DOI: 10.1017/psa.2022.66
This is a manuscript accepted for publication in *Philosophy of Science*.
This version may be subject to change during the production process.

Title: Unlimited Associative Learning as a Null Hypothesis

Abstract: A common strategy in comparative cognition is to require that one reject associative learning as an explanation for behavior before concluding that an organism is capable of causal reasoning. In this paper, I argue that standard causal-reasoning tasks can be explained by a powerful form of associative learning: unlimited associative learning (UAL). The lesson, however, is not that researchers should conduct more studies to reject UAL, but that they should instead focus on 1) enriching the cognitive hypothesis space and 2) testing a broader range of information processing patterns—errors, biases and limits, rather than successful problem solving alone.

Contact Information: Marta Halina**.** Email Address: mh801@cam.ac.uk**.** Telephone: +44 (0)7470404750**.** Mailing Address: Department of History and Philosophy of Science, University of Cambridge, Free School Lane, Cambridge CB2 3RH, United Kingdom

Acknowledgments: Thank you to Andrew Barron, Rachael Brown, Colin Klein and two anonymous reviewers for their valuable feedback. Warm thanks also to David Harrison, Phillip Kieval, Victor Parchment, Konstantinos Voudouris, and Laura Weidinger for their helpful comments and discussion. This research was funded in part by a grant from the Templeton World Charity Foundation.

1. Introduction

Associative learning has long been considered an appropriate “null hypothesis” against which to test claims about the cognitive capacities of human and nonhuman animals. This is particularly the case when the cognitive capacity under investigation is thought to be complex or human-like, such as in the case of causal reasoning, episodic memory, and theory of mind. As Starzak and Gray (2021) note, “Over and over again the familiar refrain is, ‘do animals have complex human-like cognitive abilities or can their behavior be explained in terms of simpler processes such as associative learning?’” (2). The general idea behind this approach is that if an observed behavior can be explained by appealing to a process like associative learning, then one should accept this as the best explanation for the behavior, rather than attributing new sophisticated capacities to an organism. The primary justification for treating associative learning as a null hypothesis is that it is phylogenetically widespread; thus, it’s reasonable to assume that the organism being tested is capable of associative learning and will use it when possible to solve the problem at hand (Sober 2012).

Over the past decade, Simona Ginsburg, Eva Jablonka and others have argued that human and nonhuman animals are capable of a particularly powerful form of associative learning: unlimited associative learning (UAL). Organisms with UAL can discriminate novel, compound stimuli and action patterns; learn associations between objects and events separated over time; and engage in cumulative learning. Researchers studying UAL maintain that it is found in almost all vertebrates, as well as some arthropods and cephalopod molluscs (Ginsburg and Jablonka 2021). If this is correct, then UAL is arguably also an appropriate null hypothesis for research in comparative psychology, at least in the case of vertebrates.

In this paper, I argue that UAL poses a problem for research in comparative psychology. Using causal reasoning as a case study, I show that claims regarding an organism's ability to engage in causal reasoning fail to reject UAL as a plausible alternative explanation for the available results. My conclusion, however, is not that researchers should conduct more studies with the aim of rejecting UAL as a null hypothesis. Instead, I argue that this adds to the growing consensus that the “null hypothesis" approach is problematic. Researchers should reject this approach as it oversimplifies the target of study. Instead we should endorse more fine-grained comparative approaches, such as ones that treat cognitive abilities as multi-dimensional (Starzak and Gray 2021) and focus on “signatures” rather than “success” (Bastos and Taylor 2020).

I begin in section 2 by briefly illustrating the null-hypothesis approach as it’s used in experiments on causal cognition. In section 3, I introduce UAL and show how it can explain successful performance on causal-cognition tasks. I then argue in section 4 that UAL provides a compelling reason for rejecting the null-hypothesis approach in comparative psychology and points us in the direction of a more fruitful approach. Section 5 concludes.

2. Testing Causal Reasoning

Nonhuman animals (hereafter animals) are capable of solving a wide range of physical problems, from a woodpecker finch using a cactus spine to prize an insect out of a crevice to chimpanzees outperforming human children on some physical cognition tasks (Herrmann et al. 2007). However, there are several competing explanations in the literature regarding how animals succeed in solving physical tasks. One explanation is that they engage in causal reasoning. Causal reasoning is understood in psychology as the ability to intervene on and make inferences about the world based on the world's causal structure (Bender 2020). Causal structure includes phenomena like heavy things fall to the ground, water is displaced by sinking (rather than floating) objects, and some objects can be used to displace or dislodge others.[1] According to standard views in psychology, an agent capable of causal reasoning should recognize the functional properties or physical affordances of a situation and use these to solve problems (a heavy object, whether made of stone or metal, will displace water). Such an agent should also be able to transfer knowledge acquired in one situation to another

[1] There is a rich exchange between philosophers and psychologists regarding how to understand causal structure and causal reasoning (e.g., see Woodward 2011). The target of my analysis here is how psychologists empirically investigate causal reasoning; thus, I will focus on accounts of causal reasoning as presented in this empirical literature.

https://doi.org/10.1017/psa.2022.66 Published online by Cambridge University Press

functionally equivalent situation, even if the two situations differ in all their non-functional properties (Seed et al. 2011).

A second dominant explanation for an agent's success in solving a physical problem is that the agent is relying on associative learning. In this case, the agent does not grasp the causal structure underlying the situation, but instead relies on some learned association between variables. This associative-learning explanation is referred to as a “null hypothesis” in the literature (Hanus 2016). Some researchers argue that only humans solve physical problems through causal reasoning—that all other animals rely on some form of associative learning for their success on causal tasks (Penn and Povinelli 2007; Povinelli 2012). However, empirical studies appear to undermine this view, suggesting that animals such as corvids and chimpanzees rely on causal knowledge to solve novel problems (Seed et al. 2006; Mulcahy and Call 2006, Girndt et al. 2008).

It's helpful to illustrate the null-hypothesis strategy with an example. A benchmark test for causal reasoning is the trap-tube task. In this task, participants are presented with a transparent tube baited with a reward, such as food. The tube contains various traps that must be avoided if the reward is to be successfully extracted. Participants must use their body (e.g., finger or beak) or a tool (such as a stick or rake) to extract (by pushing or pulling) the reward from the tube while avoiding the traps. If the reward falls into a trap, it can no longer be retrieved.

https://doi.org/10.1017/psa.2022.66 Published online by Cambridge University Press

The comparative psychologist Amanda Seed and colleagues have conducted numerous studies investigating the causal reasoning of animals (Seed et al. 2011). In one study, they examined whether rooks (*Corvus frugilegus*) used causal reasoning to solve the trap-tube task (Seed et al. 2006). To do this, they first tested whether rooks could solve two different versions of the trap-tube task (Tube A and Tube B in figure 1). Both versions included one functional trap and one “decoy” or non-functional trap that looked similar to a functional trap but did not interfere with reward retrieval. In Tube A, the reward could pass over the top of the decoy trap, while in Tube B, the reward would fall into the hole of the decoy trap, but this hole was open at the bottom, so rather than trapping the reward, the reward fell through and was obtainable by the participant. Four rooks were tested on Tube A and four on Tube B. After learning to successfully solve this problem, the rooks were then tested to see if they could transfer their knowledge to a new situation. Those that solved Tube A were tested on Tube B and vice versa. All of the birds that had successfully learned the original task (seven out of eight participants), succeeded in transferring to the new task, suggesting that they were not relying on a cue-based rule that was idiosyncratic to the original task, like “drop the reward into the hole with an opening at the bottom”.

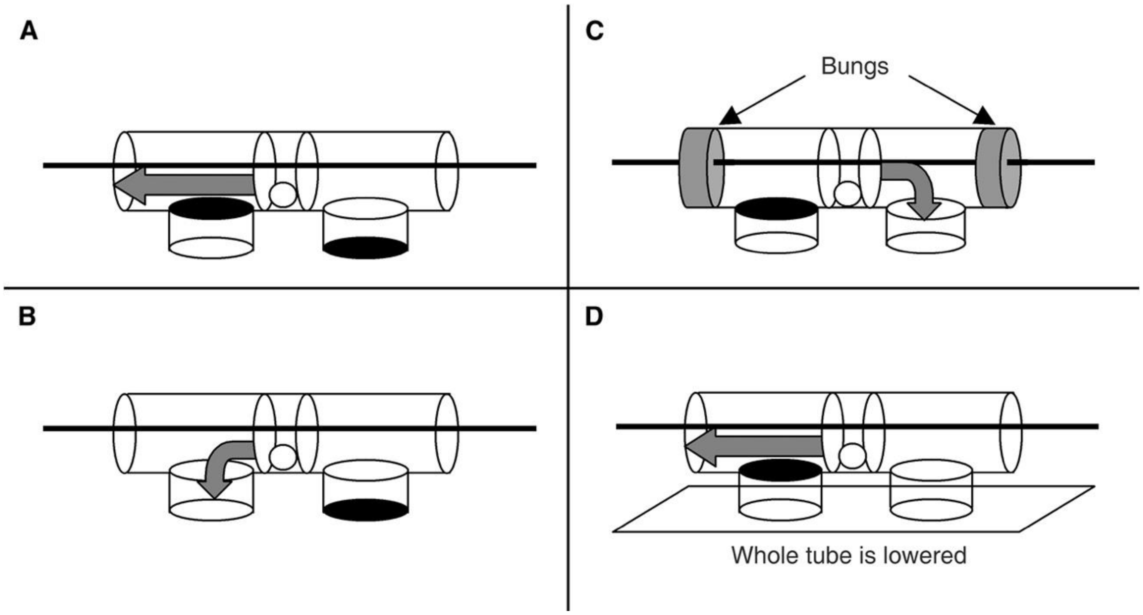

Figure 1. Trap-tube experimental apparatus from Seed et al. (2006). In Experiment 1, four rooks were tested using Tube A and four rooks were tested using Tube B. Experiment 2 tested whether those rooks who solved Tube A could solve Tube B and vice versa. Experiment 3 tested whether participants could solve Tube C and Tube D (with four birds receiving 20 trials of Tube C followed by 20 trials of Tube D and three birds receiving the same number of trials on first Tube D and then Tube C).

Although the rooks demonstrated that they were not relying on a cue that was idiosyncratic to a single task, they could still have been using a cue-based rule that was common to both Tube A and Tube B like, "avoid the hole with the black disc at the bottom" (Seed et al. 2006, 698). Thus, Seed and colleagues examined whether these birds could additionally transfer to two more novel tasks: Tube C and Tube D. Neither of these tubes contained a trap with a black disc at the bottom. Instead, the two decoy traps in Tubes A and B were converted into functional traps (see figure 1). The birds could also not rely on one single cue-based procedural rule to perform successfully on both Tubes C and D, as these tubes had no useful cues in common. For example, although a bird could use the rule “move the reward away from the black circle” to solve Tube C, this rule would fail if applied to Tube D.

One out of seven rooks successfully transferred to Tube C and Tube D. This bird (Guillem) successfully solved all four trap-tube tasks, which the authors tentatively suggest means that this individual “understood the unobservable [causal] features of the task” (Seed et al. 2006, 700).

This case study illustrates how researchers investigate causal reasoning in animals and the role associative learning plays in these investigations. Associative learning predicts that animals such as rooks can learn associative rules like “moving the object away from the black disc will get me a reward”. To determine whether rooks engage in causal reasoning, one must eliminate this alternative explanation and others like it. Doing so requires implementing control conditions that help reveal when participants might be relying on arbitrary cue-based rules, rather than relying on the underlying causal features of the system to solve the task. As Seed and colleagues write in a review of this and other causal-cognition studies: “These results suggest that the rook, chimpanzees, and New Caledonian crows did not use simple perceptual cues to solve the trap task. We propose that instead they extracted causally relevant functional information (such as surface continuity, or the solidity of barriers)" (Seed et al. 2011, 13).

Although simple associative rules might not account for the results of trap-tube tasks such as these, I argue in the next section that one can account for these results by appealing to more sophisticated forms of associative learning like unlimited associative learning (UAL). Under the current null-hypothesis approach, this finding suggests that researchers should shift their focus from eliminating simple cue-based rules to rejecting UAL as a plausible alternative explanation for successful performance on causal-reasoning tasks. In the remainder of the paper, I resist this conclusion, arguing instead that what needs rejecting is the null-hypothesis approach.

3. Unlimited Associative Learning as a Null Hypothesis

Associative learning is broadly the ability to learn associations between stimuli like a bell and food (classical condition) or between actions and outcomes like receiving food upon pressing a lever (operant conditioning). UAL is a form of associative learning in that it involves learning associations between objects, events and actions. However, an animal with UAL is capable of learning associations between a practically limitless number of stimuli and actions. Agents with UAL can learn to associate novel, compound stimuli and actions that are

https://doi.org/10.1017/psa.2022.66 Published online by Cambridge University Press

temporally separated, as well as engage in second-order conditioning (Bronfman et al. 2016; Birch et al. 2020).

Given that UAL is effectively unlimited in the range of associations that can be formed, an agent with UAL faces a problem: the world contains an endless number of covarying factors, how does such an agent learn to associate the relevant factors, while ignoring those that are irrelevant? Ginsburg and Jablonka (2019) refer to this as the "loading the dice" problem following William James (349-50). Daniel Dennett has also noted that Skinnerian creatures (i.e., those creatures that learn through associative learning) survive in virtue of making lucky first moves (Dennett 1996, 88). The loading-the-dice problem asks, how do such creatures load the dice in a way that allows them to gain useful information about the world, given the vast number of potential associations available to them?

The model of UAL advanced by Ginsburg, Jablonka and colleagues provides an answer. Inspired in part by the predictive processing literature, they argue that animals construct generative models of their environment. These models produce predictions of sensory input based on an organism's evolutionary history and prior learning. When there's a deviation between these predictions and the incoming sensory input (i.e., a prediction error), this creates an imperative to bring the predictions and sensory input into alignment. This can be done by either updating the generative model (reactive inference) or seeking out signals that agree with the model (active inference). In this way, the imperative to minimize prediction error drives learning, attention, and action. Crucially, only discrepancies between expected and actual data are registered, allowing vast amounts of incoming information to be ignored. Moreover, under the predictive processing account, some aspects of the world are dismissed as 'noise' in the sense that the associated prediction errors are not given much weight in updating the generative model. This is known as 'precision-weighting' and, as Andy Clark writes, it "delivers the system's best estimate of the trustworthiness of the sensory information itself" (Clark 2016, 60). Prediction errors that are estimated as reliable will have greater effects in terms of learning, attention, and action than prediction errors that are estimated as unreliable.

UAL can account for the successful performance of animals on causal reasoning tasks. First, it can explain how an organism learns the underlying causes or functional properties of a task more readily than arbitrary cue associations. For example, corvids (both those who routinely use tools, like New Caledonian Crows, and those who do not, like common ravens), spend a large proportion of their time manipulating objects throughout development (Kenward et al. 2011). Thus, they are exposed to numerous causal invariances, such as the invariance that solid objects do not pass through other solid objects, and that objects fall when they reach the edge of a surface and are no longer supported from below. Such invariances can be contrasted with more variable properties of the world. For example, barriers typically come in a wide range of colors—this property is a noisy signal. We should expect a rook's generative model then to update in response to causal invariances, given the reliability of the prediction errors that result. In contrast, such a model should fail to represent noisy signals like the colors of barriers. Although a rook is capable of learning to associate arbitrary cues, this should be

https://doi.org/10.1017/psa.2022.66 Published online by Cambridge University Press

more challenging than learning causal regularities. We should thus expect rooks to draw on these regularities when solving physical problems more readily than associations between arbitrary cues.

UAL can also account for successful transfers to novel problems like that exhibited by the rook Guillem in the trap-tube task. Insofar as the non-functional or decoy traps do not violate causal laws, we should expect an agent capable of UAL to generate predictions that accord with previously learned causal invariances, such as "solid objects will pass over solid surfaces, fall if unsupported, etc.". Indeed, under UAL, what is more perplexing is why subjects fail problems such as the trap-tube task, given their extensive experience with solid objects and containers. As Seed et al. note, such failures are difficult to explain, but they may be related to a lack of ecological validity (Seed et al. 2006; for examples, see Girndt et al. 2008; Mulcahy and Call 2006). In either case, UAL predicts that animals such as rooks and chimpanzees will be less likely to learn a rule relating arbitrary cues than a rule that reflects causal principles provided they have been exposed to those causal principles in the past. This is true even if those causal principles manifest in different ways—that is, if they're realized in situations that differ in their perceptual features. As Clark (2016) writes, it is the "structured probabilistic know-how distilled from prediction-driven learning that enables us to *see through the veil of surface statistics* to the world of distal interacting causes itself" (170, emphasis original). Generative models are hierarchical structures that represent latent variables at different levels of abstraction. Such models can, for example, construct a "best explanation" for a range of multimodal sensory inputs (Clark 2016, 174). Indeed, even deep convolution neural networks (with their many layers of hierarchical processing) appear to be able to perform what Buckner (2018) calls "transformational abstraction" or the ability to move between specific instances (of chairs, for example, with no perceptual features in common) to deeper representations or abstractions.

One might object that it's no surprise that UAL can explain behavior in causal-reasoning tasks, given its reliance on predictive processing. Predictive processing is well known for providing an account of how agents learn causal models of the world (Hohwy 2020; Williams 2018). Comparing UAL with predictive-processing accounts of causal reasoning is beyond the scope of this paper; however, it's worth noting that, to my knowledge, proponents of UAL have not considered UAL as a potential source of causal knowledge. The closest account of this kind is Ginsburg and Jablonka (2021), which suggests that organisms capable of UAL might serve as the basis for an evolutionary transition to "Popperian creatures" or those capable of using imagination to evaluate and select actions before trying them out in the world (Dennett 1996). However, Ginsburg and Jablonka do not discuss whether or how UAL might serve as the basis for "Pearlian creatures" or those capable of causal reasoning (Godfrey-Smith 2018).

4. From Null Hypotheses to Signature Testing

As noted, associative learning is viewed as the appropriate "null hypothesis" when investigating causal reasoning in human and nonhuman animals. As Daniel Hanus writes about the standard practice in comparative psychology, "an associative explanation should be the null hypothesis that must be rejected before any cognitive explanation [like causal reasoning] should be assumed" (2016, 242). Researchers thus design experiments to exclude associative learning as an explanation for task performance. Given this, should they also treat unlimited associative learning as a null hypothesis that must be rejected before one can conclude that an organism is engaging in causal reasoning? In this section, I argue that they should not. Although UAL is a plausible alternative explanation for performance on causal-reasoning tasks, the null-hypothesis approach is a poor method for evaluating hypotheses. We should reject the null-hypothesis approach and adopt instead a signature-testing approach for investigating causal reasoning.

There has been much discussion in philosophy of science on the null-hypothesis strategy. One recurring theme is that many hypotheses labelled "null" are in fact substantive hypotheses about the world. Given this, ceteris paribus, they should not be epistemically privileged over alternative hypotheses (Fitzpatrick 2008; Bausman 2018; Bausman and Halina 2018; Dacey 2021). Instead, like other substantive hypotheses, one should weigh the epistemic values of the purported null against the epistemic values of alternative hypotheses (e.g., internal consistency, empirical adequacy, scope, explanatory power, unification, novel prediction, etc.; see Douglas 2013). If a purported null hypothesis like associative learning is

https://doi.org/10.1017/psa.2022.66 Published online by Cambridge University Press

on a par with an alternative hypothesis with respect to its epistemic values, then there is no reason to choose the "null" over the alternative. In this way, the label "null" is misleading, as it suggests one is justified in adopting a strategy similar to statistical null hypothesis testing where one must reject the statistical null before accepting the alternative (Bausman and Halina 2018). The claim that animals engage in associative learning or UAL is a substantive claim about the world, however. Thus, unless epistemic values weigh in UAL's favor, we should not prefer this hypothesis over causal reasoning as an explanation for success on causal reasoning tasks. We should also eschew the language of "null hypotheses" altogether to avoid conflating associative learning and UAL with statistical null hypotheses (Dacey 2021).

Although UAL should not be privileged independently of evidence and other epistemic considerations, one might argue that it is the best explanation, given our current background knowledge. One justification for selecting UAL as the best explanation for an organism's performance on causal reasoning tasks is that UAL is found in that organism's evolutionary relatives (Sober 2012, 2015; Currie 2021). On the basis of cladistic parsimony, if two related species share a phenotype and it's known that the proximate mechanism causing this phenotype for one species is *M*, then this is evidence that the proximate mechanism causing the phenotype in the other species is also *M* (Sober 2012). Ginsburg and Jablonka argue on the basis of current empirical evidence that UAL can be found in almost all vertebrates (Ginsburg and Jablonka 2021). Thus, even if UAL has not been found specifically in rooks, we have

https://doi.org/10.1017/psa.2022.66 Published online by Cambridge University Press

reason to believe it's operating in this taxon, given its wide (and likely deep) phylogenetic distribution, whereas we don't have such evidence for causal reasoning.

The above considerations may justify choosing UAL over causal reasoning as the best explanation for vertebrate performance on causal reasoning tasks. However, this evaluation presumes that we must choose between two mutually exclusive hypotheses (causal reasoning and UAL), rejecting one and accepting the other (see Voudouris 2020). The hypothesis space is more complicated than this, though. First, it's unclear that UAL and causal reasoning are mutually exclusive. As we've seen, UAL provides a plausible account of how agents learn causal invariances. This diverges from accounts of causal reasoning that take causal knowledge as "core knowledge" shaped largely by evolution and changing little over ontogeny (Spelke 1994). UAL and such causal accounts, however, overlap in positing that organisms employ causal models of the world. Second, we could populate the hypothesis space with hybrid accounts where some causal structure is present in an organism via unlearned priors and other structure is associatively learned. Unlearned priors could also range from very general (e.g., a preferential orientation towards biological motion) to more specific (a principle of continuity). To choose between causal reasoning and UAL is to oversimply the hypothesis space.

Where then to go from here? I suggest we move away from "success testing" and towards what Bastos and Taylor (2020) call "signature testing". Success testing focuses on whether agents pass or fail tests, where passing a test is taken as evidence for a (usually sophisticated) cognitive capacity unless an alternative hypothesis like associative learning can explain the results. Such tests are weakly diagnostic—they minimally constrain the hypothesis space. As we have seen, success on the trap-tube task fails to distinguish between causal reasoning and UAL. In contrast, signature testing provides additional constraints on the hypothesis space by examining a wide range of information processing patterns, including errors, biases, and limitations (Bastos and Taylor 2020). For example, UAL requires selective attention (Birch et al. 2020). We should thus expect organisms with impaired selective attention (e.g., due to impairments in the midbrain superior colliculus) to exhibit limitations in their capacity to learn new causal information. An organism's ability to learn new causal invariances (e.g., in a virtual world with unusual physics) should also vary depending on their capacity to employ UAL. We should thus expect the frequency of errors on physical tasks in such a world to decrease over time insofar as an organism is relying on UAL.[2] The crucial point here is that we're evaluating the hypothesis space by looking at a range of information processing patterns, not just success in causal-reasoning tasks designed to eliminate associative learning as the main competing hypothesis or "null". Enriching the hypothesis

[2] See Starzak and Gray (2021) for additional dimensions along which causal cognition might vary and Seed et al. (2011) for a middle way (what they call "structural knowledge") between associative learning and adult human causal reasoning based on symbolic knowledge.

space together with signature testing provides a powerful alternative to the null-hypothesis approach.

5. Conclusion

Under standard methods in comparative cognition, unlimited associative learning is best understood as a null hypothesis that must be eliminated before one can conclude that an organism is capable of causal reasoning. Such an approach suggests that associative learning and causal reasoning are the only plausible cognitive explanations for performance on causal reasoning tasks. Instead, we should reject the null-hypothesis approach and evaluate cognitive hypotheses according to their epistemic values. We should also focus on enriching the space of hypotheses while adopting methods that can tightly constrain that space. This is best done using a signature-testing rather than success-testing approach.

https://doi.org/10.1017/psa.2022.66 Published online by Cambridge University Press

References

Bastos, Amalia P. M., and Alex H. Taylor. 2020. “Macphail’s Null Hypothesis of Vertebrate Intelligence: Insights From Avian Cognition.” *Frontiers in Psychology* 11 (1): 1–8.

Bausman, William C. 2018. “Modeling: Neutral, Null, and Baseline.” *Philosophy of Science* 85 (4): 594–616.

Bausman, William, and Marta Halina. 2018. “Not Null Enough: Pseudo-Null Hypotheses in Community Ecology and Comparative Psychology.” *Biology & Philosophy* 33 (3–4): 1–20.

Bender, Andrea. 2020. “What Is Causal Cognition?” *Frontiers in Psychology* 11: 1–6.

Birch, Jonathan, Simona Ginsburg, and Eva Jablonka. 2020. “Unlimited Associative Learning and the Origins of Consciousness: A Primer and Some Predictions.” *Biology & Philosophy* 35 (6): 1–23.

Bronfman, Zohar Z., Simona Ginsburg, and Eva Jablonka. 2016. “The Transition to Minimal Consciousness through the Evolution of Associative Learning.” *Frontiers in Psychology* 7: 1–16.

Buckner, Cameron. 2018. “Empiricism without Magic: Transformational Abstraction in Deep Convolutional Neural Networks.” *Synthese* 195 (12): 5339–5372.

Clark, Andy. 2016. *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford: Oxford University Press.

Currie, Adrian. 2021. *Comparative Thinking in Biology*. 1st ed. Cambridge: Cambridge University Press.

Dacey, Mike. 2021. “Evidence in Default: Rejecting Default Models of Animal Minds.” *The British Journal for the Philosophy of Science*, 1–42.

Dennett, Daniel C. 1996. *Kinds of Minds: Towards an Understanding of Consciousness*. New York: Basic Books.

Douglas, Heather. 2013. “The Value of Cognitive Values.” *Philosophy of Science* 80 (5): 796–806.

Fitzpatrick, Simon. 2008. “Doing Away with Morgan’s Canon.” *Mind & Language* 23 (2): 224–46.

Ginsburg, Simona, and Eva Jablonka. 2019. *The Evolution of The Sensitive Soul: Learning and The Origins of Consciousness*. Cambridge: MIT Press.

Ginsburg, Simona, and Eva Jablonka. 2021. “Evolutionary Transitions in Learning and Cognition.” *Philosophical Transactions of the Royal Society B: Biological Sciences* 376 (1821): 1–9.

Girndt, Antje, T. Meier, and J. Call. 2008. “Task Constraints Mask Great Apes’ Ability to Solve the Trap-Table Task.” *Journal of Experimental Psychology: Animal Behavior Processes* 34 (1): 54–62.

Godfrey-Smith, Peter. 2018. “Towers and Trees in Cognitive Evolution.” In *The Philosophy of Daniel Dennett*, ed. Bryce Huebner, 225–49. Oxford: Oxford University Press.

Hanus, Daniel. 2016. “Causal Reasoning versus Associative Learning: A Useful Dichotomy or a Strawman Battle in Comparative Psychology?” *Journal of Comparative Psychology* 130 (3): 241–48.

Herrmann, E., J. Call, M. V. Hernandez-Lloreda, B. Hare, and M. Tomasello. 2007. "Humans Have Evolved Specialized Skills of Social Cognition: The Cultural Intelligence Hypothesis." *Science* 317 (5843): 1360–66.

Hohwy, Jakob. 2020. "New Directions in Predictive Processing." *Mind & Language* 35 (2): 209–23.

Kenward, Ben, Christian Schloegl, Christian Rutz, Alexander A. S. Weir, Thomas Bugnyar, and Alex Kacelnik. 2011. "On the Evolutionary and Ontogenetic Origins of Tool-Oriented Behaviour in New Caledonian Crows (*Corvus Moneduloides*)." *Biological Journal of the Linnean Society* 102 (4): 870–77.

Mulcahy, Nicholas J., and Josep Call. 2006. "How Great Apes Perform on a Modified Trap-Tube Task." *Animal Cognition* 9 (3): 193–99.

Penn, Derek C., and Daniel J. Povinelli. 2007. "Causal Cognition in Human and Nonhuman Animals: A Comparative, Critical Review." *Annual Review of Psychology* 58 (1): 97–118.

Povinelli, Daniel J. 2012. *World without Weight: Perspectives on an Alien Mind*. Oxford: Oxford University Press.

Seed, Amanda, Daniel Hanus, and Josep Call. 2011. “Causal Knowledge in Corvids, Primates, and Children.” In *Tool Use and Causal Cognition*, ed. Teresa McCormack, Christoph Hoerl, and Stephen Butterfill, 89–110. Oxford: Oxford University Press.

Seed, Amanda M., Sabine Tebbich, Nathan J. Emery, and Nicola S. Clayton. 2006. “Investigating Physical Cognition in Rooks, Corvus Frugilegus.” *Current Biology* 16 (7): 697–701.

Sober, Elliott. 2012. “Anthropomorphism, Parsimony, and Common Ancestry.” *Mind & Language* 27 (3): 229–38.

Sober, Elliott. 2015. *Ockham’s Razors: A User’s Manual*. Cambridge: Cambridge University Press.

Spelke, Elizabeth. 1994. “Early Knowledge of Object Motion: Continuity and Inertia.” *Cognition* 51: 131–176.

Starzak, Tobias Benjamin, and Russell David Gray. 2021. "Towards Ending the Animal Cognition War: A Three-Dimensional Model of Causal Cognition." *Biology & Philosophy* 36 (2): 1–24.

Voudouris, Konstantinos. 2020. "An Interdisciplinary Approach to the Associative-Cognitive Distinction." MPhil diss., University of Cambridge.

Williams, Daniel. 2018. "Predictive Processing and the Representation Wars." *Minds and Machines* 28 (1): 141–72.

Woodward, James. 2011. "Psychological Studies of Causal and Counterfactual Reasoning." In *Understanding Counterfactuals, Understanding Causation: Issues in Philosophy and Psychology*, ed. Christoph Hoerl, Teresa McCormack, and Sarah R. Beck, 16–53. Oxford: Oxford University Press.