the 1950s should be understood as an eclectic yet singular intellectual project to reduce epistemology to code. Each man conceived of social and neural phenomena through a positivist lens, making it a smaller step for them to mirror insights from *existing* social science in their efforts to reduce the complexities of mental life down to a scale that they could manipulate. By this line of argument, I challenge the explanatory depth of discontinuity narratives that structure historical understandings of brain model development through the themes of rivalry and rupture. I then locate my account within two core genres of existing Cold War historiography.

To understand the historical and historiographical basis for the significance of this argument it will be helpful now to consider in some detail the primary intellectual frameworks within which brain-model research and AI were developed, and have been understood to date.

### Prehistories: Psychology, Cybernetics and Computing

The list of intellectual traditions that prefigured mid-century brain modelling in the United States is long.<sup>8</sup> Many of these tributaries cut out from a common source, namely the threat of authoritarian rule before, during and after the Second World War, which inspired a

---

<sup>8</sup> It includes but is not limited to: operations research, systems research, actuarial science, logical positivism, public administration, personality research, biology, automata theory, proto-computer science, proto-computer programming, computer engineering, economics, political science, first-order cybernetics, mathematics, philosophy of science and other areas of Cold War-era social and cyborg and software sciences; each with its own relevant prehistories.

kaleidoscope of social science and psychological research on how to identify, cultivate, mitigate and/or optimize notions of virtuousness in American life, be it through pedagogy, surveillance, profiling, testing and/or technology.<sup>9</sup> The lines of influence I explore stem from the particular research and collaboration networks of the four figures I study. For this reason, my focus is on the United States, with occasional reference made to developments in the United Kingdom and Europe. I begin with a brief expository glance at developments in three domains recognized in existing historiography as significant to the era I examine: modern psychology, cybernetics and digital computer programming. I do not mean to concede that these fields were of more importance to brain modelling than, say, public administration—this dissertation as a whole will show why that is not the case. Yet, given existing historical understandings of mid-century brain model research, I must revisit the relevant theoretical precedents in play in these areas during the lead up to AI's initial formalization, 1955-63.

At the turn of the twentieth century, the foundations of modern psychology took shape around scholarly debate in Europe over the boundaries between psychology and philosophy. Wilhelm Wundt established the field's first institutional experimental laboratory in Germany in 1879. Here, he trialled introspective techniques alongside measures of reaction times and sensory processes. Max Wertheimer, an Austro-Hungarian-born psychologist, founded Gestalt psychology in 1912 in opposition to Wundt's atomistic methods. Gestalt psychology favoured holistic explanations of mental dynamics. Sigmund Freud, also active in this period, reasoned that the unconscious aspects of mental life were in fact inaccessible to introspection.

JoAnne Brown's history of intelligence testing shows how the idealization of mental behaviours by turn-of-the-century psychologists intersected with large-scale social planning efforts in America.<sup>10</sup> Brown charts how this group crafted a definition of intelligence that increased the stature of their profession by leveraging the authority of quantification in other domains.<sup>11</sup> 'Control of the physical world is secondary to control of ourselves and our fellow man,' proclaimed James McKeen Cattell, a student of Wundt's, at the 1904 International

---

<sup>9</sup> Phillips, *The New Math*; Cohen-Cole, *The Open Mind*; Solovey and Cravens, *Cold War Social Science*, 2016.

<sup>10</sup> On the role that institutions have played in shaping understandings of intelligence in Anglo-American culture see: Carson, 'The Culture of Intelligence'.

<sup>11</sup> For a history of the persuasiveness of quantification across disciplines see: Porter, *Trust in Numbers*.

Congress of Arts and Sciences.<sup>12</sup> Between 1908-18, principles of mental measurement developed in nineteenth century medicine, anthropometry and phrenology were formalized into practical mental ability tests. It took only ten years after that for these tests to be fully implemented in public schools across America. After World War I, as Sarah Igo chronicles, the spread and normalization of mass surveys on political and consumer attitudes transformed notions of 'the public' around statistical measures of imagined 'majority' and 'minority' clusters, as well as cultural understandings of the 'we' who inhabited each.<sup>13</sup> By 1942, the anthropologist Margaret Mead famously asked whether democracy and social science were compatible.<sup>14</sup>

As the twentieth century progressed, the study of behaviour supplanted the study of inner mental life as the primary methodology of psychological inquiry. Jacques Loeb and his student John E. Watson positioned the mind as impenetrable to scrutiny other than by inspection of its inputs and outputs. Loeb demonstrated that exposure to light triggered muscle feedback in a moth, a point he used to argue that *all* organisms should be understood as 'chemical machines.'<sup>15</sup> In 1913, Watson distilled behaviourism down to a manifesto; 'Psychology as the behaviorist views it is a purely objective experimental branch of natural science,' he argued, 'Its theoretical goal is the prediction and control of behavior.'<sup>16</sup>

In the 1920-30s, the mind as machine metaphor grew as a fringe idea in the behaviourist school of psychology. In 1926, Clark Hull, an American behaviourist, posited a 'robot approach' to psychological inquiry. In *Science*, he claimed to have verified aspects of Ivan Pavlov's famed conditioned reflex through a synthetic test.<sup>17</sup> Hull positioned materialistic and mechanistic explanations as key to understanding the mind as an anticipatory 'psychic machine,' a term he acknowledged was still paradoxical in the eyes of his peers.<sup>18</sup> Over the 1930s, a handful of other researchers contributed to what Roberto Cordeschi has called 'the

---

<sup>12</sup> Brown, *The Definition of a Profession*, 3–17.

<sup>13</sup> Igo, *The Averaged American*.

<sup>14</sup> Mandler, *Return from the Natives*; Mead, *And Keep Your Powder Dry*. On social scientists as handmaidens for Cold War technocracy see: Rohde, *Armed with Expertise*.

<sup>15</sup> For Loeb, this explained why a moth steered towards a flame. Loeb, *Studies in General Physiology*; Cordeschi, *The Discovery of the Artificial*, 2002, 1.'

<sup>16</sup> B. Watson, 'Psychology as a Behaviorist Views It'.

<sup>17</sup> Cordeschi, 'The Discovery of the Artificial', July 1991, 220.

<sup>18</sup> As cited in: Draaisma, *Metaphors of Memory*, 140.

culture of the artificial,' in which mental processes were seen as independent of organic structures and could be verified as such.<sup>19</sup>

In 1935, intrigued by Hull's robot approach, the Russian-American mathematician Nicolas Rashevsky convened seminars on mathematical biophysics at the University of Chicago.<sup>20</sup> In a telling anecdote, Rashevsky's departmental chair mocked his ability to do meaningful work in psychology with only a desk and a pencil.<sup>21</sup> Undeterred, Rashevsky set about building networks for the creation of a new field, a 'systematic mathematical biology, similar in aim and structure to mathematical physics.'<sup>22</sup> In 1939, he founded the *Bulletin of Mathematical Biophysics*, an international journal for those who shared his view that theoretical mathematics could 'simplify' otherwise dauntingly complex biological processes, such as nerve conduction.

In 1943, the *Bulletin of Mathematical Biophysics* published 'A Logical Calculus of the Ideas Immanent in Nervous Activity' by Warren McCulloch, a neurophysiologist, and Walter Pitts, a self-taught mathematician twenty-five years his junior.<sup>23</sup> In it, the duo postulated that chains of idealised all-or-none 'psychic events' could be manipulated to represent propositional content such as calculus. Physiologists took little notice, as the study of neural networks was not new at the time.<sup>24</sup> McCulloch later credited the invention of the digital electronic computer for having popularized the paper, although the document itself made no mention of computing.<sup>25</sup> The connection came via John von Neumann, the Hungarian-American polymath, who referenced their work in 'First Draft of a Report on the EDVAC,' an early blueprint for how to build a digital computer.<sup>26</sup> The notion of an idealized neuron capable of simulating logical propositions appealed to von Neumann, who speculated on how

---

<sup>19</sup> Cordeschi, 'The Discovery of the Artificial', 218.

<sup>20</sup> Abraham, 'Nicolas Rashevsky's Mathematical Biophysics', 356.

<sup>21</sup> As retold in: Abraham, 357.

<sup>22</sup> Rashevsky, *Mathematical Biophysics*, vii; As cited in: Piccinini, 'The First Computational Theory of Mind and Brain', 181.

<sup>23</sup> Pitts, a mathematical prodigy, eventually studied under Carnap at the University of Chicago. McCulloch later claimed to have started toward this work in the mid-1920s. See: Piccinini, 'The First Computational Theory of Mind and Brain', 177–78.

<sup>24</sup> Piccinini, 175, 207.

<sup>25</sup> Boden, *Mind as Machine*, 194–95.

<sup>26</sup> Von Neumann, Aspray, and Burks, 'First Draft of a Report on the EDVAC'; Boden, *Mind as Machine*, 194–95; Haigh, Priestley, and Rope, 'Reconsidering the Stored-Program Concept'.

mathematical formalisms could map mental processes in a series of papers and talks in the late 1940s.

The success of early computing projects, and von Neumann's influence, lent legitimacy to McCulloch and Pitts' notion that neurons in the brain performed computations like bits in a computer. Literature by these authors, along with Hull, Rashevsky and others, informed the origins of cognitive science in the 1950s, which succeeded behaviourism as the dominant mode in psychology during the second half to the twentieth century. In 1943, Kenneth Craik at the University of Cambridge characterized this new horizon as the difference between 'synthetic' psychology, which examined 'the basic principles' underlying the functioning of the brain, and 'analytical' psychology, which explored anatomical, psychological and physiological phenomena.<sup>27</sup> Craik positioned *function* as more fundamental to psychology than *physicality*. He claimed that, by analogy, the laws of optics were of more fundamental significance to the study of light than the composition of a lens.

Like Rashevsky and von Neumann, Craik advocated disregarding complexity in living organisms in order to allow for mathematical symbols and mechanical componentry to reduce biological intricacies down to tractable functional relationships.<sup>28</sup> This shift away from the analytic character of behaviorism towards the synthetic character of cognitive science created an intellectual space in the United States in which mathematical metaphors between man and machine gained perceived authority.<sup>29</sup>

#### Cybernetics: The Field that Was Not a Field

It was in this milieu that cybernetics developed in the 1940s.<sup>30</sup> Broadly speaking, cybernetics was an interdisciplinary research area orientated toward the study of self-

---

<sup>27</sup> Lenzen and Craik, 'The Nature of Explanation.'

<sup>28</sup> Craik and Sherwood, 'The Mechanism of Human Action', 9.

<sup>29</sup> On this topic see: See: Gardner, *The Mind's New Science*, 140; Boden, *Mind as Machine*; Husbands, Holland, and Wheeler, *The Mechanical Mind in History*; Abraham, "'Microscopic Cybernetics'"; Piccinini, 'Computations and Computers in the Sciences of Mind and Brain'; Piccinini, 'The First Computational Theory of Mind and Brain'.

<sup>30</sup> On the history of cybernetics see: Heims, *The Cybernetics Group*; Mindell, *Between Human and Machine*; Rid, *Rise of the Machines*; Conway and Siegelman, For individual histories on select participants see: *Dark Hero of the Information Age*; Heims, *John Von Neumann and*

regulating mechanisms, like feedback and circular causality. A thermostat exemplifies this premise: it regulates temperature by incorporating feedback from its environment. Cyberneticists sought to understand an eclectic set of natural and non-natural systems using the same principles. Their work informed the design of automated artillery, ship steering and even human self-fashioning through anthropology, economics and art.<sup>31</sup> Those drawn to cybernetics sought to measure a system's state against its goal state, factoring in this information to anticipate and control change.

The term 'cybernetics,' from the Greek word for steersman, was popularized by Norbert Wiener in his 1948 book by that title. Wiener saw control and communications engineering as inextricable—in each domain a message, or sequence of messages, was distributed in time to steer behaviour by using the techniques of the 'calculus of variations' or statistical mechanics, the flow of which could be studied and understood.<sup>32</sup> Claude Shannon's development of information theory that same year formalized a mathematical measure of the amount of information encoded in a message. Much to Shannon's chagrin, this niche development in communications engineering was interpreted by many as relevant to biology, economics, psychology and other human systems in manner that suggested that each could be mapped through the prism of information management—a panacea.<sup>33</sup>

By coining new vocabulary, Wiener aimed to consolidate disparate lines of study under one field. Historians challenge the degree to which he was successful. Although cybernetics is frequently treated as if it were a coherent discipline, Ronald R. Kline argues more recently that it should be understood 'through the lens of disunity.'<sup>34</sup> Indeed, historians

---

*Norbert Wiener; Soni and Goodman, A Mind at Play; Abraham, Rebel Genius; Chaney, Runaway.* For contemporary accounts on the field as a whole see: Ashby, *An Introduction to Cybernetics*; Wiener, *Cybernetics*.

<sup>31</sup> On cybernetics as ontology, see: Galison, 'The Ontology of the Enemy'; Paidipaty, "'Tortoises All the Way Down'".

<sup>32</sup> Wiener, *Cybernetics*, 10.

<sup>33</sup> Shannon, 'The Bandwagon'.

<sup>34</sup> Similarly, Kline labels the transition from first-order cybernetics (1943-53) to second-order cybernetics (1970s-) a historiographical error. 'The two fields co-existed, rather uneasily as competitors for the diminished mantle of American cybernetics.' Kline, 'Why the Disunity of Cybernetics Matters to the History of the Human Sciences in the United States, 1940-1980', 2; Kline, *The Cybernetics Moment*, 153–65; Kline has also provided an excellent treatment of the Dartmouth workshop that gave AI its title: Kline, 'Cybernetics, Automata Studies, and the Dartmouth Conference on Artificial Intelligence' (hereafter 'Dartmouth').

discussing the genesis of cybernetics have offered diverse and inconsistent lists of its primary membership.<sup>35</sup> Similarly, the American Society for Cybernetics enumerates dozens of definitions, both from practitioners and commentators, on its website.<sup>36</sup>

After 1948, the term 'cybernetics' was used to group together almost a decade of prior research. Key texts included the aforementioned 1943 paper by McCulloch and Pitts, which formalized the mathematical study of idealized neural networks. Also central was the 1943 paper, 'Behavior, Purpose and Teleology,' published by Wiener, the Mexican physiologist Arturo Rosenblueth and Julian Bigelow, an engineer who made foundational contributions to digital computing.<sup>37</sup> At the time, 'teleological' was synonymous with 'unscientific' because it challenged the sanctity of causality by suggesting that an effect could guide its cause.<sup>38</sup> What the trio sought to isolate and understand was that class of object that exhibited behaviour that was not passive but active, and not random but guided, as if driven by a *purpose*, such as moving towards light. The trio argued that purposeful behaviour amounted to a system of difference reduction in the margin of error between a system's current state and its goal.<sup>39</sup>

Peter Galison traces how cybernetics effaced the distinctions between man and machine when equating purpose to the completion of a sequence of events. This gesture simultaneously humanized machines as capable of intention and dehumanized human beings by grading them as the product of functions.<sup>40</sup> As Steve Heims explicates, Wiener's proposition inspired a generation of social scientists to think of human connections through

---

<sup>35</sup> Andrew Pickering labels Wiener, McCulloch, W. Ross Ashby and William Grey Walter 'The Four Pioneers of Cybernetics.' With the exception of Wiener, a mathematician, each of these men worked primarily in psychiatric milieus. See Pickering, *The Cybernetic Brain*, 5. Edwards centers on both mathematicians, Wiener and Pitts, and brain scientists, Rosenblueth and McCulloch, Edwards, *The Closed World*, 240. Dupuy, following Heims, focuses on Wiener, Rosenblueth, McCulloch, Pitts, Julian Bigelow, and von Neumann, an orientation that lists mathematicians outnumbering brain scientists two to one. This list derives from the authorship of the two influential 1943 texts cited below. Dupuy, *The Mechanization of the Mind*, 67. Heims cites John Stroud and J. C. R. Licklider as other scientists working at this time 'for whom electronic engineering, psychology, and military requirements seemed to dovetail.' Heims, *The Cybernetics Group*, 75. Erwin Schrödinger, likewise, served to legitimize the use of physics and mathematics techniques in biology Schrödinger, *What Is Life?*

<sup>36</sup> 'American Society for Cybernetics - Defining Cybernetics'.

<sup>37</sup> Rosenblueth, Wiener, and Bigelow, 'Behavior, Purpose and Teleology'.

<sup>38</sup> Cordeschi, *The Discovery of the Artificial*, 2002, 117.

<sup>39</sup> Rosenblueth, Wiener, and Bigelow, 'Behavior, Purpose and Teleology', 19.

<sup>40</sup> Galison, 'The Ontology of the Enemy', 249. Visited also in: Crowther-Heyck, *Herbert A. Simon*, 187.



the lens of feedback mechanisms, including Margaret Mead, Gregory Bateson and Talcott Parsons.<sup>41</sup> By all accounts, Wiener's network was pivotal in securing a home for this interdisciplinary venture; his supper club guests Frank Fremont-Smith, of the Josiah Macy, Jr. Foundation, and Robert Morison, of the Rockefeller Foundation, funded meetings of a newly formed Teleological Society, as well as the Macy Conferences, which assembled key figures in the United States bi-annually between 1941-50. In a parallel development in the UK, The Ratio Club assembled key contributors between 1949-58.<sup>42</sup>

Despite this level of funding, Jean Pierre Dupuy describes cybernetics as a 'failure' in comparison to the study of information theory, which found a professional home in the Institute of Electrical and Electronics Engineers Society, which persists today.<sup>43</sup> Neither the Macy Conferences nor the Ratio Club had the capacity to formally train students, set a stable research agenda or grant degrees. This atrophied the domain's ability to propagate. As Andrew Pickering summarizes:

When we think of interdisciplinarity we usually think of collaborations across departments in the university,' he continued, 'The centre of gravity of cybernetics was not the university at all. Where was it then? The simplest answer is: nowhere. ...Cybernetics flourished in the interstices of a hegemonic modernity, largely lacking access to the means of reproduction: the educational system.'<sup>44</sup>

As I will show, the same argument cannot be made about mid-century AI. After World War II, MIT was America's largest non-industrial defence contractor, with nearly four times the investments of Harvard University, by then the nation's third largest non-industrial defence contractor.<sup>45</sup> This military-academic partnership was sustained during and after the late 1950s, when McCarthy and Minsky established the Artificial Intelligence Laboratory at MIT, one of three leading AI laboratories nationally during the mid-twentieth century. Similarly, Simon's research took place at the RAND Corporation, which at the time possessed perhaps

---

<sup>41</sup> Heims, *The Cybernetics Group*.

<sup>42</sup> Husbands and Holland, 'The Ratio Club'.

<sup>43</sup> Kline, *The Cybernetics Moment*, 3.

<sup>44</sup> Pickering, *The Cybernetic Brain*, 215–16.

<sup>45</sup> In 1962, the physicist Alvin Weinberg observed that MIT was either 'a university with many government research laboratories appended to it or a cluster of government research laboratories with a very good educational institution attached to it.' Cited in: Leslie, *The Cold War and American Science*, 14.

the largest computing facility in the world. Unlike cybernetics, institutional support for key researchers in what became the field of artificial intelligence was the rule, not the exception.<sup>46</sup>

### Computer Programming: Materiality, Mind and Machine

Another realm of key significance to the development of mid-century brain modelling and AI—and one I will return to often—was the emergence of digital electronic computing and, subsequently, programming. There is much to say about computing's histories, which can be understood through notions of conceptual pedigree, materiality, or labour, among other approaches. I will attempt to compress key historical developments here to provide a basis for targeted inquiry. Afterwards, I provide a brief historiography of AI and of computing.

The origins of what became computer science are often traced back the failure of the Hilbert Program to clarify the logical foundations of mathematics at the turn of twentieth century. In this scholarly crisis, the foundations of mathematics were found to be inconsistent and paradoxical. David Hilbert, a German mathematician, proposed that researchers identify a set of complete and finite axioms to which all results could be reduced. These axioms could then be proven consistent, which would resolve the crisis. Kurt Gödel's 1931 incompleteness theorems proved that a system *cannot* be both complete and consistent simultaneously. In 1936, Alan Turing resolved another aspect of Hilbert's dilemma, the 'decision problem,' which asked whether a procedure existed by which any mathematical statement could be proven true or false. Turing and Alonzo Church simultaneously proved that no such procedure could exist because some statements cannot be solved in a finite amount of time. Some cases, therefore, might run forever without halting. The method of abstraction Turing introduced to demonstrate this result is referred to as a Turing Machine. An imaginary machine of this type manipulates symbols on an infinite string of cells according to pre-programmed rules—an inventive notion that paved the way for what would become digital computing.<sup>47</sup>

---

<sup>46</sup> The conceptual boundaries between early AI and cybernetics are not well understood. AI is referenced only in passing, if at all, in various texts on cybernetics: Conway and Siegelman, *Dark Hero of the Information Age*; Heims, *The Cybernetics Group*; Pickering, *The Cybernetic Brain*; Mindell, *Between Human and Machine*; Rid, *Rise of the Machines*; Edwards, *The Closed World*, 239–74 is an exception, as is Dupuy, *The Mechanization of the Mind*, which surveys the metaphysical commitments underlying the boundaries between cybernetics and early AI.

<sup>47</sup> Gray, *Plato's Ghost*.

Between 1930-50, computing abstractions took new shape via two important advances in computer engineering. The first was electricity. In the 1930s, John Atanasoff, a physicist at Iowa State University, experimented with an electronic calculating device. In 1945, using sponsorship from the U.S. War Department, John Mauchly and J. Presper Eckert at the University of Pennsylvania developed the Electronic Numerical Integrator and Computer, or ENIAC, at the U.S. military's Aberdeen Proving Grounds. While the first 'computers' were human, often highly educated young women hired to solve advanced equations, often using rote transcription, ENIAC provided a glimpse at what might replace them—lightning fast *electronic* machines.

The second innovation was the advent of the stored-program—a concept analogous to Turing's notion of an abstract computational machine.<sup>48</sup> The stored-program was the portion of an electronic computer that held and reflexively adjusted data in a memory bank. This capacity allowed for vastly more sophisticated computation than a mechanical calculator. Whereas Harvard's electromechanical device, the Mark I, could perform two to three additions per second, for instance, ENIAC could perform five thousand.<sup>49</sup> Taken in combination with the added processing speed enabled by electrification, this step precipitated interest from a generation of researchers and engineers agog to know what an *electronic stored-program* computer could be made to do.

From the mid-to-late 1940s to the end of the 1950s, the craft and cultural status of 'setting up' or programming a computer changed dramatically.<sup>50</sup> On the ENIAC project, between 1943-45, this task was considered glorified clerical work. Convinced of its rote irrelevance, senior engineering staff tasked six young women to 'set up' the device using its cable-and-plug system, which resembled a telephone switchboard. 'We learned to diagnose troubles as well as, if not better than, the engineer,' remembered Betty Jean Jennings, the youngest of them.<sup>51</sup> In 1946, the mixed gender staff of Maurice Wilkes' Mathematical

---

<sup>48</sup> The concept of 'stored-program' has attracted controversy. It is demonstrated to be a historiographic artefact in: Haigh, Priestley, and Rope, 'Reconsidering the Stored-Program Concept'.

<sup>49</sup> See: Ensmenger, *The Computer Boys Take Over*, 33; For more on ENIAC see: Bergin, 'ENIAC: Development and Early Days'; Fritz, 'The Women of ENIAC'; Stern, 'Computers'.

<sup>50</sup> About 'firsts' in the history of computing, see: Rojas and Hashagen, *The First Computers*. On the craft practices involved in programming in this era, see: Haigh, Priestley, and Rope, *ENIAC in Action*. On the development of programming see: Priestley, *A Science of Operations*.

<sup>51</sup> Fritz, 'The Women of ENIAC', 13.

Laboratory at the University of Cambridge began construction on EDSAC.<sup>52</sup> A university mandate to avail the device to all staff and students led to the rapid innovation of a system of programmable rules, or sub-routines, that transformed the raw computing device into a precision calculating machine. In 1951, this registry of sub-routines was published as the first textbook on programming, which became ‘part of the air that researchers breathed’ in the 1950s.<sup>53</sup>

In 1950, only two electronic computers were in use in the United States. By 1955, this number grew to two hundred and forty-four.<sup>54</sup> In pioneering laboratories across the United States and Britain, coterie of mathematicians, physicists, electrical engineers and their hired staff—human computers, clerical administrators, shop boys, insightful spouses and students—engaged in an ongoing series of shop-floor debates and formal dialogues over what form stored-program computing could take and what functions it could serve. One distinctive characteristic of this period was that those who ‘set up’ a machine were still in close contact with those who engineered the hardware on which it ran.<sup>55</sup> As F. J. Gruenberger of RAND remembered, ‘In 1952, this was a give-and-take process of programmers saying “Can you do this?” and us replying “Yes, but this is easier.”’<sup>56</sup> By the late 1960s, such direct exchange was rare.

Paul Edwards has proposed that the core philosophical difference between cyberneticists and early AI researchers can be seen in the distinctions between hardware and software. Cybernetics modeled brains in hardware, AI modeled minds in software.<sup>57</sup> This

---

<sup>52</sup> EDSAC stood for: Electronic Delay Storage Automatic Calculator

<sup>53</sup> The ‘Planning and Coding’ reports were the only major account of programming prior to Wilkes, Wheeler, and Gill, *The Preparation of Programs for an Electronic Digital Computer*; Campbell-Kelly, ‘Programming the EDSAC’, 63. See: Campbell-Kelly, Jones, and Lloyd, ‘From Theory to Practice’, 35; Penn, ‘On the Logic of Labor’.

<sup>54</sup> In five year increments from 1960-70, this number grew from 5,400, to 25,000 and then 75,000. Ensmenger, *The Computer Boys Take Over*, 28.

<sup>55</sup> Heyck, ‘Defining the Computer’, April 2008, 57.

<sup>56</sup> RAND ran one of the largest installations of scientific computing in the world at the time. Gruenberger, ‘The History of the JOHNNIAC’, 50; Gruenberger, ‘Memorandum, RM-5654-PR: The History of the JOHNNIAC’, 12.

<sup>57</sup> To cite Edwards in full, ‘Cyberneticians tried to design self-organizing machines that would achieve complex behaviour through encounters with their environments. The subject... of cybernetics was always the embodied mind... The next intellectual generation... placed the emphasis of formal-mechanical modeling on the side of the formal, the disembodied, the

binary can be misleading. We have seen significant evidence that cybernetics was never a unified discipline. Nor, in the 1955-60 period I am interested in, was AI. Even the term 'software' did not come into use until around 1960.<sup>58</sup> It is worth reconsidering then, what mid-century attempts to formalise man-machine metaphors had in common. Craft knowledge is one answer. Computing in this period required careful use of paper tapes, programming cards, vacuum tubes and other unique instrumentation. Access to diverse communities of labour is another commonality that remains to be dealt with historically.<sup>59</sup> To provide a basic taxonomy for the diversity of practices and commitments that surrounded man-machine metaphors in this period, Margaret Boden differentiates between a 'strong' computer program and a 'weak' one. A strong program was one that a computer could operate. A weak program was closer to a mathematical framework; a conceptual outline not yet rendered in a manner that a digital computer could process as code. I borrow Boden's dichotomy in my own analysis.

The early 1950s brought a groundswell of activity on mind-as-computer formalisms. In 1950, Turing and Shannon both published influential papers on the topic. Each author outlined a weak program based in the tenets of behaviourism. In 'Computing Machinery and Intelligence,' Turing introduced a now famous sufficiency test for machine intelligence. To pass his imitation game, a concealed computer had to *behave* in a manner that a human judge deemed intelligent. Turing described his offering as a 'convenient fiction' because he assumed imaginary conditions like infinite computational power and perfectly discrete binary states. He conceded that he did not know how to engineer this imaginary concept, but he supported the creation of sense organs that could be taught, like a child, to perform abstract behaviours like playing chess.<sup>60</sup>

---

abstract—on the side of the mind rather than that of the brain.' Edwards, *The Closed World*, 240–41.

<sup>58</sup> Campbell-Kelly and Garcia-Swartz, 'Pragmatism, Not Ideology', 234.

<sup>59</sup> Unrelated to AI, Hicks illuminates the numerous unrecognized roles that women technologists played in the advance of the post-war computing industry in Britain, and how this group's alienation from the industry contributed to its eventual collapse. Hicks, *Programmed Inequality*.

<sup>60</sup> By Edwards' binary, Turing was both a cybernetician and an AI researcher. Turing's peculiar standing may have been a result of his relatively exceptional access to early electronic computing projects in Britain. Wiener, in comparison, had no interest in electronic computing. Turing, 'Computing Machinery and Intelligence', 460.

As computer programming matured during the 1950s, and access to electronic stored-program computers expanded across the United States and Britain, the perceived value of ‘weak’ programs began to wane in favour of ‘strong’ programs. In what follows, I examine this transition closely. Digital computing also underwent its first generational change: Turing died in 1954 and von Neumann in 1957. The Macy Conferences, Ratio Club and, more generally, first-order cybernetics declined by the early 1950s. Meanwhile, Minsky, McCarthy, Simon, Rosenblatt and their peers arrived at a horizon of untested possibility in stored-program electronic digital computing. My dissertation untangles this moment by clarifying the political and intellectual atmosphere within which this group worked.

#### Historiography from Within: Popular Accounts of AI History by Practitioners and their Kin

Having surveyed the early histories of digital computing and programming, I will now consider the historiography on mid-century brain modelling and AI in particular. One reason to focus on AI here is that the term ‘artificial intelligence’ has been used—assertively at times—to represent a broad swath of *other* research areas, such as the study of neural networks. As has been the case with practitioners, historical commentators have noted, and sometimes exploited, the term’s ambiguity. I offer a brief sketch of several influential popular accounts here in order to characterise this plasticity.<sup>61</sup> I focus on the methodological tendencies behind this conflation—as well the historiographical correctives put forward to deepen understandings of the history of AI.

Edwards describes two ‘quasi-official’ histories of AI: Pamela McCorduck’s *Machines Who Think*, published in 1979 and 2004, and Donald Crevier’s *AI: The Tumultuous History*,

---

<sup>61</sup> Edwards summarizes, about practitioners’ accounts, ‘The founders of AI have been much concerned to document their own history, which they (quite naturally) view mainly in intellectual terms.’ Edwards, *The Closed World*, 239. See: Newell, ‘Intellectual Issues in the History of Artificial Intelligence’; McCarthy, *Defending AI Research*; McCarthy, ‘John McCarthy’s Home Page’; Simon, *Models of My Life*; Solomonoff, ‘Ray Solomonoff and the Dartmouth Summer Research Project in Artificial Intelligence, 1956’ (hereafter, ‘Dartmouth’); Solomonoff, ‘Ray Solomonoff (1926-2009)’; Solomonoff, ‘Ray Solomonoff and the New Probability’; As a later example, see: Nilsson, *The Quest for Artificial Intelligence*. Oral Histories from various relevant figures are also available at the Charles Babbage Institute and MIT archives.

published in 1993.<sup>62</sup> According to these texts, the 1956 Dartmouth Summer Research Project on Artificial Intelligence marked the field's definitive origin. This two-month, ten-man gathering at Dartmouth College in Hanover, New Hampshire was convened to bring together a carefully selected group of academics and industrialists to, 'Proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.'<sup>63</sup> Participants such as Simon, McCarthy and Minsky later became prominent in the field.<sup>64</sup> In the most popular textbook on AI, Dartmouth and 1956 are described as the time and place of 'The birth of artificial intelligence.'<sup>65</sup>

McCorduck and Crevier's accounts staked the development of AI around conceptual rather than cultural, political, military, economic or colonial contingencies.<sup>66</sup> 'Even more than that of cognitive psychology,' Edwards summarizes, 'AI's story has been written as a pure history of ideas.'<sup>67</sup> This result may be due to training and methodology. McCorduck and Crevier relied heavily on interviews with key researchers. Neither took significant consideration of archival materials. Crevier, an engineering professor at the University of Quebec, relied on interviews with Minsky, Simon, Newell and others to construct his broad

---

<sup>62</sup> Edwards also listed work by Rheingold and Rose, each intended for a popular audience (Rose's book began as a GQ article). Edwards, *The Closed World*, 415; Rheingold, *Tools for Thought*; Rose, *Into the Heart of the Mind*. These articles' status as 'quasi-official' histories has dimmed since Edwards' assessment in 1997 due to the publication of new accounts of 1950-60s AI including: Skinner, *Building the Second Mind*; Boden, *Mind as Machine*; Cordeschi, 'AI Turns Fifty: Revisiting Its Origins'; Kline, 'Dartmouth'; Wilson, *Affect and Artificial Intelligence*; Olazaran, 'A Historical Sociology of Neural Network Research'; Olazaran, 'A Sociological Study of the Official History of the Perceptrons Controversy' (hereafter 'A Sociological'); Guice, 'Controversy and the State'.

<sup>63</sup> Minsky et al., 'A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence', 2 (hereafter, 'Proposal'). McCorduck, *Machines Who Think*; Crevier, *AI*.

<sup>64</sup> Along with Turing and Shannon, these figures are often described as the 'founding fathers' of AI. The term 'founding fathers' was likely introduced into AI's historiography in 1979 by McCorduck, *Machines Who Think*, xii; which 'forged the template for subsequent histories' according to: Mirowski, 'Book Review', 135. 'Founding fathers' is rehashed in: Nilsson, *The Quest for Artificial Intelligence*, 80; Gardner, *The Mind's New Science*, 30; Sabanovic, Milojevic, and Kaur, 'John McCarthy [History]', 99. It is also used in 4000+ other books related to AI, according to a search of those terms on Google Books in late 2017.

<sup>65</sup> Russell, Norvig, and Davis, *Artificial Intelligence*, 17.

<sup>66</sup> The same is largely true of Garnham, *Artificial Intelligence*; Haugeland, *Artificial Intelligence*; Copeland, *Artificial Intelligence*.

<sup>67</sup> Edwards, *The Closed World*, 239.

but relatively shallow history of AI research between 1950-90s. McCorduck's comparatively focused account of AI's genesis as a discipline in the 1950-70s drew from her evolving friendships with Simon and Newell, who were colleagues of her partner Joseph F. Traub at Carnegie Mellon University. Newell offered funding to McCorduck for early interviews and travel expenses after she had been turned down by a number of foundations and government agencies for being, in her words, 'Merely a writer and not a trained historian of science.'<sup>68</sup>

In 2004, Simon initiated a second edition of *Machines Who Think* out of frustration over the influence of other 'unreliable' historical sources.<sup>69</sup> Philip Mirowski criticizes McCorduck for equating the origins of AI with the origins of symbolic-AI, which Simon helped to establish. 'The text is redolent of Simon's exuberant opinions and personality,' Mirowski wrote.<sup>70</sup> One result was the obfuscation of AI's conceptual debts to operations research, a field developed during World War II to exact precision warfare through cost-benefit analysis of possible actions related to man-machine systems.<sup>71</sup> McCorduck had noted that her text was not meant to be exhaustive; its subtitle is, 'A Personal Inquiry into the History and Prospects of Artificial Intelligence.' Still, due to a lack of more rigorous competing narratives, the book 'forged the template for subsequent histories' on the origins of AI.<sup>72</sup>

Terminology plays an important role in this historiography. In other popular and practitioner accounts, terms of art from within the study of AI, such as 'symbolic AI,' 'sub symbolic AI' and 'Good Old Fashioned AI (GOFAI),' are often used to structure clusters of research into coherent chronologies and trends. In many ways, these terms are invaluable. 'Symbolic AI' describes an influential stream of research premised on the assumption that intelligence can be reduced to a string of symbols. 'GOFAI' is the 1950-60s period in which symbolic AI was first formalized.<sup>73</sup> 'Sub-symbolic AI' denotes a competing paradigm to

---

<sup>68</sup> McCorduck, *Machines Who Think*, 2004, xii.

<sup>69</sup> McCorduck, *Machines Who Think*, xiii–xiv; Newell eventually wrote his own history of the field: Newell, 'Intellectual Issues in the History of Artificial Intelligence'.

<sup>70</sup> Mirowski, 'Book Review', 135.

<sup>71</sup> Elsewhere Mirowski laments that neither historians of science nor economics have meaningfully engaged with the legacy of operations research. Definition paraphrased from Ellis Johnson via: Mirowski, *Machine Dreams*, 2002, 177–78.

<sup>72</sup> Mirowski, 'Book Review', 135; Gardner, *The Mind's New Science*, 138.

<sup>73</sup> 'Good Old Fashioned Artificial Intelligence' was coined in 1985 by the philosopher John Haugeland to characterise what has elsewhere been called 'symbolic,' 'classical' or 'traditional' AI. 'Symbolic' in Haugeland, *Artificial Intelligence*, 112–23; 'classical' in McCorduck, *Machines Who Think*, 389; and 'traditional' in Boden, *Mind as Machine*, 701.



symbolic AI in which research is based around the simulation of learning. The latter, known alternatively as ‘machine learning’ or ‘connectionism,’ applies mathematics to idealized neural networks to draw out fine grained statistical patterns in data.<sup>74</sup> Sub-symbolic AI has been historicized as the tradition to have succeeded GOFAI in the 1980-90s, although this framing has been disputed.<sup>75</sup>

Use of these terms raises important questions about the reification of anachronistic narratives. For instance, it is well known to historians of computing and cybernetics that research on idealized neural networks predated the introduction of the term ‘AI’ by more than a decade. The expression ‘sub-symbolic AI’ thus serves to rhetorically and anachronistically annex decades of neural network research under the banner of ‘AI.’ Also troubling is that these terms of art chronicle events *backwards* rather than forwards. None of scientists I study used terms like GOFAI during the 1950-60s. Nor, for that matter, did they all see themselves as contributing to ‘AI’ prior to around 1960, although that changed as elite organizations like MIT provided AI with an institutional base.

David F. Noble advocates for the history of automation to be told ‘in the present tense,’ meaning *forward*—as it emerged.<sup>76</sup> That this practice has been overlooked by AI researchers interested in the field’s history is exhibited in the opening line of Nils J. Nilsson’s popular 2010 book *The Quest for Artificial Intelligence*. ‘Artificial intelligence (AI) may lack an agreed-upon definition, but someone writing about its history must have some kind of definition in mind. For me...’<sup>77</sup> It is from this perspective that present terminology gets written back into the past. In his own historical survey, the influential researcher Allen Newell opined, ‘Ultimately, we will get real histories of Artificial Intelligence... written with as much

---

<sup>74</sup> The origins of the term connectionism have been traced to Hebb, *The Organization of Behavior*, xix.

<sup>75</sup> The histories of machine learning and neural networks have been explored on their own terms, separate from AI, in: Plasek, ‘On the Cruelty of Really Writing a History of Machine Learning’; Guice, ‘Controversy and the State’; Olazaran, ‘A Sociological Study’; The term artificial intelligence re-entered the cultural zeitgeist in 2010 following the mass adoption of smartphone technology, technological and exploitative reductions in the cost of labeling large data sets via platform ‘clickwork,’ breakthroughs in deep learning and heavy marketing investments by major technology firms keen to capitalize on public curiosity. See: Katz, ‘Manufacturing an Artificial Intelligence Revolution’.

<sup>76</sup> Noble, *Progress without People*, 6. On ‘the Automation Movement’ of the 1950s see: Brock, ‘From Automation to Silicon Valley’.

<sup>77</sup> Nilsson, *The Quest for Artificial Intelligence*, 13.

objectivity as the historians of science can muster.<sup>78</sup> In 2006, Minsky suggested that this indeterminacy was integral to AI when he famously described the term as a ‘suitcase word’ due to its many divergent interpretations.<sup>79</sup>

Lucy Suchman has suggested that AI is not a suitcase but a mirror, one that ‘works as a powerful disclosing agent for assumptions about the human.’<sup>80</sup> Most histories of AI as a concept do not consider this reflexivity; they treat the field’s imbrications in structural contingencies as extraneous.<sup>81</sup> Proposals for how to improve historical understandings of AI repeat this tendency. Jürgen Schmidhuber, an AI researcher turned historian, nominates Gödel’s 1931 incompleteness theorem as a more natural starting point for the field than the 1956 workshop because the theorem laid the foundations for theoretical computer science and, ‘It is the thing that counts, not its name.’<sup>82</sup> Boden, whose formidable 2006 tome on the history of cognitive science provides many helpful windows into the history of AI, cites the 1958 Symposium on the Mechanisation of Thought Processes as its own ‘Memorable catalyst for the growth of an intellectual community.’<sup>83</sup> Still, the Dartmouth origin story—and the vocabulary and timeline it is seen to assert—endures.

Reading the history of AI forward requires identifying the complex political and socio-technical nuances that contributed to the field’s rapid emergence in the mid-century. In what follows, I flag, as I hope others will too, instances in which influential scientists made small or large assumptions, either inadvertently or without reflection, about the *universal* nature of their research, meaning its reliable bearing across distinct settings, from haloed matters of human cognition (across all cultures...) to practical considerations of how to organize a research library or manufacturing business using their sought-after theory of intelligence. I shed light on how their context lent authority to this premise. By interrogating Simon, Rosenblatt, McCarthy and Minsky’s ground-level assumptions—and by treating those assumptions as claims—I aim to connect these men to the imaginaries they helped to propagate, to the network of labour they drew upon to advance their ideas, and to the materiality of their experimental processes; core considerations that have been stripped

---

<sup>78</sup> Newell, ‘Intellectual Issues in the History of Artificial Intelligence’, 1.

<sup>79</sup> Minsky, *Emotion Machine*, 12, 88.

<sup>80</sup> Suchman, *Human-Machine Reconfigurations*, 226.

<sup>81</sup> Cohen-Cole, ‘Review of Mind as Machine’. In Chapter Six, I speculate on why this is the case.

<sup>82</sup> Schmidhuber, ‘Celebrating 75 Years of AI’, 30.

<sup>83</sup> Boden, *Mind as Machine*, 336.

away and forgotten by a field awkwardly premised on an appeal to the relevance of simulation.

Since no account of AI's history has yet earned broad acceptance among scholarly historians, it is neither radical nor alternative for me to suggest that, ultimately, AI may come to be understood as a branch of political science postured and practiced as a branch of computer science. I suggest this because, from a historical perspective, the field has emerged from and relied on the particular social, industrial and supply chain infrastructures of digital computing (primarily but not exclusively) since the two industrial and academic areas emerged in the mid-twentieth century. AI is unique in that the discipline's formalisms are also used to *shape* these infrastructures by offering specific logics for their organization. Treating AI as a branch of political science might deepen scholar's interrogation of how these infrastructures serve as *desired* modes of human organization rather than the incidental ecosystem of an abstract conceptual project.

The groundwork for a reframing of this kind has been laid by prior writers whose research orbits the history of artificial intelligence, broadly construed—and it must engage thoroughly with the power and limitations of earlier research on computing.<sup>84</sup> James Fleck writes about how nepotism served to construct and reconstruct multiple generations of an institutionally homogenous AI Establishment at research centres in the U.S. and UK.<sup>85</sup> In various in-depth case studies, Stephanie Dick explicates how the materiality of mid-century computing projects shaped American scientists' notions of what constitutes mathematical proof, an approach that undermines symbolic AI's presumptions of universality.<sup>86</sup> Mikel Olazaran renders legible the dramatic contest for authority held between neural network and symbolic AI researchers in the 1970s, demonstrating how power over disciplinary narratives was deployed to sustain artificial intelligence amidst crisis.<sup>87</sup> Each scholar's work illuminates significant contingencies that await careful consolidation. I contribute my own work to this horizon and return to these author's writings throughout.

---

<sup>84</sup> Although not historical per se, an argument to this effect was recently taken up in: Lanier and Weyl, 'AI Is an Ideology, Not a Technology'; and Katz, 'Manufacturing an Artificial Intelligence Revolution'.

<sup>85</sup> Fleck, 'Development and Establishment in Artificial Intelligence'.

<sup>86</sup> Dick, 'After Math'.

<sup>87</sup> Olazaran, 'A Sociological Study'; Olazaran, 'A Historical Sociology of Neural Network Research'.

Having considered prominent works in the historiography of AI and addressed some of the provocative interpretative questions that remain open, let me now turn to the historiography of computing, to which AI research has been tightly coupled. As with cybernetics, digital computing has been treated as a coherent and connected unit rather than a heterogenous series of developments. This is largely due to historiographical biases. Histories of computing have tended to track progress along two axes: the history of ideas and the history of engineering.<sup>88</sup> The first genre links developments such as Plato's investigations on knowledge, Leibniz's study of rationalism and Ada Lovelace's musings on Charles Babbage's Analytical Engine. During and after the Second World War, these precedents are taken to have been transmuted into machinery by Turing, von Neumann, Wiener, cyberneticists and others, including the figures I discuss herein. This is canon according to computer scientists, cognitive scientists and historians of the computer as a concept.

The history of engineering narrative, in contrast, traces progress in terms of design specifications. This genre became popular in the 1980-90s as the history of computing found its footing as a discipline. Ancient myths of handcrafted automatons, such as Hephaestus's golden handmaidens or Talos colossus, were (and continue to be) linked to seventeenth century calculating machines by Pascal or Leibniz, to nineteenth century artefacts from Babbage or Hollerith, twentieth century digital devices by Bush, Turing and von Neumann and twenty-first century outputs like AlphaGo or the Terminator.<sup>89</sup> A portion of these accounts center on the agency of laboratories, industrial workshops and 'the buccaneers who built the companies whose logos grace our laptops.'<sup>90</sup> In this history of engineering tradition, debate surrounds who came first across a particular technological finish line, an approach that has

---

<sup>88</sup> I borrow this dichotomy from: Edwards, *The Closed World*, ix–xvi.

<sup>89</sup> For a recent survey of the pre-history of modern automation narratives and techniques that originated in the west, see: Cave and Dihal, 'Ancient Dreams of Intelligent Machines'.

<sup>90</sup> Quote from Mirowski, *Machine Dreams*, 77; Examples from this genre include: Cohen, *Howard Aiken*; Copeland, 'The Manchester Computer'; Aspray, 'The Stored Program Concept'; Goldstine, *The Computer*; Ferry, *A Computer Called LEO*; Bashe, *IBM's Early Computers*; Stern, 'Computers'; Austrian, 'The Machine That Carried IBM into the Electronics Business'; Isaacson, *The Innovators*; Levy, *Hackers*.

since been challenged as antiquated and pedantic.<sup>91</sup> In relation to AI, scholars in this genre like Herbert Stoyan have productively traced the development of tools like the LISP programming language, which became the *lingua franca* of U.S. based AI research in the 1960-70s.<sup>92</sup>

Significant overlap between these two genres undermines the usefulness of the dichotomy. Each chronicle changes primarily through an overlapping set of individual figures, such as Leibniz, von Neumann and Turing. Cordeschi and Kline subvert this trajectory when arguing, in relation to AI and cybernetics, that disciplinary boundaries are of limited explanatory value to historical understandings of brain theory. Cordeschi locates AI and cybernetics within an arc that spans the entire twentieth century, which he calls the 'discovery of the artificial.'<sup>93</sup> While his analysis centres on individual contributors, it incorporates many whose work has so far been overlooked, such as Hull's notion of psychic machines. Cordeschi argues that, once change is examined across a broader set of contributors, one sees that advances in large-scale computer power and memory determined disciplinary outcomes, such as the success of symbolic AI over neural networks. In making this case, Cordeschi advances a form of technological determinism. He simultaneously asserts, however, that *multiplicity* characterizes the sciences of the artificial, as Michael Mahoney and Thomas Haigh do when they question whether the history of computing is better approached as *histories* of computing(s).<sup>94</sup>

Multiplicity of purpose and heterogeneity of practice do not preclude the notion that computing's histories have uniform ideological elements that imply and impose social assumptions. Kate M. Miltner surveys a generational turn in the history of computing after the late 1990s led by Sadie Plant, Jennifer Light, Janet Abbate and others whose research exposes how patriarchal conceptions of progress have caused masculine assumptions around expertise, technological skill and contributed labour to be understood teleologically in historical writings.<sup>95</sup> These accounts 'implicitly endorse an ahistorical fiction of technological

---

<sup>91</sup> See: Haigh, Priestley, and Rope, 'Reconsidering the Stored-Program Concept'.

<sup>92</sup> On the history of LISP see: Stoyan, 'LISP History' and other work by that author.

<sup>93</sup> Cordeschi, *The Discovery of the Artificial*. For a summary of key themes see: Cordeschi, 'AI Turns Fifty: Revisiting Its Origins'.

<sup>94</sup> Haigh, 'Masculinities in the Histories of Computing(s)', 12–15; Haigh points to parallel insights in: Mahoney, 'The Histories of Computing(s)', 119–35.

<sup>95</sup> Miltner, 'Girls Who Coded'.

meritocracy,' writes Marie Hicks, whose work explicates how digital computing was used to marginalize already disenfranchised peoples like women technologists and transgender citizens in the United Kingdom in the second half of the twentieth century.<sup>96</sup> Indeed, gendered labour and oppression are consistent trends across multiple branches in the development of computing systems over the twentieth century.

A second and closely related point of continuity in the histories of computing is the rapid expansion of industrial, military and government bureaucracy. Various scholars position aspects of computer history within the histories of public administration and, in the case of digital computing, its overlap with the bureaucratization of the U.S. military. Edwards explores how the Cold War both shaped and was shaped by the computer metaphor.<sup>97</sup> He argues that the electronic computer came to be seen as a tool and symbol of U.S. military aims; a sophisticated means to close and control branches of geopolitical risk such the Soviet nuclear threat. Mirowski charts how this military infrastructure imprinted notions of information processing on American Cold War economics through institutional entanglements between academics centers like the Cowles Commission at the University of Chicago, and military funded research institutes like the RAND Corporation and MIT.<sup>98</sup> Mirowski examines, for instance, how Friedrich Hayek used the term 'information' in his description of market coordination, and how critics like Oskar Lange described the market process as 'a computing device of the pre-electronic age.'<sup>99</sup>

Agar charts the pre-history of these trends in the entangled trajectories of information technology and the British Civil Service during the nineteenth and twentieth centuries. He shows how the attribution of mechanical characteristics to the Civil Service served to legitimize projects of mechanization such as new information technologies. For instance, Agar shows that Charles Babbage's Analytical Engine was as much a political machine as a mathematical one; Babbage's Difference Engine was subsidized by government funding on the expectation that it would serve that sponsor. As the Civil Service ballooned from 16,267 to 460,000 employees between 1797-1999, shifting workplace demands compelled the

---

<sup>96</sup> Hicks, *Programmed Inequality*, 16; Hicks, 'Hacking the Cis-Tem', 1 January 2019.

<sup>97</sup> Edwards, *The Closed World*.

<sup>98</sup> Mirowski, *Machine Dreams*, 2002.

<sup>99</sup> Rogan, 'Know-How'; Mirowski, *Machine Dreams*, 2002, 236.

adoption of new tools, such as the use of Hollerith punched-card machinery to process the 1920 British colonial census of Egypt, a task that required fifteen million census cards.<sup>100</sup>

Agar reasons that in the 1950-60s, the trend towards state-managed control of information escalated into a low-profile culture of deference to technical expertise, a tradition he terms, 'discrete modernism.'<sup>101</sup> He contrasts his view that the computer was a tool and symbol of bureaucracy to Edwards' thesis that the computer was a tool and symbol of U.S. military aims. Their arguments are not mutually exclusive, he reasons, because the U.S. military became heavily bureaucratized in the mid-twentieth century.<sup>102</sup>

Agar's multi-century study of the origins of modern computing technologies points to a perceived equivalence between the functionality of the British Civil Service and the general-purpose computer. Each system processed large quantities of information using a hierarchy of pre-set but adaptable rules.<sup>103</sup> Agar could not decide whether this isomorphism was a coincidence or something more profound. That the bureaucratic system *preceded* its technological doppelganger made this isomorphism all the more compelling. In contrast, Dick insists that the replacement of existing paper-based practices and techniques with computers necessitated the development and adoption of definitively *new* practices, meaning that Agar's analogy may fail if judged at a granular level.<sup>104</sup>

Given idiosyncrasies of this type, the politics of computational practices receive sustained scrutiny across numerous historical sites. Matthew Jones elucidates how hubristic attempts by Pascal, Leibniz, Robert Hooke, Babbage and numerous others to mechanize calculation in Europe during the seventeenth to nineteenth centuries motivated important acts of emulation (and poaching), even as they failed. Consistent across these meme-like attempts was a reliance on artisanal knowledge and manufacturing prowess, rather than on the calibre of the abstractions of any individual 'genius,' a shift of historical focus that Jones uses to draw out that materiality has *long* mattered to the development of computational systems.<sup>105</sup> Elizabeth Yale shows that the archive has also long been a vehicle for complex political authority, decision making and modes of governance; from the Venetian empire's

---

<sup>100</sup> Agar, *The Government Machine*, 46, 161.

<sup>101</sup> Agar, 424–30.

<sup>102</sup> Agar, 419.

<sup>103</sup> Agar, 391.

<sup>104</sup> Dick, 'After Math', 12.

<sup>105</sup> Jones, *Reckoning with Matter*.

use of indexing in the mid-fifteenth century, to fears about Big Data today.<sup>106</sup> Ensmenger, building on work by Light, Abbate and Plant, charts how professionalization was used to construct gendered expectations in computing as modes of practice shifted from an esoteric military and academic research ventures to mainstream commercial settings with standardized devices, programming languages, analyst roles and jargon. He argues that it was *people*, after all, that proffered a computer revolution and then commandeered it, to their own benefit, as necessarily masculine in nature—a birthright of the ‘computer boys.’

Given these historiographical trends, my decision to trace developments through the work of four individual men may seem of limited explanatory value. While a primary focus on individual figures can limit the explication of structural change, it also has virtues, two of which are important here. The first, which I pursue, is a close reading of the metaphors and conceptual analogies used by key scientists to seed, advance and legitimize their research. It is well known (but not well documented, with notable exceptions) that grandstanding and poor citation practices have been longstanding vices in the AI community.<sup>107</sup> Hubert L. Dreyfus calls attention to AI’s persistently unacknowledged conceptual debts to four hundred years of rationalist philosophy and the individual men who championed it.<sup>108</sup> Sophie Smith inquires into how Thomas Hobbes, whom George Dyson calls ‘the patriarch of artificial intelligence,’ used his talents as a salesman to lay the foundations for Western statecraft in the seventeenth century—another node in a suggestive historical trend between scholarly men of power and their naturalization of social/cognitive organization through boosterish appeals to the self-evidence of abstraction and autonomy.<sup>109</sup> With this dissertation, I provide a sustained analysis of how Simon, Rosenblatt, McCarthy and Minsky framed their generative early-career research to different audiences. My vignettes inform understandings of the

---

<sup>106</sup> Yale, ‘The History of Archives’, 337, 349. AI researcher Edward A. Feigenbaum is credited by some for popularising the Francis Bacon line, ‘knowledge is power.’ An Interview with Edward A. Feigenbaum.

<sup>107</sup> Garvey, ‘The “General Problem Solver” Doesn’t Exist’.

<sup>108</sup> See: Dreyfus, ‘Why Heideggerian AI Failed and How Fixing It Would Require Making It More Heideggerian’, 331–32; Dreyfus, ‘Alchemy and Artificial Intelligence’; Dreyfus, *What Computers Can’t Do*; Dreyfus, *What Computers Still Can’t Do*.

<sup>109</sup> Dyson, *Darwin Among the Machines*, 7; Hobbes famously stated that civil philosophy was ‘no older’ than his own book on the topic. See: Hobbes, *Leviathan*; Smith, ‘The Language of “Political Science” in Early Modern Europe’, 203. Similarly, Prony’s appetite for government-approved spectacle is outlined in: Daston, ‘Enlightenment Calculations’, 189.



complex interplay between salesmanship and patronage in Cold War brain model research at elite American institutions.

A second, related advantage of this methodology is clarity around how key scientists fit in relation to the intellectual, political and social systems in which they were embedded. Simon Schaffer, in 1994, called for historians to interrogate the *sites* of intelligence. He showed how Charles Babbage leveraged deft rhetoric, showmanship and shifting norms in workplace culture to equate (to his own advantage) the philosophy of manufacturing with the philosophy of mind. By claiming to have endowed his Analytical Engine with notions of ‘memory’ and ‘insight,’ Babbage embodied his own control in the device while disembodying, and camouflaging, the skills of the work force on which its locomotion depended.<sup>110</sup> Schaffer underlines that automation has long been a matter of perspective. In nineteenth century London, some, like Karl Marx, focused on the role of workers in automation; others, like Babbage, emphasized the role of ‘automation’ alone.

The site I explore is the United States. Specifically, I examine faculty at elite universities on the eastern seaboard and their imbrications in industrial and military projects during the early Cold War. These institutions are their own window into continuities in the histories of AI and computing. Hunter Heyck shows how a high modernist faith in the power of science to reorder nature altered American research methodologies in economics, political science, sociology and other social sciences in the post-war years; this faith led researchers to shape each domain’s subject matter, regardless of its complexity, into a set of interconnected ‘systems’ fit for unrestricted deconstruction and analysis.<sup>111</sup> This trend has geopolitical implications. James C. Scott deconstructs how a similar blind faith in science to ‘modernize’ other cultures compelled disastrous western attempts to structure and optimize social dynamics in Tanzania, Russia and China in the second half of the twentieth century— attempts that failed due to their hubristic neglect, or denial, of local knowledge.<sup>112</sup> Similarly, and in a matter akin to the hubris surrounding positivist branches of AI, Jennifer Karns

---

<sup>110</sup> Schaffer, ‘Babbage’s Intelligence’, 214. For more on the interplay between scientific discourse and the consolidation of political power, see: Shapin and Schaffer, *Leviathan and the Air-Pump*.

<sup>111</sup> Crowther-Heyck, *Age of System*. Heyck examines Simon's driving role in this transformation in Crowther-Heyck, *Herbert A. Simon*; and the two-part paper: Heyck, ‘Defining the Computer’, April 2008; Heyck, ‘Defining the Computer’, April 2008.

<sup>112</sup> Scott, *Seeing Like a State*.

Alexander cites 1950-70s modernization theory as an embodiment of the west's obsession with 'a *single* route to social and economic advance,' in that case via a fast track to 'modernity.'<sup>113</sup>

A common point in historical accounts of the intersections between computing, bureaucracy and state ambitions is that computation is a highly collaborative rather than a distinctly individual craft and activity. In titling my dissertation 'Inventing Intelligence,' I seek to reintegrate the *institutional* actors whose agendas influenced the design of artificial intelligence, complex information processing, heuristic programming and surrounding notions of machine and natural intelligence in the mid-century.<sup>114</sup> I focus on individual sites of practice to separate, as best as is possible, genuine invention from speculative rhetoric. I aim to isolate the role that disingenuous claims and oversimplified historical accounts have had in shaping the field(s) pursued by specific individuals. In doing so I also take seriously the agency that institutions exerted as a significant form of 'intelligence' that helped shape the mathematical and metaphorical 'artificial intelligence' these men aimed to articulate, and liberate.

### Summary

In the 1950s, researchers crystallized a set of compelling mathematical metaphors between man and machine and set a research agenda for how to demonstrate the legitimacy of these modes of explanation. The explanatory limits of this rhetoric were inconsistently acknowledged by key researchers, their funders and many of those who have dealt with their work historically. The historiography of adjacent fields like cybernetics and digital computing suggests a need for refined interpretations of how these inquiries intersected with political outcomes such as post-war social engineering in the United States. As of yet, no comprehensive historical narrative clarifies how the evolution of AI intersected with ever larger modern bureaucracies, ever faster digital computers and ever more ruthless cultural ideals of efficiency and idealized mental behaviour in the United States in the second half of the twentieth century. My research makes a start in that direction by contextualizing a set of

---

<sup>113</sup> Emphasis mine. Alexander, *The Mantra of Efficiency*, 8.

<sup>114</sup> For concise comparisons between corporations as artificial decision systems and AI, see: Runciman, 'Diary'; Penn, 'AI Thinks like a Corporation—and That's Worrying'.

disciplinary, institutional and geopolitical influences on core conceptual precedents offered by Simon, Rosenblatt, McCarthy and Minsky in the 1950s.

My account questions the margins between early theory and performance in the history of artificial intelligence. In the following chapters, I revisit the work of four researchers to understand their distinct perspectives and the manners in which institutional affiliations served to mobilize, conform or throttle certain research outcomes and not others. I question how AI was influenced by norms, assumptions and neighboring work in management science in considering the work of Herbert Simon and others in Chapter Two; by experimental psychopathology and Hayekian economics in the work of Frank Rosenblatt examined in Chapter Three; and by automatic coding through the efforts of John McCarthy studied in Chapter Four. In Chapter Five, I examine the development of Minsky's 1961 paper, 'Steps Toward Artificial Intelligence,' which consolidated elements of these traditions into what became a popular research agenda for that new field.

## Chapter Two – On the Logic Behind the Logic Theory Machine

So called ‘intelligent’ machinery took many forms over the course of the twentieth century. Examples include the 1947 Westinghouse Network Calculator, which could solve intricate calculations in only an hour, and ‘Bessy the Bussell Engine,’ Harvard’s first sequence-controlled calculator. Only one device, the Logic Theory Machine, has been designated ‘The First AI Program’ by existing histories of that field.<sup>115</sup> The Logic Theory Machine was a virtual machine designed at the RAND Corporation in 1955-56 by Herbert A. Simon, Allen Newell and J. Clifford Shaw to discover and construct proofs from *Principia Mathematica*. In 1956, after using the Machine to prove thirty-eight theorems from that text, the trio described their results as evidence of having modelled human problem solving using an electronic digital computer.

In this chapter, I pay particular attention to articulating how Newell, Simon and Shaw developed and justified their comparison between the operations of the Logic Theory Machine and the capacities of a human being. My account focuses on Herbert A. Simon, whose contributions to the project via administrative science have received limited historical treatment.<sup>116</sup> I begin by introducing Simon’s training as a political scientist at the University of Chicago in the 1930-40s. I illustrate how that experience shaped his positivist view of the social sciences and subsequently, in the 1950s, his design of the Logic Theory Machine. Simon’s training under Rudolf Carnap, a member of the Vienna Circle of logical positivists, encouraged him to develop a ‘scientific’ account of administrative organization, a field that

---

<sup>115</sup> For a survey of instances see: Boden, *Mind as Machine*, 705.

<sup>116</sup> With some exceptions. Agre, ‘Hierarchy and History in Simon’s “Architecture of Complexity”’ connects Simon’s initial work on organizational structures to his *later* writings on complexity and cognition in the 1960s; Boden, *Mind as Machine*, 317–23, 710–13 positions Simon’s entire career as an extension of his initial work on management science, following comments from Simon himself; Daston, ‘Simon and the Sirens’ explores the same line of inquiry, but in relation to Simon’s later work on bounded rationality; Heyck, ‘Defining the Computer’, April 2008 assess Simon’s ‘bureaucratization of the mind’ through the broader lens of systems sciences (these are two articles, not one) ; Dick, ‘Of Models and Machines’ provides a condensed account, and Dick, ‘After Math’ an in-depth account, of the materiality of the data and logical architectures used by Simon and his colleagues to construct the virtual machine. I cite these accounts throughout.

he believed lacked empirical rigor.<sup>117</sup> In 1947, Simon published *Administrative Behavior*, an influential text developed to equip practical administrators and their students with linguistic and conceptual tools to analyse organizations scientifically. The book posited that all activity in an organization could be reduced to explainable decision-making processes.<sup>118</sup>

Following this, I show how Simon transmuted the linguistic and conceptual frameworks developed in *Administrative Behavior* into a new medium—electronic digital computing—after he joined the RAND Corporation as a consultant in 1952. Between 1952-54, Newell and Shaw struggled to create a chess playing program because they lacked a coherent ‘language’ with which to translate between complex human thought processes and the operations of a machine. In 1955, the trio elected to model human problem solving using the system of formal logic put forward in *Principia Mathematica*, which Simon had studied in Chicago. This text provided the group with axioms, rules of inference and methods that blended the enumerated, iterative ‘language’ of a machine with the more adaptive, heuristic logic exhibited in human problem-solving behaviour. To process this logic in a machine, they borrowed heuristics like means-ends analysis and decision premises from *Administrative Behavior*.

In this chapter, I demonstrate that the Logic Theory Machine was a model of human behaviour, not biology. Furthermore, it was a model of a certain type of behaviour, namely problem solving. More precisely, in 1955-56 the Machine modelled a sub-class of human problem solving that was a composite of early twentieth century British symbolic logic and the American administrative logic of a hyper-rationalized organization. I call attention to the fact that, at the RAND Corporation during the rise of Cold War era civilian science in the mid-1950s, these logics were celebrated. The group proudly cast the Logic Theory Machine as a contribution to the study of operations research, the science of military and business logistics and combative precision. These pronouncements invite reflection on the larger social and economic systems that Simon saw himself contributing to—namely the perpetuation of modern capitalism and U.S. military supremacy. His interlocutors at RAND valued the same

---

<sup>117</sup> Administration and organization are used interchangeably in Simon’s writing on the topic. I follow this convention.

<sup>118</sup> By Simon’s account, notions like means-ends analysis, which compared initial conditions to desired end conditions, and decision premises, which set the frame of reference in which decisions were made, were empirically valid tools one could employ to solve internal organizational problems such as poor efficiency and communication.

social order. A breakthrough, for the project's funders, would resolve mounting anxieties about exposure to new and ambiguous levels of risk amidst a deepening Cold War. All parties thus sought a means to reduce the world's complexity down to a scale they could aspire to control.

#### 1930-40s: Simon's Scientific Account of Administration

Herbert A. Simon's conception of what a social scientist ought to be was shaped by his education at the University of Chicago in the 1930-40s, where he enrolled as a political science undergraduate in 1933. At the time, the political science curriculum at Chicago followed a 'New Plan' built around the pragmatist ideals of figures like John Dewey, an educational reformer who valued students' joint moral and intellectual education.<sup>119</sup> Dewey advocated that there are many ways to know the world and that students should learn to *relate* to their environment via training that did not reduce knowledge to simply being an instrument. In the late nineteenth century, even the study of mathematics was 'valued less as a tool than as a provider of a healthy mental discipline.'<sup>120</sup> During the 1930-40s, however, a shift occurred in the department away from the discipline's origins in moral philosophy and towards a more instrumental view of knowledge. 'Knowledge was seen as a tool, not a state of being,' summarizes Hunter Heyck, "'facts" and "values" were thought to govern separate spheres.'<sup>121</sup> During this period, the curriculum shifted away from the study of political philosophy and towards the empirical study of human behaviour. By 1936, the Chicago Political Science Department became the recognized leader in its field.<sup>122</sup>

It was in this atmosphere that Simon later claimed to have evaluated his own relationship to empiricism. He opted to major in undergraduate political science and not

---

<sup>119</sup> My account of Simon's early life borrows from the work of Hunter Heyck in: Crowther-Heyck, *Herbert A. Simon*, 31–96; Simon, *Models of My Life*, 36–175; Heyck has also written extensively on the development on the modern social sciences: Crowther-Heyck, *Age of System*.

<sup>120</sup> Crowther-Heyck, *Herbert A. Simon*, 37. For more on this see: Roberts, 'Mathematics and Pedagogy'; On Dewey's pedagogy and the scientific method: Cowles, *The Scientific Method*.

<sup>121</sup> Crowther-Heyck, *Herbert A. Simon*, 38.

<sup>122</sup> Crowther-Heyck, 41.

economics, his first choice, because he disliked accounting.<sup>123</sup> He then ‘terminated’ his formal education in mathematics and pursued self-guided education to avoid having to attend accounting training.<sup>124</sup> He attended graduate classes and sought out practical experience to complement his studies. In 1938, while employed by the International City Managers’ Association, Simon learned of the status of classical administration theory.<sup>125</sup> In his 1996 autobiography, he claimed that it bothered him that this theory was built around common sense rather than empiricism, the latter of which to him carried a uniquely positivist bent.<sup>126</sup> Simon attributed the decision to build his own theoretical framework for administration theory to Chester I. Barnard’s 1938 book, *Functions of the Executive*, which explored human organizations in terms of decision making. He hung a photo of Barnard in his office.

Between 1938-43, while a PhD student in the Political Science Department at the University of Chicago, Simon set out to develop a theory of administration that was amenable to systematic observation and experimentation. In addition to Barnard, Simon studied the work of John R. Commons, a labour economist whose 1934 text *Institutional Economics* stressed that all transactions necessitated person-to-person interactions, the nature of which could not be explained fully using a model designed to explain individuals’ actions alone, such as the notion of a perfectly rational Economic Man.<sup>127</sup> *Institutional Economics* attempted to pivot the analytic focus of economics from a relation between man-and-nature to a relation between man-and-man. To explain the nuances of person-to-person interactions as they unfolded in practice, Commons developed the notion of ‘working rules,’ which were the frame of reference which every organization implicitly set for its membership. He wrote:

Working rules are continually changing in the history of an institution, including the state and all private associations, and they differ for different

---

<sup>123</sup> Simon, Transcript of Interview of Herbert Simon by Pamela McCorduck, January 01, 1975, 11.

<sup>124</sup> As cited in: Crowther-Heyck, *Herbert A. Simon*, 39.

<sup>125</sup> Classical theory asserted that orderliness in organizations was evidenced in a clear chain of command and division of labour. The first comprehensive text in public administration was White, *Introduction to the Study of Public Administration*; White was on Simon's degree committee. For analysis of the rise and decline of White's influence see: Storing, ‘Leonard D. White and the Study of Public Administration’. Another landmark volume was: Gulick, Urwick, and Pforzheimer, *Papers on the Science of Administration*. For a history of this volume and era of administration research, see: Urwick, ‘Papers in the Science of Administration’.

<sup>126</sup> Simon, *Models of My Life*, 72–73.

<sup>127</sup> Simon, *Models of My Life*, 74.

institutions. They are sometimes known as maxims of conduct. Adam Smith names them *canons* ...whatever their differences and different names, they have this similarity, that they indicate what individuals can, must, or may, do or not do.<sup>128</sup>

With help from Commons' account of group behaviour, Simon developed his own framework to explain the science of decision making in organizations. At the heart of his 1943 PhD thesis was the assumption that human thinking was inherently limited and thus susceptible to close empirical modelling and formalisation.<sup>129</sup> He argued that any two people given the same alternatives, values and knowledge would rationally reach the same decision. 'Administrative theory must be concerned with the limits of rationality,' he summarized, 'and in the manner in which organization affects these limits for the person making the decision.'<sup>130</sup> In effect, Simon sought to substitute *homo administrativus*, the Administrative Man, for *homo economicus*, the perfectly rational Economic Man, as the unit of agency in models of human decision making. This switch would allow him to deal empirically with the mess and minutia of decision making as it occurred *in reality*.

Simon's impression of human potential was influenced by Rudolf Carnap, under whose supervision he developed the basis of his dissertation.<sup>131</sup> Carnap was a central figure in the Vienna Circle of logical positivists, a community committed to the view that all meaningful philosophical inquiry can be reduced to and explained by logical analysis. Carnap built on the radical empiricism of Ernst Mach, as well as Bertrand Russell and Alfred North Whitehead's work on the foundations of mathematics. He rejected the distinction between natural and social science on the grounds that human behaviour and the behaviour of physical objects could both be explained with logical scrutiny and empirical observation. '[I] embraced a logical positivism that I have never relinquished,' Simon reflected in his 1996 autobiography, 'I would prefer to call it *empiricism* now.'<sup>132</sup>

---

<sup>128</sup> Commons, *Institutional Economics*, 71; Simon credits Commons with informing his work on bounded rationality, for which he won a Nobel Prize in 1978. See: Simon, *Models of My Life*, 87; On the genesis of bounded rationality see: Crowther-Heyck, *Herbert A. Simon*, 184–214.

<sup>129</sup> Boden, *Mind as Machine*, 319.

<sup>130</sup> Simon's PhD thesis as cited in: Simon, *Models of My Life*, 88.

<sup>131</sup> Crowther-Heyck, *Herbert A. Simon*, 70–72. His dissertation began as a final paper for Carnap's advanced research seminar on logic: Dick, 'After Math', 65.

<sup>132</sup> Emphasis his. Simon, *Models of My Life*, 44.



In 1947, Simon published *Administrative Behavior: A Study of Decision Making in Administrative Organization*, a scientific account of administrative theory built from his doctoral research.<sup>133</sup> Simon's account challenged the orthodoxy of Frederick W. Taylor's theory of scientific management, which relied on the hypothetical model of a purely rational Economic Man determined to maximize its self-interest at every intersection. Simon's theory introduced, in its place, 'Administrative Man,' a more refined version of Taylor's base-unit which pursued its self-interest only to the extent that it knew what its self-interest was. If Administrative Man could not identify what an optimal outcome was, it would settle for any outcome that aided, at that stage, in the ultimate search for an optimum result. *Administrative Behavior* thus posited that humans are rational up to a point, as Commons' working rules had implicitly suggested.

Simon framed the book as a catalogue of tools he had found effective in his own studies of the mechanics of public administration. His aim was twofold. First, he wanted to surface the contradictions latent in existing dogma and maxims of administration, such as 'diverging calls for specialization' and 'unified command.' Second, he intended to establish a new vocabulary of terms and concepts that would allow for the description of the immutable principles of administration. According to Simon, these universal principles were what his predecessors had aspired to but never articulated. For example, Simon distinguished between value judgements and factual judgements. The former decided an organization's final goals. The latter decided how those goals would be implemented.<sup>134</sup>

Given his ambitions, it is ironic that the system Simon developed in *Administrative Behavior* was not deeply rooted in empirical evidence.<sup>135</sup> Nor were its primary influences. Barnard's research, Simon acknowledged in 1996, was 'based on everyday observation, not on esoteric experimental or observational techniques.'<sup>136</sup> *Institutional Economics*, likewise, was based on the author's anecdotal observation of how collective actions influenced individual actions.<sup>137</sup> Reviewers noted the shortcomings of a speculative approach, which

---

<sup>133</sup> Simon, *Administrative Behavior*, 1997.

<sup>134</sup> Simon had developed these ideas in a 1944 paper, Simon, 'Decision-Making and Administrative Organization', 19.

<sup>135</sup> See: Sent, 'Book Reviews', 136; Simon, *Models of My Life*, 73. Simon later defended his early research by claiming that it lay a foundation for experimental work.

<sup>136</sup> Simon, *Models of My Life*, 73.

<sup>137</sup> Commons, *Institutional Economics*, 1.

remained abstract and positivist. A review of *Administrative Behavior* in the *American Journal of Nursing* lamented that the principles offered ‘very little help in making a choice’ for those operating on the ground.<sup>138</sup> Others noted his ‘foursquare’ reliance on logical positivism.<sup>139</sup> Another questioned, as did many, his ‘serious... misguided conviction that a true social science must eschew values.’<sup>140</sup> Simon’s positivist approach seemed to predetermine its connection to empirics, rather than being grounded in them.

All the same, in a post-war context in which American state administration grew rapidly, spurring derivative concerns like cost-benefit analysis, the book found a receptive audience.<sup>141</sup> The first edition sold thousands of copies per year after its release, and *Administrative Behavior* was soon considered ‘the most systematic field research on the behaviour of organizations’ by Simon’s contemporaries.<sup>142</sup> A 1956 review of the second edition described it as the ‘most notable book on the science of administration.’<sup>143</sup> In 2001, the book was voted the fifth most influential management book of the twentieth century.<sup>144</sup>

The primary technique from *Administrative Behavior* that merits our attention here is Simon’s notion of ‘decision premises,’ an adaptation of Commons’ ‘working rules.’ By this principle, those who sat higher up in the hierarchy of an organization could influence those below them by setting the parameters within which decisions were made by those lower down. Simon used the term as a catchall for the set of facts and values that served to shape the ‘decision-fabricating process,’ a process which encompassed fact-finding, design, analysis, reasoning, negotiation, intuition and even guessing.<sup>145</sup> It was this analytical framework that later helped him to model decision-making procedures in a computer.

In his 1996 autobiography, Simon acknowledged a fundamental flaw in his framework: it could not explain behaviour in human organizations that lacked a profit motive. ‘This

---

<sup>138</sup> Simon, ‘Administrative Behavior’, 1950, 46.

<sup>139</sup> Dahl, 245.

<sup>140</sup> Cook, ‘Review of Administrative Behavior’, 406. For comment on the book’s reception see: Kerr, ‘The Development History and Philosophical Sources of Herbert Simon’s *Administrative Behavior*’, 255; Subramaniam, ‘Fact and Value in Decision Making’.

<sup>141</sup> Torgerson, ‘Policy Analysis and Public Life: The Restoration of Phronesis?’, 240.

<sup>142</sup> Taylor’s *Principles* was voted first. For more see: Simon, *Models of My Life*, 88; Bedeian and Wren, ‘Most Influential Management Books of the 20th Century’, 222.

<sup>143</sup> Dahl, 244.

<sup>144</sup> Christie, Luce, and Macy Jr., ‘Communication and Learning in Task-Oriented Groups’, 7.

<sup>145</sup> Simon, *Administrative Behavior*, 1997, 23–24.

scheme needs to be modified somewhat to fit voluntary, religious, and governmental organizations,' he conceded.<sup>146</sup> *Administrative Behavior* was thus an *economic* model of administrative decision making, not a more general analysis of organizational phenomena. Simon circumvented the dynamics of moral agency by dealing with the instrumental rather than ethical dimensions of authority in a group. By adopting this frame, his model presumed a certain degree of moral apathy amongst its subjects. To quote Heyck, Simon's model of Administrative Man conceived of humans as 'perfectly malleable, perfectly docile.'<sup>147</sup> Simon claimed that individuals could only access the fullness of human rationality by participating in a group. 'The rational individual is, and must be, an organized and institutionalized individual.'<sup>148</sup> Each individual's perspective was skewed by innate cognitive shortcomings, such as subjective and limited access to information. Thus, to operate effectively, the individuals *had* to work together, which meant they had to acquiesce to the collective rationality and direction of their group. For Simon and his followers, this model of decision making explained the rationalized incentives and logic behind administrative action.

#### 1952-55: Simon's Account of Administration is Translated into Code at RAND

The United States' zero-sum standoff with the Soviet Union in the mid-twentieth century produced a new sort of warfare, one with far reaching implications for twentieth century science.<sup>149</sup> In the pivotal years that concern us here, cold warfare had to be both invented and waged. The existence of a continuous threat, often disembodied or felt via proxies (ex. the Korean War), produced a climate of anxiety that impressed upon American military leaders the need to consider an exponential number of possible contingencies. This anxiety compelled the study of new techniques like game theory and systems analysis to make statistical sense of an otherwise ambiguous risk.<sup>150</sup> The crisis, which quickly became epistemic as well as existential, also legitimized the funding and manufacturing of new

---

<sup>146</sup> Simon, 15.

<sup>147</sup> Crowther-Heyck, *Herbert A. Simon*, 184.

<sup>148</sup> Simon, *Administrative Behavior*, 1997, 111.

<sup>149</sup> For more on this see: Erickson, *How Reason Almost Lost Its Mind*; Amadae, 'Computable Rationality, NUTS, and the Nuclear Leviathan'.

<sup>150</sup> On the rise and decline of applying systems thinking to complex phenomena in America after World War II see: Hughes and Hughes, *Systems, Experts, and Computers*.

experimental technologies, such as the electronic stored-program computer. As the stand-off deepened, the military's involvement in civilian science, an entanglement which had initially been brokered out of desperation during World War II, became a normal part of academic life for many scientists. In this section, I will examine how these forces converged at the RAND Corporation in the early 1950s to facilitate the design and construction of the Logic Theory Machine.

The thrill of victory in 1945 brought newfound excitement around scientific and technological research in the United States. Convinced, as were many U.S. military and government leaders, that these areas of study had been vital to the nation's wartime victory, General 'Hap' Arnold, commanding officer of the Army Air Forces, wrote to the Secretary of War to petition for an extension of funding to these projects outside of active war. 'The conclusion is inescapable that we have not yet established the balance necessary to insure the continuance of teamwork among the military, other government agencies, industry, and the universities,' Arnold reasoned.<sup>151</sup> Within a year, the Navy had established the Office of Naval Research (ONR), 'the only government agency able to maintain connections between academia, industry, and the military services in the pursuit of new science and technology.'<sup>152</sup> In response, the Air Force established Project RAND, a think tank in Santa Monica designed to make the Air Force's 'operation, its organization, its present and future weapons, even its place in American society the subject of a new research domain, a distinct category of scientific inquiry.'<sup>153</sup>

As demonstrated by the shift in curriculum Simon experienced at Chicago, this trend towards the study of human organizations as scientific systems was not solely a military concern. Between 1939 and 1972, funding for Behavioral and Social Science research in the United States grew from approximately twenty million to over six hundred million dollars.<sup>154</sup> American social science entered an era of what Heyck refers to as 'high modernism' in which human organizations were studied as complex, hierarchic systems defined by their structure more than by their components.<sup>155</sup> At Harvard, Princeton and Berkeley in the 1940-50s, the

---

<sup>151</sup> Hounshell, 'The Cold War, RAND, and the Generation of Knowledge, 1946-1962', 241.

<sup>152</sup> Babb, 'The Genesis of the Office of Naval Research'.

<sup>153</sup> Collins, 'Planning for Modern War: Rand and the Air Force, 1945-1950', 5.

<sup>154</sup> Crowther-Heyck, *Age of System*, 61.

<sup>155</sup> Crowther-Heyck, 18-50, 214-15; Scott, 'Authoritarian High Modernism', 87-102.

undergraduate curriculum was reformulated to emphasize statistical thinking and ‘the logical framework of science’ rather than its ‘factual content.’<sup>156</sup> A 2015 survey of 1800+ journal articles published in the social sciences between 1930-70 found that the number of articles which employed the ‘core concepts and methods of high modern social science (system, structure, function, modeling)’ increased from 7% to 60%, an eightfold increase over forty years.<sup>157</sup> As with the positivist paradigm that Simon adopted in *Administrative Behavior*, the promise of high modernism was that natural and social phenomena could be catalogued—and then optimized—if modeled correctly.

As the needs of warfare changed, new theory was required to keep pace. In active ‘hot’ warfare, it was clear when efficiencies had been gained: casualties were reduced, or munitions were saved. In the ever-dynamic haze of the emergent Cold War, however, the quantification of efficiency became, at times, intractable; there were too many variables to account for. Whereas active military engagement simplified war down to a scorecard, the scope of potential outcomes during the Cold War skyrocketed. To confront this new reality, RAND researchers pioneered the field of systems analysis.<sup>158</sup> A 1956 RAND report defined systems analysis as, ‘A systematic examination of a problem of choice in which each step of the analysis is made explicit wherever possible.’<sup>159</sup> This definition was prefaced by an apology that no *precise* definition of systems analysis techniques yet existed. The only definitive hallmark of systems analysis was a reliance on heterogeneous professionals and heterogeneous organizations operating as components in a system.<sup>160</sup> The consequence, to cite Paul Edwards, was ‘A whole way of thinking: a systems philosophy of military strategy.’<sup>161</sup>

---

<sup>156</sup> Bode et al., ‘The Education of a Scientific Generalist’, 553; Fortun and Schweber, ‘Scientists and the Legacy of World War II’, 608.

<sup>157</sup> Crowther-Heyck, *Age of System*, 2.

<sup>158</sup> Within this milieu, other RAND researchers developed linear programming (George B. Dantzig), network-flow analysis (Lester R. Ford, Raymond Fulkerson), dynamic programming (Richard Bellman), Monte Carlo simulation (Herman Kahn), and improvements on game theory (Lloyd S. Shapley). See: Ware, *RAND and the Information Evolution*, 138–39.

<sup>159</sup> Hoag, ‘An Introduction to Systems Analysis’, 1.

<sup>160</sup> Fortun and Schweber, ‘Scientists and the Legacy of World War II: The Case of Operations Research (OR)’, 607. Amadae suggests that game theoretic rationality was another hallmark of systems analysis in: Amadae, ‘Computable Rationality, NUTS, and the Nuclear Leviathan’, 4–5; see also: Isaacs, *Differential Games*, vii.

<sup>161</sup> Edwards, *The Closed World*, 116. Scoblic chronicles how attempts to qualify measures of uncertainty at RAND, via Frank Knight then Herman Kahn, also legitimized that community’s

The style of interdisciplinary collaboration operationalized in systems analysis was not just a theoretical ideal; at RAND, it was also cemented into the group's physical headquarters. 'RAND represents an attempt to exploit mixed teams, and that [sic] to the extent its facility can promote this effort it should do so,' wrote John D. Williams, the building's designer, in a 1950 report. 'This implies that it should be easy and painless to get from one point to another in the building; it should even promote chance meetings of people.'<sup>162</sup> Of course, this free mobility was not afforded to all employees: RAND's physical structure separated research staff from non-research staff. Williams deemed it unfeasible to have RAND's headquarters arranged so that 'Elaine in Electronics and Ethel in Publications have optimum physical communications.'<sup>163</sup> He added, 'Nor is it especially useful that they have it.' In Williams' view, it was sufficient to prioritize the optimal arrangement of only the research staff.

The culture at RAND was also crafted to generate interdisciplinarity. Over its first fifteen years, RAND sold itself to academics as 'A university without students.'<sup>164</sup> The Santa Monica based think tank adopted the structure of a conventional university complete with departments of Engineering, Economics, Mathematics, Physics, Psychology, Chemistry and Aerodynamics. But RAND was not a university. For starters, all employees were hired on the basis of a top-secret Department of Defense security clearance. They were also required to complete a Primary Mental Abilities test, which checked for seven building block traits of 'intelligence' (ex. inductive reasoning, verbal comprehension) but entirely omitted measures of emotional, empathetic and social ability.<sup>165</sup> This dual requirement guaranteed, in the words of one employee, 'That the recipient is judged... to be personally and politically reliable, emotionally stable, patriotic and not subject to blackmail, and that he is trusted with classified military secrets.'<sup>166</sup> In short, RAND staff were groomed to be reliably compliant. Employees

---

view that warfare was a science, not an art. See: Scoblic, 'The Postwar Development of Tools to Mitigate Uncertainty', 2. On the study of uncertainty at RAND see: Erickson, *The World the Game Theorists Made*.

<sup>162</sup> Williams, 'Comments on RAND Building Program'.

<sup>163</sup> Williams.

<sup>164</sup> Hounshell, 'The Cold War, RAND, and the Generation of Knowledge, 1946-1962', 242.

<sup>165</sup> On RAND's use of the PMA testing see: Armer, 'Problems in Administrative Leadership in the Computer Field', 94-96. For more on what the PMA tested for, see: Bendig and Meyer, 'A Longitudinal Study of the Primary Mental Abilities Test'; on its methodological errors see: Eysenck, *A Model for Intelligence*, 18.

<sup>166</sup> Boulding, *Peace and the War Industry*, 194.

who lacked such clearance worked in a separate facility next-door.<sup>167</sup> Within RAND, those who worked as computer programmers were described in an internal memo as uniquely prone to conformity: 'Programmers, as a class, seem to me to be non-leaders... the programmer usually has his goals set for him by someone else. The problem is handed to him, and he is told, "Here, do this."' <sup>168</sup> This group, in the numerical analysis department, were located in the basement.

RAND sought to replicate and optimize the knowledge-producing capacity of the traditional university. But that's as far as the similarities went; if RAND were a university, it was an unusual and ideological one. Unlike non-military American universities, RAND's annual budget was ten-million dollars a year, which allowed for a pay rate fifty percent higher than at equivalent government or academic positions.<sup>169</sup> In the early years, their researchers flew first class.<sup>170</sup> Moreover, embedded within the workplace culture was a pronounced militaristic ideology. For some researchers, 'RAND's mission was nothing short of the salvation of the human race.'<sup>171</sup> In 1947, the corporation's sixth quarterly report read, 'If modern weapons have wiped out the sharp distinction between the military and civilian in time of war, so in time of peace such a differentiation has become outdated. RAND is in line with this development.'<sup>172</sup>

This brief introduction to RAND's purpose, structure and culture sets the stage for Simon's arrival in Santa Monica in 1952. In the intervening years since leaving Chicago, Simon moved from the Illinois Institute of Technology to Carnegie Technical Institute to help establish their new Graduate School of Industrial Administration in 1949, where he also served as a professor of business administration.<sup>173</sup> About the pedagogical orientation at

---

<sup>167</sup> Collins, 'Planning for Modern War: Rand and the Air Force, 1945-1950', 1.

<sup>168</sup> Armer, 'Problems in Administrative Leadership in the Computer Field', 97.

<sup>169</sup> Between 80-600 million USD per year today according to: Officer and Williamson, 'Conversion', 2015; Edwards, *The Closed World*, 115.

<sup>170</sup> *Building Computers in 1953: The Johnniac*.

<sup>171</sup> Hounshell, 'The Cold War, RAND, and the Generation of Knowledge, 1946-1962', 243.

<sup>172</sup> Collins, 'Planning for Modern War: Rand and the Air Force, 1945-1950', 1.

<sup>173</sup> Simon read widely in cybernetic literature in this period. See: Simon, *Models of My Life*, 114. Ashby's work in particular firmed his conviction that the world was a system to be understood functionally, such that the functions of organizations and organisms could be seen as equivalent. In letters between them, Ashby equated the principles that governed the organization of nerves in the brain with those that governed workers in a factory. See: Crowther-Heyck, *Herbert A. Simon*, 185.

Carnegie, Simon later recalled, 'We perceived American business education at that time as a wasteland of vocationalism that needed to be transformed into science based professionalism, as medicine and engineering have been transformed a generation or two earlier.'<sup>174</sup> This directive is echoed in a proposal sent by Simon to G.L. Bach, Dean of the Graduate School of Industrial Administration, in February 1952. The proposal was for a five-year study to explore a perceived 'intimate connection' between 'organizational structure and the learning of frames of reference and roles by members of organizations.'<sup>175</sup> The proposal, as was explicitly stated, was an extension of the theoretical framework put forward in *Administrative Behavior*.

Simon did not mention electronic computing in his 1952 Carnegie proposal. His proposed 'Program of Activity' listed field studies, interviewing techniques, observation, theory development and laboratory experimentation—but no study using computers. In contrast, Simon's paper, 'A Behavioral Model of Rational Choice,' which he wrote that same year while serving as a consultant at RAND, speculated on how computers could be used to inform his research.<sup>176</sup> One explanation for this omission was that, at the time, digital electronic stored-program computing was a prohibitively expensive enterprise. Carnegie Tech did not install its first IBM computer until 1956. Another explanation is that Simon had co-authored his proposal with Harold Guetzkow, a psychologist and friend who had introduced him to the work of major figures in psychology such as Jean Piaget, Max Wertheimer and Karl Duncker. Simon himself had had no formal training in psychology.<sup>177</sup> Reflecting Guetzkow's background, their 'Program of Activity' offered means of addressing the structures and methods of organizations that were conventional to psychology, which digital computing was not at the time.

The RAND Corporation, as opposed to Carnegie Tech, owned and operated one of the largest installations of scientific computing in the world by the early 1950s.<sup>178</sup> In addition to the various machines that the corporation rented from IBM, RAND took the pioneering step to greenlight the construction of the custom-built JOHNNIAC computer in 1950, one of only

---

<sup>174</sup> Simon, *Models of My Life*, 105.

<sup>175</sup> Simon and Guetzkow, 'Memorandum', 1.

<sup>176</sup> Simon, 'A Behavioral Model of Rational Choice', 114.

<sup>177</sup> Simon, Transcript of Interview of Herbert Simon by Pamela McCorduck, January 01, 1975, 11.

<sup>178</sup> Gruenberger, 'The History of the JOHNNIAC', 50.



a handful of general-use electronic computers in operation worldwide. In Santa Monica, Simon gained access to these devices. In 1952, he was hired to consult in RAND's new Systems Research Laboratory (SRL), one of the first experimental laboratories in management science, which was founded to explicitly study 'the problems and prospects of the "man-machine interface."' <sup>179</sup>

It was at the Systems Research Laboratory, housed behind a Santa Monica billiard hall, that Simon encountered Allen Newell and J. Clifford Shaw, with whom he would later develop the Logic Theory Machine.<sup>180</sup> In 1952, together with Newell, he tested a technique called 'protocol analysis,' which formalized all aspects of communication and decision making amongst a group of college students hired to simulate a military crisis in a replica of an Air Defense Direction Center in Tacoma, Washington.<sup>181</sup> The students were brought in to read out radar data then analyze it to determine whether or not to scramble fighter jets.<sup>182</sup> Since even a slight misjudgement could in principle allow for a successful nuclear attack, each step in the tense decision process was seen as critical. By analysing data like tape-recordings of communication between radar operators, air controllers and crew members, Simon and his colleagues sought to improve efficiency and reduce stress through the creation of new training methods.<sup>183</sup> It was during this research program that these men began to think about how machines might be able to fulfill human-based tasks. In effect, the Systems Research Laboratory enabled Simon to pilot his notion of 'decision premises' in a cutting-edge man-machine working environment, an experience that brought him one step closer to modeling his early career research on an electronic computer.

Newell, at age twenty-four, shared Simon's enthusiasm for the study of decision making. When he first joined RAND in 1950, following a physics degree in 1949 at Stanford University and an aborted year of graduate study at Princeton, Newell worked on the application of game theory and operations research techniques to the study of administrative organizations. At the Systems Research Laboratory, he and Simon bonded over a mutual

---

<sup>179</sup> Mirowski, *Machine Dreams*, 2002, 350.

<sup>180</sup> Simon had been hired as a consultant at RAND but retained his position at Carnegie under a grant from the Ford Foundation. Simon, Newell, and Shaw, 'Untitled Memo'.

<sup>181</sup> For a detailed look at this experimentation and the SRL, see: Klein, 'Implementation Rationality', 209.

<sup>182</sup> Boden, *Mind as Machine*, 322.

<sup>183</sup> McCorduck, *Machines Who Think*, 1979, 139–40.

appreciation of the work of George Pólya, a Stanford mathematician who had written extensively about the use of *heuristics*, or rules of thumb, in mathematics.<sup>184</sup> In 1983, Newell joked that, as an undergraduate physicist at Stanford, he had ‘majored in Pólya,’ because he bought all of his books and attended multiple classes he had offered.<sup>185</sup> Like systems analysis, heuristic techniques leveraged a toolset that was *in statu nascendi*, meaning in the process of being invented. In 1944’s *How to Solve It*, Pólya set about teaching new students like Newell how mathematicians *actually* went about solving a complex problem, rather than what the rule books said. *How to Solve It* outlined straightforward but effective strategies like ‘Work Backwards’ and ‘Solve a Simpler Problem’ when a more difficult problem evaded solution. At the root of this approach was an important analytical distinction: the difference between a rule (or algorithm) and a heuristic (a rule of thumb). This distinction sensitized Simon and Newell to the manner in which humans actually solved mathematical problems, a set of pragmatic behaviours that they later imagined could be simulated by a machine.

In an interview conducted decades later, Newell claimed to have experienced an epiphany at RAND in 1954 on how computers could be used to manipulate non-numerical information. He realized that complex behaviour could be generated from simpler sub processes using a computer. Newell characterised the breakthrough as ‘a case of the prepared mind.’<sup>186</sup> For reasons I have alluded to, this was only half the story. His insight was also the product of a prepared corporation.<sup>187</sup> In keeping with its model as a military-funded pseudo-university, RAND organized frequent presentations for their employees by top academics from across the country. During one such seminar, hosted by MIT mathematicians G. P. Dinneen and Oliver Selfridge, Newell learned about that group’s computer-based techniques to recognize visual patterns.<sup>188</sup> Their techniques seeded Newell’s realization that digital computers could be used to manipulate symbols and not just numbers.

Determined as ever to model complex behaviour for himself, Newell set to work on the ‘Chess Learning Machine 1’ with J. Clifford Shaw, an insurance actuary turned RAND

---

<sup>184</sup> Pólya, *How to Solve It*.

<sup>185</sup> Newell, ‘The Heuristic of George Pólya and Its Relation to Artificial Intelligence’, 196.

<sup>186</sup> McCorduck, *Machines Who Think*, 2004, 157.

<sup>187</sup> On the standardization of management techniques and international communications practices in corporations between 1870-1920 see: Yates, *Control Through Communication*.

<sup>188</sup> Dinneen, ‘Programming Pattern Recognition’; Selfridge, ‘Pattern Recognition and Modern Computers’; McCorduck, *Machines Who Think*, 1979, 133.

programmer who had designed radar displays for the Systems Research Laboratory.<sup>189</sup> Their challenge was to translate the complexity of chess-strategy into a language that was intelligible to the computer. 'This is the first attempt to define a language that the machine can use to talk about Chess,' their 1955 report read.<sup>190</sup> To accomplish the task, they patterned a 'language' on propositional calculus, using operators such as 'there exists an x such that' that allowed for the formalization of chess strategy.

The Chess Learning Machine project required Newell and Shaw to examine the nature and limits of human language when rendered in an inorganic entity. Since their goal was to produce a decision system for a *machine* to process and use, traditional biological restrictions no longer applied. Success in this enterprise, they discovered, required the careful manipulation of a computer's capacity for sophisticated expressions but never so much that it would alienate a human operator from understanding the procedures that had occurred. Their symbolic language had to satisfy both the operator's need to grasp each computational step and the material affordances of the JOHNNIAC computer at RAND. The needs of these two entities were not the same. For instance, the JOHNNIAC used sequential syntax, which read expressions in propositional calculus in one direction (such as left-to-right) as opposed to in either direction, like a human operator could comfortably manage. The resulting Chess 'language' Newell and Shaw produced was thus a composite of the idiosyncratic capacities of a machine and those of its human operator.

In November 1955, Newell presented his chess research at the Western Joint Computer Conference. Despite having made some headway, the problem of how to make a machine 'play good chess' remained unsolved. 'One can estimate the man-hours necessary to draw up the detailed flow diagrams from which machine coding follows,' Newell griped, 'These mechanisms are so complicated that it is impossible to predict whether they will work.'<sup>191</sup> Newell reasoned that a new methodology would be necessary to compress the level of complexity involved in such systems down to a workable scale. The game of chess was an 'ultra-complicated problem' for both men and machines because it involved complex information processing, which required new techniques to parse. One candidate solution he

---

<sup>189</sup> Skinner, *Building the Second Mind*, 89.

<sup>190</sup> Newell, 'Notes on Chess Learning Machine I', 1.

<sup>191</sup> Newell, 'The Chess Machine', 101.

mentioned was to equip the machine with the ability to manage this complexity on its own, such as by using Pólya's heuristics to have a device 'pull itself up by its bootstraps.'<sup>192</sup>

In the fall of 1955, Newell visited Simon at Carnegie Tech, still determined to model human problem solving on a computer. Simon's years spent formalizing the nature of decision making in administrative organizations inclined him to the challenge. To proceed, as with the Chess program, the men required a coherent schema or 'language' with which to structure their desired computational system. In experimental notes from January 1955 entitled 'Logic Theory Machine I,' Newell wrote, 'A certain language is necessary... so we could 'talk' to the machine.'<sup>193</sup> After failing to model both geometry problems and chess, a candidate solution came to them in the form of Russell and Whitehead's *Principia Mathematica*—a copy of which Simon claimed to have serendipitously 'pulled off the shelf.'<sup>194</sup>

As with Newell's encounter with Selfridge, it would be disingenuous to characterise this development as a happy accident or one that occurred quite as serendipitously as Simon himself believed. Simon had, of course, studied *Principia Mathematica* closely in Chicago under Carnap. He had based his work on decision theory within the highly formalised system used in the book, claiming in retrospect that it was Russell and Whitehead's text that scaffolded *Administrative Behavior*, providing him with a 'framework for my thinking about administrative decision making.'<sup>195</sup> In some sense, Simon had *already* simulated the logical ethos of *Principia Mathematica* in an embodied entity: the organization. He had used early twentieth century British symbolic logic to structure procedures for how to control the actions of a complex decision system.

In 1955, Simon ventured a parallel transmutation, this time to embody formal logic in an electronic digital computer with Newell and Shaw. Stephanie Dick credits *Principia Mathematica* as a critical resource for the RAND group because it made 'explicit the basic and primitive rules according to which deductive reasoning proceeded.'<sup>196</sup> In order to render

---

<sup>192</sup> Newell, 108.

<sup>193</sup> As cited in: Dick, 'After Math', 78.

<sup>194</sup> The duo assessed how to solve geometric proofs using a computer but found the description of diagrams too difficult to render in code. McCorduck, *Machines Who Think*, 2004, 161; Skinner, *Building the Second Mind*, 93–94; Simon, *Models of My Life*, 205.

<sup>195</sup> Simon, *Models of My Life*, 193. In another paper from 1944, also on decision making in administrative organizations, Simon equated rationality with logical inference. Simon, 'Decision-Making and Administrative Organization', 19.

<sup>196</sup> Dick, 'Of Models and Machines', 626.

these rules in code, however, Simon and his interlocutors had to first limit and structure the search space within which such rules would operate. This is what Simon's theory of decision premises allowed. In an administrative setting, the behaviour of a rational person could only be controlled, Simon posited, 'if the value and factual premises upon which he bases his decisions are specified for him.'<sup>197</sup> These premises could now be specified in the medium of code.

Simon, Newell, and Shaw explored how the step-by-step behaviours of a rational machine could be controlled by structuring a computer's internal decision systems as a flexible, yet ultimately closed, environment. Although they did not recognize themselves as doing so, to make a machine 'think' like a human being, they made it 'think' like one of Simon's abstracted 'organizations,' be it a model fire department, government agency, or corporation. Although they did not formally acknowledge that this is what they intended, the effect of their design was to orphan the simulation of human problem solving from the sort of complex moral decision making that humans face, but that Simon's framework could not account for. Once technically and conceptually separated from these concerns, the factual premises upon which their closed system based its decisions could be specified and operationalized. It was this process of specification that the system and notation put forward in *Principia Mathematica* allowed.

#### 1955-56: Building the Logic Theory Machine

The elementary rules outlined in *Principia Mathematica* underwent heavy revision to become instantiated in the Logic Theory Machine.<sup>198</sup> To fully convey this process of transmutation from paper to code, it is useful to run through the Machine's operation step by step. The Logic Theory Machine was a *virtual* machine, which contemporary readers might think of as a proto computer program. It was developed with the express aim to discover and construct proofs for theorems in propositional calculus. Newell, Simon and Shaw sought to demonstrate not only that a computer could be used to prove theorems, but also that it could

---

<sup>197</sup> Simon, 'Decision-Making and Administrative Organization', 19.

<sup>198</sup> For a detailed study of this transmutation see: Dick, 'Of Models and Machines'; Dick, 'After Math'.

do so in a manner ‘as similar as possible to that of human beings.’<sup>199</sup> Their adaptation of Russell and Whitehead’s system thus equated human problem solving with formal logic.

To solve a theorem using the Logic Theory Machine required three initial steps. First, an interpreted version of *Principia Mathematica*’s primitive propositions and definitions were loaded into JOHNNIAC’s finite memory. These included the following five axioms:

1.2	$p \vee p$	$\rightarrow$	$p$	$(p \text{ or } p) \text{ implies } p$
1.3	$p$	$\rightarrow$	$q \vee p$	$p \text{ implies } (q \text{ or } p)$
1.4	$p \vee q$	$\rightarrow$	$q \vee p$	$(p \text{ or } q) \text{ implies } (q \text{ or } p)$
1.5	$p \vee q \vee r$	$\rightarrow$	$q \vee p \vee r$	$[p \text{ or } (q \text{ or } r)] \text{ implies } (q \text{ or } (p \text{ or } r))$
1.6	$p \rightarrow q$	$\rightarrow$	$r \vee p \rightarrow r \vee q$	$(p \text{ implies } q) \text{ implies } [(r \text{ or } p) \text{ implies } (r \text{ or } q)]$

From these axioms, other true expressions could be derived as theorems. Following this step, various programs would be loaded into memory to perform operations on new, candidate theorems consistent with the rules of inference admitted in Russell and Whitehead’s text. Not every rule of inference was included, just the rules for substitution, detachment and replacement.<sup>200</sup> The last ingredient was a program to discover or construct proofs. For this task, the Logic Theory Machine called various ‘methods’ into action to accomplish the following:

- Substitute new variables for the variables in other true theorems, such as from the five axioms listed above.
- Follow the inference rules set out in *Principia Mathematica*, also mentioned above.
- Construct chains of theorems out of transitive syllogisms, such as the expression ‘A implies D’ from the product of ‘A implies B’, ‘B implies C,’ and ‘C implies D.’<sup>201</sup>

In an internal RAND report from 1956, Newell and Simon made explicit their debt to Russell and Whitehead. ‘These three methods correspond, in fact, to the procedures for constructing proof-chains that are legitimate in the system of *Principia*,’ they wrote, ‘Only [the first] two methods... are stated formally in *Principia* but [the third] can be shown to be a

---

<sup>199</sup> Newell and Simon, ‘Abstract: The Logic Theory Machine’, 1.

<sup>200</sup> Newell and Simon, ‘The Logic Theory Machine: A Complex Information Processing System’, 15 June 1956, 26.

<sup>201</sup> Newell and Simon, ‘Abstract: The Logic Theory Machine’, 1.

special case of [the second].'<sup>202</sup> In later publications, they did not outline their conceptual debts to the text as explicitly.<sup>203</sup> Russell and Whitehead, in comparison, were careful to situate their system within the intellectual genealogy from whence it came. They recognized that this step was necessary because the system of notation they had introduced obscured how much they had borrowed from the work of past contributors. 'Detailed acknowledgments of obligations to previous writers have not very often been possible, as we have had to transform whatever we have borrowed in order to adapt it to our system and our notation,' they wrote.<sup>204</sup> This lineage was re-concealed by Simon and Newell's transmutation of logic into code and their decision not to adequately signpost the change through sustained citation.

Also obscured by this process of translation was their subtle but important application of Simon's administrative logics. Unlike formal logics, which existed only as abstractions, the Logic Theory Machine operated in the physical world, and was thus subject to the passage of time. This constraint made efficient procedural operations a crucial priority for the system's designers. The machine's decision premises had to be adequately contained to avoid the possibility of endless search. The central difference between the Logic Theory Machine and a hypothetical algorithmic system that simply generated valid deductive sequences was that the Machine avoided trying to scan the *entire* available search space for a possible solution. Doing so would cap out the machine's limited memory or doom it to test methods *ad infinitum* if a solution did not exist. To avoid this trap, the Machine was designed to, 'Move in the right direction,' according to experimental notes.<sup>205</sup>

This ill-defined constraint, to move in the 'right' direction,' mirrored the basic orientation of Simon's administrative logics. 'Right' in this context did not mean toward the type of sometimes irrational decision-making procedures witnessed in acts of charity or faith. It meant, instead, to move toward one of two highly rationalized procedures, the first heuristic and the second mathematical in nature. The first was means-ends analysis, a sociological concept that Simon had called upon frequently to describe the logic of an organization in *Administrative Behavior*. Means-ends analysis compared an initial state to a

---

<sup>202</sup> Newell and Simon, 1.

<sup>203</sup> Newell and Simon, 'The Logic Theory Machine: A Complex Information Processing System', September 1956; Newell and Shaw, 'Programming the Logic Theory Machine'; Newell, Shaw, and Simon, 'Empirical Explorations of the Logic Theory Machine'.

<sup>204</sup> Russell and Whitehead, *Principia Mathematica Volume I*, viii.

<sup>205</sup> Newell and Simon, 'Abstract: The Logic Theory Machine', 1.

goal state. This technique could be used to appraise the desirability of one option versus another. In 1947, Simon wrote:

In the process of decision those alternatives are chosen which are considered to be appropriate means for reaching desired ends. Ends themselves, however, are often merely instrumental to more final objectives. We are thus led to the conception of a series, or hierarchy, of ends. Rationality has to do with the construction of means-ends chains of this kind.<sup>206</sup>

In *Administrative Behavior*, Simon cited *The Structure of Social Action* by the cybernetically inclined sociologist Talcott Parsons for introducing him to means-ends analysis.<sup>207</sup> Seven years later, at RAND in 1954, Simon adapted the technique for use in stored-program electronic computing. Use of means-ends analysis reduced the number of required steps in a search space to a fraction of its previous size. Once employed in practice, for example, the group found that no theorem from Chapter Two of *Principia Mathematica* required more than four steps to prove.<sup>208</sup>

The second technique used by the Logic Theory Machine to decide how to evaluate a new theorem was factoring. Factoring, from mathematics, involved breaking one large problem down into several smaller problems. If a solution to any of these sub-problems could be found, this heightened the chance of solving the initial, larger problem. To consider the effectiveness of this technique, consider a safecracker. If faced with a safe equipped with ten dials, each running from zero to ninety-nine, a safecracker would have to test one hundred billion billion possible dial settings to guarantee a result.<sup>209</sup> If, however, a heuristic could be discovered, such as an audible 'click' each time the dial moved past its correct setting, then the safecracker would only need to attempt five hundred steps to unlock the correct combination (fifty for each of the ten dials). Thus, with one simple clue, the possible search space of that challenge collapsed from  $10^{20}$  possible solutions to 500. The Logic Theory Machine was designed to automate this technique.

A summary memo written by the RAND trio about their project divided research into three areas: psychology of mental processes, heuristic programming, and information

---

<sup>206</sup> See: Simon, *Administrative Behavior*, 1997, 73.

<sup>207</sup> Parsons, *The Structure of Social Action*.

<sup>208</sup> If it could be proved using the Logic Theory Machine in 1955-56. Some theorems could not. Newell, Shaw, and Simon, 'Empirical Explorations of the Logic Theory Machine', 228.

<sup>209</sup> Simon, *Models of Thought*, 151.



processing languages. The categories reflected a rough division of labour between the group. Simon, often dialling in from Carnegie, theorized about higher mental processes in collaboration with Newell who, on the ground at RAND, worked with both Simon and Shaw to structure a suitable heuristic program. Shaw, based out of the low-ceilinged 'sweatshop' of RAND's computer heavy Numerical Analysis Department, was responsible for rendering an information processing language on the JOHNNIAC device.<sup>210</sup>

A moment's pause is needed to clarify how the first of these categories—the psychology of mental processes—reflexively shaped the latter two: heuristic programming and the creation of an information processing language. The memo outlined the group's plan to develop a program that 'effectively simulate[d] the behavior of a thinking adult.'<sup>211</sup> Use of the term 'effectively' belied their simulation's limited fidelity to that which they called 'thinking.' The group deemed heuristic programming and information processing languages sufficient in this respect. That is, I argue, because the psychology of mental processes they had adopted to shape the Logic Theory Machine—the *decision premises* they imposed (to borrow Simon's earlier language for this sort of move) to set the boundaries of what that entity could 'think'—had already sacrificed significant depth and complexity. Put precisely, means-ends analysis and factoring worked as heuristic methods because the environment they manipulated was closed, not open—as it was in physical reality. Factoring could only effectively reduce a large search space, not an infinite one. Cognition had to be rendered as a problem space prior to there being a conceivable safe to ingeniously crack with mathematics. Each narrowing offers a glimpse into the motivating assumption underlying the RAND group's work. It was not just that they assumed that human mental life was coherent in their experimentation, but also that they assumed that it was provably *restricted*.

The study of bounded rationality characterised a later portion of Simon's career. At this stage, his thinking in this respect was largely informed by his research on administrative theory. That domain, as discussed, cohered for him the virtues of imposing a limitation on human behaviour. His efforts to formalize an administrative science dealt pragmatically with the optimal management of resources across domains. 'Administrative theory must be interested in the factors that will determine with what skills, values, and knowledge the

---

<sup>210</sup> As quoted in: Dick, 'Of Models and Machines', 629.

<sup>211</sup> Simon, Newell, and Shaw, 'Untitled Memo', 1.

organization member undertakes his work,' he argued in *Administrative Behavior*, 'These are the "limits" to rationality with which the principles of administration must deal.'<sup>212</sup>

The material affordances of the JOHNNIAC device introduced an analogous economy of resource management to the development of the Logic Theory Machine. 'It is clear that each additional test, search, description, and the like, has its costs in computing effort as well as its gains in performance,' the men explained at the 1957 Western Joint Computer Conference, 'The costs must always be balanced against the performance gains, since there are always alternative heuristics which could be added to the system in place of those being used.'<sup>213</sup>

Insofar as the Logic Theory Machine was a manifestation of Simon's 'Administrative Man', as I argue it was, it was also, by his own reasoning, an instance of 'Economic Man.' Conjoining the two were principled notions of efficient resource allocation that Simon, in his early work, treated as foundational to both. In a section from *Administrative Behavior* entitled 'The Diagnosis of Administrative Situations,' republished as 'The Proverbs of Administration,' his 1946 article for *Public Administration Review*, he wrote, in a manner worth quoting in full:

The theory of administration is concerned with how an organization should be constructed and operated in order to accomplish its work efficiently. A fundamental principle of administration, which follows almost immediately from the rational character of "good" administration, is that among several alternatives involving the same expenditure that one should always be selected which leads to the greatest accomplishment of administrative objectives; and among several alternatives that lead to the same accomplishment that one should be selected which involves the least expenditure. Since this "principle of efficiency" is characteristic of any activity that attempts rationally to maximize the attainment of certain ends with the use of scarce means, it is as characteristic of economic theory as it is of administrative theory. The "administrative man" takes his place alongside the classical "economic man."

Actually, the "principle" of efficiency should be considered a definition rather than a principle: it is a definition of what is meant by "good" or "correct" administrative behavior. It does not tell how accomplishments are to be maximized, but merely states that this maximization is the aim of administrative activity, and that administrative theory must disclose under what conditions the maximization takes place.

Before turning to the next subject, slightly more can be said about the resemblance I am point to between Simon's administrative theory, summarized above, and his theory of

---

<sup>212</sup> Simon, *Administrative Behavior*, 1997, 47.d

<sup>213</sup> Newell, Shaw, and Simon, 'Empirical Explorations of the Logic Theory Machine', 226.

complex information processing, epitomized in the Logic Theory Machine. In the 1950s, Simon, Newell, and Shaw took pains to define the terms of their analysis. By 'complex' information processing they meant a class of system characterized by a large number of different processes. These processes accomplished similar functions across different contexts, despite their being highly contingent on both internal and environmental outcomes. Neither a computer nor an algorithm was complex because neither operated with a high number of variable conditions. This style of componentry was complicated, certainly, but not complex. A hallmark of a complex information processing system was that it could process the dynamism inherent to complexity even within limited logical architecture, such as that of a digital electronic computer.

Organizations, by Simon's prior account, followed a similar planning structure. Each ranked procedural needs into a unique hierarchy to avoid inconsistencies. *Administrative Behavior* argued for a hierarchy for organizations that optimized for an efficient mix of deference to purpose, process, clientele, and place.<sup>214</sup> The virtual Machine optimized for an efficient mix of trial-and-error search, systematic use of experience, cues in the total problem-solving process, a reasonable level of machine performance, and some minimum threshold of operational intelligibility.<sup>215</sup> Success in each domain was to be measured by the system's *internal interrelations* and the behaviours that resulted from their design rather than by, say, the system's materiality or relationality to ancestry. The rational character of an information system that could play 'good chess' thus resembled the rational character of "'good" administration.'<sup>216</sup> A 'good' system was one that exhibited a self-evident level of efficiency, capitalizing on internal dynamics that made it possible to accomplish similar functions across different contexts.

In 1956, despite having borrowed heavily from the esoteric domains of administrative and formal logics, Newell and Simon unabashedly compared the behaviour exhibited by their Machine to that of a human problem solver. In an internal report for the RAND Corporation, they justified the analogy via reference to four central capacities: the first was that the

---

<sup>214</sup> Simon, *Administrative Behavior*, 1997, 36.

<sup>215</sup> Newell, Shaw, and Simon, 'Elements of a Theory of Human Problem Solving.', 162.

<sup>216</sup> Simon and Newell, 'Current Developments in Complex Information Processing', 5; Simon, 'The Proverbs of Administration', 64.

Machine leveraged, ‘A liberal use of search processes.’<sup>217</sup> This distinguished it from machines that processed algorithms in a rote fashion, as one would do if one were to try every possible dial-setting from a safecracker. The Logic Theory Machine simulated *heuristics* not algorithms. Thus, it could ‘search’ new information in what Simon and his collaborators saw as an informed and adaptable manner, worthy of comparison to a human problem solver.

Second, the group saw further justification in the Machine’s use of ‘descriptions simulating... certain perceptual processes and its use of notions of “similarity.”’<sup>218</sup> That these perceptual tasks—identifying characteristics such as likeness and identity—could now be automated by a computer, suggested, to them, a formal equivalence between the two type of entities. That this type of perception, and the judgements it rendered, were exclusively functional was extraneous to the system’s designers.

The trio also claimed that the Logic Theory Machine’s ‘procedures for working backward from the goal to be attained [to the position it began at]’ provided ‘a means for guiding its selections of steps.’<sup>219</sup> Of course, this simply meant that the Machine leveraged means-ends analysis—the decision technique Simon had trialled in relation to organizations in *Administrative Behavior*. The fourth and final capacity that the group at RAND called upon to justify their comparison between man and machine was based on the latter’s use of ‘memory.’ The Logic Theory Machine’s ability to ‘store the theorems it proves and to use them in subsequent proofs,’ was, to their eyes, a human-like behaviour. It suggested that the system could do something functionally equivalent to learning from experience. (In fact, the use of adaptable machine memory was a core characteristic of stored-program computing. As such, it was not unique to the Logic Theory Machine, at least not in the way that the prior three capabilities were at that time.)

By the winter of 1956, Newell, Simon and Shaw had used the Machine to prove thirty-eight of fifty-two theorems from Chapter Two of *Principia Mathematica*. This impressive result raised the stakes of their man-machine metaphor and broadcast their results for others to consider. When Bertrand Russell heard the news, he exclaimed, ‘I wish Whitehead and I had known about this possibility before we wasted ten years doing it by hand.’<sup>220</sup> Russell’s

---

<sup>217</sup> Newell and Simon, ‘Abstract: The Logic Theory Machine’, 2.

<sup>218</sup> Newell and Simon, 2.

<sup>219</sup> Newell and Simon, 2.

<sup>220</sup> Simon, *Models of My Life*, 208.

quip implied that his and Whitehead's labour had been wasted. On the contrary, as I have shown, the conceptual tools they developed and refined in *Principia Mathematica* were essential contributions to the Logic Theory Machine.<sup>221</sup> Simon, Newell and Shaw had transmuted, rather than displaced, his and Whitehead's efforts; they had simply not cited the conceptual debts carefully. Russell and Whitehead's system of axioms, rules of inference and methods provided the RAND trio with an internally coherent system with which to model complex emergent behaviour. No equivalent system then existed for chess, which is why the Chess Learning Machine had failed. *Principia Mathematica*, in contrast, provided the American group with a workable starting point with which to simulate human problem-solving behaviour. *Administrative Behavior* then helped to transmute that logic into code.

#### 1956-58: Framing the Logic Theory Machine

Between 1956-58, Simon and Newell searched for how best to frame their results.<sup>222</sup> Their various accounts in published journal articles, in conference proceedings and in unpublished manuscripts from these years speak to the nuanced commitments implied by their formal analogy between human and machine behaviour. During this period, a few points remained consistent. First, the term 'artificial intelligence' was never used. Simon and Newell saw their contribution as a theory of problem solving, not intelligence. This was underlined by the fact that they both knew about the term following their participation in the 1956 Dartmouth workshop, which I return to in [Chapter Four](#). As such, the fabled 'First Prototype of AI' was not a contribution to AI, at least not in the eyes of those who had developed it. A second consistent point was that their result was not tied to biology. The words 'brain' and 'mind' do not figure in the early papers, except when they are mentioned as explicitly unrelated to their system in an article for *Psychology Review*. Their model was of human behaviour, not biology. They wrote:

We wish to emphasize that we are not using the computer as a crude analogy to human behavior—we are not comparing computer structures with brains,

---

<sup>221</sup> The *Journal of Symbolic Logic* elected not to publish the result with the Machine listed as an author because the theorems had already been proven in *Principia*. McCorduck, *Machines Who Think*, 1979, 142.

<sup>222</sup> Shaw's involvement slowed after this stage. He was thanked for having 'realized' their work in a computer but was not listed as an author in future papers.

nor electrical relays with synapses. Our position is that the appropriate way to describe a piece of problem-solving behavior is in terms of a program... Digital computers come into the picture only because they can, by appropriate programming, be induced to execute the same sequences of information processes that humans execute when they are solving problems.<sup>223</sup>

Other aspects of the Logic Theory Machine's significance were treated more ambiguously. One such issue—the system's reification of intellectual conformity—is simultaneously subtle and contentious, largely because it was implied and never explicitly considered by Simon or his collaborators in writing. In an initial trial of the virtual Machine held in January 1956, prior to the system's successful implementation in the JOHNNIAC device, Simon had his three children, his partner Dorothea Pye and a handful of his graduate students act out the sequence of functions and calls to memory that Shaw would later translate into machine code.<sup>224</sup> Each participant received a set of index cards that stated either a sub-routine, component rule, or logical axiom from the system's 'memory.' In his autobiography, Simon cast the step-by-step trial as an important milestone in the development of his machine, as well as in the history of science generally. 'We invented a computer capable of thinking non numerically, and thereby solved the venerable mind/body problem,' he claimed.<sup>225</sup>

As Boden points out, Simon would spend the next three decades of his career trying to justify this ambitious claim.<sup>226</sup> I want to draw specific attention to one aspect of his characterization. In his 1996 autobiography, where he published the claim, Simon equated the actions of his family and students in the test to those of the slave boy in Plato's *Meno*.<sup>227</sup> In this Socratic dialogue, Plato leads a young slave through an exercise in geometry. That the boy eventually solves a question revealed, for Plato (and later for Simon), that such knowledge was dormant within him all along; a capacity already possessed, not acquired. Simon used the metaphor to call attention to the genius of his Machine. *It too*, he had proved, could act out dormant intellectual capacities.

---

<sup>223</sup> Newell, Shaw, and Simon, 'Elements of a Theory of Human Problem Solving.', 153.

<sup>224</sup> Simon, *Models of My Life*, 207.

<sup>225</sup> Simon, 190.

<sup>226</sup> Boden, *Mind as Machine*, 327.

<sup>227</sup> Simon, *Models of My Life*, 207.

Simon's deliberate Meno reference also spoke to an aspect of relative shallowness present in his attempt to model human decision-making capabilities. A slave, like an organization optimized for a predetermined metric like profit, existed within a highly restricted field of agency, one that Simon glorified by showing how it could be manipulated to realize novel results. Neither entity, nor Simon's obliging friends and family, could truly question the premises assigned to them, as a free-thinking agent would. In this way, Simon's theory, and the system he operationalized, were each a model and celebration not just of cognition, but also of conformity.

He was not alone in reifying this characteristic. Gaspard de Prony, whose work had inspired Simon, had subverted the high regard computation held in French culture in the early nineteenth century when writing that a definitive *lack* of intelligence was desirable amongst the workers who would compute portions of his mathematical tables.<sup>228</sup> Moved by this insight, Babbage resolved to eliminate what he called the 'inattention, idleness... [and] dishonesty of human agent[s]' by designing engines capable of dividing labour according to workers' 'natural capacity and acquired habits,' a lawlike structure that, to Babbage's benefit, *he* oversaw and could manipulate.<sup>229</sup> Schaffer details how an emerging class of intellectual aristocrats in Victorian Britain sought social capital through analogous means by collectively theorizing mechanical equivalences for the measure of mental life.<sup>230</sup> Allusions to mechanization naturalized their newly christened 'brain-work' at the high end of the mechanized economic order of the day; a testament to maintaining the political status quo. Daston chronicles the rise of 'managerial intelligence' from this period through to the first half of the twentieth century, as epitomized by the lucrative insurance trade, growing government bureaucracies, and the actuarial sciences.<sup>231</sup> Horan, similarly, dubs the second half of the twentieth century the Actuarial Age.<sup>232</sup> In each era, elites in the U.S. and Britain gained and preserved power by manipulating the cultures of conformity bred by their idealisations of niche measures of mental life.

---

<sup>228</sup> As cited in: Daston, 'Enlightenment Calculations', 195. See also: Grier, 'Human Computers', 34–38.

<sup>229</sup> As cited in: Schaffer, 'Babbage's Intelligence', 210.

<sup>230</sup> Schaffer, 'OK Computer'.

<sup>231</sup> Daston, 'Calculation and the Division of Labor, 1750-1950', 12; Daston, 'Enlightenment Calculations'.

<sup>232</sup> Horan, 'Actuarial Age'.

It is worth pausing here to dwell on the implications of Simon's sophisticated reification of intellectual conformity, one that went on to inspire the intellectual concerns of generations of researchers. Although this nuance was not of explicit interest to his contemporaries in 1956, likely because its existence was not yet widely known, reviewers later returned to the subject in print. In 1989, Rupert F. Chisholm, a professor of management, critiqued Simon's theory of administrative behaviour for being unable to account for the distribution of authority within organizations. He questioned how dynamic human judgment could be reduced to the more static or mechanical outcomes of a bureaucracy, which stood vulnerable to exploitation from above. Chisholm wrote:

Simon fails to deal adequately with ways of protecting individuals against potential oppression by organizations. He implies that achieving rationality requires pursuing organizational goals to the maximum degree regardless of the propriety of these goals. This view, at the extreme, reduces individuals to automotons [sic] who blindly follow the orders of superiors. Albert Speer's (1970) description of the bizarre internal culture of the Third Reich gives a vivid example of carrying this logic to the extreme.<sup>233</sup>

Speer's book, *Inside the Third Reich*, had detailed why it had taken him so long to appreciate his complicity in terror as Minister of Armaments for the Nazis between 1942-45. He attributed the oversight partially to the tremendous 'administrative assignment' he had been given.<sup>234</sup> Speer had joined the Nazis to benefit his architectural career but came to realize, all too late, that his decision making had conformed to the decision premises shaped for him.<sup>235</sup> He did not 'think' while in this institution so much as agree to conform to the system he had entered.

This is a contentious and anachronistic parallel to draw given that Chisholm's critique was not made until the late 1980s. Depending on one's perspective, however, it could be equally contentious to leave the possibility of such dynamics plainly unacknowledged given that the Logic Theory Machine made use of and re-formalized, without clear citations, notions of restricted agency analogous to those theorized in *Administrative Behaviour*. This transmutation perpetuated Simon's orientation to goal-seeking behaviour in an entirely new

---

<sup>233</sup> Chisholm's critique built on critique by Herbert J. Storing. See: Chisholm, 'The Storing Critique Revisited', 424.

<sup>234</sup> Speer et al., *Inside the Third Reich*, xiv.

<sup>235</sup> I note that one of AI's most vocal critics, Joseph Weizenbaum, escaped from Nazi Germany with his family in 1936. Nilsson, *The Quest for Artificial Intelligence*, 393-98.



medium: digital computing. Hannah Arendt's reflections on the banality of evil capture the nuanced moral possibilities presented by the diffusion of responsibility across bureaucratic, or in this case, techno-bureaucratic structures.<sup>236</sup> Rob Nixon, similarly, coins the term 'slow violence' to challenge characterizations of harm as expedient, visceral or obvious. He writes, 'By slow violence I mean a violence that occurs gradually and out of sight, a violence of delayed destruction that is dispersed across time and space, an attritional violence that is typically not viewed as violence at all.'<sup>237</sup> Thom Davies asks, by extension, 'out of sight to whom?'<sup>238</sup> By glorifying conformity and a form of reduced agency dressed as machine-coded genius, Simon and Newell reinscribed the slow violence of administrative logics into the Logic Theory Machine.<sup>239</sup>

Between 1956-58, Simon and Newell abstained from commenting on either the political or philosophical significance of their Machine, even as they cast their results as a profound break in the history of scientific thought. In 1956, they published a series of papers under the title, 'The Logic Theory Machine: A Complex Information Processing System.' As I address in Chapter Four, their findings made waves at the 1956 Dartmouth Summer Research Project on Artificial Intelligence, which gave the field its name and initial membership.<sup>240</sup> In print, they framed their report as no more than a catalogue of the Machine's technical capacities, such as its ability to process complex information.

Interspersed in the groups' descriptions, however, were significant claims about how their work should be interpreted. The duo stated that their system was not a 'simulation' of complex behaviour but a 'realization' of it.<sup>241</sup> They argued that the term 'simulation' implied the deliberate imitation of a physical system as opposed to the simulation of an abstract set of characteristics. This point laid out the central premise behind what would later come to known as symbolic AI: that human behaviours could be stripped of their physical instantiation

---

<sup>236</sup> Arendt, *Eichmann in Jerusalem*.

<sup>237</sup> Nixon, *Slow Violence and the Environmentalism of the Poor*, 2.

<sup>238</sup> Davies, 'Slow Violence and Toxic Geographies'.

<sup>239</sup> In this vein, Simon argued in 1962 that complex systems were *characteristically* hierarchical. Simon, 'The Architecture of Complexity'.

<sup>240</sup> Simon and Newell, 'Plans for the Dartmouth Summer Research Project on Artificial Intelligence'.

<sup>241</sup> Newell and Simon, 'The Logic Theory Machine: A Complex Information Processing System', September 1956, 62.

yet retain the rhetorical weight of their initial embodiment when ‘realized’ in a new physical system.

The nature of this manoeuvre merits a moment’s attention as it represents the crux of their broader project. This transfer of abstract characteristics between physical forms required, in Simon and Newell’s formulation, that each entity first be conceived of as an ‘information processing system.’ This type of system involved two basic components acting in tandem: a set of memories and a set of information processes. A memory was a location where information was stored over time in the form of symbols, such as the register of a digital computer. An information process was a function that manipulated that information as its inputs and its outputs, such as a set of commands sent to a computer’s memory. In effect, the framework required that all sophisticated decision systems be reimagined as, at some level, rooted in symbolic logic.

A programming language was unequivocally an information processing system. Simon, Newell and Shaw dubbed their custom programming language for the JOHNNIAC, the ‘Logic Language’ or ‘Information Processing Language.’ They built the Logic Theory Machine using this language, just as contemporary programs are built using languages like Python or C++. The Logic Language gave structure to the memory inside the JOHNNIAC device. It modelled formalisms from *Principia Mathematica* into ‘linked lists’ that represented information as reflexive lists in a database. ‘Linked lists were one new form of materiality and representation designed to make logical propositions into digital things,’ writes Dick, who chronicles this re-formalism.<sup>242</sup>

Newell and Simon’s nuanced treatment of materiality exposed the metaphysical commitments underlying their idiosyncratic transmutation of formal logic into programming code. Like Alan Turing, they distinguished between a digital computer in *principle*, which could realize any degree of complexity, and a digital computer in *practice*, the realization of which was unambiguously limited in speed and memory due to the presence of natural laws. ‘In the real world, the appealingly egalitarian abstractions of the Church-Turing thesis quickly break down in the face of the temporal and spatial constraints of the physical universe,’ writes

---

<sup>242</sup> Dick, ‘Of Models and Machines’, 630; Dick, ‘After Math’, 47. This list-based logical architecture became influential for subsequent languages, like John McCarthy’s LISP language, which became the *lingua franca* of U.S. based AI research in the 1960 and 1970s.

Ensmenger.<sup>243</sup> These limitations beget novel solutions that put a premium on compression and expressive flexibility. ‘Talk’ to a computer had to be efficient.

Another obvious way to maximize efficiencies in simulation, besides negotiating the material affordances of the JOHNNIAC, was to mirror rational behaviours like problem solving. Doing so narrowed the proverbial range of discussion with a computer by bounding its analysis to instances of behavioural conformity and control, be they militarily inspired or economic. In a 1957 lecture to the Operations Research Society of America, a discipline born from and in service of optimal military logistics and strategy, Simon and Newell proudly positioned their Machine within this legacy. Their paper, ‘Heuristic Problem Solving: The Next Advance in Operations Research,’ positioned the Logic Theory Machine as a milestone in the entangled histories of capitalism and manufacturing.<sup>244</sup> Simon unironically credited Adam Smith, whose foundational contributions to Western economics served to orient modern capitalism, as the inventor of the computer. While these systems were not computers in practice, they were computational in principle, he argued. Simon explained that thinkers like Gaspard de Prony and Charles Babbage had simply translated Smith’s ideas into hardware through iterative stages of development.<sup>245</sup> In demoting physicists and electrical engineers from the list of significant contributors, he positioned the study of *social organization* as the primary source of inspiration for computation. As an appeal to the overlapping and increasingly evident value of management science and operations research, which applied ‘intelligence’ to administration, Simon stated:

For an appropriate patron saint for our profession, we can most appropriately look back a full half century before Taylor to the remarkable figure of Charles Babbage... He was one of the strongest mathematicians of his generation, but he devoted his career to the improvement of manufacturing arts, and – most remarkable of all – to the invention of the digital computer in something very close to its modern form.<sup>246</sup>

---

<sup>243</sup> Ensmenger, *The Computer Boys Take Over*, 31.

<sup>244</sup> As did Newell, who in 1957 outlined heuristic methods as appropriate for chess, theorem proving and management problems ‘like balancing a production line.’ Newell, ‘Tentative Outline of Michigan Summer Session Lectures’.

<sup>245</sup> Simon also credited the French weavers and mechanics responsible for the Jacquard loom, which he positioned as having led to punched card computing. Simon and Newell, ‘Heuristic Problem Solving’, 2–3.

<sup>246</sup> Simon and Newell, 1–2.

By aligning the Logic Theory Machine with the legacy of Taylor, the founder of management science, as well as the legacies of Babbage and Smith, Simon and Newell endorsed the broader social order that these figures had helped to bring about and sustain.<sup>247</sup> Like them, Simon and Newell sought to re-orient human social order around tractable notions of efficiency rendered through mathematics and other modes of quantification. RAND sought this end too; their Systems Research Laboratory was among the world's first experimental laboratories in management science. Simon and Newell's contribution was the notion that fine-grained human behaviours such as problem solving could be rendered in a symbolic language fit for 'realization' in a digital computer. Furthermore, they positioned the digital computer program as perhaps the *only* means that scientists had to gain insight into complex behaviour, since in their view programs could be taken as equivalent to theories.

This lineage casts the Logic Theory Machine in a different light than other histories have so-far offered. The Machine has been read primarily as a mathematical distillation of human problem-solving techniques. I have argued that it was more than that: it was also an endorsement of the broader political systems that centred those techniques. Recent histories of U.S. social science during the Cold War have surfaced conservative currents underlying modes of scientific research positioned, during that period, as value-neutral. Wolfe recounts how U.S. propaganda popularized a cultural ideal abroad of science as free and unencumbered by government interference.<sup>248</sup> Cohen-Cole demonstrates how personality traits like autonomy, creativity, and the use of reason were packaged into the concept of 'open mindedness' to function, ironically, as a shorthand for the virtuous mid-century democratic citizen, unaffected by Communist conformity. Bernstein explores how statism 'dramatically configured' American economists' faith in the functional independence of the market.<sup>249</sup> These histories revisit Schaffer's histories of expert 'intellectuals,' be they scientists, policy makers, or university administrators, advancing the specific normative aims of their own collectives.

In their 1957 talk to the Operations Research Society of America, Simon and Newell proudly endorsed the capitalist and military social orderings to which their techniques (on

---

<sup>247</sup> On the pre-history of this trend see: Agar, *The Government Machine*, 41; Schaffer, 'Babbage's Intelligence', 216.

<sup>248</sup> Wolfe, *Freedom's Laboratory*.

<sup>249</sup> Bernstein, *A Perilous Progress*, 4.

their account) contributed. Regrettably, to my knowledge, the duo did not comment in this period on another parallel mode of normativity, this time via what I have identified as the second key ingredient of their system: formal logic. Prior luminaries in that tradition has been more forthcoming. The logician Gottlob Frege, from whose account of logic Russell and Whitehead had borrowed to generate *Principia Mathematica*, had positioned formal logic as an idealized *approximation* to human problem solving and not as human problem solving itself, as Simon and Newell implied. In 1893, Frege stipulated that formal logic was ‘the way in which one *ought* to think,’ rather than how we actually think, a declaration that Leech argues was representative of his position on the matter.<sup>250</sup> Similarly, Immanuel Kant stated, ‘In logic we do not want to know how the understanding is and does think and how it has previously proceeded in thought, but rather how it *ought* to proceed in thought.’<sup>251</sup> These commitments, while strong in their own way, lacked the RAND group’s adherence to a functional equivalence between formal logic and routine human problem solving.

As I have shown, Simon’s high opinion of logic was definitive for the intellectual culture of the RAND Corporation in that period and in the wings of the University of Chicago in which he had trained in the 1930-40s. In Chicago, Carnap, who had been a student under Frege, had explicitly opted out of the question of whether or not logic was normative by 1937—while Simon was still a student there. In Carnap’s view, all logical systems were *prima facie* equal.<sup>252</sup> ‘Everyone is at liberty to build his own logic, i.e. his own form of language, as he wishes. All that is required of him is that, if he wishes to discuss it, he must state his methods clearly, and give syntactical rules,’ the philosopher insisted.<sup>253</sup> Thus, on whether or not logic was the substance of thought or its ideal mode, Carnap diverged from Frege, his instructor. Simon seems to have followed Carnap, either by implicitly rejecting the normativity of logic or by opting out of the question altogether as his instructor had.

At RAND in the 1950s, the rigorous application of formal analysis to the social sciences was a guiding principle, as evidenced by their development of game theory, systems analysis and other quantification techniques intended to rationalize behaviours as strictly logical or analytically tractable. Rodrigo Ochigame, who has studied mathematical logics developed in

---

<sup>250</sup> As cited in: Leech, ‘Logic and the Laws of Thought’, 1.

<sup>251</sup> As cited in: Leech, 1.

<sup>252</sup> Steinberger, ‘Frege and Carnap on the Normativity of Logic’, 154.

<sup>253</sup> Carnap, *The Logical Syntax of Language*, Sect 17.

Poland, Brazil and the Soviet Union in the same time period, situates the logical tradition at RAND as distinctly American. Scholars in other contexts abandoned commitments to absolute consistency and tractability, such as dialectical logic.<sup>254</sup> These non-classical logics reveal the normative status of contradiction in Frege and Russell's theories, Ochigame argues.

Simon's presumption of having crafted *universal* decision procedures, meaning those that would apply across a multiplicity of domains, was informed by his own idiosyncratic geographical and temporal context. He believed that classical theory bore universally, as did many people in the institutions in which he worked. Although RAND researchers were promised latitude to develop their own initiatives, they did so within an operational environment in which conformity to military needs was explicit, as evidenced by the corporation's thorough security regime, Primary Mental Abilities testing, university-esque recruitment techniques and architectural layout. In this way, the decision premises set out in the Logic Theory Machine aligned with the basic 'decision premises' of RAND and Simon's quarters of the University of Chicago; in each venue, the search space for new theory was confined to notions of strategy believed to be absolute.

In *The Closed World*, Edwards characterises these techniques (ex. systems thinking), and related technologies (ex. digital computers) and social systems (ex. U.S. military bureaucracies), as constitutive elements of a 'closed-world' discourse extending from America's strategic containment of Communism during the Cold War. Digital computing advanced as ideology, he argues, because belief in the possibility and desirability of a quantified world fed techniques like automatic control mechanisms that encoded conformity into human behaviour. I argue that through complex information processing and heuristic programming, Simon and Newell fused a RAND-funded Cold War desire for assurance through technological modes of quantification with longer-standing, value laden—and, in this case, uncited—assumptions about the relative worth of different aspects of mental life such as highly formalized modes of problem solving and derivative conformities, as well as the social orders that, historically, naturalized such orderings. The Logic Theory Machine thus marks an intersection between Edwards' closed world and Cohen-Cole's open mind; its creators championed the deference to technology and bureaucracy chronicled by Edwards but did so through a veil of neutrality and the apparent self-evidence of their claim that the Logic Theory

---

<sup>254</sup> Ochigame, 'Paraconsistent Society'.

Machine embodied problem solving. Simon and Newell celebrated computerized logistics as tantamount to thought itself, proselytizing a culture of deference to technology as an ideal of human agency.

That the RAND team were not criticized for this assumption speaks to the hyper-rational character of the community they inhabited. Other communities remained unconvinced. ‘The Logic Theory approach is closed-system thinking,’ summarised Robert Lechner in a 1959 letter distributed to colleague across Sylvania Electric Systems, a major U.S. electronics and computer-parts manufacturer.<sup>255</sup> Lechner, like the Oxford logician Hao Wang, took issue with Simon’s rhetoric, claiming that Simon had modelled a closed system yet treated it as open.<sup>256</sup> Indeed, the RAND group purposefully conflated the two systems. At the 1958 *Symposium on Creative Thinking* at the University of Colorado, they began, ‘We ask first whether we need a theory of creative thinking distinct from a theory of problem solving. Subject to minor qualifications, we conclude there is no such need.’<sup>257</sup> By equating problem solving with creativity, and then encoding it into the Logic Theory Machine, Simon and Newell valorized the same closed world ethos that surrounded them at RAND—only this time as cognition.

Simon and Newell resisted the term ‘artificial intelligence’ in the mid-1950s on the grounds that it had been *their* results, not Minsky and McCarthy’s, that validated the notion that human cognitive processes could be simulated. By the late 1950s, for no clear reason other than that artificial intelligence had gained traction elsewhere, Simon and Newell began to use ‘AI,’ ‘heuristic programming’ and ‘complex information processing’ interchangeably. Dick argues that Simon and Newell believed that minds and computers were ‘the same kind of thing,’ meaning ‘species of the genus information processor,’ as the duo put it in 1972.<sup>258</sup> To this I add that Simon viewed *organizations* through that same lens, meaning as a decision system whose inner workings could be rationalized, ordered, and encoded scientifically.

Simon’s contributions to administrative theory prefigured his work in computing in at least one important way: it taught him how to model complex behaviour in a closed

---

<sup>255</sup> Lechner to Cooperstein, ‘The Mechanization of Human Thought Processes I’, 21 April 1959.

<sup>256</sup> Lechner borrowed these categories from Frederic Bartlett, who in 1958 had distinguished between free ‘adventurous thinking’ and ‘closed thinking,’ the latter of which involved deduction of a limited number of elements. Bartlett, *Thinking*.

<sup>257</sup> Newell, Shaw, and Simon, ‘The Processing of Creative Thinking’.

<sup>258</sup> Dick, ‘Of Models and Machines’, 623; Newell and Simon, *Human Problem Solving*.

environment. Each project lionized conformity by equating closed systems with open ones.<sup>259</sup> My account enmeshes with that of Dick, who argues persuasively that for Simon and Newell, along with contemporaries like Wang, engagement with digital computing didn't just alter their intuitions about automated proofs, it altered their understanding about intuition itself, along with their understanding of formality. When implementing their ideas in a digital computer, these men re-formalised materiality and language, shifting, in the case of Newell and Simon, from paper and pencil to symbolic notation and linked-list storage in the JOHNNIAC's serial memory. My account offers new insight by showing how the materiality of computing re-formalised the men's *ideological* commitments by providing a new vessel for their administrative conception of human 'problem solving.'

## Conclusion

In this chapter, I have worked to reclaim the conceptual origins of Newell, Simon and Shaw's virtual program, the Logic Theory Machine. I locate precedents in administrative theory and formal logic that provided this trio with a set of novel conceptual tools. These frameworks helped them to formalize a sophisticated man-machine metaphor and to render it in code. I explore how, in this process of transmutation, elements from *Administrative Behavior*, such as decision premises and means-ends analysis, were repurposed alongside axioms from *Principia Mathematica* in a manner that, once translated into a machine language, recast formal logic as more universally applicable than its co-designers, like Frege, had taken it to be. This translation was acceptable in environments like the RAND Corporation in 1955-56, where the application of formal logic was a fertile part of the organization's vision and military infused operating principles.

By my account, institutional momentum played an important role in shaping the conception, legitimization and development of the Logic Theory Machine. The unique conditions in place at the RAND Corporation in 1952-56, afforded by the U.S. military to advance American Cold War priorities, brought Simon and Newell together and encouraged their exploratory venture with funding, time and access to world-class computing technology.

---

<sup>259</sup> On the politically fraught interplay between cognitive science and notions of 'open-mindedness' in the mid-century see: Cohen-Cole, *The Open Mind*.



The think-tank's university-like internal structure also connected Newell to Selfridge, which Newell experienced as a happy accident. While these episodes might have seemed like serendipity to these participants, they were the product of a well-funded ideological orientation toward systems of control capable of containing the cataclysmic risks posed by the Cold War.

## Chapter Three – On Frank Rosenblatt’s Perceptron Theory, 1945-61

While Simon, Newell and Shaw worked feverishly to mimic heuristic problem-solving techniques on the JOHNNIAC computer in Los Angeles, Frank Rosenblatt, a young psychologist based at the Cornell Aeronautical Lab in Buffalo, NY, developed his own ambitious theories of how to simulate niche aspects of mental behaviour. From the 1950s onwards, Rosenblatt examined how to mimic biological perception and memory storage using an IBM 704. His result was perceptron theory, which he developed until his untimely passing in 1971.

There is a relative dearth of historical scholarship on Frank Rosenblatt.<sup>260</sup> Popular accounts have tended to frame his contributions in relation to the non- or sub-symbolic school of AI, which interprets knowledge as *learned*, from the bottom-up, via complex interactions with an environment, as epitomized by human perception. Non-symbolic AI is contrasted with symbolic AI, which takes knowledge to be a top-down system of *rules* or procedures, epitomized by human problem-solving techniques, like those advanced by Simon and Newell. I eschew this dichotomy on the basis that emphasis on conceptual differences between these schools has overshadowed historical understandings of the significant commonalities between Rosenblatt and the RAND group. These include overlapping imbrications in military research priorities, radical faith in the explanatory potential of digital computing, an avowed commitment to computational understandings of cognition, and an unwillingness to identify as ‘AI’ researchers in the mid 1950s.

In this chapter, I draw on new archival evidence to situate Rosenblatt’s early career research, 1945-62, in relation to the U.S Navy’s coordinated post-war campaign to fund and retain domestic civilian scientists during peacetime. When Rosenblatt graduated from Cornell University as an undergraduate in 1950, faith in the transformative potential of new medical and statistical techniques to fortify civic life ran high amongst government and military research funding bodies. Within this milieu, Rosenblatt began his PhD research on the explanatory limits of statistics in personality science. I contrast his bold claims and research methodologies with developments made in psychotherapy and the U.S. insurance industry to

---

<sup>260</sup> Olazaran, ‘A Historical Sociology of Neural Network Research’; Boden, *Mind as Machine*; Seising, ‘A Brain Model for the Perception of the Outside World’.

normalize the scientific observation of families and marginalized peoples under the guise of patriotism, scientism and risk reduction. I argue that Rosenblatt's early faith in statistics, as well as his presumed entitlement to the prying research infrastructures needed to test and substantiate new computer-based statistical techniques, helps us understand his work on brain modelling after the mid-1950s.

In published work from the late 1950s and early 1960s, Rosenblatt characterised perceptron theory in opposition to AI. He positioned research on symbol-based notions of cognition as 'logical contrivances' that failed to honour biological phenomena with fidelity. At the same time, he turned to Friedrich Hayek's notion of decentralized market behaviour to justify his own theory of mechanised perception in a formal logical framework. I argue that Rosenblatt should be understood as part of a cohort of mid-century brain model researchers, along with Minsky, McCarthy and Simon, rather than as the biologically oriented defector he believed himself to be. I further show that, as for Simon, U.S. military funding, in this case from the Office of Naval Research and the Cornell Aeronautical Lab, helped to secure Rosenblatt's placement at Cornell University, even as his interests earned him an 'oddball' disciplinary status that did not fit well into existing structures.

#### 1945-56: Rosenblatt and Civilian-Led Military Science in Post-War America

For senior officials in the U.S. government and military, the end of World War II put at risk a key strategic advantage: the nation's scientific and technological supremacy. This advantage had been sustained, in part, by contributions to radar technology, the atomic bomb and numerous other research efforts made by civilian scientists nationwide. In 1941, President Roosevelt established the Office of Scientific Research and Development to coordinate and fund these efforts. Declarations of peace in 1945 threatened to end such patronage indefinitely. In response, a tenacious group of technocrats, including twenty-six of the nation's leading scientists, worked to 'sell' and 'educate' authorities in the Navy Department, Executive Branch and Congress on the notion that continued U.S. military

supremacy would require long-term investment in peacetime civilian-led research on 'weapons and weapons systems.'<sup>261</sup>

The Office of Naval Research (ONR) was established in 1946 to carry this torch. Per the Vinson Bill, the organization was to coordinate industry, government, the military and academia in the mould of *active war* partnerships established by the Office of Scientific Research and Development.<sup>262</sup> Gone, however, was the motivating force of war. This entangled American security needs with the development of American science and technology.<sup>263</sup> Even in peacetime, ONR's mandate was unequivocally militaristic. In accordance with federal law, it would 'plan, foster and encourage scientific research in recognition of its paramount importance as related to the maintenance of future naval power, and the preservation of national security.'<sup>264</sup> Colin Babb, ONR's official historian, characterizes the establishment of the group as a signal of a 'new kind of "peacetime," one where the boundaries separating the beginning and ending of war had become blurred or indistinguishable.'<sup>265</sup>

In a string of speeches and consultations with university legal teams held during the winter of 1945, Commander R. D. Conrad of the U.S. Navy persuaded administrators at the University of Chicago, University of California, Harvard University, Massachusetts Institute of Technology and the California Institute of Technology to sign research contracts with the military.<sup>266</sup> Since no other national agency was yet in place to fund and coordinate such work, the ONR was able to capitalize on exclusive access to top civilian researchers. Within its first three months, the organization awarded an estimated two hundred research contracts,

---

<sup>261</sup> This account is provided by a group of technocrats recounting their experience in a self-described 'authoritative history' of the ONR written in 1961. See: Old, 'The Evolution of the Office of Naval Research', 30.

<sup>262</sup> Babb, 'The Genesis of the Office of Naval Research'.

<sup>263</sup> A great deal has been written about this intersection. See: Slayton, *Arguments That Count*; Wolfe, *Freedom's Laboratory*; Kaiser, 'Nuclear Democracy'; Erickson, *How Reason Almost Lost Its Mind*; Agar, *Science in the Twentieth Century and Beyond*, 263–402; Carson, *Reappraising Oppenheimer*; Lowen, *Creating the Cold War University*; Edwards, *The Closed World*; Leslie, *The Cold War and American Science*; Galison, 'The Ontology of the Enemy'.

<sup>264</sup> National Research Council (U.S.), *Naval Engineering in the 21st Century*, 18.

<sup>265</sup> Babb, 'The Genesis of the Office of Naval Research'.

<sup>266</sup> Old, 'The Evolution of the Office of Naval Research', 35.

seventy five percent of which went to universities and colleges.<sup>267</sup> The total investment topped twenty-two million dollars.

ONR funding was not uniformly contingent on fulfilling military aims. In 1946, Cornell University (where Rosenblatt spent his career) was awarded ONR grants for a total of twelve of twenty-eight proposed projects. Cornell reported its obligations in return for this largesse as being limited to a 'report from time to time', in what amounted to a 'new departure for the Navy.'<sup>268</sup> Still, despite this comparatively hands-off approach, the ONR's sizeable budget and ability to select projects afforded it discretion over the direction of basic research. From 1946-53, Deputy Science Director Mina Rees made the study of special-purpose digital machines and high-speed memory an ONR priority. Rees made explicit her appetite for experimentalism when stating an intention to keep ONR, 'Two years ahead of the time.'<sup>269</sup> For the U.S. Air Force, she explained in 1950, this meant investment in 'statistical control' technologies and systems designed to solve extremely large matrix problems.<sup>270</sup>

By the 1950s, statistics had long developed in parallel with aspirations for social control. Francis Galton, Karl Pearson and other proponents of eugenics refined and peddled statistical measures of population control and selective breeding during the late nineteenth and early twentieth centuries.<sup>271</sup> Statistics gained institutional status as a standalone discipline following the creation of the *Annals of Mathematical Statistics* in the United States in 1930.<sup>272</sup> While the atrocities of the Second World War contributed to a rejection of hereditary explanations of differences within groups of the population, other statistical explanations for social orderings, such as the measure of an individual's mental capacity via IQ testing, blossomed.<sup>273</sup>

As statistics matured, it also fractured. Debate in the *Journal of the American Statistical Association* in 1926 calcified two divergent camps, one committed to the theoretical aspects of mathematical statistics, and the other to the less rarefied concerns of

---

<sup>267</sup> Babb, 'The Genesis of the Office of Naval Research'.

<sup>268</sup> Kramnick and Altschuler, *Cornell: A History, 1940–2015*, 12.

<sup>269</sup> Rees, 'The Federal Computing Machine Program', 735.

<sup>270</sup> Rees, 736.

<sup>271</sup> For an overview of the history of eugenics see: Kevles, *In the Name of Eugenics*, 1995.

<sup>272</sup> Hunter, 'Drawing the Boundaries', 7–8.

<sup>273</sup> Carson, 'The Culture of Intelligence', 645.

practical applied statistics, free from advanced mathematics or statistical theory.<sup>274</sup> Matthew L. Jones identifies Rees as a primary driver of both funding and symbolic support for theoretical statistics, along with the mathematical statistician Harold Hotelling.<sup>275</sup> Ironically, Jones notes, Rees championed these investments off the back of the success of *applied* statistics during World War II.

As with information technology, scientific medicine garnered a significant influx of government investment in the United States in the post-war period. The first civilian treatment with penicillin in 1942, along with new therapies for infectious diseases like malaria, pneumonia and tuberculosis, buoyed public expectations in a revolutionary era of medical cures, treatments and discovery.<sup>276</sup> Members of Congress faced pressure to fulfil on this promise. Laboratory and office space at the National Institute of Health tripled in size by 1958. *Collier's* magazine celebrated the Institute for its ambition to 'put medicine back together again,' estimating an average life span of ninety by the turn of the century. A new era of medical science was seen to be afoot.

It was into this nexus of ambitious state investments in health and information technologies that Frank Rosenblatt entered academia in the late 1940s. Born to Katherine Rose Rosenblatt and Frank F. Rosenblatt, a publisher, in New Rochelle, N.Y. in 1928, Frank went on to attend the reputed Bronx High School of Science in the year below Marvin Minsky.<sup>277</sup> He pursued social science and biology at Cornell, which was overcrowded with post-war applicants at the time, and graduated in 1950 with a major in social psychology. In a 1969 interview, he recalled feeling disdain for having to memorize biological nomenclature as an undergrad, preferring instead to engage with contemporary theory.<sup>278</sup> He excelled in

---

<sup>274</sup> Hunter, 'Drawing the Boundaries', 13–14.

<sup>275</sup> Hotelling supervised the PhD of Kenneth Arrow, whom Amadae credits with developing the 'science' of decision theory; a reductionist, axiomatic and set-theoretic treatment of human rationality that paralleled Hotelling's reductionist bent on theoretical statistics. Amadae, *Rationalizing Capitalist Democracy*, 3.

<sup>276</sup> Mandel, 'Beacon of Hope', 1–8.

<sup>277</sup> 'Dr. Frank Rosenblatt Dies at 43; Taught Neurobiology at Cornell'; 'Faculty Information: Frank Rosenblatt'.

<sup>278</sup> Frank Rosenblatt, Interview with Frank Rosenblatt, interview by Micheal Wright, July 15, 1969, 112.

sociology, economics and philosophy but struggled with hard sciences.<sup>279</sup> In 1950, he graduated and set out on a PhD in experimental psychopathology.

In a 1958 article, Rosenblatt claimed that it was during his PhD that he first considered the mathematics of brain modelling.<sup>280</sup> The bulk of his writing that survives from the early 1950s, however, pertained to the use of statistics in social psychology, not neurology. From 1951-53, Rosenblatt studied schizophrenia as a Fellow at the U.S. Public Health Service.<sup>281</sup> He credited this research for instilling in him the conviction that, 'the problems of measurement and data analysis would prove fundamental to scientific progress in psychopathology.'<sup>282</sup> He was not alone in approaching schizophrenia through the lens of data analysis. In 1956, Gregory Bateson, a driving figure in the cybernetic movement, published a theory of schizophrenia based on communication theory, notions of learning and the successful transmission of information in social settings.<sup>283</sup>

In his 1956 dissertation, Rosenblatt set out to design a new method to apply statistics to social psychology. 'The k-Coefficient: Design and Train Application of a New Technique for Multivariate Analysis' outlined a statistical technique for multiple correlations.<sup>284</sup> The paper positioned psychology as too computationally complex to permit a continued use of classical statistical techniques like multiple regression equations and multiple correlation coefficients. In practice, the limits of the techniques were quickly reached. Solutions derived from multiple regression equations became too approximate once the number of independent variables in use grew sufficiently large, or once other relevant contingencies, like interactive relationships between variables, were taken into account.<sup>285</sup>

Rosenblatt's solution was not to improve computers but to improve statistics *using* computers. He explained:

All research psychologists are familiar with problems in which the simultaneous working of a large number of variables seems to determine a

---

<sup>279</sup> 'Student Transcripts'.

<sup>280</sup> Rosenblatt, 'The Design of an Intelligent Automaton', 7.

<sup>281</sup> Rosenblatt, 'Two Theorems of Statistical Separability in the Perceptron', 1959, 420.

<sup>282</sup> Rosenblatt, 'The Design of an Intelligent Automaton', 7.

<sup>283</sup> Bateson based this work on Bertrand Russell's theory of logical types. See: Bateson et al., 'Toward a Theory of Schizophrenia'.

<sup>284</sup> Rosenblatt, 'The K-Coefficient'.

<sup>285</sup> Multiple regression equations ascertain the joint effects of a number of independent variables.

piece of behaviour, or a personality trait, or the outcome of an experiment. Such complex relationships are not peculiar to psychology; they are equally true, for example, of the gas laws in physics. However, psychology more than the physical sciences must deal with these relationships statistically, rather than as perfect mathematical functions... if we are going to take account of all the curvilinear and interaction relationships that might exist between our variables.<sup>286</sup>

By the 1950s, the use of psychological evidence to inform statistical techniques was commonplace. At the turn of the century, the British psychologist Charles E. Spearman, influenced by the eugenicist Francis Galton, had developed factor analysis to identify a hypothesized logical structure behind various tests of intercorrelated mental phenomena, such as different types of intelligence. Spearman hypothesized that underneath all expressions of intelligence, be they musical, scientific, or otherwise, lay a single common intelligence factor, '*g*.' This figure was intended to serve as a complete description for a certain pattern of phenomena across all cases. Factor analysis was also used as a heuristic, or incomplete theory, when no better interpretations were available.

In 1935, the American psychologist Louis Leon Thurstone, a pioneer of psychometrics, explored the explanatory potential of *multi* factor analysis. Thurstone challenged Spearman's unitary theory of intelligence by devising a measure for results *between* psychological factors. He derived this approach from matrix theory, a branch of mathematics in which rectangular arrays of expressions, meaning rows and columns of quantities, are treated as single entities. This placed his work within the more mathematical branch of statistics. That multi factor analysis required considerable computational resources is prefigurative of Rosenblatt's aims. Thurstone himself had to hire twelve expert human computers to complete his book on the topic.<sup>287</sup>

Rosenblatt's PhD questioned how *digital* computers, rather than human computers, would alter the design and use of statistics in psychology. He set out to develop and benchmark his k-coefficient technique against conventional methods like Thurstone's multi factor analysis and Paul Lazarsfeld's latent structure analysis, testing considerations like

---

<sup>286</sup> Emphasis his. Rosenblatt, 'The K-Coefficient', 1.

<sup>287</sup> Thurstone thanked a Leone Chesire for integral support to the project, an allusive citation of women's underappreciated intellectual labour in the histories of computation. Thurstone, *The Vectors of Mind*, x; Miltner, 'Girls Who Coded'.



computability, accuracy and reliability between variables.<sup>288</sup> Rosenblatt hoped that his k-coefficient would recapture complexity that had been discarded in order to reduce the number of variables to something computable. 'There is no doubt but that some information is lost in a factor-analysis,' he challenged, 'The "factors" can in any case only be estimated from other related measurements – they can not be measured directly.'<sup>289</sup> In essence then, his k-coefficient was to improve on the resolution of an iteratively more complex yet philosophically similar line of prior statistical theories.

With funding and technical support from Cornell's Department of Psychology, and several hundred hours of donated computer time and advice from the staff of the Cornell Computing Center, Rosenblatt co-developed the digital Electronic Profile Analyzing Computer (EPAC) between 1951-53, a custom-built device based, in part, on the ENIAC and funded, in part, by the U.S Public Health Service.<sup>290</sup> With use of war-surplus vacuum tubes, EPAC could input twenty-five items per second from paper marked in electrographic pencil. A panel of neon lights then displayed one of forty-two possible outputs, each symbolizing a non-zero sum of the item's squared differences. With the push of a button, the evaluation process reset.

The purpose of EPAC was to process data collected from a paid, six-hundred item survey of more than two-hundred Cornell undergraduates. The survey was conducted on paper and then fed into the computer. These data would be used to test the viability of Rosenblatt's k-coefficient against a novel baseline: the complex relations between personality type and familial relationships during a student's first twelve years of life.<sup>291</sup> Survey questions pertained to personal relationships between a student, their parents and their siblings, as well as other permutations of that roster. Rosenblatt's ambitious study generated results requiring two and a half million arithmetic operations to model and test, a task that necessitated use of an IBM Card-Programmed Electronic Calculator in addition to EPAC.

A 1953 article in the *Ithaca Journal* entitled 'Student Designs "Idiot Brain" to Measure Answer Patterns' captured the cultural and emotional climate that the twenty-five year old

---

<sup>288</sup> Rosenblatt, 'The K-Coefficient', 4; Thurstone, *Multiple-Factor Analysis*; Lazarsfeld, 'The Logical and Mathematical Foundations of Latent Structure Analysis'.

<sup>289</sup> Rosenblatt, 'The K-Coefficient', 4.

<sup>290</sup> Colleagues called EPAC 'Frank's Machine.' Thanks to Bruce W. Knight for comment.

<sup>291</sup> Rosenblatt, 'The K-Coefficient', 48–49, 205–9.

psychologist had tapped into by exploring computer-driven psychology research in the early 1950s. 'Rosenblatt is testing the idea that personalities could be classified in a scientific and objective way,' the article read, 'If the answers fall into clusters which are very similar to each other or very different from each other, the chances are that personalities can be classified objectively.'<sup>292</sup> This characterization blurred the lines between absolute and heuristic measurements, the latter being a useful approximation to, but not an acceptable measure of, complex psychological phenomena.

The *Journal's* readiness to deride the capabilities of Rosenblatt's 'idiot brain' device while simultaneously granting it authority as a 'scientific' and 'objective' view of personality research is indicative of the liminal cultural status of digital computing at the time. Writes Edwards, 'In the 1950s, when computers were still very new and rather awe-inspiring, any application automatically inherited their aura of almost erotic scientificity.'<sup>293</sup> Edmund Berkeley, author of the popular 1949 book *Giant Brains, Or, Machines That Think*, which introduced a generation of Americans to the field, outlined the risks of this scientificity in a chapter on computing and 'Social Control.' 'It is not right nor proper for a scientist, a man who is loyal to truth as an ideal, to have no regard for what his discoveries may lead to,' he wrote, in regards to how computing tools would be used.<sup>294</sup> Berkeley cautioned that human prejudice occurred 'before judgment' and thus could not be swayed by information.<sup>295</sup> He lamented that scientific developments amplified the consequences of prejudice, citing the deaths of 70,000 people via 'a single weapon' at Hiroshima four years earlier. Berkeley's implication was that computing would amplify prejudice and harms in new ways as well.

Given this mix of cultural fascination and fear, it is no surprise that Rosenblatt's blend of personality research and computation found public attention in the *Ithaca Journal*. At the time, personality research carried on its own aura of scientificity. In 1953, the Bernreuter Personality Inventory, a 125-item personality test developed by Stanford PhD student Robert Bernreuter, sold more than one million units.<sup>296</sup> Caley Horan explicates how Americans'

---

<sup>292</sup> 'Student Designs "Idiot Brain" to Measure Answer Patterns'.

<sup>293</sup> Edwards, *The Closed World*, 121.

<sup>294</sup> On the book's status, see: Boden, *Mind as Machine*, 1074; Berkeley, *Giant Brains or Machines That Think*, 205.

<sup>295</sup> Berkeley, *Giant Brains or Machines That Think*, 205–6.

<sup>296</sup> Gibby and Zickar, 'A History of the Early Days of Personality Testing in American Industry', 172.

fascination with the quantification of social behaviours in the post-war period both fostered and was fostered by a concerted push on the part of the private insurance industry to introduce neo-liberal notions of risk and governance into family life and the provision of social welfare. Indeed, since 1870, consumer credit bureaus in America had quantified fundamental aspects of civic identity like trustworthiness, while simultaneously normalising large-scale surveillance infrastructures.<sup>297</sup> In Horan's account, risk-based 'actuarial logics' normalized archetypal social profiles at the expense of tacit intricacies of character, lending credence to the notion that a person (or certain people) could—and in some sense should—be understood through the prism of statistics. Feeding into this trend was an ascendant and well-resourced insurance industry as well as racist and classist cultural inclinations to deem certain groups a statistical or probabilistic 'risk' to society.<sup>298</sup>

Notions of family life were also transformed by Cold War risk logics. Deborah Weinstein charts how the family unit came to be seen by a set of psychological researchers as a brooding source of perversion when proper home-life conditions were not met. Using concepts drawn from cybernetics and systems theory, including Bateson's research, psychotherapists reframed the family unit as a locus of pathology in the post-war period.<sup>299</sup> The family came to be seen as a 'system' amenable to close observation and measurement, with different patterns of interaction and communication revealing instances of perceived vulnerabilities.<sup>300</sup>

Weinstein's account explains elements of Rosenblatt's PhD research. About his survey's results, Rosenblatt wrote that parents' overprotectiveness during a male child's first twelve years of life, 'appears to be strongly related to a pessimistic outlook.'<sup>301</sup> His experimental design probed for additional correlations in this vein. He asked, for instance,

---

<sup>297</sup> Of note is that large-scale commercial surveillance often preceded large-scale state surveillance in America. See: Lauer, *Creditworthy*. On the financialization of U.S. culture after World War II see: Krippner, *Capitalizing on Crisis*. On the normalization of public survey techniques and findings see: Igo, *The Averaged American*. On the history of privacy as an idea and ideal see: Igo, *The Known Citizen*.

<sup>298</sup> Horan, 'Actuarial Age', 6; Bouk, *How Our Days Became Numbered*. On blackness as a longstanding site for surveillance see Browne, *Dark Matters: On the Surveillance of Blackness*.

<sup>299</sup> Weinstein, *The Pathological Family*, 47–81.

<sup>300</sup> On the entanglements of bodies and disembodied techno-logics, see: Wernimont, *Numbered Lives*.

<sup>301</sup> Rosenblatt, 'The K-Coefficient', 67.

whether or not the subject's mother / father had tended to (a) sulk to make them feel sorry for something they did, (b) make and act on threats to their partner or individual children, (c) aim to be the centre of attention, (d) worry about their appearance, (e) follow a strict moral code, (f) act depressive or moody, or (g) take to bed at the slightest complaint, even when there was nothing seriously wrong.<sup>302</sup> The list continued. Participants were expected to rate their response on a five-point scale between extremes such as:

- My mother always seemed ready to blame my father for anything that went wrong; she seemed to think that everything that happened was his fault.
- My mother would never blame my father for anything, even when it was clearly his fault; she would never admit that he was in the wrong about anything.<sup>303</sup>

Undergirding this survey design was the assumption that the complexities of human behaviour could be made legible, sortable and, presumably, medically and/or politically rectifiable through the development of new statistical techniques and the modes of social engineering they might enable. Rosenblatt's EPAC, the 'idiot brain,' would accelerate progress toward this horizon by fulfilling the labour required to parse multi-million-point arithmetic operations. Alongside this hope, however, Rosenblatt cautioned that advances in computing would not guarantee progress for the field. Classical statistical techniques such as multiple regression equations and multiple correlation coefficients offered considerable conceptual flexibility, but their use would inevitably exhaust 'Even the most revolutionary developments in the field of computing machines.'<sup>304</sup>

Due to limited time and funding, Rosenblatt was unable to trial all of his proposed techniques.<sup>305</sup> He failed to prove that his k-coefficient was better than conventional statistics at estimating multiple correlations for large numbers of variables. 'The possibility still remains that we might obtain equally good results through a factor analysis of the independent

---

<sup>302</sup> Rosenblatt, 90.

<sup>303</sup> Rosenblatt, 99.

<sup>304</sup> Multiple correlation coefficients measured the proportion of variance between complex equations. Rosenblatt noted that curvilinear interactions between numerous variables would quickly become computationally intractable after effects were taken to their fourth or fifth order. Rosenblatt, 3.

<sup>305</sup> He was only able to test those relevant to linear rather than curvilinear relationships, for instance.

variables,' he concluded.<sup>306</sup> In a twist, given his high ambitions, Rosenblatt also cautioned that the direction of his research might, in fact, create new technical challenges, such as low correlations that came out spuriously high. He could not tell whether spurious inaccuracies were inherent to his k-coefficient or whether irrelevant data simply caused patterns to be found in statistics that did not relate to the real world.

Despite these experimental shortcomings, Rosenblatt neglected to acknowledge a fundamental limit, or upper bound, on the explanatory potential of statistical techniques in psychology. Instead, he pointed rosily to a departure point for the field, meaning the line of experimentation in which results could no longer be managed via desk calculations alone. Centring the computer in psychology was, he championed, inevitable. Rosenblatt simultaneously neglected to consider, or at least explicitly acknowledge, any ethical limitations to this line of study, in defiance of those like Berkeley who advocated that scientists like him be mindful of such things. Rosenblatt characterized his highly-sensitive survey results as 'incidental' to the experiment's larger aim of developing his k-coefficient.<sup>307</sup> With an eye to the future, he echoed Simon and Newell in advocating for complex problems to be solved using analytical techniques like factor analysis, which broke large problems into simple ones, as Lazarsfeld and Thurstone had. The future of psychology for Rosenblatt lay in *mathematical* rather than practical statistics.

With his PhD, Rosenblatt sought to measure the strength of the relationships that existed between variables in psychology. He assumed these relationships would conform to empirical measurement and turned to digital computing as a new means with which to test this hypothesis. EPAC greatly reduced computational labour, a precedent that enabled the development of new statistical techniques like the k-coefficient. Yet this promising horizon remained just that—a horizon not yet reached. Waning time and money forced the young scientist to conclude his project before he could demonstrate that his technique meaningfully improved on classical methods. He nonetheless completed the project convinced that the complexities of personality research, social psychology and psychopathology were amenable to statistical methods and characterization via digital computation.

---

<sup>306</sup> Rosenblatt, 'The K-Coefficient', 66.

<sup>307</sup> Rosenblatt, 7.

Hints of an alignment between this aim and the general aims of the U.S. military can be found in the project's seven-item bibliography. In it, Rosenblatt cited Lazarsfeld's research on latent structure analysis, which he returned to throughout the project as an example of theory he hoped to improve upon. Lazarsfeld had published this work in *Studies in Social Psychology in World War II, Volume IV, Measurement and Prediction*, a volume produced by the Research Branch, Information and Education Division of the War Department, whose responsibility it was to study the role of ideology in World War II and to 'stiffen the ideological supports of the men.'<sup>308</sup> Even if Rosenblatt's project was not directly funded by the ONR or another military sponsor, its orientation brought him close to those institutions in ways that became symbiotic as his career progressed.

#### 1955-58: From the Statistics of Personality Research to the Statistics of Brain Modelling

While neither memory nor cognition received direct treatment in his dissertation, Rosenblatt claimed that his transition from personality research into brain modelling was seamless.<sup>309</sup> A look at his early publication and employment record challenges that claim while also revealing how his career intersected with U.S. Air Force prerogatives. In 1955, as Minsky and McCarthy prepped for Dartmouth workshop, Rosenblatt joined the Cornell Aeronautical Laboratory in Buffalo, NY, a hotbed for military aircraft testing since its establishment in 1946. As was the case for Simon and Newell, who operated out of RAND's Systems Research Laboratory, Rosenblatt was employed by the Aeronautical Laboratory's Systems Research Department. Due to his computer experience, he was tasked with the study of systems planning and systems analysis under the title of Research Psychologist.<sup>310</sup> In his first three years, he made contributions to undisclosed 'information processing and weapons control systems.'<sup>311</sup> He recalled in a 1969 interview feeling like an outsider, a psychologist

---

<sup>308</sup> 'Studies in Social Psychology in World War II', 3.

<sup>309</sup> Rosenblatt, 'The Perceptron: A Perceiving and Recognizing Automaton (Project PARA)', Preface, ii.

<sup>310</sup> Rosenblatt, Interview with Frank Rosenblatt, 114.

<sup>311</sup> Rosenblatt, 'The Design of an Intelligent Automaton', 7.

hired into an aeronautical laboratory because ‘everybody else was getting one and they thought that they should have one too.’<sup>312</sup>

In that same year, Rosenblatt published ‘Parallax and Perspective During Aircraft Landings,’ in the *American Journal of Psychology*—a paper that indicated how his research aims and the ideological aims of the U.S. military had overlapped.<sup>313</sup> In it, he interpreted the mathematics underlying subjective components at work in the psychology of locomotion. The paper explored the spatial judgements required to successfully land an aircraft, a line of inquiry that would influence his perspective on human intelligence. Rosenblatt’s turn toward perception was, in part, a consequence of this aspect of the military’s agenda. Research for ‘Parallax’ was funded by the U.S. Air Force and co-authored with Paul Olum, a Cornell mathematician, and James J. Gibson, a former U.S. Air Force researcher whose work posited that human perception occurred without mediation via sense-data or other intermediary cognitive processes.<sup>314</sup> In their account, the trio provided a mathematical analysis of what they termed ‘motion perspective,’ a generalized description of various parallax phenomena, meaning the tell-tale clues—size, shading, superposition—that the brain received to gauge the distance of a given object. As was the case in Rosenblatt’s dissertation, the goal of this research was to consolidate complex phenomena down to simpler patterns. In this case, the phenomena involved were *perceptual* dynamics, not the dynamics of childhood experience and adult personality profiles.

Rosenblatt’s research over the next five years elaborated on the dynamics of perception, the topic that would define his career. In January 1957, he prepared an internal report entitled, ‘The Perceptron: A Perceiving and Recognizing Automaton (Project PARA).’<sup>315</sup> Project PARA was launched at the Cornell Aeronautical Laboratory in July 1956 with sponsorship from the Mathematical Sciences Division and, after July 1957, the Information

---

<sup>312</sup> Rosenblatt, Interview with Frank Rosenblatt, 114.

<sup>313</sup> Contract No. AF33(038)-22373. See: Gibson, Olum, and Rosenblatt, ‘Parallax and Perspective during Aircraft Landings’.

<sup>314</sup> Large, *Ecological Philosophy*, 131.

<sup>315</sup> Rosenblatt, ‘The Perceptron: A Perceiving and Recognizing Automaton (Project PARA)’.

Systems Branch of the ONR, whom Rosenblatt thanked for financial and intellectual support.<sup>316</sup>

As had been the case with his thesis, Rosenblatt's aim with Project PARA was twofold: to demonstrate the feasibility of a new statistical technique and to trial the 'workability' of that technique on custom-built hardware. He aimed to establish the technical and economic feasibility of formulating a 'brain analogue' capable of recognizing similar patterns of optical, electrical or tonal information. By January 1957, research on this front was limited to a statistical model only. Within a year, however, he had run simulations of new techniques on an IBM 704 to test to what extent probability theory could represent certain factual aspects of biological memory storage, such as equipotentiality, meaning the notion that the brain could redistribute and carry out memory functions even after an injury.

Rosenblatt's military interlocutors saw a strategic value in this initiative. Project PARA was approved by Alexander Stieber, Head of the Air Defense Section of the Systems Research Department. A funding extension request for a computer and three members of staff emphasized that 'the collation of military intelligence' was one of PARA's four central uses.<sup>317</sup> Of the hundred or so entities included on a 1958 distribution list for a follow-up report, approximately eighty percent were either academics, primarily in psychology and engineering, or military organizations, including the ONR, the Armed Services Technical Information Agency, National Security Agency and Air Force Cambridge Research Center.<sup>318</sup>

Military applications did not figure in Rosenblatt's account of perceptron theory in January 1958. Of importance, in his view, was that the theory advanced a robust alternative to the 'paper exercises' of contemporaries like Pitts, McCulloch, S. C. Kleene and J. T. Culbertson, whose deterministic logical notations belied understandings of biological fact.<sup>319</sup> Neural dynamics operated like a complex switching function, he argued, not like the serial

---

<sup>316</sup> I have not yet corroborated the ONR's rationale for this grant. As stated, high-speed memory and special-purpose digital machines were core foci. Babb notes that primary source documentation may no longer exist. In 2016, the archivist at the National Archives at College Park Maryland stated that the ONR destroyed early records in the 1970s. Some duplicate records may exist in unsorted portions of their extensive collection.

<sup>317</sup> Rosenblatt, 'The Perceptron: A Perceiving and Recognizing Automaton (Project PARA)', Preface.

<sup>318</sup> Rosenblatt, 'The Perceptron: A Theory of Statistical Separability in Cognitive Systems (Project PARA)', 263–68.

<sup>319</sup> Rosenblatt, 2.



‘memory’ of a computer. Rather than interpret information using strings of symbols and the taxing logical depth of a digital stored-program computer, which demanded ever-more storage and processing power to compute each sequential step, perceptron theory was to identify or categorize objects ‘directly,’ like a reflex.<sup>320</sup> Reflexes appeared to require a relatively shallow layer of logical depth (i.e. number of logical steps). If tapped, a human knee would jerk—no significant processing would be needed. Rosenblatt proposed to train a machine to operate in a similar fashion, via the statistical logic of association.<sup>321</sup> He saw this avenue as a means to avoid the combinatorial explosion that accompanied probabilistic methods, which lacked the coherence and computational tractability of a pre-determined conceptual schema provided by, say, formal logic or calculus.

Despite his appeals, Rosenblatt’s account was not so dissimilar from that of his peers. Like Simon and Newell, he centred on language as a ‘first requirement’ for adequate formal analysis of a neural network.<sup>322</sup> Like the RAND group, he developed a highly formalized language; perceptron theory simply leveraged statistics rather than, say, in the case of Pitts and McCulloch, Boolean algebra. Like Minsky, whose work I explore in [Chapter Five](#), Rosenblatt justified his theory in reference to notions of economy. He set out to pragmatically ‘economize’ otherwise irreducibly complex stimulus patterns by only storing information required for the categorization of inputs.<sup>323</sup> He called Project PARA, ‘the theory of *statistical separability*’ because perceptron theory only distilled the statistical behaviour that an artificial network used to represent an observed object.

A step-by-step walkthrough of perceptron theory in action will help to clarify these procedures. Rosenblatt’s statistical model was designed to identify divergent patterns, like an ‘X’ or a ‘Y’, as rendered in momentary iterative snapshots of a given stimuli. As the system witnessed enough of these snapshots, it ‘learned’ to recognize, or model, a representation of their structure. First, a class of stimuli was chosen. A *photoperceptron* involved learned visual patterns. A *phonoperceptron* learned tonal patterns. Next, a sequence of stimuli was provided for the perceptron to learn from. Each new stimulus fed into the perceptron’s three

---

<sup>320</sup> Rosenblatt, ‘The Perceptron: A Perceiving and Recognizing Automaton (Project PARA)’, 2.

<sup>321</sup> Minsky explored but ultimately abandoned associative learning as a primary building block in his own PhD research. He chose reinforcement learning instead. See: [Chapter Five](#).

<sup>322</sup> Rosenblatt, ‘The Perceptron: A Theory of Statistical Separability in Cognitive Systems (Project PARA)’, v.

<sup>323</sup> Rosenblatt, 2.

sets of model neurons: its Sensory System, Association System and Response System (Figure 1.0).




Figure 1.0 – General organization of a perceptron.<sup>324</sup>

In this three-phase schematic, initial inputs were generated via the Sensory System (S-System), which in the Cornell lab in 1958 meant the raster produced by a TV camera or photocell, meaning a *photoperceptron*. This information was next fed to the Association System (A-System) as a series of either excitatory (+) or inhibitory (-) signals. The A-System served as a ‘register,’ ‘counter’ or ‘memory’ between S-System inputs (ex. photocell data) and the outputs generated by the Response System (R-System), which would then be fed back to inhibit certain connections depending on their strength in relation to a threshold.<sup>325</sup> In practice, the A-System performed like a set of switching functions, balancing weighting internally between signals from the S-System and feedback from the R-System.

As had been the case on EPAC, the R-System of a perceptron accommodated a display feature—a series of lights—that an onlooker could use to tell the R-System whether the A-System’s internal weightings had judged an input correctly or not. A perceptron’s internal ‘memory’ would thus ‘learn’ via a process of statistical refinement informed by a series of sensory and discriminatory values from the S-System and R-System.<sup>326</sup> The first 1957 PARA report summarized a perceptron’s performance as, ‘A process of learning to give the same output signal (or print the same word) for all optical stimuli which belong to some arbitrarily constituted class.’<sup>327</sup>

<sup>324</sup> Rosenblatt, ‘The Perceptron: A Perceiving and Recognizing Automaton (Project PARA)’, 6.

<sup>325</sup> Rosenblatt, ‘The Perceptron: A Theory of Statistical Separability in Cognitive Systems (Project PARA)’, 46, 58, 231.

<sup>326</sup> The Response System feeds back which pulses to inhibit, which also refines the system.

<sup>327</sup> Rosenblatt, ‘The Perceptron: A Perceiving and Recognizing Automaton (Project PARA)’, 3.

Consider one hundred pieces of paper, each with an 'X' written at its centre. When each 'X' is held in front of the photocell, the S-System feeds this data to the A-System, which attempts to decipher, statistically from its 'memory,' whether the object it is viewing is an 'X' or not. Its determination is then fed to the R-System, where a series of lights go off and an onlooker can judge whether the system was correct or not. This determination is fed back to the A-System to inform its internal switching functionality, or 'memory,' which then reorients its weightings to improve the likelihood of deciphering the next 'X' correctly. After a sufficiently high number of successful instances, such as one hundred or more correct guesses, a perceptron was designed to have captured in its 'memory' a pattern of key features for the particular object it had been trained to recognize. An 'X,' for instance, might be recognized as having '/' and '\' as intersecting lines or '··' and '··' as intersecting points. Of key importance was that a perceptron would determine this distinguishing signature for itself. It could be a '··' a '\' or some other Mercurial feature set.<sup>328</sup>




Figure 2.0 – A more advanced perceptron schematic from 1957.<sup>329</sup>

Perceptron theory posited that the basis for associative learning lay in the power of the relatively simple memory function outlined above. Rosenblatt spent much of the next five years developing this theory. He conceded in each report that a major shortcoming of using a statistical approach to memory was that emergent errors would be *inherent* to the system, given the lack of any additional adjudicating system like one informed by formal logic or some other decision language. Exact reproducibility and errorless retention had to be sacrificed to

<sup>328</sup> In more complex cases, the quirks of this signature pattern would not be obvious to a human onlooker, since each result was the product of a statistical weighting.

<sup>329</sup> Rosenblatt, 'The Perceptron: A Perceiving and Recognizing Automaton (Project PARA)', 7.

maximize mathematical purity and minimize, or ‘economize,’ computational demands. Statistical error thus became a probability to reduce but not eliminate, in the same way that manufacturing errors could be reduced but never eliminated.

Jones situates Rosenblatt’s theory within one of two distinct camps in statistics to emerge during the second half of the twentieth century. The first originated in mathematical statistics. It was oriented around a generative or realist model of induction, in which ‘an understanding of how data are generated reflects an understanding of the corresponding law of nature.’<sup>330</sup> This was not the camp that Rosenblatt worked within. The second, which perceptron theory exemplified, pursued an instrumentalist or predictive model of induction, seeking functions that fit to the data, rather than functions that fit to a corresponding law of nature.<sup>331</sup> Jones cites Rosenblatt’s theory as the best known case in the Anglophone world of the instrumentalist school from the 1960-70s, with researchers in Japan and France rejecting this hard-line data-driven approach.<sup>332</sup>

In the opening pages of his first PARA report, Rosenblatt positioned his offering as ‘a black box.’<sup>333</sup> In adopting this frame, he emphasized the *internal* logics of statistical theory at the expense of whatever *social* logics would inevitably accompany his theory’s application in the real world. The all-important role of human-curated input data and the discriminatory capacity of the system’s human trainer were treated as *outside* his schematic(s). These inputs, and their role in shaping outcomes, were thus abstracted *out* of perceptron theory. This obfuscation, in turn, complicated the epistemic and ontological status of a perceptron by clouding how a human contributor directed circumstantial results. Ironically, it had been the social logic of employment, in this case derived from research for the U.S. Air Force, that initially led Rosenblatt to consider the mathematics of perception in the first place.<sup>334</sup> As I will now show, these nuances gained visibility as Rosenblatt engaged with the public and his peers in 1958 and afterwards.

---

<sup>330</sup> Jones, ‘How We Became Instrumentalists (Again)’, 678–79.

<sup>331</sup> Jones states that this instrumentalist branch of statistics developed independently in the U.S. and the Soviet Union but came together after the 1990s.

<sup>332</sup> Jones, ‘How We Became Instrumentalists (Again)’, 678–79.

<sup>333</sup> Rosenblatt, ‘The Perceptron: A Perceiving and Recognizing Automaton (Project PARA)’, 3.

<sup>334</sup> Küçük argues that historians of science remain needlessly uncritical of Western scholars’ university salaries, which permit unproductive labour and the modes of discovery it enables. ‘The Scientific Revolution follows not geniuses and academies but high professor salaries.’ Küçük, *Science without Leisure*, 90–94, 229.

## 1958: Rosenblatt's Statistical Approach to Brain Modelling

In early 1957, perceptron theory was primarily a theory of perception and efficient memory storage. The mechanization of related capacities, such as 'recognition,' 'concept formation' and the ability to generalize from 'experience' were also listed as foci.<sup>335</sup> Neither intelligence nor problem solving, however, were included on this list. This differentiated Rosenblatt's approach from related inquiries undertaken by McCarthy and Minsky. After the autumn of 1957, as Rosenblatt began simulations of his theory on an IBM 704, this point of difference dissolved.<sup>336</sup> In summer 1958, due perhaps to the validation of preliminary results and his new status as Principal Psychologist of the Aeronautical Laboratory, Rosenblatt elected to describe his project in relation to a more audacious capacity: mechanical 'intelligence.' In the laboratory's *Research Trends* magazine, he posited, as Simon and Newell had, that transformative scientific discovery lay in wait for those who could render neural activity using mathematics.

In this period, Rosenblatt's research also began to reach a wider public and scholarly audience. The two communities received his theories differently. In 'The Design of an Intelligent Automaton,' his article for *Research Trends*, Rosenblatt compared brain model research to physics prior to Newton; awaiting an integrated set of principles that, he boldly implied, existed.<sup>337</sup> Statistics, mathematics and digital computing would render these principles legible and knowable. Although Rosenblatt made no reference to 'artificial intelligence' in the article, his use of 'intelligent' and 'automaton' in the article's title indicated his broader alignment with radically empirical epistemological projects undertaken by Simon, Newell, Minsky and McCarthy. An automaton could be intelligent, the article declared.

Additional coverage of Rosenblatt's research in venues outside of scholarly journals and conferences did little to help to build sober understanding of his accomplishments. In July 1958, following an ONR-backed press conference, *The New York Times* wrote:

The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its

---

<sup>335</sup> Rosenblatt, 'The Perceptron: A Perceiving and Recognizing Automaton (Project PARA)', 1.

<sup>336</sup> Rosenblatt, 'The Design of an Intelligent Automaton', 4, 6.

<sup>337</sup> Rosenblatt, 1.

existence. Later perceptrons will be able to recognize people and call out their names and instantly translate speech in one language to speech and writing in another language, it was predicted.<sup>338</sup>

A subsequent article entitled 'Electronic "Brain" Teaches Itself' quoted Rosenblatt as having claimed that a perceptron could 'perceive, recognize and identify its surroundings without human training or control' and would one day be 'conscious' of their existence.<sup>339</sup>

As historians have chronicled elsewhere in connection to what has been called the Perceptrons Controversy of the late 1960s, this hyperbolic introduction to Rosenblatt's work left a sour taste in the mouth of his contemporaries.<sup>340</sup> Rosenblatt addressed the risk of confusion when interviewed for *The New Yorker* that fall. He criticized the 'loose talk' surrounding mechanical brains and provided, instead, a 'safe definition' of his theory and results. 'Our success in developing the perceptron means that for the first time a nonbiological object will achieve an organization of its external environment in a meaningful way,' he stated.<sup>341</sup>

In late 1958, Rosenblatt published his first two treatments of perceptron theory for a scholarly audience. Unlike in *The New York Times* and *Research Trends* magazine, he was careful in these venues to relate his work to the wider research community. In 'The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain,' published in *Psychological Review*, Rosenblatt juxtaposed his brain model to others developed by Rashevsky, McCulloch, Pitts, J.T. Culbertson, Kleene and Minsky. Once again, he characterised contemporaries' work as 'logical contrivances' in comparison to his model's fidelity to biological fact.<sup>342</sup> Unabashed, he argued that a fissure existed at the philosophical foundations of brain model theory:

---

<sup>338</sup> 'New Navy Device Learns by Doing'.

<sup>339</sup> 'Electronic "Brain" Teaches Itself'.

<sup>340</sup> In 1969, Minsky and Seymour Papert published *Perceptrons*, a mathematical take down of Rosenblatt's theory. They argued that a single-layer neural network could not be used to calculate parity and that the explanatory potential of such systems could not scale. In the late 1980s, neural network research underwent a resurgence. It became apparent to the community that Minsky and Papert had misrepresented the field as 'sterile.' See: Minsky and Papert, *Perceptrons*, 232. Minsky stated, 'Everybody seems to think I'm the devil.' As quoted in: Olazaran, 'A Sociological Study', 652. Guice and Olazaran relate Rosenblatt's lofty claims about his early work to this later scandal. See: Guice, 'Controversy and the State', 111–13; Olazaran, 'A Sociological Study', 621–23.

<sup>341</sup> Mason, Stewart, and Gill, 'Rival', 45.

<sup>342</sup> Rosenblatt, 'The Perceptron', 387. Published under ONR Contract Nonr-2381(00).

The proponents of this [opposing, deterministic] line of approach have maintained that, once it has been shown how a physical system of any variety might be made to perceive and recognize stimuli, or perform other brain like functions, it would require only a refinement or modification of existing principles to understand the working of more realistic nervous systems... The writer takes the position... that these shortcomings are such that a mere refinement or improvement of the principles already suggested can never account for biological intelligence; a difference in principle is clearly indicated.<sup>343</sup>

Despite his stated commitment to the epistemic primacy of biological fact in brain modelling, it is difficult to surmise from Rosenblatt's critique what exactly marked, for him, the threshold between fidelity and infidelity to biological phenomena. The paper listed a number of ways in which competing models failed to respect the complexity of biological phenomena: synchronization requirements, excessive specificity of connections, absence of equipotentiality, a lack of neuroeconomy; the list continued. Rosenblatt did not, however, structure these failings into a taxonomy that formalized a minimum threshold for justified biological analogy. Nor did he deeply interrogate competing frameworks in print.

Rather than deconstruct the work of others in detail, Rosenblatt turned to where he felt the field of brain modelling *should* head next. He called upon work by the psychologist Donald Hebb, the neuroscientist Peter Milner (Hebb's student), the neurophysiologist John Eccles and, most curiously given his subject matter, the economist and philosopher Friedrich Hayek. Each figure, he reasoned, conceived of the activity of the nervous system in relation to its natural environment, an orientation that steered toward biological realities, which Rosenblatt held as sacrosanct, and away from the trappings of the abstract logical 'contrivances' he laboured to gain distance from.<sup>344</sup>

Of this list of figures, Rosenblatt isolated Hebb and Hayek as having provided the richest accounts of this mind-environment relationship. In *The Organization of Behavior*, published in 1949, Hebb, a Canadian psychologist, outlined his theory of cell-assemblies; the

---

<sup>343</sup> Rosenblatt, 388.

<sup>344</sup> Core ideas included: that every nervous system was plastic, unique, generated largely at random, prone to self-organization around similarity and difference, and responsive to positive and negative reinforcement. Hebb, *The Organization of Behavior*; Milner, 'The Cell Assembly: Mark II'; Hayek, *The Sensory Order*; Hebb and Hayek are again cited in: Rosenblatt, 'The Design of an Intelligent Automaton', 7; Rosenblatt, 'A Probabilistic Model for Visual Perception', 296.

notion that reoccurring behaviours became hard-wired into the brain via repeated activity, or that ‘neurons wire together if they fire together,’ as another scholar later quipped.<sup>345</sup> Hayek had developed a similar perspective while studying law and psychology as an undergraduate at the University of Vienna. In an unfinished 1919 paper entitled ‘Contributions to a Theory of How Consciousness Develops,’ he hypothesized that external stimuli set neural linkages that strengthened or weakened in relation to use. Hayek saw consciousness as the expression of these trained connections. ‘One might almost say that each individual thinks with his past,’ he wrote.<sup>346</sup>

Bruce Caldwell characterizes this 1919 paper as the origin of Hayek’s long-brewing attack on the philosophical foundations of behaviourism and physicalism. Three decades later, in 1952, Hayek’s initial ideas received full treatment as *The Sensory Order: An Inquiry Into the Foundations of Theoretical Psychology*, which he described to a friend as ‘the most important thing I have yet done.’<sup>347</sup> The book landed with little fanfare in psychology; Rosenblatt’s interest was a notable exception in this respect.<sup>348</sup>

In *Sensory Order*, Hayek used market operations as a template to characterize neural activity. Both systems were of a scale of complexity that defied centralization. In the marketplace, neither buyers nor sellers could structure associations as competently as the decentralized distribution of processes across nodes. In the brain, classifications for a flood of stimuli formed and reformed passively and constantly, strengthening and weakening in the process in a manner akin to Hebb’s cell assemblies.

For Hayek, and later for Rosenblatt as well, behaviourism had failed to adequately characterize neural complexity. Real world stimuli-response relationships were defined by multiplicity and reflexivity, not simplicity or segregation. Rosenblatt challenged that brain model theory based in deterministic logics neglected the very complexity its proponents claimed to want to distil. Hayek’s theorization provided a convenient justification for this

---

<sup>345</sup> Hebb, *The Organization of Behavior*; Klein, ‘The Hebb Legacy.’; Lowel and Singer, ‘Selection of Intrinsic Horizontal Connections in the Visual Cortex by Correlated Neuronal Activity’, 211.

<sup>346</sup> Hayek had traced fibre bundles of the human brain while working in the laboratory of brain anatomist Constantin von Monakow. See: Caldwell, ‘Some Reflections on F.A. Hayek’s *The Sensory Order*’, 241.

<sup>347</sup> As cited in: Caldwell, 239.

<sup>348</sup> Hayek’s isolation from psychology is attributed to a lack of recognized expertise and his distaste of then-fashionable Keynesian ambitions to apply scientific rationales to social theory. See: Mirowski, *Machine Dreams*, 2002, 237.



interpretation. To Hayek, the human mind was not, 'A problem of the responses of the individual to an independently existing or objectively given phenomenal world.'<sup>349</sup> On the contrary, the mind existed in relation to the difference between the phenomenal world and the physical world.<sup>350</sup> Hayek and Rosenblatt assumed that this difference must be managed through relation and distribution, not centralized regulation. Rosenblatt's notion of neural efficiencies thus built on Hayekian economics, just as Simon's theory of problem solving had built on his positivist theory of administration.

In November 1958, Rosenblatt presented his theory at the 'Mechanisation of the Thought Processes' conference at the National Physical Laboratory (NPL) in Teddington, London.<sup>351</sup> Teddington's composition, far more than Dartmouth, revealed the plurality of state and industrial stakeholders keen to realize the mechanization of thought in the mid-century. The symposia convened two hundred delegates split evenly across industry, government and academia.<sup>352</sup> The Mathematics Division at the NPL had been established in 1945 to provide computational services to all departments of government, a mandate that had enabled Alan Turing's development of the Pilot ACE computer, 1945-48. This mandate extended the NPL's role, dating back to 1900, as a coordination body for the standardisation and regulation of scientific measurements across the Empire.<sup>353</sup> At the event, cyberneticists like W. Ross Ashby and Donald MacKay gave papers alongside psychologists like Horace Barlow and Richard Gregory, and leading computer engineers such as Grace Hopper and John Backus, author of FORTRAN, the only programming language to predate LISP, the *lingua franca* of AI in the 1960-70s. Minsky and McCarthy also spoke. Foci encompassed academic and industrial subjects, such as artificial thinking, automatic programming, industrial planning and clerical mechanization.<sup>354</sup>

---

<sup>349</sup> As quoted in: Caldwell, 'Some Reflections on F.A. Hayek's The Sensory Order', 247.

<sup>350</sup> As quoted in: Caldwell, 'Some Reflections on F.A. Hayek's The Sensory Order', 247. This conviction was informed by Ernst Mach's view that science formed the economy of thought. Later, Hayek identified von Neumann's unfinished theory of automata as a far more mathematical corroboration of this general perspective. See: Mirowski, *Machine Dreams*, 238-39.

<sup>351</sup> Rosenblatt, 'Two Theorems of Statistical Separability in the Perceptron', 1959.

<sup>352</sup> Sutherland, 'The Mechanization of Thought Processes'.

<sup>353</sup> *The Future of Artificial Intelligence: Views from History*. See Simon Schaffer: 12:15-15:10.

<sup>354</sup> Agar notes that operations research fell out of favour in the UK in the 1950s due to practitioner's vague multi-disciplinary claims to expertise, which were judged with some

At Teddington, Rosenblatt posited that mental behaviour was fundamentally intuitive rather than analytical. He cited *The Computer and the Brain*, von Neumann's final text prior to his death in 1957.<sup>355</sup> Von Neumann had laid the blueprint for digital electronic computing in 1945 and intended this book to do the same for artificial thought. He argued that logic and mathematics were 'secondary' languages of the brain generated by some yet-to-be-understood 'primary' language that physiological evidence suggested would be probabilistic in nature.<sup>356</sup> Rosenblatt seized on this argument, which echoed his view that symbolic logic and Boolean algebra were too inflexible to convincingly model brain activity 'directly,' as if by reflex.<sup>357</sup>

To realize the potential of this avenue, Rosenblatt proposed to alter the base-unit of a neural network from a simple binary building block to a complex one: the perceptron. He claimed that a network comprised of perceptrons as its base-units afforded sophisticated associative learning by mapping correlations between *series* of stimuli. The paper presented two theorems to this effect. The first outlined in mathematical terms how a human operator simulated a perceptron on an IBM 704 to recognize differences between letters of the alphabet, geometric forms and positions on the retina.<sup>358</sup> This demonstrated an 'intuitive' theory of cognition in action, since decisions required no human-crafted logical circuitry in order to function.

The second theorem demonstrated mathematically how the same system accomplished this result *on its own* by trending its memory towards responses that grouped in similar ways. An unsupervised perceptron, he explained, would 'Tend... to arrive at a "useful" division of its environment, without human intervention.'<sup>359</sup> This amounted to a coherent new approach to the design of neural networks.

Rosenblatt did not clarify who would decide what was 'useful.' In his conclusion, he stated that he would avoid speculation on how his theory, and its derivative uses, would enter

---

suspicion in government. Agar, *The Government Machine*, 250. In America, operations research was embraced to bolster neo-classical economics. See: Mirowski, *Machine Dreams*, 177–90.

<sup>355</sup> Von Neumann, *The Computer and the Brain*.

<sup>356</sup> As cited in: Rosenblatt, 'Two Theorems of Statistical Separability in the Perceptron', 1959, 421.

<sup>357</sup> Rosenblatt, 422.

<sup>358</sup> Rosenblatt, 432–33.

<sup>359</sup> Rosenblatt, 441.

the world. This was a disingenuous position to take given that the collation of military intelligence had been named explicitly as one of theory's proposed applications, along with language translation, induction and concept formation.<sup>360</sup> Each of these uses would benefit PARA's military underwriters. Rudolf Seising adds that perceptron theory was used by the U.S. Air Force to locate ships in aerial photography, a domain that Rosenblatt was not unaware of, given his 1955 'Parallax' paper.<sup>361</sup> Perceptron theory's immediate use was thus not neutral but malicious; it provided military underwriters with a means to divide and categorize the world in keeping with how they already viewed it.

Rosenblatt's hesitation to acknowledge potential applications was not accompanied by any hesitation about his theory's scientific utility. He concluded his Teddington paper by asserting that, 'As a concept, it would seem that the perceptron has established, beyond doubt, the feasibility in principle of nonhuman systems which may embody human cognitive functions at a level far beyond that which can be achieved through present day automata.'<sup>362</sup> He framed the principles underlying the future of information processing as definitively *statistical*, not logical.

Conference attendees questioned this enthusiasm. Edward Newman, who shaped the ACE computer at NPL during the 1940s, noted that he had co-developed a related system years earlier, which undermined Rosenblatt's claim of a lack of past precedents other than from von Neumann. Stafford Beer, likewise, pointed to similar, early work by A. M. Uttley, then in attendance. Beer pressed Rosenblatt for proposing a false dichotomy between Boolean algebra and statistical mathematics and for posturing an 'elaborate mathematical edifice' via 'exaggerated' claims about the 'ethos of potency' surrounding perceptron theory.<sup>363</sup> John McCarthy contested that there was a clear difference between discrimination and description in perception and that perceptron theory was incapable of the latter. Description involved having a system *describe* something that it has seen; discrimination involved the ability to *learn* from a series of stimuli. Perceptron theory only explored discrimination, not description. Rosenblatt conceded each of these claims.

---

<sup>360</sup> Indeed, this was the report's only mention of 'intelligence.'

<sup>361</sup> Seising, 'A Brain Model for the Perception of the Outside World'.

<sup>362</sup> Rosenblatt, 'Two Theorems of Statistical Separability in the Perceptron', 1959, 449.

<sup>363</sup> Rosenblatt, 463.

Boden describes Teddington as a ‘Memorable catalyst for the growth of an intellectual community.’<sup>364</sup> At the event, McCarthy, Minsky, Selfridge *and* Rosenblatt (who had not been invited to Dartmouth) all presented their earliest pivotal research efforts—papers that would come to define the field—alongside and in debate with McCulloch, MacKay, Pask, Ashby, Backus, Hopper and other recognized figures in American information science. While popular historical accounts emphasize a contrast between Rosenblatt’s early approaches and the work of others in the early symbolic-AI programme, I argue that the substance of his peers’ critique at Teddington reflects, conversely, that perceptron theory was largely in line with the *broader* community then interested in how to mechanize thought processes. Although Rosenblatt sought to distance himself from determinist logics and the beginnings of the symbolic programme, he did not challenge the premise that the brain *could* be modelled, statistically or otherwise, using mathematics. Nor, like his peers, did he explicitly consider the myriad ways in which his work could eventually be used by industrial and state actors in attendance.

#### 1958-71: On the ‘Second Direction’ of Brain Modelling and its Military Suitors

The composition of Teddington speaks to Rosenblatt’s alignment with a community of industrial, government and academic actors oriented around similar goals. In print, Rosenblatt identified only with the last of these groups. For Rosenblatt, the study of machine intelligence remained an academic rather than industrial vocation, one in which he protected the primacy of biological evidence against those who would trivialize it. By the late 1950s, this dichotomy was repeated in colleagues’ work. In 1959, McCarthy published a paper in the *New York Herald Tribune Engineers’ News Supplement* arguing that research into machine intelligence progressed in *two* directions, one into ‘nerve nets’ trained to discriminate between mosaics of dots on receptor cells (à la Rosenblatt), and the other into ‘artificial intelligence,’ meaning the techniques used to ‘solve problems “requiring human intelligence” without trying to imitate the structure of the brain.’<sup>365</sup>

---

<sup>364</sup> Boden, *Mind as Machine*, 336. For a survey of neural network research in this period see: Olazaran, ‘A Sociological Study’, 618.

<sup>365</sup> McCarthy, ‘Getting Closer to Machines That Think’.

Rosenblatt championed this polarity. In *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*, published as a technical report for the Cornell Aeronautical Laboratory in 1961 and then in 1962 as his first book, Rosenblatt positioned those in AI research as ‘the loyal opposition.’<sup>366</sup> In a 1962 paper for the American Institute of Electrical Engineers, he framed brain model theory as a choice between Galilean experimental methods and Aristotelian rationalism.<sup>367</sup> Twentieth century engineering had leveraged the Aristotelian postulate to great effect, he argued, citing abstract logical representations as having delivered data processing networks. Yet pure reason alone would not clarify the principles of *neural* dynamics. Statistical models were uniquely able to inform understandings of the dynamics of the mind, particularly those derived from biological evidence.

Rosenblatt’s book emphasized, from its outset, the intellectual distance he perceived between his work and that of his ‘opposition.’ He used the distinction to bolster the significance of perceptron theory, which he argued was motivated by a want to advance scientific understandings of natural intelligence, rather than to toy with engineering inventions like his contemporaries in AI. This supposed dichotomy obscured the two domain’s similarities. Tellingly, Rosenblatt concluded *Neurodynamics* by conceding that perceptron theory would merge, eventually, with Simon and Newell’s work on heuristic programming.<sup>368</sup> The two approaches were not, in fact, incommensurable; perceptron models clarified the ‘eyes and ears’ of psychological systems, while heuristic programming clarified goal-motivated behaviour. Nascent learning and problem-solving methodologies could thus, in his view, be rectified and integrated as they matured.

This conclusion reveals an Aristotelian dimension to Rosenblatt’s perceptron project. While he held up biological fidelity as an ideal, his perceptron theory permitted gross simplifications in order to yield an analysable brain model, as seen in his reductive use of flow-charts and instrumentalist statistics. His justification for these simplifications was that, without them, ‘the number of possible connection diagrams becomes, for all practical

---

<sup>366</sup> Rosenblatt, ‘Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms’; Rosenblatt, *Principles of Neurodynamics*, viii, 10.

<sup>367</sup> Rosenblatt, *Principles of Neurodynamics*, vii–viii; Rosenblatt, ‘Analytical Techniques for the Study of Neural Nets’. Rosenblatt neglected continuities *between* these scholastics, as approached in: Küçük, *Science without Leisure*.

<sup>368</sup> Rosenblatt, *Principles of Neurodynamics*, 584.

purposes, infinite.<sup>369</sup> Only statistical techniques could lead scientists to the edge of discovery by confirming or falsifying ‘at just the point where the complexity of the known systems begins to make further theoretical speculation impractical.’<sup>370</sup> As promising as it sounded, he did not clarify where exactly this inflection point fell.

This omission speaks to a deterministic aspect of perceptron theory. In keeping with Hayek’s notion of market dynamics, Rosenblatt’s statistical methodology swapped the authority of centrally determined routines like Aristotelian formal logic for a different sort of authority, namely a set of decentralised determined routines. Whatever sensory inputs fed the S-System of a perceptron determined, albeit probabilistically, its output. Rosenblatt’s justification for this switch was that it improved fidelity to biological fact. Yet, even in aggregate, these routines were still beholden to a certain systemic subjectivity. ‘It is likely that we must always be satisfied with approximate answers obtained by a careful selection of experiments, rather than by either a theoretical analysis or an exhaustive examination of the alternatives,’ Rosenblatt explained.<sup>371</sup>

Hayek had advocated for a similar approach in *Sensory Order*. He argued that complex systems marked a challenge to the advance of scientific inquiry, one that *necessitated* the use of new methodologies. ‘While it is certainly desirable to make our theories as falsifiable as possible, we must also push forward into fields where, as we advance, the degree of falsifiability necessarily decreases. This is the price we have to pay for an advance into the field of complex phenomena,’ he argued.<sup>372</sup> Rather than formalising a top-down deterministic system, the two figures advocated developing a bottom-up deterministic system, one decided by the content of its inputs.

Philip Agre, a practitioner turned critic of AI, explicates that, in Hayek’s case, this ideal of autonomous decentralization rested upon the existence of unacknowledged institutional structures.<sup>373</sup> Hayek’s libertarian economy prefigured and thus implicitly required a robust institutional framework so that ‘the miracles of self-organized complexity’ would reliably

---

<sup>369</sup> Rosenblatt, ‘Two Theorems of Statistical Separability in the Perceptron’, 1959, 422.

<sup>370</sup> Rosenblatt, ‘Analytical Techniques for the Study of Neural Nets’, 290.

<sup>371</sup> Rosenblatt, 291.

<sup>372</sup> As cited in: Caldwell, ‘Some Reflections on F.A. Hayek’s *The Sensory Order*’, 250.

<sup>373</sup> On autonomy and technology, see: Winner, *Autonomous Technology*.

occur.<sup>374</sup> Similarly, Rosenblatt's system presumed and thus implicitly required the existence of a human operator and a set of novel input stimuli. Each was structurally integral to his theory, yet both were treated as ancillary to the validity of the 'black box' statistical mechanisms.

Rosenblatt had overlooked these considerations during his PhD as well. He had collected large stores of deeply sensitive information about Cornell undergraduates in order to refine his k-coefficient. He considered this infrastructure 'incidental' in relation to the scientific insights it could supply. This sense of entitlement was not out of place in the heady post-war research context, in which social psychology research was direct funded and developed by the U.S. military to, on the whole, 'stiffen the ideological supports of the men.' By treating this experimental infrastructure as fundamentally distinct from his theory, Rosenblatt misrepresented the character of his statistical work. He cast perceptron theory as perception without judgement; as if statistics had an eye of its own (when in fact it had to eye at all). His use of flow-charts gave form to this pledge; Kathryn Henderson describes flow-charts as their own 'conscription devices,' a boundary object capable of simplifying complex phenomena down the point of presumed consensus, even if mutual understanding is not met.<sup>375</sup>

These decisions occurred against a backdrop in which Rosenblatt's career benefited from military patronage. In 1959, after a failed bid to join the Brookhaven National Laboratory, he reassessed his status at Cornell. To his good fortune, U.S. military support helped to secure his continued placement there. In June, the Institute for Defense Analysis awarded his group a \$10,000 contract upon his relocation from the lab in Buffalo to Cornell's central campus in Ithaca. By September 1961, the ONR provided his project an additional \$153,000 worth of contracts, with \$108,000 committed for 1962.<sup>376</sup> 'It was largely through the pressure of the ONR that the university acknowledged that such a program did have a place at Cornell, that a course should be taught in the general theory of brain mechanisms,'

---

<sup>374</sup> Agre, 'Hierarchy and History in Simon's "Architecture of Complexity"'. On the role of institutions in state building: Acemoglu and Robinson, *Why Nations Fail*.

<sup>375</sup> As cited in: Ensmenger, 'The Multiple Meanings of a Flowchart', 324, Ensmenger shows how flowcharts developed in industry in 1960-70s were often overly simplistic and optimistic in their intent, see: 335.

<sup>376</sup> Wright to Atwood, 'Perceptron and Allied Matters - Frank Rosenblatt', 24 November 1959.

Rosenblatt recalled, equating the help to ‘aiding and abetting’ his eventual admission in Ithaca under ‘oddball status.’<sup>377</sup>

In 1959, Rosenblatt became Director of the Cognitive Systems Research Program and Lecturer in the Department of Psychology as well as the head of the Cognitive Systems Section at the Aeronautical Laboratory.<sup>378</sup> He described his status as that of a ‘free-floating administrative entity’ that no single department wanted yet many recognized as important.<sup>379</sup> A 1959 memo by Cornell’s Vice President of research cited Rosenblatt as listing overlap with numerous disciplines: neurochemistry, electrophysiology, radiation biology, genetics and embryology.<sup>380</sup> A 1964 memo to Cornell’s Provost revisited how best to integrate him into the university.<sup>381</sup> Cornell’s Vice President for Research deemed Rosenblatt worthy of tenure consideration but cautioned that there was ‘no obviously suitable home for him’ within existing departments. Psychology, where he had trained, would need to broaden its interests to integrate his mechanical interests; electrical engineering and systems analysis lacked his biological interests; the Veterinary College and Applied Mathematics Center were considered but not decisively. Ultimately, in 1966, Rosenblatt became an associate professor in the Neurology and Behaviour section of the newly formed Division of Biological Sciences, a section he then chaired until his death in a boating accident in 1971.<sup>382</sup>

As he transitioned from early to mid-career status, Rosenblatt reflected on scientists’ accountability to the wider world. He opened *Neurodynamics* with an apology for the manner in which perceptron theory had initially been publicized in 1958.<sup>383</sup> In a private 1963 letter to the *Chicago Tribune*, he reprimanded a journalist for embellishing that an attached red light signified that the perceptron was blushing, as if emotionally aware. No similar act of self-awareness about accountability seemed to have followed the two embellished *New York Times* articles in 1958.<sup>384</sup> Into the 1960s, Rosenblatt also questioned the materiality behind his statistical methods. In 1962, he reflected on whether computation would be sufficient for

---

<sup>377</sup> Rosenblatt, Interview with Frank Rosenblatt, 114.

<sup>378</sup> Rosenblatt and Fuch, ‘Proposal to the National Science Foundation for Joint Support of the Cognitive Systems Research Program, 1962-1967’, 41.

<sup>379</sup> Rosenblatt, Interview with Frank Rosenblatt, 115.

<sup>380</sup> Wright to Atwood, ‘Perceptron and Allied Matters - Frank Rosenblatt’, 24 November 1959.

<sup>381</sup> Long to Corson and Keast, ‘Procedures for Incorporating Frank Rosenblatt’, 11 March 1964.

<sup>382</sup> ‘Cornell Professor Dies’.

<sup>383</sup> Rosenblatt, *Principles of Neurodynamics*, vii.

<sup>384</sup> Rosenblatt, ‘To Publisher, Chicago Tribune’, 21 June 1963.



progress in psychology, given that the complexity of observed neural activity surpassed  $10^{12}$  to  $10^{14}$  parameters, which no machine could expect to parse.<sup>385</sup>

Despite this reflection, Rosenblatt appears not to have questioned whether statistical inquiry was itself insufficient to account for the material differences between natural and artificial systems in the 1960s. He was not alone. Jones characterizes the take up of computational statistics in pattern recognition in the 1960-70s as 'less an academic discipline than a cluster of like-minded practitioners oriented around common sets of goals.'<sup>386</sup> This cluster favoured practical uses of predictive capabilities over techniques that explained causal neural dynamics or some other empirically-motivated insight.<sup>387</sup> Although Rosenblatt would go on to engineer a large-scale perceptron, the Tobermory system, between 1962-64, he also turned his focus directly to materiality in the mid-to-late 1960s when exploring the transfer of learned behaviours in rats via the extraction and transfer of neural chemistry.<sup>388</sup> Olazaran explains that this turn was in part due to funders' diminished faith in neural networks; by the late 1960s, DARPA had committed to symbolic AI and the ONR had come to doubt perceptron theory.<sup>389</sup> Rosenblatt was reported

On campus at Cornell, Rosenblatt became a well-known anti-war and anti-racism activist.<sup>390</sup> He welcomed five or six students at a time to live with him at his farm home.<sup>391</sup> Following the armed occupation of the Cornell student union in 1969, a product of the prolonged disenfranchisement and glib appeasement of black students, he helped to establish Constituent Assemblies at Cornell, a governance structure that increased student representation in university decision making. At the national level, he contributed to Eugene McCarthy's 1968 presidential bid for the Democratic Party. He organized faculty support at Cornell, helped to publish McCarthy's first national ad and took months off of his academic research to volunteer alongside the politician's national staff, including at the famously

---

<sup>385</sup> Rosenblatt, 'Analytical Techniques for the Study of Neural Nets', 291.

<sup>386</sup> Jones, 'How We Became Instrumentalists (Again)', 677.

<sup>387</sup> Jones contrasts this highly mathematical orientation to statistics with practices in Japan that refused high levels of abstraction in favour of preserving the significance of more localized, holistic notions of statistical evidence.

<sup>388</sup> See, for example: Rosenblatt, Farrow, and Herblin, 'Transfer of Conditioned Responses'.

<sup>389</sup> Olazaran, 'A Historical Sociology of Neural Network Research', 174.

<sup>390</sup> 'Prof. Rosenblatt Dies'. See also: Rosenblatt, Interview with Frank Rosenblatt.

<sup>391</sup> 'Tribute to Dr. Frank Rosenblatt'. Personal correspondence with George Nagy.

violent Democratic national convention in Chicago.<sup>392</sup> These connections were familial; his brother, Maurice Rosenblatt, a lobbyist in Washington D.C., has been credited with having instigated and orchestrated the successful censorship of Senator Joseph McCarthy in 1954.<sup>393</sup>

When Rosenblatt died in 1971, three congressmen made a rare tribute to his memory from the floor of the U.S. House of Representatives, including former Senator Eugene McCarthy.<sup>394</sup> Curiously, his obituary credited him for having developed a new technique for the application of computer programming to political statistics, a development that might be regarded as galvanizing the latent power that he had held, implicitly, after amassing hundreds of personality profiles for his PhD.<sup>395</sup>

## Conclusion

The available historical literature on Frank Rosenblatt has tended to interpret the significance of his research as *he* framed it. Rosenblatt is cast as having epitomized one of two camps of mid-century brain model theory, the other being an ‘in group’ that would go on to study AI at the Massachusetts Institute of Technology (Minsky), Stanford University (McCarthy) and Carnegie Mellon University (Simon).<sup>396</sup> Rosenblatt supported this dichotomy when characterizing his approach in contrast to a ‘loyal opposition’ working on formal logic and heuristics. As Plasek rightly argues, there is good reason to decouple historical understandings of AI from neural network and machine learning research; if measured by volume of published material, AI would be judged as a ‘disciplinary backwater’ of machine learning.<sup>397</sup> I agree and add to this call for a thorough re-evaluation of the field that the fashions of disciplinary categories have their limits as historiographical frames. A pivot from AI to machine learning could simply make Rosenblatt the ‘in’ crowd, and machine learning

---

<sup>392</sup> Rosenblatt, Interview with Frank Rosenblatt, 117.

<sup>393</sup> Maurice Rosenblatt also co-founded The National Committee for an Effective Congress with Eleanor Roosevelt in 1948: Scates, *Maurice Rosenblatt and the Fall of Joseph McCarthy*.

<sup>394</sup> Simonds, ‘Cornell Professor Eulogized in House’; ‘Tribute to Dr. Frank Rosenblatt’.

<sup>395</sup> I have been unable to determine the nature of this technique, nor the substance of his likely related involvement in research on the Indochina War at the Cornell Center for International Studies. ‘Frank Rosenblatt Dies; Researcher, Politician’.

<sup>396</sup> Gardner, *The Mind’s New Science*, 30–31.

<sup>397</sup> Plasek, ‘On the Cruelty of Really Writing a History of Machine Learning’, 6.

the disciplinary frame, rather than centering on Minsky, McCarthy and Simon and their pursuit of AI.

I have argued that this dichotomy unnecessarily restricts understandings of what the two fields had in common. Rosenblatt challenged AI on the grounds that it failed to inform theories of natural intelligence by suggesting that neural activity was deterministic in some important sense, which biological evidence did not support. He simultaneously cast natural intelligence as susceptible to an instrumentalist brand of *statistical* modelling, a proposal he rooted not just in biological evidence on the dynamics at work in human visual perception, but also in Hayek's notion of decentralized market activity, which provided a corollary for the otherwise intractable behaviour of another complex phenomenon: economies. This basis for perceptron theory aligns Rosenblatt with Simon, McCarthy, Minsky and other mid-century researchers who used social metaphors to model neural behaviour. In addition to this commonality, each of these men believed that neural dynamics could be made legible with abstract mathematics and digital computing machinery, albeit in different ways. Historical focus on disciplinary boundaries alone can draw scrutiny away from the origins of this profound proposition. Each man's early career also benefitted from significant military and industrial patronage, which helped to sustain their theoretical inquiries in the 1950s and afterwards.

## Chapter Four – Applied Epistemology: On John McCarthy’s Account of Programmed Language as Intelligence, 1952-59

The advent of digital electronic computing in the United States and Britain in the mid-to-late 1940s and 1950s provided industrial, academic and government actors with the means to organize both information and people in radically new ways. Between 1952-57 in particular, new computational mechanisms like ‘assemblers’ and ‘compilers’ consolidated the prohibitively laborious instruction sets used to ‘program’ a digital computer into ever more accessible and coherent programming ‘languages.’ In the U.S., industrial actors like IBM advanced and leveraged these consolidations to usher new paying audiences into standardized modes of computing.<sup>398</sup> The medium’s diverse capabilities were iterated upon in an eclectic yet cumulative fashion by an ever-growing network of contributors connected through a common attendance to code. So began a formidable cycle, either vicious or virtuous depending on one’s perspective, in which increasingly sophisticated tools developed, in part, by for-profit and military research bodies became deeply infused into social practices. While historians like Nathan Ensmenger and Mar Hicks have explored how computing was used to reconfigure workplace structures in ways that disproportionately benefitted men over women, connections between the maturation of programming languages and understandings of the human mind remain less well understood.<sup>399</sup>

Amidst the initial wave of experimentation between 1952-57, John McCarthy, an American mathematician, co-introduced the term artificial intelligence (AI) with Marvin Minsky, Claude Shannon and Nathaniel Rochester. In 1955, under that banner, the group collectively advanced the conjecture that ‘every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.’<sup>400</sup> In this chapter, I argue that the maturation of computer programming after 1952 lent plausibility to AI that shored up its perceived legitimacy between 1955-1961. To show this, I distinguish between McCarthy’s uniquely positivist vision for artificial intelligence as a

---

<sup>398</sup> On the labour required from repairmen to maintain the visage of standardization in computing see: Orr, *Talking about Machines*.

<sup>399</sup> Ensmenger, *The Computer Boys Take Over*; Hicks, *Programmed Inequality*.

<sup>400</sup> Minsky et al., ‘Proposal’.

*discipline*—informed by his background in mathematics, I show, and his desire to reduce epistemology to a branch of applied mathematics—and the pragmatics of mid-century programming as an emerging *profession*, which military, commercial and educational backers nurtured with technological consolidation and standardization.

My story follows McCarthy, the figure commonly cited with coining AI and initiating formative events along with foundational theory. I bring unexamined archival material to light from McCarthy's early career in order to connect his intellectual aspirations for AI to entangled developments in automatic programming. During the mid-to-late 1950s, McCarthy and his interlocutors colloquially described automatic programming and AI in functionally similar terms. McCarthy simultaneously claimed that emerging programming techniques would reduce epistemology to a branch of applied mathematics. Popular accounts of McCarthy's research assess his legacy through the prism of this aspiration. He is memorialized as a co-founder of what is now called symbolic AI; a subfield of contemporary AI premised on the assumption that intelligence stems from the manipulation of symbolic representations. I contend that these accounts rationalise McCarthy's desired *ends* at the expense of the complex *means* he used to advance this audacious proposition. Such accounts overlook the significant degree to which AI benefitted from developments in programming undertaken for altogether different purposes.

The structure of this chapter is as follows. To begin, I consider McCarthy's formal training (1944-51) and publication record (1951-58) to situate his definition of intelligence in relation to his early career research on logic, partial differential equations and automata theory. I contrast McCarthy's positivist inclinations with the views of Claude Shannon and John von Neumann, two authoritative figures in post-war American information science who took pains to clarify the limitations of their own analogies between natural and artificial information processing systems. I weigh these competing views alongside iterative developments in programming in the early 1950s, such as the introduction of compilers, assembly languages and the first commercial stored-program computers, which reduced the labour and expertise required to operate a computer and standardized techniques that industry, academia and government could take up in earnest.

In the mid-1950s, McCarthy benefitted from this boom in innovation. After surveying state of the art automata research as co-editor of *Automata Studies* with Shannon between 1953-55, he worked to build a research community around a more audacious vision of

machine intelligence in 1955 and afterwards. He believed that the key to progress lay in the study and formalization of language. I relate this aspiration to those of industrial organizations like IBM, whose financial and infrastructural support both advanced the professionalization of programming through acts of commercial standardization and helped McCarthy at various major milestones in his career, including at the Dartmouth workshop that formalized the term AI and helped to galvanize its earliest membership—a process I chronicle in new detail. I argue that AI's fabled origin moment takes on less significance when seen in relation to these broader facilitating trends in industrial information and human management.

### 1944-55: The Road to Automata

John McCarthy was born in Boston in 1927 to Ida Glatt, a Lithuanian Jewish immigrant trained in political economy, and John Patrick McCarthy, a literary Irish Catholic immigrant with a fourth-grade education.<sup>401</sup> Both championed the rights of the oppressed through formal participation in the Communist Party. It was in this vein of Irish, Jewish, and Marxist radicalism that McCarthy was first introduced to math and science. When he developed a life-threatening sinus condition as an adolescent, the family relocated to Los Angeles, where he taught himself advanced calculus and applied to study mathematics at the California Institute of Technology (Caltech) in 1944.

At Caltech, McCarthy honed his understanding of logic, mathematics and the nascent field of digital electronic computing. In 1946, at age nineteen, he outlined a book on, 'Dialectical Materialism for the Scientist.'<sup>402</sup> Dialectical materialism is the philosophy of science that underpins Marxism. It holds that matter is the essence of all reality, including the mind. Within this tradition, McCarthy saw reason to pursue the 'application of formal logic to science,' a positivist intuition he would foster throughout his career. Remarkably, he included in his plans hand-drawn schematics for an electronic computer. The design included cathode

---

<sup>401</sup> Biographical details derive from McCarthy's personal papers and Hilts, *Scientific Temperaments*, 197–287; Cohen, *The Know-It-Alls*; Markoff, *Machines of Loving Grace*; Nilsson, 'John McCarthy, 1927-2011'; Andresen, 'John McCarthy'; Cherry, 'Remembering John McCarthy [Spectral Lines]'; McCarthy and Lifschitz, *Formalizing Common Sense: Papers by John McCarthy*; McCarthy, *Oral History of John McCarthy*; McCarthy, *Oral History: John McCarthy*; O'Dowd, 'This Son of an Irish Emigrant Invented Artificial Intelligence'.

<sup>402</sup> McCarthy, 'Dialectical Materialism for the Scientist'.

ray tubes and a quantum amplifier. In 1946, the field of computing remained esoteric; John von Neumann's classified 'First Draft of a Report on the EDVAC' had only just offered a blueprint for *digital* electronic computing to a small set of experts. Even by the early 1950s, there were no more than a dozen electronic computers in operation in the United States.<sup>403</sup>

Two events in 1948 provided McCarthy with insight into the emerging field. First, as a graduate student in mathematics, he attended lectures about the creation of the Standards Western Automatic Computer at the University of California.<sup>404</sup> Since the machine remained under construction until 1950, at which point he had moved on to Princeton, McCarthy was not able to witness the device in operation. He could, however, learn about its internal architecture.<sup>405</sup> The second event was the Hixon Symposium on Cerebral Mechanisms in Behavior, an influential conference convened at Caltech in September 1948 to debate theories of human behaviour.<sup>406</sup> The Symposium convened leading psychologists such as Heinrich Klüver, Wolfgang Köhler and Karl Lashley, as well as Warren McCulloch, von Neumann and others. McCarthy later credited von Neumann's presentation for having 'triggered' his interests in the analogy between natural and artificial information systems.<sup>407</sup>

In his Hixon talk, entitled 'The Logic of Analogue Nets and Automata,' von Neumann made explicit the limitations of his mathematical analogy between brain-like and computer-like information processing systems.<sup>408</sup> He outlined an axiomatic, black-box interpretation of biological fact that allowed for mathematical analogies to be made between the two domains, which he called automata. He positioned his theoretical work as that of an 'outsider' prohibited from meaningful speculation on anything other than the dynamics of *idealised* versions of elementary physiological units.<sup>409</sup> He conceded that in order for his methodology to exemplify rather than omit biological fact, logic and mathematics must inevitably undergo

---

<sup>403</sup> Nofre, Priestley, and Alberts, 'When Technology Became Language', 50.

<sup>404</sup> Cohen, *The Know-It-Alls*.

<sup>405</sup> Huskey, 'SWAC-Standards Western Automatic Computer', 53; Alexander, 'The National Bureau of Standards Eastern Automatic Computer', 84; Hutchins, *First Steps in Mechanical Translation*, 4.

<sup>406</sup> 'First Hixon Symposium on "Cerebral Mechanisms and Behavior"'.

<sup>407</sup> Nilsson, *The Quest for Artificial Intelligence*, 77.

<sup>408</sup> Published as, 'The General and Logical Theory of Automata.' See: Von Neumann and Jeffress, 'The General and Logical Theory of Automata'.

<sup>409</sup> Von Neumann and Jeffress.

‘a pseudomorphosis to neurology.’<sup>410</sup> In von Neumann’s framework, formal logic would come to embody the logic seen in natural systems—less rigid, combinatorial and all-or-none—in order to account for realities of nature such as the passing of time, which formal logic existed outside of and thus made no affordances for.<sup>411</sup>

In 1949, McCarthy followed von Neumann to Princeton to undertake a PhD but diverged with him over the limits of logic. An elder McCarthy claimed in interviews to have spent this period testing how a finite automaton representing the brain could be made to interact with a finite automaton representing its environment. This subject would have led him close to von Neumann’s unfinished theory of automata, which examined how a finite automaton could self-reproduce to create something more complex than itself. McCarthy claimed that, despite encouragement from von Neumann, he soon swore off this project due to reservations about whether it would show how the brain represented knowledge *explicitly* as he believed a theoretical computer language could one day be trained to do.<sup>412</sup> Since McCarthy did not publish his work on automata from Princeton, it is difficult to corroborate this claim.

What McCarthy’s early interest in automata had in common with his PhD research was that each mathematical domain could be used to probe and analyse the nature of otherwise intractably large systems. His twenty-eight-page dissertation ‘Projection Operators and Partial Differential Equations,’ completed in May 1951, developed a method for approximating solutions of differential equations based on the intersection of spaces with simpler properties.<sup>413</sup> The method examined how a large problem could be solved through the use of division and iteration. Solomon Lefschetz, a specialist in algebraic topology, supervised.

In 1952, McCarthy was invited to spend the summer at Bell Labs working on communication theory under Claude Shannon and alongside Marvin Minsky.<sup>414</sup> ‘Bell Labs had

---

<sup>410</sup> Von Neumann and Taub, ‘The General and Logical Theory of Automata’, 311.

<sup>411</sup> He cited Ludwig Boltzmann’s theories on thermodynamics as representative of this course. Von Neumann and Jeffress, ‘The General and Logical Theory of Automata’, 16.

<sup>412</sup> Nilsson, ‘John McCarthy, 1927-2011’, 3.

<sup>413</sup> McCarthy, Oral History: John McCarthy; McCarthy, ‘Projection Operators and Partial Differential Equations’.

<sup>414</sup> Unknown, ‘Letter to Mr. Marvin L. Minsky from Bell Laboratories, June 6, 1952’, 6 June 1952; McCorduck, *Machines Who Think*, 2004, 122; Nilsson, ‘John McCarthy, 1927-2011’, 3.



the best machine shop supply room available,’ recalled Gloria Rudisch, Minsky’s widow. ‘The three of them used to chat over a wild idea in the morning and then build it into a working prototype in the afternoon.’<sup>415</sup> Following this internship, McCarthy was asked by Shannon to co-develop a volume on the study of automata.<sup>416</sup> Shannon had by then developed an introductory survey of mathematical research on the topic, which he published as ‘Computers and Automata’ in the *Proceedings of the Institute of Radio Engineers* in 1953.<sup>417</sup>

Like von Neumann, Shannon in 1953 stressed obvious differences between natural and artificial automata, such as size, input-output mechanisms and structural organization. With these limitations acknowledged, he highlighted Turing’s universal machines, W. Ross Ashby’s Homeostat and von Neumann’s work on self-reproducing automata as exemplary research. Shannon’s paper simultaneously pushed and questioned the limits of *non-numerical* computing, in which computers were used to manipulate things other than just numbers, such as words. ‘Can we program a digital computer so that (eventually) 99 per cent of the orders it follows are written by the computer itself? Can a machine be constructed which will design other machines?’ he wondered.<sup>418</sup>

Anthropomorphic rhetoric remained contentious in the early 1950s, a period in which the language used to describe computing techniques was still in flux. The 1953 book *Automatic Digital Calculators* apologized for using ‘memory’ in relation to computing and acknowledged that terminology in the field had ‘not yet stabilized.’<sup>419</sup> Delegates to the 1953 Conference on High Speed Digital Computers at the National Physical Laboratory had only recently agreed to set norms like spelling the activity ‘program’ and not ‘programme.’<sup>420</sup>

Indeterminacy over anthropomorphic terms was, in part, a product of uncertainty over computing’s potential to mimic such abilities. In the same 1953 collection as Shannon’s paper, Maurice Wilkes, Director of the Mathematical Laboratory at the University of Cambridge and by then a recognized expert on what would come to be called software,

---

<sup>415</sup> Personal communication with Gloria Rudisch.

<sup>416</sup> Nilsson, ‘John McCarthy, 1927-2011’, 3. Guice and Fleck misstate that the automata volume began as a 1952 event organized by McCarthy, in: Guice, ‘Controversy and the State’, 108–9; Fleck, ‘Development and Establishment in Artificial Intelligence’, 172.

<sup>417</sup> Shannon, ‘Computers and Automata’.

<sup>418</sup> Shannon, 1241.

<sup>419</sup> Booth and Booth, *Automatic Digital Calculators*, 1953, v.

<sup>420</sup> Booth and Booth, v.

cautioned that a “generalized “learning” program’ capable of this sort of behaviour was a long way off.<sup>421</sup> Anthony Oettinger, a member of Wilkes’ lab, had only one year earlier made a case for the future of learning machines. In his 1952 paper ‘Programming a Digital Computer to Learn,’ Oettinger outlined two operational programs that showed how a digital computer could be made to exhibit ‘animal-like behaviour.’<sup>422</sup> To justify this claim, he assumed a functional equivalence between the behaviour of a program and the behaviour of a human being. His Shopping Programme ‘visited’ different ‘shops,’ meaning locations in memory, to verify the location of different ‘items,’ also stored in memory. Oettinger argued that the program ‘learned’ because it could identify an item’s whereabouts if queried about that item a second time.

McCarthy himself had not had the opportunity to program a machine by 1952. In 1953, he moved to Stanford as an acting associate professor of mathematics and to deepen his study of differential equations. McCarthy’s outlook on sophisticated machine behaviour, and his role in substantiating its promise, shifted between 1953-55 as he fielded contributions for his and Shannon’s co-edited automata volume. Despite entries from von Neumann and leading cybernetics researchers like Donald MacKay and W. Ross Ashby, McCarthy grew disappointed by the state of the art. He wrote to Shannon in 1954, ‘The collection as a whole does not represent great progress but is certainly representative of current thought.’<sup>423</sup> While the contributions from British cyberneticists like Ashby and MacKay offered novel heuristics for future study, no submissions ‘Lead directly to a solution of the problem of thinking automata.’<sup>424</sup> Even von Neumann’s contribution, the seed of which had so dazzled McCarthy in Pasadena six years earlier, was now simply, ‘An important result, perhaps slightly off the main track.’<sup>425</sup> Wilkes replied from the UK that he would be pleased to contribute a short piece about learning programs but would rather share his thoughts on, ‘The *limitations* of

---

<sup>421</sup> Wilkes, ‘Can Machines Think?’. Wilkes also gave this message in a 1951 BBC broadcast. See: Jones, ‘Five 1951 BBC Broadcasts on Automatic Calculating Machines’.

<sup>422</sup> Oettinger, ‘Programming a Digital Computer to Learn.’ Oettinger has received uneven historical treatment. He is almost entirely absent from: Boden, *Mind as Machine*; Cordeschi, ‘AI Turns Fifty: Revisiting Its Origins’; Cordeschi, *The Discovery of the Artificial*, Ch. 5-6.

<sup>423</sup> Kline, ‘Dartmouth’, 8.

<sup>424</sup> Kline, 8.

<sup>425</sup> McCarthy to Shannon, 11 August 1954.

programs of this type.<sup>426</sup> Like McCarthy, his intuition was that, ‘Some entirely new ideas are necessary if further progress is to be made.’<sup>427</sup>

McCarthy’s own contribution, ‘The Inversion of Functions Defined by Turing Machines,’ examined how ‘intellectual problems’ could be solved computationally if they were first rendered in the form of an abstract Turing machine.<sup>428</sup> His method improved the efficiency of a test designed to show whether or not a proposed solution was indeed solvable, or whether the attempt would run on, computationally, forever. Enumerated search techniques lacked the efficiency of rational structure, he argued.<sup>429</sup> To improve on them, his technique prioritized certain enumerated procedures over others by referencing the abstract machine’s prior ‘experience’ of which methods were most likely to terminate. If, for example, A terminated and C was like A, then C would be a better candidate for the systems’ next step than B would be.

One challenge with this technique was that there did not yet exist any clear way to corroborate even its partial success. ‘It is a great advantage in conducting a search to be able to know when one is close,’ McCarthy conceded.<sup>430</sup> Simon, Newell and Shaw had used heuristic ‘rules of thumb’ to distinguish desirable paths from undesirable ones.<sup>431</sup> Rosenblatt had used biological evidence. McCarthy sought an idealized solution. He proposed to evaluate what he called the ‘properties’ and ‘relevant concepts’ about new computational enumerations to see which would prove to be relevant for future decisions. These interpretations could then themselves be enumerated to determine how to prioritize new actions.<sup>432</sup> Roughly speaking, McCarthy proposed to structure a Turing Machine with a predisposition to organize *itself*.

McCarthy and Shannon debated how to characterize their new volume. In February 1954, McCarthy prepared a talk on ‘The Intelligence of Automata’ for the Scientific Research

---

<sup>426</sup> Emphasis mine. Letter to M. Wilkes of University of Cambridge, May 19, 1953

<sup>427</sup> Letter to M. Wilkes of University of Cambridge, May 19, 1953

<sup>428</sup> McCarthy, ‘The Inversion of Functions Defined by Turing Machines’.

<sup>429</sup> McCarthy, 177–78.

<sup>430</sup> McCarthy, 180.

<sup>431</sup> From 1953-55, McCarthy was based in the same Mathematics Department at Stanford as George Pólya, by then retired, whose work on heuristics had inspired Simon and Newell. See Chapter Two.

<sup>432</sup> McCarthy, ‘The Inversion of Functions Defined by Turing Machines’, 180.

Society Sigma Xi at Stanford University.<sup>433</sup> McCarthy proposed to Shannon the title ‘Towards Intelligent Automata’ and asked if the collection, ‘should be regarded as a step in the direction of the design of intelligent automata’ or if that was ‘too presumptuous.’<sup>434</sup> That the volume was published under the title *Automata Studies* suggests the senior researcher’s answer. A similar debate played out in revised drafts of the volume’s introduction. McCarthy’s first draft made four references to intelligence.<sup>435</sup> The published version mentioned it only once, and then only in reference to a submitted paper.<sup>436</sup> Elsewhere, the term ‘intelligent machines’ was replaced by ‘automata theory.’<sup>437</sup> Decisions over terminology implied a difference in the two men’s intuitions of what their mathematics represented. While Shannon had taken steps in his 1953 paper to explicate the ways in which the brain-computer analogy failed to align with biological fact, McCarthy believed that to find a solution, the nature of thought had to be rendered at a higher level of abstraction, one with little tie to biological fact.

McCarthy’s ardent faith in logic was not just intellectual. In unpublished scientific notes, book-reviews and short stories from this period, he returned repeatedly to questions that concerned how logic could penetrate the unjust, top-down, power structures of the established social hierarchy. In ‘Secession,’ he described a group of scientists who develop the first ever human colony on the moon to escape Earth’s red-tape and enjoy unfettered academic pursuits.<sup>438</sup> In ‘Prescription for a Utopian Colony,’ he outlined how to build a real utopian colony in the United States, complete with details on employment, population, yearly cost, basic rules and other considerations. In *Levels of Theory in Engineering and Politics*, he explored non-fiction scenarios, questioning how computers could be used to improve decision procedures in political systems in the U.S. and Russia. ‘There should be a semi-mathematical theory for decision procedure for formulating policy,’ he contented, inadvertently echoing Simon’s conflation of social governance with computing.<sup>439</sup> McCarthy’s

---

<sup>433</sup> McCarthy to Minsky, December 1953.

<sup>434</sup> Kline, ‘Dartmouth’, 8; Kline, *The Cybernetics Moment*, 157.

<sup>435</sup> McCarthy, ‘Draft of Towards Intelligent Automata, Introduction’, 1–4.

<sup>436</sup> Shannon and McCarthy, *Automata Studies*, viii.

<sup>437</sup> McCarthy, ‘Draft of Towards Intelligent Automata, Introduction’, 1; Shannon and McCarthy, *Automata Studies*, vi.

<sup>438</sup> McCarthy, ‘Secession’, 71–75.

<sup>439</sup> McCarthy, ‘Levels of Theory in Engineering and Politics’.

various accounts presumed that computing technologies could be productively applied to governance without contradiction.<sup>440</sup>

Given the state of anti-Communist rhetoric in the mid-century, it is noteworthy that McCarthy, who had been raised in a Communist household and had participated in a Communist cell while at Princeton, chose to develop these subjects in writing.<sup>441</sup> In 1949, the Board of Regents at the University of California had voted to impose a 'loyalty oath' on all staff, which resulted in the firing of thirty-one non-signers in 1950.<sup>442</sup> Other American university staff faced similar threats. Had his notes been shared or publicized, critics might have cast McCarthy's faith in logic as dangerous, especially the material that gravitated towards how computers could be used to centralize government decision procedures.<sup>443</sup> This ideological orientation distinguished him from Shannon and von Neumann, who made the caveat that logic and mathematics were insufficient for holistic inquiry into neural dynamics. For a young McCarthy, the promise of logic and computation extended far further, bordering on emancipatory political change.

#### 1955-56: From Automata to Automatic Programming and AI

The advent of automatic programming techniques in the mid-1950s provided McCarthy with a new medium with which to test his ambitions to simulate cognition using tools like formal logic. In what follows, I consider how the automation and standardization of programming techniques preceded and intersected with the 1956 Dartmouth Summer

---

<sup>440</sup> Analogous macho techno-emancipatory themes became the cult fantasies that spawned Silicon Valley in the 1990s and 2000s, according to Barbrook and Cameron, 'The Californian Ideology'. AI's image of emancipatory free labour, they state, inspired science fiction authors like Asimov and Heinlein to fabricate poverty and racism free techno-utopias. See: Turner, *From Counterculture to Cyberculture* on the arc of computing from emblem of conservatism to symbol of rebellion (for the elite).

<sup>441</sup> McCarthy claimed to have remained sympathetic after quitting the left in 1953, in: Hilts, *Scientific Temperaments*, 261. His niece claimed that he disavowed communism to become a conservative Republican after Russia's invasion of Czechoslovakia in 1968. O'Dowd, 'This Son of an Irish Emigrant Invented Artificial Intelligence and Changed the World'.

<sup>442</sup> Kaiser, 'Nuclear Democracy', 245–46.

<sup>443</sup> In 1949, while at Princeton, McCarthy joined his local Communist Party cell, which he remained in for a few years before quitting due to poor turnout. Nilsson, 'John McCarthy, 1927-2011', 5.

Research Project on Artificial Intelligence, which gave the field its name and initial membership. I relate the early development of AI directly to these trends to show how McCarthy's vision for the field was entangled with industrial, administrative and pedagogical needs unfolding around him in response to labour shortages in the field of programming.

To begin, I provide a brief expository look at the history of digital electronic programming leading up to the mid-1950s. Two entangled innovations from the 1930s and mid-to-late 1940s altered the field of mechanical computing irreversibly by 1950. The first was electricity. In the 1930s, John Atanasoff, a physicist at Iowa State University, experimented with an electronic calculating device. Others attempted similar projects but failed to manifest meaningful results. In 1945, using sponsorship from the U.S. War Department, John Mauchly and J. Presper Eckert at the Moore School for Electrical Engineering at the University of Pennsylvania developed ENIAC, or Electronic Numerical Integrator and Computer, for the Ballistics Research Laboratory at Aberdeen Proving Grounds. While the first 'computers' were often young highly educated women hired to solve complex equations using rote transcription, ENIAC was a glimpse of what would replace them: error-prone yet lightning fast electrical machines whose rapid improvement was also a military priority.

In addition to being electrical, ENIAC was also programmable. This was the second transformative innovation: the advent of 'stored-program' techniques, which set the foundations for the emergence of 'software' in the 1960s. In 1946, ENIAC could hold and reflexively adjust data in a memory bank. This capacity allowed for a far broader computational range than had been possible with mechanical calculators.<sup>444</sup> Whereas Harvard's Mark I, an electromechanical device that did not have stored-program capacity, could perform two to three additions per second, ENIAC could perform five thousand.<sup>445</sup> As a handful of pioneering research teams in the U.S. and UK experimented with stored-program techniques into the late 1940s and early 1950s, the craft and cultural status of programming slowly began to shift. The procedural labour required to make a machine operate was no

---

<sup>444</sup> For a critical take on the stored-program concept in the historiography of computing see: Haigh, Priestley, and Rope, 'Reconsidering the Stored-Program Concept'. On the false equivalence between stored-program computing and Turing's notion of a universal machine see: Agar, *The Government Machine*, 7.

<sup>445</sup> See: Ensmenger, *The Computer Boys Take Over*, 33. On ENIAC see: Bergin, 'ENIAC: Development and Early Days'; Fritz, 'The Women of ENIAC'; Stern, 'Computers'.

longer dismissed by engineers, often chauvinistically, as glorified clerical work. Computing gained an aura of ‘erotic scientificity’ in the mid-1950s, due in part to the seemingly protean nature of stored-program capabilities.<sup>446</sup>

Bit by bit, the accumulation of automated programming techniques in different labs across the U.S. and UK reduced the rote drudgery and expertise required to ‘get programming right.’<sup>447</sup> In September 1947, Kathleen Booth introduced the notion of Contracted Notation, which represented programming commands as mnemonics, such as ‘cR’ for ‘Clear Register.’<sup>448</sup> Booth abstracted onerously dense machine code instructions such as, say, 100101... into intuitive symbols. She also invented the first assembler and assembly language, metaprograms that usefully assembled comprehensible instruction sets into obtuse machine code, which saved considerable labour.<sup>449</sup> In 1952, Grace Hopper implemented the first compiler, which pushed this progress further by consolidating sets of translated instructions on the UNIVAC I device.<sup>450</sup> Each structural refinement reduced the otherwise significant time required to identify and debug logical errors.

Prior to the mid-1950s, individual digital stored-program computers required distinct programming techniques. Individual material affordances like a machine’s clock speed, serial or parallel processor or plugboard design influenced the logical architecture of that device. As a by-product of this heterogeneity, programming terminology had still not stabilized by 1953, eight years after ENIAC first became operational.<sup>451</sup> This began to change following the emergence of commercially fabricated digital electronic computers in the mid-1950s. The existence of carbon-copy machines made the standardization of automatic programming

---

<sup>446</sup> For a concise survey of the literature on gender and computing see: Miltner, ‘Girls Who Coded’; Edwards, *The Closed World*, 121.

<sup>447</sup> Automatic coding/programming carried various meanings in this period. From Nofre et al.: ‘the use of the computer itself to take over routine, mechanizable aspects of the programming process, such as the conversion between binary and decimal representations of numbers, the assembly of subroutines taken from a library into complete programs, or the translation of various forms of pseudo-code into machine code.’ See: Nofre, Priestley, and Alberts, ‘When Technology Became Language’, 49. Quote as cited in: Booth and Booth, *Automatic Digital Calculators*, 213.

<sup>448</sup> On Kathleen Booth’s foundational work on machine translation and machine learning, see: Priddy, ‘Kathleen Booth’.

<sup>449</sup> Norberg, O’Neill, and Freedman, *Transforming Computer Technology*, 75.

<sup>450</sup> Meyr, ‘Grace Hopper and the Marvelous Machine’.

<sup>451</sup> Booth and Booth, *Automatic Digital Calculators*, 1965, Preface to the First Edition.

techniques possible across numerous devices. IBM's FORTRAN system (for FORMula TRANslation), for example, is commonly viewed as the first commercial programming language.<sup>452</sup> It would run on any of the one hundred and twenty-three IBM 704s produced between 1955-60.<sup>453</sup>

The landscape for interoperability changed as digital electronic computing expanded exponentially after 1950. While only two digital electronic computers were in operation in the United States in 1950, 243 machines were operational in 1955, with 5,400 in 1960, 25,000 in 1965 and 75,000 in 1970.<sup>454</sup> Amidst this proliferation, the consolidation of automatic coding techniques into programming language and language families like FORTRAN, ALGOL and COBOL expanded computing's client base by lowering the expertise required to operate a device. 'No computer manufacturer would succeed in selling his machines without an auto-programming system,' commented Kathleen Booth in her pathbreaking 1953 book of programming.<sup>455</sup> IBM developers claimed to have invested eighteen man years of work into FORTRAN between 1954-57.<sup>456</sup> They estimated that ninety percent of the time required for scientific use of a computer elapsed in coding and debugging and that FORTRAN could reduce this investment of time by eighty percent.

The commercial case for automatic coding had not been obvious to IBM, a company I focus on here due to their role in later AI research. Nathaniel Rochester, an IBM employee and co-convenor of the Dartmouth workshop, claimed in personal records that he and his team had 'sold the idea' to management as a patriotic act after IBM reactivated its military product unit in response to the Korean War.<sup>457</sup> Having been hired by the Navy in 1947 to construct the arithmetic unit of the Whirlwind I, among the first digital computers in the U.S., Rochester was particularly well positioned to recognize the prospects of the field. He spent

---

<sup>452</sup> FORTRAN was developed by John Backus, a mathematician and IBM employee, to allow mathematical notation, such as algebraic formulas, to be included within the programs used by a 704. Nofre, Priestley, and Alberts, 'When Technology Became Language', 50; 'IBM: Preliminary Report'; 'IBM: The FORTRAN Automatic Coding System'; Stoyan, 'Early LISP History (1956 - 1959)', 301.

<sup>453</sup> 'History of IBM, 1950s, 1954'.

<sup>454</sup> Ensmenger, *The Computer Boys Take Over*, 28.

<sup>455</sup> Booth and Booth, *Automatic Digital Calculators*, 1965, 215.

<sup>456</sup> Backus et al., 'The FORTRAN Automatic Coding System', 188.

<sup>457</sup> Goldstein, Oral History: Nathaniel Rochester; IBM CEO Thomas J. Watson Sr. was amazed to receive twenty orders for the device: Brock, *The Second Information Revolution*, 98-99; On IBM's decision see: Edwards, *The Closed World*, 381.



the rest of his career at IBM and recalled his initial surprise to learn that the company had ‘no intention of venturing into stored-program machines.’<sup>458</sup> Once approved, responsibility fell to him to design and construct IBM’s first large-scale scientific computer, the Defense Calculator, released in 1952 as the IBM 701. By 1954, he oversaw a staff of 450 people.<sup>459</sup> Under Rochester, they designed the IBM 702, IBM’s first commercial automatic programming calculator, and the IBM 704 and 705, two influential scientific and commercial machines.



Figure 3.0 – An IBM 704 Electronic Data-Processing Machine<sup>460</sup>

This brief preamble on the history of programming languages leading up the mid-1950s provides context for McCarthy’s views on computing prior to the Dartmouth workshop. In February 1955, he left a brief teaching position at Stanford to become an assistant professor of mathematics at Dartmouth College. Despite a longstanding interest in computing, he had never actually tried to program a computer until that year.<sup>461</sup> This was partly beyond his control; neither Stanford, where he had been based, nor Dartmouth, to which he had just arrived, owned a stored-program computer for him to experiment with. To his good fortune, Philip M. Morse, the ‘founding father of operations research,’ had just

---

<sup>458</sup> Gruenberger, ‘The History of the JOHNNIAC’, 50.

<sup>459</sup> Rochester, ‘Biographical Data, Nov 3, 1957’, 3.

<sup>460</sup> Courtesy of International Business Machines Corporation, © 1954 International Business Machines Corporation. ‘IBM 704 Electronic Data-Processing Machine: Manual of Operation’.

<sup>461</sup> Nilsson, ‘John McCarthy, 1927-2011’, 5.

convinced James Killian, then president of MIT, to open the institute's new IBM-funded Computation Center to sister colleges in the New England area.<sup>462</sup> As the chosen representative of Dartmouth College, McCarthy was introduced to Rochester, who invited him to spend the summer of 1955 at IBM.<sup>463</sup>

In April 1955, due in part to frustration over the lack of path-breaking research in *Automata Studies*, McCarthy appealed to Warren Weaver at the Rockefeller Foundation to finance a ten-man, six-week summer research session on automata and brain models.<sup>464</sup> He enlisted Shannon in this project, since Weaver had been a nominal co-author on Shannon's landmark paper on information theory.<sup>465</sup> In an internal memo, Weaver expressed misgivings about the personal nature of McCarthy's query, which lacked formal institutional backing. He directed McCarthy to Robert Morison, head of Rockefeller's Biological and Medical Research Division, but told Morison privately, 'I am very doubtful that the RF ought to do anything about this.'<sup>466</sup>

Morison, too, expressed qualms. He questioned whether a seminar format was the right mechanism for discovery given that mathematical theories of brain functions remained, 'Pre-Newtonian.'<sup>467</sup> Despite hesitations, the trio entertained who might participate in such an event. The majority of nominees were either mathematicians, such as Rochester, Minsky, Oliver Selfridge and George Pólya, whose work on heuristics had so inspired Simon and Newell, or cyberneticists, broadly construed, such as von Neumann, Ashby and MacKay.<sup>468</sup>

---

<sup>462</sup> Aker, *Calculating a Natural World*, 286–88; Little, 'Philip M. Morse and the Beginnings', 146; McCarthy, 'History of LISP', 217.

<sup>463</sup> Stoyan argues that McCarthy likely did not learn about FORTRAN during this period of his work at IBM, although its development played a significant role in his development of LISP: Stoyan, 'Early LISP History (1956 - 1959)', 300–301; Stoyan, 'The Influence of the Designer on the Design—J. McCarthy and LISP'.

<sup>464</sup> Weaver, 'WW Diary Entry Re: John McCarthy'.

<sup>465</sup> Kline, 'Dartmouth', 8.

<sup>466</sup> Weaver, 'Inter-Office Correspondence, WW to RSM'. Morison was an acquaintance of Wiener: Kline, *The Cybernetics Moment*, 80.

<sup>467</sup> Morison, 'New York'.

<sup>468</sup> Selfridge's 1954 research on learning systems inspired Newell to begin work on the Logic Theory Machine. Selfridge, 'Pattern Recognition and Modern Computers'.

If, as McCarthy later claimed, he coined the term ‘artificial intelligence’ to get away from cybernetics, this desire was not obvious in his pitch for funding.<sup>469</sup> Pickering credits Ashby as one of the four pioneers of British cybernetics. Edwards, Dupuy and Heims credit von Neumann as another foundational contributor. McCarthy’s pitch was more tactful; he combined established theorists with a cast of rising mathematicians. Morison, in response, urged he and Shannon to include psychologists like Hans-Lukas Teuber and Karl H. Pribram ‘if only for the purpose of keeping the group from speculating too wildly on how the brain might work.’<sup>470</sup>

In August 1955, a committee of four—Shannon, McCarthy, Rochester and Minsky—submitted their proposal to Morison.<sup>471</sup> At a price tag of \$13,500, the ‘Dartmouth Summer Research Project on Artificial Intelligence’ was to assemble ten men for two months at Dartmouth College to proceed informally ‘on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.’<sup>472</sup> The group posited that ‘a significant advance’ could be made if a curated group of scientists could work on relevant subtopics together in a shared location.

Automatic programming techniques anchored this first collective research agenda for the newly dubbed artificial intelligence problem. ‘The major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have,’ the co-convenors summarized.<sup>473</sup> They presumed that the rules that governed the craft art of programming could be made to resemble whatever rules governed natural language.

---

<sup>469</sup> ‘One of the reasons for inventing the term ‘artificial intelligence’ was to escape association with “cybernetics,”...I wished to avoid having either to accept Norbert... Wiener as a guru or having to argue with him.’ McCarthy, ‘Review by Brian P. Bloomfield (Ed.), The Question of Artificial Intelligence (1987)’. McCarthy also claimed that *Automata Studies* had attracted only mathematical work, and that ‘intelligence’ would invite broader proposals: McCorduck, *Machines Who Think*, 96.

<sup>470</sup> Morison, ‘New York’, 1.

<sup>471</sup> Minsky credited McCarthy with doing the bulk of the organising. Minsky, An Interview with Marvin L. Minsky, 11.

<sup>472</sup> Minsky et al., ‘Proposal’, 2. This oft-cited mission statement was drafted by Minsky, who suggested it to replace allusion to there being ‘nothing that the human brain can do that [a machine can’t be made to simulate].’ Minsky, ‘To Drs. J McCarthy and Nat Rochester, IBM, Poughkeepsie, N.Y.’, n.d.

<sup>473</sup> Minsky et al., ‘Proposal’, 1.

Rochester claimed, for instance, that neural networks could be used to ‘form concepts’ in the mold of Hebb’s assembly theory. To realize this vague goal, co-conveners first sought baselines to measure the efficiency of their progress, such as how to characterize and measure computational self-improvement, how to judge between different levels and types of abstraction and how to set a criterion for efficiency of computation.

A brief sub-proposal from each co-convenor developed notions of programming and efficiency further. Shannon put off grandiose gestures of machine intelligence to pursue adaptation in automata, studying how exceedingly simple mathematical mechanisms gave rise to complex behaviours. He proposed to use the conceptual tools he had formalized as information theory to clarify how reliable transmissions carried through unreliable elements in an artificial information network. Minsky, too, doubled down on his own ongoing work. He proposed to revisit his PhD research on reinforcement learning in a computer, which I discuss in [Chapter Five](#). Minsky hoped, by the end of that summer, to have designed the framework for a program that integrated abstract sensory and motor capacities into a goal-seeking system, such that the program could successfully manipulate its environment and exhibit higher order behaviour.<sup>474</sup>

Rochester, of IBM, mixed technical and non-technical terminology in a manner that belied his project’s audacious vagaries. He aimed to formalize the role of imagination in mechanical thought by developing a computer program that expressed originality in its discovery of a solution. His sub-proposal made blunt use of psychology; he equated human beings to stored-program calculators that used game theory to predict successful behaviours in relation to their surrounding ‘culture.’ He singled out the Monte Carlo method as the best mathematical approximation to the brain’s presumed use of randomness in acts of creativity, since reliance on pure randomness alone would produce chaotic human behaviour.<sup>475</sup> With passing references to Craik, Hebb and Lashley, Rochester proposed that a similar mathematical ‘hunch’ was needed to leverage randomness for original mechanical

---

<sup>474</sup> Minsky et al., 8.

<sup>475</sup> The Monte Carlo Method is a random sampling technique used to generate approximate solutions to intractably complex systems. In 1957, Herman Kahn, a luminary of RAND, said, ‘A Monte Carlo problem done in a completely straightforward fashion is almost prima facie evidence of insufficient thought.’ See: Scoblic, ‘The Postwar Development of Tools to Mitigate Uncertainty’. Draft shared privately.

thought.<sup>476</sup> Like his co-convenors, Rochester positioned the hard work in AI as understanding *how* the brain used computation, not *if* it did.

Morison responded with scepticism. 'This new field of mathematical models for thought... is still difficult to grasp very clearly.'<sup>477</sup> He approved a \$7,500 grant, half of the requested \$13,500.<sup>478</sup> His hesitation was not purely subject based; Kline identifies that Morison and Weaver had by then co-funded research on mathematical biology undertaken by Wiener and Rosenblueth, which led to the book *Cybernetics*.<sup>479</sup> Morison offered a 'modest gamble for exploring a new approach.'<sup>480</sup> Supportive parties fell into rank. Provost Donald Morrison and John Kemeny, Chairman of the Mathematics Department, agreed to host the event in Hanover and to facilitate payments. IBM, Bell Laboratories and other home-organizations agreed to cover staff salaries. From June 18, 1956 to August 17, 1956, visitors came and went from the top floor of the Math Department at Dartmouth College for the 1956 Dartmouth Summer Research Project on Artificial Intelligence.

Since the summer meeting figures centrally in existing histories of AI, I will include more detail here than may be strictly necessary to understand McCarthy's work. My aim is not to center on the event, but simply to preserve relevant details for posterity. For instance, Ronald Kline's illuminating account of the Dartmouth meeting mentions that no formal record of attendance survived.<sup>481</sup> I have uncovered a three-page report authored by McCarthy that lists twenty-one invitees, twenty attendees and each attendee's approximate duration of stay.<sup>482</sup> Core participants were Simon, Newell, Ray Solomonoff (a friend of Minsky's), Trenchard More (a masters student of Shannon's), Julian Bigelow (Institute of Advanced Study), David Sayre (IBM), Oliver Selfridge (MIT) and the co-convenors.<sup>483</sup> Nine others visited

---

<sup>476</sup> For Minsky, sensory data provided this hunch. Minsky et al., 'Proposal', 15.

<sup>477</sup> As cited in: Kline, 'Dartmouth', 10.

<sup>478</sup> Grant: GA BMR 5550. Approximately \$120,000 today.

<sup>479</sup> Kline, *The Cybernetics Moment*, 161.

<sup>480</sup> Morison, 'Letter from Robert Morison to John McCarthy'.

<sup>481</sup> Kline notes Morison never received a summary, which informed his decision not to fund the Teddington conference. Kline, *The Cybernetics Moment*, 163–64.

<sup>482</sup> McCarthy, 'Report on Dartmouth Summer Research Project'.

<sup>483</sup> McCarthy, Minsky and Solomonoff for the full period; Shannon and Rochester for two weeks at both the beginning and end; Simon and Newell for the first one and a half weeks; More for all but a period of two weeks; Bigelow for two weeks at the end; Sayre for one-week; Selfridge intermittently. See: McCarthy, 2. IBM funded More's participation. The Rockefeller grant covered Solomonoff.

for one to two days.<sup>484</sup> Margaret Andrews served as secretary.<sup>485</sup> Surviving notes by Solomonoff list twenty participants,<sup>486</sup> which corroborates those I cite in footnotes.<sup>487</sup> Notes from More list thirty-two attendees, a handful of whom remain uncorroborated.<sup>488</sup>

This attendee list clarifies that psychological and neurophysiological expertise was in short supply amongst those who participated in the 1956 Dartmouth event, even while concepts like learning, imagination and intelligence were dealt with in earnest. The group was comprised primarily of mathematicians, physicists and engineers. Teuber and Pribram, the psychologists Morison had suggested to include to avoid speculation on psychological processes, did not attend, although McCarthy and Rochester visited Pribram prior to the meeting.<sup>489</sup> All involved had hoped that Hebb would participate, but it is unclear if he did. The view of 'intelligence' developed in the first collective workshop on AI was thus informed primarily by a set of specialists working *outside* the human sciences, which dulled basic

---

<sup>484</sup> McCulloch (MIT), W. Ross Ashby (Barnwood House, England), Abraham Robinson (University of Toronto), Peter Milner (McGill University), Tom Etter (New York), Arthur Samuel (IBM), Kenneth 'Kent' Shoulders (MIT), John Nash (MIT), Donald MacKay (University of London; invited but did not attend due to his partner's pregnancy) and Bernard Widrow (MIT) Misspelled as 'Woodrow' in Solomonoff, 'Dartmouth', 9; Widrow's participation corroborated in: Widrow, Bernard Widrow, an oral history. List from: McCarthy, 'Report on Dartmouth Summer Research Project', 3.

<sup>485</sup> I have been unable to locate Andrews or any notes she may have from the meeting. McCarthy, 'Report on Dartmouth Summer Research Project', 3. Dr. Gloria Rudisch, Minsky's partner, also attended for periods of time.

<sup>486</sup> See: Solomonoff, 'Dartmouth', 9.

<sup>487</sup> With the exception of Milner, a neuroscientist and associate of Hebb's, all three lists confirm the attendance of those I outlined in footnotes. Solomonoff's list does not.

<sup>488</sup> Uncorroborated attendees: Backus (IBM), 'Frankel' (likely Stanley Frankel, Manhattan Project advisor turned consultant due to Red Scare), David W. Hagelbarger (Bell Labs, profiled in 'Computers and Automata'), John Holland (IBM, Rochester's partner on Hebbian cell-assembly simulation), E.F Moore (Bell Labs, *Automata Studies* contributor), Walter Pitts (MIT, loosely), 'Rappert' (likely Prof. Anatol Rapoport, Mathematical Biology, University of Michigan), Norman Shapiro ('55 Princeton grad under Church, *Automata Studies* contributor), Albert M. Uttley (NPL, *Automata Studies* contributor, Teddington organizer) and Norbert Wiener (MIT). Solomonoff, 'Dartmouth', 9, 12; adds W.A. Clark, B.G. Farley, R. Culver, Bill Shutz. Notes from More and Solomonoff suggest that Alex Bernstein (IBM) participated and presented to the group on August 8<sup>th</sup>. A letter from Simon lists a F.B. Fitch. Ronald Kline adds Herbert Gelernter (IBM, printed as 'Gerlertner'). Kline, *The Cybernetics Moment*, 163.

<sup>489</sup> Neither did Wilkes or Oettinger attend for reasons that remain unclear. Geographical distance and a lack of budget offer a potential explanation. Morison, 'New York', 17 June 1955, 1; Kline, *The Cybernetics Moment*, 158; McCarthy, 'Letter from John McCarthy to Dr. Robert Morison'.

scrutiny of the groups' heady re-articulation of human faculties through the prism of mathematics.



Figure 4.0 – ‘The Artificial Intelligencers.’ Attendees at Dartmouth.<sup>490</sup>

*Rochester, Minsky, McCarthy*

*Selfridge, Solomonoff, Milner, Shannon*

#### 1956: The Dartmouth Summer Research Project for McCarthy

Like the Macy Conferences that propelled cybernetic theory, proceedings of the Dartmouth meeting can be analysed from multiple perspectives. My focus in this chapter is on McCarthy's contributions, not on the event as a whole. In Hanover, McCarthy sought to develop a language that blended the deductive power of formal logic with the representative

---

<sup>490</sup> Used with permission from Gloria Rudisch.

flexibility of natural language. This goal was not without precedent. In 1950, Shannon had argued that due to factors like the high frequency of the letter 'E' and the tendency for 'H' to follow 'T', the English language was approximately 50% redundant.<sup>491</sup> In his sub-proposal to Morison, McCarthy proposed to design a new 'formal,' 'logical' or 'artificial' language that corresponded to English.<sup>492</sup> This language would be capable of expressing conjectural arguments and notions of physical objects or events, along with other benefits of natural languages.<sup>493</sup> In principle, this would make it possible to program a machine to perform simple tasks like playing a game. The printouts of such a program would lend insight into the perceived relation of language to intelligence.

In 'When Technology Became Language,' David Nofre, Mark Priestley and Gerard Alberts argue that the increasing heterogeneity of commercial computer installations in this period contributed to a shift in the metaphorical use of the term 'language' in connection with such machines. Commercial machines led computer educators and managers, 'To draw on the disciplines of symbolic logic and linguistics to develop models of intelligibility that would enable abstraction away from the machine and toward the development of free-standing notations.'<sup>494</sup> In the process, programming languages became epistemic objects that required new conceptual tools to understand and make use of them, such as standardized terminology and pedagogical techniques. Between 1955-60, the metaphor of 'language' in computing, 'Lost its anthropomorphic connotation' and took on a more abstract meaning related to formal languages like logic and linguistics.<sup>495</sup>

In 'An Approach to the Artificial Intelligence Problem,' an unpublished manuscript likely drafted that summer, McCarthy did as Nofre and colleagues describe; he advocated for the elevation of computer programs to the status of epistemic object. He classified the program as 'the agent in artificial intelligence' to normalize phrasing such as 'The program does this' rather than 'The machine does this.'<sup>496</sup> By affording agency to the program and not the computer, McCarthy projected a form of Cartesian dualism onto computing; a projection

---

<sup>491</sup> Shannon, 'Prediction and Entropy of Printed English'.

<sup>492</sup> Minsky et al., 'Proposal', 60–61.

<sup>493</sup> Minsky et al., 16.

<sup>494</sup> Nofre, Priestley, and Alberts, 'When Technology Became Language', 42.

<sup>495</sup> Nofre, Priestley, and Alberts, 41.

<sup>496</sup> McCarthy, 'An Approach to the Artificial Intelligence Problem', 3.



that mirrored his own timidity with hardware.<sup>497</sup> His approach rendered computing immaterial by focusing agency on the instruction sets, not the machines they ran on. Left unresolved was the boundary between hardware and software, to use contemporary terms. Also unresolved was the question of what minimum amount of computation constituted an act of intelligence on the part of a program.

On this point, McCarthy hedged. His objective for AI as of 1956 was to, 'Write a calculator program which can solve intellectual problems as well as or better than a human being' in areas like program writing, theorem proving or game play.<sup>498</sup> By this definition, new automatic coding techniques like assembly languages and compilers would count as AI too, as would perhaps all programming languages. Minsky later acknowledged this appropriative excess, which he argued would also encompass time-sharing, a system co-developed by McCarthy for coordinated multi-party access to a single device that revolutionized computing and laid the groundwork for personal computing.<sup>499</sup>

The closest McCarthy came to a defined threshold between programming and AI hinged on a subjective view of scientific discovery. He deemed a program that revealed nothing that its developers could not have guessed prior to running it to be a waste of machine time.<sup>500</sup> AI was meant to encompass *novel* decision procedures, not routine procedures. His definition of novelty, however, remained open. He proposed, as an exemplar, a 'transition' event in which a program would improve *itself* iteratively in a manner that signified a conversion into further self-improvement.<sup>501</sup> This proposal echoed Shannon's musings about whether a computer could one day write 99 per cent of its own orders.

It was from within this highly mathematical framing that McCarthy positioned intelligence. In keeping with his *Automata Studies* paper on how to efficiently order a Turing Machine, his conception of intelligence in computer programs resembled a measurable version of self-refining computational efficiency. 'The only real problem [for McCarthy] is the

---

<sup>497</sup> 'I was very shy of proposing hardware modification, especially as I did not understand electronics well enough to read the logic diagrams.' Lee, McCarthy, and Licklider, 'The Beginnings at MIT', 20.

<sup>498</sup> McCarthy, 'An Approach to the Artificial Intelligence Problem'.

<sup>499</sup> Minsky, An Interview with Marvin L. Minsky, 17–18.

<sup>500</sup> McCarthy, 'An Approach to the Artificial Intelligence Problem', 3.

<sup>501</sup> On the myth of automation see: Mindell, *Our Robots, Ourselves*; Taylor, 'The Automation Charade'; Winner, *Autonomous Technology*.

search problem—how to speed it up,’ noted Ray Solomonoff, whose handwritten remembrances provide perhaps the only on-site record of the workshop’s proceedings.<sup>502</sup> As with the threshold between AI and automatic coding, exact terms for this enterprise remained unclear. In ‘Approach,’ for instance, McCarthy claimed that a program had to be given access to ‘enough facts,’ derived from observation and deduction, to determine its own procedures. He did not specify his criterion for what constituted a ‘fact’; only that a computer could be made to manipulate one.

McCarthy’s appetite for indeterminacy had its limit: he refused statistical methods like those pursued by Rosenblatt, whose exclusion from Dartmouth was likely due to relative professional obscurity. Statistical componentry lacked the modularity McCarthy sought to selectively manipulate and understand high-level mental behaviour, at least in the short term. He was not alone in this inclination amongst workshop participants at Dartmouth. Upon arrival, attendees purportedly looked up ‘heuristic’ in the dictionary to seek out common conceptual ground. While the term had been absent from the event’s initial proposal, a sub-proposal from Simon and Newell had hinted at their progress using heuristic programming to solve theorems from *Principia Mathematica* on the Logic Theory Machine.<sup>503</sup> Attendees More and Samuel, along with uncorroborated attendee David W. Hagelbarger of Bell Laboratories, had also developed their own heuristic and symbolic programs in the prior year; Samuel’s had even been profiled on television six months earlier.<sup>504</sup> Simon and Newell’s arrival that summer with results in hand solved what Solomonoff later called, ‘the demo to sponsor problem,’ meaning the felt need to oblige a hesitant Morison with a proof of concept.<sup>505</sup> In notes from Hanover, Minsky noted that ‘the inspiring progress’ of Newell, Simon and More had had ‘considerable effect on the direction of our work.’<sup>506</sup> McCarthy later designated the RAND group as, ‘the stars of the show.’<sup>507</sup>

---

<sup>502</sup> Solomonoff, ‘Dartmouth’, 18.

<sup>503</sup> For analysis on the distinction between heuristic and algorithmic approaches in this period, see: Boden, *Mind as Machine*, 711–12.

<sup>504</sup> Boden notes that Samuel did not bring a printout of his code to Dartmouth. Boden, *Mind as Machine*, 707; It is not clear whether Hagelbarger attended. Shannon claims he was working on a calculus machine at this time. Shannon, ‘Letter from Dr. Claude Shannon to Dr. Herbert A. Simon’; on More see: Kline, *The Cybernetics Moment*, 162.

<sup>505</sup> Solomonoff, ‘Untitled Notes Re: Wendy Conquest’; Solomonoff, ‘Dartmouth’, 19.

<sup>506</sup> Minsky, ‘A Framework for Artificial Intelligence’, 4 July 1956.

<sup>507</sup> McCarthy, ‘The Dartmouth Workshop--as Planned and as It Happened’.

Even with this milestone met, Kline notes that Dartmouth was not the breakaway success McCarthy had hoped it would be. Funding was halved, participants dropped out or left after a day and the group's diverse research foci sprawled.<sup>508</sup> The report outlines no formal intellectual output, nor do surviving records suggest that a research agenda was distilled and shared (as I return to in the next chapter, it took Minsky five years to publish his agenda). High hopes for the collective identification of a criterion of success were put off, and visions of a group project developing a chess or draughts (checkers) program puttered out. Minsky, McCarthy and Rochester did agree to develop specifications for 'a program to solve problems in plane geometry using both syntactic and semantic methods,' but the participants' list on that project was confined to Rochester, Minsky, and Herb Gelernter, with McCarthy joining as a consultant.<sup>509</sup>

To summarize, the Dartmouth workshop did not establish AI per se. After the event, the domain's viability, membership, research foci and methodologies remained plastic. This changed as McCarthy doubled down on the need for subsequent organization; a fitting tribute to two parents who had steeped his childhood in labour organizing. At the 1956 Symposium on Information Theory, held three weeks after the Dartmouth workshop, a special session entitled 'Automata' introduced pre-scheduled papers from Rochester, Newell and Simon, as well as impromptu remarks from Minsky, McCarthy and Solomonoff.<sup>510</sup> Kline characterizes the symposium as a foundational moment for other reasons: key papers by Noam Chomsky, George Miller, and Simon and Newell (on the Logic Theory Machine) helped to formalize cognitive science.<sup>511</sup> Miller himself described the meeting as a milestone in the decline in use of the term cybernetics.<sup>512</sup> It was also the point at which the term 'artificial intelligence' gained its first formal airing outside of Hanover. McCarthy used the opportunity to establish

---

<sup>508</sup> Kline, *The Cybernetics Moment*, 163.

<sup>509</sup> McCarthy, 'Report on Dartmouth Summer Research Project', 2.

<sup>510</sup> McCarthy, 'Letter to Dr. Robert S. Morison, The Rockefeller Foundation', 3 September 1956; Kline, *The Cybernetics Moment*, 162.

<sup>511</sup> Having missed a submission date, McCarthy was permitted to present briefly on events at Dartmouth but not on Simon and Newell's work, since they pushed back firmly on the implication that the Logic Theorist had come out of the summer meeting. See: Kline, *The Cybernetics Moment*, 164.

<sup>512</sup> Kline, 165.

a mailing list comprised of forty-seven receptive parties.<sup>513</sup> The event provided him at least one satisfaction that *Automata Studies* had not; in correspondence with a colleague he stated that ‘interest in automata’ had become, ‘as we call it now, the artificial intelligence problem.’<sup>514</sup> His organizing had begun to pay off.

Nofre et al. note that standardized terminology and pedagogical techniques played an important role in clarifying the epistemic status of programming languages in the mid-to-late 1950s. Accounts from Dartmouth participants after the summer workshop provide a window into this process of standardization around AI. Once back to IBM, for example, Rochester added artificial intelligence to the list of core foci in his Information Research department.<sup>515</sup> A May 1957 memo cites ten of thirty members of Selfridge’s division at MIT’s Lincoln Laboratory assigned to the study of pattern recognition and artificial intelligence.<sup>516</sup> In a 1956 letter to McCarthy, Solomonoff shared that he might be able to get funding for unspecific work on artificial intelligence under the banner of ‘Information Retrieval.’<sup>517</sup> Bernard Widrow, then a temporary faculty member at MIT working on computer circuits, recalled ‘When I came back to MIT, what I wanted to work on was artificial intelligence.’<sup>518</sup> ‘The Dartmouth Conference made that phrase dominate the others,’ McCarthy claimed later.<sup>519</sup>

Not all participants supported this imposition. The handful of researchers who had already demonstrated their results were hesitant to accept the new moniker. ‘The word artificial makes you think there’s something kind of phony about this, [like] there’s nothing real about this work at all,’ remembered Samuel.<sup>520</sup> In March 1956, and for years afterwards, Simon and Newell positioned their research on chess and theorem proving programs as *heuristic programming* and *complex information processing*. In their Dartmouth sub-proposal, they stated that only a portion of their research fell under the banner of AI.<sup>521</sup> ‘They didn’t

---

<sup>513</sup> McCarthy, ‘People Interested in the Artificial Intelligence Problem’. That November, McCarthy sent around a draft of Solomonoff’s ‘An Inductive Inference Machine.’

<sup>514</sup> McCarthy, ‘Letter to 2nd Lieut. Matthew Kabrisky’, 20 September 1956.

<sup>515</sup> Rochester, ‘Biographical One-Sheet’.

<sup>516</sup> Morison, ‘Memo Re: Oliver Selfridge, Wednesday, May 8, 1957’.

<sup>517</sup> Solomonoff, ‘Letter to John McCarthy’, 28 December 1956.

<sup>518</sup> Widrow, Bernard Widrow, an oral history.

<sup>519</sup> McCorduck, *Machines Who Think*, 1979, 96.

<sup>520</sup> McCorduck, *Machines Who Think*, 97; On his heuristic program see: Boden, *Mind as Machine*, 706.

<sup>521</sup> Simon and Newell, ‘Plans for the Dartmouth Summer Research Project on Artificial Intelligence’.

like the term AI really,’ opined Minsky later, ‘but I think that was sort of justification for the labs and the way they got funded,’ he speculated.<sup>522</sup>

Accusations of clannishness, professional nepotism and the perpetuation of cliques grew to surround AI in the years that followed.<sup>523</sup> Samuel recalled:

I’ve always objected to this in-group running things when you’re on the outside. And that was fostered by that meeting, I think. Not deliberately, but meetings of that sort tend to do that, and that’s my one objection to what’s been done in the field of artificial intelligence. It’s always been run as a sort of closed group.<sup>524</sup>

Indeed, of the twenty participants named by Solomonoff, twelve were affiliated with MIT, seven with IBM and three with Bell Laboratories, with some overlap.<sup>525</sup> In a manuscript drafted that July, Minsky described attendees as ‘The Artificial Intelligence Group.’<sup>526</sup> Rochester titled a photo (Figure 4.0) of core participants, ‘The Artificial Intelligencers.’ The historian James Fleck traces the genealogy of a first-generation ‘AI Establishment’ orbiting institutional centres in the U.S. and UK in the 1960s through to subsequent establishments in following decades populated largely by their students, with subsequent generations mirroring this homogeneity.<sup>527</sup>

Different explanations have been offered to contextualize this clannishness. McCorduck attributes insularity in U.S. based AI research to the field’s single major funding source, ARPA, in the 1960-70s.<sup>528</sup> Minsky attributed it to professional insecurity. ‘AI people were a sort of beleaguered minority, one that computer scientists, more than humanists, were prone to distrust given proponent’s claims to computationally derived novelty.’<sup>529</sup> Neither explanation is fulfilling. That the Dartmouth workshop existed at all was evidence that the earliest seeds of AI took root in existing social capital. Minsky’s remembrance downplays

---

<sup>522</sup> Minsky, An Interview with Marvin L. Minsky, 8.

<sup>523</sup> The most high profile being Joseph Weizenbaum’s 1976 critique of the ‘artificial intelligentsia’ in Weizenbaum, *Computer Power and Human Reason*, 179. McCarthy responded with, McCarthy, ‘An Unreasonable Book’.

<sup>524</sup> McCorduck, *Machines Who Think*, 1979, 130.

<sup>525</sup> Of the uncorroborated attendees, two were affiliated with IBM, two with Bell Labs, five with academia and one with the UK government.

<sup>526</sup> Minsky, ‘A Framework for Artificial Intelligence’, 4 July 1956, 2.

<sup>527</sup> Fleck, ‘Development and Establishment in Artificial Intelligence’, 209.

<sup>528</sup> McCorduck, *Machines Who Think*, 1979, 109–10.

<sup>529</sup> Minsky, An Interview with Marvin L. Minsky, 30.

the group's considerable access to institutional support even prior to ARPA's involvement. Ray Solomonoff had tried and failed in the early 1950s to bring his own team of scientists together around similar ideas.<sup>530</sup> In the UK, Donald Michie, who had worked under Turing at Bletchley Park, pursued his interest in machine thought as a hobby until 1962 due to limited AI opportunities domestically.<sup>531</sup>

It is not my goal to explicate this emerging hierarchy in detail—just to show that it existed. In broad strokes, Dartmouth served as a clearing house for relevant research and researchers, and an imperfect one at that. In Boden's view, the Logic Theory Machine was only '*so far as most of the participants knew*, the first functioning program devoted to a task normally thought of as requiring significant intelligence.'<sup>532</sup> Others had already considered and in some cases even operationalized heuristic or symbolic programs—although these projects did not command the attention of the Logic Theory Machine in Hanover, either because attendees were not present (Oettinger, Gerald Dinneen, Hagelbarger, the Booths) or, I speculate, because they did not arrive with compelling evidence in hand (Samuel, More, Selfridge), as the RAND group had.<sup>533</sup> In 'Machine Learning and Intelligence,' from their 1953 book *Automatic Digital Calculators*, Kathleen Booth and her husband outlined the basics of learning theory, symbolic reasoning, the role of reinforcement learning and the design of sensory organs, also outlining a shopping program as Oettinger had.<sup>534</sup>

Interpretations of the workshop's composition and proceedings need not draw attention away from larger trends then impacting the spread of computing in academia and elsewhere. Talent shortages for programming jobs were a chronic problem by 1956, referred to by employers as the 'persistent personnel problem.'<sup>535</sup> A host of the 1953 Eastern Joint Computer Conference lamented this challenge as 'one of the most difficult bottlenecks in the expansion of our field.' At the 1954 Conference on Training Personnel for the Computing Machine Field, which united industry (IBM, RAND, Bell), academia (Harvard, MIT), military

---

<sup>530</sup> Solomonoff, 'Dartmouth', 14, 9; Solomonoff, 'Ray Solomonoff and the New Probability', 40.

<sup>531</sup> Fleck, 'Development and Establishment in Artificial Intelligence', 183.

<sup>532</sup> Emphasis hers. Boden, *Mind as Machine*, 705.

<sup>533</sup> For more on some names in this list see: Boden, 701.

<sup>534</sup> Booth and Booth, *Automatic Digital Calculators*, 1953, 211–16.

<sup>535</sup> Ensmenger, *The Computer Boys Take Over*, 59.

(ONR) and government (the Census Bureau), one attendee acknowledge the ‘universal feeling’ among leaders about the need to solve this bottleneck.<sup>536</sup>

The first published use of the term ‘artificial intelligence’ occurred in this context. At the 1956 Symposium on the Impact of Computers on Science and Society, organized by the Institute of Radio Engineers Professional Group on Electronic Computers, attendees fielded proposals for how to address labour shortages. Most proposals, but not all, involved changes to educational norms, such as secondment programs to allow engineers to teach high school students.<sup>537</sup> One proposal called to simplify the use of a computer by automating the intellectual labour required to program it. By investing in research on automatic coding techniques, untrained students could participate without so high a degree of specialist knowledge. ‘Automatic coding techniques will unquestionably do much to simplify the problem of communication with machines and therefore to relieve the shortage of programmers,’ argued David Sayre, of IBM’s Programming Research Group and a core participant at Dartmouth.<sup>538</sup>

John Mauchly, of Remington Rand, responded to Sayre by stating, ‘It is certainly true that many of us are interested in what has been given the name “artificial intelligence.” ...It may be that ultimately the “artificial intelligence” which we have been discussing will be able to reduce some of that [programming work] load, but it will be quite a while before that’s done.’<sup>539</sup> Mauchly’s comment suggests he may have identified AI as an instance of automatic coding. In a September 1957 letter, McCarthy aligned the two domains as well. To Saul I. Gass of IBM’s Applied Science Division, he wrote, ‘My main field of interest in the artificial intelligence problem and my interest in automatic programming stems from a belief that automatic programming is a step in the direction of making machines behave in ways that would be regarded as intelligent were the behaviour human.’<sup>540</sup>

These passing references to functional similarities between automatic coding and artificial intelligence point to how the success of the former could underwrite, and perhaps even be confused with the plausibility of the latter, despite significant distinctions. AI

---

<sup>536</sup> Grosch, ‘After-Luncheon Remarks’, 124; Ensmenger, *The Computer Boys Take Over*, 56.

<sup>537</sup> ‘Symposium on the Impact of Computers on Science and Society’, 157.

<sup>538</sup> ‘Symposium on the Impact of Computers on Science and Society’, 149.

<sup>539</sup> Mauchly co-created ENIAC. ‘Symposium on the Impact of Computers on Science and Society’, 155; As chronicled in: Cordeschi, ‘AI Turns Fifty: Revisiting Its Origins’.

<sup>540</sup> McCarthy, ‘Letter to Dr. Saul I Gass, IBM Corporation’, 9 September 1957.

researchers had demonstrated little evidence of the viability of their nascent field. As discussed in Chapter Two, Simon and Newell rejected the option to classify their results as ‘AI’ in this period. Automatic coding, in comparison, was booming. During a speech at the 1957 Automatic Coding symposium at the Franklin Institute in Philadelphia, Charles Katz of Remington Rand stated, ‘In five short years, the field of automatic coding has grown so tremendously, that today there are more compilers than there are computers... Whereas just a few years ago it was unusual to have one paper on automatic coding presented at a computer symposium, today entire symposiums such as this are dedicated to the discussion of new developments in the field.’<sup>541</sup>

Automatic coding increased accessibility to digital stored-program computers by rendering functions in familiar language. ‘The pseudo-codes of new compiling systems more and more closely approach the languages familiar to the user,’ Katz explained.<sup>542</sup> FORTRAN drew on algebraic symbology; B-Zero, a business compiler, used a modified form of English. Russell C. McGee of the General Electric Company provided a definition of automatic programming that was premised on this notion of accessibility. ‘Automatic programming’ meant: a programmed mechanism that (a) resembled the language in which the initial problem was formulated (ex. English, algebra), (b) translated that language, via an assembly routine, into a machine language and (c) did so with minimal expenditure of time, money and manpower from the moment of problem origination to its solution. A cartoon from the event equated this process to human-to-human translation. It depicted a French man and an American speaking to each other in their non-native tongues. Between them was an interpreter who spoke each of the two men’s non-native languages. This interpreter symbolised the role of an automatic coding system.

---

<sup>541</sup> ‘Automatic Coding’, 17.

<sup>542</sup> ‘Automatic Coding’, 17.






Figure 5.0 – A cartoon depicts the function of automatic programming.<sup>543</sup>

Automatic coding, broadly speaking, decoupled programming from the peculiarities of idiosyncratic machines, which varied in word length, storage capacity and other modes of versatility. It simulated a more comprehensible ‘language’ with which non-specialists could configure instruction sets for the machine. This decoupling allowed accountants, statisticians, mathematicians and other uninitiated professionals to make new use of the medium. In correspondence with *Science* in 1957, McCarthy acknowledged this horizon, ‘I am convinced that really widespread use of automatic calculators by scientists depends on the development of compilers, etc.’<sup>544</sup>

For John W. Carr of the University of Michigan, AI was a natural extension of this uptick in accessibility. At the 1957 symposium, he defined automatic coding in similar terms to those that Katz had used. His definition of automatic programming included not only that the system could add useful symbols to its own vocabulary or compare between elements or structures within its own memory, but also that it could, eventually, leverage inductive manipulation in this process.<sup>545</sup> Carr named artificial intelligence explicitly when explaining the latter—meaning he ranked AI as a category of automatic programming. He cited Solomonoff’s still unpublished paper on an inductive inference machine as well as Simon and

<sup>543</sup> From the Historical and Interpretive Collections of The Franklin Institute, Philadelphia, PA. ‘Automatic Coding’, 57.

<sup>544</sup> Compiling techniques are a subset of automatic coding. McCarthy, ‘Letter to Mr. Earl J. Scherago, *Science Magazine*’, 29 April 1957.

<sup>545</sup> ‘Automatic Coding’, 67.

Newell's progress with the Logic Theory Machine as examples. Like McCarthy, Carr saw meaningful overlap between automatic coding and AI.

Popular accounts of symbolic AI tend to chronicle the field's initial development in relation to the 'transition' event McCarthy then sought—one in which the machine would act with perceived sentience. McCarthy and others hoped to develop a program that could accomplish such things, such as by improving itself iteratively in a manner that signified the looming consolidation of further *self*-improvements. These accounts interpret the historical significance of technical developments through the prism of their desired ends, not their actual day-to-day means. By contextualizing means rather than ends, one finds that McCarthy faced pressure from his mentor, Shannon, and original funding body, via Morison, to justify his heightened rhetoric. Dartmouth changed little in this regard, since Simon and Newell categorized their results as a separate enterprise. McCarthy benefitted, I argue, from the computing community's presumptions that AI fitted naturally within the rapidly advancing field of automatic coding. While this may have been true, those committed to the latter did not share McCarthy's radical vision of a machine epistemology, nor did they undertake their work for that goal. Histories of AI that center only on gestures towards a machine's potential ability for self-improvement miss and mask the role of the growing library of techniques (e.g., programming languages and the derivative mechanisms upon which they rely) undertaken below that threshold in the 1950s and prior, such as assembly languages and compilers. These advances, collectively won, put meat on the bone for McCarthy's comparatively personal and positivist dream of engineering a language for thought.

#### 1956-59: Programming as Applied Epistemology

During and after 1956, McCarthy benefitted from access to expanding networks of funding and institutional support that had been created and maintained for the development of automatic programming techniques. This support encompassed teaching placements and access to computing devices, which parlayed into new career opportunities and foundational new theory. After 1956, perhaps as a result of this new access, McCarthy began to challenge traditional approaches to epistemology using strongly positivist language and assumptions about the promise of AI. In this period, he developed the Advice Taker, a hypothetical

programming system that combined English with predicate calculus. His hybrid framework laid the basis for his List Processing Language, called LISP, which became the *lingua franca* of U.S. based AI research in the 1960-70s.

Whereas interest in AI spread slowly through select centres on the Eastern seaboard between 1956-58, interest in computing grew rapidly in the United States between 1955-60. In July 1956, halfway through the Dartmouth workshop, James Killian, president of MIT, opened access to the new, three-story, IBM-funded MIT Computation Center to twenty-five sister colleges and universities across New England.<sup>546</sup> A press release from December 1956 described the centre as, 'The largest and most versatile data processing facility yet to be made available primarily for education and basic research.'<sup>547</sup> As an appointee from Dartmouth, McCarthy was invited to make use of the 18,000-square foot facility and 25-member staff.<sup>548</sup> There he was given the chance to work with both IBM 702 and IBM 704 devices.

Tasked by his employers at Dartmouth College to educate new students in the major features of computing, McCarthy set about trying to build a curriculum. Personal notes from his efforts reveal how challenging it was—both intellectually and logistically—to program devices such as an IBM 704 successfully. McCarthy asked students to prepare a set of programming 'cards' that would be run, batch by batch, on a 704 located one hundred and thirty miles away from Dartmouth at the MIT Computation Center. Grading for the course was straightforward. 'Getting the program to work on the first try will ensure a high grade on the exercise and not getting it to work at all [after three attempts] will ensure a low grade,' McCarthy summarized.<sup>549</sup> Even a basic competency with such a complex system was worthy of a distinguished grade.

In a lengthy memorandum to Dartmouth staff, McCarthy jockeyed to convince his colleagues that the administrative burden involved in learning to program was not just beneficial but necessary. 'I think that sooner or later programming will become as basic a part of a scientific education as calculus, and strongly advise learning it even if you don't see immediate application to your problems,' he wrote, 'Admittedly, this is an extreme view.'<sup>550</sup>

---

<sup>546</sup> See: Aker, *Calculating a Natural World*, 286–88; Little, 'Philip M. Morse and the Beginnings', 146; McCarthy, 'History of LISP', 217.

<sup>547</sup> Aker, *Calculating a Natural World*, 287.

<sup>548</sup> Nilsson, 'John McCarthy, 1927-2011', 4.

<sup>549</sup> McCarthy, 'Untitled Note on Teaching Computing at Dartmouth'.

<sup>550</sup> McCarthy, 'Dartmouth Use of the IBM 704 to Be Located at MIT', 9.

He made plain his intention to ‘sell the idea’ to his peers in academia.<sup>551</sup> McCarthy positioned learning to program like learning a new language. It was a ‘capital investment’ that would allow staff to execute faster, cheaper and more accurate calculations than those generated by-hand. Computation also made calculations easier to reproduce and re-use. New England academics were receptive; forty percent of his students at the MIT Computation Center were professors.

McCarthy encouraged his colleagues to make use of the programming techniques refined by users elsewhere. He highlighted the work of SHARE, a computer user group launched in September 1955 as a coordinating body for, initially, IBM 701 users. A manual produced by the MIT Computation Center in 1957 described one of this group’s outputs, the SHARE Assembly Program, as a ‘common language for [IBM] 704 users.’<sup>552</sup> The Assembly Program provided a library of pre-existing programs that followed standards of nomenclature and mnemonics agreed upon by members.<sup>553</sup> This onerous process of standardization enabled *all* members—including new users like the staff at Dartmouth—to benefit from the network effect of a popular language. The more who used the SHARE Assembly Program, the more refined, expressive and intuitive its procedures became.

Amidst this organising, McCarthy’s career accelerated. In September 1957, with help from Kemeny, chair of the Department of Mathematics at Dartmouth, he became a Sloan Fellow in the physical sciences. The position allowed him to relocate to the MIT Computation Center full-time to develop new theory. In September 1958, he and Minsky co-founded the MIT Artificial Intelligence Project with two programmers, a secretary, a typewriting machine and six graduate students, a milestone I return to in [Chapter Five](#). In that same period, McCarthy became an assistant professor in communication sciences within MIT’s Electrical Engineering Department.

At the 1958 Mechanisation of Thought Processes Conference in Teddington, London, McCarthy presented his first major attempt at a goal he had first laid out at the Dartmouth workshop: to design a hybrid logical/natural language for AI—one that would endow a machine with a basic understanding of the world. His conference biography listed AI,

---

<sup>551</sup> McCarthy, 1.

<sup>552</sup> Arden et al., ‘Coding for the MIT-IBM 704 Computer’, Preface; Lorenzo, *Endless Loop*.

<sup>553</sup> Arden et al., ‘Coding for the MIT-IBM 704 Computer’, 140.

automatic coding and mathematical logic as his primary interests.<sup>554</sup> His paper, ‘Programs with Common Sense,’ was a collaboration with Minsky.<sup>555</sup> In it, McCarthy introduced his hypothetical computer program, The Advice Taker.

As the Advice Taker became foundational to symbolic AI, it is worth pausing briefly to summarize it. The system was designed to account for, model and manipulate high level ‘common sense’ abstractions, such as deciding how a user would go from location A to B. ‘A program has common sense if it automatically deduces for itself a sufficiently wide class of immediate consequences of anything it is told and what it already knows,’ McCarthy explained.<sup>556</sup> Schematically, the system worked as follows:

1.  $\text{at(I,desk)} \rightarrow \text{can}(\text{go}(\text{desk,car,walking}))$
2.  $\text{at(I,car)} \rightarrow \text{can}(\text{go}(\text{home,airport,driving}))$
3.  $\text{did}(\text{go}(\text{desk,car,walking})) \rightarrow \text{at(J,car)}$
4.  $\text{did}(\text{go}(\text{home,airport,driving})) \rightarrow \text{at(J,airport)}$ <sup>557</sup>

By fusing English and calculus, McCarthy believed he could eventually describe *all* possible mental abstractions in formal language. To model this language, he repurposed the ‘linked list’ structure used to represent information in the Logic Theory Machine, which I outlined in [Chapter Two](#). A central benefit of this system was that it equipped a computer with recursive control over its inputs and procedures. In principle, the machine could teach itself rather than relying on the contributions of a human theorist at each iterative step. To improve initial functionalities, the Advice Taker’s planning function, or memory, was to be freighted with a base level of ‘common sense’ that the executory function could then manipulate to draw conclusions from.<sup>558</sup> What made the system unique was that in operation it would, in principle, articulate both the heuristics it used *and* the formal system it followed to execute

---

<sup>554</sup> McCarthy, ‘Programs with Common Sense’, 1959, 75.

<sup>555</sup> Minsky, An Interview with Marvin L. Minsky, 11.

<sup>556</sup> McCarthy, ‘Programs with Common Sense’, 1959, 78.

<sup>557</sup> McCarthy’s Advice Taker was designed to use a combination of list structures and recursive expressions, such as declarative and imperative sentences, to represent objects such as ‘the year 1776’ and their individual property lists, such as ‘the year that the American Revolution started.’ McCarthy, 82.

<sup>558</sup> McCarthy, ‘An Approach to the Artificial Intelligence Problem’.

a decision. This would make visible, via a print-out, how the machine manipulated abstractions to solve everyday problems, revealing the recipe for intelligence in a machine.<sup>559</sup>

At Teddington, Yehoshua Bar-Hillel, an Israeli mathematician embedded in cybernetic thought, balked. 'McCarthy wants to say no more than that a machine, in order to behave like a human being, must have the knowledge of a human being.'<sup>560</sup> Bar-Hillel challenged a specific failing: the Advice Taker's ability to designate the relation 'at' as transitive.<sup>561</sup> Without this ability, every object in McCarthy's system landed in the same spatial neighborhood as every other object in existence, an outcome that clearly missed the relational logic of English that McCarthy so desired to make explicit.<sup>562</sup> Roughly speaking, Bar-Hillel's challenge was that representations of reality require tethering in some meta-level organizing principle. In physical reality, natural laws provide such a tether; the laws of motion and energy dictate broad classes of physical relationships for instance. To *simulate* this ordering, symbolic representations, such as those used in symbolic AI, required an equally resolving meta-frame to clarify the endless permutations of relations upon relations upon relations, ad nauseum.

McCarthy countered on strong terms. He conceded that his proposal rested on unspecified assumptions, but claimed that philosophy, in comparison, lacked the clarity and expressive potential to render the logic of thought legible. He stated:

Whenever we program a computer to learn from experience, we build into the programme a sort of epistemology. It might be argued that this epistemology should be made explicit before one writes the programme, but epistemology is in a foggier state than computer programming even in the present half-baked state of the latter. I hope that once we have succeeded in making computer programs reason about the world, we will be able to reformulate epistemology as a branch of applied mathematics no more mysterious or controversial than physics.<sup>563</sup>

This was not an isolated claim. In unpublished manuscripts from this period, McCarthy returned to this mathematical ideal repeatedly. At times, he used the term 'applied epistemology' and 'artificial intelligence' interchangeably. In these manuscripts, he

---

<sup>559</sup> McCarthy, 'Programs with Common Sense', 1959, 77.

<sup>560</sup> Bar-Hillel had attended the Macy Conferences. McCarthy, 85.

<sup>561</sup> If A was close to B and B was close to C, this would not mean that A was close to C.

<sup>562</sup> This critique is known as 'the frame problem.' McCarthy and Hayes, 'Some Philosophical Problems from the Standpoint of Artificial Intelligence', 31.

<sup>563</sup> McCarthy, 'Programs with Common Sense', 1959, 90.

positioned programming as the most rigorous methodology available to illuminate the true dynamics of epistemology. He wrote:

The sceptic has two courses open to him. First, he can doubt that a computer program can be constructed along the lines about to be proposed which will exhibit intelligent behavior. Second, he can propose a program constructed along the lines of his philosophy and try to convince us that it will behave intelligently. Once one system of epistemology is programmed and works no other will be taken seriously unless it also leads to intelligent programs.<sup>564</sup>

In this excerpt, McCarthy wagered that intelligence would inevitably submit to formal modeling. This is an important proposition to single out because it is a radical example of the nuanced criteria for success that machine intelligence researchers in the 1950s crafted and traded upon as they sought to establish measures of progress in a nascent domain. McCarthy cast epistemology as a zero-sum spectator sport: to be won or lost through feats of programming rather than via philosophical ruminations. He belittled the presence of complexity as a methodological error on the part of philosophers. Most perplexingly, he suggested that the validation of his enterprise would be uncontroversial, as if physics was uncontroversial. With meager evidence, he advanced the notion that intelligent behavior was simpler than it appeared, and that computing would prove this claim to be true.

Rhetoric was an important lever here, as it was for all the men I profile in this dissertation. In 1958, McCarthy predicated his ambitious vision for AI on the success of *future* discoveries. The Advice Taker was a hypothetical program, not an operational one. This was unexceptional in some sense, since the paper purposefully served as a statement on a fruitful direction, not a set of results. Complicating matters, however, was the sensational rhetoric that accompanied this update. The margin between high rhetoric and limited results, and between pen and paper proposals and actual machine activity, gives the impression that the Advice Taker was both a paper and a pitch. McCarthy appropriated backbone cultural processes, such as ‘advice’ and ‘common sense,’ *without citing evidence*; a peculiar closing-of-worlds that would have likely been received differently by a community less sympathetic to his desired ends.<sup>565</sup> Bar-Hillel, unconvinced, called it ‘half-baked.’<sup>566</sup>

---

<sup>564</sup> McCarthy, ‘Physical and Mental Events and Intelligent Machines’, 5.

<sup>565</sup> English has proven easier to model than many indigenous languages that defy reduction to Western logics. Wolf and Wolf, ‘Sacred Waveforms’; Archer et al., ‘Making Kin with the Machines’.

<sup>566</sup> McCarthy, ‘Programs with Common Sense’, 1959, 85.

Preliminary results and changes in the composition of the mid-to-late 1950s brain model community helped McCarthy to make this temporal lien. Simon and Newell's 'demo to sponsor' solution was one contribution. Von Neumann's untimely passing in 1957 and Shannon's loss of interest in automata was another—two senior, critical voices on the limits of machine intelligence had fallen quiet. The growing sophistication of automatic coding techniques, which McCarthy held up as synonymous with AI, offered a third source of perceived validation, as I have argued. By the 1960s, AI had also gained institutional legitimacy at reputable research centers like MIT, Stanford and Carnegie Institute of Technology, cementing the sort of disciplinary validation that had eluded first-order cybernetics.

Despite these developments, strong critiques persisted into the 1960s. In 1961, Jerome Wiesner, chairman of President John F. Kennedy's Science Advisory Committee (PSAC), hired Herbert Teager (MIT), Oettinger (Harvard) and John Griffith (IBM) to examine a request to fund an Institute for Non-Numerical Studies.<sup>567</sup> The report distinguished between *numerical* computations of well-understood physical phenomena and *non-numerical* computations of 'symbols, meanings, and decisions,' the latter encompassing AI, pattern recognition, simulation and information retrieval. In 1958, Wiesner had helped Minsky and McCarthy to establish their AI Group at MIT while serving as Director of the Research Laboratory of Electronics. The initial draft of his 1961 PSAC report took a harder line. The draft sparked controversy by critiquing the proliferation of military contracts undertaken to prove 'vague' and, in the case of AI, 'shallow' and 'wishful' theories using large-scale hardware procurements. Reviewers of an early draft were split over expectations for tractable theory. McCarthy, Minsky and Newell pushed back against its 'negative' and 'fantastic' tone.<sup>568</sup>

The revised and final 1962 report ultimately recommended against funding the Institute but not against continuing efforts in the area of 'man-machine cooperation.' This uncertain endorsement from Kennedy's office speaks to the liminal status of digital information processing techniques generally at that time. In 1958, the Communications of the Association for Computing Machinery (ACM), a new journal, fielded letters from practitioners

---

<sup>567</sup> The source of this request remains unclear to me. My account summarizes: Slayton, *Arguments That Count*, 74–77. I've not had access to the report itself.

<sup>568</sup> Newell, 'Comments on Ad Hoc PSAC Panel on Non-Numerical Information Processing', 15 December 1961; Slayton, *Arguments That Count*, 77.



around the country over what to name their emerging and often inconsistently titled professions. Suggestions included synnoetics ('science of the mind' in Greek), computer science and the generalizable notion of comptology (ex. nuclear comptologist, logistics comptologist).<sup>569</sup> A separate batch of letters to the editor cited: Turingineer, Turologist, Flow-Chartsman, Applied Meta-Mathematician and Applied Epistemologist.<sup>570</sup>

Amidst national deliberation over the uncertain status of non-numerical information processing, McCarthy developed a strongly positivist case for how AI would legitimize the last candidate title in this list, applied epistemology. 'Work on the artificial intelligence problem will settle the main problems of epistemology in a scientific way. Therefore, traditional epistemology should be abandoned as an intellectual discipline,' he summarized in unpublished notes.<sup>571</sup> McCarthy deemed classical formulations naïve, indefinite and uninteresting in comparison to programmed techniques—commitments not seen in ACM's letters to the editor, nor supported universally by colleagues. Hao Wang at the University of Oxford, who in 1958 proved over two-hundred theories of logic in three minutes without Simon and Newell's heuristic method, flatly rejected the notion that human mental processes reduced to finite rules that could be simulated.<sup>572</sup> McCarthy's desire to displace philosophy with programming—to ask 'How does it know?' rather than 'How do I know?'—was uniquely audacious among colleagues interested in mechanized logic and within the broad field we know as computer science.<sup>573</sup>

Of importance here is that McCarthy's vision would have been even more radical to propose and Sisyphean to realize if unaccompanied by the rapid maturation of automatic coding techniques between 1952-57. These developments allowed McCarthy's claims to *seem* plausible. Consider that, as mentioned, IBM claimed to have invested eighteen man years of work into creating FORTRAN in 1954-57.<sup>574</sup> Commercial programming languages and the user networks who refined them, like SHARE, converged in a manner that enabled decentralized

---

<sup>569</sup> See: Correll, Khodr, and Vanderburgh, 'Letters to the Editor'; Ensmenger, *The Computer Boys Take Over*.

<sup>570</sup> Weiss and Corley, 'Letters to the Editor', 7.

<sup>571</sup> McCarthy, 'Methodology of Work on the Artificial Intelligence Problem', 81.

<sup>572</sup> On Wang: Dick, 'After Math', 153–92; Lechner to Cooperstein, 'What Is a Self-Organizing System', 18 March 1959.

<sup>573</sup> McCarthy, 'Sophisticated Epistemology', 90.

<sup>574</sup> Backus et al., 'The FORTRAN Automatic Coding System', 188.

contributors to cooperate through a *centralized* medium. Progress for SHARE, the MIT Computation Center and IBM was also, broadly speaking, progress for McCarthy. All involved were interested in advancing refined and expressive programming techniques, albeit for different purposes.

This interplay can be traced at a granular technical level as well. In a 1959 article for the *New York Herald Tribune Engineers' News Supplement*, McCarthy lamented, 'The present bottleneck in writing heuristic programs is not in inventing heuristics but in writing the programs that make the computer use them. It may take six months to write a program incorporating an idea that can be described to another person in five minutes. Progress in artificial intelligence depends on better "languages"—i.e. symbolism—for communicating with computers.'<sup>575</sup> From 1956-62, McCarthy worked to translate his hypothetical Advice Taker system into an operational language, LISP. During this period, he borrowed techniques developed for other programming languages like IPL and FORTRAN. Between 1956-58, he implemented key ideas for LISP in a FORTRAN-based language called FLPL, inspired in part by Simon and Newell's Information Processing Language.<sup>576</sup>

Behind this process of language development was institutional support that tends to be treated as incidental in existing accounts.<sup>577</sup> It was not. In 1959, McCarthy credited IBM employees with inventing a key feature of LISP, the *cons* subroutine, which was used to create a new list from an existing list in memory.<sup>578</sup> Industrial investment in and refinement of programming techniques brought nuance to the interface language McCarthy sought between man and computer; indeed, his wish was to see epistemology programmed first so that it could *then* be used to remake that field as a branch of applied mathematics. In his 1978 paper 'History of LISP,' he credited his decision to write LISP on an IBM 704 to that company's financial support of AI through the Dartmouth workshop, MIT Computation Center and

---

<sup>575</sup> McCarthy, 'Getting Closer to Machines That Think'.

<sup>576</sup> McCarthy to Minsky, 'Comments on History of LISP Draft', 25 January 1978, 1; McCarthy, 'History of LISP'.

<sup>577</sup> On the manner in which students and civic organizations shaped the direction of personal computing, see: Rankin, *A People's History of Computing in the United States*.

<sup>578</sup> The *cons* subroutine, which is short for 'construct,' was and is a fundamental function in LISP. See 'Acknowledgements' in: McCarthy, 'Recursive Functions Symbolic Expressions and Their Computation by Machine, Part I', 195; McCarthy et al., *LISP 1.5 Programmer's Manual*, 2; Backus et al., *Fortran*.

Minsky's attempts to simulate plane geometry.<sup>579</sup> That IBM is not typically considered a 'founding father' of AI speaks to how great-man narratives erase structural contingencies by channeling decentralized developments through individual actors.

The importance of these distributed contributions to AI research can be seen in the lead up to the 1963 publication of *Computers and Thought*, an influential early textbook in AI that I explore in the next chapter. McGraw-Hill, the volume's publisher, was sceptical of its reception and asked to cut the size of the volume by a third prior to its publication. Co-editor Julian Feldman recalled, in relation to their wish to publish the book in the first place, 'Part of the motivation was an anti-AI movement claiming that AI was all smoke and mirrors. We had a collection of papers representing real accomplishments—a potent counter-argument.'<sup>580</sup> Newell later described operationalized programs as 'the coin of the realm' in AI research; a baseline for success that warded off those critical of the field's legitimacy.<sup>581</sup>

McCarthy's legacy should be understood through this lens. The history of AI in the 1950s is to no small extent a history of the rhetoric used by a handful of scientists and industrialists to position the shifting epistemic status of programming languages in that period in heightened terms. By design, industrial acts of commercial standardization and diverse research on technical consolidation lowered the bar to entry into the programming profession. McCarthy was an acolyte for *this* change while simultaneously advancing his own audacious claims about epistemology. McCarthy's tenacity as a community organiser intersected with both trends: it helped him to nurture the establishment of a discipline (AI) while simultaneously promoting aspects of a new profession (programming).

In relation to the first, between 1955-1959 McCarthy helped to initiate *Automata Studies* (1956), the Dartmouth workshop (1956), the AI mailing list (1956), the AI Group (1958) and the AI Laboratory (1959), co-formalizing both a controversial disciplinary title and an elite initial membership in the process. Over time, in relation to the second, these initial efforts toward centralization began to change in character. Each of the five milestones listed above required some degree of centralized coordination *from McCarthy*. Subsequent interventions

---

<sup>579</sup> McCarthy to Minsky, 'Comments on History of LISP Draft', 25 January 1978, 1; McCarthy, 'History of LISP'.

<sup>580</sup> See: Feldman, 'Computers and Thought—The Back Story', 56–57.

<sup>581</sup> Feigenbaum, Feldman, and Armer, *Computers and Thought*, iii. On the transition between 'weak' programs (e.g. mathematical frameworks) and 'strong' programs (e.g. operating programs) in this era, see: Boden, *Mind as Machine*, 704–5.

did not require the same degree of handholding primarily because they coordinated participants around a different medium: programming. This medium was highly collaborative, despite being decentralized. In the 1960s, McCarthy helped to initiate LISP and time-sharing, each of which leveraged computing as its *own* extraordinary coordination mechanism. Here, he shifted from developing a supposed language of thought to developing tools to equip researchers to reason about what they believed was the language of thought—and to do so in a manner that enabled them to draw on each other's work more directly. These developments have been conflated in accounts that treat symbolic AI in isolation from the history of programming.

After 1962, LISP development had become so 'multi-stranded' that McCarthy confessed that he struggled to keep up with it.<sup>582</sup> At the same time, in effect, AI became a standalone discipline, operating beyond the control of its earliest advocates. As it developed, critics began to challenge incongruencies they saw between the pragmatics of programming and the suggestiveness of AI rhetoric. Drew McDermott, a graduate-turned-critic of MIT's AI program in the 1970s, lamented his colleagues 'contagious' use of what he called, *wishful mnemonics*, meaning words and phrases that served as 'incantations' for a desired result rather than sober descriptions of a particular mechanism or function.<sup>583</sup> McDermott argued that these conventions had warped researchers' relationship to the epistemic significance of their designs. In the 1980s, Philip Agre, who also graduated from MIT's AI program as a critic, described a popular 1960 AI paper as trying to 'hypnotize its reader.'<sup>584</sup> Boden traces this phenomenon in relation to LISP when writing, 'LISP's major advantage—that it appeared to be using English words—was a mixed blessing... Too often, programmers unthinkingly assumed that the word like symbols in the program meant much the same as the

---

<sup>582</sup> McCarthy to Minsky, 'Comments on History of LISP Draft', 25 January 1978, 1. On the 'bricolage' of components that made LISP see: Priestley, 'AI and the Origins of the Functional Programming Language Style'.

<sup>583</sup> Special thanks to Momin M. Malik for this reference. McDermott lamented that AI programs were given aspirational titles like 'The General Problem Solver' rather than 'The Local Feature Guided Network Searcher.' The same was true of functions in various programming languages, which were titled after their intended purpose (i.e. GOAL) rather than their actual function (i.e. TRY-NEXT). McDermott, 'Artificial Intelligence Meets Natural Stupidity', 4–6. McDermott, 4–6.

<sup>584</sup> Agre, 'Phil Agre's Mind'.

corresponding words in natural language. As a result, they deceived themselves and others about the power and psychological relevance of their research.<sup>585</sup>

## Conclusion

Retracing McCarthy's early career reveals how he used the notion of an 'artificial intelligence' as its own wishful mnemonic. In the mid-to-late 1950s, AI was an aspiration, not just an activity. McCarthy leveraged provocative rhetoric to stake out his vision for a nascent field; a vision that some contemporaries resisted. Shannon dismissed 'intelligence' as a relevant descriptor in *Automata Studies*, in line with cautions made by von Neumann. Both Morison and President Kennedy's Science Advisory Committee remained tentative about the viability of mathematical languages of thought. Still, in the space of only five years, McCarthy managed, with help, to formalize a community and institutional home for the study of 'artificial intelligence.' In the late 1970s and mid 1980s, pupils of the MIT AI program like McDermott and Agre voiced criticism of the wishful language used to substantiate such claims, which they argued had distorted practitioners' sense of reality.

In this chapter, I have provided a detailed historical account of McCarthy's early career in order to situate his intellectual trajectory within the emergence of commercial programming languages, programming pedagogy, programming user groups, and most of all, automatic coding techniques. I have argued that the history of early AI should be read with these trends in mind, particularly since McCarthy and others colloquially equated AI to automatic coding techniques. This historiographical reframing helps to reveal the material culture active behind symbolic AI in the 1950s and behind McCarthy's unabashed vision for applied epistemology. Omitting this contextualization obscures the decentralized systems of knowledge creation and the complex built environments that such claims left implicit, as if the viability of programming techniques was self-evident, unconnected or predestined.

The 1950s provide a unique window into these imbrications because terminology remained amorphous. Matti Tedre characterizes electronic computing in the late 1940s and early 1950s as a Kuhnian pre-science.<sup>586</sup> In shop-floor and published debates, some of which

---

<sup>585</sup> Boden, *Mind as Machine*, 802.

<sup>586</sup> Tedre, 'The Development of Computer Science', 237.

I chronicle here, talk of ‘automatic calculating machines’ transitioned into talk of ‘electronic computers.’ By 1960, mention of ‘automatic programming’ and ‘automatic coding’ techniques became discussions about ‘software.’<sup>587</sup> ‘Automata’ and ‘brain models’ became ‘the artificial intelligence problem’ and then just ‘artificial intelligence.’ Other terms, like ‘machine learning,’ ‘complex information processing,’ and ‘applied epistemology’ were assimilated into this tradition as well, both by Rosenblatt and McCarthy and by those who have so far treated their work historically, as discussed in Chapter One.

During and after the 1970s, the momentum of symbolic AI slowed as researchers found diminishing returns in an exclusive focus on symbolic reasoning in the exploration of intelligence. In 1985, John Haugeland introduced the term ‘Good Old Fashioned AI’ or ‘GOFAI’ to characterize the take-up and subsequent downturn in symbolic AI research in the 1950-60s.<sup>588</sup> As I have shown, this characterization emphasizes fashionable research trends within AI but masks the ways in which the professionalization of computer programming also accelerated brain model research and development. By shifting from histories of scientific personality and research fashions to histories of language propagation, commercial interests and pedagogy, structural contingences come one step closer into view.<sup>589</sup>

---

<sup>587</sup> Campbell-Kelly and Garcia-Swartz, ‘Pragmatism, Not Ideology’, 234.

<sup>588</sup> Haugeland, *Artificial Intelligence*.

<sup>589</sup> Accounts of how English became the de facto language in contemporary world science may reveal parallels for how programming became a baseline norm in various institutional contexts. See: Gordin, *Scientific Babel*.

## Chapter Five – Contextualising Marvin Minsky’s Agenda for Artificial Intelligence, 1950-61

The best-selling 1963 anthology *Computers and Thought* has been described as the most influential volume in AI’s early years.<sup>590</sup> Anchoring the collection were two 1961 papers by Marvin Minsky, one that provided an overview of approaches and attitudes to AI—which served as a research agenda for new students—and the other an authoritative bibliography.<sup>591</sup> Kuhn argues that textbooks form the base of normal science by standardizing the theories that each new member to a scientific community must master in order to learn the trade.<sup>592</sup> Mahoney, similarly, argues that a research agenda speaks to consensus among practitioners over key problems, their priority and what constitutes a solution.<sup>593</sup> With this in mind, I chronicle key milestones in Minsky’s early career to contextualize his development of each paper. These provide a window into the pre-history of *Computers and Thought* and the structural and conceptual forces that shaped the crystallization of artificial intelligence into a discipline between 1955-63.

This chapter brings new archival materials to light. I consider unpublished papers, letters and unclassified reports from Minsky’s expansive yet largely unexplored personal archives to broaden historical understandings of his early intellectual trajectory. Existing

---

<sup>590</sup> As ‘a kind of early text book’ in Dick, ‘After Math’, 38; ‘The most important survey of early AI research’ in Olazaran, ‘A Historical Sociology of Neural Network Research’, 93–94; a ‘classic volume’ in Cordeschi, *The Discovery of the Artificial*, xix; ‘A good view of the state of affairs... [in] 1961’ in Dreyfus, *What Computers Can’t Do*, 43. The volume has been published in four languages and four editions, with 1313 citations on Google Scholar (fall 2019); Feigenbaum, Feldman, and Armer, *Computers and Thought*. For background by an author see: Feldman, ‘Computers and Thought—The Back Story’. Simon was Feigenbaum’s PhD supervisor. The volume began as a reading list for students at the Business School at the University of California, Berkley.

<sup>591</sup> The first, Minsky, ‘Steps Toward Artificial Intelligence,’ had 1932 citations on Google Scholar (fall 2019); it is described as ‘one of the most important papers in early AI’ by Olazaran, ‘A Historical Sociology of Neural Network Research’, 122; and as ‘seminal’ in Boden, *Mind as Machine*, 299. The second, Minsky, ‘A Selected Descriptor-Indexed Bibliography to the Literature on Artificial Intelligence,’ (hereafter ‘Bibliography’) has 89 citations (fall 2019) but would not have been cited for the same reasons as a paper. The bibliography is identified as a guide to research in that period in Cordeschi, *The Discovery of the Artificial*, xix

<sup>592</sup> Kuhn and Hacking, *The Structure of Scientific Revolutions*, 143, 43.

<sup>593</sup> Hashagen et al., *History of Computing*, 28.

accounts are primarily popular in nature, with notable exceptions.<sup>594</sup> I integrate primary source materials from his career prior to 1961.<sup>595</sup> My account deepens understandings of the role Minsky played in centring AI research around symbolic reasoning and notions of industrial efficiency.

My story proceeds in three parts. I first consider Minsky's initial training and his commitment to developing mathematics with a basis in neurology. He developed SNARC, for Stochastic Neural-Analogue Reinforcement Computer, the first mechanical neural network, as a doctoral researcher in mathematics at Princeton University in 1950. Minsky read heavily in cybernetic theory and mathematical biophysics in this period and was tentative about neural metaphors. In the mid-1950s, he shifted away from these domains in search of 'universal' decision procedures. Prior accounts attribute this shift to Simon and Newell's results with the Logic Theory Machine in 1956. I argue that his intellectual transition toward totalising disembodied logics had been in motion since at least 1954; early demonstrations of heuristic programming (from the RAND group and others, see [Chapter One](#), [Chapter Four](#)) simply accelerated Minsky's transition to symbolic reasoning.

The second piece of my story traces Minsky's five-year development of 'Steps Toward Artificial Intelligence'. This 1961 paper outlined an initial research agenda for AI. The paper had its basis in a version Minsky developed during the Dartmouth workshop, which he also co-convened, as, 'A Framework for Artificial Intelligence.' I chronicle various iterations of this framework between 1955-61 to reveal how Minsky came to see notions of mathematical efficiency as crucial to understanding what he believed to be the axiomatic properties of the dynamics of the human mind. Minsky's desire to formalise optimal search and storage techniques aligned him with leaders in industry, government and the military, who had their own purposes for developing related techniques.

In my final section, I explore the earliest implications of this entanglement between brain model theory and military-industrial administrative concerns. I characterise Minsky as having been influenced by the institutional contexts in which he worked. As a Junior Fellow

---

<sup>594</sup> In addition to those outlined in [Chapter One](#), the following works have dealt historically with aspects of Minsky's career: Olazaran, 'A Historical Sociology of Neural Network Research'; Boden, *Mind as Machine*; Seising, 'Marvin Lee Minsky (1927–2016)'; For a recent collection of familial essays about Minsky see: Minsky, Solomon, and Xiao, *Inventive Minds*.

<sup>595</sup> A special thank you to Gloria Rudisch, Margaret Minsky, Charlotte Minsky and their extended family for such generosity during archival visits to Boston.



at Harvard University between 1954-57 and as an Assistant Professor of Mathematics at the Massachusetts Institute of Technology after 1958, Minsky was embedded in environments that prized the design of archival and information-service infrastructures like research libraries. He engaged with these efforts and equated their design directly to AI. I argue that it is important to recognise that his two foundational contributions to the field—a research agenda and a bibliography—each attempted to structure and synthesize *existing* research. Whereas Simon’s theories can be understood via his engagements in administrative logics, Minsky’s theories can be understood through the lens of archival logics.

Revisiting Minsky’s early career from this perspective brings to light his entanglements with the aims of the U.S. joint services, whom he thanked for having contributed ‘in too many ways to cite individually’ in work leading up to ‘Steps Toward Artificial Intelligence.’<sup>596</sup> This entanglement has not received sustained treatment. I compare Minsky’s relaxed attitude about funding sources to musings by U.S. Air Force investigators whom, by the early 1960s, had come to see artificial intelligence and its sister domains as the basis for an epistemology of control. In the view of Air Force administrators, this epistemology could be scaffolded and *built* using computing technologies, circumventing perceived shortcomings in military logistics like the unpredictable influence of human emotions. Illuminating AI’s roots in information services rather than neurophysiology will help to orient research towards this feature of AIs continued engagement with archives and inventories. I end by introducing the strongly militaristic bent of the first job advertisements for AI researchers in *The New York Times* in 1962. I challenge that the expression Good Old Fashioned Artificial Intelligence misrepresents and obscures this legacy as apolitical or twee.

#### 1944-54: Neural Networks as a ‘Plausible Analogy’ for Mental Behaviour

Marvin Minsky was born in New York City in 1927. His parents were progressive members of the city’s influential Jewish intellectual community. In notes for an unpublished autobiography, he attributed his fascination in the dynamics of cognition to their erudite social network. ‘Adults were very interested in intelligence,’ he recalled.<sup>597</sup> Harvard

---

<sup>596</sup> Minsky, ‘Some Methods of Artificial Intelligence and Heuristic Programming’, 5.

<sup>597</sup> Minsky, ‘#Biogmm-A-G’, 7.

criminologist Sheldon Glueck, who studied whether juvenile delinquency could be linked to genetics, is said to have pointed a young Minsky to readings on intelligence testing by Alfred Binet and Lewis M. Terman. Binet, a French psychologist, co-invented the first practical IQ test to identify learning difficulties in children. Terman, a 'pivotal' figure in the American eugenicist movement, expanded on and radicalized Binet's framework by arguing that its results were generalizable and definitive, a bold assumption that Binet had resisted.<sup>598</sup> Minsky found IQ uninspiring; he would not return to it in print during his early career despite his active pursuit of related concerns.

As a physics undergraduate at Harvard University in 1946, Minsky encountered mathematical biophysics and cybernetics. In a 1989 oral history, he claimed to have read the Macy Conference proceedings 'word for word,' growing enamoured of the neurophysiologist and cybernetician Warren McCulloch in the process.<sup>599</sup> He pored over a half century of research in psychology and was drawn to debate over whether the dynamics of the nervous systems were continuous or discrete.<sup>600</sup> In Nicholas Rashevsky's 1938 book *Mathematical Biophysics*, which called for a 'systematic mathematical biology, similar in aim and structure to mathematical physics,' Minsky saw how axioms of neurophysiological behaviour might explain gross psychological phenomena.<sup>601</sup> Rashevsky's journal had opened the 'intellectual space' for McCulloch and Walter Pitts to publish, 'A Logical Calculus of Ideas Immanent in Nervous Activity' in 1943 after it had been rejected elsewhere.<sup>602</sup> The paper formally equated the behaviour of idealized neural networks to mathematical logic.

Minsky, a lifelong tinkerer, used his time at Harvard to pursue not just theory, but also hands-on experimentation in neurophysiology, psychology and computer programming. Under John Henry Welsh, a physiologist, he examined how to manipulate the claw of a

---

<sup>598</sup> Terman's framework, conventionally used in military recruitment, was also championed by those who fought to pass the xenophobic 1924 John Reed Immigration Act. Stern, *Eugenic Nation*, 19; Terman, *The Measurement of Intelligence*; on Binet see: Gould, *The Mismeasure of Man*, 181.

<sup>599</sup> Minsky, An Interview with Marvin L. Minsky, 21.

<sup>600</sup> Minsky, 'The Discrete Approach: Neural Net and Related Theories', 1.

<sup>601</sup> Rashevsky was a 'maverick' in relation to his contemporaries at the University of Chicago in the 1930s. See: Abraham, 'Nicolas Rashevsky's Mathematical Biophysics', 333; as cited in: Piccinini, 'The First Computational Theory of Mind and Brain', 181.

<sup>602</sup> See Chapter One for this paper's connections to cybernetics. Abraham, '(Physio)Logical Circuits', 22; McCulloch and Pitts, 'A Logical Calculus of the Ideas Immanent in Nervous Activity'.

crayfish by exciting and inhibiting certain nerves; under Marcus Singer, a neuroanatomist, he dissected a human brain—an experience with *actual* human neuro-anatomy that distinguished him from Simon and McCarthy.<sup>603</sup> In 1949, Minsky enrolled in a course on computer logic under Howard Aiken, designer of Harvard's Mark I computer, that taught him techniques for how to simplify Boolean circuitry.<sup>604</sup> In 1993, he described his undergraduate experience as having formed in him the opinion that the structure of brain fibres was akin to the circuits of a radio or a television set; a system of patterns to be rendered legible by acts of reduction and formalization, just as humans had done with patterns in music, mathematics and social affairs.<sup>605</sup>

Minsky set to work on his own mathematical brain-machine metaphor as a doctoral student at Princeton University in 1950-54, alongside John McCarthy.<sup>606</sup> His supervisor, Albert William Tucker, was a specialist in game theory known for having formalized the Prisoner's Dilemma.<sup>607</sup> In 1984, Tucker described Minsky's PhD as the most unusual thesis supervision he had ever conducted.<sup>608</sup> For one, the subject Minsky wanted to examine—automata—fell outside the purview of the mathematics department. Tucker arranged for the chairman of Princeton's biology department to check that Minsky's assumptions about the nervous system were reasonable from a physiological point of view. The chairman 'took no responsibility for what came out of those assumptions,' but claimed they were reasonable enough to proceed. 'It was really far-out at that time,' Tucker recalled.<sup>609</sup>

Minsky's project refined the notion that neural phenomena could be modelled with mathematics. His 1954 dissertation, 'Theory of Neural-Analog Reinforcement Systems and its Application to the Brain Model Problem,' questioned how a physical system could be designed to model and automate acts of memorization, recognition, attention or reasoning.

---

<sup>603</sup> Seising, 'Marvin Lee Minsky (1927–2016)', 26.

<sup>604</sup> Minsky claimed to have developed programs for use on Harvard's MARK I. He did not specify why he elected not to run them. Minsky, '#Biogmm-A-G'; Bernstein, 'Marvin Minsky's Vision of the Future', 14 December 1981.

<sup>605</sup> Minsky, '#Biogmm-N-Z'; Minsky, Solomon, and Xiao, *Inventive Minds*, 155.

<sup>606</sup> Both men shared Solomon Lefschetz, a specialist in algebraic topology, as their initial supervisor. They also shared a summer internship at Bell Labs under Claude Shannon in 1952. Unknown, 'Letter to Mr. Marvin L. Minsky from Bell Laboratories, June 6, 1952', 6 June 1952.

<sup>607</sup> Tucker became chair of the mathematics department after Lefschetz, inheriting Minsky as his supervisee. See: Seising, 'Marvin Lee Minsky (1927–2016)', 26.

<sup>608</sup> Tucker, The Educational Program at Princeton in the 1930s [Transcript no. 31], 10.

<sup>609</sup> Tucker, 10.

He sought an audacious goal: to ‘duplicate the activities of so-called “sentient” organisms’ by designing a system that was capable of complex behaviour yet straightforward enough that its structure could be understood, such as by generating complexity from simpler elements.<sup>610</sup>

Automata attracted interest from the top American mathematicians in the early to mid 1950s, including John von Neumann and Claude Shannon. Both developed mathematical metaphors between activity in the brain and activity in artificial information processing systems like a telephone network. Of special interest to von Neumann, a member of Minsky’s doctoral advisory committee, was how the brain managed redundancy to maintain resilience amidst injury—a capacity that promised rich applications if it could be understood. That these early concerns of Minsky’s had a military application was made clear in a 1953 letter he received from a researcher at the Jet Propulsion Laboratory at the California Institute of Technology, who asked how such techniques could be used to improve missile and control systems.<sup>611</sup>

Like von Neumann, but unlike McCarthy and Simon, a young Minsky took pains to qualify the neurophysiological basis for his mathematical postulates. In unpublished notes from 1950, he juxtaposed theory from two camps interested in the mechanics of logical nets: ‘the Chicago school’ (McCulloch, Pitts, Rashevsky) and ‘the Wiener school,’ for Norbert Wiener.<sup>612</sup> He reasoned that the promise of the Wiener school had already been exhausted; Wiener had proven an equivalence between discrete feedback mechanisms and reinforcement theories of learning posited by the behaviourist B. F. Skinner. In his view, this was not enough to account for higher-order levels of complexity, such as the mathematics of a goal-seeking machine possessed by a changing goal.<sup>613</sup> While papers by cybernetic scholars made up the bulk of his PhD bibliography, he hesitated to develop existing theory.<sup>614</sup>

Minsky built on cybernetic concerns by reassessing which biological evidence to try and model and which to ignore. He contrasted Rashevsky’s work on mathematical biophysics

---

<sup>610</sup> Minsky, ‘Theory of Neural-Analog Reinforcement Systems’, 1–1.

<sup>611</sup> Benesch to Minsky, 30 December 1953.

<sup>612</sup> Minsky, ‘The Discrete Approach: Neural Net and Related Theories’, 4.

<sup>613</sup> Minsky, 8.

<sup>614</sup> With the exception of ‘A Logical Calculus,’ published in 1943, each of his PhD citations was less than five years old by 1953. Minsky, ‘Theory of Neural-Analog Reinforcement Systems’, Bibliography.

with the research of Stephen Cole Kleene, a RAND mathematician and contributor to Shannon and McCarthy's 1956 volume *Automata Studies*. Kleene explored how stimuli altered the states of a neural network.<sup>615</sup> His model assumed that time in a network could be quantized into a sequence of discrete binary moments and that the total number of relevant cells could be treated as finite. Rashevsky's model, in contrast, required *no* time quantization and allowed a continuously variable (i.e. non-binary) internal state function. Rashevsky's assumptions, in short, made for a more plausible analogy to natural systems.

In his dissertation, Minsky eschewed the need for discrete time quantization in favour of an approximation to that effect. By his model, cells could fire at any time as opposed to in a mandatory sequence of discrete 'moments' indexed as integers. He preferred this formulation because he believed that biological cells did not operate in a binary fashion, as either 'quiet' or 'firing.' The internal state of his artificial cells was quantified as a *continuous* rather than discrete function, measured via a 'pulse history' like blood pressure. He described the binary approach on its own as an 'unnatural imposition' and titled his system as a hybrid 'neural-analog network.'<sup>616</sup> Within this framework, cells followed a *probabilistic* firing condition, meaning that whether or not one would fire was a condition of its pulse history rather than a binary internal state.<sup>617</sup>

Minsky's core conjecture was that a system designed in this way could be trained to learn via the mechanics of reinforcement theory. This theory—that an organism's previous experiences or history conditioned its response to environmental stimuli—had its roots in behaviourism, by then the dominant school of thought in American psychology. Behaviourism assumed that passive stimulus-response dynamics formed the basis for all mental activity. In 1926, Ivan Pavlov famously demonstrated that a dog would salivate in the presence of the person who fed them, not just in the presence of food.<sup>618</sup> Pavlov's findings epitomized a wave

---

<sup>615</sup> Later published as: Kleene, 'Representations of Events in Nerve Nets and Finite Automata'; Minsky to Kleene, Undated.

<sup>616</sup> Minsky, 'Theory of Neural-Analog Reinforcement Systems', 2–2.

<sup>617</sup> Minsky purportedly drafted a monograph entitled *Analogue of a Nervous System* in this period. A Princeton University Press reviewer deemed it a 'contribution to knowledge of the very first rank.' I have not yet located this document and cannot discern its relation to his dissertation. Bailey Jr. to Minsky, 12 June 1951; Bailey Jr. to Minsky, 8 April 1954.

<sup>618</sup> Pavlov and Anrep, *Conditioned Reflexes*.

of research over the first half of the twentieth century into how environmental stimuli conditioned reflexes in organic beings.

To test whether reinforcement theory could be operationalised in a computer to make a *machine* learn, Minsky developed the world's first neural network simulator in 1951, called SNARC, for Stochastic Neural-Analogue Reinforcement Computer.<sup>619</sup> SNARC was comprised of three hundred vacuum tubes that simulated forty artificial 'neurons' powered by 'a lot of motors' and coordinated via 'a surplus gyropilot [an automatic steering device] from a B-24 bomber.'<sup>620</sup> The prototype was paid for via a four-thousand-dollar grant from the U.S. Air Force secured by George Miller, a Harvard psychologist under whom Minsky had studied how language shaped human cognition.<sup>621</sup>

Minsky summarized how SNARC functioned in a 1981 profile in *The New Yorker*.<sup>622</sup> He equated the system's electrical signals to the behaviour of a rat in a maze. First, the 'rat' (i.e. an electrical signal) was entered into a random node in the neural network, or 'maze.' It would then 'learn' its way to a specified end point via random behaviour informed by references to traces of past activities. Successful attempts reinforced path dependencies, which increased the probability of future success. Weightings between each of the forty atomised nodes in the network (or 'Snarcs' to borrow Minsky's terminology) changed in a manner that, to him, resembled learned behaviour, seen as one Snarc moving to (or toward) another.<sup>623</sup> A system of hardwired lights chronicled this process of self-reinforcement in action.

SNARC introduced 'Chicago School' abstractions into the physical world. Minsky was simultaneously assured and tentative about his results. On the one hand, the analogue device demonstrated how a sophisticated electronic network could manage redundancy in a manner analogous to the human brain. The device continued to work towards its goal via reinforcement even with fuses blown and wires pulled, an achievement that provided an

---

<sup>619</sup> Contract AF33(038)14343. Minsky, 'Theory of Neural-Analog Reinforcement Systems', 4–33.

<sup>620</sup> A gyropilot is a non-magnetic compass used for the automatic steering of an aircraft or ship. Cited in: Bernstein, 'Marvin Minsky's Vision of the Future', 14 December 1981.

<sup>621</sup> Approximately \$31,500.00 USD in 2020. Officer and Williamson, 'Conversion', 2020; Minsky, An Interview with Marvin L. Minsky, 21.

<sup>622</sup> Bernstein, 'Marvin Minsky's Vision of the Future', 14 December 1981, Page unknown.

<sup>623</sup> Minsky's 'Snarc' echoed McCulloch's idealized notion of a 'psychon,' meaning a 'least psychic event.' Abraham, '(Physio)Logical Circuits', 7; Abraham, 'Microscopic Cybernetics', 8; Seising, 'Marvin Lee Minsky (1927–2016)', 26.

empirical basis for a theory that von Neumann had postulated but never operationalised.<sup>624</sup> On this count, results were thrilling; SNARC legitimized a key aspect of the mathematical metaphor between natural and artificial information processing systems. Heartened, Minsky wrote, 'There is no evident limit to the degree of complexity of behaviour that may be acquired by such a system.'<sup>625</sup>

When describing the character and significance of other achievements, Minsky was less assured. He made liberal use of scare quotes throughout his dissertation to convey the 'behavior' of his 'brain model' and its 'plausible' ability to 'consider' or 'learn' actions etc.<sup>626</sup> In some sense, this grammatical hedge was necessary; while Minsky's aim was to simulate 'learning,' 'memorization,' 'recognition,' 'attention' and 'reasoning,' the bulk of his research examined the mathematics of reinforcement theory alone. He conceded that this theory would be insufficient on its own to explain *all* behaviours of a complicated organism.<sup>627</sup> Models of prominent learning theories, such as associative learning (the notion that organisms could acquire related behavioural roles from related experiences or stimuli), would also be needed.<sup>628</sup>

Theoretical limitations and ill-defined rhetorical implications were ultimately deemed by Minsky to be inconsequential in relation to what he perceived to be an exigent need for new brain model theory. Medical professionals, by his account, lacked both the biological data and appropriate mathematical methodologies required to even begin to explain how synaptic properties connected to observable mental disorders. Neural networks were to bridge this gap.<sup>629</sup> Minsky's characterisation positioned the neglect of neural networks as imprudent, given medical needs, which added a subtle pressure to accept the tentative language of his theory.

---

<sup>624</sup> Minsky, 'Theory of Neural-Analog Reinforcement Systems', 157.

<sup>625</sup> Minsky, 1–5.

<sup>626</sup> Minsky, 1–3, 1–1, 1–5.

<sup>627</sup> Minsky, 4.3–4.4.

<sup>628</sup> Per Chapter Three, Frank Rosenblatt used this technique to develop perceptron theory.

<sup>629</sup> Minsky equated his 'geometric' approach to neural networks to Gestalt field theory. Field theory explored the holistic relations between an individual and its environment. Minsky described it as akin to 'network theory,' a term that left open the possibility that the two could be rectified mathematically, linked by the study of networks. How, exactly, remained unclear. Minsky, 'Theory of Neural-Analog Reinforcement Systems', 2–10.

To resolve this latent tension between pragmatism and plausibility, and to justify his ‘far out’ approach, Minsky returned frequently in his dissertation to the raw irreducibility of the brain’s physiological complexity. He acknowledged the futile prospect of accounting for *all* evidence-based considerations, which ranged from reinforcing biochemical agents to non-uniform time quantization between nodes, along with other motley chemical, mechanical and biological factors that, individually and together, defied consolidation into mathematical formula.<sup>630</sup> ‘If the transmission laws [of synaptic events] were completely known, they may be so complex that a precise computation of neural activity in a net would still be impractical,’ he concluded.<sup>631</sup> Neural networks were to resolve this dilemma by providing a plausible yet provisional first approximation to biological reality.

As his research matured during the second half of the 1950s, Minsky paid less attention to soberly framing his work against the unostentatious realities of neural complexity. Between 1955 and 1956, he slowly abandoned his initially liberal use of scare quotes around concepts like ‘learning’ along with the tentative commitments they signified. ‘Learning’ became learning—a shift that helped to redefine the term. By removing his quotation marks, Minsky implicitly advanced the claim (still with little to show for it) that such vocabulary was legitimate. He jettisoned his reliance on metaphor, in which a proposed action is purposefully not literal, and through this redefinition, asserted that his nascent area of experimentation was substantive of learning.

Two competing motivations lay behind this slow and invariable re-definition, both of which can be seen in Minsky’s PhD research. First, to be *pragmatic* about the development of new neural theory given an urgent medical need; a strategy that he presumed results would eventually vindicate. Second, to have his research sustain *plausible* ties to biological fact, which required that some phenomena be respected at the expense of others. Minsky’s view of learning as a mathematical process was to be the nexus between biological systems and their non biological correlates.

Inspired by Rashevsky’s unorthodox mathematical biophysics in the late 1930s and by McCulloch and Pitts’ theoretical neural networks in 1940s-era cybernetics, Minsky pressed forward with the development of SNARC in 1951 and then with the design of functional brain

---

<sup>630</sup> Minsky, 6–19, 6–24, 3–16.

<sup>631</sup> Minsky, 3–2.



models. SNARC sensitized him to the improbability of achieving accuracy in his line of experimentation. Progress would require concessions. Faced with an insurmountable neurophysiological complexity, Minsky turned to analogue computing to substantiate the viability of neural network theory and to reduce the number of candidate considerations to something more manageable. He concluded his PhD with the claim that, in the future, high speed digital computing would be an appropriate tool for new research given the enormous number of calculations involved in verifying neural schematics and in realizing networks capable of higher order behaviour.<sup>632</sup>

#### 1954-56: A Change of Mind – Minsky's Turn to Heuristic Programming

As Minsky's career advanced in the decade that followed, his research interests shifted further from the material to the immaterial. Concerns about neurophysiological evidence gave way to questions regarding mathematical efficiencies and heuristic techniques akin to those developed by Simon and Newell ([Chapter Two](#)). While this shift was informed, in part, by his experience at the formative 1956 Dartmouth Summer Research Project on Artificial Intelligence, which introduced him to various results in heuristic programming, it was not simply these results that changed his mind. Unpublished papers from the period 1954-56 reflect a similar trajectory in his thinking. In this period and afterwards, Minsky's research was informed by questions about search and information management in the adjacent context of industrial and military organizational processes.

Minsky graduated from Princeton University in 1954. He had considered an application to medical school but opted instead to join a new department at Tufts College in Boston devoted to systems analysis. Researchers at the RAND Corporation had developed this field to characterise the arbitrary risk landscape of the Cold War.<sup>633</sup> RAND attempted to recruit Minsky that February. They told him they were a more stimulating research environment than 'most universities' and that they had already recruited the top twenty-five young mathematicians in the country, along with other luminaries in the field. According to

---

<sup>632</sup> Minsky, 5–71.

<sup>633</sup> Singer to Minsky, 13 February 1953.

RAND, this group purportedly struggled to ‘resist’ the ‘charms’ of the organization’s military research projects, which served as the ‘breadwinning’ pillar of its research.<sup>634</sup>

At Tufts, Minsky sought out generalized applications for the sort of network dynamics he had developed during his PhD. In ‘Discrete Selection Processes,’ a 1954 report funded by the Office of Naval Research (ONR), he explored both the potential and recurring limitations of what he called ‘universal’ decision procedures.<sup>635</sup> As in the design of SNARC, Minsky devised various schemes that used tentative trial and error techniques to proceed toward a predetermined goal. He described his procedures as ‘universal’ because they exhibited features that he claimed were typically seen in a wide selection of systems, including fully automatic weapons systems, operations research procedures, industrial processes and, remarkably, in physiological processes.<sup>636</sup> As with Simon ([Chapter Two](#)), Minsky saw explicit commonalities between industrial selection processes and biological selection processes.

Ultimately, the inquiry was short lived. Six months after Minsky’s arrival, Senator Joseph McCarthy shuttered the Tufts department due to fears of communist subversion. From 1954-57, Minsky studied mathematics and neurology as a Junior Fellow at the Harvard Society of Fellows.<sup>637</sup> Letters of endorsement from Shannon, von Neumann and Wiener spoke to his rising status within a community of researchers interested in mathematical models of neural activity. Events at the 1955 Institute of Radio Engineers national convention lend detail to the concerns of that community at the time.<sup>638</sup> In a symposium on ‘The Design of Machines to Simulate the Behavior of the Human Brain,’ panellists McCulloch (MIT) and Anthony Oettinger (University of Cambridge) marvelled at the brain’s inimitable yet suggestive self-regulating structure. Oettinger predicted that this structure would continue to inform the design of control systems but that a second divergent path, divorced from biology, would inevitably take shape alongside it. Panellist Otto H. Schmitt (University of Minnesota) agreed that biological metaphors were being misused to describe a computer’s behaviour without

---

<sup>634</sup> Mood to Minsky, ‘Invitation to Join the Rand Corporation’, 26 February 1954.

<sup>635</sup> Contract: Nonr-494(03). Miles, ‘Operations Research: An ASTIA Report Bibliography’, 18.

<sup>636</sup> Minsky, ‘Discrete Selection Processes’, 1.

<sup>637</sup> Seising, ‘Marvin Lee Minsky (1927–2016)’, 26.

<sup>638</sup> Minsky is said to have organized a session on machine intelligence at the 1955 Western Joint Computer Conference, but proceedings do not mention him nor the term ‘machine intelligence.’ From: Husbands, Holland, and Wheeler, *The Mechanical Mind in History*, 399; revisited in Kline, *The Cybernetics Moment*, 159. See: Joint Computer Conference (1955: Los Angeles, 1955 JCC; Selfridge, ‘Pattern Recognition and Modern Computers’.

evidence. He equated this disingenuous rhetoric to ‘using a hypothesis to prove a hypothesis.’<sup>639</sup>

Panelist Nathaniel Rochester (IBM), with whom Minsky would go on to co-convene the 1956 Dartmouth workshop, took a stronger position on the potential behind such analogies, albeit in an indirect fashion. He shared how he and his team had used an IBM 701 to simulate and subsequently rework an aspect of Donald Hebb’s influential 1949 theory of cell assemblies.<sup>640</sup> Hebb posited that the arousal of neural pathways strengthened their interconnections, a framing that gestured toward a possible bridge between psychology and neurophysiology—the same basic premise that Minsky sought to validate using SNARC.<sup>641</sup> Rochester’s simulation of a Hebbian assembly netted a novel result. It indicated that overlapping networks must compete against and inhibit each other’s influence in order to function. Having simulated Hebb’s theory, Rochester’s team found that not doing so would overstimulate the relevant synapses with positive feedback loops.<sup>642</sup> This finding, deliberated on in collaboration with Hebb’s team, asserted the role of mathematicians and computer engineers in the development of brain theory.

Minsky, in attendance at the Symposium as an official panel questioner, pressed the group for insight on how to develop *more* specific models in this area. His colleagues cautioned him against the unfounded presumption that phenomena like memory could be sharply localized like the network of a telephone switchboard or conventional computer. Per Schmitt, the topic illustrated a fissure between two lines of experimental inquiry—one biologically inspired and the other not. It also revealed the delicacy with which brain model theorists had to deal with claims about neurophysiology, given a lack of consensus over the significance and generalizability of initial results like Rochester’s.

By 1955, as profiled in [Chapter Four](#), plans for the 1956 Dartmouth Summer Research Project on Artificial Intelligence were underway with Minsky at their centre. In correspondence with Rochester and McCarthy about the workshop’s proposal, Minsky pushed to adapt the group’s mission statement to reflect elements of the two diverging paths

---

<sup>639</sup> ‘Symposium The Design of Machines to Simulate the Behavior of the Human Brain’, 244.

<sup>640</sup> Minsky had been aware of Hebb’s work since at least 1950: Minsky, ‘The Discrete Approach: Neural Net and Related Theories’.

<sup>641</sup> Hebb, *The Organization of Behavior*.

<sup>642</sup> Boden, *Mind as Machine*, 279.

mentioned above. He foregrounded, as themes, learning theory and the need for precise descriptions of the *principles* behind the brain's physiological structure. Rather than proceed on the basis that there was 'nothing that the human brain can do' that a machine could not be made to simulate it, as the initial proposal read, he proposed to proceed on the basis that there existed, 'no aspect of *learning* or other features of intelligence that cannot in principle be *described* so precisely' that a machine could not be made to simulate it.<sup>643</sup> The final Dartmouth proposal reflected this change, edited to convey its inverse, such that *every* aspect of learning could in principle be precisely described.

To meet this target, Minsky doubled down on the explanatory power of mathematics. He explained to co-convenors that he had interfaced with the Rockefeller Foundation about having event participants demonstrate some reasonable degree of mathematical acumen to partake, a requirement that indicated the type of expertise he aimed to engage. It is not clear that any such requirement was ever instituted, but the suggestion indicated an inclination to travel down the second of Schmitt's two diverging paths, away from strict allegiances to biological expertise and evidence and toward mathematical understandings of neural behaviours.

This inclination was reflected in his surrounding work. In his proposed contribution to the 1955 Dartmouth workshop, Minsky set out to generalize his PhD research on reinforcement learning by modelling an idealized 'motor abstraction' operating in an abstract environment.<sup>644</sup> By the end of the summer, he hoped to model a correspondence between a 'sensory situation' and a 'motor situation,' which together worked towards a conceptual goal, such as 'imaginative' behaviour. In 'Some Universal Elements for Finite Automata,' published in McCarthy and Shannon's 1956 *Automata Studies* volume, he showed how mathematical objects could be ordered to build up complicated 'machinery,' meaning finite automata, capable of 'universal' mathematical functions.<sup>645</sup> This writing drew on neurophysiology for inspiration only. Both projects examined how simple abstract elements could be configured to build upwards toward complex behaviours, as he had explored in his PhD.

---

<sup>643</sup> Emphasis mine. Minsky, 'To Drs. J McCarthy and Nat Rochester, IBM, Poughkeepsie, N.Y.', n.d.

<sup>644</sup> Minsky et al., 'Proposal', 9.

<sup>645</sup> Minsky, 'Some Universal Elements for Finite Automata'.

During the Dartmouth workshop, Minsky resolved this ongoing work into ‘A Framework for an Artificial Intelligence,’ which served as the basis for his influential 1961 paper, ‘Steps Toward Artificial Intelligence.’<sup>646</sup> It was in this paper, drafted in Hanover, that he first connected his work explicitly to the notion of intelligence. His interlocutors’ influence echoed in the decision. Minsky stated that research on heuristic programming by Simon, Newell and Trenchard More (a masters student of Shannon’s) had had considerable effect on the direction of his thinking, as well as on the group as a whole.<sup>647</sup> Indeed, Minsky described the framework itself as a *heuristic*—a pivotal term that had appeared in neither his proposal nor in the initial workshop proposal. Introduction of the term ‘intelligence’ was likely due to McCarthy, who had used it to describe his own earlier work on automata.

In ‘A Framework for an Artificial Intelligence,’ Minsky speculated that intelligent behaviour emerged from the interplay of numerous overlapping mechanisms, which he called ‘blocks’ or ‘boxes.’ By describing their interplay, he sought to isolate their basic operations and to understand how the dynamics involved could be brought to bear on a broader class of problems. The resulting framework, which he defended as too tentative and primitive for him to strongly endorse, functioned as follows. ‘Abstractions’ were formed in the ‘Characterizer’ Box. After that, the ‘Method Box,’ ‘Clean-up Box’ and the ‘Evaluator’ determined how the system would process each input, such as the inputs needed for the discovery of proofs for theorems in mathematical logic, as in Simon and Newell’s virtual Logic Theory Machine.

Efficiency was central to Minsky’s framework. He designated abstractions that could not be represented precisely as ‘poor’ or ‘useless’ to the overall system. A good abstraction (the term Minsky used to mean computer program) would suppress information that was irrelevant to its goal. A good problem, similarly, would be ‘well-defined’ so that a program could solve it. Minsky gave the example of executing a search for a specific piece of semantic content in an expression—this problem was well defined. In not so many words, he equated the art of problem-solving to the efficient search of a mathematical matrix. He concluded that the most efficient system possible would come freighted with an ‘endowment’ of pre-existing models that indicated the favourability of certain actions or paths. This ‘Model Box’ of ‘conceptual and operational vocabularies’ would ‘represent a large part of that which is

---

<sup>646</sup> Minsky, ‘A Framework for Artificial Intelligence’, 4 July 1956, 1; Minsky, ‘Steps Toward Artificial Intelligence’.

<sup>647</sup> Minsky, ‘A Framework for Artificial Intelligence’, 4 July 1956, 2.

laboriously pumped into each child by its culture,’ he explained.<sup>648</sup> Even a small Model Box, or bank of explanatory models, would render an ‘astonishingly high intelligence,’ he speculated.

Unlike the flowcharts Rosenblatt developed for perceptron theory (see [Chapter Three](#)), which abstracted away the influence of a human operator, Minsky’s formulation integrated the role of a human overseer. Atop all the aforementioned boxes ran The Master Program, which decided when a problem had been definitively solved or if enough time had passed to stop its operation. Minsky suggested that if a mathematical mechanism could not be found then ‘some human is given the job’ to ‘be in charge of whatever evolutionary processes are to be applied to the sets of Methods and Characters (and perhaps Models).<sup>649</sup> This person would make ‘executive decisions’ on considerations like how to balance the number of search procedures available with the desire for efficiency.<sup>650</sup> In granting this possibility, Minsky unknowingly echoed aspects of Simon’s theory of *decision premises*, meaning the principle that those who sat higher up in the hierarchy of an organization could influence those below them by setting the parameters within which decisions were made by those lower down.<sup>651</sup>

Near the end of the Dartmouth workshop, attendees gathered to discuss undertaking a group project. Minsky resisted calls to study chess and pressed instead to outline a program to solve problems in plane geometry.<sup>652</sup> This approach satisfied his wish to link judgements about the outside world with measurable internal states. In ‘Notes on the Geometry Problem’ he stated, ‘One of the reasons Plane Geometry might be a rewarding domain for artificial intelligence is that there is a good chance that we could find a language that was simultaneously suitable for machine use and human use.’<sup>653</sup> Minsky cited Simon’s Logic Theory Machine as a precedent for how the propositional calculus used in Euclidean deduction could be leveraged ‘without great difficulty’ to execute high-level behaviours.

---

<sup>648</sup> Minsky, 15.

<sup>649</sup> Minsky, 19.

<sup>650</sup> Minsky, 13.

<sup>651</sup> McCarthy, in contrast, proposed to avoid such needs by having a system improve *itself*. See [Chapter Four](#).

<sup>652</sup> McCarthy et al., ‘Proposal’, 2.

<sup>653</sup> Minsky, ‘Notes on the Geometry Problem 1’.

Rochester and Herbert Gelernter, also from IBM, agreed to help with the project, with McCarthy serving as a consultant.<sup>654</sup>

After the Dartmouth workshop concluded, Solomonoff checked-in via a series of letters to see how Minsky's organization of 'The Brain Club' and a 'Brain Theory Group' was proceeding.<sup>655</sup> Minsky later credited Solomonoff with convincing him during the workshop that neural networks were 'not much good' compared to symbolic representations—that building mechanisms up toward complexity using simple, abstract elements would take too long to be worthwhile.<sup>656</sup> Solomonoff, a student of Carnap's at the University of Chicago, had tried similar techniques and found them lacking.<sup>657</sup> At Dartmouth, he developed a statistical system entitled 'An Inductive Inference Machine' that turned Minsky's attention toward how to describe abstract machines to manage well-defined analytical tasks.

Since Solomonoff's paper was pivotal for Minsky, it is worth summarizing it briefly. In 'An Inductive Inference Machine,' Solomonoff presumed that all possible or actual 'space-time configurations' could be reduced into finite 'words' that could be manipulated statistically to understand and operationalize their interrelations. 'Gasoline' was one such word, as was 'automobile.' His system was designed to classify all manner of events and their related outcomes using automated references to past events and outcomes.<sup>658</sup> On paper, the system could only learn elementary arithmetic operations but Solomonoff gestured toward far more sophisticated behaviours. The crux of this intervention, for Minsky, was procedural: Solomonoff had advocated to design theories of inferences *first* and *then* ask, 'How would I make a machine do exactly that?'<sup>659</sup> This shifted Minsky from studies of neural behaviour to

---

<sup>654</sup> Simon later wrote to Minsky that he had drawn heuristics from his high school textbook when he attempted to model geometry. Simon to Minsky, 27 September 1956; Wells and Hart, *Modern Plane Geometry*.

<sup>655</sup> Solomonoff to Minsky, September 28, 1956; Solomonoff recommended that Tom Etter join Minsky's 'Brain Theory Group.' Solomonoff to Minsky, Undated.

<sup>656</sup> Crevier, *AI*, 37; Boden, *Mind as Machine*, 894; corroborated in Solomonoff, 'Untitled Notes Re: Wendy Conquest'; Conquest, Drake, and Rockmore, *Mind in the Machine: The Discovery of Artificial Intelligence*.

<sup>657</sup> Solomonoff wrote two papers about the statistical analysis of neural networks with Anatol Rapoport in 1950-51. Solomonoff, 'Structure of Random Nets'; Solomonoff and Rapoport, 'Connectivity of Random Nets'; For more on Solomonoff's career see: Dowe, 'Introduction to Ray Solomonoff 85th Memorial Conference'.

<sup>658</sup> Solomonoff, 'An Inductive Inference Machine', 14 August 1956. Published as: Solomonoff, 'An Inductive Inference Machine', 1957.

<sup>659</sup> Crevier, *AI*, 37.

studies of something that resembled how to formalize strategic tasks and embody them in a machine.

This intellectual project aligned with the aims of the military, within which theories of learning and pattern recognition were of considerable interest. Minsky's participation in the Dartmouth workshop was paid for by the MIT's Lincoln Laboratory, a research centre jointly funded by the U.S. Army, Navy and Air Force.<sup>660</sup> In 1957, he joined the lab under Dartmouth collaborator Oliver Selfridge, a student of McCulloch's.<sup>661</sup> A May 1957 memo stated that ten of thirty members of Selfridge's division were assigned to the study of pattern recognition and artificial intelligence.<sup>662</sup> Selfridge himself was then developing a machine for the U.S. Air Force that could read unencrypted Morse code signals with an accuracy purported to be just below that of a human being. The learning system Selfridge developed for this challenge became 'part of the standard model of human cognition' in the West, galvanizing a branch of machine intelligence that later became known as supervised learning.<sup>663</sup>

Of special note here, given Minsky's involvement in the lab, was that Selfridge designed his learning framework in distinctly *administrative* terms, downplaying metaphors to the brain. Selfridge's paper, entitled 'Pandemonium: A Paradigm for Learning,' equated learning with pattern recognition.<sup>664</sup> Pandæmonium was the capital of Hell in John Milton's *Paradise Lost*—a place 'full of demons,' per the term's meaning in Greek.<sup>665</sup> Selfridge described his framework as an organization of 'demons,' meaning sub-routines, shrieking commands through a four-level administrative hierarchy in a manner that caused some orders to get through while others failed. Clusters of shrieks revealed weighed connections, which revealed a pattern.<sup>666</sup>

---

<sup>660</sup> Pay attributed in: Minsky, 'Heuristic Aspects of the Artificial Intelligence Problem', 6. Minsky remained in Selfridge's group until at least 1959: Selfridge, 'Quarterly Progress Report'.

<sup>661</sup> Minsky, An Interview with Marvin L. Minsky, 12.

<sup>662</sup> Morison, 'Memo Re: Oliver Selfridge, Wednesday, May 8, 1957'.

<sup>663</sup> Bjork, 'Selfridge's Milton', 9.

<sup>664</sup> Selfridge used multi-layered networks processing information in a parallel rather than serial fashion to detect features in inputs, such as weightings that suggested which codes translated to which atomised words. Selfridge, 'Pandemonium: A Paradigm for Learning'.

<sup>665</sup> Husbands, Holland, and Wheeler, *The Mechanical Mind in History*, 399.

<sup>666</sup> Selfridge, 'Pandemonium: A Paradigm for Learning', 516. Selfridge's architecture resembled Prony's pyramidlike 'monument of calculation,' designed for 'manufacturing' calculations. See: Daston, 'Enlightenment Calculations', 188.



Bjork attributes Selfridge's appeal with both cognitivists and connectionists to the ambiguity of this metaphor. It was neither organic nor synthetic, neither computer based, nor brain related, which meant that proponents of symbolic AI and of neural networks could each see their work in its premise. Selfridge, whose grandfather had founded the eponymous UK department store, had effectively modelled an administrative organization, although it is not clear that this was his source of inspiration. Bjork suggests that Selfridge's multilevel architecture was inspired by Milton's frequent revisiting of angelic ranks and titles. 'Selfridge's faith that computers, given the proper architecture, could learn without the guiding hand of a programmer parallels Milton's faith that readers, given a free press, could self educate without the regulation of a licenser, much less a king,' he argues.<sup>667</sup>

During his stint at Selfridge's Group 34 laboratory, Minsky drafted an unclassified technical report on the design of intelligent machines. This December 1956 paper iterated on his framework from Hanover. It was entitled, 'Heuristic Aspects of the Artificial Intelligence Problem'.<sup>668</sup> Minsky claimed that the report was part of his preparations for a book entitled *Learning Systems and Artificial Intelligence*, a project outlined in July 1957 but never finished.<sup>669</sup> The unclassified report, which speaks to the direction of his thinking at that time, began with an extended reflection on the meaning of intelligence, a concept that Minsky elected not to pin down with any absolute and concise definition. 'The things we are trying to accomplish are always related to some set of ad hoc ground rules, problems, and resources.'<sup>670</sup> Stated differently, what might seem to be intelligent at first glance was treated as banal once its inner mechanisms were understood. Intelligence was a moving target.

In a bid to identify reliable structures despite this indeterminacy, Minsky once again conceptualized intelligence in terms that foregrounded mathematical notions of efficiency. He wrote:

In exploring the 'artificial intelligence problem,' we are not searching for any kind of isolated solution to the question of 'what is intelligence and how can it be embodied in material systems?' Instead, we are searching for new and better ways of achieving performances that, at the moment, command our respect... A

---

<sup>667</sup> Bjork, 'Selfridge's Milton', 14.

<sup>668</sup> Minsky, 'Heuristic Aspects of the Artificial Intelligence Problem'.

<sup>669</sup> Minsky, 'Learning Systems and Artificial Intelligence'.

<sup>670</sup> Emphasis his. Minsky, 'Heuristic Aspects of the Artificial Intelligence Problem', i.

central requirement for a 'better' way to do something seems, in general, to involve some notion of efficiency.<sup>671</sup>

Minsky did not qualify *whose* preferences or 'respect' such a system would be configured for, just that any worthy judgement on the topic would be self-evident given the lofty design of the underlying schematics. This brazen rhetorical manoeuvre was a claim dressed as a hypothesis. Minsky presumed to speak on *behalf* of a universal perspective, one in which 'We could all agree [on what behaviours] embody, or reflect, intelligence.'<sup>672</sup> This curious position revealed both his training and milieu; for a small group of American mathematicians operating within elite institutions during the early Cold War, intelligence was tantamount to efficiency.

Minsky's framework for AI rendered 'intelligence' in the language of a mathematician. At first glance, it appeared to humbly concede that intelligence was inescapably defined by the subjective view of its observer. This concession appeared to lighten the philosophical commitments upon which his theory rest, since intelligence could not be modelled in any commensurable language, nor understood objectively as a single 'thing.' Simultaneously, however, Minsky posited that intelligence could be modelled as a matrix, or search space, populated by constellations of distributed skills endowed to individuals in varied measure. Beholding a familiar constellation of skills provided one means to judge another's intelligence or judge one's own analytic inadequacies. The more intelligent one was, the more familiar one became with the breadth of possible constellations.

At the heart of this framework was the claim that a mathematician's analytical toolkit—which dealt in matrixes, search space, analytical modelling, etc.—provided that expert with a privileged sense of the hallowed landscape upon which intelligence could be judged. Minsky equated a mathematician's trained intuition for what constituted an elegant proof to an AI researcher's trained intuition or 'sense' of what constituted the componentry underlying intelligent behaviour.<sup>673</sup> This could mean recognizing a new constellation of skills or finding intelligent behaviour compressed in an unanticipated constellation. As with mathematical proofs, the theory held, *efficiencies* in the process of compression were a reliable proxy for an actor's depth of insight.

---

<sup>671</sup> Minsky, iii.

<sup>672</sup> Minsky, i.

<sup>673</sup> Minsky, II–2.

Jennifer Karns Alexander chronicles the pre-history behind this type of abstract claim-making around efficiency. Alexander clarifies that pre-modern notions of efficiency were non-generalized.<sup>674</sup> This changed during and after the Industrial Revolution as efficiency shifted from its use in nineteenth century thermodynamic laws of energy transfer into being a household word in the United States and Europe. Due in no small part to the emancipatory narrative(s) of modernity itself, notions of 'efficiency' were deployed to reconfigure social and political ideals around industrial values and capacities. Modernity, with industry and technology as its vessels, promised epistemological certainty, unified foundational truths and limitless progress through constant refinement and discipline (along with technological conditioning). In this light, efficiency came to represent something more; a 'divine simplicity, economy, and power' born from the hyper-rational assignment of perceived resources.<sup>675</sup> Alexander cites Charles Darwin's theory of natural selection, which Minsky appealed to explicitly in his AI framework, as an extended metaphor of efficiency in nature. Alexander also cited Joseph Conrad's fictional character Marlow as epitomizing the term's recurring political content. Marlow mused that efficiency separated savages from the civilized.

Minsky acknowledged the cultural basis of efficiency in his framework but did not expand on it. He reasoned that an intelligent system should be programmed with some sort of initial 'black box' endowment equivalent to what a human child received from their 'culture.' He suggested, as a starting point, Euclidean plane geometry. In his view, plane geometry provided a language of accessible semantic content that both humans and machines could manipulate—a yardstick with which to benchmark iterative improvements in performance in both domains. Heuristic techniques were another endowment. A third option was to have a machine generate an endowment for itself via interactions with its environment.

These proposals were premised on Minsky's belief by 1956 that, 'Human beings are instances of certain kinds of very complicated machines.'<sup>676</sup> He equated his proposal to have a machine self-generate its own learning infrastructure to the success of natural selection, which he also conceptualized of in mathematical terms. In his view, natural selection was an efficient search function, one in which evolution granted survival to whichever organism

---

<sup>674</sup> Alexander, *The Mantra of Efficiency*, 9.

<sup>675</sup> Alexander, 7.

<sup>676</sup> Minsky, 'Heuristic Aspects of the Artificial Intelligence Problem', III–17.

assembled the best heuristics from experience.<sup>677</sup> The task for an AI researcher was to design models that could generate analogous *artificial* search functions—Rosenblatt’s perceptron theory self-generated heuristics for visual recognition. Self-generating automata were another candidate area of study.

Minsky’s commitment to heuristic programming techniques deepened after the Dartmouth workshop, but not enough to end his hopes for neural networks. He came to believe that the two techniques would prove commensurable eventually.<sup>678</sup> Bottom up discrimination techniques made possible using neural networks—such as an electronic ‘rat’ finding its way through an electronic ‘maze’—would eventually intersect with the descriptive techniques made possible by heuristic programming, which in principle might allow that ‘rat’ to characterize and describe its surroundings.

Between 1950-56, and particularly in between 1954-56, Minsky transitioned from pursuing cybernetically-inspired theories of mathematical brain models to new theories of problem-solving techniques oriented around abstract notions of efficiency. Lost in transition was his initial fidelity to neurophysiological ‘plausibility.’ New papers neglected to incorporate past concerns like the role of reinforcing biochemical agents or metabolic mechanisms. Medical application was also no longer mentioned. The development of heuristic programming engaged other sources of inspiration. Minsky named one plainly: culture. According to his theory, cultural cues endowed humans with the techniques needed to prioritize and order intelligent behaviours, a convenient yet ill-defined standing reserve for his intellectual project. If neural networks could not model these cues efficiently, then perhaps heuristic techniques could. Minsky thus turned his attention toward developing ‘*a better theory of the task*,’ rather than a better theory of the brain.<sup>679</sup>

#### 1956-61: Minsky Outlines a Research Agenda for AI

Minsky was far from alone in wanting to formalise effective search and retrieval procedures in the late 1950s. After becoming an Assistant Professor of Mathematics at MIT

---

<sup>677</sup> Minsky, III–5.

<sup>678</sup> Minsky, IV–10, IV–17.

<sup>679</sup> Emphasis hers. Boden, *Mind as Machine*, 894.

in 1958, he engaged more frequently with questions of efficiency in information retrieval. He wrote a paper on the structure of scientific publishing and spoke to industrial, academic and government managers about AI and heuristic programming. As one of America's largest non-industrial defence contractors at the time, MIT encouraged these engagements, which have gone unreported in histories of AI.<sup>680</sup> The mid-century is instead remembered as a folksy period of 'Good Old Fashioned Artificial Intelligence,' or GOFAI, following terminology introduced by the philosopher John Haugeland in 1985.<sup>681</sup> GOFAI locates symbolic AI as the product of a certain era, i.e. the 'good old fashioned' one, but it glosses over and thus shrouds the actual character of that era.

In the late 1950s, AI research found an institutional home at MIT. In 1957-58, Minsky and McCarthy engaged a handful of graduate students to establish the Artificial Intelligence Group there. McCarthy had created a mailing list for researchers interested in the topic after Dartmouth.<sup>682</sup> In 1958, he joined MIT as Assistant Professor of Communication Science. From this post, the two men engaged various academic and industrial parties interested in AI. In 1957, John W. Carr of the University of Michigan invited Minsky to speak about AI to academics and computer manufacturers at an event entitled, 'Where do Computers – Hardware and Programming – Go from Here?'<sup>683</sup> In prepared notes, Minsky outlined the group's vision. 'The language of programming is better than that of learning theory to describe learning theory,' he stated, echoing the sentiment of McCarthy's provocation that epistemology reduced to applied mathematics.<sup>684</sup>

Minsky's desire for new search techniques aligned him with scientific communities otherwise disinterested in the mathematics of abstract neural architectures. In August 1957, he submitted a paper entitled 'A Proposal for Improvement of Scientific Publication and Reference Services' to the 1958 International Conference on Scientific Information, which explored effective storage and retrieval mechanisms for scientific research.<sup>685</sup> The paper outlined an information management system in which scientists subscribed to 'follow' work by other scientists they respected. Minsky argued that the requisite technology for such a

---

<sup>680</sup> Leslie, *The Cold War and American Science*, 14.

<sup>681</sup> Haugeland, *Artificial Intelligence*.

<sup>682</sup> McCarthy, 'People Interested in the Artificial Intelligence Problem'.

<sup>683</sup> Carr to Backus et al., 'Memo', 17 June 1957.

<sup>684</sup> Minsky, 'Re: Course Material', Estimated 1957.

<sup>685</sup> Minsky, 'Re: Int. Conf. on Scientific Information', 16 August 1957.

system already existed; data simply needed to be centralized into a processing faculty and tested with powerful computers and ‘a great deal of research into the automatic programming language problem.’<sup>686</sup>

A scribbled note in Minsky’s copy of the event’s programme connected these impressions to AI. It stated that ‘the artificial intelligence work fits here’ next to the conference’s suite of offerings on ‘Possible development of a general theory of storage and search.’<sup>687</sup> The development of such theories was one of the conference’s seven research pillars; all seven overseen by representatives from government, industrial or military organizations.<sup>688</sup>

Contrast this group’s composition and goal with those of the RAND Corporation’s 1958 ‘Simulation of Cognitive Processes’ summer school, where Minsky taught. An event brochure outlined that, ‘Participants will study the rationale and technique of using computer programs as theories of human problem solving, concept formation, and social interaction.’<sup>689</sup> At the International Conference, AI-related theories of storage and search were judged by administrators in the context of scientific publishing infrastructure(s). At the RAND summer school, AI-related theories about the simulation of information processing structures in the central nervous system were judged by social scientists in the context of human cognition. Simon and Newell, who hosted and instructed the RAND event, celebrated such parallels between administration and cognition. In a 1959 talk, Newell described AI, complex information processing and heuristic programming as synonymous.<sup>690</sup> The duo held that complex processes were functionally equivalent regardless if they were performed in administrative or neurophysiological systems. Minsky inherited this equivalence as he embraced symbolic reasoning.

The parallel between neural and infrastructural procedures was, of course, not as straightforward as these men suggested. In the context of scientific publishing, for instance, the design and implementation of search procedures was *managed* by human

---

<sup>686</sup> Minsky, ‘A Combined Publication, Retrieval, and Research Facility’, 8.

<sup>687</sup> ‘The 1958 International Conference on Scientific Information’, 3.

<sup>688</sup> Sheppard to Minsky, ‘Re: International Conference on Scientific Information’, 20 March 1958.

<sup>689</sup> ‘Summer Research Training Institute on Simulation of Cognitive Processes, 30 June to 18 July 1958’, 1.

<sup>690</sup> Newell, ‘Talk on “Heuristic Programs”’.

administrators—secretaries, directors, boards, scientists—held accountable for their bureaucratic decisions via heterogeneous acts of deliberation, voting and, realistically, pure chance.<sup>691</sup> At the RAND summer school, researchers developed theories of cognition using programs to ‘simulate’ human problem solving, collapsing search procedures into a *seemingly* disembodied state. That the two endeavours lived side by side—with AI researchers in between—suggested not only that bureaucracy could inform understandings of neural activity, but that neural activity could inform notions of bureaucracy. This touch of scientificity lent an aura to disembodied forms of social organization, the perceived value of which fluctuated with AI’s indeterminate purchase on scientific authority. That AI decision tools are now used to structure content for billions of users speaks to the missing histories of bureaucracy, automation and power that connect this formative mid-century period to the digital infrastructure guiding core aspects of contemporary life.

Tension between information-as-administration in one setting and information-as-psychology in the other remained unresolved in the years that followed. It is beyond the scope of my inquiry to deconstruct this broader phenomenon, but the portion of Minsky’s career I examine herein provides a window into evolving norms. Consider the 1958 ‘Mechanisation of Thought Processes’ conference at the National Physical Laboratory (NPL) in Teddington, London, which Boden describes as a ‘memorable catalyst for the growth of an intellectual community.’<sup>692</sup> Here, Minsky presented his research alongside leading figures in cybernetics (Donald MacKay, W. Ross Ashby, McCulloch), computer programming (Grace Hopper, John Backus), physiology (Horace Barlow, Rosenblatt), AI and brain modelling (McCarthy, Selfridge) and a bulk of delegates from industry and government.<sup>693</sup> The event’s programme reflected this pairing between academia and industry; topics ranged from mechanical language translation and artificial thinking to industrial planning and clerical mechanization.

This pairing was not unique to Teddington. In November 1959, MIT hosted an Industrial Liaison Symposia on Artificial Intelligence, where Minsky gave a paper on heuristic

---

<sup>691</sup> ‘International Conference on Scientific Information’, Area 7.

<sup>692</sup> Boden, *Mind as Machine*, 336.

<sup>693</sup> The NPL itself existed as a coordinating body for the standardisation and regulation of scientific measurements across the British Empire. Sutherland, ‘The Mechanization of Thought Processes’; *The Future of Artificial Intelligence: Views from History*. See Simon Schaffer: 12:15-15:10.

programming.<sup>694</sup> One month later, he was invited to give five lectures at the University of Michigan on 'Programming Concepts, Automata, and Adaptive Systems' to students from industrial, government and university laboratories.<sup>695</sup> The published version of his final Teddington talk stated, 'The work leading up to this paper was supported, in too many ways to cite individually, by the joint services of the U.S.A.'<sup>696</sup>

Minsky's Teddington paper, 'Some Methods of Artificial Intelligence and Heuristic Programming,' refined, once again, on the framework he had first developed at the Dartmouth workshop and at Selfridge's lab at MIT. While the bulk of the paper improved on prior theory, this iteration marked a departure from Minsky's roots in cybernetics. By 1958, citations of papers by Selfridge, Newell and Solomonoff displaced references by Wiener, von Neumann and Rashevsky. Minsky alluded to growing consensus amongst his interlocutors that 'hierarchies of organization,' meaning symbolic reasoning, would be required to express higher order concepts in physical representations such as neural networks.<sup>697</sup>

The community Minsky alluded to took on additional shape that next year. In 1959, he and McCarthy's AI Group at MIT became the Artificial Intelligence Laboratory.<sup>698</sup> In a 1981 profile for *The New Yorker*, Minsky recounted the promotion as a product of serendipity. Jerome Wiesner, director of the Research Laboratory of Electronics at MIT, purportedly happened on he and McCarthy in a hallway. Intrigued by what he heard about AI, Wiesner offered support and they named a price: fifty thousand dollars. Wiesner, who had recently secured a new joint-services contract, obliged in exchange for their taking responsibility for training six new graduate students.

The serendipitous arrangement speaks to MIT's lucrative position as a top military contractor. An abundance of military resources meant that neither Minsky nor McCarthy would have to write a grant proposal for years to come. In 1963, J. C. R. Licklider, a colleague turned director at the U.S. Advanced Research Projects Agency (ARPA), endowed the lab with additional financial support related to Project MAC. After 1962-63, 'We had a million dollars a year, sort of forever,' Minsky claimed in a 1993 interview, echoing comments he had also

---

<sup>694</sup> 'Third of the Fall Series on Industrial Liaison Symposia', 3.

<sup>695</sup> Holland to Minsky, 23 December 1959.

<sup>696</sup> Minsky, 'Some Methods of Artificial Intelligence and Heuristic Programming', 5.

<sup>697</sup> Minsky returned to this subject two years later in: Minsky and Selfridge, 'Learning in Random Nets'.

<sup>698</sup> Bernstein, 'Marvin Minsky's Vision of the Future', 14 December 1981.



made in an 1989 oral history.<sup>699</sup> As an indication of the relative intellectual freedom they enjoyed, Minsky later attributed his decision to quit the AI Laboratory in the 1970s to change in policy that required researchers to justify their work to funders.<sup>700</sup> Minsky stepped down as the Lab's co-director in 1974.<sup>701</sup>

Despite having secured an institutional footing by the late 1950s, the foci of AI research remained broadly amorphous amongst related work on self-organizing systems, automata, complex information processing and other areas of learning theory. This outlook slowly changed after Teddington when Minsky was approached to prepare a paper for a special 'Computer Issue' of the *Institute of Radio Engineers Proceedings*, the second of its kind since 1913.<sup>702</sup> Guest editor Harry T. Larson requested that Minsky draw together state-of-the-art and independent efforts of 'higher-level information processing functions such as learning and problem solving.'<sup>703</sup> 'Higher' in this context, Larson clarified, meant 'moving toward the processes human intelligence can perform,' a characterization that once again blurred the boundaries between theory and application by suggesting that the former could accomplish the latter.<sup>704</sup>

In 'Steps Toward Artificial Intelligence,' the resulting paper, Minsky delivered on the framework he had developed since 1956. First came a disclaimer: there did not yet exist a generally accepted theory of intelligence. Second, a proposal based in pragmatism: why not attempt to derive a theory of intelligence, nonetheless? Third, an ingredient list of methodological areas to pursue: search, pattern-recognition, learning, planning and induction. This time, Minsky also included a bibliography of roughly one hundred recommended readings—a testament to the field's rapid growth. He explained that the overview provided would outline an admittedly rough and subjective 'flavor' of projects

---

<sup>699</sup> I have not yet been able to confirm these figures through other sources. Minsky, An Interview with Marvin L. Minsky, 9. A partial transcription of the 1993 interview appeared in: Minsky, '#Biogmm-N-Z'. By 1981, they were reported to have operated on a budget of two and a half million dollars: Bernstein, 'Marvin Minsky's Vision of the Future'.

<sup>700</sup> This mandate was born of the second Mansfield Amendment, which in 1973 required the ONR and ARPA to show direct military applications for any projects they funded. Minsky, An Interview with Marvin L. Minsky, 27.

<sup>701</sup> Seising, 'Marvin Lee Minsky (1927–2016)', 30.

<sup>702</sup> Larson, 'The Computer Issue'.

<sup>703</sup> Minsky, 'Steps Toward Artificial Intelligence', 8.

<sup>704</sup> Larson, 'The Computer Issue', 5.

‘directly concerned with getting machines to take over a larger portion of problem-solving tasks.’<sup>705</sup> The nature of problem solving, he made clear, was core to the study of artificial intelligence.

Seising states that it was at this juncture in Minsky’s career that he transitioned into philosophical analysis of brain behaviour.<sup>706</sup> I have shown how the process was in fact more gradual, and that his theory and rhetoric overlapped with other groups in this period interested in the study of search and storage procedures. I have also argued that this direction of research was used to rationalize eschewing the messier realities of biological systems.

‘Steps Toward Artificial Intelligence’ was not the only vessel for these ideas that Minsky published in 1961. In March of that year, he also published, ‘A Selected Descriptor-Indexed Bibliography to the Literature on Artificial Intelligence’ in the *Institute of Radio Engineers’ Transactions on Human Factors in Electronics*. This paper, too, would be re-published in the influential 1963 volume *Computers and Thought*. ‘An invaluable component of such a reference volume is a good bibliography,’ the editors wrote, ‘We believe that Minsky’s descriptor-indexed bibliography will be a particularly useful tool for the researcher.’<sup>707</sup> In the paper, Minsky provided a list of approximately five hundred AI-related readings, five times as many as he had recommended in ‘Steps.’ ‘This particular field is still young,’ he explained, ‘but there are already many instances in which workers have wasted much time in rediscovering... schemes already reported.’<sup>708</sup>

There are many ways to interpret the contents of a large bibliography. Minsky himself confessed that he had not had a chance to read every item he cited. Even he, then, did not have a precise means to characterise the exact ‘flavor’ of AI research he set out as authoritative. Since nearly the full list of papers that Minsky cited in ‘Steps Toward Artificial Intelligence’ were included, again, in the expanded ‘Bibliography’ paper, it is possible to isolate the specific categories of research and the authors—within the broader literature of brain model theory—Minsky prioritized in the abridged version. As stated in this chapter’s introduction, the significance of this prioritization is that the paper (both papers, in fact, but

---

<sup>705</sup> Minsky, ‘Steps Toward Artificial Intelligence’, 28.

<sup>706</sup> Seising, ‘Marvin Lee Minsky (1927–2016)’, 28.

<sup>707</sup> Feigenbaum and Feldman, *Computers and Thought*, v.

<sup>708</sup> Minsky, ‘Bibliography’, 39.

‘Steps’ especially) went on to inform a new generation of AI researchers when republished in 1963 in *Computers and Thought*.

To demonstrate Minsky’s priorities, I take an enumerated approach. The table below (Figure 6.0) examines citations included in both papers. It lists the degree of overlap between Research Categories assigned to papers in ‘Steps’ and to Research Categories assigned to papers in ‘Bibliography.’ Juxtaposing the two reveals the relative priority of different Research Categories. At the top, with the most instances of repeat papers, were Planning Schemes, Heuristics, Search Problems and Problem-Solving. Topics treated as low priority in ‘Steps’ included Servomechanisms, Cybernetics and Mathematical Theory of Computers and Automata, which received roughly a quarter as many duplicate citations.

<u>Categories in ‘Bibliography’</u>	<u>% of Overlap with Bibliography in ‘Steps Toward Artificial Intelligence’</u>
F: Planning Schemes	52.00%
J: Heuristics	49.32%
C: Search Problems	48.35%
G: Problem-solving	38.92%
I: Inductive Inference Schemes	36.17%
E: Learning Systems	33.11%
D: Pattern Recognition and Perception	32.99%
H: Languages	22.98%
K: Theories of brain function	21.59%
M: Memory and Information Retrieval	21.57%
L: Epistemological Questions	19.70%
B: Computer Structures	16.67%
N: Servomechanisms and Stability Mechanisms, Cybernetics	13.98%
A. Mathematical Theory of Computers and Automata	12.50%
P: Some Special Categories	9.00%

*Figure 6.0 – Overlap by category between bibliographical items in ‘Steps Toward Artificial Intelligence’ and ‘Bibliography.’ Eighty eight of ninety-nine items in ‘Steps’ included. Seven did not repeat.*

Minsky stressed that both documents were imperfect. ‘Steps,’ he stated, represented *his* preferences for the field.<sup>709</sup> Figure 6.0 gives shape to these preferences. It reveals the degree to which Minsky advocated a pivot away from earlier theory and towards the study of the symbolic reasoning techniques he had come to endorse. A second table, below, provides another set of characteristics to consider. Columns two and three in Figure 7.0 flag that research by Minsky’s inner circle figured prominently in his ‘Bibliography’ paper. In this case, the prominence was a matter of conceptual breadth. To summarize briefly, Minsky categorized each citation in ‘Bibliography’ with a set of labels derived from a taxonomy of sub-categories in AI research that he had created. These sub-categories derived from the fifteen core categories listed in Figure 6.0 (ex. Planning Schemes, Heuristics). Some papers were deemed to apply to many subcategories. Others were deemed to apply only to a few.

---

<sup>709</sup> See: Minsky, 39. Some categories had fewer total papers and were thus easier to represent, as a sample, in full.

<u>Article Number</u>	<u>Relevant Sub-Categories</u>	<u>Author</u>	<u>Title</u>	<u>Published</u>	<u>Sponsored by U.S. Military</u>
328	23	M. L. Minsky	'Heuristic Aspects of the Artificial Intelligence Problem'	1956	Yes
332	18	Minsky and O. G. Selfridge	'Learning in random nets.'	1960	Yes
363	18	J. C. Shaw, H. Simon, A. Newell	'A variety of intelligent learning in a general problem solver'	1960	Yes
330	15	Minsky	'Some methods of heuristic programming and artificial intelligence'	1959	Yes
451	15	Selfridge	'Pattern recognition and learning'	1956	Yes
541	15	N. Wiener and A. Rosenblueth	'Cybernetics'	1948	Unknown
333	13	Minsky	'Steps toward artificial intelligence'	1961	Yes
323	12	G. A. Miller, E. Galanter, K. Pribram	'Plans and the structure of behaviour 1'	1960	Unknown
326	11	Minsky	'Neural Nets and the Brain Model Problem'	1954	Yes
20	10	W. R. Ashby	<i>Design for a Brain</i>	1952	Unknown
352	10	Newell	'The Chess Machine'	1955	Yes
355	10	Shaw, Simon, Newell	'Empirical explorations of the Logic Theory Machine'	1957	Yes
90	9	A. N. Chomsky	'Review of B. F. Skinner's 'Verbal Behavior''	1959	Unknown
151	9	R. M. Friedberg, B. Dunham, J. H. North	'A learning machine part II'	1959	Unknown
357	9	Shaw, Simon, Newell	'Elements of a theory of human problem solving'	1958	Yes

Figure 7.0 – Items in 'Bibliography' labelled as relevant to the highest number of sub-categories. E.g., Article 323 was labelled as relevant to 'Human learning behavior,' 'Conditional probability,' and ten other sub-categories.

Minsky labelled his own work as relevant to twenty-three sub-categories, the widest number of sub-categories of any paper. In line with this persuasion, eleven of the fifteen most multifaceted (by this measure) papers were either his own or were authored by Dartmouth workshop colleagues like Selfridge, Newell and Shaw. A third of the citations included under a special condensed reading list termed 'Selected Reading' were also authored by core Dartmouth workshop participants, with a total of around half by other Dartmouth participants or invitees. Since Minsky knew this research best and had only had occasion to read approximately half the papers that were included, this bias was not unfounded. He made

explicit that he had indiscriminately admitted papers on problem-solving machinery into the 'Bibliography,' while mathematical, psychological and physiological papers had faced tougher scrutiny.

Rendering Minsky's citations in 'Bibliography' in this enumerated format calls attention to a specific way in which his platform as a spokesperson for AI in the early 1960s helped him to assert his perspective of where the field should lead thereafter. In 'Steps Toward Artificial Intelligence,' Minsky was candid about this prerogative. He included a disclaimer that the paper was a summary of his own preferences and an incomplete one at that. He also cautioned that the 'Bibliography' paper was incomplete. He did not, however, describe the latter as the product of his own preferences. This lent an air of objective authority to the most value-laden issue under review: the definition of intelligence. His emphasis on colleague's work hinted at the sort of clannishness in AI that Arthur Samuel and others later criticised.

This format permits an additional window into the foundations of AI. As shown in the right-hand column of Figure 7.0, ten out of the fifteen papers indexed as the most multifaceted works in the discipline as of 1961 were funded either directly by the U.S. military or through an intermediary such as MIT's Lincoln Laboratory or the RAND Corporation. This audit foregrounds that the sciences of pattern recognition, learning routines, problem-solving theory and other related domains had been of sustained interest to the U.S. military. This invested interest would have been easy to miss in the density of the 'Bibliography' paper. Minsky did not gloss over the military's involvement in AI. He thanked Alice M. Pierce of the Air Force Cambridge Research Center for a 'very large and well selected compilation [of papers], which accounts for most of the citations I have not seen myself.'<sup>710</sup> Pierce's initial draft included a sweeping four hundred and eleven titles, roughly one hundred short of Minsky's final bibliography.<sup>711</sup> By this measure, the U.S. Air Force, like Minsky, were invested in clarifying the foundations of artificial intelligence.

---

<sup>710</sup> Minsky, 40.

<sup>711</sup> Minsky cited compilations by Newell, Oettinger and R. A. Kirsch. Each cut the literature across different lines. Minsky and Pierce clustered work around Pattern Recognition, Problem Solving and Learning Routines. Pursuant with Minsky's advice, Pierce cut research on physiology, psychology, logic, automata theory decision theory, theory of games and neural network theory. Pierce, 'Bibliography on Artificial Intelligence: First Draft'; Pierce, 'A Concise

## 1961-1970s: AI and Military Infrastructure

With notable exceptions, prior accounts of AI's history have broadly overlooked the implications of these structural entanglements with U.S. military priorities.<sup>712</sup> Boden, for instance, is critiqued for offering a foundational account of cognitive science that omits the role that political and cultural forces had in shaping that field.<sup>713</sup> Mirowski, similarly, criticizes McCorduck for disregarding that the history of AI 'passes directly through the history of operations research.'<sup>714</sup> So far in this chapter, I have brought to light similar entanglements through the early career of Marvin Minsky in particular. Intersections include funding for SNARC (U.S. Air Force), Minsky's decision procedure research at Tufts (ONR), his participation in the Dartmouth workshop (MIT Lincoln Laboratory), various milestones in the five-year development of the 'Steps' paper (U.S. joint forces) and later career teaching and development (RAND Corporation, MIT military-industrial affiliates).

Using this group of four key researchers as a proxy for broader trends, I speculate that military funding of brain model research in the mid-to-late 1950s was not coordinated by any clandestine or coherent mission. In 'Bibliography,' for instance, Minsky stated that he had not included literature on operations research due to unfamiliarity with its contents, a challenge to Mirowski's critique. Minsky also claimed, however, to have made the omission despite feeling 'sure that that is a valuable source of heuristic analysis and techniques.'<sup>715</sup> In the early 1960s, however, as theory coalesced and AI matured into a discipline with secure institutional backing, connections between brain model theory and military aims became more obvious to

---

Bibliography of the Literature on Artificial Intelligence'; Hobbs to Minsky, 14 May 1959. Pierce, as both a woman and military researcher in the 1950s, would have lacked the professional network then available to an MIT professor, which prevented her gaining due credit for this considerable undertaking.

<sup>712</sup> Edwards, *The Closed World*; Fleck, 'Development and Establishment in Artificial Intelligence'.

<sup>713</sup> Cohen-Cole, 'Review of Mind as Machine'.

<sup>714</sup> Operations research was developed during World War II to exact precision warfare through cost-benefit analysis of possible actions related to man-machine systems. Mirowski, 'Book Review', 135.

<sup>715</sup> Minsky, 'Bibliography', 40.

the patrons of that domain, enabling them to see for themselves how AI could be brought to bear on military aims.

A 1963 report by the Cornell Aeronautical Laboratory in Buffalo, New York, where Rosenblatt was based, provides a window into this alignment. The report, entitled, 'Information Processing Relevant to Military Command,' surveyed Minsky's bibliography along with bibliographies by Pierce, Newell, Oettinger, Kirsch and others. In the eyes of its authors, the connections between brain model theory and U.S. Air Force command technologies ran deep. The report concluded:

The link between language and epistemology defines the single most important front for an advance in information processing technology... Natural language breaches the interface between conscious reasoning and the underlying mechanisms and serves as the medium for the conscious organization, transmission, storage and retrieval of information.<sup>716</sup>

While there existed few field applications for such theory in 1963, the report recommended that automation of specific functions in military command situations be undertaken, as it would improve various command, management and control systems in the future.<sup>717</sup>

Programmed epistemologies were cast as a new front in military strategy.

This endorsement was not just speculative. The report recommended that assessments be made of the material costs of adopting such systems, including their space consumption and power requirements. AI and related techniques were treated as a new source of superior decision-making in military systems, one that would allow workaround for human vulnerabilities like 'fatigue, saturation, confusion, and emotions.' The fast, cheap and iterative application of even crude approximations to human decision-making methods were estimated to outperform 'mediocre' decision systems if judged in aggregate. The report proposed that brain model research was to substitute hierarchies of human operators with methods that used simulation and automated heuristics instead. The earlier parallels, discussed above, between search procedures in scientific publishing and in search procedures in psychological at RAND, had become intersecting lines.

In the report, this techno-bureaucratic trajectory was painted as a long-term strategic horizon. Using language that echoed Rosenblatt's positivist orientation to theoretical

---

<sup>716</sup> Murray, 'Information Processing Relevant to Military Command', 138–39.

<sup>717</sup> On later efforts to automate international diplomacy see: Rohde, 'Pax Technologica'.



statistics, the authors speculated that human values could be automated and routed using the yet-to-be-realised powers of statistical mechanisms and information management. ‘Subjective human values may be assigned a measure or at least ordered, and used in “objective” mechanized decisions,’ it read.<sup>718</sup> They advocated engaging with management science, which held ‘obvious implications for the management of military affairs.’<sup>719</sup> Self-organizing systems were to reveal the basis for learning across myriad subject areas, from morphogenesis in ecology to group dynamics in economics and socio-economics. Like Minsky—and perhaps following him—the authors assumed that such systems could be mathematically ordered ‘in that special way which we call intelligent.’<sup>720</sup>

These overlaps between military priority and mathematicians’ designs for a simulated intelligence did not just play out in esoteric military reports and hybrid conferences. After 1960, mainstream advertisements for AI researchers began to be published in *The New York Times*. Their content speaks to AI’s embeddedness in Cold War military-industrial prerogatives in the United States. A handful of ads printed between 1960-64 foregrounded the tensions underlying AI research. In one light, questions of information management structures appeared routine, obscure or even banal. Titles like ‘Systems Implementation Engineer’ and ‘Information Systems Engineers’ were listed side-by-side. In another light, new techniques conjured feelings of rapid technological development and raw power. An ad for the International Electric Corporation (Figure 8.0) captures this tension. It shows a white hand reigning in computer tapes like the steer of a chariot next to the title, ‘Securing man’s control of the complex weapons he has invented.’<sup>721</sup> The ad makes a macho appeal to the prospective candidate to come aboard and help military intermediaries reign in the sheer power of large-scale command and control systems. Under ‘Programmers/Analysts’ is listed a call for experts in artificial intelligence.

Another ad (Figure 9.0) revisits these tensions. This time, the reader peers through the peaked cap of a military commander to find a team of computer operators deliberating inside, as if their judgements and expertise were the basis for that commander’s thinking and authority. The ad somersaults military hierarchy by implying that new recruits would guide

---

<sup>718</sup> Murray, ‘Information Processing Relevant to Military Command’, 141.

<sup>719</sup> Murray, 121.

<sup>720</sup> Murray, 128.

<sup>721</sup> ‘Display Ad 701 -- No Title’.

the judgement of senior authorities as a result of their technological expertise. By joining, they would learn 'How to Inform the Military Commander.' Once again, under 'Programmers/Analysts' is listed a call for experts in artificial intelligence. A 1962 ad for the same organization depicted figures running to prepare a jet for take-off. Next to them read the word 'Decision' in all-capitalized bold lettering. In this instance, the speed at which computations were made in a computer was used a proxy for decisive leadership. 'Command of global air forces, both manned bomber and missile, leans heavily upon instant decision-making,' it read.<sup>722</sup> In each ad, motifs of macho earned authority were rehearsed in an appeal to apply artificial intelligence to the real world.


---

<sup>722</sup> 'Display Ad 552 -- No Title'.

**ENGINEERS, PROGRAM ANALYSTS:**

## securing man's control of the complex weapons he has invented

Factoring into large scale command and control systems the latest advances in all relevant technologies is a continuing responsibility of ITT's International Electric Corporation.



Currently, IEC engineers and scientists are adding a remarkable new capability to the SAC Command and Control System 465-L: data presentation in color.

Operating at speeds that appeared incredible only a short time ago, the system enables computer outputs to be converted to alpha-numeric form . . . photographed . . . developed and projected on control center screens in as many as 7 colors in a matter of seconds.

This new capability opens up a whole new field of data format techniques to be explored. An obvious and immediate value is the enhancement of human perception through color changes denoting differing degrees of situation criticality.

### OPPORTUNITIES IN MANY COMMAND AND CONTROL AREAS NOW OPEN TO SYSTEMS ENGINEERS AND SENIOR PROGRAMMERS

Many of these positions are on 465-L, for which IEC is providing systems management, development and design. Other opportunities relate to large-scale commercial digital communication systems, oceanic systems, and satellite control. Your inquiry about any of the positions listed below will receive immediate attention.

**PROGRAMMERS/ANALYSTS.** For real-time programming analysis and development. Broad activities encompass advanced programming systems, including special color display routines; diagnostic programs; automatic recovery; problem-oriented language; artificial intelligence.

**OPERATIONS ANALYSTS.** To establish systems requirements in satellite control, air traffic control, ASW and command/control. Also, assignments in man/machine communications and information retrieval.

**SYSTEMS IMPLEMENTATION ENGINEERS.** Electronic engineers to develop tests for stressing and evaluating communication-display-computer systems. Recommend improvement and refinements. Also, field positions for installation and integration of digital command/control systems.

**INFORMATION SYSTEMS ENGINEERS.** For design of command/control and advanced communications systems. Experience in traffic, antenna and propagation theory, and mathematics as applied to communications and space technology.

**DIGITAL SYSTEMS ENGINEERS.** Engineers with management ability to direct sub-systems engineering effort on a global command/control system. Experience is desired in message traffic control, data processing systems, data display and multi-sequencing techniques.

Write fully in confidence to Mr. E. A. Smith, Rm. 23-H, ITT-International Electric Corporation, Route 17 & Garden State Parkway, Paramus, New Jersey.

An Equal Opportunity Employer


**INTERNATIONAL ELECTRIC CORPORATION**

**ITT**

Figure 8.0 – 14 October 1962 – Artificial intelligence appears under 'Programmers/Analysts'.<sup>723</sup>

<sup>723</sup> 'Display Ad 701 -- No Title'.

## CRUCIAL SYSTEMS CHALLENGES OF THE '60's



### How to Inform the Military Commander

Today's decisions at the highest level of military command require a range, precision and speed of communication and information processing beyond virtually anything conceivable in the past. Further, optimization of the electronic portion of a command control system cannot be considered independently of the capabilities of the ultimate, human decision-maker in the chair of command.

A good case in point is the SAC global command and control system 465-L, for which ITT International Electric Corporation carries systems development, design and management responsibilities. In order to further multiply the effectiveness of the military commander, faced with the crucial task of assimilating vast quantities of information projected on the

screens before him, ITT engineers and scientists recently added a remarkable new capability to 465-L: data presentation in color.

Operating at speeds that appeared incredible only a short time ago, the system enables computer outputs to be converted to alpha-numeric form...photographed...developed and projected on control center screens in as many as 7 colors in a matter of seconds.

This new capability opens up a whole new field of data format techniques to be explored. An obvious and immediate value is the enhancement of human perception through color changes denoting differing degrees of situation criticality.

#### OPPORTUNITIES IN MANY COMMAND AND CONTROL AREAS NOW OPEN TO SYSTEMS ENGINEERS AND SENIOR PROGRAMMERS

Many of these positions are on 465-L. Other opportunities relate to large-scale commercial digital communication systems, oceanic systems, and satellite control. Your inquiry about any of the positions listed below will receive immediate attention.

**PROGRAMMERS/ANALYSTS.** For real-time programming analysis and development. Broad activities encompass advanced programming systems, including special color display routines; diagnostic programs; automatic recovery; problem-oriented language; artificial intelligence.

**OPERATIONS ANALYSTS.** To establish systems requirements in satellite control, air traffic control, ASW and command/control. Also, assignments in man/machine communications and information retrieval.

**SYSTEMS IMPLEMENTATION ENGINEERS.** Electronic engineers to develop tests for stressing and evaluating communication-display-computer systems. Recommend improvement and refinements. Also, field positions for installation and integration of digital command/control systems.

**INFORMATION SYSTEMS ENGINEERS.** For design of command/control and advanced communications systems. Experience in traffic, antenna and propagation theory, and mathematics as applied to communications and space technology.

**DIGITAL SYSTEMS ENGINEERS.** Engineers with management ability to direct sub-systems engineering effort on a global command/control system. Experience is desired in message traffic control, data processing systems, data display and multi-sequencing techniques.

Write fully in strict confidence to Mr. E. A. Smith, Manager of Employment, Box 95-M, ITT-International Electric Corporation, Route 17 and Garden State Parkway, Paramus, New Jersey.

# ITT

INTERNATIONAL ELECTRIC CORPORATION

Figure 9.0 – 4 November 1962 – Artificial intelligence appears under 'Programmers/Analysts.'<sup>724</sup>

<sup>724</sup> 'Display Ad 798 -- No Title'.

The term Good Old Fashioned Artificial Intelligence severs AI from these aspects of its Cold War history. GOFAI casts symbolic AI as antiquated without paying credence to the structural reasons *why* such theory initially became dominant in the 1960s and why it eventually sputtered into decline. GOFAI treats AI's origins as apolitical or even ahistorical by equating a methodology with an era, the latter of which it positions as contentless. Symbolic AI and GOFAI are often used synonymously. The term's stubborn popularity speaks to practitioners' alienation from, and perhaps ignorance of, their own history. Clarifying these entanglements helps to identify the differences between clandestine worldviews and clandestine infrastructures.

As I have argued, there were many parties interested in sophisticated new tools to manage information, and not all of them for the same reasons that attracted a mathematician. The term 'artificial intelligence' cast information management through the value-laden lens of human intelligence, which involved Minsky, his interlocutors and the field in deeply political questions of cultural priority that existing histories have failed to acknowledge.

## Conclusion

In between 1950 and 1961, Marvin Minsky developed pathbreaking work on neural networks and heuristic programming. In two 1961 papers, he helped to forge a research agenda for AI that, through four editions of *Computers and Thought*, has influenced the field considerably. One output of this period—symbolic AI—has been memorialized as Good Old Fashioned Artificial Intelligence. I have argued that this label fails to account for the Cold War and industrial contexts that Minsky and his colleagues worked within. The mathematicians involved in AI's development were embedded in dialogue about how to order information across various domains. These diverse application areas, which included scientific publication management and industrial administration, allowed them to test the assumption that universal solutions could be found to such dynamics—that 'systems' reduced to collections of 'information processing' techniques.

AI quickly found a receptive audience. At RAND, alongside Simon and Newell, Minsky trained a new generation of scholars in the idea that psychological mechanisms could reduce

to code; that epistemology could be programmed. In other teaching material, Minsky argued, 'The language of programming is better than that of learning theory to describe learning theory.'<sup>725</sup> These strong formulations, and the symbolic program generally, bolstered military observers, who poured over bibliographic material and recommended to superiors that the link between language and epistemology was critical to the future of information processing technology.<sup>726</sup> Derivative service vendors like the International Electrical Corporation soon followed suit, crafting means to recruit AI researchers to join the battle in the early 1960s. When military grants purported to be worth millions arrived at the MIT AI Laboratory unencumbered, the search for these principles deepened even further. In my view, the presumed dichotomy between cybernetics and 'AI' has failed to capture these consistencies, in part because 'AI' is only part of the story.

---

<sup>725</sup> Minsky, 'Re: Course Material', Estimated 1957.

<sup>726</sup> Murray, 'Information Processing Relevant to Military Command', 138–39.

## Chapter Six – Conclusion

As I have shown here, the 1950s offer a novel vantage into a set of brain model theories developed in the United States during the Cold War with the emergence of the unpolished rhetoric, techniques, and hopes of a handful of relatively unknown scientists eager to test the viability of mechanized intelligence and problem solving. In the hands of most commentators, the 1956 Dartmouth Summer Research Project on Artificial Intelligence in which they participated, along with others, has usually been treated as a hyperbolic genesis moment for symbolic AI, initiating the era of Good Old Fashioned Artificial Intelligence, and a paradigmatic rivalry between symbolic and sub-symbolic methods. This chronology of rupture and rivalry typically scaffolds conceptions of brain model development in subsequent decades, anchoring the relevance of research fashions, figures, and events in what amounts to a historical vacuum.

My dissertation challenges the explanatory depth of this discontinuity narrative in relation to brain model development in the 1950s. In the preceding chapters, I have questioned the paradigmatic status of the rivalry between symbolic and sub-symbolic methods by surfacing commonalities, both theoretical and structural, between Herbert Simon's development of complex information processing and heuristic programming, John McCarthy and Marvin Minsky's development of artificial intelligence, and Frank Rosenblatt's development of machine learning. I argue that the varied procedures these men used to simulate aspects of cognition should be approached as a diversified yet singular intellectual project to reduce epistemology to code.

In addition, I have worked to recapture instances in which each man borrowed from strands of *existing* social science to advance their research. These continuities anchor brain modelling in the context of postwar American social science, be it public administration and operations research (Simon), instrumentalist statistics and Hayekian economics (Rosenblatt), notions of 'common sense' and their relation to programming languages (McCarthy) or cybernetics and archive management (Minsky). The picture that emerges is not one of rupture but of remaking—a recalibrated 'origin' of AI that re-contextualizes research fashions in relation to localized contingencies. That contemporary machine learning researchers *still* turn

to subjective social and non-neural metaphors to model neural phenomena like learning suggests paths for further exploration.<sup>727</sup>

### An Emerging Discipline

The argument outlined above unfolds over the course of my thesis as follows. After a brief historiographical review, Chapter Two demonstrated why a touchstone of the discontinuity narrative—that the Logic Theory Machine was ‘The First AI Program’—is an instance of history told in reverse.<sup>728</sup> I show that the Machine was purposefully *not* a contribution to artificial intelligence when developed in 1955-1956. It was, instead, the first operationalized instance of what Simon and Newell called *heuristic programming*, or ‘The Next Advance in Operations Research.’<sup>729</sup> The duo had no interest in intelligence or even biology. Their goal was to model problem solving behaviour on the JOHNNIAC, such as how military operators on an Air Force base made decisions in a staged crisis.

To model human problem solving techniques, this group borrowed from two lines of theory derived from American and British philosophy and management science in the 1930s and 1940s. The first was Simon’s positivist account of administrative theory, distilled in his 1947 book *Administrative Behaviour*. Simon’s training at the University of Chicago under Rudolf Carnap encouraged him to develop a ‘scientific’ account of administrative organization. That book, an extension of Simon’s PhD research, cohered the heuristics of limited rationality for use in an applied setting. Simon theorized that to rationalize the behavior of an interactive decision system—in this case, an organization—one must first assume and impose a decision hierarchy upon that system. Agents in the upper portions of that hierarchy set the ‘decision premises’ governing the boundaries of possibility for agents below. Once conformed to this top-down enclosure, techniques like means-ends analysis and

---

<sup>727</sup> They leverage notions of zero-sum competition into General Adversarial Networks (GANs), a popular class of machine learning system premised, as the name suggests, on techniques used to model adversarial human relationships. Castelle juxtaposes GANs with the theories of the sociologist Pierre Bourdieu in: Castelle, ‘The Social Lives of Generative Adversarial Networks’. Before his passing in 2019, Patrick Winston, a student of Minsky’s and his successor as director of MIT’s AI Laboratory, told me that he, too, was surprised that social logics were not recognised as a foundational source of metaphor in the field.

<sup>728</sup> For a survey of instances in which this term is used, see: Boden, *Mind as Machine*, 705.

<sup>729</sup> Simon and Newell, ‘Heuristic Problem Solving’.



the logical notation and axioms of Russell and Whitehead's *Principia Mathematica* could be implemented to steer the organization's behaviour in certain ways. By my account, it was a small step for Simon to later bring Russell and Whitehead's system to structure the decision premises of a new Logic Theory Machine, intended to simulate human problem solving.

These formal and administrative logics gelled with the culture of conformity in place at the RAND Corporation in the mid-1950s. As prior employees reported, RAND's intellectual objectives, mandatory security clearances, and embrace of Primary Mental Abilities testing imposed a culture of conformity upon their practices, as did, I argue, their Cartesian architectural layout and operations. While Simon and Newell experienced their success with the Machine as serendipitous, I contextualize it as in line with investments made by RAND to develop control theory and technologies able to characterise the uncertain risks of cold warfare.

Symbolic AI, the research tradition that this portion of Simon's research helped to formalize, per [Chapter Five](#), is often juxtaposed with another paradigm, sub-symbolic AI, as epitomized by machine learning. In his first book in 1962, Rosenblatt certified this budding rivalry by describing AI researchers like Simon, McCarthy, and Minsky as his 'loyal opposition.'<sup>730</sup> He contrasted their 'paper exercises' ordering deterministic logical notations with perceptron theory, his own statistical approach to the mechanization of perception, which he developed for Cornell Aeronautical Laboratory's Systems Research Department under Project PARA.<sup>731</sup> As demonstrated in [Chapter Three](#), Rosenblatt's positivist inclinations undermined the distance he placed between symbolic AI and perceptron theory. In his PhD, Rosenblatt had gestured toward an equivalence between the use of statistics in psychology and the use of statistics in the discovery of gas laws in physics. This strongly empiricist view of human behaviour informed his later experimentation in psychology. He conceded no fundamental limit, or upper bound, on the explanatory potential of statistical techniques transformed by the computational capacity of digital technologies. Determined to realize this prospect, he developed the Electronic Profile Analyzing Computer (EPAC) between 1951 and 1953 to pilot and test a new statistical technique for multiple correlations, the 'k-coefficient.'

---

<sup>730</sup> Rosenblatt, 'Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms'; Rosenblatt, *Principles of Neurodynamics*, viii, 10.

<sup>731</sup> Rosenblatt, 'The Perceptron: A Theory of Statistical Separability in Cognitive Systems (Project PARA)', 2.

He hoped it would outperform classical methods like Lazarsfeld's latent structure analysis, and Thurstone's multi factor analysis.<sup>732</sup> Failure in this respect did not discourage him from seeking a mathematical, rather than practical, statistics in psychology.

Perceptron theory, which became his life's work, revisited this agenda. With the help of computing, Rosenblatt sought to transform statistics to capture the pattern—or rather, the code—that he presumed was latent in otherwise irreducibly complex neural stimuli patterns. Rosenblatt believed that there *existed* a code behind such activities, and that digital computers were the appropriate tool with which to isolate it. To justify this hypothesis, he referenced Hayek's idealization of decentralized market behaviour. The marketplace epitomized, for both men, the virtues of the decentralized distribution of processes across complex nodes. This framework, in their view, structured associations more competently than a centralized alternative. Both believed that the brain, by analogy, classified complex stimuli through a similarly decentralized mode of relation and distribution—in Rosenblatt's case, strengthening and weakening synapses in a manner akin to Hebb's cell assemblies. It did not rely on centralized regulation, Rosenblatt argued, let alone deterministic centralized regulation. As Simon had done with *Administrative Behaviour*, Rosenblatt reached to the dynamics of *social* phenomena to bring the complexities of mental life down to a scale he could manipulate.

Over the course of the 1950s, Minsky shifted his focus away from neurophysiology as an empirical basis for brain modelling and toward the abstract social logics and psychology of problem-solving other, per Chapter Five. This pivot, I argue, demonstrates his sustained commitment to articulating cognition through mathematics. During his PhD and development of SNARC, Minsky detailed the limits to mathematical metaphors for neural activity, honouring the complexities of brain-computer fidelity as he advanced his own aspirational vision of scientific medicine. He championed neural network research, however speculative, as essential to undertake because it would provide medical professionals with a mathematical means to bridge an otherwise open gap between biological data and observable mental disorders such as depression.<sup>733</sup> This view changed between 1955 and 1956. Minsky deemed the insights gleaned from artificial neural networks too slow and painstaking to realize in

---

<sup>732</sup> Rosenblatt, 'The K-Coefficient', 4; Thurstone, *Multiple-Factor Analysis*; Lazarsfeld, 'The Logical and Mathematical Foundations of Latent Structure Analysis'.

<sup>733</sup> Minsky, 'Theory of Neural-Analog Reinforcement Systems', 2–10.

comparison to the style of mimicry of problem solving that lay within reach of a competent mathematician. Rather than derive logics from an empirical basis, Minsky joined McCarthy in pursuit of what he called ‘universal’ decision procedures.<sup>734</sup> ‘Steps Toward Artificial Intelligence’ cemented this pivot. In that landmark paper, he, too, conflated problem solving and intelligence, seeking ‘a better theory of the task,’ rather than a better theory of the brain.<sup>735</sup> In formalizing that theory, Minsky moved between archival and cognitive understands of search and storage, and between the expert groups who examined each.

As a central component of the research agenda outlined in *Computers and Thought*, the popular first textbook for AI, published in 1963, ‘Steps Toward Artificial Intelligence’ firmed the intellectual concerns and orientation of a subsequent generation of AI researchers. It is telling that the editors of *Computers and Thought* omitted neural network research completely on the grounds that ‘Intelligent performance by a machine is an end difficult enough to achieve without “starting from scratch.”’<sup>736</sup> That, to them, starting from ‘scratch’ meant employing mathematics and not empirical insights drawn from biology, chemistry or physics showed the AI community’s commitments to abstraction.

The omission of neural networks from *Computers and Thought* is a candidate milestone in any rivalry narrative of brain model development in the 1960s. In regard to the 1950s, however, I speculate that infighting over the nuances of deterministic versus statistical approaches to the simulation of mental life ultimately helped researchers to overlook and normalize a more fundamental claim: that mathematics could be used to model brain activity *at all*. Debate over methodology drew focus away from that broader claim and the anxieties entailed in it and toward a derivative practical concern: if the mind could be modelled, how would one proceed? For the four researchers I study, the answer to this concern was conspicuously consistent: through programming. That *this* assumption drew no controversy or infighting between them speaks to the underappreciated line of solidarity I have emphasized by re-clustering these figures as co-contributors to a single intellectual project characterised by the radically empiricist aspiration to reduce epistemology to code.

To my knowledge, only McCarthy dealt explicitly in the 1950s with this horizon as his purposeful aim. ‘Once one system of epistemology is programmed and works no other will

---

<sup>734</sup> Contract: Nonr-494(03). Miles, ‘Operations Research: An ASTIA Report Bibliography’, 18.

<sup>735</sup> Boden, *Mind as Machine*, 894.

<sup>736</sup> Feigenbaum and Feldman, *Computers and Thought*, v–vi.

be taken seriously unless it also leads to intelligent programs,' he wrote, 'The artificial intelligence problem will settle the main problems of epistemology in a scientific way.'<sup>737</sup> While Simon, Rosenblatt, and Minsky's accounts differed in degree to McCarthy's, they did not, I argue, differ in kind. In the 1950s, each man elevated programmatic methods to a level tantamount to thought itself. In 1955, Simon implemented the Logic Theory Machine on his children, partner, and graduate students as a proof of concept for its later implementation in the JOHNNIAC. From this experiment, he claimed that he had solved the mind/body problem by inventing a system for non-numerical computation.<sup>738</sup> In short, Simon elevated rule following to the status of thinking, claiming that cognition existed outside of its material instantiation. In fact, I argue, he had made a virtue out of his participants' conformity to his creative organizational framework, meaning that of a multi-human computer.

The novelty of digital computation as a medium, for Simon's framework, was the extent to which it *could* be programmed. Whereas McCarthy wanted a computer to explain the mysteries of epistemology step-by-step as it processed decisions comparable to human judgements, Simon sought to encode his own conception of epistemology in order to demonstrate its usefulness, particularly to the fields of operations and systems research. He and his collaborators sought to condition and steer behaviours to certain outcomes and were among the first (alongside More, Samuel, and Hagelbarger) to demonstrate that a digital computer's affordances were suitable to heuristic methods, in this case: the exercise of adapted axioms from *Principia Mathematica*.<sup>739</sup> By my account, this claim reified a complacent acceptance of Simon's radically empiricist view of behaviour by devaluing human 'thinking' to the status of mechanized conformity, or operationalized rule following. He treated evidence of success in a closed system (e.g., the Logic Theory Machine) as equivalent to success in an open one (e.g., the human mind), and thereby normalized a limited vision of the latter.

Rosenblatt and Minsky's theories reified restricted agency in similar ways. In his schematics of perceptron theory, Rosenblatt abstracted the role of the system's human operator out of existence. His flowcharts charted how a Perceptron would pass judgement

---

<sup>737</sup> McCarthy, 'Physical and Mental Events and Intelligent Machines', 5.

<sup>738</sup> Simon, *Models of My Life*, 190.

<sup>739</sup> Boden, *Mind as Machine*, 707.

on a new stimuli.<sup>740</sup> They made no reference, however, to the intellectual contribution afforded by a human operator, meaning the person tasked with choosing and inputting which stimuli to input into the machine. The system's field of judgement was thus both highly restricted (in that a human necessarily guided what it would judge) and glorified as autonomous. These nuances were not dealt with carefully in the 1950s. 'Electronic "Brain" Teaches Itself' wrote *The New York Times* in July 1958, following an ONR-backed press conference.<sup>741</sup> Like Babbage and Simon, Rosenblatt conceived of mechanized cognition using a system that excused the role of its chief administrator: in this case, himself.<sup>742</sup>

Minsky, in his 1956 paper 'Heuristic Aspects of the Artificial Intelligence Problem,' which formed the basis for 'Steps' in 1961 and an unfinished 1957 book, hedged that the concept of intelligence required no 'isolated' character so long as its operations improved iteratively toward efficiency.<sup>743</sup> This was a tenuous definition based in substitution and an appeal to self-evidence. Minsky anchored the sufficiency of his model's problem solving ability to the question of whether or not it performed in a manner that commanded, in his words, 'our respect.'<sup>744</sup> In pursuing what he called 'universal' decision procedures, Minsky presumed not only that such a category existed, and that it could be articulated using mathematics, but that its character was self-evident. 'We could all agree [on what behaviours] embody, or reflect, intelligence,' he wrote.<sup>745</sup> It was by this path that Minsky's attempt to account for the impossibly fluid cultural definition of intelligence reduced, instead, to a convoluted deference to method. In this case, that method remained an ill-defined appeal to the virtue of 'efficiency,' the historical character of which Jennifer Karns Alexander has chronicled as similarly fluid.<sup>746</sup>

One reason to take stock of the rhetorical and philosophical commitments these men made in their search for disciplinarity in the 1950s is that digital computer engineering and

---

<sup>740</sup> Per Henderson and Ensmenger, flow charts of this type also served as 'conscription devices,' even if mutual understanding went unmet. As cited in: Ensmenger, 'The Multiple Meanings of a Flowchart', 324, 335.

<sup>741</sup> 'Electronic "Brain" Teaches Itself'.

<sup>742</sup> Schaffer, 'Babbage's Intelligence', 210.

<sup>743</sup> Minsky, 'Heuristic Aspects of the Artificial Intelligence Problem'.

<sup>744</sup> Remaking, again, an image of cognition rooted in the agency of its chief administrator. Minsky, iii.

<sup>745</sup> Minsky, i.

<sup>746</sup> Alexander, *The Mantra of Efficiency*, 8.

programming each matured into professions in the same period. Matti Tedre describes electronic computing in the late 1940s and early 1950s as a Kuhnian pre-science.<sup>747</sup> Within the 1950s, use of the term ‘electronic computers’ replaced ‘automatic calculating machines’ and ‘software’ replaced ‘automatic programming’ and ‘automatic coding.’<sup>748</sup> In the same decade, ‘automata’ and ‘brain models’ became ‘the artificial intelligence problem’ and then just ‘artificial intelligence.’ In ways that remain to be treated historically, a host of other terms were also assimilated into this tradition.<sup>749</sup> Like the Victorian aristocrats who earlier offered competing accounts of the dynamics of mental life via appeal to the behaviour of clocks, pianos, and telegraphs, the elite American researchers behind these terms indulged a ‘metaphorical menagerie’ distinct to their own locale; the increasingly computerized American empire.<sup>750</sup>

Simon, Rosenblatt, McCarthy and Minsky contributed to this shift in customs of civility and sociability, particularly in relation to the digitalization of cognition in American scholarly life. In the 1930s, Rashevsky had been lampooned by colleagues at the University of Chicago for presuming that he could conduct meaningful research in psychology with only a pencil. In the early 1950s, professional etiquette compelled Minsky to use scare quotes around mathematical metaphors to notions like ‘learning.’ Senior researchers who otherwise welcomed such comparisons, like Shannon and von Neumann, judged claims of strict fidelity to be premature and even bombastic. Both Morison and President Kennedy’s Science Advisory Committee remained tentative about the viability of mathematical languages of thought. Only a decade later, however, the MIT Artificial Intelligence Laboratory—scare quote free—was fully operational with a budget purported to be in the millions.

The astounding pace of this shift from provisional to professional theory can be understood through three large scale and entangled contributing factors. They are: a broad alignment between brain model theory and the diverse aims of wealthy patronage networks in the U.S., the liminal status of brain model metaphors in the 1950s, and the maturation of

---

<sup>747</sup> Tedre, ‘The Development of Computer Science’, 237.

<sup>748</sup> Campbell-Kelly and Garcia-Swartz, ‘Pragmatism, Not Ideology’, 234.

<sup>749</sup> This list includes ‘engineering psychology,’ ‘applied epistemology,’ ‘neural cybernetics,’ ‘non-numerical computing,’ ‘neuraldynamics,’ ‘advanced automatic programming,’ ‘automatic coding,’ ‘fully automatic programming,’ ‘hypothetical automata,’ and ‘machine intelligence’

<sup>750</sup> Schaffer, ‘OK Computer’.

programming in that same period. These structural advantages should not be overlooked since, as George A. Reisch writes, ‘Cold War intellectual life was no meritocracy. Winners and losers over the long term were not always determined according to intellectual talent.’<sup>751</sup> While I do not mean to suggest that brain model theory was undeserving of its rapid success, I do aim to contextualize why it fared better than other related areas, such as cybernetics.

I have returned to the first factor—patronage—at various points in this dissertation. I argue that mid-century brain model development aligned with government, industrial and military priorities in complex ways that, despite their nuances, availed generous funding opportunities to the figures I examine. Still, this patronage was neither uniform nor wholly coordinated. As such, it is perhaps best characterized, in aggregate, as a kludge. Daniel Volmar makes this case in relation to the U.S. military’s embrace of digital computing more generally. He challenges that ‘America does not have *a* military organization at all, but rather, *many*.’<sup>752</sup> Stated differently, Volmar reasons that lines of command in U.S. military structures were ad hoc rather than uniform—that accounts like Edwards’ *Closed World* concede too readily that the U.S. Air Force exercised ‘willful’ or ‘rational’ bureaucratic decision making. He argues that a long series of independently overseen ventures and ‘boondoggles’ only gained coherence as ideology *after* the fact, since there existed no single uniform set of parameters for success during their development.<sup>753</sup>

I interpret the sum of 1950s brain modelling patronage along similar lines. To show this, I must revisit key intersections in brief detail. First, consider Rosenblatt and Minsky’s early-career aspirations for statistics. These aspirations aligned, as discussed in Chapter Three, with post-war confidence in the United States over the transformative potential of scientific medicine to improve civic life.<sup>754</sup> Rosenblatt’s EPAC device was funded, in part, by the U.S. Public Health Service. He credited that body with instilling in him the conviction that, ‘the problems of measurement and data analysis would prove fundamental to scientific progress in psychopathology.’<sup>755</sup>

---

<sup>751</sup> Reisch, *How the Cold War Transformed Philosophy of Science*, xv.

<sup>752</sup> Emphasis his. Volmar, ‘The Computer in the Garbage Can’, 391.

<sup>753</sup> Volmar, 14.

<sup>754</sup> Mandel, ‘Beacon of Hope’, 1–8.

<sup>755</sup> Rosenblatt, ‘The Design of an Intelligent Automaton’, 7.

Rosenblatt's next major patron was the U.S. Air Force, which advanced its own ambitions for 'statistical control' technologies in the 1950s, as well as for systems designed to solve extremely large matrix problems.<sup>756</sup> After this PhD, Rosenblatt was hired by the Cornell Aeronautical Laboratory, a hotbed for military aircraft testing. His co-authored paper 'Parallax and Perspective During Aircraft Landings,' in the *American Journal of Psychology*, marked a pivot in his research trajectory from the dynamics of childhood experience and adult personality profiles to the dynamics of perception, the topic that would come to define his career.<sup>757</sup> This support was sustained. Near the end of his life, Rosenblatt described the U.S. military as 'aiding and abetting' his 'oddball status' at Cornell.<sup>758</sup> Administrative records corroborate this support. 'It was largely through the pressure of the ONR that the university acknowledged that such a program did have a place at Cornell, that a course should be taught in the general theory of brain mechanisms,' Rosenblatt recalled.

At RAND, Simon also benefited in a circuitous fashion from monies directed by the U.S. Air Force. In this case, military support of Simon's research through RAND allowed him liberal, complimentary access to the largest computing facility in the world, as well as to J. Clifford Shaw, whose competency in computer programming was not just useful, but, in my opinion, decisive to the Logic Theory Machine—a contingency that remains underappreciated in existing literature.<sup>759</sup> Carnegie Tech did not install its first IBM computer until 1956. Even then, it is unlikely that Simon and his collaborators would have enjoyed the same large-scale technological resources afforded to them by RAND.

For McCarthy, it was industrial rather than military support that enabled his research. In the early 1950s, executives at IBM, coaxed by employee Nathaniel Rochester, who would later co-convene the Dartmouth event, came to recognize the commercial case for large-scale scientific computing and automatic coding. To meet increasing demand for stored program machines after the outbreak of Korean War, the company invested in the MIT Computation Center, which Philip M. Morse, the 'founding father of operations research,' opened,

---

<sup>756</sup> Rees, 'The Federal Computing Machine Program', 736.

<sup>757</sup> Contract No. AF33(038)-22373. See: Gibson, Olum, and Rosenblatt, 'Parallax and Perspective during Aircraft Landings'.

<sup>758</sup> Rosenblatt, Interview with Frank Rosenblatt, 114.

<sup>759</sup> Ware, *RAND and the Information Evolution*, 62.



incidentally, midway through the Dartmouth workshop.<sup>760</sup> This resource proved pivotal to McCarthy, who had not yet had access to a functional computer through his academic employment at Dartmouth or Stanford. McCarthy did not program a computer until 1955, the same year he coined the term artificial intelligence.<sup>761</sup>

These strands of alignment between 1950s brain modelling and American industrial aims can also be seen in the composition of the Dartmouth meeting and subsequent events like it. Of the twenty participants named by Solomonoff as present in Hanover, twelve were affiliated with MIT, seven with IBM and three with Bell Laboratories, with some overlap—affiliations that broadly subsidized their participation.<sup>762</sup> A dozen analogous events in the late 1950s and early 1960s, such as the RAND Corporation's 1958 'Simulation of Cognitive Processes' summer school, congregated a similar set of industrial and academic leaders to devise new theories of search and storage, often in ways that did not strictly delineate between neural and social or logistical influences, as discussed in [Chapter Five](#). The 1958 Mechanisation of Thought Processes symposium at NPL, for example, convened a delegation comprised of two-thirds government and industrial representatives.

A second way to understand the pace of this transition into disciplinary status is through the liminal status of brain model metaphors in the mid-century. In 1961, Kennedy's Science Advisory Committee denied funding for an Institute for Non-Numerical Studies, meaning computations of 'symbols, meanings, and decisions,' on the grounds that such theory was 'shallow' and 'wishful.'<sup>763</sup> McCarthy, Minsky and Newell rejected this characterization as 'negative' and unfairly 'fantastic'.<sup>764</sup> Ultimately, the committee endorsed 'man-machine cooperation' while simultaneously rejecting non numerical studies.

That the differences underlying these areas were broadly amorphous at the turn of the decade was captured in 1958 by the new journal *Communications of the Association for Computing Machinery*. The ACM questioned its readers about the professional title they used

---

<sup>760</sup> Aker, *Calculating a Natural World*, 286–88; Little, 'Philip M. Morse and the Beginnings', 146; McCarthy, 'History of LISP', 217.

<sup>761</sup> Nilsson, 'John McCarthy, 1927-2011', 5.

<sup>762</sup> Of the uncorroborated attendees, two were affiliated with IBM, two with Bell Labs, five with academia and one with the UK government.

<sup>763</sup> Slayton, *Arguments That Count*, 74–77.

<sup>764</sup> Newell, 'Comments on Ad Hoc PSAC Panel on Non-Numerical Information Processing', 15 December 1961; Slayton, *Arguments That Count*, 77.

to describe themselves or their field. Answers ranged from synnoetics, computer science, and comptology to Flow-Chartsman, Applied Meta-Mathematician and Applied Epistemologist.<sup>765</sup> Shannon's refusal to include 'intelligence' in the title of *Automata Studies* but simultaneous willingness to co-convene the Dartmouth workshop under the title of artificial intelligence provides another glance into the indeterminacies active at the heart of the new research area. These indeterminacies availed breadth for manoeuvring toward plausibility in the face of complaints, such as from Morison, that mathematical theories of brain functions remained, 'Pre-Newtonian.'<sup>766</sup>

A third and final way to understand AI's rapid formalization into a discipline, which connects to the second, arises when judging the maturation of programming in the 1950s in parallel with the formalization of brain model theory. In the same decade that computational 'learning' became computational learning, 'pseudo-code' became code. I argue that the standardization and professionalization of programming aided in the rapid development of AI theory into a discipline in at least two important ways. First, it vastly improved, at little to no intellectual cost to brain model researchers, the expressive range and accessibility of programming techniques. Second, largely as a result of those gains, it perpetuated the viability of algorithmic cultures generally. I will assess each factor in turn.

At the 1956 Dartmouth Summer Research Project, the event's co-convenors lamented the expressive range of existing programming techniques. 'The major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have,' they wrote.<sup>767</sup> Prior this, between 1952 and 1954, Newell and Shaw had struggled to create a coherent 'language' with which to translate between the operations of a machine and complex human thought processes.<sup>768</sup> Elsewhere, as noted, IBM developers claimed to have invested eighteen man years of work into FORTRAN between 1954-57, the result of which was an estimated eighty to ninety percent reduction in the time spent coding and debugging. Propelling these consolidations was an exponential growth of digital computing in America. The number of operational digital computers in the country rose from 2 in 1950, to 243 in

---

<sup>765</sup> Correll, Khodr, and Vanderburgh, 'Letters to the Editor'; Weiss and Corley, 'Letters to the Editor'; Ensmenger, *The Computer Boys Take Over*.

<sup>766</sup> Morison, 'New York'.

<sup>767</sup> Minsky et al., 'Proposal', 1.

<sup>768</sup> As cited in: Dick, 'After Math', 78.

1955, to 5,400 in 1960.<sup>769</sup> The consolidation of automatic coding techniques using assemblers, compilers and, later, programming languages like FORTRAN, ALGOL and COBOL drastically reduced the expertise required to operate a device. ‘No computer manufacturer would succeed in selling his machines without an auto-programming system,’ wrote Booth in 1953.<sup>770</sup>

Standardization reduced the significant costs of research and development in brain modelling projects, substantiating, in the process, an entire class of languages equal in pedigree to the Information Processing Language that Simon, Newell, and Shaw had labored to create as the logical architecture on which the Logic Theory Machine ran. They positioned this ‘language’ as the most difficult part of their experimental procedure. It is perhaps unsurprising then that during the 1950s, Simon, Rosenblatt, McCarthy and Minsky each used their distinct platforms within major U.S. research centers to champion programming techniques as the way of the future. They advocated for the adoption of programming methods to influential audiences like tenured faculty (McCarthy at the MIT Computation Center), up and coming social scientists (Simon, Minsky at the RAND Corporation), and interdisciplinary civilian and military researchers (Rosenblatt at Cornell).

A second reason that the maturation of programming in the 1950s advanced the formalization of brain model theory was visibility. Digital technologies in public life lent an air of viability to the brain model project, which benefitted from the perceived authority of that medium. This is why an operational program was treated as the solution to what Solomonoff called ‘the demo to sponsor problem’ at Dartmouth.<sup>771</sup> Whether or not the results were significant, digital tools and programs commanded authority via what Edwards describes as ‘their aura of almost erotic scientificity.’<sup>772</sup> This is why Simon and Newell’s print out of their work on heuristic programs generated more excitement in Hanover than similar results (without documentation) by More and Samuel. This also explains why, by the early 1960s, per [Chapter Five](#), operationalized programs were treated as a hedge for AI practitioners

---

<sup>769</sup> Ensmenger, *The Computer Boys Take Over*, 28.

<sup>770</sup> Booth and Booth, *Automatic Digital Calculators*, 1965, 215.

<sup>771</sup> Solomonoff, ‘Untitled Notes Re: Wendy Conquest’; Solomonoff, ‘Dartmouth’, 19.

<sup>772</sup> Edwards, *The Closed World*, 121.

against critiques of developing pseudoscience.<sup>773</sup> In each case, brain model research gained credibility via appeal to the viability of algorithmic cultures more generally.

This national techno-cultural trend influenced Simon, Rosenblatt, McCarthy and Minsky as well. In notes, journal entries, and published papers from the 1940s and 1950s, each speculated on how computing technology could be leveraged to reconfigure all matter of daily life, including the country's social orderings. McCarthy saw computational logic as a liberating force from hierarchical rule in national governance. 'There should be a semi-mathematical theory for decision procedure for formulating policy,' he contended.<sup>774</sup> Simon, a graduate of the Chicago School of political science, had been trained to believe that 'Societies needed to be led and economies needed to be regulated for there to be progress.'<sup>775</sup> For him, the Logic Theory Machine amounted to the opening salvo of a grand reconfiguration of daily routine around machine-driven problem solving, including 'discovering new scientific laws, understanding English prose... making investment decisions... the full array of intelligent processes.'<sup>776</sup> In 1958, he claimed heuristic programming as a victory for the 'soft' sciences, as 'something we can point to when the superior accomplishments of the natural sciences become too embarrassing for us.'<sup>777</sup> Rosenblatt, similarly, compared brain modelling to physics prior to Newton, suggesting that an integrated set of principles awaited discovery.<sup>778</sup> In his obituary, he was credited for applying computer programming to political statistics in campaigning.<sup>779</sup> While his methods differed from Simon's, his faith in the social role of method did not; both saw the systematic application of computation to social routines as legitimate and desirable. McCarthy, in a move that Americans nationwide would replicate a half century later (although not strictly by his influence) elevated programs to the status of thinking object, positioning the program as 'the agent in artificial intelligence' in order to normalize framings like, 'The program does this' and not 'The machine does this.'<sup>780</sup>

---

<sup>773</sup> Also explored in: Fleck, 'Development and Establishment in Artificial Intelligence'.

<sup>774</sup> McCarthy, 'Levels of Theory in Engineering and Politics'.

<sup>775</sup> Crowther-Heyck, *Herbert A. Simon*, 44.

<sup>776</sup> Simon, Newell, and Shaw, 'Untitled Memo'.

<sup>777</sup> Simon and Newell, 'Heuristic Problem Solving', 2.

<sup>778</sup> Rosenblatt, 'The Design of an Intelligent Automaton', 1.

<sup>779</sup> 'Frank Rosenblatt Dies; Researcher, Politician'.

<sup>780</sup> McCarthy, 'An Approach to the Artificial Intelligence Problem', 3.

Despite ongoing debate about the merits of deterministic versus statistical methods in brain model development, each man, eventually, came to model his early experiments on the same device, an IBM 704. The uniformity of their approaches calls attention to the role that digital computing as a medium had on scholarly customs in the United States in the mid-century. While they debated the relative merits of their competing techniques, they did not acknowledge or contest this deeper line of continuity. That is, I argue, because it was a core assumption for each man that computing programming was an adequate tool for their communal goal of articulating aspects of the dynamics of thought.

AI's history has so far been distanced from the material history of programming, but that is changing. As Dick has shown, the idiosyncratic materiality of individual devices led mathematicians to rethink the very meaning of intuition, not just the threshold of what constituted mathematical proof in automated decision systems.<sup>781</sup> I have charted how this medium connected figures who simultaneously treated each other as rivals. I have also surfaced in new detail the patronage networks that supported various types of brain model theory, all thought not in an overarching, coordinated manner.

#### Fitting AI into Histories of American Social Science Research in the Mid-Twentieth Century

In the early 1960s, the U.S. Air Force positioned accounts of epistemology-as-information by Minsky, Rosenblatt, and their colleagues as a persuasive strategy for knowledge management infrastructure within the U.S. military. They acknowledged, as Rosenblatt had, that theories on automated reason and automated perception could be combined in a modular fashion. That these competing research fashions were taken as cumulative, rather than mutually exclusive, shows the limits of historiographical frames that trace their development along the lines of GOFAI versus neural networks.

In the above, I have outlined various reasons to consider the formalization of artificial intelligence into a discipline as continuous with and embedded within a broader intellectual project to render epistemology in code. I relate this project to developments in programming to clarify the nature of the influence that patronage had on conceptions of mechanical minds in the United States in the 1950s. As stated, much of the brain-model theory that I chronicle

---

<sup>781</sup> Dick, 'After Math'.

here was funded primarily by one or all of the U.S. joint forces. Support trickled down through subsidiary institutions like the RAND Corporation, the Cornell Aeronautical Laboratory and the Office of Naval Research. It was provisioned, initially, for a handful of loosely connected or unrelated reasons: to gain access to civilian talent during peacetime (in the cases of Simon, Rosenblatt, Minsky); to clarify and tactfully insure against the amorphous dynamics of ‘cold’ wartime risk (Simon); to optimize efficient team communication on Air Force bases during crisis (Simon); and to automate aspects of aerial surveillance (Rosenblatt).

Given the sensitivities involved in these alignments, I will conclude with a comment on how my research fits into existing historiography on Cold War social science. To do so, I draw on parallel surveys of three decades of Cold War historiography by David C. Engerman and Joel Isaac.<sup>782</sup> These surveys identify two distinct genres of scholarship on how the conditions of that conflict intersected with the development of the social sciences in the United States. That there are two genres, rather than one, speaks to a sustained rift at the center of contemporary scholarship over the foundations of American social science.

The first genre treats ‘Cold War’ as an adjective, exploring evidence that corroborates it as the active consolidating force in the transformative development of professional research norms in disciplines like economics, psychology, and sociology. Depending on the setting, according to Isaac, this pressure was either *constructive*, as in S. M. Amadae and Philip Mirowski’s accounts of how the U.S. military promoted and sponsored game theory research in the post-war period to advance warfare planning, or *repressive*, as in John McCumber’s study of how McCarthyism led philosophy departments to censor socially engaged existentialist and pragmatist inquiry in favour of arcane, overtly value neutral topics like semantics and mathematical logic.<sup>783</sup>

The second genre considers analogous pressures on intellectual trends in the social sciences but with a crucial difference—it does not go as far as to attribute those changes solely to the conditions of the Cold War.<sup>784</sup> Prefigurative and external factors such as

---

<sup>782</sup> Engerman, ‘Social Science in the Cold War’; Isaac, ‘The Human Sciences in Cold War America’.

<sup>783</sup> Amadae, *Rationalizing Capitalist Democracy*; Mirowski, *Machine Dreams*; McCumber, *The Philosophy Scare*; See: Isaac, ‘The Human Sciences in Cold War America’, 737; On the rise and fall of the Unity of Science movement see: Reisch, *How the Cold War Transformed Philosophy of Science*.

<sup>784</sup> Rohde, *Armed with Expertise*; Solovey, *Shaky Foundations*.

industrial priorities also explain and contextualize postwar trends. Bruce Seeley's critique of Rebecca S. Lowen's *Creating the Cold War University* provides an example. Seeley contends that the university-as-a-service model that Lowen argues was cemented in U.S. academic culture due to the Cold War was already in use in nineteenth century agricultural experimentation and early twentieth century engineering. 'The military-industrial complex did not always call the tune,' summarizes Isaac.<sup>785</sup> To clarify the explanatory boundaries at stake, Engerman distinguishes 'Cold War Social Science' from social science that occurred during the Cold War.<sup>786</sup>

These boundaries remain contentious.<sup>787</sup> Engerman and Isaac structure their surveys to act as appeals to colleagues to monitor hyperbole and preserve complexity, nuance, and, ultimately, precision in deepening research on the influence of American geopolitics on mid-century social science. Isaac calls for 'middle-range contextualizations,' especially for research areas that lack sufficient historical scholarship to justify larger claims on the road toward consensus.<sup>788</sup> My dissertation examines three such areas when assessing artificial intelligence, complex information processing, and, to a lesser extent, machine learning. I argue that these branches of the mathematical sciences should be reinterpreted as a unified whole grafted, in part, from—and best seen in continuity with—the social sciences of the postwar era.

This is a novel argument; and it relies on multiple lines of evidence to understand the particular form of generality involved. Existing literature tends not to include AI and machine learning as offshoots of Cold War human sciences. Both go broadly unacknowledged in recent compendiums, nor do Engerman nor Isaac consider the fields in their larger surveys.<sup>789</sup> Lacking sufficient context, one could mistakenly swing analysis in the other direction, indulging too readily in assumptions of enmeshment between the two domains. At first glance, postwar social science asserted similar aspirational characteristics to the research I

---

<sup>785</sup> Isaac, 'The Human Sciences in Cold War America', 740.

<sup>786</sup> Engerman, 'Social Science in the Cold War', 393.

<sup>787</sup> Footnote 10 in Solovey and Cravens, *Cold War Social Science*, 21 provides an overview of debate and related literature from recent conferences. The Cold War's end date is also debated.

<sup>788</sup> Isaac, 'The Human Sciences in Cold War America', 728.

<sup>789</sup> Solovey and Cravens, *Cold War Social Science*, 2012; Ross and Porter, *The Cambridge History of Science, Volume 7: The Modern Social Sciences*.

have profiled: neutral objectivity, universal applicability, over confidence in scientific maturity, and faith in systematized rationalization through professionalization.<sup>790</sup> Both areas received substantial funding from the U.S. establishment in the postwar period, often through subsidiaries of federal agencies like RAND on behalf of the U.S. Air Force.

One reason I adopted a biographical approach to my inquiry was to preserve these lines of nuance while attempting, simultaneously, to paint a broader picture of brain model development as a whole. This ‘middle-range’ framing is appropriate since each of the four men I assess fit into the U.S. defence apparatus, and its related historical genres, in a different way. Simon’s Logic Theory Machine and Rosenblatt’s Project PARA come closest to fulfilling the use of ‘Cold War’ as adjective. Absent that conflict, the idiosyncratic material resources, institutional leeway, and career stability they leveraged to pursue those projects would have been prohibitively difficult to access in combination. For these men, each operating out of a Systems Research laboratory, the Cold War served as a *constructive* rather than *repressive* pressure; the type that would justify the use of that term as an adjective. I caveat, however, that this pressure was not necessarily the product of any deliberate meddling by military or government administrators—I will return to that topic in a moment.

McCarthy and Minsky’s ventures into brain modeling between 1955-1959 did not rely on military support to such a degree that the adjective ‘Cold War’ adequately contextualizes the character of their research. Still, both men unequivocally gained from the avenues availed to them by the scale of post-war defense funding for civilian science. The four-thousand-dollar grant from the U.S. Air that Minsky received for SNARC, along with his participation fees for the Dartmouth workshop, and his invitations to engagements with RAND, MIT’s military-industrial affiliates and others, increased his career mobility. Yet each boost reached him in a smaller sum, dispersed at a staggered interval. These securities did not accrue in the steady and sizeable allowance afforded to Simon and Rosenblatt for analogous projects. McCarthy’s career, by and large, progressed without direct military funding in the 1950s.

The same cannot be said of Minsky and McCarthy’s AI research during the 1960s and afterwards. Although this question extends beyond my period of focus (1955-1961) and beyond the archives I have had occasion to consult for this project, the duo established the AI Laboratory at MIT with a fifty thousand dollar joint-services grant. Grants of this type, as

---

<sup>790</sup> Solovey and Cravens, ‘Cold War Social Science’, 2012, x.



well as the pragmatic inventions they were intended to fund, were formative to MIT's research culture in the mid-century. In the 1930s and 1940s, MIT President Karl Compton had emphasized practical research for industrial and federal patrons. During the war, MIT was the largest academic contractor of Roosevelt's National Defense Research Committee and Office of Scientific Research and Development, which set the mold for the military's funding of academic research during peacetime.<sup>791</sup> Lowen captures—using Cold War as an adjective—how entrepreneurial administrators like Compton and Stanford University Provost Frederick Terman, retooled the incentives in play at leading U.S. research centres to shift research from an 'ivory tower' model, wherein knowledge was pursued for its own sake, into a knowledge-as-a-'service' model, in which academics who pursued industrial and government research priorities benefited. This entrepreneurial model of the 'multiversity' became the norm in the United States, she argues.

My archival subject (scientists) differed from Lowen's (administrators). This complicates my giving any final say on whether or not the development of brain model research faced direct steering or interference by administrators at MIT, Cornell, RAND, or ONR. I have found no evidence to support such a possibility, at least not in the 1950s. McCumber cautions that administrators at private institutions like Pomona College navigated McCarthyism so discretely that faculty were 'totally unaware' of their interference.<sup>792</sup> I leave the question open, then, of whether the military had premeditated plans for the mathematics of the mind. I suspect not.

Corroboration of deeper intrigue would not greatly remake my thesis. As stated in Chapter Two, adjacent histories of Cold War diplomacy and social science by Wolfe, Cohen-Cole and Heyck, as well as pre-figurative histories of calculation and automation in the nineteenth century by Daston and Schaffer, have chronicled the Janus faced conservatism embedded in related modes of experimentation cast, duplicitously, as centrist or value-neutral. Wolfe revisits how the cultural ideal of science as free and unencumbered by government interference was, ironically, popularized via U.S. propaganda with the direct *and* inadvertent support of prominent U.S. scientists.<sup>793</sup> Cohen-Cole charts how the concept of 'open-mindedness' was used as shorthand for the virtuous mid-century democratic citizen,

---

<sup>791</sup> Lowen, *Creating the Cold War University*, 52.

<sup>792</sup> McCumber, *The Philosophy Scare*, 17.

<sup>793</sup> Wolfe, *Freedom's Laboratory*.

unaffected, ironically, by Communist conformity. Heyck demonstrates how a detached claim to universal applicability in the decision sciences and systems analysis served to expand and consolidate the cultural hegemony of the modern liberal state, rather than break from it.<sup>794</sup>

In accounts that prefigure such trends, Lorraine Daston charts how the separation of calculation from intelligence during the rapid industrialization of the early nineteenth century steered cultural value in the west away from rote calculation and toward forms of ‘managerial intelligence’ that rewarded those with the executive capacity to harness (e.g. manage or command) human computers, such as insurance agencies and government bureaucracies.<sup>795</sup> Gaspard de Prony famously wrote that a *lack* of intelligence was desirable amongst the class of workers minding his mathematical tables.<sup>796</sup> Schaffer adds that, after 1870, a new class of ‘intellectual aristocracy’ or public intellectual in Victorian England set out to curate the measures of mental life by formalizing insights from their own social economies. Babbage, the industrialist, had earlier seen cognition as akin to a factory.<sup>797</sup> Aldous Huxley, as metaphysician, mused on a measure of consciousness tantamount to the mechanical equivalent of heat. Metrics of this type helped to popularize notions like brain waves and brain power, along with the men who promoted them.

Each of these accounts pivots on the themes of authority and the application of conformity via misdirection. My thesis turns this critical lens toward the foundations of early decision technologies, toward *infrastructure* that belies a deeper set of philosophical entanglements than centrist accounts have entertained, particularly given its ability ‘to radically transcend the circumstances and locality of its production.’<sup>798</sup> Where my intervention differs from other Cold War historians like Wolfe, McCumber, Cohen-Cole and Lowen is in emphasizing how the theme of authority intersected with certain modes of materiality (e.g., digital computing tools). Solovey argues that mid-century decision sciences like game theory boosted the rational choice model of social science during the Cold War because they offered researchers a special type of ‘scientific’ results.<sup>799</sup> In chronicling the

---

<sup>794</sup> Heyck, ‘Producing Reason’.

<sup>795</sup> Daston, ‘Calculation and the Division of Labor, 1750-1950’, 12; Daston, ‘Enlightenment Calculations’.

<sup>796</sup> Daston, ‘Enlightenment Calculations’, 195. See also: Grier, ‘Human Computers’, 34–38.

<sup>797</sup> As cited in: Schaffer, ‘Babbage’s Intelligence’, 210.

<sup>798</sup> Radin, “‘Digital Natives’: How Medical and Indigenous Histories Matter for Big Data”, 45.

<sup>799</sup> Solovey, *Shaky Foundations*.

precious step from machine ‘learning’ to machine learning, I ask how computer programs arrived at their own special technological status, not only as a legitimate basis for scientific evidence, but as a proxy for mental behaviour and as a mode of authority distinct from its programmers; a disembodied ‘public intellectual’ of the twentieth and twenty first centuries.

In their analysis of ‘a change in what it *meant* to be rational in the age of nuclear brinkmanship,’ the authors of *How Reason Almost Lost Its Mind* use Cold War as adjective to chart the spread of a special form of rationality that premised method above content.<sup>800</sup> I argue that in mid-century brain model research, one finds an initial blurring in the boundaries between method and content. In AI, programming served as both method and content, complicating the hierarchy between natural and social systems of knowledge. McCarthy longed for a technologically *programmed* epistemology to settle the main problems of that field in ‘a scientific way.’<sup>801</sup> *Computers and Thought*, the first major textbook on AI, advanced this goal. Allen Newell reflected that, due in part to that text, operationalized computer programs became ‘the coin of the realm’ in AI research in the 1960s. The cultural dynamics informing this threshold are revealed in the early career decisions of Simon, Rosenblatt, McCarthy, Minsky and their colleagues, who sought collectively to formalize a special new type of ‘technological’ results, and a discipline capable of maturing them.<sup>802</sup>

Forman argues that the mid-1970s to mid-1980s witnessed a shift in which ‘technology’ assumed the cultural primacy that ‘science’ had held for millennia. This inversion is judged by him to have marked the shift from modernity to postmodernity. Whereas technology derived from science in modernity, in postmodernity the opposite was true. Forman argues that it was ‘beyond the capacity of a modernist intellectual,’ including cultural theorists like Marx, Veblen, and Dewey, ‘to unthink the primacy of science to and for technology’ in the nineteenth and twentieth centuries.<sup>803</sup>

Now, he continues, the reverse is true: technology is universally accepted as ‘the principal model for all those “ordering” activities that constitute culture.’<sup>804</sup> The transition into postmodernity has come at the expense of disciplinarity, ‘modernity’s final and

---

<sup>800</sup> Emphasis theirs. Erickson, *How Reason Almost Lost Its Mind*, 21.

<sup>801</sup> McCarthy, ‘Physical and Mental Events and Intelligent Machines’, 5.

<sup>802</sup> Feigenbaum, Feldman, and Armer, *Computers and Thought*, iii.

<sup>803</sup> Forman, ‘The Primacy of Science in Modernity, of Technology in Postmodernity, and of Ideology in the History of Technology’, 24.

<sup>804</sup> Forman, 53.

distinctive mode of knowledge production and curation.<sup>805</sup> Behind this shift, Forman contends, is a transition from the valuation of means to the valuation of ends. The development of mid-century brain model research provides a unique window into this proposed inversion. The figures I examine set out to manufacture scientific authority *using* technology. Even when they failed, their visions caught the attention of funders and the public alike. Further study is needed to contextualize the legacy of these *decision technologies*, including their relations to the social sciences that informed their design.

By shifting from histories of scientific personality and research fashions to histories of language propagation, commercial interests and pedagogy, these structural contingences come one step closer into view.<sup>806</sup> Narratives of rupture and rivalry, the era of Good Old Fashioned Artificial Intelligence, Dartmouth mythologies, and the presumed dichotomy between cybernetics and ‘AI,’ have so far failed to account for these continuities, in part because ‘AI’ is only part of the story. Examining brain model history through the lens of continuity rather than discontinuity allows for new modes of interpretive analysis into the fabled futures of the past.

---

<sup>805</sup> Forman, ‘On the Historical Forms of Knowledge Production and Curation’, 60.

<sup>806</sup> Accounts of how English became the de facto language in contemporary world science may reveal parallels for how programming became a baseline norm in various institutional contexts. See: Gordin, *Scientific Babel*.

## Bibliography

- Abraham, Tara H. “‘Microscopic Cybernetics’: Mathematical Logic, Automata Theory, and the Formalization of Biological Phenomena, 1936-1970’. PhD Dissertation, University of Toronto, 2000.
- . ‘Nicolas Rashevsky’s Mathematical Biophysics’. *Journal of the History of Biology* 37, no. 2 (2004): 333–85. <https://doi.org/10.1023/B:HIST.0000038267.09413.0d>.
- . ‘(Physio)Logical Circuits: The Intellectual Origins of the McCulloch-Pitts Neural Networks’. *Journal of the History of the Behavioral Sciences* 38, no. 1 (2002): 3–25. <https://doi.org/10.1002/jhbs.1094>.
- . *Rebel Genius: Warren S. McCulloch’s Transdisciplinary Life in Science*. Cambridge, Massachusetts: The MIT Press, 2016.
- . ‘Review of Mind as Machine: A History of Cognitive Science by Margaret A. Boden’. *The British Journal for the History of Science* 41, no. 4 (December 2008): 623–24. <https://doi.org/10.1017/S0007087408001775>.
- Acemoglu, Daron, and James A. Robinson. *Why Nations Fail: The Origins of Power, Prosperity and Poverty*. 1st ed. New York: Crown Publishers, 2012.
- Adam, Alison. *Artificial Knowing: Gender and the Thinking Machine*. London ; New York: Routledge, 1998.
- Agar, Jon. ‘How to Kill AI’. Presented at the Origin Myths of Artificial Intelligence: Histories of Technology and Power, Pavilion Room, Hughes Hall, Cambridge, 30 November 2018. <http://lcfi.ac.uk/events/origin-myths-artificial-intelligence-histories-tec/>.
- . *Science in the Twentieth Century and Beyond*. History of Science. Cambridge, UK ; Malden, MA: Polity Press, 2012.
- . *The Government Machine: A Revolutionary History of the Computer*. History of Computing. Cambridge, Mass: MIT Press, 2003.
- Agre, Philip E. ‘Hierarchy and History in Simon’s “Architecture of Complexity”’. *Journal of the Learning Sciences* 12, no. 3 (July 2003): 413–26. [https://doi.org/10.1207/S15327809JLS1203\\_4](https://doi.org/10.1207/S15327809JLS1203_4).
- . ‘Notes on Critical Thinking, Microsoft, and EBay, along with a Bunch of Recommendations and Some URL’s.’ Home Page, Philip E. Agre. Accessed 11 October 2019. <https://pages.gseis.ucla.edu/faculty/agre/notes/00-7-12.html>.
- . ‘Phil Agre’s Mind’. This is the website of Six Silberman., 12 July 2000. <http://wtf.tw/ref/agre.html>.
- Akera, Atsushi. *Calculating a Natural World: Scientists, Engineers, and Computers during the Rise of U.S. Cold War Research*. Inside Technology. Cambridge, Mass: MIT Press, 2007.
- Akman, Varol. ‘Formalizing Common Sense: Papers by John McCarthy’. *Artificial Intelligence* 77, no. 2 (September 1995): 359–69. [https://doi.org/10.1016/0004-3702\(95\)90018-7](https://doi.org/10.1016/0004-3702(95)90018-7).
- Alexander, Jennifer Karns. *The Mantra of Efficiency: From Waterwheel to Social Control*. Baltimore: The Johns Hopkins University Press, 2008.
- Alexander, S. N. ‘The National Bureau of Standards Eastern Automatic Computer’, 84–89. ACM Press, 1951. <https://doi.org/10.1145/1434770.1434784>.
- Amadae, S. M. ‘Computable Rationality, NUTS, and the Nuclear Leviathan’. In *The Decisionist Imagination Democracy, Sovereignty and Social Science in the Twentieth Century*, edited by Nicolas Guilhot and Daniel Bessner. Berghahn Books, 2018.

- . *Rationalizing Capitalist Democracy: The Cold War Origins of Rational Choice Liberalism*. Chicago: University of Chicago Press, 2003.
- ‘American Society for Cybernetics - Defining Cybernetics’, 15 January 2020. <http://www.asc-cybernetics.org/foundations/definitions.htm>.
- An Interview with Edward A. Feigenbaum. Interview by Mie-Sophia Augier, Morten Vendelø, and Edward A. Feigenbaum, 17 October 2002. [https://www.researchgate.net/publication/236904576\\_An\\_Interview\\_with\\_Edward\\_A\\_Feigenbaum](https://www.researchgate.net/publication/236904576_An_Interview_with_Edward_A_Feigenbaum).
- Andresen, S.L. ‘John McCarthy: Father of AI’. *IEEE Intelligent Systems* 17, no. 5 (September 2002): 84–85. <https://doi.org/10.1109/MIS.2002.1039837>.
- Archer, Pechawis, Jason Edward Lewis, Noelani Arista, and Suzanne Kite. ‘Making Kin with the Machines’. *Journal of Design and Science*, no. 3.5 (16 July 2018). <https://jods.mitpress.mit.edu/pub/lewis-arista-pechawis-kite>.
- Arden, D., S. Best, F. Corbato, F. Helwig, J. McCarthy, A. Siegel, F. Verzuh, and M. Watkins. ‘Coding for the MIT-IBM 704 Computer’. Massachusetts Institute of Technology, October 1957.
- Arendt, Hannah. *Eichmann in Jerusalem: A Report on the Banality of Evil*. Penguin Classics. New York, N.Y.: Penguin Books, 2006.
- Armer, Paul. ‘Problems in Administrative Leadership in the Computer Field’, 93–101. ACM Press, 1964. <https://doi.org/10.1145/1142635.1142643>.
- Ashby, W. Ross. *An Introduction to Cybernetics*. London: John Wiley & Sons, 1963.
- . ‘Design for a Brain’. *Electronic Engineering* 20 (December 1948): 379–83.
- Aspray, W. ‘The Stored Program Concept’. *IEEE Spectrum* 27, no. 9 (September 1990): 51. <https://doi.org/10.1109/6.58457>.
- Austrian, Geoffrey D. ‘The Machine That Carried IBM into the Electronics Business’. *Annals of the History of Computing* 3, no. 2 (April 1981): 186–88.
- ‘Automatic Coding: Proceedings of the Symposium Held January 24-25, 1957 at the Franklin Institute in Philadelphia’. Philadelphia, PA: Journal of the Franklin Institute, April 1957. Minsky 6AR. Personal Archives of Gloria Rudisch.
- B. Watson, John. ‘Psychology as a Behaviorist Views It’. *Psychological Review* 20 (1913): 158–77.
- Babb, Colin E. ‘The Genesis of the Office of Naval Research’. *Naval Science and Technology: Future Force* (blog). Accessed 1 March 2017. <http://futureforce.navylive.dodlive.mil/2016/08/the-genesis-of-the-office-of-naval-research/>.
- Backus, J. W., R. J. Beeber, S. Best, R. Goldberg, H. L. Herrick, R. A. Hughes, L. B. Mitchell, et al. *Fortran: Automatic Coding System of the IBM 704*. New York: IBM’s Applied Science Division and Programming Research Dept, 1956.
- Backus, J. W., H. Stern, I. Ziller, R. A. Hughes, R. Nutt, R. J. Beeber, S. Best, et al. ‘The FORTRAN Automatic Coding System’. In *Papers Presented at the February 26-28, 1957, Western Joint Computer Conference: Techniques for Reliability on - IRE-AIEE-ACM ’57 (Western)*, 188–98. Los Angeles, California: ACM Press, 1957. <https://doi.org/10.1145/1455567.1455599>.
- Bailey Jr., Herbert S. Letter to Marvin Minsky, 12 June 1951. Minsky 2AR. Personal Archives of Gloria Rudisch.
- . Letter to Marvin Minsky, 8 April 1954. Minsky 2AR. Personal Archives of Gloria Rudisch.

- Barbrook, Richard, and Andy Cameron. 'The Californian Ideology'. *Science as Culture* 6, no. 1 (January 1996): 44–72. <https://doi.org/10.1080/09505439609526455>.
- Barry, Andrew, Georgina Born, and Andrew Pickering, eds. 'Accountability Integration and Ontology'. In *Interdisciplinarity: Reconfigurations of the Social and Natural Sciences*. Abingdon, Oxon; New York, NY: Routledge, 2013.
- Bartlett, F.C. *Thinking: An Experimental and Social Study*. Unwin University Books. Basic Books, 1958.
- Bashe, Charles J., ed. *IBM's Early Computers*. MIT Press Series in the History of Computing. Cambridge, Mass: MIT Press, 1986.
- Bateson, Gregory, Don D. Jackson, Jay Haley, and John Weakland. 'Toward a Theory of Schizophrenia'. *Behavioral Science* 1, no. 4 (1956): 251–64. <https://doi.org/10.1002/bs.3830010402>.
- Bedeian, Arthur G., and Daniel A. Wren. 'Most Influential Management Books of the 20th Century'. *Organizational Dynamics* 29, no. 3 (December 2001): 221–25. [https://doi.org/10.1016/S0090-2616\(01\)00022-5](https://doi.org/10.1016/S0090-2616(01)00022-5).
- Bemer, Bob. 'Origins of Time Sharing', 30 June 2012. <http://archive.is/gRSr>.
- Bendig, A. W., and William J. Meyer. 'A Longitudinal Study of the Primary Mental Abilities Test'. *Journal of Educational Psychology* 52, no. 1 (February 1961): 50–60.
- Benesch, Samuel E. Letter to Marvin Minsky, 30 December 1953. Minsky 2AR. Personal Archives of Gloria Rudisch.
- Benjamin, Ruha. *Race after Technology: Abolitionist Tools for the New Jim Code*. Medford, MA: Polity, 2019.
- Bergin, Thomas J., ed. 'ENIAC: Development and Early Days'. In *50 Years of Army Computing: From ENIAC to MSRC: A Record of a Symposium and Celebration, November 13 and 14, 1996, Aberdeen Proving Ground*, 24–37. Army Research Lab Technical Publishing Branch, 2000.
- Berkeley, Edmund. *Giant Brains or Machines That Think*. New York: Science Editions, Inc, 1961.
- Bernstein, J. A. I. 'Marvin Minsky's Vision of the Future'. *The New Yorker*, 14 December 1981. <http://www.newyorker.com/magazine/1981/12/14/a-i>.
- . 'Marvin Minsky's Vision of the Future'. *The New Yorker*, 14 December 1981. <https://www.newyorker.com/magazine/1981/12/14/a-i>.
- Bernstein, Michael A. *A Perilous Progress: Economists and Public Purpose in Twentieth-Century America*. Princeton, NJ: Princeton Univ. Press, 2001.
- Berrien, F. K. Review of *Principles of Industrial Psychology*, by Thomas A. Ryan and Patricia C. Smith. *The American Journal of Psychology* 67, no. 4 (December 1954): 739. <https://doi.org/10.2307/1418506>.
- Bjork, Olin. 'Selfridge's Milton: The Legacy of Pand(a)Emonium in Cognitive Science and Artificial Intelligence'. *Journal of Literature and Science* 11, no. 2 (2018). <http://www.literatureandscience.org/wp-content/uploads/2018/12/BJORK-FINAL.pdf>.
- Blair, Ann. *Too Much to Know: Managing Scholarly Information before the Modern Age*. New Haven London: Yale University Press, 2010.
- Bobzien, Susanne. 'Ancient Logic'. In *Stanford Encyclopaedia of Philosophy*, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University, Winter 2016. <https://plato.stanford.edu/entries/logic-ancient/>.

- Bode, H., F. Mosteller, J. Tukey, and C. Winsor. 'The Education of a Scientific Generalist'. *Science* 109, no. 2840 (3 June 1949): 553–58. <https://doi.org/10.1126/science.109.2840.553>.
- Boden, Margaret A. *Mind as Machine: A History of Cognitive Science*. 2 vols. Oxford : New York: Clarendon Press ; Oxford University Press, 2006.
- Bolter, J. David. *Turing's Man: Western Culture in the Computer Age*. Chapel Hill: University of North Carolina Press, 1984.
- Booth, Kathleen, and Andrew Booth. *Automatic Digital Calculators*. Third. London: Butterworths, 1965.
- Booth, Kathleen H. V., and Andrew D. Booth. *Automatic Digital Calculators*. First. London: Butterworth Scientific Publications, 1953.
- Bouk, Daniel B. *How Our Days Became Numbered: Risk and the Rise of the Statistical Individual*. Chicago and London: The University of Chicago Press, 2015.
- Boulding, Kenneth Ewart. *Peace and the War Industry*. 2 ed. Transaction/Society Book Series, 11. New Brunswick, N.J: Transaction Books, 1973.
- Brock, David C. 'From Automation to Silicon Valley: The Automation Movement of the 1950s, Arnold Beckman, and William Shockley'. *History and Technology* 28, no. 4 (December 2012): 375–401. <https://doi.org/10.1080/07341512.2012.756236>.
- Brock, Gerald W. *The Second Information Revolution*. Cambridge, Mass: Harvard University Press, 2003.
- Broussard, Meredith. *Artificial Unintelligence: How Computers Misunderstand the World*. Cambridge, Massachusetts: The MIT Press, 2018.
- Brown, JoAnne. *The Definition of a Profession: The Authority of Metaphor in the History of Intelligence Testing, 1890-1930*. Princeton, N.J: Princeton University Press, 1992.
- Browne, Simone. *Dark Matters: On the Surveillance of Blackness*. Durham: Duke University Press, 2015.
- Building Computers in 1953: The Johnniac*. Computer History Museum, n.d. [https://www.youtube.com/watch?v=cVnOCT\\_r3so](https://www.youtube.com/watch?v=cVnOCT_r3so).
- Buolamwini, Joy, and Gebru. 'Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification'. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, Vol. 81, 2018.
- Bush, Robert R. Letter to Marvin Minsky. 'Reprint Request', 21 October 1963. Minsky 6AR. Personal Archives of Gloria Rudisch.
- Caldwell, Bruce. 'Some Reflections on F.A. Hayek's The Sensory Order'. *Journal of Bioeconomics* 6, no. 3 (2004): 239–54. <https://doi.org/10.1007/s10818-004-5505-9>.
- 'California Institute of Technology Commencement Exercises, June 11, 1948'. California Institute of Technology, 11 June 1948. <http://resolver.caltech.edu/CaltechCampusPubs:20111213-152042892>.
- Campbell-Kelly, Martin. 'Programming the EDSAC: Early Programming Activity at the University of Cambridge'. *IEEE Annals of the History of Computing* 20, no. 4 (1998): 46–67.
- Campbell-Kelly, Martin, and Daniel D. Garcia-Swartz. 'Pragmatism, Not Ideology: Historical Perspectives on IBM's Adoption of Open-Source Software'. *Information Economics and Policy* 21, no. 3 (August 2009): 229–44. <https://doi.org/10.1016/j.infoecopol.2009.03.006>.
- Campbell-Kelly, Martin, Cliff B. Jones, and John L. Lloyd. 'From Theory to Practice: The Invention of Programming, 1947-51'. In *Dependable and Historic Computing: Essays*



- Dedicated to Brian Randell on the Occasion of His 75th Birthday*, 1st ed., 23–38. Lecture Notes in Computer Science 6875. New York, NY: Springer-Verlag Berlin Heidelberg, 2011.
- Camus, Albert, and Justin O'Brien. *Resistance, Rebellion, and Death*. New York: Vintage International, 1995.  
<http://banq.lib.overdrive.com/ContentDetails.htm?id=DED5F7EB-4FA9-442C-83FB-E223E40DE336>.
- Carnap, Rudolf. *The Logical Syntax of Language*. London: Routledge, 1937.
- Carr, John. Letter to J. Backus, R. Eberhart, D. Sayre, A. Cohen, G. Hopper, W. Clark, C. Eldert, et al. 'Memo Re: Where Do Computers – Hardware and Programming – Go from Here?', 17 June 1957. Unlisted - Letters 1960; 1960 Letters 2-3. Personal Archives of Gloria Rudisch.
- Carr, John W. Letter to B. A. Galler, Saul Gorn, A. J. Perlis, A. W. Burks, N. R. Scott, H. Garner, M. L. Minsky, et al., 18 April 1957. Unlisted - Letters 1960; 1960 Letters 4-6. Personal Archives of Gloria Rudisch.
- . Letter to A. J. Perlis, M. L. Minsky, Saul Gorn, and Allan Newell, 20 June 1957. Unlisted - Letters 1960; 1960 Letters 7-8. Personal Archives of Gloria Rudisch.
- Carson, Cathryn, ed. *Reappraising Oppenheimer: Centennial Studies and Reflections*. Berkeley Papers in History of Science 21. Berkeley, Calif: Office for History of Science and Technology, Univ. of California, 2005.
- Carson, John. 'The Culture of Intelligence'. In *The Cambridge History of Science*, edited by Theodore M. Porter and Dorothy Ross, 1st ed., 635–48. Cambridge University Press, 2003. <https://doi.org/10.1017/CHOL9780521594424.038>.
- Castelle, Michael. 'The Social Lives of Generative Adversarial Networks'. New York, NY: ACM, 2020.
- Cave, Stephen. 'Intelligence: A History'. *Aeon*, 21 February 2017. <https://aeon.co/essays/on-the-dark-history-of-intelligence-as-domination>.
- . 'The Problem with Intelligence: Its Value-Laden History and the Future of AI'. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 29–35. New York NY USA: ACM, 2020. <https://doi.org/10.1145/3375627.3375813>.
- Cave, Stephen, and Kanta Dihal. 'Ancient Dreams of Intelligent Machines: 3,000 Years of Robots'. *Nature* 559, no. 7715 (July 2018): 473–75. <https://doi.org/10.1038/d41586-018-05773-y>.
- Chaney, Anthony. *Runaway: Gregory Bateson, the Double Bind, and the Rise of Ecological Consciousness*. Chapel Hill: University of North Carolina Press, 2017.
- Cheney, Jim, and Center for Environmental Philosophy, The University of North Texas. 'Postmodern Environmental Ethics: Ethics of Bioregional Narrative'. *Environmental Ethics* 11, no. 2 (1989): 117–34. <https://doi.org/10.5840/enviroethics198911231>.
- Cherry, Steven. 'Remembering John McCarthy [Spectral Lines]'. *IEEE Spectrum* 48, no. 12 (December 2011): 10–10. <https://doi.org/10.1109/MSPEC.2011.6085765>.
- Chisholm, Rupert F. 'The Storing Critique Revisited: Simon as Seen in the Science of Politics'. *Public Administration Quarterly* 12, no. 4 (1989): 411–36.
- Christie, Lee, R. Duncan Luce, and Josiah Macy Jr. 'Communication and Learning in Task-Oriented Groups'. Cambridge, Mass.: Massachusetts Institute of Technology, 13 May 1952.

- Clark, Wesley, and Belmont Farley. 'Generalization of Pattern Recognition in a Self-Organizing System'. In *Proceedings of the 1955 Western Joint Computer Conference*, 86–91. New York: Institute of Radio Engineers, 1955.
- . 'Simulation of Self-Organizing Systems by Digital Computer'. In *IRE Transactions on Information Theory*, 4:76–84. New York: Institute of Radio Engineers, 1954.
- Coady, David, and James Chase, eds. *The Routledge Handbook of Applied Epistemology*. Routledge Handbooks in Philosophy. London ; New York: Routledge, Taylor & Francis Group, 2019.
- Cohen, I. Bernard. *Howard Aiken: Portrait of a Computer Pioneer*. History of Computing. Cambridge, MA: MIT Press, 1999.
- Cohen, Noam. *The Know-It-Alls: The Rise of Silicon Valley as a Political Powerhouse and Social Wrecking Ball*. New York: The New Press, 2017.
- Cohen-Cole, Jamie. 'Review of Mind as Machine: A History of Cognitive Science by Margaret Boden'. *Isis* 99, no. 4 (December 2008): 811–12. <https://doi.org/10.1086/597681>.
- . *The Open Mind: Cold War Politics and the Sciences of Human Nature*. Chicago London: The University of Chicago Press, 2014.
- Collins, Martin J. 'Planning for Modern War: Rand and the Air Force, 1945-1950'. PhD Dissertation, University of Maryland, 1998.
- Commons, John Rogers. *Institutional Economics. Vol. I: Its Place in Political Economy*. Transaction Publishers, 1990.
- Conquest, Wendy, Bob Drake, and Dan Rockmore. *Mind in the Machine: The Discovery of Artificial Intelligence*, 2006.  
<https://www.cs.dartmouth.edu/~rockmore/MindMachine.html>.
- Conway, Flo, and Jim Siegelman. *Dark Hero of the Information Age: In Search of Norbert Wiener, the Father of Cybernetics*. New York: Basic Books, 2006.
- Cook, Thomas I. 'Review of Administrative Behavior, a Study of Decision-Making Processes in Administrative Organization'. *The Review of Politics* 12, no. 3 (1950): 405–7. <https://doi.org/10.1017/S0034670500047082>.
- Copeland, B. 'The Manchester Computer: A Revised History Part 1: The Memory'. *IEEE Annals of the History of Computing* 33, no. 1 (January 2011): 4–21. <https://doi.org/10.1109/MAHC.2010.1>.
- Copeland, B. Jack. *Artificial Intelligence: A Philosophical Introduction*. Oxford, UK ; Cambridge, Mass: Blackwell, 1993.
- Cordeschi, Roberto. 'AI Turns Fifty: Revisiting Its Origins'. *Applied Artificial Intelligence* 21, no. 4–5 (25 April 2007): 259–79. <https://doi.org/10.1080/08839510701252304>.
- . *The Discovery of the Artificial: Behavior, Mind, and Machines Before and Beyond Cybernetics*. Studies in Cognitive Systems, v. 28. Dordrecht ; Boston: Kluwer Academic Publishers, 2002.
- . 'The Discovery of the Artificial. Some Protocybernetic Developments 1930–1940'. *AI & Society* 5, no. 3 (July 1991): 218–38. <https://doi.org/10.1007/BF01891917>.
- 'Cornell Professor Dies'. *Journal (Unspecified)*. 13 July 1971. Cornell Faculty Bio files, Box 1 - Frank Rosenblatt Publications 1956-58. Division of Rare and Manuscript Collections, Cornell University Library.
- Correll, Quentin, Otto Matty Khodr, and Alexander Vanderburgh. 'Letters to the Editor'. *Communications of the ACM* 1, no. 7 (1 July 1958): 2–3. <https://doi.org/10.1145/368873.368877>.

- Couldry, Nick, and Ulises Ali Mejias. *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Culture and Economic Life. Stanford, California: Stanford University Press, 2019.
- Cowles, Henry M. *The Scientific Method: An Evolution of Thinking from Darwin to Dewey*. Cambridge, Massachusetts: Harvard University Press, 2020.
- Craik, Kenneth James Williams, and Stephen L Sherwood. 'The Mechanism of Human Action'. In *A Selection of Papers, Essays and Other Writings by the Late Kenneth J.W. Craik*, 7–92. Cambridge: C.U.P., 1966.
- Crevier, Daniel. *AI: The Tumultuous History of the Search for Artificial Intelligence*. New York, NY: Basic Books, 1993.
- Crowther-Heyck, Hunter. *Age of System: Understanding the Development of Modern Social Science*. Baltimore: Johns Hopkins University Press, 2015.
- . *Herbert A. Simon: The Bounds of Reason in Modern America*. Baltimore: Johns Hopkins University Press, 2005.
- Culbertson, J. T. *Consciousness and Behavior*. Dubuque, Iowa: Wm. C. Brown Co., 1950.
- Dahl, Robert A. Edited by Herbert Simon. *Administrative Science Quarterly* 2, no. 2 (1957): 244–48. <https://doi.org/10.2307/2390693>.
- Daston, Lorraine. 'Calculation and the Division of Labor, 1750-1950'. *Bulletin of the German Historical Institute* 62, no. Spring (2018): 9–30.
- . 'Enlightenment Calculations'. *Critical Inquiry* 21, no. 1 (1994): 182–202.
- . 'Simon and the Sirens: A Commentary'. *Isis* 106, no. 3 (September 2015): 669–76. <https://doi.org/10.1086/683531>.
- Davies, Thom. 'Slow Violence and Toxic Geographies: "Out of Sight" to Whom?' *Environment and Planning C: Politics and Space*, 10 April 2019, 239965441984106. <https://doi.org/10.1177/2399654419841063>.
- Dick, Stephanie. 'After Math: (Re)Configuring Minds, Proof and Computing in the Postwar United States'. University of Harvard, 2014.
- . 'AfterMath: The Work of Proof in the Age of Human–Machine Collaboration'. *Isis* 102, no. 3 (September 2011): 494–505. <https://doi.org/10.1086/661623>.
- . 'Of Models and Machines: Implementing Bounded Rationality'. *Isis* 106, no. 3 (September 2015): 623–34. <https://doi.org/10.1086/683527>.
- Dinneen, Gerald P. 'Programming Pattern Recognition'. In *Proceedings of the 1955 Western Joint Computer Conference*, 94–100. New York: Institute of Radio Engineers, 1955. <https://doi.org/10.1145/1455292.1455311>.
- 'Display Ad 190 -- No Title'. *The New York Times*, 2 October 1960. New York Times (1923–Current file) pg. F22. ProQuest Historical Newspapers: The New York Times with Index.
- 'Display Ad 552 -- No Title'. *The New York Times*, 1 July 1962. New York Times (1923–Current file) pg. 101. ProQuest Historical Newspapers: The New York Times with Index.
- 'Display Ad 701 -- No Title'. *The New York Times*, 14 October 1962. New York Times (1923–Current file) pg. 171. ProQuest Historical Newspapers: The New York Times with Index.
- 'Display Ad 752 -- No Title'. *The New York Times*, 6 May 1962. New York Times (1923–Current file) pg. 196. ProQuest Historical Newspapers: The New York Times with Index.
- 'Display Ad 798 -- No Title'. *The New York Times*, 4 November 1962. New York Times (1923–Current file) pg. 171. ProQuest Historical Newspapers: The New York Times with Index.

- Index. Retrieved from  
<https://ezp.lib.cam.ac.uk/login?url=https://search.proquest.com/docview/116258670?accountid=9851>.
- Douglass, Frederick. 'The Meaning of July Fourth for the Negro'. Corinthian Hall, Rochester, New York, 5 July 1852.
- Dowe, David L. 'Introduction to Ray Solomonoff 85th Memorial Conference'. In *Algorithmic Probability and Friends. Bayesian Prediction and Artificial Intelligence*, edited by David L. Dowe, 7070:1–36. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013. [https://doi.org/10.1007/978-3-642-44958-1\\_1](https://doi.org/10.1007/978-3-642-44958-1_1).
- 'Dr. Frank Rosenblatt Dies at 43; Taught Neurobiology at Cornell'. *The New York Times*, 13 July 1971. <https://www.nytimes.com/1971/07/13/archives/dr-frank-roosenblatt-dies-at-43-taught-neurobiology-at-cornell.html>.
- Draaisma, D. *Metaphors of Memory: A History of Ideas about the Mind*. Cambridge, U.K. ; New York: Cambridge University Press, 2000.
- Dreyfus, Hubert L. 'Alchemy and Artificial Intelligence'. Santa Monica, CA: RAND Corporation, December 1965. <https://www.rand.org/pubs/papers/P3244.html>.
- . *What Computers Can't Do; a Critique of Artificial Reason*. 1st ed. New York: Harper & Row, 1972.
- . *What Computers Still Can't Do: A Critique of Artificial Reason*. Cambridge, Mass: MIT Press, 1992.
- . 'Why Heideggerian AI Failed and How Fixing It Would Require Making It More Heideggerian'. In *The Mechanical Mind in History*, edited by Phil Husbands, Owen Holland, and Michael Wheeler, 331–72. Cambridge, Mass: MIT Press, 2008.
- Dupuy, Jean-Pierre. *The Mechanization of the Mind: On the Origins of Cognitive Science*. New French Thought. Princeton: Princeton University Press, 2000.
- Dyson, George. *Darwin Among the Machines*. London: Allen Lane, The Penguin Press, 1998.
- Edwards, Paul N. *The Closed World: Computers and the Politics of Discourse in Cold War America*. Inside Technology. Cambridge, Mass. London: MIT, 1997.
- 'Electronic "Brain" Teaches Itself'. *The New York Times*, 13 July 1958.
- Engerman, David C. 'Social Science in the Cold War'. *Isis* 101, no. 2 (June 2010): 393–400. <https://doi.org/10.1086/653106>.
- Ensmenger, Nathan. *The Computer Boys Take Over: Computers, Programmers, and the Politics*. History of Computing. Cambridge, Mass: MIT Press, 2010.
- . 'The Multiple Meanings of a Flowchart'. *Information & Culture* 51, no. 3 (August 2016): 321–51. <https://doi.org/10.7560/IC51302>.
- Erickson, Paul. *How Reason Almost Lost Its Mind: The Strange Career of Cold War Rationality*. Chicago ; London: The University of Chicago Press, 2013.
- . *The World the Game Theorists Made*. Chicago : London: The University of Chicago Press, 2015.
- Eubanks, Virginia. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. First Edition. New York, NY: St. Martin's Press, 2017.
- Eysenck, Hans J., ed. *A Model for Intelligence*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1982. <https://doi.org/10.1007/978-3-642-68664-1>.
- 'Faculty Information: Frank Rosenblatt'. Kroch Library, n.d. Frank Rosenblatt. Cornell University Faculty Biographical files, #47-10-3394. Division of Rare and Manuscript Collections, Cornell University Library.

- Feigenbaum, Edward A., and Julian Feldman, eds. *Computers and Thought*. New York: McGraw-Hill, 1963.
- Feigenbaum, Edward A., Julian Feldman, and Paul Armer, eds. *Computers and Thought*. 1st AAAI Press ed., 1995. Menlo Park : Cambridge: AAAI Press ; MIT Press, 1995.
- Fein, Louis. 'The Role of the University in Computers, Data Processing, and Related Fields'. In *Papers Presented at the March 3-5, 1959, Western Joint Computer Conference on XX - IRE-AIEE-ACM '59 (Western)*, 119–26. San Francisco, California: ACM Press, 1959. <https://doi.org/10.1145/1457838.1457859>.
- Feldman, Julian. 'Computers and Thought—The Back Story'. *IEEE Annals of the History of Computing* 33, no. 4 (April 2011): 56–61. <https://doi.org/10.1109/MAHC.2011.69>.
- Ferry, Georgina. *A Computer Called LEO: Lyons Teashops and the World's First Office Computer*. London: Fourth Estate, 2003.
- Fetzer, J. H. 'Program Verification: The Very Idea'. *Communications in ACM* 31, no. 9 (1988): 1048–63.
- 'First Hixon Symposium on "Cerebral Mechanisms and Behavior"'. *Engineering and Science*, October 1948.
- Flamm, Kenneth. *Creating the Computer: Government, Industry, and High Technology*. Washington, D.C: Brookings Institution, 1988.
- Fleck, James. 'Development and Establishment in Artificial Intelligence'. In *Scientific Establishments and Hierarchies*, edited by Norbert Elias, Herminio Martins, and Richard Whitley, 169–217. Dordrecht: Springer Netherlands, 1982. [http://dx.doi.org/10.1007/978-94-009-7729-7\\_7](http://dx.doi.org/10.1007/978-94-009-7729-7_7).
- 'For Nortronics, a Division of Northrop. Classified Ad 682 -- No Title'. *The New York Times*, 12 August 1962. New York Times (1923-Current file) pg. W21. ProQuest Historical Newspapers: The New York Times with Index.
- Forman, Paul. 'On the Historical Forms of Knowledge Production and Curation: Modernity Entailed Disciplinarity, Postmodernity Entails Antidisciplinarity'. *Osiris* 27, no. 1 (January 2012): 56–97. <https://doi.org/10.1086/667823>.
- . 'The Primacy of Science in Modernity, of Technology in Postmodernity, and of Ideology in the History of Technology'. *History and Technology* 23, no. 1–2 (1 March 2007): 1–152. <https://doi.org/10.1080/07341510601092191>.
- Fortun, M., and S. S. Schweber. 'Scientists and the Legacy of World War II: The Case of Operations Research (OR)'. *Social Studies of Science* 23, no. 4 (1993): 595–642.
- 'Frank Rosenblatt Dies; Researcher, Politician'. *Post Standard*. 14 July 1971. Cornell Faculty Bio files, Box 1 - Frank Rosenblatt Publications 1956-58. Division of Rare and Manuscript Collections, Cornell University Library.
- Friedberg, R. M. 'A Learning Machine: Part I'. *IBM Journal of Research and Development* 2, no. 1 (January 1958): 2–13. <https://doi.org/10.1147/rd.21.0002>.
- Fritz, W.B. 'The Women of ENIAC'. *IEEE Annals of the History of Computing* 18, no. 3 (Fall 1996): 13–28. <https://doi.org/10.1109/85.511940>.
- Galison, Peter. 'The Ontology of the Enemy: Norbert Wiener and the Cybernetic Vision'. *Critical Inquiry* 21, no. 1 (Autumn 1994): 228–66.
- Gardner, Howard. *The Mind's New Science: A History of the Cognitive Revolution*. Paperback ed., Nachdr. New York: Basic Books, 1998.
- Garnham, Alan. *Artificial Intelligence: An Introduction*. Introductions to Modern Psychology. London ; New York: Routledge & Kegan Paul, 1987.

- Garvey, Colin. 'The "General Problem Solver" Doesn't Exist: Mortimer Taube & The Art of AI Criticism'. St. Louis, 2018.
- Gibby, Robert E., and Michael J. Zickar. 'A History of the Early Days of Personality Testing in American Industry: An Obsession with Adjustment.' *History of Psychology* 11, no. 3 (August 2008): 164–84. <https://doi.org/10.1037/a0013041>.
- Gibson, James J., Paul Olum, and Frank Rosenblatt. 'Parallax and Perspective during Aircraft Landings'. *The American Journal of Psychology* 68, no. 3 (September 1955): 372. <https://doi.org/10.2307/1418521>.
- Goldstein, Andrew. Oral History: Nathaniel Rochester, 6 June 1991. [http://ethw.org/Oral-History:Nathaniel\\_Rochester](http://ethw.org/Oral-History:Nathaniel_Rochester).
- Goldstine, Herman Heine. *The Computer: From Pascal to von Neumann*. Princeton, N.J: Princeton University Press, 1993.
- Gordin, Michael D. *Scientific Babel: How Science Was Done before and after Global English*. Chicago ; London: The University of Chicago Press, 2015.
- Gould, Stephen Jay. *The Mismeasure of Man*. New York: Norton, 1996.
- Gray, Jeremy. *Plato's Ghost: The Modernist Transformation of Mathematics*. Princeton, N.J: Princeton University Press, 2008.
- Grier, David Alan. 'Human Computers: The First Pioneers of the Information Age'. *Endeavour* 25, no. 1 (March 2001): 28–32. [https://doi.org/10.1016/S0160-9327\(00\)01338-7](https://doi.org/10.1016/S0160-9327(00)01338-7).
- Grosch, H. R. J. 'After-Luncheon Remarks', 124–25. ACM Press, 1953. <https://doi.org/10.1145/1434878.1434880>.
- Gruenberger, F. J. 'Memorandum, RM-5654-PR: The History of the JOHNNIAC'. The Rand Corporation, October 1968.
- . 'The History of the JOHNNIAC'. *Annals of the History of Computing* 1, no. 1 (July 1979): 49–64.
- Gugerty, Leo. 'Newell and Simon's Logic Theorist: Historical Background and Impact on Cognitive Modeling'. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 50, no. 9 (October 2006): 880–84. <https://doi.org/10.1177/154193120605000904>.
- Guice, Jon. 'Controversy and the State: Lord ARPA and Intelligent Computing'. *Social Studies of Science* 28, no. 1 (1998): 103–38.
- Gulick, Luther, Lyndall Urwick, and Carl Pforzheimer. *Papers on the Science of Administration*. New York: Institute of Public Administration, 1937.
- Guthrie, Stewart. *Faces in the Clouds: A New Theory of Religion*. New York: Oxford University Press, 1993.
- H. Leahey, Thomas. 'The Mythical Revolution of American Psychology'. *American Psychologist* 47, no. 2 (February 1992): 308–18. <https://doi.org/10.1037/0003-066X.47.2.308>.
- Hacking, Ian. *The Taming of Chance*. Ideas in Context. Cambridge [England] ; New York: Cambridge University Press, 1990.
- Haigh, Thomas. 'Masculinities in the Histories of Computing(s)', 2008. <http://www.tomandmaria.com/Tom/Writing/Haigh%20Masculinity%20for%20Workshop.pdf>.
- Haigh, Thomas, Mark Priestley, and Crispin Rope. *ENIAC in Action: Making and Remaking the Modern Computer*, 2018.
- . 'Reconsidering the Stored-Program Concept'. *IEEE Annals of the History of Computing* 36, no. 1 (2014): 4–17. <https://doi.org/10.1109/MAHC.2013.56>.

- Hall, Stuart, and B. Gielben, eds. 'The West and the Rest: Discourse and Power'. In *Formations of Modernity*, 276–320. Understanding Modern Societies 1. Cambridge: Polity Press and Open University, 1992.
- Hamming, R. W., and C. Berkeley. 'Giant Brains, or Machines That Think.' *The American Mathematical Monthly* 58, no. 4 (April 1951): 276. <https://doi.org/10.2307/2306984>.
- Haraway, Donna Jeanne. *Staying with the Trouble: Making Kin in the Chthulucene*. Experimental Futures: Technological Lives, Scientific Arts, Anthropological Voices. Durham: Duke University Press, 2016.
- Hashagen, Ulf, Reinhard Keil-Slawik, Arthur L. Norberg, and Heinz Nixdorf MuseumsForum, eds. *History of Computing: Software Issues: International Conference on the History of Computing, ICHC 2000, April 5-7, 2000, Heinz Nixdorf MuseumsForum, Paderborn, Germany*. Berlin ; New York: Springer, 2002.
- Haugeland, John. *Artificial Intelligence: The Very Idea*. 6. print. Bradford Books. Cambridge, Mass.: MIT Press, 1993.
- Hayek, F. A. *The Sensory Order: An Inquiry into the Foundations of Theoretical Psychology*. Chicago: University of Chicago Press, 1952.
- Hebb, Donald O. *The Organization of Behavior: A Neuropsychological Theory*. New York: John Wiley, Inc, 1949.
- Heims, Steve J. *John Von Neumann and Norbert Wiener: From Mathematics to the Technologies of Life and Death*. Cambridge, Mass: MIT Press, 1980.
- . *The Cybernetics Group*. Cambridge, Mass: MIT Press, 1991.
- Heyck, Hunter. 'Defining the Computer: Herbert Simon and the Bureaucratic Mind--Part 1'. *IEEE Annals of the History of Computing* 30, no. 2 (April 2008): 42–51. <https://doi.org/10.1109/MAHC.2008.18>.
- . 'Defining the Computer: Herbert Simon and the Bureaucratic Mind--Part 2'. *IEEE Annals of the History of Computing* 30, no. 2 (April 2008): 52–63. <https://doi.org/10.1109/MAHC.2008.19>.
- . 'Producing Reason'. In *Cold War Social Science: Knowledge Production, Liberal Democracy, and Human*, edited by Mark Solovey and Hamilton Cravens. New York: Palgrave Macmillan US, 2012. <https://doi.org/10.1057/9781137013224>.
- Hicks, Marie. 'Hacking the Cis-Tem'. *IEEE Annals of the History of Computing* 41, no. 1 (1 January 2019): 20–33. <https://doi.org/10.1109/MAHC.2019.2897667>.
- . 'Hacking the Cis-Tem'. *IEEE Annals of the History of Computing* 41, no. 1 (1 January 2019): 20–33. <https://doi.org/10.1109/MAHC.2019.2897667>.
- . *Programmed Inequality: How Britain Discarded Women Technologists and Lost Its Edge in Computing*. History of Computing. Cambridge, MA: MIT Press, 2017.
- Hilts, Philip J. *Scientific Temperaments: Three Lives in Contemporary Science*. New York: Simon and Schuster, 1982.
- IBM. 'History of IBM, 1950s, 1954', 1 December 2019. [https://www.ibm.com/ibm/history/history/year\\_1954.html](https://www.ibm.com/ibm/history/history/year_1954.html).
- Hoag, Malcolm W. 'An Introduction to Systems Analysis'. Research Memoranda, 1956.
- Hobbes, Thomas. *Leviathan*. Edited by Christopher Brooke. Penguin Classics. Harmondsworth, Meddlesex: Penguin Books, 2017.
- Hobbs, Charles F. Letter to Marvin Minsky, 14 May 1959. Minsky 6AR. Personal Archives of Gloria Rudisch.
- Holland, John H. Letter to Marvin Minsky, 23 December 1959. Unlisted - Letters 1960; 1960 Letters 1. Personal Archives of Gloria Rudisch.

- Horan, Caley Dawn. 'Actuarial Age: Insurance and the Emergence of Neoliberalism in the Postwar United States'. Dissertation, University of Minnesota, 2011.  
<http://hdl.handle.net/11299/115896>.
- Hounshell, David. 'The Cold War, RAND, and the Generation of Knowledge, 1946-1962'. *Historical Studies in the Physical and Biological Sciences* 27, no. 2 (1997): 237–67.  
<https://doi.org/10.2307/27757779>.
- Hughes, Agath C., and Thomas Parke Hughes. *Systems, Experts, and Computers: The System Approach in Management in Management and Engineering, World War II and After*. Cambridge, Massachusetts; London, England: MIT Press, 2000. Doi: 10.2307/j.ctt5vjs3h.
- Hughes, Thomas P. *Rescuing Prometheus*. 1. Vintage Books ed. New York: Vintage Books, 2000.
- Hunter, Lynette. 'Rhetoric and Artificial Intelligence'. *Rhetorica* 9, no. 4 (November 1991): 317–40. <https://doi.org/10.1525/rh.1991.9.4.317>.
- Hunter, Patti W. 'Drawing the Boundaries: Mathematical Statistics in 20th-Century America'. *Historia Mathematica* 23, no. 1 (February 1996): 7–30.  
<https://doi.org/10.1006/hmat.1996.0002>.
- Husbands, Phil, and Owen Holland. 'The Ratio Club: A Hub of British Cybernetics', 91–148, 2008.
- Husbands, Phil, Owen Holland, and Michael Wheeler, eds. *The Mechanical Mind in History*. Cambridge, Mass: MIT Press, 2008.
- Huskey, H.D. 'SWAC-Standards Western Automatic Computer: The Pioneer Day Session at NCC July 1978'. *IEEE Annals of the History of Computing* 19, no. 2 (June 1997): 51–61.  
<https://doi.org/10.1109/85.586073>.
- Hutchins, John. *First Steps in Mechanical Translation*, 1997.  
<http://www.hutchinsweb.me.uk/MTS-1997.pdf>.
- 'IBM 704 Electronic Data-Processing Machine: Manual of Operation'. New York: International Business Machines Corporation, 1955. Minsky 6AR. Personal Archives of Gloria Rudisch.
- 'IBM: Preliminary Report: Specifications for the IBM Mathematical FORMula TRANSLating System, FORTRAN'. Progr. Res. Group, Appl. Sci. Div. IBM Corp., 1954.
- 'IBM: The FORTRAN Automatic Coding System for the IBM704 EDPM, Programmers Manual'. IBM Corp., 1956.
- Igo, Sarah E. *The Averaged American: Surveys, Citizens, and the Making of a Mass Public*. Cambridge: Harvard University Press, 2008.  
<http://VH7QX3XE2P.search.serialssolutions.com/?V=1.0&L=VH7QX3XE2P&S=JCS&C=TC0000108126&T=marc&tab=BOOKS>.
- . *The Known Citizen: A History of Privacy in Modern America*. Cambridge, Massachusetts: Harvard University Press, 2018.
- Ilin, M. *100,000 Whys: A Trip Around the Room*. Translated by Beatrice Kinkad. Philadelphia, PA: J.B. Lippincott Company, 1933.
- 'International Conference on Scientific Information', n.d. Minsky 3AR. Personal Archives of Gloria Rudisch.
- Irvine, M.M. 'Early Digital Computers at Bell Telephone Laboratories'. *IEEE Annals of the History of Computing* 23, no. 3 (September 2001): 22–42.  
<https://doi.org/10.1109/85.948904>.



- Isaac, Joel. 'The Human Sciences in Cold War America'. *The Historical Journal* 50, no. 3 (September 2007): 725–46. <https://doi.org/10.1017/S0018246X07006334>.
- Isaacs, Rufus. *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*. Mineola, N.Y: Dover Publications, 1999.
- Isaacson, Walter. *The Innovators: How a Group of Hackers, Geniuses, and Geeks Created the Digital Revolution*. New York: Simon & Schuster, 2014.
- Jebens, Holger. 'The Crisis of Anthropology'. *Paideuma* 56 (2010): 99–121.
- Jerábek, Hynek. *Paul Lazarsfeld and the Origins of Communications Research*. Milton Park, Abingdon, Oxon ; New York, NY: Routledge, 2017.
- Joint Computer Conference (1955 : Los Angeles, Calif.). *Proceedings of the Western Joint Computer Conference: Held March 1-3, 1955 in Los Angeles, Calif.* Western Joint Computer Conference. New York: Institute of Radio Engineers, 1955.
- Jones, A. 'Five 1951 BBC Broadcasts on Automatic Calculating Machines'. *IEEE Annals of the History of Computing* 26, no. 2 (April 2004): 3–15. <https://doi.org/10.1109/MAHC.2004.1299654>.
- Jones, Matthew L. 'How We Became Instrumentalists (Again): Data Positivism since World War II'. *Historical Studies in the Natural Sciences* 48, no. 5 (November 2018): 673–84. <https://doi.org/10.1525/hsns.2018.48.5.673>.
- . *Reckoning with Matter: Calculating Machines, Innovation, and Thinking about Thinking from Pascal to Babbage*. Chicago ; London: The University of Chicago Press, 2016.
- Kaiser, David. 'Nuclear Democracy: Political Engagement, Pedagogical Reform, and Particle Physics in Postwar America'. *Isis* 93, no. 2 (June 2002): 229–68. <https://doi.org/10.1086/344960>.
- Katz, Yarden. 'Manufacturing an Artificial Intelligence Revolution'. *SSRN Electronic Journal*, 2017. <https://doi.org/10.2139/ssrn.3078224>.
- Kay, Lily E. *The Molecular Vision of Life: Caltech, the Rockefeller Foundation, and the Rise of the New Biology*. New York Oxford: Oxford University Press, 1993.
- Kellogg, Oliver Dimon. *Foundations of Potential Theory*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1967. <https://doi.org/10.1007/978-3-642-86748-4>.
- Kemeny, John. Letter to E. B. Newman, 23 April 1956. Minsky 2AR. Personal Archives of Gloria Rudisch.
- Kerr, Gerry. 'The Development History and Philosophical Sources of Herbert Simon's *Administrative Behavior*'. *Journal of Management History* 13, no. 3 (3 July 2007): 255–68. <https://doi.org/10.1108/17511340710754707>.
- Kevles, Daniel J. *In the Name of Eugenics: Genetics and the Uses of Human Heredity*. 1st Harvard University Press pbk. ed. Cambridge, Mass: Harvard University Press, 1995.
- . *In the Name of Eugenics: Genetics and the Uses of Human Heredity*. 1st Harvard University Press pbk. ed. Cambridge, Mass: Harvard University Press, 1995.
- Keyes, Os. 'The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition'. *Proceedings of the ACM on Human-Computer Interaction* 2, no. CSCW (1 November 2018): 1–22. <https://doi.org/10.1145/3274357>.
- Kleene, S. C. 'Representations of Events in Nerve Nets and Finite Automata'. In *Automata Studies*, edited by Claude Elwood Shannon and J. McCarthy, 3–43. Annals of Mathematics Studies 34. Princeton, NJ: Princeton University Press, 1956.

- Klein, Judy L. 'Implementation Rationality: The Nexus of Psychology and Economics at the RAND Logistics Systems Laboratory, 1956–1966'. *History of Political Economy* 48, no. suppl 1 (2016): 198–225. <https://doi.org/10.1215/00182702-3619274>.
- Klein, Raymond M. 'The Hebb Legacy.' *Canadian Journal of Experimental Psychology/Revue Canadienne de Psychologie Expérimentale* 53, no. 1 (1999): 1–3. <https://doi.org/10.1037/h0087295>.
- Kline, Ronald. 'Cybernetics, Automata Studies, and the Dartmouth Conference on Artificial Intelligence'. *IEEE Annals of the History of Computing* 33, no. 4 (April 2011): 5–16. <https://doi.org/10.1109/MAHC.2010.44>.
- . *The Cybernetics Moment: Or Why We Call Our Age the Information Age*. New Studies in American Intellectual and Cultural History. Baltimore: Johns Hopkins University Press, 2015.
- . 'Why the Disunity of Cybernetics Matters to the History of the Human Sciences in the United States, 1940-1980'. Unpublished. Prepared for Cybernetics and Society Research Group Seminar, University of Cambridge, 7 April 2017.
- Kohler, Robert E. *Partners in Science: Foundations and Natural Scientists, 1900-1945*. Chicago: University of Chicago Press, 1991.
- Kramnick, Isaac, and Glenn C. Altschuler. *Cornell: A History, 1940–2015*. New York: Cornell University Press, 2014.
- Krippner, Greta R. *Capitalizing on Crisis: The Political Origins of the Rise of Finance*. First Harvard University Press paperback edition. Cambridge, Massachusetts London, England: Harvard University Press, 2012.
- Küçük, Harun. *Science without Leisure: Practical Naturalism in Istanbul, 1660-1732*, 2020. <https://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=2325414>.
- Kuhn, Thomas S., and Ian Hacking. *The Structure of Scientific Revolutions*. Fourth edition. Chicago ; London: The University of Chicago Press, 2012.
- Lanier, Jaron, and Glen Weyl. 'AI Is an Ideology, Not a Technology'. *Wired*, 15 March 2020. <https://www.wired.com/story/opinion-ai-is-an-ideology-not-a-technology/>.
- Large, David. *Ecological Philosophy*, 2003.
- Larson, Harry. 'The Computer Issue'. *Proceedings of the IRE* 49, no. 1 (January 1961): 4–7. <https://doi.org/10.1109/JRPROC.1961.287774>.
- Lauer, Josh. *Creditworthy: A History of Consumer Surveillance and Financial Identity in America*. Columbia Studies in the History of U.S. Capitalism. New York: Columbia University Press, 2017.
- Lazarsfeld, Paul F. 'The Logical and Mathematical Foundations of Latent Structure Analysis'. In *Measurement and Prediction, Volume IV of The American Soldier: Studies in Social Psychology in World War II.*, edited by S. A. Stouffer. Princeton University Press, 1950.
- Lechner, R. Letter to M. Cooperstein. 'The Mechanization of Human Thought Processes I: The Newell & Simon Approach, Conclusions and Alternatives', 21 April 1959. Unlisted - What is a Self Organizing System. Personal Archives of Gloria Rudisch.
- . Letter to M. Cooperstein. 'What Is a Self-Organizing System', 18 March 1959. Unlisted - What is a Self Organizing System. Personal Archives of Gloria Rudisch.
- Lee, J.A.N., J. McCarthy, and J.C.R. Licklider. 'The Beginnings at MIT'. *IEEE Annals of the History of Computing* 14, no. 1 (1992): 18–54. <https://doi.org/10.1109/85.145317>.

- Leech, Jessica. 'Logic and the Laws of Thought'. *Philosophers' Imprint* 15, no. 12 (March 2015). <http://hdl.handle.net/2027/spo.3521354.0015.012>.
- Lenzen, V. F., and K. J. W. Craik. 'The Nature of Explanation.' *The Philosophical Review* 53, no. 5 (September 1944): 503. <https://doi.org/10.2307/2181361>.
- Leslie, Stuart W. *The Cold War and American Science: The Military-Industrial-Academic Complex at MIT and Stanford*. New York: Columbia Univ. Press, 1993.
- Levy, Steven. *Hackers*. 1st ed. Sebastopol, CA: O'Reilly Media, 2010.
- Little, John D. C. 'Philip M. Morse and the Beginnings'. *Operations Research* 50, no. 1 (February 2002): 146–48. <https://doi.org/10.1287/opre.50.1.146.17799>.
- Loeb, Jacques. *Studies in General Physiology*. Chicago: University of Chicago Press, n.d.
- Long, F. A. Letter to D. R. Corson and W. R. Keast. 'Procedures for Incorporating Frank Rosenblatt into a Cornell Department and for Considering Him for Promotion to a Tenure Faculty Position', 11 March 1964. Dale R Corson papers, 1963-1977; Box 20 Folder 13. Division of Rare and Manuscript Collections, Cornell University Library.
- Longo, B. 'Edmund Berkeley, Computers, and Modern Methods of Thinking'. *IEEE Annals of the History of Computing* 26, no. 4 (October 2004): 4–18. <https://doi.org/10.1109/MAHC.2004.28>.
- Lorenzo, Mark Jones. *Endless Loop: The History of the BASIC Programming Language*. CreateSpace Independent Publishing Platform, 2017.
- Lowel, S, and W Singer. 'Selection of Intrinsic Horizontal Connections in the Visual Cortex by Correlated Neuronal Activity'. *Science* 255, no. 5041 (10 January 1992): 209–12. <https://doi.org/10.1126/science.1372754>.
- Lowen, Rebecca S. *Creating the Cold War University: The Transformation of Stanford*. Berkeley: University of California Press, 1997.
- Mahoney, Michael S. 'The Histories of Computing(s)'. *Interdisciplinary Science Review* 30, no. 2 (2005): 119–35.
- . 'The Histories of Computing(s)'. *Interdisciplinary Science Reviews* 30, no. 2 (June 2005): 119–35. <https://doi.org/10.1179/030801805X25927>.
- Mandel, Richard. 'Beacon of Hope: The Clinical Center Throughout Forty Years of Growth and Change in Biomedicine'. U.S. Department of Health and Human Services, Public Health Service, National Institutes of Health, 1993. [https://history.nih.gov/exhibits/beacon/Beacon\\_of\\_Hope.pdf](https://history.nih.gov/exhibits/beacon/Beacon_of_Hope.pdf).
- Mandler, Peter. *Return from the Natives: How Margaret Mead Won the Second World War and Lost the Cold War*. New Haven: Yale University Press, 2013.
- Marchbanks, R. M., F. Rosenblatt, and R. D. O'Brien. 'Serotonin Binding to Nerve-Ending Particles of the Rat Brain and Its Inhibition by Lysergic Acid Diethylamide'. *Science* 144, no. 3622 (29 May 1964): 1135–37. <https://doi.org/10.1126/science.144.3622.1135>.
- Markoff, John. *Machines of Loving Grace: The Quest for Common Ground between Humans and Robots*. New York: ecco, an imprint of HarperCollinsPublishers, 2015.
- Marsh, Leslie. 'Hayek: Cognitive Scientist Avant La Lettre'. In *The Social Science of Hayek's 'the Sensory Order'*. Volume 13, edited by William N Butos. Bingley: Emerald Group Pub., 2010.
- Masís, Jethro. 'Making AI Philosophical Again: On Philip E. Agre's Legacy'. *Continent.*, 2014.
- Mason, Harding, D Stewart, and Brendan Gill. 'Rival'. *The New Yorker*, 6 December 1958.

- Massarik, Fred. Review of *Principles of Industrial Psychology*, by Thomas Arthur Ryan and Patricia Cain Smith. *Industrial and Labor Relations Review* 8, no. 4 (July 1955): 617. <https://doi.org/10.2307/2520114>.
- McCarthy, John. 'An Approach to the Artificial Intelligence Problem'. The Dartmouth AI Archives; Box A. Personal Archives of Grace Solomonoff. Accessed 8 January 2018. <http://raysolomonoff.com/dartmouth/boxa/mccaraiapproach.pdf>.
- . 'An Approach to the Artificial Intelligence Problem'. The Dartmouth AI Archives; Box A. Personal Archives of Grace Solomonoff. Accessed 26 January 2017. <http://raysolomonoff.com/dartmouth/boxa/mccaraiapproach.pdf>.
- . 'An Unreasonable Book'. Professor John McCarthy: Father of AI. Accessed 8 November 2016. <http://jmc.stanford.edu/artificial-intelligence/reviews/weizenbaum.pdf>.
- . 'Assorted'. Professor John McCarthy: Father of AI. Accessed 8 November 2016. <http://jmc.stanford.edu/>.
- . Letter to Marvin Minsky. 'Comments on History of LISP Draft', 25 January 1978. Storage - Cabinet 6 Drawer 2 - McCarthy LISP History. Personal Archives of Gloria Rudisch.
- . 'Dartmouth Use of the IBM 704 to Be Located at MIT', n.d. SC0524 ACCN 2013-247 Box 1; Miscellaneous from Notebook. Department of Special Collections and University Archives, Stanford University Libraries.
- . 'Decisions'. Stanford, Calif, n.d. John McCarthy Papers (SC0524); Notes, 1946-1956. Department of Special Collections and University Archives, Stanford University Libraries. <http://www.oac.cdlib.org/view?docId=kt3b69r8ms;developer=local;style=oac4;doc.view=items>.
- . *Defending AI Research: A Collection of Essays and Reviews*. CSLI Lecture Notes, no. 49. Stanford, Calif: CSLI Publications, 1996.
- . 'Dialectical Materialism for the Scientist'. Stanford, Calif, n.d. John McCarthy Papers (SC0524); Notes, 1946-1956. Department of Special Collections and University Archives, Stanford University Libraries. <http://www.oac.cdlib.org/view?docId=kt3b69r8ms;developer=local;style=oac4;doc.view=items>.
- . 'Draft of Towards Intelligent Automata, Introduction'. Stanford, Calif, n.d. John McCarthy Papers (SC0524); Notes, 1946-1956. Department of Special Collections and University Archives, Stanford University Libraries. <http://www.oac.cdlib.org/view?docId=kt3b69r8ms;developer=local;style=oac4;doc.view=items>.
- . 'Getting Closer to Machines That Think'. *Herald Tribune Engineers' News Supplement*, 24 May 1959. Storage - Cabinet 6 Drawer 2 - McCarthy. Personal Archives of Gloria Rudisch.
- . 'History of LISP'. *ACM SIGPLAN Notices* 13, no. 8 (1 August 1978): 217–23. <https://doi.org/10.1145/960118.808387>.
- . 'John McCarthy's Home Page'. Personal blog. Accessed 22 November 2016. <http://www-formal.stanford.edu/jmc/>.
- . 'Letter from John McCarthy to Dr. Robert Morison', 2 September 1955. Record group 1.0002, series 200, box 26, folder 219. Rockefeller Archive Center, Sleepy Hollow, NY.

- . 'Letter to 2nd Lieut. Matthew Kabrisky', 20 September 1956. SC0524 ACCN 2013-247 Box 1; Letters. Department of Special Collections and University Archives, Stanford University Libraries.
- . 'Letter to Dr. Detlev W. Bronk, National Academy of Sciences, National Research Council', 31 August 1956. SC0524 ACCN 2013-247 Box 1; Letters. Department of Special Collections and University Archives, Stanford University Libraries,.
- . 'Letter to Dr. Robert S. Morison, The Rockefeller Foundation', 3 September 1956. SC0524 ACCN 2013-247 Box 1; Letters. Department of Special Collections and University Archives, Stanford University Libraries,.
- . 'Letter to Dr. Saul I Gass, IBM Corporation', 9 September 1957. SC0524 ACCN 2013-247 Box 1; Letters. Department of Special Collections and University Archives, Stanford University Libraries,.
- . 'Letter to Mr. Earl J. Scherago, Science Magazine', 29 April 1957. SC0524 ACCN 2013-247 Box 1; Letters. Department of Special Collections and University Archives, Stanford University Libraries,.
- . 'Levels of Theory in Engineering and Politics'. Stanford, Calif, n.d. John McCarthy Papers (SC0524); Notes, 1946-1956. Department of Special Collections and University Archives, Stanford University Libraries.  
<http://www.oac.cdlib.org/view?docId=kt3b69r8ms;developer=local;style=oac4;doc.view=items>.
- . 'Mailout', 18 November 1956. Record group 1.0002, series 200, box 26, folder 219. Rockefeller Archive Center, Sleepy Hollow, NY.
- . 'Methodology of Work on the Artificial Intelligence Problem'. Stanford, Calif, n.d. John McCarthy Papers (SC0524); Mathematics, 1955-? Department of Special Collections and University Archives, Stanford University Libraries.  
<http://www.oac.cdlib.org/view?docId=kt3b69r8ms;developer=local;style=oac4;doc.view=items>.
- . 'Note to Self'. Stanford, Calif, n.d. John McCarthy Papers (SC0524); Mathematics, 1955-? Department of Special Collections and University Archives, Stanford University Libraries.  
<http://www.oac.cdlib.org/view?docId=kt3b69r8ms;developer=local;style=oac4;doc.view=items>.
- . 'On Paying the Weatherman', n.d. Record group 1.0002, series 200, box 26, folder 219. Rockefeller Archive Center, Sleepy Hollow, NY.
- . 'On the Meaning of "The Reader Will Easily Verify"', n.d. Record group 1.0002, series 200, box 26, folder 219. Rockefeller Archive Center, Sleepy Hollow, NY.
- . Oral History: John McCarthy. Interview by Peter Asaro and Selma Šabanović, 8 June 2011. [http://ethw.org/Oral-History:John\\_McCarthy](http://ethw.org/Oral-History:John_McCarthy).
- . Oral History of John McCarthy. Interview by Nils J. Nilsson, 12 September 2007. <https://www.youtube.com/watch?v=KuU82i3hi8c>.
- . 'Outline for Book 19 Nov 1946'. Stanford, Calif, n.d. John McCarthy Papers (SC0524); Notes, 1946-1956. Department of Special Collections and University Archives, Stanford University Libraries.  
<http://www.oac.cdlib.org/view?docId=kt3b69r8ms;developer=local;style=oac4;doc.view=items>.

- . ‘People Interested in the Artificial Intelligence Problem’, 1956. SC0524 ACCN 1995-247 Box 25; 1956-1959 (articles); Summer Project Memoranda. Department of Special Collections and University Archives, Stanford University Libraries.
- . ‘Physical and Mental Events and Intelligent Machines’. Stanford, Calif, n.d. SC0524 ACCN 2013-247 Box 1; Old Notes. Department of Special Collections and University Archives, Stanford University Libraries,.
- . ‘Programs with Common Sense’. In *Mechanisation of Thought Processes: Proceedings of a Symposium Held at the National Physical Laboratory on 24th, 25th, 26th and 27th November 1958*, edited by National Physical Laboratory (Great Britain), 75–91. Symposium (National Physical Laboratory), v. 1. H.M. Stationery Office, 1959. <https://books.google.co.uk/books?id=rE0IAAAAMAAJ>.
- . ‘Programs with Common Sense’. In *Proceedings of a Symposium Held at the National Physical Laboratory on 24th, 25th, 26th and 27th November 1958*, 1:75–92. London: Her Majesty’s Stationery Office, 1959.
- . ‘Projection Operators and Partial Differential Equations’. University of Princeton, 1951.
- . ‘Recursive Functions Symbolic Expressions and Their Computation by Machine, Part I’. *Communications of the ACM* 3, no. 4 (1 April 1960): 184–95. <https://doi.org/10.1145/367177.367199>.
- . ‘Reminiscences on the Theory of Time-Sharing’. *Professor John McCarthy: Father of AI* (blog), 1983. <http://jmc.stanford.edu/computing-science/timesharing.html>.
- . ‘Report on Dartmouth Summer Research Project’, August 1956.
- . ‘Review by Brian P. Bloomfield (Ed.), The Question of Artificial Intelligence (1987)’, 1987. [www-formal.stanford.edu/jmc/reviews/bloomfield.pdf](http://www-formal.stanford.edu/jmc/reviews/bloomfield.pdf).
- . ‘Review of C. Wright Mills’ The Power Elite’. Stanford, Calif, n.d. John McCarthy Papers (SC0524); Notes, 1946-1956; Department of Special Collections and University Archives, Stanford University Libraries. <http://www.oac.cdlib.org/view?docId=kt3b69r8ms;developer=local;style=oac4;doc.view=items>.
- . ‘Secession’. Stanford, Calif, n.d. John McCarthy Papers (SC0524); Notes, 1946-1956. Department of Special Collections and University Archives, Stanford University Libraries. <http://www.oac.cdlib.org/view?docId=kt3b69r8ms;developer=local;style=oac4;doc.view=items>.
- . ‘Sophisticated Epistemology’. Stanford, Calif, n.d. John McCarthy Papers (SC0524); Mathematics, 1955-? Department of Special Collections and University Archives, Stanford University Libraries. <http://www.oac.cdlib.org/view?docId=kt3b69r8ms;developer=local;style=oac4;doc.view=items>.
- . ‘The Dartmouth Workshop--as Planned and as It Happened’. John McCarthy’s Home Page. Accessed 1 November 2019. <http://www-formal.stanford.edu/jmc/slides/dartmouth/dartmouth/node1.html>.
- . ‘The Inversion of Functions Defined by Turing Machines’. In *Automata Studies*, edited by Claude Elwood Shannon and J. McCarthy, 5th printing., 177–82. Annals of Mathematics Studies 34. Princeton, NJ: Princeton University Press, 1972.
- . ‘Untitled’. Stanford, Calif, n.d. John McCarthy Papers (SC0524); Notes, 1946-1956. Department of Special Collections and University Archives, Stanford University

- Libraries.  
<http://www.oac.cdlib.org/view?docId=kt3b69r8ms;developer=local;style=oac4;doc.view=items>.
- . ‘Untitled Note on Teaching Computing at Dartmouth’. Stanford, Calif, n.d. John McCarthy Papers (SC0524); Notes, 1946-1956. Department of Special Collections and University Archives, Stanford University Libraries.  
<http://www.oac.cdlib.org/view?docId=kt3b69r8ms;developer=local;style=oac4;doc.view=items>.
- . Letter to Marvin Minsky, December 1953. Minsky 2AR. Personal Archives of Gloria Rudisch.
- . Letter to Claude E. Shannon, 11 August 1954. Box 1. Library of Congress, Claude Elwood Shannon papers. <https://lccn.loc.gov/mm2001084831>.
- McCarthy, John, Paul W. Abrahams, Daniel J. Edwards, Timothy P. Hart, and Michael I. Levin. *LISP 1.5 Programmer’s Manual*. Second. Massachusetts Institute of Technology, 1985.
- McCarthy, John, and Patrick Hayes. ‘Some Philosophical Problems from the Standpoint of Artificial Intelligence’. In *Machine Intelligence 4*, edited by B. Meltzer and Donald Michie. Edinburgh University Press, 1969.
- McCarthy, John, and Vladimir Lifschitz. *Formalizing Common Sense: Papers by John McCarthy*. Exeter: Intellect, 1998.
- McCarthy, John, Marvin Minsky, Nathaniel Rochester, and Claude E. Shannon. ‘A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence’. John McCarthy’s Home Page, 31 August 1955. <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>.
- McCorduck, Pamela. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. San Francisco: W. H. Freeman, 1979.
- . *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. 25th anniversary update. Natick, Mass: A.K. Peters, 2004.
- McCulloch, W. S., and Walter Pitts. ‘A Logical Calculus of the Ideas Immanent in Nervous Activity’. *Bulletin of Mathematical Biophysics* 5 (1943): 115–33.
- McCumber, John. *The Philosophy Scare: The Politics of Reason in the Early Cold War*. Chicago: The University of Chicago Press, 2016.
- McDermott, Drew. ‘Artificial Intelligence Meets Natural Stupidity’. *ACM SIGART Bulletin*, no. 57 (1 April 1976): 4–9. <https://doi.org/10.1145/1045339.1045340>.
- McLaughlin, Francis M. ‘John Rogers Commons: Are His Insights Important in Teaching Modern Labor Economics?’ Presented at the 37th Annual Lonergan Workshop, Boston College, Boston, 20 June 2010. <http://fmwww.bc.edu/EC-P/wp739.pdf>.
- Mead, Margaret. *And Keep Your Powder Dry: An Anthropologist Looks at America*. Margaret Mead—Researching Western Contemporary Cultures, v. 2. New York: Berghahn Books, 2000.
- Meyr, Jessica. ‘Grace Hopper and the Marvelous Machine: Lessons for Modern Technical Communicators from the Mark I ASCC Manual’, 2017.  
<https://stars.library.ucf.edu/etd/5396>.
- Mignolo, W.D., and A. Escobar. *Globalization and the Decolonial Option*. Taylor & Francis, 2013. <https://books.google.co.uk/books?id=VbmMAQAAQBAJ>.
- Miles, H.W. ‘Operations Research: An ASTIA Report Bibliography’. ASTIA Comprehensive Bibliographies. Arlington, Virginia: Defense Documentation Center, January 1962.

- Miller, George A. 'The Cognitive Revolution: A Historical Perspective'. *Trends in Cognitive Sciences* 7, no. 3 (March 2003): 141–44. [https://doi.org/10.1016/S1364-6613\(03\)00029-9](https://doi.org/10.1016/S1364-6613(03)00029-9).
- Miller, George A. 'The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information'. *Psychological Review* 63, no. 2 (1956): 81–97. <https://doi.org/10.1037/h0043158>.
- Milner, P. M. 'The Cell Assembly: Mark II'. *Psychological Review* 64 (1957): 242–52.
- Miltner, Kate M. 'Girls Who Coded: Gender in Twentieth Century U.K. and U.S. Computing'. *Science, Technology, & Human Values*, 7 May 2018. <https://doi.org/10.1177/0162243918770287>.
- Mindell, David A. *Between Human and Machine: Feedback, Control, and Computing Before Cybernetics*. Johns Hopkins Studies in the History of Technology. Baltimore: The Johns Hopkins University Press, 2002.
- . *Our Robots, Ourselves: Robotics and the Myths of Autonomy*. New York, New York: Viking, 2015.
- Minsky, Marvin. 'A Combined Publication, Retrieval, and Research Facility', 16 August 1957. Minsky 3AR. Personal Archives of Gloria Rudisch.
- . 'A Framework for Artificial Intelligence', 4 July 1956. Box A. Paper of Ray Solomonoff. <http://raysolomonoff.com/dartmouth/boxa/dart56minskyai.pdf>.
- . 'A Framework for Artificial Intelligence', 4 July 1956. SC0524 ACCN 1995-247 Box 25; 1956-1959 (articles); Summer Project Memoranda. Department of Special Collections and University Archives, Stanford University Libraries.
- . 'A Selected Descriptor-Indexed Bibliography to the Literature on Artificial Intelligence'. *IRE Transactions on Human Factors in Electronics* HFE-2, no. 1 (March 1961): 39–55. <https://doi.org/10.1109/THFE2.1961.4503297>.
- . An Interview with Marvin L. Minsky. Interview by Arthur L. Norberg, 1 November 1989. <https://conservancy.umn.edu/bitstream/handle/11299/107503/oh179mlm.pdf?sequence=1&isAllowed=y>.
- . 'Discrete Selection Processes'. Navy Contract Number Nonr-494(03). Medford, MA: Tufts College, July 1954.
- . 'Exploration Systems and Synaptic Procedures', 1956. SC0524 ACCN 1995-247 Box 25; 1956-1959 (articles); Summer Project Memoranda. Department of Special Collections and University Archives, Stanford University Libraries.
- . 'Heuristic Aspects of the Artificial Intelligence Problem'. Boston, MA: Lincoln Laboratory, Massachusetts Institute of Technology, 17 December 1956.
- . 'Learning Systems and Artificial Intelligence', 20 July 1957. Minsky 3AR. Personal Archives of Gloria Rudisch.
- . 'Notes for an Unpublished Autobiography, #Biogmm-A-G'. Brookline, MA, n.d. Personal Archives of Gloria Rudisch.
- . 'Notes for an Unpublished Autobiography, #Biogmm-H-M'. Brookline, MA, n.d. Personal Archives of Gloria Rudisch.
- . 'Notes for an Unpublished Autobiography, #Biogmm-N-Z'. Brookline, MA, n.d. Personal Archives of Gloria Rudisch.
- . 'Notes on the Geometry Problem 1', 16 August 1956. Box H01-H56. The Herbert Stoyan collection on LISP programming Collection (Number: X5687.2010), Computer History Museum.



- . ‘Re: Course Material’, Estimated 1957. Minsky 2AR. Personal Archives of Gloria Rudisch.
- . ‘Re: Int. Conf. on Scientific Information’, 16 August 1957. Minsky 3AR. Personal Archives of Gloria Rudisch.
- . ‘Some Methods of Artificial Intelligence and Heuristic Programming’. In *Proceedings of a Symposium Held at the National Physical Laboratory on 24th, 25th, 26th and 27th November 1958*, 1:3–27. London: Her Majesty’s Stationery Office, 1959.
- . ‘Some Universal Elements for Finite Automata’. In *Automata Studies*, edited by Claude Elwood Shannon and J. McCarthy, 5th printing., 117–28. Annals of Mathematics Studies 34. Princeton, NJ: Princeton University Press, 1972.
- . ‘Steps Toward Artificial Intelligence’. *Proceedings of the IRE* 49, no. 1 (January 1961): 8–30. <https://doi.org/10.1109/JRPROC.1961.287775>.
- . ‘The Discrete Approach: Neural Net and Related Theories’, 26 April 1950. Minsky 2AR. Personal Archives of Gloria Rudisch.
- . *The Emotion Machine: Commonsense Thinking, Artificial Intelligence and the Future of the Human Mind*. New York: Simon & Schuster, 2007.
- . ‘Theory of Neural-Analog Reinforcement Systems and Its Application to the Brain-Model Problem’. Princeton University, 1954. Photocopy. Ann Arbor, Michigan, University Microfilms, 1962.
- . ‘To Drs. J McCarthy and Nat Rochester, IBM, Poughkeepsie, N.Y.’, n.d. Personal Archives of Gloria Rudisch.
- . Letter to S. C. Kleene, Undated. Minsky 2AR. Personal Archives of Gloria Rudisch.
- . Letter to W. Grey Walter, 22 August 1954. Minsky 2AR. Personal Archives of Gloria Rudisch.
- Minsky, Marvin, John McCarthy, Claude E. Shannon, and Nathaniel Rochester. ‘A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence’, 31 August 1955. Dart564props. Online Papers of Ray Solomonoff. <http://raysolomonoff.com/dartmouth/boxa/dart564props.pdf>.
- Minsky, Marvin, and Seymour Papert. *Perceptrons: An Introduction to Computational Geometry*. Expanded ed. Cambridge, Mass: MIT Press, 1988.
- Minsky, Marvin, and Oliver Selfridge. ‘Learning in Random Nets’. London: Butterworths, 1960.
- Minsky, Marvin, Cynthia Solomon, and Xiao Xiao. *Inventive Minds: Marvin Minsky on Education*. Cambridge, MA: MIT Press, 2019.
- Mirowski, Philip. ‘Book Review: McCorduck’s Machines Who Think after Twenty-Five Years - Revisiting the Origins of AI’. *AI Magazine*, 2003.
- . *Machine Dreams: Economics Becomes a Cyborg Science*. Cambridge ; New York: Cambridge University Press, 2002.
- Mood, A. M. Letter to Marvin Minsky. ‘Invitation to Join the Rand Corporation’, 26 February 1954. Minsky 2AR. Personal Archives of Gloria Rudisch.
- Morison, Robert. ‘Letter from Robert Morison to John McCarthy’, 30 November 1955. Record group 1.0002, series 200, box 26, folder 219. Rockefeller Archive Center, Sleepy Hollow, NY.
- . ‘Memo Re: Oliver Selfridge, Wednesday, May 8, 1957’, 18 November 1956. Record group 1.0002, series 200, box 26, folder 219. Rockefeller Archive Center, Sleepy Hollow, NY.

- . 'New York', 17 June 1955. Record group 1.0002, series 200, box 26, folder 219. Rockefeller Archive Center, Sleepy Hollow, NY.
- . 'New York', 17 June 1955. Dartmouth Conference File. Rockefeller Foundation.
- Murray, A. E. 'Information Processing Relevant to Military Command: Survey, Recommendations and Bibliography'. Volumes 1-2 of Information Processing Relevant to Military Command: Survey, Recommendations and Bibliography. Buffalo, N.Y.: System Design Laboratory, Electronic Systems Division, Air Force Systems Command, U. S. Air Force, 1963.
- Nagy, George. 'Frank Rosenblatt (1928-1971) – List of Publications (Partial and Unverified)', 30 March 2011.
- National Research Council (U.S.), ed. *Naval Engineering in the 21st Century: The Science and Technology Foundation for Future Naval Fleets*. Transportation Research Board Special Report 306. Washington, D.C: Transportation Research Board, 2011.
- 'New Navy Device Learns by Doing'. *The New York Times*, 8 July 1958, sec. 25:2.
- Newell, A., and J. C. Shaw. 'Programming the Logic Theory Machine', 230–40. ACM Press, 1957. <https://doi.org/10.1145/1455567.1455606>.
- Newell, A., J. C. Shaw, and H. A. Simon. 'Empirical Explorations of the Logic Theory Machine: A Case Study in Heuristic', 218–30. ACM Press, 1957. <https://doi.org/10.1145/1455567.1455605>.
- Newell, A., and H. Simon. 'The Logic Theory Machine: A Complex Information Processing System'. *IEEE Transactions on Information Theory* 2, no. 3 (September 1956): 61–79. <https://doi.org/10.1109/TIT.1956.1056797>.
- Newell, Allen. 'Comments on Ad Hoc PSAC Panel on Non-Numerical Information Processing', 15 December 1961. Unlisted - Newell Letter 1961. Personal Archives of Gloria Rudisch.
- . 'Intellectual Issues in the History of Artificial Intelligence'. In *Artificial Intelligence: Critical Concepts*, edited by Ronald Chrisley and Sander Begeer. London ; New York: Routledge, 2000.
- . 'Notes on Chess Learning Machine I'. Santa Monica, CA, 16 January 1955. <http://doi.library.cmu.edu/10.1184/pmc/newell/box00002/fld00108/bdl0001/doc0001>.
- . 'Talk on "Heuristic Programs"', 19 October 1959. Box POD storage Marvin 1952\_1959. Personal Archives of Gloria Rudisch.
- . 'Tentative Outline of Michigan Summer Session Lectures', 18 May 1957. Unlisted - Letters 1960; 1960 Letters 10. Personal Archives of Gloria Rudisch.
- . 'The Chess Machine: An Example of Dealing with a Complex Task by Adaptation', 101–8. ACM Press, 1955. <https://doi.org/10.1145/1455292.1455312>.
- . 'The Heuristic of George Pólya and Its Relation to Artificial Intelligence'. In *Methods of Heuristics*, edited by M Groner, R Groner, and W. F. Bischoof, 195–244. New Jersey: Lawrence Erlbaum Associates, Inc. Publishers, 1983.
- Newell, Allen, J. C. Shaw, and Herbert Simon. 'The Processing of Creative Thinking'. 1958 Symposium on Creative Thinking at the University of Colorado: RAND Corporation, 16 September 1958. Storage - Minsky 6 AR. Personal Archives of Gloria Rudisch.
- Newell, Allen, J. C. Shaw, and Herbert A. Simon. 'Elements of a Theory of Human Problem Solving.' *Psychological Review* 65, no. 3 (1958): 151–66. <https://doi.org/10.1037/h0048495>.

- Newell, Allen, and Herbert Simon. *Human Problem Solving*. Upper Saddle River, NJ: Prentice Hall, 1972.
- Newell, Allen, and Herbert A. Simon. 'Abstract: The Logic Theory Machine'. Series IV: RAND Corporation -- 1949-1973 (1991, 1994). Carnegie Mellon University, 15 May 1956. <http://doi.library.cmu.edu/10.1184/pmc/simon/box00006/fld00404/bdl0001/doc0001>.
- . 'The Logic Theory Machine: A Complex Information Processing System'. P-868. Santa Monica, CA: RAND Corporation, 15 June 1956.
- 'News and Notices'. *SIAM Review* 3, no. 3 (July 1961): 270–92. <https://doi.org/10.1137/1003058>.
- Nilsson, Nils J. 'John McCarthy, 1927-2011'. National Academy of Sciences, 2012. [http://ai.stanford.edu/~nilsson/John\\_McCarthy.pdf](http://ai.stanford.edu/~nilsson/John_McCarthy.pdf).
- . *The Quest for Artificial Intelligence: A History of Ideas and Achievements*. Cambridge ; New York: Cambridge University Press, 2010.
- Nixon, Rob. *Slow Violence and the Environmentalism of the Poor*. 1. Harvard Univ. Press paperback ed. Cambridge, Mass.: Harvard Univ. Press, 2013.
- Noble, David F. *Progress without People: New Technology, Unemployment, and the Message of Resistance*. Toronto, Ont: Between the Lines, 1995.
- Nofre, David, Mark Priestley, and Gerard Alberts. 'When Technology Became Language: The Origins of the Linguistic Conception of Computer Programming, 1950–1960'. *Technology and Culture* 55, no. 1 (2014): 40–75. <https://doi.org/10.1353/tech.2014.0031>.
- Norberg, Arthur L., Judy E. O'Neill, and Kerry J. Freedman. *Transforming Computer Technology: Information Processing for the Pentagon, 1962-1986*. Johns Hopkins Studies in the History of Technology, new ser., no. 18. Baltimore: Johns Hopkins University Press, 1996.
- Ochigame, Rodrigo. 'Paraconsistent Society: Mathematical Logic and Social Paradox'. Workshop presented at the Towards a History of Artificial Intelligence, Columbia University, New York, 24 May 2019.
- O'Dowd, Niall. 'This Son of an Irish Emigrant Invented Artificial Intelligence and Changed the World'. *Irish Central*, 17 January 2020.
- Oettinger, Anthony G. 'Programming a Digital Computer to Learn'. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 43, no. 347 (December 1952): 1243–63. <https://doi.org/10.1080/14786441208520256>.
- Officer, Lawrence H., and Samuel H. Williamson. 'Computing "Real Value" Over Time With a Conversion Between U.K. Pounds and U.S. Dollars, 1774 to Present'. *MeasuringWorth* (blog), 2015. [www.measuringworth.com/exchange/](http://www.measuringworth.com/exchange/).
- . 'Computing "Real Value" Over Time With a Conversion Between U.K. Pounds and U.S. Dollars, 1774 to Present'. *MeasuringWorth* (blog), 2020. [www.measuringworth.com/exchange/](http://www.measuringworth.com/exchange/).
- Olazaran, M. 'A Historical Sociology of Neural Network Research'. University of Edinburgh, 1991. <http://hdl.handle.net/1842/20075>.
- . 'A Sociological Study of the Official History of the Perceptrons Controversy'. *Social Studies of Science* 26, no. 3 (1 August 1996): 611–59. <https://doi.org/10.1177/030631296026003005>.
- Old, Bruce S. 'The Evolution of the Office of Naval Research'. *Physics Today* 14, no. 8 (August 1961): 30–35. <https://doi.org/10.1063/1.3057690>.

- O'Neil, Cathy. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. First edition. New York: Crown, 2016.
- Orr, Julian E. *Talking about Machines*. Cornell University Press, 1996.  
[www.jstor.org/stable/10.7591/j.ctt1hhfnkz](http://www.jstor.org/stable/10.7591/j.ctt1hhfnkz).
- Paidipaty, Poornima. "'Tortoises All the Way down': Geertz, Cybernetics and 'Culture' at the End of the Cold War". *Anthropological Theory* 20, no. 1 (March 2020): 97–129.  
<https://doi.org/10.1177/1463499619899747>.
- Parsons, Talcott. *The Structure of Social Action*. New York, NY: McGraw-Hill, 1937.
- Pavlov, Ivan Petrovich, and G. V. Anrep. *Conditioned Reflexes*. Mineola, N.Y: Dover Publications, 2003.
- Penn, Jonnie. 'AI Thinks like a Corporation—and That's Worrying'. *The Economist*, 26 November 2018. <https://economist.com/open-future/2018/11/26/ai-thinks-like-a-corporation-and-thats-worrying>.
- . 'On the Logic of Labour: EDSAC and the History of Open-Source Software (OSS)', Forthcoming.
- Phillips, Christopher J. *The New Math: A Political History*. Chicago: The University of Chicago Press, 2015.
- Piccinini, Gualtiero. 'Computations and Computers in the Sciences of Mind and Brain'. PhD Dissertation, Università di Torino, 2003.
- . 'The First Computational Theory of Mind and Brain: A Close Look at McCulloch and Pitts's "Logical Calculus of Ideas Immanent in Nervous Activity"'. *Synthese* 141, no. 2 (August 2004): 175–215. <https://doi.org/10.1023/B:SYNT.0000043018.52445.3e>.
- Pickering, Andrew. *The Cybernetic Brain: Sketches of Another Future*. Paperback Ed. Chicago, Ill: University of Chicago Press, 2011.
- Pierce, Alice M. 'A Concise Bibliography of the Literature on Artificial Intelligence'. Bedford, MA: Electronics Research Directorate, Air Force Cambridge Research Center, Air Research and Development Command, September 1959.
- . 'Bibliography on Artificial Intelligence: First Draft'. Bedford, MA: Communication Sciences Laboratory, Air Force Cambridge Research Center, 5 May 1959.
- Pitts, Walter, and Warren S. McCulloch. 'How We Know Universals the Perception of Auditory and Visual Forms'. *The Bulletin of Mathematical Biophysics* 9, no. 3 (1 September 1947): 127–47. <https://doi.org/10.1007/BF02478291>.
- Plasek, Aaron. 'On the Cruelty of Really Writing a History of Machine Learning'. *IEEE Annals of the History of Computing* 38, no. 4 (October 2016): 6–8.  
<https://doi.org/10.1109/MAHC.2016.43>.
- Pollan, Michael. 'The Intelligent Plant'. *The New Yorker*, 23 December 2013.
- Pólya, George, and John Horton Conway. *How to Solve It*. Princeton, NJ: Princeton University Press, 1945.
- Porter, Theodore M. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. 2. print., and 1. paperback printing. History and Philosophy of Science. Princeton, N.J: Princeton Univ.Press, 1996.
- Poznanski, Roman R, Kenneth A Lindsay, Jay R Rosenberg, Olaf Sporns, and George R Reeke. *Modeling in the Neurosciences, Second Edition*. London: CRC Press, 2004.
- Priddy, Jason. 'Kathleen Booth', 6 May 2019.  
<https://medium.com/@BatmanPriddy/kathleen-booth-7bb303fb15be>.
- Priestley, Mark. *A Science of Operations*. History of Computing. London: Springer London, 2011. <https://doi.org/10.1007/978-1-84882-555-0>.

- . ‘AI and the Origins of the Functional Programming Language Style’. *Minds and Machines* 27, no. 3 (September 2017): 449–72. <https://doi.org/10.1007/s11023-017-9432-7>.
- ‘Prof. Rosenblatt Dies’. *The Cornell Daily Sun*, 9 January 1971. Cornell Faculty Bio files, Box 1 - Frank Rosenblatt Publications 1956-58. Division of Rare and Manuscript Collections, Cornell University Library.
- Pursell, Carroll W. ‘A Preface to Government Support of Research and Development: Research Legislation and the National Bureau of Standards, 1935-41’. *Technology and Culture* 9, no. 2 (April 1968): 145. <https://doi.org/10.2307/3102148>.
- Radin, Joanna. “‘Digital Natives’: How Medical and Indigenous Histories Matter for Big Data”. *Osiris* 32, no. 1 (1 September 2017): 43–64. <https://doi.org/10.1086/693853>.
- Rankin, Joy Lisi. *A People’s History of Computing in the United States*. Cambridge, Massachusetts: Harvard University Press, 2018.
- Rashevsky, N. *Mathematical Biophysics: Physicomathematical Foundations of Biology*. Chicago: University of Chicago, 1938.
- Rees, Mina. ‘Digital Computers - Their Nature and Use’. *American Scientist* 40, no. 2 (1952): 328–35.
- . ‘Mathematics and Federal Support’. *Science* 119, no. 3099 (1954): 3A-3A.
- . ‘The Federal Computing Machine Program’. *Science* 112, no. 2921 (1950): 731–36.
- . ‘The Impact of the Computer’. *The Mathematics Teacher* 51, no. 3 (1958): 162–68.
- Reisch, George A. *How the Cold War Transformed Philosophy of Science: To the Icy Slopes of Logic*. Cambridge ; New York: Cambridge University Press, 2005.
- Rheingold, Howard. *Tools for Thought: The History and Future of Mind-Expanding Technology*. 1st MIT Press ed. Cambridge, Mass: MIT Press, 2000.
- Rid, Thomas. *Rise of the Machines: A Cybernetic History*. First edition. New York: W. W. Norton & Company, 2016.
- Roberts, David. ‘Mathematics and Pedagogy: Professional Mathematicians and American Educational Reform, 1893-1923’. PhD Dissertation, 1997.
- Rochester, N., J. Holland, L. Haibt, and W. Duda. ‘Tests on a Cell Assembly Theory of the Action of the Brain, Using a Large Digital Computer’. *IEEE Transactions on Information Theory* 2, no. 3 (September 1956): 80–93. <https://doi.org/10.1109/TIT.1956.1056810>.
- Rochester, Nathaniel. ‘Biographical Data, Nov 3, 1957’, n.d. Nathaniel Rochester Papers, Box 1, Folder 6, Biographical File 1956-60. Library of Congress, Washington D.C. Accessed 1 March 2017.
- . ‘Biographical One-Sheet’, n.d. Nathaniel Rochester Papers, Box 1, Folder 7, Biographical File 1961-68. Library of Congress, Washington D.C. Accessed 1 March 2017.
- Rogan, Tim. ‘Know-How: Market Systems Have Made Better Use of More Information than Economic Planners. What If AI and Machine Learning Changed That?’ *Aeon*, 25 March 2019. <https://aeon.co/essays/big-data-ai-and-the-peculiar-dignity-of-tacit-knowledge>.
- Rohde, Joy. *Armed with Expertise: The Militarization of American Social Research during the Cold War*. American Institutions and Society. Ithaca: Cornell University Press, 2013.
- . ‘Pax Technologica: Computers, International Affairs, and Human Reason in the Cold War’. *Isis* 108, no. 4 (2 December 2017): 792–813. <https://doi.org/10.1086/695679>.

- Rojas, Raúl, and Ulf Hashagen, eds. *The First Computers: History and Architectures*. History of Computing. Cambridge, Mass: MIT Press, 2000.
- Rolfe, John M. 'Two Cambridge Inventors'. Royal Aeronautical Society. *Flight Simulation Group* (blog), n.d. [http://www.raes-fsg.org.uk/18/the\\_cambridge\\_cockpit/](http://www.raes-fsg.org.uk/18/the_cambridge_cockpit/).
- Rose, Frank. *Into the Heart of the Mind*. Century, 1985.
- Rosenblatt, F., and R. G. Miller. 'Behavioral Assay Procedures for Transfer of Learned Behavior by Brain Extracts.' *Proceedings of the National Academy of Sciences* 56, no. 5 (1 November 1966): 1423–30. <https://doi.org/10.1073/pnas.56.5.1423>.
- Rosenblatt, Frank. 'A Probabilistic Model for Visual Perception'. *Acta Psychologica* 15 (1959): 296–97. [https://doi.org/10.1016/S0001-6918\(59\)80143-8](https://doi.org/10.1016/S0001-6918(59)80143-8).
- . 'Analytical Techniques for the Study of Neural Nets'. *IEEE Trans. Applications and Industry* 83, no. 74 (September 1964): 285–92.
- . Interview with Frank Rosenblatt. Interview by Micheal Wright, 15 July 1969.
- . *Principles of Neurodynamics*. Washington, DC: Spartan Books, 1962.
- . 'Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms'. Buffalo, N.Y.: Cornell Aeronautical Laboratory, Inc., 15 March 1961.
- . 'The Design of an Intelligent Automaton'. *Research Trends, Cornell Aeronautical Laboratory, Inc., Buffalo 21, New York*, Summer 1958. Personal Archives of George Nagy.
- . 'The K-Coefficient: Design and Train Application of a New Technique for Multivariate Analysis'. 1956. Cornell University Faculty Biographical files, #47-10-3394. Division of Rare and Manuscript Collections, Cornell University Library.
- . 'The Perceptron: A Perceiving and Recognizing Automaton (Project PARA)'. Buffalo, N.Y.: Cornell Aeronautical Laboratory, Inc., January 1957. <https://blogs.umass.edu/brain-wars/files/2016/03/rosenblatt-1957.pdf>.
- . 'The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain.' *Psychological Review* 65, no. 6 (1958): 386–408. <https://doi.org/10.1037/h0042519>.
- . 'The Perceptron: A Theory of Statistical Separability in Cognitive Systems (Project PARA)'. Buffalo, N.Y.: Cornell Aeronautical Laboratory, Inc., January 1958.
- . 'To Publisher, Chicago Tribune', 21 June 1963. Cornell Faculty Bio files, Box 1 - Frank Rosenblatt Publications 1956-58. Division of Rare and Manuscript Collections, Cornell University Library.
- . 'Two Theorems of Statistical Separability in the Perceptron'. *Proceedings of a Symposium Held at the National Physical Laboratory, November 1958* 1, no. Mechanisation of Thought Processes (1959): 421–56.
- . 'Two Theorems of Statistical Separability in the Perceptron'. In *Proceedings of a Symposium Held at the National Physical Laboratory on 24th, 25th, 26th and 27th November 1958*, 1:419–72. London: Her Majesty's Stationery Office, 1959.
- Rosenblatt, Frank, John T. Farrow, and William F. Herblin. 'Transfer of Conditioned Responses from Trained Rats to Untrained Rats by Means of a Brain Extract'. *Nature* 209, no. 5018 (January 1966): 46–48. <https://doi.org/10.1038/209046a0>.
- Rosenblatt, Frank, and Wolfgang Fuch. 'Proposal to the National Science Foundation for Joint Support of the Cognitive Systems Research Program, 1962-1967', n.d. Personal Archives of George Nagy.
- Rosenblueth, Arturo, Norbert Wiener, and Julian Bigelow. 'Behavior, Purpose and Teleology'. *Philosophy of Science* 10, no. 1 (1943): 18–24.

- Ross, Dorothy, and Theodore M. Porter, eds. *The Cambridge History of Science, Volume 7: The Modern Social Sciences*. Cambridge ; New York: Cambridge University Press, 2003.
- Runciman, David. 'Diary'. *London Review of Books*, 25 January 2018.
- Russell, Bertrand, and Alfred North Whitehead. *Principia Mathematica Volume I*. Second Edition. London: Cambridge University Press, 1963.  
<https://archive.org/details/PrincipiaMathematicaVolumeI>.
- Russell, Stuart J., Peter Norvig, and Ernest Davis. *Artificial Intelligence: A Modern Approach*. 3rd ed. Prentice Hall Series in Artificial Intelligence. Upper Saddle River: Prentice Hall, 2010.
- Sabanovic, Selma, Stasa Milojevic, and Jasleen Kaur. 'John McCarthy [History]'. *IEEE Robotics & Automation Magazine* 19, no. 4 (December 2012): 99–106.  
<https://doi.org/10.1109/MRA.2012.2221259>.
- Scates, Shelby. *Maurice Rosenblatt and the Fall of Joseph McCarthy*. Seattle: University of Washington Press in association with History Ink, 2006.
- Schaffer, Simon. 'Babbage's Intelligence: Calculating Engines and the Factory System'. *Critical Inquiry* 21, no. 1 (1994): 203–27.
- . 'OK Computer'. In *Ecce Cortex: Beiträge Zur Geschichte Des Modernen Gehirns*, edited by Michael Hagner, 254–85, 1999.
- Schmidhuber, Jürgen. '2006: Celebrating 75 Years of AI - History and Outlook: The Next 25 Years'. In *50 Years of Artificial Intelligence: Essays Dedicated to the 50th Anniversary of Artificial Intelligence*, edited by Max Lungarella, 36–48. Lecture Notes in Computer Science Lecture Notes in Artificial Intelligence 4850. Berlin: Springer, 2007.
- Schrödinger, Erwin. *What Is Life? The Physical Aspect of the Living Cell*. Cambridge: Cambridge University Press, 1944.
- Schwartz, L. *Theorie des Distributions*. Paris: Hermann, 1978.
- Scoblic, J. Peter. 'The Postwar Development of Tools to Mitigate Uncertainty: Frank Knight, the RAND Corporation, and the Pursuit of "Ersatz Experience"'. Harvard Business School, 16 March 2019.
- Scott, James C. 'Authoritarian High Modernism'. In *Seeing like a State: How Certain Schemes to Improve the Human Condition Have Failed*, Nachdr., 87–102. Yale Agrarian Studies. New Haven, Conn.: Yale Univ. Press, 2008.
- . *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. New Haven, Conn.: Yale Univ. Press, 2008.
- Seising, Rudolf. 'A Brain Model for the Perception of the Outside World'. Workshop presented at the Towards a History of Artificial Intelligence, Columbia University, New York, 24 May 2019.
- . 'Marvin Lee Minsky (1927–2016)'. *Artificial Intelligence in Medicine* 75 (January 2017): 24–31. <https://doi.org/10.1016/j.artmed.2016.12.001>.
- Selfridge, O. G. 'Quarterly Progress Report: Pattern Recognition - Group 54 (December-March)'. Boston: Massachusetts Institute of Technology, 1960 1959. Minsky 3AR. Personal Archives of Gloria Rudisch.
- Selfridge, Oliver. 'Pandemonium: A Paradigm for Learning'. In *Mechanisation of Thought Processes: Proceedings of a Symposium Held at the National Physical Laboratory on 24th, 25th, 26th and 27th November 1958*, edited by National Physical Laboratory (Great Britain), 511–31. Symposium (National Physical Laboratory), v. 1. Her Majesty's Stationary Office, 1959.

- . 'Pattern Recognition and Modern Computers'. In *Proceedings of the 1955 Western Joint Computer Conference*, 91–93. New York: Institute of Radio Engineers, 1955.  
<http://portal.acm.org/citation.cfm?doid=1455292.1455310>.
- Sent, Esther-Mirjam. 'Book Reviews; A Review of Herbert A. Simon's Administrative Behavior'. *Journal of Economic Methodology* 6, no. 1 (March 1999): 133–39.  
<https://doi.org/10.1080/13501789900000008>.
- Shannon, Claude E. 'Prediction and Entropy of Printed English'. *Bell System Technical Journal* 30, no. 1 (January 1951): 50–64. <https://doi.org/10.1002/j.1538-7305.1951.tb01366.x>.
- . 'Computers and Automata'. *Proceedings of the Institute of Radio Engineers* 41, no. 10 (October 1953): 1234–41.
- . 'Letter from Dr. Claude Shannon to Dr. Herbert A. Simon', 24 April 1956. Record group 1.0002, series 200, box 26, folder 219. Rockefeller Archive Center, Sleepy Hollow, NY.
- . 'Programming a Computer for Playing Chess'. In *Computer Chess Compendium*, edited by David Levy, 2–13. New York, NY: Springer New York, 1988.  
[https://doi.org/10.1007/978-1-4757-1968-0\\_1](https://doi.org/10.1007/978-1-4757-1968-0_1).
- . 'The Bandwagon'. *Information Theory, IRE Transactions* 2, no. 1 (March 1956): 3.
- . 'XXII. Programming a Computer for Playing Chess'. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 41, no. 314 (March 1950): 256–75. <https://doi.org/10.1080/14786445008521796>.
- Shannon, Claude E., and John McCarthy. Letter to Marvin Minsky. 'Invitation to Marvin Minsky to Contribute to Automata Studies', 18 May 1953. Den 2A - Yellow Folder - Unlabelled Materials from Steps; Yellow Folder 19. Personal Archives of Gloria Rudisch.
- Shannon, Claude Elwood, and J. McCarthy, eds. *Automata Studies*. Annals of Mathematics Studies 34. Princeton, NJ: Princeton University Press, 1956.
- Shapin, Steven, and Simon Schaffer. *Leviathan and the Air-Pump: Hobbes, Boyle, and the Experimental Life*. Princeton, N.J: Princeton University Press, 2011.
- Shapiro, Stewart, and Teresa Kouri Kissel. 'Classical Logic'. In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University, Spring 2018. <https://plato.stanford.edu/entries/logic-ancient/>.
- Shell-Gellasch, Amy. 'Mina Rees and the Funding of the Mathematical Sciences'. *The American Mathematical Monthly* 109, no. 10 (December 2002): 873.  
<https://doi.org/10.2307/3072455>.
- Sheppard, Mary. Letter to Marvin Minsky. 'Re: International Conference on Scientific Information', 20 March 1958. Minsky 3AR. Personal Archives of Gloria Rudisch.
- Simon, H. *Models of Thought*. New Haven, Ct: Yale Univ. Pr, 1979.
- Simon, Herbert A. 'A Behavioral Model of Rational Choice'. *The Quarterly Journal of Economics* 69, no. 1 (1955): 99–118. <https://doi.org/10.2307/1884852>.
- . 'Administrative Behavior'. *AJN The American Journal of Nursing* 50, no. 2 (1950).  
[https://journals.lww.com/ajnonline/Fulltext/1950/02000/Administrative\\_Behavior.71.aspx](https://journals.lww.com/ajnonline/Fulltext/1950/02000/Administrative_Behavior.71.aspx).
- . *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organizations*. 4th ed. New York: Free Press, 1997.
- . 'Decision-Making and Administrative Organization'. *Public Administration Review* 4, no. 1 (1944): 16–30. <https://doi.org/10.2307/972435>.



- . *Models of My Life*. 1st MIT Press ed. Cambridge, Mass: MIT Press, 1996.
- . 'The Architecture of Complexity'. *Proceedings of the American Philosophical Society* 106, no. 6 (n.d.): 467–82.
- . 'The Proverbs of Administration'. *Public Administration Review* 6, no. 1 (1946): 53. <https://doi.org/10.2307/973030>.
- . Transcript of Interview of Herbert Simon by Pamela McCorduck, January 01, 1975. Interview by Pamela McCorduck, 1 January 1975. Series VIII: Publications -- (1938) 1949-2000; Document 38145. Carnegie Mellon University Libraries Digital Collections. <http://doi.library.cmu.edu/10.1184/pmc/simon//box00079/fld06444/bdl0001/doc0002>.
- . Letter to Marvin Minsky, 27 September 1956. Minsky 2AR. Personal Archives of Gloria Rudisch.
- Simon, Herbert A., and Harold Guetzkow. 'Memorandum to G. L. Bach, Dean of Graduate School of Industrial Administration; February 28, 1952'. Pittsburgh, PA, 28 February 1952. <http://doi.library.cmu.edu/10.1184/pmc/simon/box00016/fld01138/bdl0002/doc0001>.
- Simon, Herbert A., and Allen Newell. 'Current Developments in Complex Information Processing'. The RAND Corporation, 1 May 1956. Series IV: RAND Corporation -- 1949-1973 (1991, 1994). Carnegie Mellon University Libraries Digital Collections. <http://doi.library.cmu.edu/10.1184/pmc/simon/box00006/fld00396/bdl0001/doc0001>.
- . 'Heuristic Problem Solving: The Next Advance in Operations Research'. *Operations Research* 6, no. 1 (1958): 1–10.
- Simon, Herbert A., Allen Newell, and J. C. Shaw. 'Untitled Memo', n.d. Study -- N.D. (RAND -- Correspondence -- 1949-1957). <http://doi.library.cmu.edu/10.1184/pmc/simon/box00005/fld00367/bdl0039/doc0001B>.
- Simon, Herbert, and Allen Newell. 'Plans for the Dartmouth Summer Research Project on Artificial Intelligence', 6 March 1956. The Dartmouth AI Archives; Box A. Personal Archives of Grace Solomonoff. <http://raysolomonoff.com/dartmouth/boxa/dart564props.pdf>.
- Simonds, John. 'Cornell Professor Eulogized in House'. *Journal (Unspecified)*. 29 July 1971. Cornell Faculty Bio files, Box 1 - Frank Rosenblatt Publications 1956-58. Division of Rare and Manuscript Collections, Cornell University Library.
- Simpkins, C. Alexander, and Annellen M. Simpkins. 'Review of: An Introduction to Cybernetics'. *IEEE Robotics & Automation Magazine* 20, no. 1 (March 2013): 96–96. <https://doi.org/10.1109/MRA.2013.2247707>.
- Singer, Marcus. Letter to Marvin Minsky, 13 February 1953. Minsky 2AR. Personal Archives of Gloria Rudisch.
- Skinner, Rebecca E. *Building the Second Mind: 1956 and the Origins of Artificial Intelligence Computing*. UC Berkeley Previously Published Works, 2012. <http://escholarship.org/uc/item/88q1j6z3>.
- Slayton, Rebecca. *Arguments That Count: Physics, Computing, and Missile Defense, 1949-2012*. Inside Technology. Cambridge, Massachusetts: The MIT Press, 2013.

- Smith, Sophie. 'The Language of "Political Science" in Early Modern Europe'. *Journal of the History of Ideas* 80, no. 2 (2019): 203–26. <https://doi.org/10.1353/jhi.2019.0012>.
- Solomon, Cynthia, Margaret Minsky, and Brian Harvey, eds. *LogoWorks: Challenging Programs in Logo*. Byte Books. New York: McGraw-Hill, 1986.
- 'Solomon Lefschetz'. School of Mathematics and Statistics, University of St Andrews, October 2003. <http://www-history.mcs.st-andrews.ac.uk/Biographies/Lefschetz.html>.
- Solomonoff, Grace. 'Ray Solomonoff (1926-2009)'. *Algorithms* 3, no. 3 (20 July 2010): 255–59. <https://doi.org/10.3390/a30302555>.
- . 'Ray Solomonoff and the Dartmouth Summer Research Project in Artificial Intelligence, 1956'. Oxbridge Research. The Dartmouth AI Archives. Personal Archives of Grace Solomonoff. Accessed 1 December 2016. <http://www.raysolomonoff.com/dartmouth/dartray.pdf>.
- . 'Ray Solomonoff and the New Probability'. In *Algorithmic Probability and Friends: Bayesian Prediction and Artificial Intelligence: Papers from the Ray Solomonoff 85th Memorial Conference, Melbourne, VIC, Australia, November 30 - December 2, 2011*, edited by David L. Dowe, 37–54. Lecture Notes in Computer Science 7070. Berlin: Springer, 2013.
- Solomonoff, Raymond. 'An Inductive Inference Machine', 14 August 1956. The Dartmouth AI Archives. Personal Archives of Grace Solomonoff. <http://raysolomonoff.com/publications/indinf56.pdf>.
- . 'An Inductive Inference Machine' IRE Convention Record, Section on Information Theory, Part 2 (1957): 56–62.
- . Letter to Marvin Minsky, Undated. Minsky 2AR. Personal Archives of Gloria Rudisch.
- . 'Letter to John McCarthy', 28 December 1956. SC0524 ACCN 1995-247 Box 22; 1948-56 Reprints, Articles. Department of Special Collections and University Archives, Stanford University Libraries.
- . 'Structure of Random Nets'. In *Proc. Int. Cong. Mathematicians*, 674–75. Providence, RI: American Mathematical Society, 1950.
- . 'Untitled Notes Re: Wendy Conquest', 2005. The Dartmouth AI Archives; Box A. Personal Archives of Grace Solomonoff. <http://raysolomonoff.com/dartmouth/boxa/raywendypoints.pdf>.
- . Letter to Marvin Minsky, 28 September 1956. Minsky 2AR. Personal Archives of Gloria Rudisch.
- Solomonoff, Raymond, and Anatol Rapoport. 'Connectivity of Random Nets'. *Bulletin of Mathematical Biophysics* 13, no. 2 (June 1951): 107–17. <https://doi.org/10.1007/BF02478357>.
- Solovey, M, and H Cravens, eds. *Cold War Social Science: Knowledge Production, Liberal Democracy, and Human*. New York: Palgrave Macmillan, 2016.
- Solovey, Mark. *Shaky Foundations: The Politics-Patronage-Social Science Nexus in Cold War America*. Studies in Modern Science, Technology, and the Environment. New Brunswick, N.J: Rutgers University Press, 2013.
- Solovey, Mark, and Hamilton Cravens, eds. *Cold War Social Science: Knowledge Production, Liberal Democracy, and Human*. New York: Palgrave Macmillan US, 2012. <https://doi.org/10.1057/9781137013224>.

- , eds. 'Foreward'. In *Cold War Social Science: Knowledge Production, Liberal Democracy, and Human*. New York: Palgrave Macmillan US, 2012.  
<https://doi.org/10.1057/9781137013224>.
- Soni, Jimmy, and Rob Goodman. *A Mind at Play: How Claude Shannon Invented the Information Age*. New York: Simon & Schuster, 2017.
- Sparrow, B., J. Liu, and D. M. Wegner. 'Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips'. *Science* 333, no. 6043 (5 August 2011): 776–78. <https://doi.org/10.1126/science.1207745>.
- Spearman, Mary Dale. Letter to Marvin Minsky, 25 October 1960. Minsky 3AR. Personal Archives of Gloria Rudisch.
- Speer, Albert, Richard Winston, Clara Winston, and Eugene Davidson. *Inside the Third Reich: Memoirs*. New York: Touchstone, 1997.
- Staley, Richard. 'The Interwar Period as a Machine Age: Mechanics, the Machine, Mechanisms, and the Market in Discourse'. *Science in Context* 31, no. 3 (September 2018): 263–92. <https://doi.org/10.1017/S0269889718000236>.
- Steinberger, Florian. 'Frege and Carnap on the Normativity of Logic'. *Synthese* 194, no. 1 (January 2017): 143–62. <https://doi.org/10.1007/s11229-015-0880-4>.
- . 'The Normative Status of Logic'. In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University, Spring 2017. <https://plato.stanford.edu/archives/spr2017/entries/logic-normative/>.
- Stern, Alexandra Minna. *Eugenic Nation*. 1st ed. University of California Press, 2005.  
[www.jstor.org/stable/10.1525/j.ctt1pn5jp](http://www.jstor.org/stable/10.1525/j.ctt1pn5jp).
- Stern, Nancy. 'Computers: From Eniac to Univac: As Inventors, Eckert and Mauchly Were Clear Successes, but as Entrepreneurs They Had Some Difficult Times'. *IEEE Spectrum* 18, no. 12 (1981): 61–69. <https://doi.org/10.1109/MSPEC.1981.6369705>.
- Stonier, Tom. *Beyond Information: The Natural History of Intelligence*. London ; New York: Springer-Verlag, 1992.
- Storing, Herbert J. 'Leonard D. White and the Study of Public Administration'. *Public Administration Review* 25, no. 1 (March 1965): 38.
- Stoyan, Herbert. 'Early LISP History (1956 - 1959)', 299–310. ACM Press, 1984.  
<https://doi.org/10.1145/800055.802047>.
- . 'LISP History'. *ACM Lisp Bulletin*, no. 3 (1 December 1979): 42–53.  
<https://doi.org/10.1145/1411829.1411837>.
- . 'The Influence of the Designer on the Design—J. McCarthy and LISP'. In *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, edited by John McCarthy and Vladimir Lifschitz, 409–26. Boston: Academic Press, 1991.
- 'Student Designs "Idiot Brain" to Measure Answer Patterns'. *Ithaca Journal*. 20 October 1953.
- 'Student Transcripts'. Kroch Library, n.d. Student Transcripts; Rosenblatt, Frank - Class of 1950. Cornell University Faculty Biographical files, #47-10-3394. Division of Rare and Manuscript Collections, Cornell University Library.
- 'Studies in Social Psychology in World War II: The Work of the War Department's Research Branch, Information and Education Division'. *Social Science Research Council* 3, no. 1 (March 1949).
- Subramaniam, V. 'Fact and Value in Decision Making'. *Public Administration Review* 23, no. 4 (December 1963): 232. <https://doi.org/10.2307/973898>.

- Suchman, Lucille Alice. *Human-Machine Reconfigurations: Plans and Situated Actions*. 2nd ed. Cambridge ; New York: Cambridge University Press, 2007.
- ‘Summer Research Training Institute on Simulation of Cognitive Processes, 30 June to 18 July 1958’, 3. Santa Monica, CA: The Rand Corporation, 1958.
- Sutherland, G. B. B. M. ‘The Mechanization of Thought Processes’. In *Mechanisation of Thought Processes: Proceedings of a Symposium Held at the National Physical Laboratory on 24th, 25th, 26th and 27th November 1958*, edited by National Physical Laboratory (Great Britain), ix. Symposium (National Physical Laboratory), v. 1. H.M. Stationery Office, 1959.
- ‘Symposium on the Impact of Computers on Science and Society’. *IEEE Transactions on Electronic Computers* EC-5, no. 3 (September 1956): 142–58.  
<https://doi.org/10.1109/TEC.1956.5219938>.
- ‘Symposium The Design of Machines to Simulate the Behavior of the Human Brain’. *IEEE Transactions on Electronic Computers* EC-5, no. 4 (December 1956): 240–55.  
<https://doi.org/10.1109/TEC.1956.5219960>.
- Taylor, Astra. ‘The Automation Charade’. *Logic Magazine*, August 2018.  
<https://logicmag.io/05-the-automation-charade/>.
- Tedre, Matti. ‘The Development of Computer Science: A Sociocultural Perspective’. University of Joensuu, 2006.
- Temple, Helen, and Marvin Minsky. ‘Request to Meet Dr. Hebb While in Montreal’, 11 May 1956. Minsky 2AR. Personal Archives of Gloria Rudisch.
- Terman, Lewis M. *The Measurement of Intelligence*. Cambridge, MA: The Riverside Press, 1916.
- ‘The 1958 International Conference on Scientific Information’. Washington, D.C, 1958.
- The Future of Artificial Intelligence: Views from History*. University of Cambridge: The Leverhulme Centre for the Future of Intelligence, 2018.  
<https://www.youtube.com/watch?v=hjB35dRUhi4>.
- ‘Third of the Fall Series on Industrial Liaison Symposia’. Boston, MA: Massachusetts Institute of Technology, 1959.
- Thurstone, L. L. *Multiple-Factor Analysis; a Development and Expansion of The Vectors of Mind*. Chicago, Ill: University of Chicago Press, 1947.
- . *The Vectors of Mind: Multiple Factor Analysis for the Isolation of Primary Traits*. The University of Chicago Press, 1935.  
<https://archive.org/details/vectorsofmindmul010122mbp>.
- Todd, Zoe. ‘An Indigenous Feminist’s Take On The Ontological Turn: “Ontology” Is Just Another Word For Colonialism’. *Journal of Historical Sociology* 29, no. 1 (March 2016): 4–22. <https://doi.org/10.1111/johs.12124>.
- Torgerson, D. ‘Policy Analysis and Public Life: The Restoration of Phronesis?’ In *Political Science in History: Research Programs and Political Traditions*, edited by James Farr, John S. Dryzek, and Stephen T. Leonard. Cambridge ; New York: Cambridge University Press, 1995.
- ‘Tribute to Dr. Frank Rosenblatt’. Congressional Record, Proceedings and Debates of the 92nd Congress, First Session. U.S. House of Representatives, 28 July 1971.
- Tucker, Albert. The Educational Program at Princeton in the 1930s [Transcript no. 31], 12 April 1984. Department of Mathematics Oral History Project records, Box 2. Princeton University Archives, Department of Rare Books and Special Collections,

- Princeton University Library.  
<https://findingaids.princeton.edu/collections/AC057/c29>.
- Turing, Alan. 'Computing Machinery and Intelligence'. *Mind* 49 (1950): 433–60.
- Turner, Fred. *From Counterculture to Cyberculture: Stewart Brand, the Whole Earth Network, and the Rise of Digital Utopianism*. 1. paperback ed. Chicago, Ill.: Univ. of Chicago Pr, 2008.
- Turner, Raymond. 'Programming Languages as Technical Artifacts'. *Philosophy & Technology* 27, no. 3 (September 2014): 377–97. <https://doi.org/10.1007/s13347-012-0098-z>.
- Unknown. 'Letter to Mr. Marvin L. Minsky from Bell Laboratories, June 6, 1952', 6 June 1952. Personal Archives of Gloria Rudisch.
- Urwick, L. F. 'Papers in the Science of Administration'. *Academy of Management Journal* 13, no. 4 (1 December 1970): 361–71. <https://doi.org/10.2307/254827>.
- Volmar, Daniel. 'The Computer in the Garbage Can: Air-Defense Systems in the Organization of US Nuclear Command and Control, 1940–1960'. PhD Dissertation, Harvard University, 2018.
- Von Neumann, John. *Functional Operators (AM-22), Volume 2*. Princeton University Press, 1950. <http://www.jstor.org/stable/j.ctt1bc543b>.
- . *The Computer and the Brain*. Edited by Klara Von Neumann. New Haven: Yale University Press, 1958.
- Von Neumann, John, William Aspray, and Arthur W. Burks, eds. 'First Draft of a Report on the EDVAC'. In *Papers of John von Neumann on Computing and Computer Theory*. Charles Babbage Institute Reprint Series for the History of Computing, v. 12. Cambridge, Mass. : Los Angeles: MIT Press ; Tomash Publishers, 1987.
- Von Neumann, John, and Lloyd A. Jeffress. 'The General and Logical Theory of Automata'. In *Cerebral Mechanisms in Behavior*, 1–41. London: Chapman & Hall, 1951.
- Von Neumann, John, and A. H. Taub. 'The General and Logical Theory of Automata'. In *John von Neumann: Collected Works; Design of Computers, Theory of Automata and Numerical Analysis*, 5:288–328. London: Pergamon Press, 1961.
- Ware, Willis H. *RAND and the Information Evolution: A History in Essays and Vignettes*. Santa Monica, CA: Rand Corp, 2008.
- Weaver, Warren. 'Inter-Office Correspondence, WW to RSM', 14 June 1955. Record group 1.0002, series 200, box 26, folder 219. Rockefeller Archive Center, Sleepy Hollow, NY.
- . 'WW Diary Entry Re: John McCarthy', 4 April 1955. Record group 1.0002, series 200, box 26, folder 219. Rockefeller Archive Center, Sleepy Hollow, NY.
- Weber, Max. 'Science as a Vocation'. In *From Max Weber: Essays in Sociology*, edited by H. H. Gerth and C. Wright Mills, 129–56. New York: Oxford University Press, 1946.
- Weinstein, Deborah. *The Pathological Family: Postwar America and the Rise of Family Therapy*. Cornell University Press, 2013.
- Weiss, E. A., and Henry P. T. Corley. 'Letters to the Editor'. *Communications of the ACM* 1, no. 4 (1 April 1958): 5. <https://doi.org/10.1145/368796.368802>.
- Weizenbaum, Joseph. *Computer Power and Human Reason: From Judgment to Calculation*. San Francisco: Freeman, 1976.
- Wells, Webster, and Walter Wilson Hart. *Modern Plane Geometry, Graded Course*. D. C. Heath and Company, 1926.
- Wernimont, Jacqueline. *Numbered Lives: Life and Death in Quantum Media*. Media Origins. Cambridge, MA: The MIT Press, 2018.

- White, Leonard. *Introduction to the Study of Public Administration*. New York: Macmillan Co., 1926.
- Widrow, Bernard. Bernard Widrow, an oral history. Interview by Andrew Goldstein, 14 March 1997. [http://ethw.org/Oral-History:Bernard\\_Widrow](http://ethw.org/Oral-History:Bernard_Widrow).
- Wiener, Norbert. *Cybernetics or Control and Communication in the Animal and the Machine*. Cambridge, Mass: MIT Press, 2007.
- . Letter to Marvin Minsky, 14 December 1953. Minsky 2AR. Personal Archives of Gloria Rudisch.
- Wilkes, M. 'Can Machines Think?' *Proceedings of the Institute of Radio Engineers* 41, no. 10 (October 1953): 1230–34. <https://doi.org/10.1109/JRPROC.1953.274272>.
- Wilkes, Maurice V., David J. Wheeler, and Stanley Gill. *The Preparation of Programs for an Electronic Digital Computer: With Special Reference to the EDSAC and the Use of a Library of Subroutines*. Cambridge, MA: Addison-Wesley Press, 1951.
- Williams, John D. 'Comments on RAND Building Program'. Santa Monica, CA: RAND Corporation, 26 December 1950. <https://www.rand.org/pubs/classics/building.html>.
- Wilson, Elizabeth A. *Affect and Artificial Intelligence*. In *Vivo : The Cultural Mediations of Biomedical Science*. Seattle: University of Washington Press, 2010.
- Winner, Langdon. *Autonomous Technology: Technics-out-of-Control as a Theme in Political Thought*. 9. printing. Cambridge, Mass.: MIT Pr, 2001.
- Wolf, Micheal Running, and Caroline Running Wolf. 'Sacred Waveforms: An Indigenous Perspective on the Ethics of Collecting and Usage of Spiritual Data for Machine Learning'. Presented at the NeurIPS | 2019 Thirty-third Conference on Neural Information Processing Systems, Vancouver, Canada, 14 December 2019. <https://slideslive.com/38922118/joint-workshop-on-ai-for-social-good-1>.
- Wolfe, Audra J. *Freedom's Laboratory: The Cold War Struggle for the Soul of Science*. Baltimore: Johns Hopkins University Press, 2018.
- Wright, T. P. Letter to S. S. Atwood. 'Perceptron and Allied Matters - Frank Rosenblatt', 24 November 1959. Dale R Corson papers, 1963-1977; Box 16 Folder 37. Division of Rare and Manuscript Collections, Cornell University Library.
- Yale, Elizabeth. 'The History of Archives: The State of the Discipline'. *Book History* 18, no. 1 (2015): 332–59. <https://doi.org/10.1353/bh.2015.0007>.
- Yarmush, David. Letter to Marvin Minsky, 8 January 1951. Minsky 2AR. Personal Archives of Gloria Rudisch.
- Yates, JoAnne. *Control through Communication: The Rise of System in American Management*. Johns Hopkins paperbacks ed. Studies in Industry and Society 6. Baltimore, Md.: Johns Hopkins Univ. Press, 1993.
- Yockey, Hubert P. *Symposium on Information Theory in Biology, Gatlinburg, Tennessee, October 29-31, 1956*. New York: Pergamon Press, Symposium Publications Division, 1958.