

Iterative and directed exploration of self-structuring embodied agents

Toby Ralphs Howison

Department of Engineering
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Declaration

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. 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 relevant Degree Committee.

Toby Ralphs Howison
May 2021

Abstract

Title: Iterative and directed exploration of self-structuring embodied agents

Author: Toby Ralphs Howison

Abstract: Roboticists increasingly look to biological systems for inspiration when designing and improving robotic systems. Simultaneously, building bio-inspired robotic systems can aid in the understanding of fundamental biological principles. In this context, the concept of *embodied intelligence* hypothesises that intelligent behaviours are the result of complex interactions between the brain, body (or morphology) and environment, rather than being driven purely by computational power in the brain. At the most basic level, embodied intelligence is driven by real-world physical interactions. By harnessing these interactions, unconventional ‘brainless’ robotic systems have demonstrated complex behaviours driven purely by passive interactions.

This thesis explores how complex behaviours emerge from interactions in two different low-level physical systems: falling paper and Bernoulli-balls. In falling paper systems, different paper shapes exhibit a range of behaviours when released into free fall. By altering morphological properties such as shape and weight, different behavioural modes can be triggered. In Bernoulli-ball systems, a ball is placed into a vertical airflow. If the morphological properties of the ball, for example size and density, and the environmental properties of the airflow, for example speed and width, are combined appropriately, the ball exhibits self-stable hovering within the airflow.

In Part I, I investigate falling paper systems. I introduce the novel V-shaped falling paper system. The relationship between morphology and system behaviours is explored and a data-driven modelling approach is developed to understand this. I explore the nature of behaviour transitions in the system. Certain behaviour transitions appear random, while others are more deterministic, and this variability is linked to morphology. Different methods are developed to represent this. I investigate generalised falling paper systems via the development of an automatic experimental platform capable of fabricating, dropping, observing and modelling hundreds of different paper shapes. Since falling paper systems are challenging to model using conventional methods, combining a data-driven approach with automatic experimentation is powerful.

In Part II, I investigate Bernoulli-ball systems. I explore the behaviour of a single Bernoulli-ball. A reduced-order model is derived to represent the main dynamics, and a minimalistic control policy is developed to modulate the ball hovering height by changing the airflow properties. I introduce the novel concept of a *collective* Bernoulli-ball: multiple hovering balls in a single airflow. This collective system exhibits a range of agent- and population-level behaviours, and these are investigated. The stability of, and relationship between, different behaviours is shown to be dependent on the balloon morphology and the environmental properties of the airflow.

In summary, the work in this thesis relates to the emergence of non-trivial behaviours from low-level embodied physical systems. The main contributions are the investigation of novel dynamics in these systems and the development of methods for understanding, representation and design. Ultimately, the work represents a small step toward the goal of creating artificial lifeforms with increasingly complex behaviours.

Acknowledgements

I would like to thank my supervisor Dr Fumiya Iida for all his help and advice with this PhD. I must also thank the past and present members of BIRL, particularly Dr Josie Hughes and Dr Fabio Giardina, whose help has been invaluable. I thank the EPSRC and The Mathworks for their financial support. I also thank my thesis examiners, Professor Josh Bongard and Professor Rodolphe Sepulchre. Finally, I would like to thank my family and friends, who have supported me throughout this PhD.

Preface

Chapters 2–7 of this thesis are based on seven peer-reviewed publications. Chapter 8 is based on a publication currently under review. Each publication has been edited to fit the narrative of this thesis; however, some overlap may occur. The contribution of each author is summarised at the beginning of the corresponding chapters. The publications are:

- **Chapter 2:** Howison, T., Hauser, S., Hughes, J., & Iida, F. (2021). Reality-assisted evolution of soft robots through large-scale physical experimentation: a review. *Artificial Life*, 26(4), 484–506.
- **Chapter 3:** Howison, T., Hughes, J., Giardina, F., & Iida, F. (2019). Physics driven behavioural clustering of free-falling paper shapes. *PLoS One*, 14(6), e0217997.
- **Chapter 4:** Howison, T., Hughes, J., & Iida, F. (2021). Morphological sensitivity and falling behaviours of paper V-shapes. *Artificial Life* [Forthcoming: accepted May 2021].
- **Chapter 5:** Howison, T., Hughes, J., & Iida, F. (2020). Morphologically programming the interactions of V-shaped falling papers. In *Artificial Life Conference Proceedings*, pages 359–366.
- **Chapter 6:** Howison, T., Hughes, J., & Iida, F. (2020). Large-scale automated investigation of free-falling paper shapes via iterative physical experimentation. *Nature Machine Intelligence*, 2(1), 68–75.
- **Chapters 3–6:** Howison, T., & Iida, F. (2020). Automatically designing the behaviours of falling paper: the emergence of non-trivial behaviours via interaction with the physical world. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion (GECCO)*, pages 1385–1386.
- **Chapter 7:** Howison, T., Giardina, F., & Iida, F. (2020). Augmenting self-Stability: height control of a Bernoulli-ball via bang-bang control. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3974–3980.

- **Chapter 8:** Howison, T., Crisp, H., Hauser, S. & Iida, F. (2020). On the stability and behavioural diversity of single and collective Bernoulli-balls. [under review as of 01/2022].

Throughout this thesis I try, as far as is practicable, to use the following pronoun convention. For work carried out solely by me, discussion that is my own, and general structural prose I use the pronoun ‘I’. For work and analysis completed in collaboration with my publication co-authors I use the pronoun ‘we’. I also use ‘we’ when discussing results with the reader.

Table of contents

List of figures	1
List of tables	5
1 Introduction	7
1.1 Research framework	9
1.1.1 Self-structuring embodied agents	9
1.1.2 System parameters	10
1.1.3 System observations	11
1.1.4 Iterative and directed exploration	11
1.2 Target systems	12
1.2.1 Falling paper	12
1.2.2 Bernoulli-balls	13
1.2.3 Underlying principles	14
1.3 Thesis objectives	14
1.4 Thesis contributions and structure	15
1.4.1 Conceptual	15
1.4.2 Novel dynamics	16
1.4.3 Methodological	18
2 Reality-assisted evolution for embodied intelligence	21
2.1 Embodied intelligence and soft robots	26
2.2 Modelling and representation	27
2.2.1 Conventional modelling	27
2.2.2 Physics engines and simulators	28
2.2.3 Data-driven modelling	29
2.2.4 Transferability, robustness and sim2Real	29
2.3 Large-scale physical experimentation	30

2.3.1	Materials and fabrication	30
2.3.2	Sensing and actuation	31
2.3.3	The fabrication gap	32
2.3.4	Scaling up physical experimentation	32
2.3.5	Modular and distributed systems	33
2.4	The search for novelty	34
2.4.1	Bio-inspired encoding	34
2.4.2	Optimisation algorithms	36
2.4.3	Quality-diversity and illumination algorithms	37
2.4.4	Developmental robotics	38
2.5	Discussion and conclusion	38
I	Falling paper	41
3	Physics-driven behavioural clustering of free-falling paper shapes	43
3.1	Falling paper and related problems	45
3.1.1	Flight, insects, wings	47
3.1.2	Seeds	47
3.1.3	Snowflakes	47
3.2	Physics-driven behavioural clustering	48
3.2.1	Data acquisition and processing	48
3.2.2	Formulation of generic PDQs (1)	50
3.2.3	Exponent search policy (2)	51
3.2.4	PDQ clustering (3–4)	52
3.2.5	Evaluation of clustering performance (5)	52
3.2.6	Optimal exponent vector selection (6-7)	53
3.3	V-shaped falling paper system	54
3.3.1	System morphology	54
3.3.2	System behaviours	54
3.3.3	System outputs	54
3.3.4	Manufacturing	56
3.3.5	Testing	56
3.4	VSFP experimental results	57
3.4.1	Behavioural diversity and structure	59
3.4.2	Behavioural outputs	60
3.5	PDBC results	61

3.5.1	Optimal exponent vector	61
3.5.2	Solution landscape	62
3.5.3	Physical significance	64
3.6	Discussion and conclusion	66
4	Transition behaviours in the V-shaped falling paper system	69
4.1	Stochastic behaviour transitions	70
4.1.1	Stochastic transition sequences	71
4.2	Deterministic behaviour switching	76
4.2.1	Rotational experimental setup	77
4.2.2	Behavioural hysteresis	79
4.2.3	Permanent behaviour changes	80
4.2.4	Behaviour transitions without hysteresis	81
4.2.5	Sequential logic representation	82
4.3	Discussion and conclusion	83
5	Bayesian optimisation of falling paper	87
5.1	Bayesian optimisation	88
5.2	Optimisation results	90
5.3	Discussion and conclusion	90
6	Large-scale automated investigation of free-falling paper shapes via iterative physical experimentation	93
6.1	Methods	94
6.1.1	Iterative physical experimentation system	95
6.1.2	Fabrication	95
6.1.3	Experimentation and data capture	97
6.1.4	Data analysis	99
6.1.5	Searching and exploration	102
6.2	Results	103
6.2.1	Effect of morphological parameters on observable characteristics	106
6.2.2	Behavioural grouping	108
6.3	Discussion and conclusion	110
II	Bernoulli-balls	113
7	Height control of a Bernoulli-ball via bang-bang control	115

7.1	Model and control	117
7.1.1	Dynamics modelling	117
7.1.2	Parameter identification	120
7.2	Global asymptotic stability	122
7.3	Augmented self-stability	126
7.3.1	Bang-bang control algorithm	126
7.3.2	Simulation analysis of idealised system	129
7.3.3	Simulation analysis of latency-included system	129
7.4	Experimental results	130
7.4.1	Experimental set-up	130
7.4.2	Dynamics model validation	131
7.4.3	Control demonstration	132
7.4.4	Analysis	136
7.5	Additional behaviours: bifurcations via environmental modulation	138
7.6	Discussion and conclusion	141
8	Stability and behavioural diversity of collective Bernoulli-balls	143
8.1	Experimental setup	145
8.2	Behaviours of collective Bernoulli-balls	146
8.2.1	Behavioural primitives	146
8.2.2	Automatic behavioural classification	147
8.2.3	Behaviours in isotropic populations	148
8.2.4	Stochastic behaviour transitions	151
8.2.5	Behaviours in anisotropic populations	154
8.3	Altering other mechanical properties	155
8.4	Discussion and conclusion	158
III	Discussion and conclusion	161
9	Discussion and conclusion	163
9.1	Discussion of contributions	163
9.1.1	Thesis objective one: systematic exploration of dynamics in novel self-structuring embodied agents	163
9.1.2	Thesis objective two: development of methods for iterative and directed exploration	164
9.2	Outlook and future work	166

9.2.1 Final remarks 167

References

List of figures

1.1	Thesis framework: iterative and directed exploration of self-structuring embodied agents	9
2.1	Reality-assisted evolution framework	24
3.1	PDBC algorithm flow chart	49
3.2	Parameterised design of VSFP system	55
3.3	Snapshots of falling behaviours in VSFP system	55
3.4	Schematic of behavioural modes	56
3.5	VSFP experimental setup	57
3.6	VSFP experimental procedure	58
3.7	Examples of VSFP morphologies	59
3.8	Dominant behaviours across the VSFP design space	60
3.9	PDBC solutions for the VSFP system	63
3.10	PDBC solution landscape for VSFP system	65
4.1	Characterisation of VSFP system to steady and transitional behaviours	70
4.2	Behaviour transition snapshots from VSFP system	71
4.3	Behaviour switching characteristics across design parameter space	73
4.4	Highlighted behaviour switching diagrams	74
4.5	Transition diagrams across morphological parameter space	75
4.6	Stochastic switching diagram for full VSFP system	76
4.7	Rotational experimental setup	78
4.8	Behaviour hysteresis of VSFP system	81
4.9	Helicopter to undulating behavioural transitions	82
4.10	Sequential logic circuit representation of falling behaviours	83
5.1	Bayesian optimisation of VSFP system	91
5.2	Bayesian optimisation sampling strategy	92

5.3	Comparison of optimised performance to natural systems	92
6.1	Schematic of iterative physical experimentation approach	95
6.2	Diagram of the experimental setup	96
6.3	Block diagram of experimental setup	97
6.4	Plan view schematic of the laser cutter and paper feed system	97
6.5	Flow chart of system operation	98
6.6	Distribution of randomly sampled design parameters for each shape	102
6.7	Representative sample of manufactured shapes	104
6.8	Confusion matrices showing the agreement with human observers when classifying behaviours	105
6.9	Time-lapse images of the three falling behaviours observed in each shape .	106
6.10	Trajectories and observable area profiles of falling shapes for different auto- matically classified behavioural groups	107
6.11	Variation in measured falling characteristics with respect to design parame- ters and automatically classified behaviour groups	109
6.12	Effect of initial condition on horizontal displacement	110
6.13	Automatically classified falling behaviours in the Reynolds number Re and non-dimensional moment of inertia I^* parameter space	111
7.1	Experimental set-up for hovering a sphere in a vertical airflow	118
7.2	Airflow model	119
7.3	Experimentally measured mean flow velocity	122
7.4	Experimentally measured mean outlet nozzle velocity V_0 as a function of converged control demand u^* compared with the fitted model.	123
7.5	Mean steady state hovering height z_{mean} as a function of converged control demand u^*	124
7.6	Phase portrait and sample trajectory of self-stable vertical dynamics	127
7.7	Example of a bang-bang control policy	128
7.8	Bang-bang control switching curves	130
7.9	Schematic of experimental architecture	131
7.10	Example trajectory of dynamics model and real-world system	133
7.11	Controlled and open-loop response of real-world system	135
7.12	Snapshot images of controlled system	136
7.13	Controlled and open-loop response to consecutive trajectory following . . .	137
7.14	Model and experimental control cost function	138

7.15	Bifurcation behaviours of a hovering ping-pong ball in an airflow oscillating between a high and low power	140
7.16	Fourier fast transform of hovering ping-pong ball	141
8.1	Collective Bernoulli-ball experimental setup	145
8.2	Behavioural primitives in one, two and three ball collective Bernoulli-ball systems	147
8.3	Time series data and automatically classified behaviours for a two minute snapshot of experimental data	149
8.4	Mean corrected height for two Bernoulli-balls	150
8.5	Average proportion of time spent in different behavioural modes across different population sizes and airflows	151
8.6	DTMC behavioural switching diagrams for isotropic populations	153
8.7	Time series data and DTMC switching diagrams for anisotropic population	155
8.8	Extending behaviours by increasing morphology and environmental complexity using two fans and an adhesive balloon	157
8.9	Time series of data of adhesion experiments	157

List of tables

2.1	Highlighted literature for model-based, model-free and reality-assisted design of physically embodied robots	25
3.1	PDBC algorithm input parameters	50
3.2	Suggested exponent ranges and increments for PDBC algorithm	52
3.3	Qualitative description of behavioural modes in VSFP system	56
4.1	Switching probabilities in VSFP system	75
4.2	Logical representation of behaviours in rotational experimental system. . .	83
6.1	List and description of extracted features	100
6.2	List of shape parameters and their corresponding ranges	103
6.3	Number of experiments carried out for each shape	103
6.4	Classification agreement with human observers for each shape	105
7.1	Measured system parameters	121
7.2	Unknown system parameters.	121
7.3	Comparison of model and reality	132
7.4	Open-loop and controlled settling time and improvement factor	137

Chapter 1

Introduction

From simple cellular organisms to the most complex lifeforms, biological systems exhibit extraordinary levels of diversity and adaptability. In contrast, state-of-the-art robotic systems are generally considered to be significantly limited in the scope of their behavioural diversity (Moravec, 1988). Roboticists are increasingly looking to biological systems as a source of inspiration for designing better robotic systems (Iida and Ijspeert, 2016). In parallel, building robots inspired by biological principles can help us to understand some of the fundamental forces that govern natural living systems (Taylor and Jefferson, 1993).

An increasingly popular framework for understanding and characterising the complex behavioural diversity seen in biological systems is that of *embodied intelligence* (Clark, 2008; Pfeifer and Bongard, 2006; Pfeifer et al., 2007). Simply put, the perspective hypothesises that biological intelligence can only be fully understood, and replicated, by acknowledging the fact that an agent is always embodied within a physical body and embedded within an environment. Hence, rather than being driven purely by computational power in the brain, intelligent behaviours are the result of the complex coupling between the brain, body and environment. This surprisingly simple concept offers an elegant explanation for the limitations of many modern robotics technologies. A popular example is legged locomotion. On the one hand, the well known ASIMO humanoid robot (Sakagami et al., 2002; Shigemi et al., 2018) uses a complex control system and a range of actuated joints to achieve a range of walking behaviours. Yet, its walking appears unnatural and inefficient (of course, the ASIMO has many other functionalities on top of locomotion). On the other hand is the so-called passive dynamic walker (PDW) (Collins et al., 2005; McGeer et al., 1990), a bipedal robot with a specific configuration of passive joints and rigid links, but no power source or controller. When placed on an inclined surface the robot proceeds to walk downhill with an incredibly life-like gait. Of course, this behaviour is reliant on a very specific environmental configuration. If the PDW is placed on a flat surface it will show no locomotion behaviours,

and will do nothing. Nevertheless, by using a well designed morphology and carefully structured environmental interactions, the PDW demonstrates that behaviours commonly associated with brain computation can actually emerge from a purely physical system.

Current research into embodied intelligence is rooted in a long history of exploring the fundamental principles of biological and man-made systems over the last century. In the mid 1900s, the field of cybernetics (Heylighen and Joslyn, 2001; Wiener, 1965) arose at the intersection between disciplines such as mathematics, neuroscience, and engineering. Driven in part by the accelerating development of computer science, cybernetics focuses primarily on understanding regulatory feedback in living and man-made systems. A particularly relevant concept is the idea of *ultrastability* (Ashby, 1961), a property of a system that allows it to change its internal structure in response to environmental stimuli. Natural systems are ultrastable both at the individual level, e.g. agents can adapt to their environment, and at the evolutionary level, e.g. species evolve as the evolutionary pressures change. The concept of ultrastability is analogous with the idea of embodied physical interactions being used as a stabilising property in man-made systems. Cybernetics researchers demonstrated the idea of ultrastability in practice with the *homeostat* (Ashby, 1960), a mechanical device built with the single purpose of regaining stability when perturbed by the environment. The homeostat was one of the first demonstrations of how a relatively simple set of physical components and electrical circuits could be configured to exhibit feedback mechanisms that reliably drove the system toward equilibrium. More recently, the work of Rodney Brooks in the 1980's and 90's was highly influential (Brooks, 1990). In contrast to the prevailing view of classical AI at the time that increasingly complex computational power was the most effective approach for designing intelligent systems, Brooks proposed a simpler behaviour-based approach. A key concept was that high-level intelligence emerges from the interactions and behaviours of smaller elements in a system, which are, in their own right, intelligent. In parallel, the physical grounding hypothesis was proposed, stating that intelligent systems must be embedded in a physical environment and learn, in a bottom-up manner, the causal relationships between perception and action (Brooks, 1991).

Embodied intelligence is particularly relevant to the fields of *artificial life* (Steels and Brooks, 2018) and *soft robotics* (Iida and Laschi, 2011). Motivated by the goal of investigating, understanding and ultimately creating life-like systems, artificial life research is often characterised by a rejection of conventional approaches to robot design and control. Instead, concepts such as self-organisation (Pfeifer et al., 2007), emergence and open-ended evolution (Eiben et al., 2012) are used to explain and design for the complexity we see in the natural world. Meanwhile, research into soft robotics focuses on how soft materials can be used to improve the performance of robotic systems. Unlike conventional robots, whose basis

lies in rigid body mechanics (Angeles, 2002), soft robots are defined by highly deformable materials and sustained environmental interaction. By embracing this non-linearity, soft roboticists often aim to outsource computational and control tasks to the intrinsic properties of soft materials and their environmental interactions. The research in this thesis lies at the intersection of these two fields.

1.1 Research framework

In this thesis I propose a framework of *iterative and directed exploration of self-structuring embodied agents* (see Figure 1.1) to investigate aspects of embodied intelligence in physical systems.

1.1.1 Self-structuring embodied agents

As discussed, the concept of embodied intelligence is based on the coupling of the brain, body and environment. We can conclude that, at the most basic level, embodied intelligence is driven by real-world physical interactions. Hence, by exploring how structure emerges

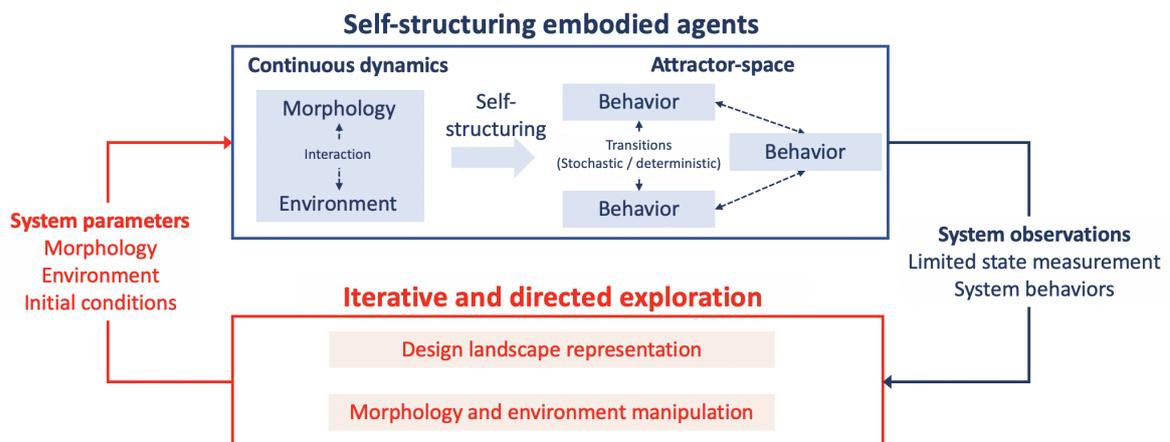


Figure 1.1: Thesis framework: iterative and directed exploration of self-structuring embodied agents. Through morphology–environment interaction, embodied agents self-structure their continuous interaction-driven dynamics into an attractor space from which discrete behaviours emerge. From the system we can observe a limited amount of state measurements, but typically much of the dynamics are hidden from observation. We can also assign behaviours to different patterns in the system. System behaviours capture the full dynamics in a reduced-order symbolic representation. System parameters have an effect on these observations. Morphological and environmental parameters can be altered to change how the discretised attractor space is structured. Initial conditions are also important in determining short- and long-term behaviours. We can influence the emergence of different behaviour mechanisms via iterative (i.e. repeating toward more complex behaviours) and directed (i.e. with some goal in mind) exploration).

from low-level physical systems we can build up our understanding of embodied intelligence from the bottom up.

I focus on the low-level interactions of *self-structuring* embodied agents, i.e. low-level physical systems that exhibit complex and seemingly structured behaviours. Such systems are initially characterised in terms of the continuous dynamics generated via the interaction between morphology and the environment. These continuous dynamics may contain a structured attractor space from which discrete *behaviours* emerge. A system can exhibit multiple behaviours, and can transition between them predictably (deterministic) or unpredictably (apparently stochastic). The concept of self-structuring, therefore, refers to the phenomena that the continuous dynamics of a system are structured into discrete behaviours through a discretised attractor space, purely via morphology–environment interactions. It is, of course, also possible that agents are not structured, and the system remains in a random or chaotic state. The existence of a structured attractor space in which basins of attraction correspond to stable behaviours is thought to be a key requirement for embodied intelligence (Pfeifer and Bongard, 2006).

An elegant example of a system with these properties is that of chemical oil droplets (Horibe et al., 2011). In this system a single oil droplet is placed floating on a base liquid. Driven by the chemical reaction with the surrounding fluid, the droplet moves over time. This self-movement is driven purely by the interaction between morphology and environment, i.e., there is no external inputs. Four distinct behavioural modes emerge from the continuous dynamics of the droplet self-movement (fluctuating, circular, direction and vibrating). The system exhibits stochastic switching between the basic behaviour modalities. Hence, from the seemingly simple system of a floating oil droplet, continuous dynamics and discrete behaviours and transitions emerge. It can even be argued that these behaviours form the basis for minimal perception cognition (Hanczyc and Ikegami, 2010).

1.1.2 System parameters

We may induce different behaviours in an embodied agent by manipulating the parameters that define it. Morphological properties such as shape, mass and constituent materials can be varied. Alternatively, and just as importantly, the environmental conditions in which the morphology is embedded can be manipulated. When describing the environment, we need to clearly delineate where the environment of interest begins and ends, and what features this environment includes. For example, in the case the PDW the environment extends to the incline on which it walks and the surface properties of the ground, but not necessarily to the fluid properties of the air in which it moves. Initial conditions can also be varied, and this may have a significant impact over the evolution of observed behaviours. For example,

a system could be initialised such that it is likely to converge to a certain attractor state, or with initial states that lead to stochastic switching between attractors.

In general, the number of available system parameters is far lower than the degrees-of-freedom in the realised dynamics, making the parameters-to-dynamics mapping severely underdetermined. Further still, for many systems we may only partially understand the mechanisms driving behaviour. From a negative perspective this means that we are only in control, or even aware, of a small fraction of the total degrees-of-freedom in the system. From a positive perspective, however, we see that by selecting appropriate system parameters we can modulate highly complex dynamics with a minimal set of input variables. For example, the behaviours of interacting oil droplets are driven by a complex set of chemical reactions that would require a huge number of states to be described. However, by varying a single parameter, such as droplet size, the structure and emergence of discrete patterns in the system can be modulated easily.

1.1.3 System observations

The observation of self-structuring embodied agents is challenging. As described in 1.1.2, much of the dynamics can be highly non-linear, can have large degrees of freedom or are completely unknown. Hence, full state measurement can be difficult and we can only rely on a *limited* amount of state measurement. For example, the only states easily measurable in the interacting oil droplet system are the position, velocity and size of each droplet over time. There are multiple hidden states driving the underlying dynamics that cause these measurable states to change.

While direct state observation can be limited, we also have the option of observing the behaviours that emerge through the discretised attractor space. In doing so, the full system dynamics can be abstracted to the level of a small number of behavioural modes, offering a symbolic representation of the system that is intrinsically driven by its dynamics, but does not require a full observation or understanding of these dynamics. However, observing behaviours involves an element of subjectivity insofar that behaviours must be defined. Different behavioural definitions may have a significant influence on the quality of analysis and interpretations of the system.

1.1.4 Iterative and directed exploration

Understanding, exploring and ultimately designing self-structuring embodied agents requires iterative and directed feedback mechanisms to link system parameters and system observations. A key component of this feedback is representation of the design landscape. In many cases, the dynamics cannot be effectively captured in a simulation environment so we must rely on representations generated from real-world observations. As discussed in

1.1.3, we cannot rely on complete state measurements. Hence, representation methods that make use of output behaviours are useful, for example behavioural clustering to categorise and interpret the different behaviours in the system. Self-structuring embodied agents may be stochastic or deterministic in their behaviours, so different approaches for modelling and understanding are required. Representation may also take the form of a fitness metric that allows us to compare different agents with each other. A second key component is manipulation of the morphology and environment. By changing certain system parameters we can capture diverse dynamics that aid in our understanding of the system. Morphology is the conventional target for design-space exploration, for example evolving morphologies for locomotion (Rieffel et al., 2014; Vujovic et al., 2017). Since self-structuring embodied agents are driven by their interactions, it is also vital to explore how environmental manipulation influences the system (Auerbach and Bongard, 2012). Representations can be iteratively updated as more real-world evaluations are carried out.

1.2 Target systems

In this thesis, I investigate two low-level physical systems that can be characterised as self-structuring embodied agents, and that require iterative and directed exploration in order to understand their behaviours: falling paper and Bernoulli-balls. This section introduces these systems and explains their relevance to the research framework (Figure 1.1).

1.2.1 Falling paper

The first target system in this thesis is falling paper. First proposed by James Clerk Maxwell in the 1800s (Maxwell, 1854), the ‘falling-paper problem’ is to understand the rich variety of behaviours exhibited by simple laminar shapes (made of paper or other materials) as they fall freely through a fluid. Though conceptually a very simple system, falling paper manifests a range of interesting properties. Paper as a material is easily defined by its mechanical properties such as density, elasticity and shape. However, when released into the environment, highly complex fluid-mechanics interactions are induced between the paper and air. Depending on morphological and environmental factors, the style in which paper falls varies considerably. The stochastic nature of this behavioural emergence means that we cannot *guarantee* certain behaviours. Instead, we see certain attractor states to which shapes are highly likely to converge. We can introduce almost infinite richness into the environment by simply adding various airflows. Hence, we can explore the relationship between environmental attractor states and behavioural emergence. Paper shapes can be fabricated, dropped and observed with minimal equipment. Synthetic methodologies are not

only feasible, but preferable due to the difficulty in modelling Pfeifer and Bongard (2006). This makes it an ideal system for studying *real-world* robot development and evolution.

In the well-studied case of falling disks four behaviours (steady falling, periodic oscillation, tumbling and chaotic falling) are observed as system properties such as diameter, density and fluid viscosity are varied (Field et al., 1997; Lee et al., 2013; Stringham et al., 1969; Willmarth et al., 1964; Zhong et al., 2011, 2013). Other simple shapes have been studied, including rectangles (Mahadevan et al., 1999; Skews, 1990; Wang et al., 2013), annuli (Vincent et al., 2016) and parallelograms (Varshney et al., 2013); behavioural transitions have also been studied (Andersen et al., 2005; Belmonte et al., 1998; Fernandes et al., 2005). A number of heavily simplified numerical solutions have also been presented (Jin and Xu, 2008; Pesavento and Wang, 2004); however, the problem has evaded a comprehensive solution. As a result, many studies focus on using experimental behavioural observations to understand the driving physical phenomena. These approaches tend to characterise behaviours using dimensionless quantities such as Reynolds number Re or dimensionless moment of inertia I^* (Belmonte et al., 1998; Chrust et al., 2013; Field et al., 1997; Mahadevan et al., 1999; Mittal et al., 2004). Using these quantities allows the construction of a dimensionless parameter space in which different regions correspond to different falling behaviours. Similar approaches using other dimensionless quantities such as the Froude Fr or Strouhal St numbers have been used in the analysis of behavioural diversity in other systems (Goto and Tanaka, 2015; Iversen, 1979; Lauga and Nadal, 2017; Smith, 1971; Willmarth et al., 1964).

1.2.2 Bernoulli-balls

The second target system is the so-called Bernoulli-ball system, an elegant fluid dynamics phenomena in which spherical objects self-stabilise and hover in an airflow. The process, which works for many shaped objects, is driven by fluidic interactions between the ball and its environment (Anderson and Eberhardt, 2001; Kamela, 2007; Waltham et al., 2003). The Bernoulli-ball exhibits two forms of self-stabilising behaviours. First, it is self-stable in the horizontal plane: when perturbed horizontally the ball returns to its original position. This is commonly explained using Bernoulli's principle, which states that as the speed of a fluid increases, the fluid pressure decreases. Hence, the pressure within an airflow is lower than the surrounding environment, creating a self-stabilising force about the jet centerline. Second, it is self-stable in the vertical direction: when released into the airflow, the ball will eventually settle around a nominal height. This is due to the balance between drag forces and the ball mass, and the dissipative effect of moving in the airflow. Bernoulli-ball systems have a number of interesting features in the context of self-structuring embodied agents. We can scale up complexity in the system with ease. For example, properties such as

ball mass and surface properties can be changed. Similarly, the environmental conditions can be modulated easily by changing the type of airflow in which the balls hover to create complex and anisotropic conditions. Multiple balls can be added to the same system, creating a *collective* Bernoulli-ball system with complex population-level behaviours that are driven by the basic dynamics of agents interacting with the environment and each other.

There has been surprisingly little research into Bernoulli-ball systems (Nudehi et al., 2017). However, there is a rich history of investigating behavioural diversity in similar systems. For example, significant research has looked at the bouncing behaviours of balls, for example on a sinusoidally vibrating table (Holmes, 1982). Similar work has been carried out into the ‘blind-juggler’, a platform that uses a parabolic dish to indefinitely bounce a ball without any feedback control (hence, blind) (Reist and D’Andrea, 2012; Ronsse et al., 2007). A key feature of these studies is that, depending on the conditions of the environment and the ball, different behaviours could be induced such as stability, period doubling and chaotic modalities. Collective systems have also been researched. Simulations of hard spheres have shown how mixtures of large and small spheres will usually exhibit spatial clustering based on morphology (Dickman et al., 1997). Similar results have been shown in systems of floating objects (Nakajima et al., 2012).

1.2.3 Underlying principles

At the basic level, these systems are interesting to us because their behaviours are driven by their underlying Newtonian physics. This is also a key feature of the PDW, which is driven by gravitational potential coupled with a well constrained leg morphology. The falling paper and Bernoulli-ball system are similarly driven by gravitational forces, yet have behaviours that are far less constrained than PDWs. Falling paper behaviours are driven by morphological-environment interactions that we can primarily modulate by changing morphology, e.g. the paper shape. The amount of complexity we can easily add to the system via environmental changes is minimal before the system behaviours become too chaotic to systematically study. On the other hand, behaviours in the Bernoulli-ball are primarily modulated by changing environmental conditions, since the complexity we can easily add to the morphological structure is fairly limited. Hence, these two systems allow us to explore morphology-environment interactions from two different perspectives.

1.3 Thesis objectives

In this thesis I use a bottom-up approach to explore the behaviours of falling paper and Bernoulli-ball systems. There are two main objectives that this work aims to achieve:

1. **Systematic exploration of dynamics in novel self-structuring embodied agents.** To systematically explore dynamics in the falling paper and Bernoulli-ball systems from the perspective of self-structuring embodied agents.
2. **Development of methods for iterative and directed exploration.** To develop methodologies for the iterative and directed exploration of the behaviours in falling paper and Bernoulli-ball systems, including behavioural representation and design-space exploration.

1.4 Thesis contributions and structure

The main contributions of this thesis can be categorised as *conceptual* (new frameworks and ideas), *novel dynamics* (new examples of self-structuring embodied agents) and *methodological* (new methods, tools and hardware). Each contribution is linked to one or more chapters in which the bulk of the contribution can be found.

1.4.1 Conceptual

Reality-assisted evolution of soft robots through large-scale physical experimentation (Chapter 2)

Key contribution: literature review summarising developments in the design of physically embodied soft robots under the unifying framework of reality-assisted evolution.

Design approaches for robotic systems have long been dominated by two opposing ideologies. Model-based approaches that utilise simulation tools to evaluate system behaviours *in silico*, and model-free approaches in which physical experiments are conducted *in reality*. A key problem with model-based approaches is overcoming the so-called ‘reality-gap’: the disparity between simulated and realised behaviours (Mouret and Chatzilygeroudis, 2017). A key problem with model-free approaches is the time and cost constraints involved when constructing physical systems. These challenges have become particularly apparent in the design of soft robotic systems, which are highly challenging to simulate and generally costly to fabricate (Laschi et al., 2016).

In Chapter 2, I introduce the framework of *reality-assisted* evolution to summarise a growing trend towards combining *in silico* model-based and *in reality* model-free approaches to improve the design of physically embodied soft robots. *In silico*, data-driven models build, adapt and improve representations of the target system using real-world experimental data. By simulating huge numbers of virtual robots using these data-driven models, optimisation

algorithms can illuminate multiple design candidates for transference to the real world. *In reality*, large-scale physical experimentation facilitates the fabrication, testing and analysis of multiple candidate designs. Automated assembly and reconfigurable modular systems enable significantly higher numbers of real-world design evaluations than previously possible. Large volumes of ground-truth data gathered via physical experimentation can be returned to the virtual environment to improve data-driven models and guide optimisation. Grounding the design process in physical experimentation ensures that the complexity of virtual robot designs does not outpace the model limitations or available fabrication technologies.

1.4.2 Novel dynamics

V-shaped falling paper system (Chapters 3, 4 and 5)

Key contribution: analysis of behaviours in the V-shaped falling paper system. Dominant behaviours in the design-space are explored through systematic experiments. Stochastic and deterministic behavioural transitions are investigated and analysed.

As discussed in 1.2.1, falling paper systems exhibit a range of discrete behavioural modes driven by morphology-environment interaction. The range of systems found in the literature is generally limited to simple shapes with low levels of deformation. In this thesis I present the novel V-shaped falling-paper (VSFP) system, which contains a more complex morphology and exhibits large levels of deformation. It is a paper V-shape defined by two morphological parameters: wing angle θ and wing length l . It exhibits four distinct falling behaviours (plummeting, undulating, helicopter rotation and asymmetric rotation) whose emergence is strongly dependent on morphology. In Chapter 3, I introduce the system and demonstrate how dominant behaviours vary across the design space. In Chapter 4, I analyse the nature of behavioural transitions in the VSFP system. I investigate stochastic behaviour transitions, showing how morphology influences the likelihood of different transitions, with certain morphologies leading to a wide range of possible paths through the attractor space. Deterministic transitions are also investigated using a second experimental platform. I demonstrate how behaviour transitions can be induced by modulating the energy input to the system, and that the system has a basic form of memory. In Chapter 5, I demonstrate how morphology in the system can be optimised to minimise falling speed and compare the biological equivalent of the VSFP system, flying seeds.

Single Bernoulli-ball system (Chapter 7)

Key contribution: analysis of dynamics in a single Bernoulli-ball system. A reduced order model is developed and the system stability is analysed. A low-bandwidth bang-bang controller is used to demonstrate how height can be controlled in the system by altering the airflow characteristics.

Self-structuring embodied agents exhibit natural and reliable behaviours based on their environmental and morphological properties. In Chapter 7, I analyse the dynamics of a single Bernoulli-ball system and present a strategy for height control that utilises the system self-stability. In the single Bernoulli-ball system, once active we can only influence behaviour by altering the environment, e.g., the airflow properties. There is a significant amount of latency between a change in control input and the airflow settling to a new power. Additional latency is introduced during system observation, which in this case uses computer vision to detect the ball height. Given that, conventional feedback control is a poor choice for height control. Instead, a low-bandwidth minimalistic approach to control strategy that utilises model-based, feedforward bang-bang control at a global level and self-stabilising dynamics at a local level is proposed. A reduced-order model of the system is presented, the control algorithm is developed and hardware validation is performed.

Collective Bernoulli-ball system (Chapter 8)

Key contribution: analysis of the collective Bernoulli-ball system, showing how environmental and morphological factors influence behaviour in the system and drive stochastic behaviours at the agent and population level.

As discussed in 1.2.2, the complexity of Bernoulli-ball systems can be increased by adding additional balls to the airflow. In Chapter 8, I focus on the behaviour of these collective Bernoulli-ball systems. I demonstrate how behavioural diversity can be induced in the case of a single ball via modulation of the environment. I then show how more diverse behaviours are triggered in the collective system. I discuss this in the context of embodied intelligence and open-ended evolution, suggesting that the system exhibits a rudimentary form of evolutionary dynamics in which balls compete for favourable regions of the environment and exhibit intrinsic 'alive' and 'dead' states based on their position in or outside of the airflow.

1.4.3 Methodological

Physics-driven behavioural clustering (Chapter 3)

Key contribution: development of the physics-driven behavioural clustering algorithm. The algorithm generates a set of symbolic equations that define a parameter space to represent the design landscape of system parameters, state measurements and system behaviours.

As discussed in 1.1.4, a key challenge when exploring self-structuring embodied agents is representing the system dynamics, especially when these are highly non-linear and cannot be fully captured in simulation. There is a growing trend for using machine learning methods to generate symbolic equations based on real-world experimental data of dynamical systems. One of the most well known approaches (Schmidt and Lipson, 2009) can, for example, generate equations of motion for simple dynamical systems (e.g., a pendulum) and use this learnt knowledge to find equations for more complex systems. While highly effective, in many systems (including falling paper) the number of measurable states can be too low for effective equation generation.

In Chapter 3, I present physics-driven behavioural clustering (PDBC), a novel algorithm for generating physically meaningful equations for systems with discrete behavioural modes and limited state measurements. The method transforms the system inputs, (e.g., morphological parameters) and measurable system states (e.g., falling speed) into a set of equations that describe a parameter space in which the different system behaviours are represented by distinct regions. By testing large numbers of candidate equations, the method finds an optimal set of equations to represent the behavioural diversity in the system most distinctly. The PDBC algorithm is demonstrated on the V-shaped falling paper system, which exhibits four distinct behavioural modes that depend on a few morphological parameters. Using 49 experimental observations, the method discovered a set of candidate functions that distinguish behaviours with an error of 2.04%, while also aiding insight into the physical phenomena driving each behaviour.

Large-scale physical experimentation on falling paper (Chapter 6)

Key contribution: development of a hardware and software platform for large-scale physical experimentation on generalised falling paper systems, with validation against preexisting data on falling disks, and extension to other morphologies.

As discussed in 1.1.4, exploring the design space of self-structuring embodied agents is key for understanding how physical interactions shape behaviour, and for discovering new behaviours and features of the system. Exploration in simulation is easy, for example evaluating millions of soft robotic morphologies as part of an evolutionary algorithm. However, if the underlying real-world dynamics cannot be modelled and effectively simulated then we must rely on physical experimentation, i.e., a model-free or reality-assisted approach. In this context, a powerful approach is to utilise technologies such as robotic arm automation and 3D printing to scale up the number of real-world evaluations. Large-scale experimental investigations have been demonstrated a number of times, such as for the real-world evolution of modular robots (Brodbeck et al., 2015), and in other settings such as automated experimentation on fluid dynamics problems (Fan et al., 2019).

In Chapter 6, I present a hardware and software platform for large-scale physical experimentation on generalised falling paper systems. The platform combines robotic automation, computer vision and machine learning to autonomously fabricate, drop, analyse and classify the behaviours of hundreds of shapes. The system is validated by reproducing results for falling discs, which exhibit four falling styles: tumbling, chaotic, steady and periodic. A previously determined mapping from a non-dimensional parameter space to behaviour groups is shown to be consistent with these new experiments for tumbling and chaotic behaviours. However, steady or periodic behaviours are observed in previously unseen areas of the parameter space. More complex hexagon, square and cross shapes are investigated, showing that the non-dimensional parameter space generalises to these shapes. The system highlights the potential of robotics for the investigation of complex physical systems, of which falling paper is one example, and provides a template for future investigation of such systems.

Chapter 2

Reality-assisted evolution for embodied intelligence¹

From simple cellular organisms to the most complex lifeforms, biological systems exhibit extraordinary levels of diversity over a large spatio-temporal scale. The apparently coherent behavioural repertoires observed in nature are increasingly explained using an *embodied* view of intelligence that focuses on the closely coupled interactions between the brain, body and environment (Clark, 2008; Pfeifer et al., 2007). Understanding and harnessing these embodied interactions, therefore, is a key step in building truly lifelike artificial systems. In this context the field of *soft* robotics is particularly interesting. Soft robots can generally be defined as having highly deformable bodies or elements, often constructed using unconventional (in the robotics sense) materials with highly non-linear properties (Laschi et al., 2016; Nurzaman et al., 2014; Pfeifer et al., 2012; Rus and Tolley, 2015). They can be further characterised by their uniquely complex and sustained environmental interactions, for example in terrestrial or aquatic environments (Corucci et al., 2018). Given this, the methodologies on which conventional robotics science was developed – e.g., rigid-body kinematics and dynamics – are often inapplicable when applied to soft robots. Instead, significant research has focused on how to utilise the characteristics of soft robots in an embodied framework. By harnessing

¹This chapter is based on the following peer-reviewed publication:

- Howison, T., Hauser, S., Hughes, J., & Iida, F. (2021). Reality-assisted evolution of soft robots through large-scale physical experimentation: a review. *Artificial Life*, 26(4), 484–506.

Contributions

T. Howison – devised the main conceptual ideas, wrote the bulk of the paper and created the figures.

S. Hauser – contributed to writing of 2.3.5, general discussions and proofreading.

J. Hughes – contributed to writing of 2.2.2 and 2.3.1.

F. Iida – contributed to general discussions and proofreading.

the mechanical intelligence of soft bodies, researchers hope to develop robotic systems with comparable performance to biological systems (Kim et al., 2013).

While the promise of soft embodied intelligence and soft robotics is well documented, the design process by which to realise it is non-trivial (Howard et al., 2019). A classical process for designing robotic systems is through the use of models and simulations. *Model-based* approaches often demonstrate the design of virtual soft robots *in silico* via artificial evolution. By mimicking the processes from which embodied intelligence emerged in nature, these methods have successfully demonstrated complex, meaningful robot behaviours (Cheney et al., 2013). However, evolved robot designs frequently do not perform as expected when transferred into the real world: this is an aspect of the so-called reality gap (Mouret and Chatzilygeroudis, 2017). Further still, available fabrication technologies may prevent the realisation of virtually evolved robot body-plans altogether. Alternatively, *model-free* approaches restrict design to the physical world, for example via top-down imitation of biological systems such as octopus tentacles or elephant trunks (Nakajima et al., 2015) or using real-world evolutionary approaches (Vujovic et al., 2017). By evaluating robots *in reality* the problems associated with transference between models and the real world can be avoided. However, physical experimentation is generally time-consuming and resource intensive, restricting the effective exploration of large design spaces.

In this chapter I summarise recent developments in the design of soft robots from the perspective of a rapidly growing direction I term *reality-assisted* evolution. There has been a growing interest in design methodologies that combine model-based and model-free approaches to improve the overall design of physically embodied robots (see Table 2.1). Under this framework, as summarised in Figure 2.1, simulation-only and simulation-free design methodologies are unified via large-scale physical experimentation and data-driven modelling. *In silico*, data-driven modelling is used to build and improve representations of the target system based on real-world experimental data. Physics engines, software-based approximations of physical systems, can be tuned to maximise their predictive accuracy, and auxiliary models can be built to estimate the transferability of virtual robots to the real world (Koos et al., 2012). Alternatively, models can be built from scratch to estimate fitness across the design parameter space (Saar et al., 2018). These data-driven models can be used to simulate the behaviour of huge numbers of virtual robots and, in conjunction with optimisation and illumination (Mouret and Clune, 2015) algorithms, to discover diverse and highly performant designs that are likely to survive crossing the reality gap. *In reality*, large-scale physical experimentation systems can fabricate and test multiple robot designs. Recent developments in automated and scalable fabrication approaches facilitate an increasingly large number of real-world evaluations, e.g., (Kriegman et al., 2020b). Similarly, modular

and adaptive robots allow multiple morphologies and controllers to be tested on a single reconfigurable platform (Nygaard et al., 2018; Veenstra et al., 2018). By situating physically embodied robots within a task environment, significant amounts of useful experimental data can be gathered. Grounding the design process in physical experimentation ensures the complexity of simulated robot designs does not outpace the model accuracy or available fabrication technologies. Further still, modulating complexity in the task environment *in reality* can drive the heterogeneous improvement of data-driven models *in silico* from simple to more challenging design representations.

This chapter is structured as follows. In 2.1 I outline the open problems in embodied intelligence for soft robots. Three classes of soft robot that typify embodied intelligence in soft robots and drive the need for a reality-assisted framework are highlighted. In 2.2 I discuss advances in modelling and representation for soft robotics, starting from conventional approaches and finishing with data-driven modelling and transferability methods. In 2.3 I discuss advances in large-scale physical experimentation. It covers the range of possible fabrication methods before exploring how these can be scaled up, for example via robotic automation or modular and distributed systems. In 2.4 I discuss optimisation methods for designing soft robots. It explores design encoding and typical optimisation algorithms, then explore how novelty based search methods offer a more robust approach within the reality-assisted framework. In 2.5 I conclude the chapter.

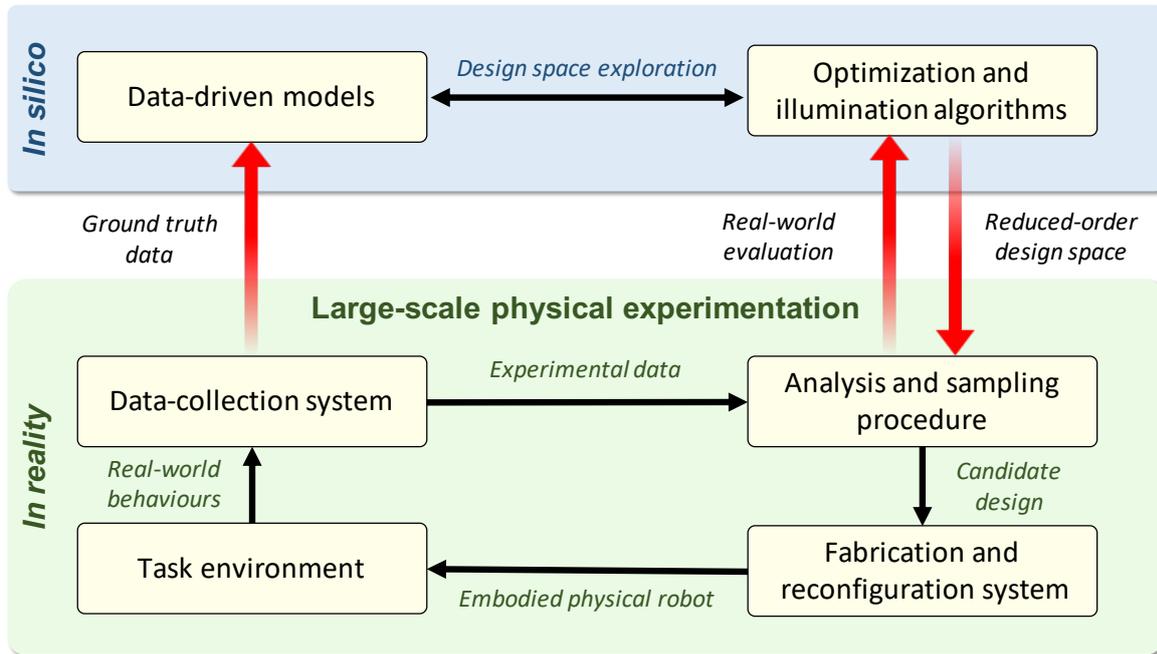


Figure 2.1: The reality-assisted evolution framework for designing physically embodied soft robots. *In silico*, data-driven models are constructed and adapted based on ground-truth data. Physics engines, software-based approximations of physical systems, can be tuned to maximise their predictive accuracy. Alternatively, machine learning can be used to build models from scratch. Auxiliary models can also be built to estimate the performance disparity of virtual robots after transference to the real world. By using data-driven models to simulate huge numbers of possible robot designs, optimisation and illumination algorithms can discover a range of diverse and highly performant designs across the parameter space. A reduced-order design space of promising candidates can be established for transference to the real-world. *In reality*, large-scale physical experimentation facilitates the fabrication and analysis of multiple candidate designs. Increasingly automated experimental platforms can be fabricated and test large numbers of candidate designs without human input. Alternatively, new modular and adaptive robotic systems facilitate the testing of multiple robot designs on one reconfigurable platform. By placing physical robots in a task environment and observing their behaviours, large volumes of useful experimental data can be gathered. This data can inform the sampling procedure for physical experimentation *in reality*, and can be used to update data-driven models and as an input to optimisation algorithms *in silico*. From the bottom up, changes to the physical task environment drive the gradual complexification of data-driven models and robot designs toward more complex embodiments and behaviours.

Author Reference	Target System	Parameter Space	Model / Simulation	Search Method	Realised Designs:	
					Morphologies	Controllers
(a) Model-based						
Hiller and Lipson (2011)	Soft robot	M + C	Voxelyze	EA	5	1
Caluwaerts et al. (2014)	Tensegrity robot	M + C	NTRT	EA	1	2
Cellucci et al. (2017)	Recyclable robot	M	ODE	MOEA	3	1
Peng et al. (2018)	Robotic arm	C	MuJoCo	DNN	1	1
(b) Model-free						
Khazanov et al. (2014)	Tensegrity robot	C	—	EA	1	250
Brodbeck et al. (2015)	Modular robot	M + C	—	EA	500	500
Mouret and Clune (2015)	Soft robotic arm	C	—	MAP-Elites	1	640
Vujovic et al. (2017)	Soft-legged robot	M + C	—	Evo-devo	75	75
Nygaard et al. (2018)	Quadruped	M + C	—	MOEA	192	192
Veenstra et al. (2018)	Soft fish robot	C	—	EA	1	200
(a) Reality-assisted						
Koos et al. (2012)	Multiple	C	Bullet physics	MOEA	2	2
Cully et al. (2015)	Hexapod	C	ODE	IT&E	10	10
Rosendo et al. (2017)	Modular robot	M + C	GP	BO	25	25
Saar et al. (2018)	Hopping robot	M + C	GP	BO	15	40
Rieffel and Mouret (2018)	Tensegrity robot	C	GP	BO	1	30
Rosser et al. (2019)	Flapping robot	M	PYROSIM	MOEA	16	1
Kwiatkowski and Lipson (2019)	Robotic arm	C	DNN	—	2	2
Howison et al. (2020a)	Falling paper	M	KMC	Random	500	—
Howison et al. (2020b)	Falling paper	M	GP	BO	40	—
Kriegman et al. (2020a)	Living cell robots	M	Voxelyze	EA	5	1
Kriegman et al. (2020b)	Modular soft robot	M	Voxelyze	Exhaustive	108	1

KeyM: *Morphology*C: *Control*GP: *Gaussian Process*KMC: *K-Means Clustering*ODE: *Open Dynamics Engine*NTRT: *NASA Tensegrity Robotics Toolkit*MAP-Elites: *Multi-Dimensional Archive of Phenotype Elites*DNN: *Deep Neural Network*EA: *Evolutionary Algorithm*MO: *Multi-Objective*MO: *Multi-Objective Evolutionary Algorithm*BO: *Bayesian optimisation*IT&E: *Intelligent Trial and Error*

Table 2.1: Highlighted literature for model-based, model-free and reality-assisted design of physically embodied robots, showing the target system, design parameters, modelling and simulation tools, design search method and the number of morphologies and/or controllers that were realized *in reality*. Note that in many cases these publications reported significant numbers of additional experiments, for example manual design comparisons. (a) Model-based approaches. Designs are evaluated using modelling and simulation tools and are not tested in the real-world until the end of the design process. (b) Model-free approaches. Designs are exclusively evaluated in the real-world, without the use simulation or modelling tools. (c) Reality-assisted approach. As categorized in this review chapter, designs are evaluated using a combination of modelling tools and real-world experimentation.

2.1 Embodied intelligence and soft robots

Conventional robotic systems can usually be described using a discrete architecture. The exact position of actuation, sensing and other morphological elements is known and can be relied upon not to change during deployment. Interaction with the environment can be actively regulated, for example in pick-and-place robotic arms. The continuum nature of soft robots means they cannot be described using this discrete architecture, but rather as a set of loosely-coupled parallel processes and environmental interactions spanning the microscopic to macroscopic scales. At the microscopic scale, molecular and material interactions induce physical properties such as elasticity or durability, as well as active behaviours such as those seen in electroactive polymers or shape memory alloys (Whitesides, 2018). At the organism level these properties manifest as distinct higher-order behaviours, with environmental interactions inducing highly non-linear dynamics, for example sustained contact forces or turbulent fluid dynamics (Katzschmann et al., 2016). Further still, at the population level soft robots can be placed in the context of modular and distributed robotics, with interaction between agents driving complexity. In terms of embodied intelligence, therefore, the structure and interdependence of these interactions is key for the emergence of useful, non-trivial behaviours. Given this, three classes of soft robot are particularly interesting:

1. **Soft robots with complex interactions.** A key area of interest is soft robots with a *simple* embodiment but *complex* environmental interactions. For example, a soft extendable tube that exhibits complex behaviours via contact with the environment (Hawkes et al., 2017), a simple water-immersed silicon tentacle whose dynamics can be utilised for real-time information processing (Nakajima et al., 2015) or liquid droplets whose properties drive sensing and motility functionalities (Cira et al., 2015). These non-trivial behaviours that emerge via interaction between morphology and the environment represent a fundamental basis on which more complex behaviours can be built built.
2. **Soft robots with complex hybrid structures** Another key area of interest is so-called soft-rigid robots. The rapid progression in additive manufacturing has enabled the construction of heterogeneous continuum structures with anisotropic elasticity profiles. For example, an anthropomorphic hand with rigid bones but soft ligaments utilises its hybrid structure to achieve complex dexterity tasks (Hughes et al., 2018). Complex soft structures can also be built using rigid components with soft connections, for example a tensegrity robot that combines the robustness of rigid components with the versatility of a deformable body (Rieffel and Mouret, 2018). Robots constructed using

living tissues have also shown the potential of hybrid structures for encoding functional behaviours (Kriegman et al., 2020a). Similar to biological systems, hybrid robots can diversify their behaviours by complexifying their morphology.

3. **Soft robots that can change their own embodiment** The final class concerns one of the most fascinating aspects of soft robots, growth and adaptation. One of the key properties separating biological systems from robots is an ability to change ones own embodiment, over the course of a lifetime or in response to the environment. Adaptation has been shown to improve soft robot evolution in real-world (Vujovic et al., 2017) and simulated robots (Kriegman et al., 2018). Morphogenesis has been demonstrated on modular robotic platforms (Vergara et al., 2017). Understanding how to harness the power of developmental processes is a key step for progressing from simple soft robots to complex and intelligent machines that manifest behaviours driven by seemingly conscious action.

2.2 Modelling and representation

Models (mathematical or otherwise logical representations) and simulations (evaluated models) of soft robotic systems facilitate *in silico* experimentation on a scale many orders of magnitude higher than can be achieved *in reality*. This is beneficial for discovering and testing candidate designs before transference to the real world, and for the systematic testing of different design methodologies, e.g., evolutionary algorithms. However, all model representations of the real world have some form of reality gap (Mouret and Chatzilygeroudis, 2017). While it may be possible to construct models of physical phenomena, it may not always be possible to evaluate these models in simulation to a sufficient granularity for the purposes of design. This section discusses approaches for modelling soft robots, from conventional methodologies to the latest physics engines and data-driven modelling techniques.

2.2.1 Conventional modelling

Popular modelling approaches for soft robots are at the level of kinematics or second-order full dynamics (Lipson, 2014). Such methods involve deriving precise geometric and dynamic expressions, which are in turn used to derive Jacobian kinematics or equations of motion. For example, elastic systems can be modelled as mass-spring lattices or arrays, providing a simple estimation of deformation. Similar approaches include traditional beam-bending and constant-curvature approximations, which often have analytical solutions. Bespoke models can combine these approaches for different modalities of soft body deformation, for example continuum structures with constant-curvature models (Della Santina and Rus, 2019; Runge et al., 2017; Webster III and Jones, 2010). Finite element methods (FEM) are also

used for modelling soft systems (Zhang et al., 2017), and have seen particular usage for exploring the design parameter spaces of soft pneumatic actuators (Connolly et al., 2015; Drotman et al., 2017) and controller verification (Zheng et al., 2019). FEM models are generally computationally expensive, requiring a trade-off between simulation time and accuracy (Pozzi et al., 2018). However, work on real-time FEM has shown some promise for the control of soft robots (Duriez, 2013; Zhang et al., 2016).

2.2.2 Physics engines and simulators

Generalised physics engines that can simulate soft robots as they interact with their environment are growing in popularity. A number of physics engines offer soft body simulators, including Bullet Physics², CryEngine³, MuJoCo⁴, Open Dynamics Engine⁵ and the NASA Tensegrity Robotics Toolkit⁶. Perhaps the best known physics engine for deformable bodies is Voxelyze⁷, a voxel-based representation that has been used for a large number of soft robotics studies (Cheney et al., 2013; Hiller and Lipson, 2014; Kriegman et al., 2017, 2020b) and has become a benchmark on which different design optimisation methods are tested. A recent gpu-accelerated re-implementation of Voxelyze, voxcraft-sim⁸, has recently been released.

There have also been recent advances in differentiable soft-simulators, which have the advantage of solvability using gradient-based optimisation algorithms. These are particularly efficient for solving optimal control and motion planning problems. However, implementation for soft systems is challenging and currently highly computationally expensive. Examples of soft-body differentiable systems include Chain Queen (Hu et al., 2019) and the differentiable cloth simulator (Liang et al., 2019). The ability to use gradient-based optimisation methods could allow for far more computationally efficient optimisation and exploration of design.

Despite developments in soft-body modelling there are still many limitations. Rigid body approximations, FEM and constant curvature models can be effective at predicting the behaviour of specific problems, but do not generalise well. Physics engines such as Voxelyze certainly offer a more generalised simulation environment. Additionally, many physics engines are developed by the video game industry, so come with certain technical and financial weight behind them. However, with generality can come a loss of accuracy when compared to problem specific modelling methods. In general, different physics engines offer

²<https://pybullet.org/>

³<https://www.cryengine.com/>

⁴<http://www.mujo.co.org/>

⁵<https://www.ode.org/>

⁶<https://github.com/NASA-Tensegrity-Robotics-Toolkit>

⁷<https://github.com/jonhiller/Voxelyze>

⁸<https://voxcraft.github.io/>

different trade-offs and capabilities, each suited to different elements of soft robot simulation (Silva and Maciel, 2012).

2.2.3 Data-driven modelling

Significant research has been carried out into data-driven modelling techniques that update their representation based on ground-truth data. Data-driven approaches offer the ability to incrementally improve predictions as more data becomes available. Some approaches construct models directly from ground-truth data, for example deriving symbolic equations to describe real-world phenomena (Brunton et al., 2016; Rudy et al., 2017). Learnt knowledge of fundamental concepts can be gradually combined for understanding more complex systems (Schmidt and Lipson, 2009). Other methods detect underlying structure in large data-sets, for example dynamic mode decomposition (Schmid, 2010).

Alternatively, increasingly black-box approaches can be used to estimate the behaviours of soft systems without explicitly discovering the underlying dynamics. Neural networks, for example, have been used for soft robot pose reconstruction (Scimeca et al., 2019), inverse kinematics (Giorelli et al., 2015) or control strategies (Choi et al., 2018), and are effective without having to explicitly model the mechanisms behind soft deformation. The rapidly growing field of neural ordinary differential equations could see particular usage in data-driven modelling for soft robots (Rubanova et al., 2019).

More abstracted methods bypass direct modelling of the underlying dynamics completely, instead learning surrogate models to map between the design space and fitness landscape. Gaussian process (GP) modelling, for example, is a supervised learning technique for solving regression problems (Snoek et al., 2012). The power of GP models lie in their leveraging of Bayesian inference to estimate black-box function behaviours in relatively few samples. GP models have been used to estimate design fitness in the reality-assisted design of a hopping (Saar et al., 2018) robot, mapping morphological and control parameters to locomotion speed, and a tensegrity robot (Rieffel and Mouret, 2018), mapping motor speeds to locomotion speeds.

2.2.4 Transferability, robustness and sim2Real

Related to data-driven modelling is the idea of transferability. Simulation-to-real (sim2real) performance can be improved by learning the disparity between robot behaviours observed *in silico* and *in reality* (e.g., their transferability). The *transferability approach* augments existing system models with supervised learning to understand the limits of the model and simulator (Koos et al., 2012). This so-called robot-in-the-loop method uses a minimal number of physical experiments to build a surrogate model of the sim2real disparity across the design space. By including this disparity measure within a multi-objective optimisation

framework, the design fitness landscape can be adjusted to reward designs that are likely to cross the reality gap with a low disparity. Incremental testing in the real world improves the sim2real disparity surrogate model accuracy, driving the design process towards performant *and* transferable solutions. A recent study on a flapping robot, whose dynamics are notoriously challenging to model, used incremental simulated and real-world experimentation to explicitly measure the sim2real disparity across the design space (Rosser et al., 2019).

Robustness filters can be also used to assess how well simulated designs perform under perturbations (Jakobi, 1997), for example in their morphology, controller or environment. If the reality-gap is viewed as a large perturbation to the design space, then designs that can withstand high perturbations in simulation and still exhibit the same behaviours are likely to also perform well in the real-world. These methods have been used for the design of living tissue robots (Kriegman et al., 2020a), in which simulated robot designs were passed through a robustness filter that injected noise into their control systems, and a transferability filter that measured their performance once transferred to reality and returned this to the model environment in the form of constraints. Transferability and robustness methods are useful tools because they maximise the utility of existing simulation tools, allowing generalised physics engines to be used in many different design contexts. However, the design process is still fundamentally limited by the accuracy of the model tool. Areas of the design space that cannot be simulated accurately may still provide useful data, but cannot be realized in the form of a physically embodied robot.

2.3 Large-scale physical experimentation

Those implementing reality-assisted evolution rely on carrying out a sufficient number of physical experiments. This section discusses methods and challenges for fabricating and testing physical soft robots on a large scale. First it discusses fabrication, sensing and actuation approaches for soft robots, drawing attention to the difficulty in realising simulated designs in the physical world. Next it discusses approaches for scaling up experimentation to enable more design evaluations in the real world, including automated fabrication and modular robotic systems.

2.3.1 Materials and fabrication

One of the great opportunities offered by soft robotics is the large pool of potential materials and fabrication technologies available to the designer (Schmitt et al., 2018). Many soft robots are fabricated using casting and moulding, typically with silicon elastomers or expanding foam (Ilievski et al., 2011; Polygerinos et al., 2017). These approaches offer a flexible and potentially scalable approach for fabrication, but can be a highly labour inten-

sive. Additive manufacturing has also recently seen widespread adoption, with 3D printing shown to be especially useful for multi-material soft-rigid robots, e.g., an anthropomorphic hand (Hughes et al., 2018), and for seamless incorporation of sensing and actuation mechanisms (Wehner et al., 2016). 3D printing can also be exploited with flexure-based rigid techniques (Hughes and Iida, 2017), with increasingly high-quality flexible filaments and stereolithography techniques offering significantly more capabilities (Gul et al., 2018).

There are a number of alternative materials and approaches which potentially offer more scalable approaches. This includes hot melt adhesives, which have seen widespread usage for soft robots (Nurzaman et al., 2013; Vujovic et al., 2017), and offer the advantage of being cheap, accessible and quick to form. Reconfigurable soft robots are another approach, and can be fabricated using cutting and folding inspired by origami (Onal et al., 2011; Rus and Tolley, 2018; Zhakypov and Paik, 2018). Similarly, laser cutting is a rapid method for fabricating simple robots (Chapter 6), or creating 3D flexure based actuators (Lipton et al., 2018). In addition to these well established techniques there are a number of exciting novel fabrication approaches emerging, including knitted or balloon-structure robots (Maziz et al., 2017; Yarbasi and Samur, 2018).

2.3.2 Sensing and actuation

Another key challenge in the development of soft robots is sensing and actuation (Pfeifer et al., 2012). Actuation methods such as pneumatically-driven systems are popular, however these often require tethers or large external compressors (Sun et al., 2013). Tendon-driven systems are also common and allow for use of motors as an actuation source, providing a higher controllability and power density (Camarillo et al., 2008). While these can be efficient, their size and rigidity makes them hard to integrate into continuum bodies. More unconventional approaches use smart materials, for example dielectric actuators or shape memory alloys (Li et al., 2019; Motzki et al., 2019). These methods have the clear advantage that actuation can be seamlessly integrated into a robot morphology and are especially useful from the perspective of embodied intelligence. However, it is currently challenging to generate large forces with these technologies (Gu et al., 2017).

Sensing is challenging due to the large deformations and potentially infinite degrees of freedom of deformation that soft-body systems exhibit (Shih et al., 2020). A number of different sensing approaches have been proposed including piezoresistive materials (Firouzeh et al., 2015), capacitive sensors (Frutiger et al., 2015), ionic metal soft sensors (Chossat et al., 2013), and also sensors which use cameras (Gilday et al., 2020) or exploit physical structures (Scimeca et al., 2019). Again, more novel methods have been presented, for example harnessing the sensing properties of liquid droplets (Cira et al., 2015).

A growing trend in soft robotic sensing and actuation is the use of bio-hybrid materials. The boundaries between biological and artificial systems can be blurred with the development of bio-hybrid robots (Romano et al., 2019). Some of the capabilities, structures and materials from biological systems can be incorporated to interact with robots (Alapan et al., 2019). For example, a bio-hybrid system could incorporate biological muscles, where the muscles are controlled artificially (Morimoto et al., 2018). This approach could allow the functional gap between current soft-robotics systems and biological systems to be closed.

2.3.3 The fabrication gap

The fabrication approaches discussed in 2.3.2 have a complex set of limitations dictating what types of design can and can't be built to a sufficient level of accuracy. Hence, if the fabrication process is not properly included in the design process we may see a *fabrication gap* (Rieffel and Pollack, 2005), i.e. the difference between the prescribed design and physically fabricated end product. The fabrication gap problem can be addressed by considering available fabrication methodologies early in the design process.

One approach is to use fabrication rules to filter designs at the simulation stage, for example using a build filter to rule out designs that cannot be physically fabricated (Kriegman et al., 2020a). Alternatively, designers can use fabrication rules – 3D printing commands, for example – as a means to describe designs (e.g., in (Rieffel and Pollack, 2005)). *Grammar*-based approaches such as this allows designs to be described via vocabulary of parts and fabrication methods (Lau et al., 2011). These ideas have recently been advanced with the concept of a *robot compiler*, e.g., a process for converting a robot design into parts, structures and fabrication instructions. Robot compilers have been used to convert simulated designs to origami robots whose fabrication is guaranteed (Mehta et al., 2014; Mehta and Rus, 2014; Schulz et al., 2017).

Each fabrication technology introduces its own characteristic to the final designs, which in turn affects their embodiment and thus their behaviour. Selecting the correct fabrication approach could vastly increase the performance of soft robots. Indeed, if fabrication is viewed as a *filter* between the virtual and physical environments then this could improve the design process by automatically ruling out designs that could not cross the fabrication gap. The recent work on robot compilers and grammars demonstrates, to some extent, that including fabrication at the representation stage can allow us to harness limited fabrication technologies to our advantage in the design process.

2.3.4 Scaling up physical experimentation

There is a growing trend for testing large numbers of robot designs *in reality*, and significant research is being carried out to enable more scalable physical experimentation

methodologies. Indeed, this is one of the driving factors in presenting the reality-assisted framework as a feasible option for unifying model-based and model-free design methodologies.

Online controller learning is relatively easy to implement with physical experimentation as the morphology remains fixed, and has been applied successfully for locomotion (Khazanov et al., 2014; Nygaard et al., 2018) and a soft fish robot (Veenstra et al., 2018). Other studies have co-optimised control and morphology in the real world. One platform uses a ‘mother robot’ to automatically construct different robot designs and test them for their locomotive behaviours (Brodbeck et al., 2015; Rosendo et al., 2017; Vujovic et al., 2017). By automating the fabrication process, large numbers of real-world robots can be evaluated with minimal human intervention. These studies reveal the effectiveness of scaling up physical experimentation, but are still fundamentally limited to orders of magnitude fewer design evaluations than is achievable *in silico*.

To address this, methods of rapidly fabricating soft robots have been developed, and robotic automation has been used to partially or fully automate the fabrication process. A number of soft robotic fabrication techniques are particularly suited, or can be adapted for rapid fabrication. Additive manufacturing using flexible thermoplastic elastomers (Gul et al., 2018; Wallin et al., 2018) is one example, as is ‘1-D’ printing techniques that can fold a single string of recyclable material into a target morphology (Cellucci et al., 2017). In this approach thermoplasters can be extruded using high precision robot arms to create structures, or to connect existing components together. This approach offers a high level of accuracy and design complexity in a relatively short time. New printing technologies are also developing where complex 3D structures can be fabricated in one step (Wehner et al., 2016), improving the speed at which robot designs can be fabricated. Alternatively, low cost soft robot construction kits (voxcraft) have been presented along with appropriate simulation tools to design and test many physical robots (Kriegman et al., 2020b).

2.3.5 Modular and distributed systems

Modular robotic systems allow significant physical experimentation without the need to continually fabricate new robots. Instead, modules can be reconfigured to make use of existing fabricated components. Not only does this speed up experimentation, but it allows investigation into bio-inspired topics such as morphogenesis. For example, emergent self-assembly approaches inspired by nature have been replicated in a number of distributed robotic systems (Rubenstein et al., 2014; Werfel et al., 2014). Here, global strategies are converted to local assembly rules via a compiler. Modular robot systems such as these are

forced to consider fabrication early in the design process, as modules can only be connected in certain configurations. In doing so, users can create an implicitly ‘buildable’ design space.

Systems where either the assembly or disassembly process is manually performed include for example the programmable-matter systems Robot Pebbles (Gilpin et al., 2010) and Mice (Gilpin et al., 2008). Even more interestingly, there is a special type of modular robots which can autonomously self-assemble and self-disassemble, i.e. self-reconfigurable modular robots. In theory, these systems could perform experiments autonomously (limited by their operation time), as they would assemble into a design, test it, and re-assemble into the next design by themselves. Within the domain of rigid robots there have been a large number of different such modular robotic technologies. Recent examples of these include M-blocks (Romanishin et al., 2013), SMORES (Davey et al., 2012), Soldercubes (Neubert and Lipson, 2016), Roombots (Hauser et al., 2020). Realising soft modular or self-assembling robots is less common, largely due to the inherent difficulties in fabricating and controlling continuum bodies. However, modular soft robots are seeing increasing use (Kriegman et al., 2020b), for example tensegrity systems (Zappetti et al., 2017), tendon-driven structures (Malley et al., 2017), and soft modular cubes for investigating morphogenetic movements of the embryo (Vergara et al., 2017).

A key related field is self-replicating robotics. Introduced half a century ago by John von Neumann, self-replication has long been a dream of roboticists. Requiring futuristic technologies, it would allow a self-sustaining pipeline of robots to be fabricated autonomously. A number of examples have been developed (Suthakorn et al., 2003), and 3D printers have provided an imperfect replicator (Bowyer, 2014). However, there has not yet been significant progress towards soft self-replicating systems.

2.4 The search for novelty

Key to the success of any design optimisation process are the methods used to search for favourable designs and the metrics by which designs are compared. As discussed in 2.3, a key process in reality-assisted evolution is discovering designs in simulation that are likely to transfer well to reality. This section summarises advances in optimisation techniques for soft-robot design.

2.4.1 Bio-inspired encoding

Robot designs are typically described via an encoding which systematically maps input parameters to design features. The design space of all possible designs is, therefore, defined by the encoding along with the valid range of input parameter values. There is a range of desirable encoding features. Encoding schemes should output feasible designs for a given

task environment, but ideally generalise to many task-environments. They should be receptive to the influence of optimisation algorithms searching for favourable designs within the design space, i.e. they should be *evolvable*. Finally, as mentioned previously, having a grounding to the available fabrication technology is advantageous.

Direct encodings map input parameters directly to design features. This can be a powerful approach as it allows a human designer to specify the key features of the output design, hence ensuring the design space is feasible and can be fabricated. For example, robot foot and leg parameters have been optimised with direct encodings (Saar et al., 2018). However, the approach performs poorly in terms of scalability and generality. Since every node in the discretised design space needs a separate, exact description, the total amount of information stored in the direct encoding is as large as the design itself. This can be problematic, especially for soft robots where a fine design discretisation can quickly make the optimisation process intractable. Furthermore, a direct encoding scheme suitable for a given task may not transfer to another.

Instead, research has focused on *indirect* or *generative* encoding systems that can describe a complex design space with relatively few parameters (Hornby and Pollack, 2001; McCormack et al., 2004). The idea is inspired by biological systems, where phenotypic expression often occurs in repeating patterns that can be described by a singular part of the encoding, in this case a genome. Lindenmayer systems (L-systems) and their variations (Jacob, 1994; von Mammen and Jacob, 2007) are generative encodings that incorporate bio-inspired aspects of morphogenesis, the development of morphological characteristics. Their particular properties make them effective for modelling natural plant systems, with directed growth and branching. However, they have not yet found a clear application in the design of soft robots.

Another more promising indirect encoding is *compositional pattern-producing networks* (CPPNs) (Clune et al., 2011). CPPNs mimic the natural ontogenetic process without the need to directly simulate the chemical mechanisms involved. Instead, describing morphology in terms of a network where each node is a mathematical function (e.g., a Gaussian or cosine function), and the final design is the result of design coordinates queried through this network. By including regularity in the encoding (Clune et al., 2011), design information can be compressed significantly while maintaining diversity within the design space. CPPNs have been combined with a various optimisation algorithms that evolve their underlying network to design soft robots in a simulation environment (Cheney et al., 2013), with some having been transferred to the real world. Growing evidence has supported the idea that generative encoding schemes using CPPNs or similar outperform other approaches (Richards and Amos, 2014; Tarapore and Mouret, 2014).

2.4.2 Optimisation algorithms

Single objective optimisation (SOO) strategies aim to maximise performance in a specific behavioural feature or task (Trianni and López-Ibáñez, 2015), for example locomotion speed. Evolutionary algorithms (EAs) mimic natural selection and offer an effective approach to searching for favourable designs, providing the desired phenotypic behaviours can be encoded within a fitness function. In simulation they have been shown to effectively optimise a range virtual agents (Sims, 1994) including for locomotion (Cheney et al., 2013; Duarte et al., 2017) or navigating confined space (Cheney et al., 2015). EAs have also been demonstrated in the real world, evolving modular robots for locomotion (Brodbeck et al., 2015; Vujovic et al., 2017) and for controlling a soft fish robot (Veenstra et al., 2018).

Bayesian optimisation (BO) (Frazier, 2018) is an alternative strategy for the global optimisation of expensive-to-evaluate black box functions, and has seen significant usage in SOO problems for robotics. BO methods sequentially build and improve a surrogate GP model of the underlying function and use this to efficiently sample the parameter space. Significantly, BO allows users to choose a range of acquisition functions to decide the next sample point, some of which include a penalty for evaluation time. The algorithm, therefore, can search for designs that maximise the predictive accuracy of the GP model and design fitness while also minimising time spent doing physical experimentation.

BO has been used for a range of physical robot design optimisation tasks including modular robots (Rosendo et al., 2017), locomotion control strategies (Rieffel and Mouret, 2018), falling paper shape morphologies (Chapter 5) and for the co-optimisation of morphology and control (Saar et al., 2018), where it was shown to outperform a human designer. In this study, physical robots were iteratively fabricated and tested. After each test, the design fitness was used to update the GP model which, in this case, mapped leg parameters to predicted locomotion distance. BO has also been used in Intelligent Trial and Error (IT&E), designed to quickly adapt robot controllers to morphological damage by using real-world experimentation to learn optimal control policies (Cully et al., 2015).

There is a growing trend for implementing multi-objective optimisation (MOO) approaches in the evolution of robotic systems. Since intelligent behaviours can rarely be characterised by a single fitness function, MOO offers a framework for characterising more complex behavioural fitness functions into the design process (Trianni and López-Ibáñez, 2015). Multi-objective evolutionary algorithms (MOEAs) use multiple objectives to drive the evolutionary process, offering the designer a set of solutions with trade-offs between the various objectives (Mouret and Doncieux, 2008, 2012). MOEAs are effective for evolving simulated soft robot behaviours, for example in aquatic environments (Corucci et al., 2018). Transferability approaches also use the MOO framework, finding Pareto optimal solutions

that balance design fitness and with an estimation of transferability (Koos et al., 2012). MOO frameworks are effective for finding a range of possible design candidates, however will not return solutions in areas of the design space that are not Pareto optimal. Hence, many potential candidate designs can be missed, especially in the case of poorly defined objectives.

2.4.3 Quality-diversity and illumination algorithms

Designers can incorporate solution novelty within the optimisation objectives. Such algorithms aim to actively create divergence in the solution space, rather than converging to one global maxima, for example by formulating the optimisation problem to search only for behavioural novelty, i.e. *novelty search* (NS) (Doncieux et al., 2019; Lehman and Stanley, 2011a). The formulation has been shown to outperform traditional objective based searches for solving task-based problems such as maze solving or obstacle avoidance (Mouret and Doncieux, 2012), and for the online morphological adaptation of a simulated underwater robot (Corucci et al., 2015). In the context of design optimisation, this offer an alternative way of evaluating designs based on their behaviour, rather than directly on an fitness function. Because many points in the design space may map to the similarly performant behaviours, this approach could significantly improve search efficiency. However, without any direct measure of solution fitness, searches based on novelty can be too divergent to ever present an optimal solution candidate.

To address this, so-called *illumination* or *quality diversity* (QD) algorithms have been presented which hybridise the concept of novelty searching and direct fitness optimisation (Pugh et al., 2016). QD algorithms aim to discover the range of behavioural niches in the system, as well as the associated optimal solution of each niche. One example, Novelty Search with Local Competition (NS+LC) (Lehman and Stanley, 2011b), lets similarly behaved solutions compete based on a fitness function. When evolving virtual agents for locomotion, NS+LC found more functional morphological diversity than a global fitness function. Similarly, the Multi-Objective Landscape Exploration algorithm (MOLE) (Clune et al., 2013) searches for diverse solutions that are as far away from each other as possible in a user-defined feature space. One of the most popular QD implementations is the Multi-Dimensional Archive of Phenotype Elites (MAP-Elites) (Mouret and Clune, 2015) algorithm. Here, a set of highly performant solutions are discovered across a discretised feature space defined by the user. The holistic nature of MAP-Elites allows the human designer to define interesting robot behaviours and use these to structure the evolutionary search, and has been reported to outperform traditional EAs in a range of tasks.

The QD approach may be particularly effective for the design of soft robots because it is more accommodating to unknown system behaviours. In the case of soft robotics, we often

do not know how a particular soft system will behave. optimising a range of solutions with different behaviours makes it more likely some of these solutions will be effective when transferred into the real world.

2.4.4 Developmental robotics

Taking inspiration from nature, there is growing evidence that mimicking ontogenesis and allowing a developmental phase, where a robot can alter its morphology or controller, can be highly beneficial (Doursat and Sánchez, 2014; Eiben et al., 2013; Rieffel et al., 2014). Incorporating a developmental process into the design optimisation framework adds a number of benefits. First, the design-space complexity can be reduced and offloaded to the developmental interaction between the robot and its environment. Second, performance disparities introduced by the reality gap can be corrected via developmental processes in the real world. Including developmental processes also means a robot can adapt to a changing task environment, or repair itself if damaged. Similarly, controller development allows robots to correct for damage or an unexpected change in task (Cully et al., 2015). Third, it could allow the design space of embodied robots to be explored more efficiently and effectively. It has been demonstrated that implementing a developmental framework combined with a novelty search is more effective for evolving simulated soft robots than using a traditional fitness driven approach (Joachimczak et al., 2015). A key study (Kriegman et al., 2018) demonstrated how phenotypic plasticity can guide and improve evolution, so-called evo-devo, and increase the overall evolvability across the design space. Large-scale developmental processes have seen limited implementation for real-world embodied soft robots, but the preliminary results indicate the potential (Vujovic et al., 2017).

2.5 Discussion and conclusion

This review chapter is written in the context of two key problems in the design of soft robots. First, although model-based methods are effective for searching large design spaces, we are currently unable to sufficiently model many physical processes within soft robots, and are also unable to physically fabricate many virtual soft-robot designs. These reality and fabrication gaps severely limited the scope of designing physically embodied soft robots with a model-based approach. If we subscribe to the view that intelligent behaviours result from complex embodied interaction within the environment, then there is a clear need for significant physical experimentation in lieu of highly accurate modelling and simulation. Second, while model-free physical experimentation does fully capture the embodied behaviours of soft robots, the approach is not scalable to the size achievable with a model-based approach. Designers cannot, for example, explore the evolutionary design of

complex soft robots since exploration of these design spaces requires a significant numbers of evaluations.

The challenges associated with designing soft robots combined with the potential benefits of utilising soft robots in the context of embodied intelligence are forcing robotiscists to rethink their design methodologies and philosophies. The literature presented here (as highlighted in Table 2.1) demonstrates that an effective strategy is a reality-assisted evolution framework in which model-based and model-free physical experimentation methods are unified to improve the design of physically embodied soft robots that exhibit useful, meaningful behaviours. *In silico*, data-driven modelling can be used to build and update models of real-world systems. Tuning physics simulations (Kriegman et al., 2020a), building models from scratch (Saar et al., 2018) and learning auxiliary transferability models (Koos et al., 2012) are all data-driven methods for reducing the reality-gap disparity. *In reality*, large-scale physical experimentation methods are facilitating significant numbers of real-world evaluations. Automated (Chapter 6) and scalable (Kriegman et al., 2020b) methods show promise, as do modular and reconfigurable systems (Nygaard et al., 2018; Vergara et al., 2017). At a high-level, novel optimisation techniques can guide design exploration *in silico* and *in reality*. QD and illumination algorithms such as MAP-Elites (Mouret and Clune, 2015) seem particularly suited to modulating the design process by discovering multiple diverse candidate designs to be tested in the real world. Bayesian optimisation (Rieffel and Mouret, 2018) methods also show promise since they allow designers to include the cost of physical experimentation directly in the objective function.

Part I

Falling paper

Chapter 3

Physics-driven behavioural clustering of free-falling paper shapes¹

Complex physical phenomena are often governed by highly non-linear, multidimensional dynamics. Hence, it can be challenging to understand these systems using traditional modelling tools, as we lack knowledge of the underlying physical phenomena required to implement these. The obvious course of action, then, is to infer these phenomena via physical experimentation. Automating this inference process, in other words automating the discovery of system physics from experimental data, has been the focus of intensive study.

Schmidt and Lipson (Schmidt and Lipson, 2009) developed an algorithm to automatically discover analytical relationships in dynamical systems, ranging from simple harmonic oscillators to more complex chaotic double pendulum systems. This was preceded by a method of non-linear model synthesis from directly observed data using co-evolution (Bongard and Lipson, 2005). Meanwhile, in the fluid dynamics community sparse regression has been used to determine the fewest terms in the dynamic governing equations required to accurately represent the data (Brunton et al., 2016). Data-driven approaches to modelling have also shown the ability to predict behaviours of dynamic systems (Kutz et al., 2016; Schmid, 2010).

¹This chapter is based on the following peer-reviewed publication:

- Howison, T., Hughes, J., Giardina, F., & Iida, F. (2019). Physics driven behavioural clustering of free-falling paper shapes. *PLoS One*, 14(6), e0217997.
- Data availability: <https://doi.org/10.1371/journal.pone.0217997>

Contributions

T. Howison – devised the PDBC algorithm and VSFP system, planned and carried out experiments, wrote the bulk of the paper and created the figures.

J. Hughes – assisted with experiments, helped write initial paper draft, proofreading and general discussions.

F. Giardina – contributed to early exploratory work and general discussions.

F. Iida – contributed to general discussions and proofreading.

Similarly, big-data has been utilised for the prediction and physical understanding of complex systems include (Gandomi and Haider, 2015; Goza and Colonius, 2018; Quiroz et al., 2018; Strogatz, 2018; Yang et al., 2015). Other studies used evolutionary algorithms with feedback from environmental interaction to optimise robotic morphologies without any system model (Brodbeck et al., 2015; Nurzaman et al., 2013; Rosendo et al., 2017).

These approaches present a few problems. First is the reliance on large data-sets. Sampling through physical experimentation typically involves searching high dimensional landscapes (Kutz, 2017). This makes data generation difficult, especially for expensive-to-evaluate functions. Second, although highly effective at identifying the inherent physical relationships of non-linear systems, they have not shown the ability to predict the boundaries of these non-linear behaviours. This is of particular importance in systems with a diverse range of discrete behavioural modes over their parameter space. In such systems, the dynamics of different behaviours may be significantly different, and the driving factors causing behavioural switches unclear. Such discrete behaviour systems are seen widely throughout nature including laminar-turbulent behaviours in fluid dynamics (Sano and Tamai, 2016), gait patterns in locomotion (Alexander, 1989; Corucci et al., 2018) or even the behaviour of flocking systems (Toner et al., 2005).

In this chapter I present Physics-driven behavioural clustering (PDBC), a novel method that automates the process of discovering functions that enable behavioural clustering and physical understanding of systems with discrete behavioural modes. The PDBC method has the potential to discover physically insightful clustering functions based on relatively few experimental observations, thus enabling breakthroughs in the understanding of expensive-to-evaluate and behaviourally diverse systems. In the PDBC method, observational data is organised and transformed into the parameter space of a set of generic functions. We hypothesise that there exists a set of functions whose parameter space is divided into distinct regions corresponding to different behavioural modes. Furthermore, we hypothesise that more physically relevant functions – such as Re and I^* in the case of falling disks (Field et al., 1997) – will cluster similar behaviours together more effectively. We propose that the predictive accuracy and clustering strength of a standard unsupervised clustering algorithm in this parameter space can be used as a direct metric for physical significance, with strongly clustered solutions with low predictive errors being more physically relevant.

We address the challenging problem of clustering and understanding the falling behaviours in the V-shaped falling paper (VSFP) system, which is a new contribution to the falling paper system class. This is inspired by the falling and fluttering behaviours observed by helicopter seeds (Lentink et al., 2009). The VSFP system is an interesting challenge because although the design parameter space is limited, it exhibits rich behavioural diversity.

Therefore it is an ideal system to demonstrate the PDBC method. We demonstrate the PDBC method can effectively cluster and help explain the VSFP behaviours.

This chapter is structured as follows. In 3.2 I describe the PDBC algorithm for a general system. In 3.3 I describe the VSFP system. In 3.4 I present experimental results of the VSFP system. In 3.5 I demonstrate the effectiveness of the PDBC method on the VSFP system. In 3.6 I conclude the chapter.

3.1 Falling paper and related problems

In this section I briefly outline the key literature relating to falling paper and other similar systems. The study of how paper, or more generally plate and disk shaped objects, fall in air and other liquids has been studied extensively from many different points of view. Maxwell (1854) presented the first explanation for the falling dynamics of such objects, famously observing that '*Everyone must have observed that when a slip of paper falls through the air, its motion, though undecided and wavering at first, sometimes becomes regular*'. Maxwell proposed that this regular motion is the result of rotation—termed *autorotation*—induced by a pressure differential between the leading and trailing paper edge. Furthermore, he suggested that the exact behaviour is highly dependent on the initial conditions and paper shape, with symmetrical shapes exhibiting a more pronounced rotation. The mathematical, fluid mechanics, and computational methods of the time meant that Maxwell was unable to quantitatively test any of his hypotheses. Techniques have developed since then, allowing a more in depth analysis of the problem.

Willmarth et al. (1964) conducted a large scale study on the steady and unsteady motions of falling disks, showing that behaviour exhibited a systematic dependence on the Reynolds number (Re) and dimensionless moment of inertia (I^*). The boundary between steady and unsteady disk oscillation was identified. Smith (1971) presented an experimental investigation into the autorotation of wings, showing that freely falling wings either autorotated, rocked or remained stable. Again, the type of motion observed was highly dependent on the wing Re and I^* . Iversen (1979) presented similar findings for the autorotation of flat plates. Hans Lugt David W Taylor (1983) contributed a comprehensive review of autorotation for various objects and described the theoretical difficulties of modelling this. Skews (1990) studied the autorotation of rectangular plates with different thickness to chord ratios in a wind tunnel, showing that autorotation can be achieved for square cylinders. Tanabe and Kaneko (1994) presented a phenomenological model (one describing the empirical relationship of phenomena to each other without derivation from first principles) of a rigid and two-dimensional falling paper problem. Five distinct falling patterns—periodic rotation, chaotic rotation, chaotic fluttering, periodic fluttering, and simple perpendicular fall—were

identified. However, Mahadevan et al. (1995) commented that this model is not physically rigorous and doesn't improve upon simpler models. Field et al. (1997) found a similar set of behaviours predicted by Tanabe and Kaneko (1994) for disks falling water, showing a clear relationship between falling behaviour, Re and I^* . Belmonte et al. (1998) observed and modelled thin rectangular strips of paper, showing that falling motion can be divided into a fluttering and a tumbling phase, and that the transition between phases occurs at a relatively consistent Froude Number. Mahadevan et al. (1999) showed that, for rectangular cards, the frequency of tumbling scales with the square root of thickness over width.

Mittal et al. (2004) numerically investigated the flutter-tumble transition for falling objects, showing a higher tendency to tumble with a higher Re and lower thickness ratio. Pesavento and Wang (2004) presented the first "*solution*" to the problem, solving the Navier-Stokes equation for a two-dimensional rigid plate. Their model helps to further understand the transition from flutter to tumble, and has the potential for generalisation to flapping flight. Fonseca and Herrmann (2005) presented a three dimensional numerical study on the dynamics of a falling ellipsoid, identifying three falling behaviours: steady-falling, periodic oscillations and chaotic oscillations. Meanwhile, Fernandes et al. (2007) studied the motion of disks *rising* in water, showing that at a certain Re and body aspect ratio, velocity and position oscillation is induced. Similar work can be found in Fernandes et al. (2005). Jin and Xu (2008) developed a numerical model of elliptic and rectangular plates, showing that translation velocity between the two shapes is similar but rotational motion depends strongly on shape.

Chrust et al. (2013) numerically simulated the dynamics of freely falling disks, presenting an state-diagram of behaviours based on two independent parameters (Different to $Re - I^*$ described previously): the Galileo number—ratio between effects of gravity and viscosity—and nondimensionalized mass. Wang et al. (2013) investigate the influence of aspect ratio on the falling behaviour of rectangular plates. Varshney et al. (2013) reported one of the few studies investigating non circular/rectangular shapes, showing that falling parallelograms exhibit both autorotation and autogyration when falling. Huang et al. (2013) studied the influence of mass distribution on fluttering plates, finding that the symmetry of oscillation is lost when the centre of mass (COM) is shifted, resulting in more horizontal motion. At a critical COM, the plate falls vertically down with the COM leading. Auguste et al. (2013) numerically showed that disk aspect ratio has a large effect on the falling path and falling regime boundaries.

Heisinger et al. (2014) investigated coins falling in water, computing a probability density function for the landing point at different Re regimes e.g. steady, fluttering, chaotic. Each regime pdf exhibit a different distribution, for instance when tumbling, the pdf resembles

a ring around the centre point. V Vincent et al. (2016) showed that falling coins that poses holes are more stable, because of inner vortex ring formation. The findings have implications for robotic systems in free flight and more generally stability of gliding animals.

3.1.1 Flight, insects, wings

Ellington et al. (1996) studied airflow over a moth wing, showing that lift forces disagree with the conventional aerodynamic theories of the time. Dickinson et al. (1999) identified three mechanisms during insect flight: delayed stall, rotational circulation, and wake capture. Berman and Wang (2007) studied insect wing kinematics using a rigid model, finding the optimal kinematics for hovering. An interesting question is what morphology is optimal, and how does it affect kinematics - falling paper is passive flight. Bergou et al. (2007) studied a similar topic, finding that the pitching motion of insect wings during flight can be passive, and that fluid forces assist rather than resist pitching. Bomphrey et al. (2010) found that bumblebees exploit a leading-edge vortices for flight.

3.1.2 Seeds

Norberg (1973) presented a wide scale investigation of Samaras—winged fruit or seed that autorotates when falling—noting that their functional significance is to reduce falling speed to allow seed dispersal. The similarity between Samaras and insect wings is described. McCutchen (1977) investigated a range of Ash and Tulip Samaras. He suggests the stability and slow descent time of *Zanonia* Samaras when falling contributes to their evolutionary success. Lentink et al. (2009) measured three dimensional flow around Samaras, showing that lift is the result of a stable leading-edge vortex. This conclusion is generalised for the lift generated by hovering insects and bats. Varshney et al. (2013) studied the helical motion of falling maple seeds using high speed video, finding three distinct phases: tumbling, tilting toward vertical axis, gyration and settling. Furthermore, seeds with torn wings gyrate like full winged seeds. Lee and Choi (2018) presented a scaling law relating lift force, seed geometry and terminal rotational and linear velocities.

3.1.3 Snowflakes

Magono and Nakamura (1965) studied the fall velocity of snowflakes, showing a dependence on density and size. An empirical equation was derived to estimate falling velocity based on snowflake density. List and Schemenauer (1971) furthered this work, looking at the falling behaviour of different shaped snow crystals and hail. Drag coefficients and Best numbers—relating drag and Reynolds number—were calculated experimentally and used to estimate terminal velocities of different shaped particles. Jayaweera and Cottis (1969) presented similar work on the relationship between length, diameter and the falling velocities

of falling ice discs and cylinders. Finally, Christiansen and Barker (1965) showed that the terminal velocities of prisms, disks, cylinders and spheres are related to object density, fluid density and characteristic length ratios.

3.2 Physics-driven behavioural clustering

The PDBC method is inspired by the idea of dynamic similarity, which uses dimensionless quantities to assess the similarity between different systems whose properties are not necessarily the same. For example, the flow of two fluids with different densities in pipes of differing diameters are said to be similar if a dimensionless quantity – the Reynolds Number (Re) – is the same for each flow (Gerhart et al., 2016). Furthermore, the value of Re indicates the flow behaviour, e.g., laminar or turbulent. Hence, dimensionless quantities can be used both for clustering and physical insight.

We hypothesise that for dynamic systems with discrete behavioural modes there exists a set of functions whose parameter space is divided into distinct regions – separated by transitional zones – corresponding to different behavioural modes. Although not strictly dimensionless, we expect these functions to represent the underlying structure of dimensionless quantities, and hence term them pseudo-dimensionless quantities (PDQs). We further hypothesise that the more effectively PDQs cluster similar behaviours together, the greater physical insight they contain.

PDBC is a formalised approach for searching through and evaluating different PDQs. Figure 3.1 shows a schematic of the process, which I explain individually in detail below.

3.2.1 Data acquisition and processing

The PDBC method is intended for use with systems that exhibit discrete and distinctive behavioural patterns as certain system variables are changed. The input of PDBC is experimental data of such systems, containing a range of variables, behavioural patterns and outputs. These behavioural patterns should be distinguished during the data acquisition phase by the user, either visually or otherwise. Table 3.1 summarises each input of the process in detail.

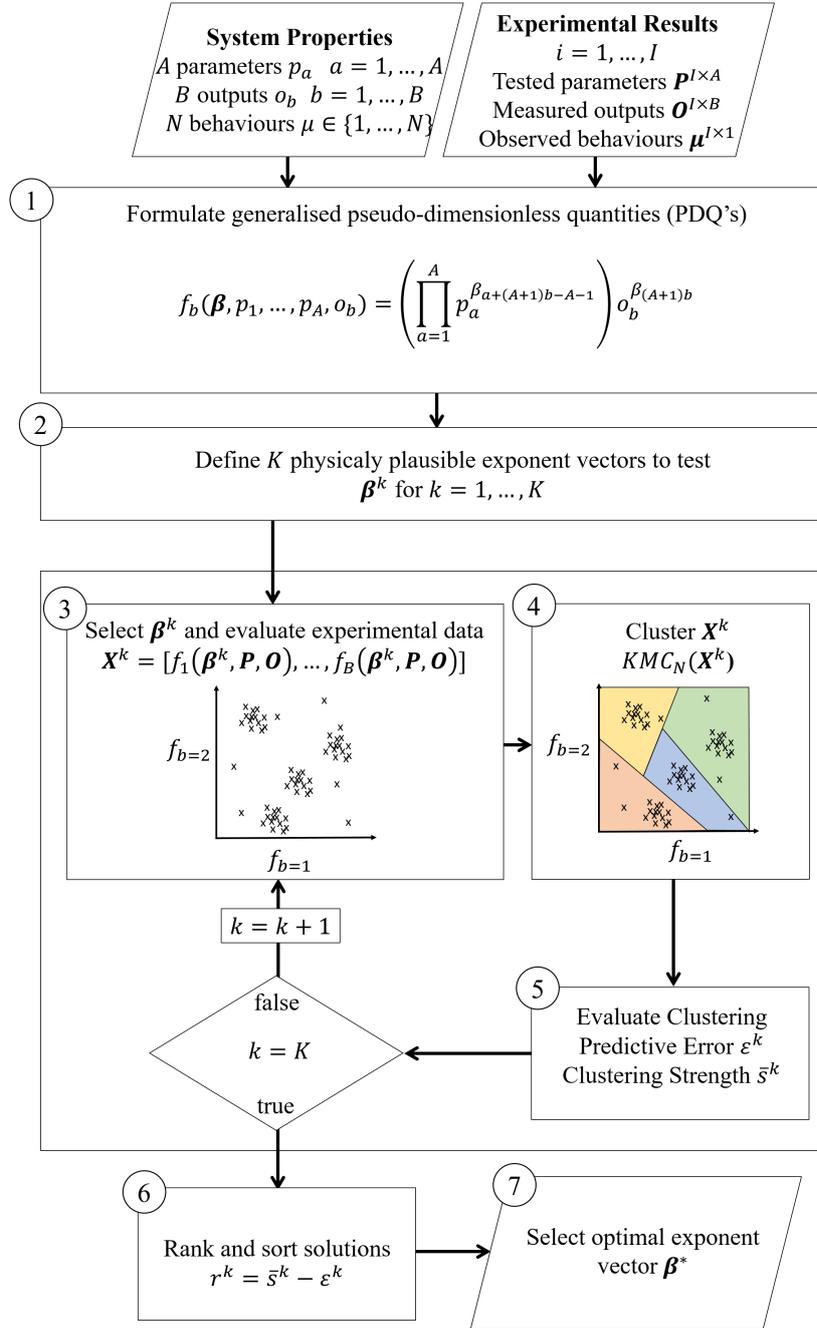


Figure 3.1: Flow chart of generalised physics-driven behavioural clustering method. Experimental observations with different system parameters along with their corresponding behaviours are the input. (1) A set of general functions called PDQs are formulated for the system. (2) Using a heuristic physics based approach, define a range of plausible values for the PDQ exponents. (3) Evaluate the PDQs with different combinations of exponents. (4–5) For each exponent combination, run a clustering algorithm in the PDQ parameter space and evaluate the predicted error of system behaviours, and the clustering strength. (6–7) Choose the exponent combination with the best trade off between minimising predictive error and maximising clustering strength.

Input	Symbol	Size	Description
System Parameters	p_a	$a = 1, \dots, A$	System parameters that can be varied, leading to different behavioural modes.
System Outputs	o_b	$b = 1, \dots, B$	System outputs that can be measured for all behavioural modes.
Behavioural Category	μ	$\mu \in \{1, \dots, N\}$	Numeric identifiers for each of the N behavioural modes observed in the system.
Experimental System Parameters	$\mathbf{P} = \begin{bmatrix} p_{1, \dots, A}^1 \\ \vdots \\ p_{1, \dots, A}^I \end{bmatrix}$	$I \times A$	Array of system parameter combinations tested over I experiments.
Experimental System Outputs	$\mathbf{O} = \begin{bmatrix} o_{1, \dots, B}^1 \\ \vdots \\ o_{1, \dots, B}^I \end{bmatrix}$	$I \times B$	Array of system outputs measured over I experiments.
Experimental System Behaviours	$\boldsymbol{\mu} = \begin{bmatrix} \mu^1 \\ \vdots \\ \mu^I \end{bmatrix}$	$I \times 1$	Matrix of observed behavioural modes observed over I experiments

Table 3.1: PDBC algorithm input parameters.

3.2.2 Formulation of generic PDQs (1)

The first step in the PDBC process is to formulate a set of generic PDQs. As previously stated PDQs are representations of dimensionless quantities, so should describe some relationship between the system parameters and outputs. A review of many common dimensionless quantities shows this relationship is usually characterised by the product of system parameters and outputs, raised to some exponent. Hence, the generic PDQs should facilitate the testing of many different combinations of the system inputs, outputs and exponents.

To satisfy this requirement, we formulate generic PDQs as exponential equations including the system parameters, outputs and generic exponents β_c for $c = 1, 2, \dots, C$, with $\boldsymbol{\beta} = (\beta_1, \beta_2, \dots, \beta_C)$ being the *exponent vector*. Each PDQ includes all system parameters

but only one output, with each term having one exponent. This allows us to specifically explore the relationship between the system parameters and each output. The total number of generic exponents, then, is $C = AB + B$ and the PDQs are formulated as follows

$$f_b(\boldsymbol{\beta}, p_1, \dots, p_A, o_b) = \left(\prod_{a=1}^A p_a^{\beta_{a+(A+1)b-A-1}} \right) o_b^{\beta_{a+(A+1)b-A}} \quad (3.1)$$

where as described in Table 3.1 $a = 1, \dots, A$ and $b = 1, \dots, B$. Using this formulation we can generate any number, say K , of exponent combinations $\boldsymbol{\beta}^k = (\beta_1^k, \beta_2^k, \dots, \beta_C^k)$ for $k = 1, 2, \dots, K$.

3.2.3 Exponent search policy (2)

Given the generic PDQs, we next define a policy to search through the possible exponent values previously described. The goal of this policy is to evaluate the physically plausible exponents for a particular parameter or output, while ignoring those which are physically unlikely. Hence, for each parameter and output the user should define an exponent range and discretisation increment that give rise to a set of plausible exponents. we denote this algebraically as follows. The a th parameter has an exponent range

$$\beta_{a+(b-1)A+(b-1)}^k \in \{-\Pi_a, -\Pi_a + \Delta_a, \dots, \Pi_a - \Delta_a, \Pi_a\} \quad (3.2)$$

and the b th output has an exponent range

$$\beta_{A+(b-1)A+(b-1)}^k \in \{-\Pi_b, -\Pi_b + \Delta_b, \dots, \Pi_b - \Delta_b, \Pi_b\} \quad (3.3)$$

where Π_a and Π_b define the minimum and maximum values for each parameter or output, and Δ_a and Δ_b the corresponding discretisation increment. The total number of exponent combinations K is therefore

$$K = \prod_{a=1}^A \left(\frac{2\Pi_a}{\Delta_a} + 1 \right) \prod_{b=1}^B \left(\frac{2\Pi_b}{\Delta_b} + 1 \right) \quad (3.4)$$

Clearly, setting Π_a , Π_b , Δ_a and Δ_b requires a heuristic approach. For example, parameters with units of length may relate to inertial terms or their inverses, so $\Pi = 4$, and could be discretised with $\Delta = 0.5$. Limiting the range or using a large increment may lead to useful PDQs being lost. However, increasing the range or using a low increment vastly increases the computational cost of the PDBC process.

The authors present Table 3.2 as a suggested guide for choosing reasonable exponent ranges for certain parameter and outputs types.

Term	Units	Π	Δ	Reasoning
Length	L	4	0.5	Includes inertial terms
Angle	—	4	0.5	Includes Inertial Terms
Linear Velocity	$L \cdot T^{-1}$	2	0.5	Includes Energy Terms
Angular Velocity	T^{-1}	2	0.5	Includes Energy Terms

Table 3.2: Suggested exponent ranges and increments for PDBC algorithm.

3.2.4 PDQ clustering (3–4)

We seek to determine how, given an exponent vector $\boldsymbol{\beta}^k$, the PDQs cluster similar behaviours together. To achieve this we apply the K-Means unsupervised clustering algorithm (Hartigan and Wong, 1979) on the PDQ parameter space. This partitions the experimental observations into N clusters – corresponding to the number of system behaviours – which can be evaluated for their predictive accuracy and clustering strength. As previously mentioned, we hypothesise that more physically meaningful PDQs will yield stronger and more accurate clustering.

We evaluate the experimental parameters and outputs \mathbf{P} and \mathbf{O} into the PDQ parameter space \mathbf{X}^k

$$\mathbf{X}^k = \begin{bmatrix} f_1(\boldsymbol{\beta}, P_1^1, \dots, P_A^1, O_1^1) & \dots & f_B(\boldsymbol{\beta}, P_1^1, \dots, P_A^1, O_B^1) \\ \vdots & \vdots & \vdots \\ f_1(\boldsymbol{\beta}, P_1^I, \dots, P_A^I, O_1^I) & \dots & f_B(\boldsymbol{\beta}, P_1^I, \dots, P_A^I, O_B^I) \end{bmatrix} \quad (3.5)$$

The K-Means clustering algorithm is applied to \mathbf{X} , yielding

$$\mathbf{v}^k = KMC_N(\mathbf{X}^k) \quad (3.6)$$

where N is the number of clusters to form, in this case the number of system behaviours, and \mathbf{v}^k is an I dimensional array of cluster assignments, with $v_i^k \in \{1, 2, \dots, N\}$. As is standard practice, the algorithm is run multiple times, three in this case, to avoid clustering anomalies.

3.2.5 Evaluation of clustering performance (5)

We introduce two measures of clustering performance; predictive error ε and clustering strength \bar{s} .

Predictive error: K-Means is an unsupervised method, so the cluster assignments in \mathbf{v}^k do not correspond to the behavioural labels in $\boldsymbol{\mu}$. In order to associate clusters assignments with behavioural labels we define $\hat{\mathbf{v}}_k$, in which we uniquely reassign cluster assignments such

that the fraction of misclassified behaviours – the *predictive error* ε^k – is minimised

$$(\hat{v}_i^k \neq v_i^k) \wedge (\exists! \hat{v}_i^k \in \{1, 2, \dots, N\}) \iff \min_{\hat{v}^k} \varepsilon^k \quad (3.7)$$

where

$$\varepsilon^k = \frac{1}{I} \sum_{i=1}^I e_i^k \quad (3.8)$$

and

$$e_i^k \begin{cases} 1 & \hat{v}_i^k \neq \mu_i \\ 0 & \hat{v}_i^k = \mu_i \end{cases} \quad (3.9)$$

Clustering Strength: We use the silhouette criterion (Rousseeuw, 1987) $s_i^k \in [-1, 1]$ to quantify the clustering strength. s_i^k is a measure of data consistency within clusters, representing how similar the i th observation is to its own cluster, relative to other clusters. The higher s , the stronger the clustering assignment for a particular observation is. We define the clustering strength as the mean of s_i^k for all observations, e.g.,

$$\bar{s}^k = \frac{1}{I} \sum_{i=1}^I s_i^k \quad (3.10)$$

3.2.6 Optimal exponent vector selection (6-7)

At this point in the PDBC process, all candidate exponent vectors have been evaluated for their predictive error ε^k and clustering strength \bar{s}^k . Hence, we must define a measure of optimality by which we sort the PDQs from the most to the least physically insightful. We denote the optimal PDQ exponent vector as β^{k*} .

Solutions with a low predictive error are desirable, as under our hypothesis these PDQs are likely to be more physically insightful. However, if the exponent search space is large there may be multiple solutions with a low predictive error; some arising from physical significance and some arising coincidentally. Hence, we must also consider the clustering strength of the solution, with stronger clustering also indicating more physical insight.

To achieve this we introduce the exponent ranking factor r^k , that rewards strongly classified solutions with low predictive error. It is simple the sum of $-\varepsilon^k$ and \bar{s}^k .

$$r^k = \bar{s}^k - \varepsilon^k \quad (3.11)$$

Hence, the optimal exponent vector β^{k*} corresponds to the highest rank r^{k*} , and represents the PDQs with the best trade-off between predictive error and clustering strength. To ease

the interpretation of the rankings we can sort the solutions in descending order in terms of r^k , defining the solution rank number $\hat{k}(r^k)$ such that

$$r^{\hat{k}(r^k)-1} \leq r^{\hat{k}(r^k)} \leq r^{\hat{k}(r^k)+1} \quad (3.12)$$

for $\hat{k}(r) \in \{1, \dots, K\}$. Hence, the optimal solution r^{k^*} corresponds to $\hat{k}(r^{k^*}) = 1$, with solutions $\hat{k}(r^k) = 2, \dots, K$ decreasing in their optimality.

3.3 V-shaped falling paper system

To test the PDBC method we created the V-shaped falling paper (VSFP) system, in which the passive falling behaviours of a V-shaped paper structure with an affixed mass are studied. The VSFP is a novel addition to the falling paper class of systems, and is to our knowledge unstudied. Here, we describe the VSFP system and experimental procedure, in the context of the PDBC method.

3.3.1 System morphology

Consider the passive falling behaviours of a V-shaped piece of paper with an affixed mass; together termed a *structure*. The morphology of a structure is fully defined by the four parameters shown in Figure 3.2: the wing length l , wing angle θ , wing width w and affixed mass m . l and θ may vary, while w and m are fixed at 10 mm and 5 g. Hence, the two system parameters to be used in the PDBC method are $p_1 = l$ and $p_2 = \theta$.

3.3.2 System behaviours

As the morphological parameters l and θ are varied, the passive falling behaviours change. During free-fall, structures exhibit a transient and steady state phase; when falling they may rapidly pass through more than one behaviour before settling on a single behaviour. In this study, we neglect the transient phase as we found it to be highly unpredictable. Hence, the output of each drop test is the steady state behaviour. Four behavioural modes are directly observable; plummeting, undulating, helicopter rotation and asymmetric rotation. Figure 3.3 shows example snapshots of each of these while Table 3.3 outlines the characteristics of each behaviour; see also supplementary video². The rotative behaviours (c,d) are easily distinguishable from each other and the non-rotative behaviours (a,b).

3.3.3 System outputs

There are many possible outputs, such as falling speed, rotation speed, rotation angle, oscillatory frequency or horizontal speed; some of these are only measurable in certain

²<https://doi.org/10.1371/journal.pone.0217997>

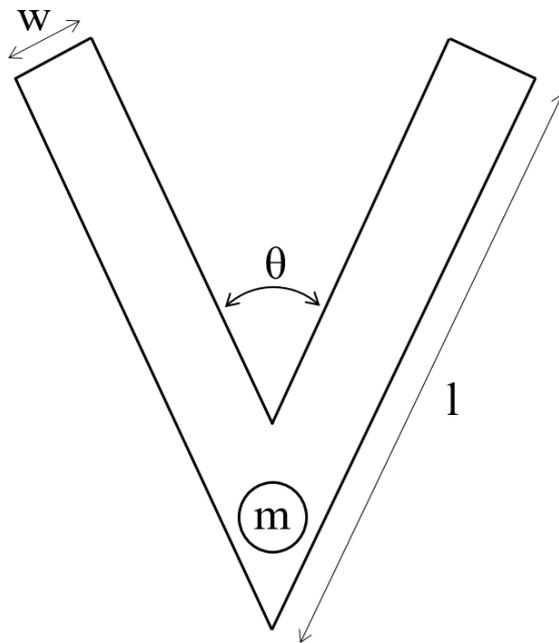


Figure 3.2: A Parameterised paper V shape with an affixed mass. The variable parameters are wing length l and wing angle θ . The fixed parameters are wing width $w = 10$ mm and affixed mass $m = 5$ g.

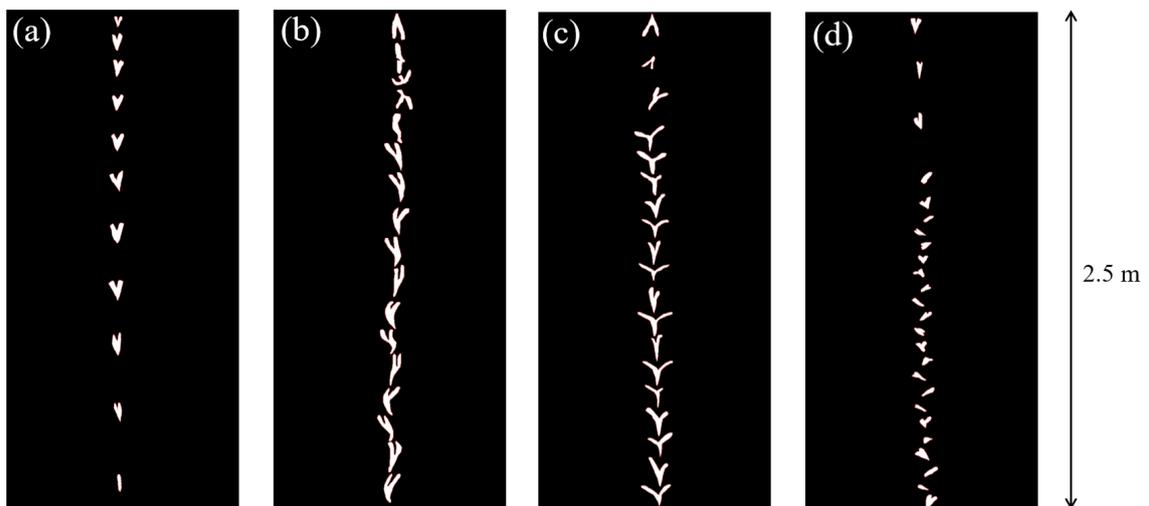
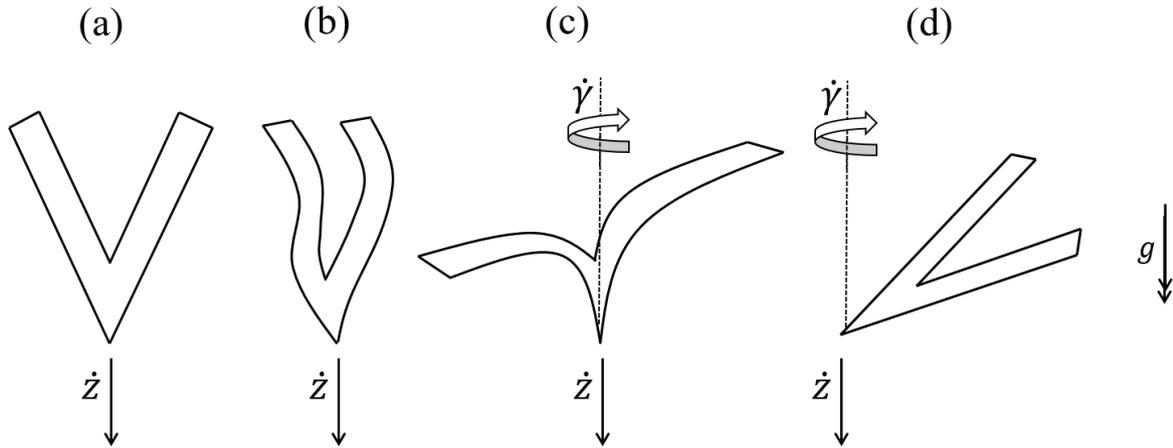


Figure 3.3: Snapshot images of the four directly observable behaviours in the VSFP system. (a) Plummeting (b) Undulating (c) Helicopter rotation (d) Asymmetric rotation. These snapshots show the structures falling height of 3 m.

Behavioural Mode	Characteristics
Plummeting	Shape falls directly to the ground with no wing movement.
Undulating	Shape falls directly to the ground, wings oscillate.
Helicopter Rotation	Wings splay in either direction, shape rotates to the ground.
Asymmetric Rotation	Shape rotates around mass.

Table 3.3: Qualitative description of behavioural modes in VSFP system

Figure 3.4: Diagrammatic representation of the four different behaviours, showing the measured outputs \dot{z} and $\dot{\gamma}$. For the non-rotative behaviours (a,b) $\dot{\gamma} = 0$.

behavioural modes. The PDBC method is designed to be used with *universally measurable outputs*, which we define as being observable in all behavioural modes. In the case of the VSFP system, this limits the outputs to falling speed \dot{z} and rotation speed³ $\dot{\gamma}$. Figure 3.4 shows a schematic of the different behaviours and universally measurable outputs.

3.3.4 Manufacturing

Figure 3.6 shows the experimental procedure. An Endurance MakeBlock XY engraving/cutting machine – as shown in Figure 3.5a – was to cut the shape out of Silvine A4 Graph Refill paper. The paper has a weight of 80 gsm (grams per square metre). The mass – for which 2 standard M4 steel washers were used – was affixed to the tip using superglue, with one washer on either side of the shape.

3.3.5 Testing

Each structure was manually dropped from a height of 3 m into still air and using a tip up initial condition, as shown in Figure 3.5c. Structures fell against a black backdrop, and were

³The rotational speed of the plummeting and undulating behavioural modes being negligible, but measurable.

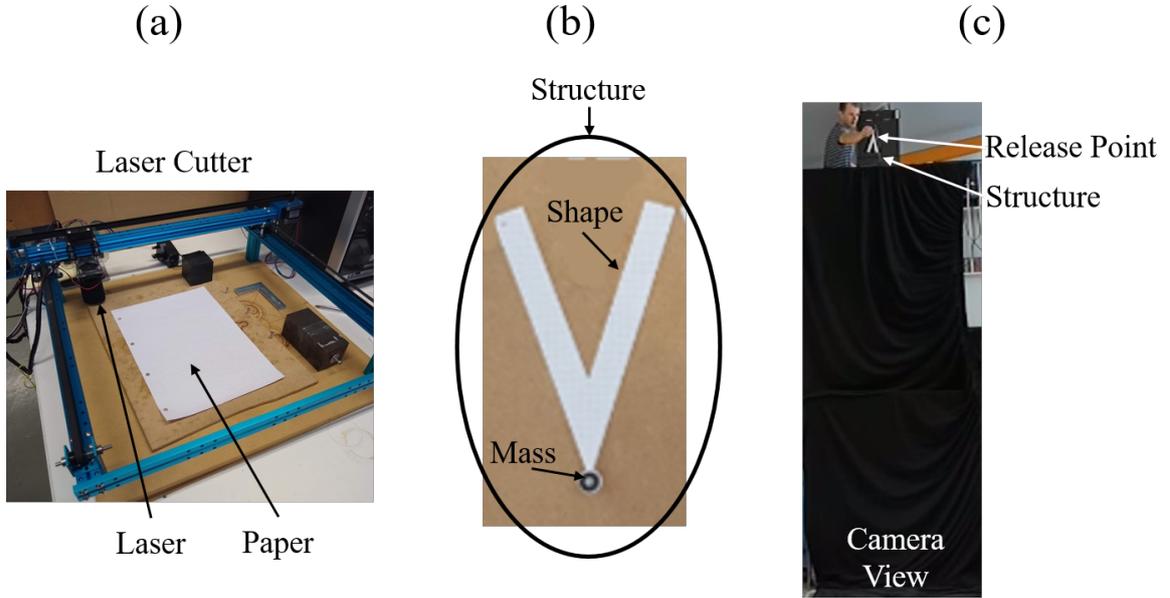


Figure 3.5: The experimental set-up, showing (a) Endurance MakeBlock XY engraving/cutting machine (b) Structure comprising of paper shape and affixed mass (c) Camera view of experimental drop zone

recorded using a Logitech BRIO camera recording at 120 fps. The system outputs $\dot{\gamma}$ and \dot{z} were manually extracted from the video data. Each structure was dropped $J = 5$ times, and the average outputs $\bar{\dot{\gamma}}$ and $\bar{\dot{z}}$ were calculated

$$[\bar{\dot{\gamma}} \ \bar{\dot{z}}] = \frac{1}{J} \sum_{j=1}^J [\dot{\gamma}_i^j \ \dot{z}_i^j] \quad (3.13)$$

3.4 VSFP experimental results

A series of structures were manufactured and their behaviours recorded, as previously described. The PDBC method was applied to these results with the aim of clustering the system behaviours and gaining physical insight into the system. In this section I describe the VSFP results, including the type of behaviours observed, their outputs and any relationship to l and θ .

The l - θ parameter search space was discretised such that

$$l \in \{75, 95, 115, 135, 155, 175, 195\} \text{ (mm)} \quad (3.14a)$$

$$\theta \in \{30, 37.5, 45, 52.5, 60, 67.5, 75\} \text{ (deg)} \quad (3.14b)$$

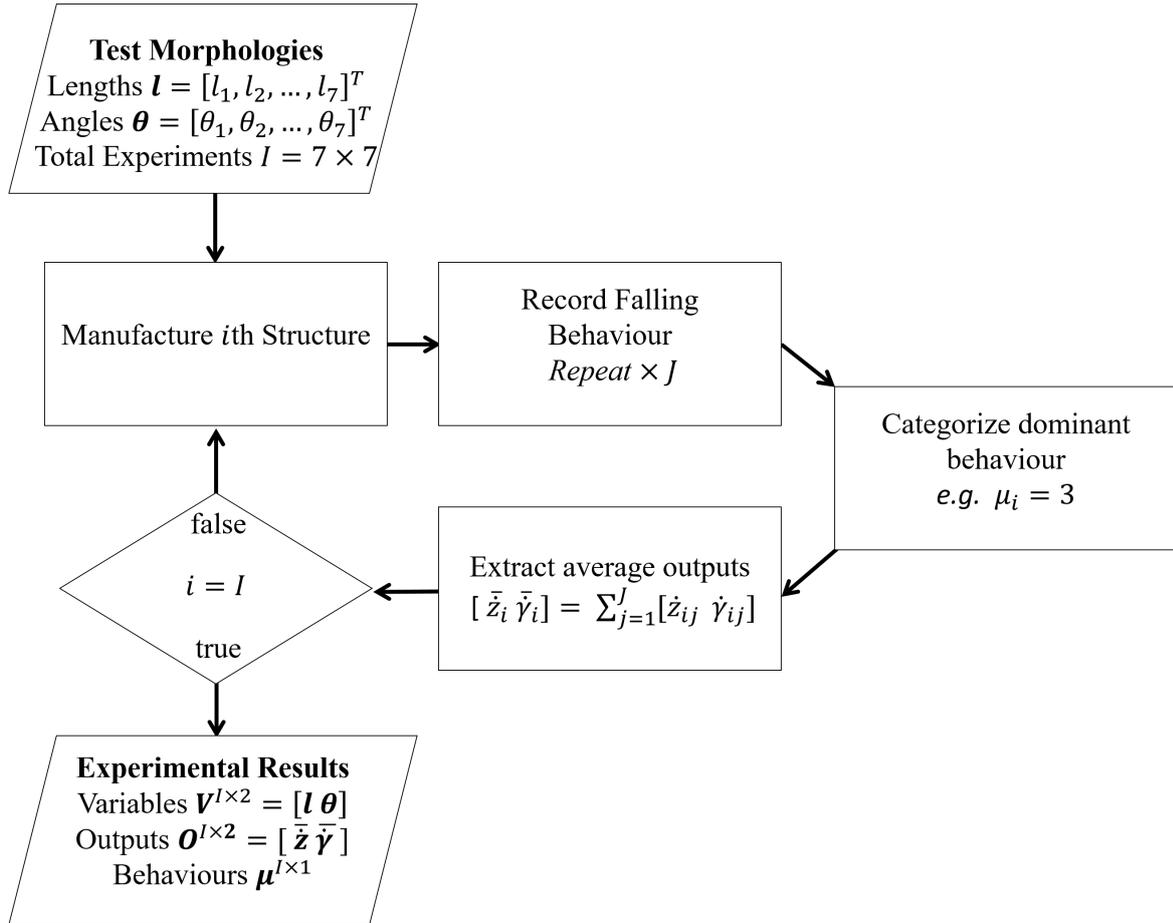


Figure 3.6: Flow chart of the experimental procedure. Structures are manufactured using the laser cutter. Structures are dropped and recorded $J = 5$ times. The dominant falling behaviour is extracted across these trials. The measured outputs are extracted and average over these trials, yielding \bar{z} and $\dot{\bar{y}}$. The process is repeated for every combination of l and θ in the search space.

Hence, a total of $I = 49$ structures were tested, some of which are shown in Figure 3.7.



Figure 3.7: Example of manufactured and tested structures.

First, I describe the results of these experiments.

3.4.1 Behavioural diversity and structure

There are four observable behavioural modes, as described in Table 3.3. Figure 3.8 shows the observed dominant behaviour as a function of l and θ , which were also stored in the behavioural ground-truth vector μ . There are five distinct behavioural regions; lines have been added by hand to indicate their approximate boundaries. Despite this apparent structure, there is no obvious rule to differentiate between behaviours based solely on l and θ . Morphologies with $l \geq 155$ mm are dominated by undulating behaviour across all angles except 30° . These morphologies have long wings with a range of angles. Morphologies with $l \leq 95$ mm are dominated by asymmetric rotation, except at the limits of $\theta \leq 37.5^\circ$ and $\theta = 75^\circ$. These morphologies have short wings with a smaller range of angles. Plummeting behaviours can be observed in morphologies with $l \leq 115$ mm and $\theta \leq 37.5^\circ$, and also morphologies with $95 \text{ mm} \leq l \leq 135 \text{ mm}$ and $\theta \geq 67.5^\circ$. Plummeting is the only behaviour observed in two distinct regions of the morphological search space, with the morphologies having either short wings with a low angle or mid-length wings with a high angle. The helicopter rotation region spans a range of l and θ . At the lower boundary l increases as

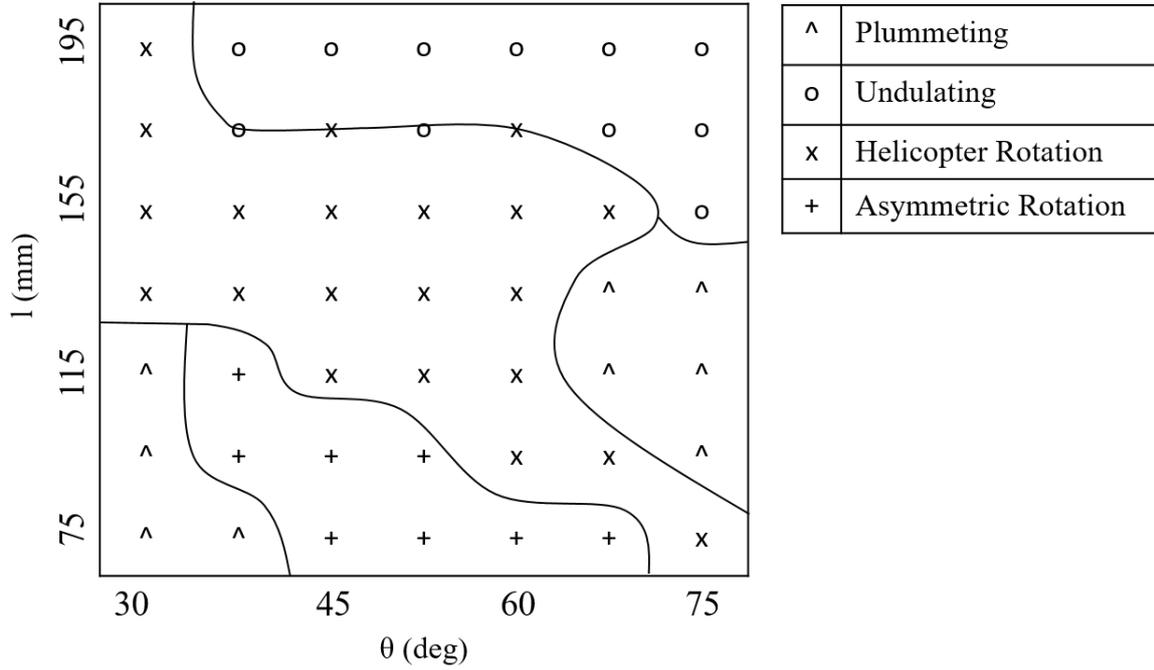


Figure 3.8: Dominant system behaviours across the morphological search space. Each behaviour is marked with a symbol, as shown in the legend. Lines have been added to estimate where the boundary between behaviours lies.

θ decreases. The upper boundary is less well defined, with a general transition from to plummeting behaviours.

3.4.2 Behavioural outputs

The behavioural outputs \bar{z} and $\bar{\gamma}$ were extracted. Figure 3.9a shows these outputs plotted against each other, as well a clustering regions which demonstrate the need for the PDBC method. There is a clear distinction between the rotating and non-rotating behaviours. The plummeting and undulating observations are non-rotating, i.e. $\bar{\gamma} = 0 \text{ rad}\cdot\text{s}^{-1}$ so the output space is one-dimensional. The helicopter and asymmetric rotation behaviours have non-zero $\bar{\gamma}$ and \bar{z} components. Plummeting behaviours range from $2.4 \text{ m}\cdot\text{s}^{-1}$ to $3.5 \text{ m}\cdot\text{s}^{-1}$ in \bar{z} and $0 \text{ rad}\cdot\text{s}^{-1}$ in $\bar{\gamma}$. Undulating behaviours range from $2.1 \text{ m}\cdot\text{s}^{-1}$ to $3.9 \text{ m}\cdot\text{s}^{-1}$ in \bar{z} and $0 \text{ rad}\cdot\text{s}^{-1}$ in $\bar{\gamma}$. Helicopter rotation behaviours range from $0.9 \text{ m}\cdot\text{s}^{-1}$ to $2.1 \text{ m}\cdot\text{s}^{-1}$ in \bar{z} and $3 \text{ rad}\cdot\text{s}^{-1}$ to $20 \text{ rad}\cdot\text{s}^{-1}$ in $\bar{\gamma}$. Asymmetric rotation behaviours range from $0.9 \text{ m}\cdot\text{s}^{-1}$ to $2.1 \text{ m}\cdot\text{s}^{-1}$ in \bar{z} and $5 \text{ rad}\cdot\text{s}^{-1}$ to $9 \text{ rad}\cdot\text{s}^{-1}$ in $\bar{\gamma}$.

In general, non-rotative behaviours almost exclusively fall faster than rotative behaviours. Within this, plummeting behaviours tend to fall faster than undulating behaviours. For rotative behaviours there is no clear behaviour that falls fastest or slowest. Nor is there an

obvious relationship between falling speed and rotation speed, with the range in \bar{z} being similar for both behaviours. However, the $\bar{\gamma}$ in helicopter rotation is around four times that of asymmetric rotation.

3.5 PDBC results

We applied to PDBC method to the VSFP system with the aim of discovering a set of functions to classify the $N = 4$ observed behaviours and infer physical significance from this clustering. The VSFP has two input variables, the length l and angle θ . It has two outputs, the mean falling speed \bar{z} and mean rotation speed $\bar{\gamma}$. Hence, $A = 2$ and $B = 2$ so we formulate two PDQs with a total of $C = 6$ exponent, yielding

$$f_1(\boldsymbol{\beta}^k, l, \theta, \bar{z}) = l^{\beta_1^k} \theta^{\beta_2^k} \bar{\gamma}^{\beta_3^k} \quad (3.15a)$$

$$f_2(\boldsymbol{\beta}^k, l, \theta, \bar{\gamma}) = l^{\beta_4^k} \theta^{\beta_5^k} \bar{z}^{\beta_6^k} \quad (3.15b)$$

where the exponent vector ranges were set following those described in Table 3.2

$$\beta_{1,2,4,5}^k \in \{-4, -3.5, -3, -2.5, -2, -1.5, -1, -0.5, 0, 0.5, 1, 1.5, 2, 2.5, 3, 3.5, 4\} \quad (3.16a)$$

$$\beta_{3,6}^k \in \{-2, -1.5, -1, -0.5, 0, 0.5, 1, 1.5, 2\} \quad (3.16b)$$

Hence, the total number of exponent vectors to test was $K = 6765201$.

3.5.1 Optimal exponent vector

After examining all exponent vectors, we ranked the solutions with respect to the criterion specified by (3.11). While most of these solutions do not cluster the experimental data well as shown in Figure 3.9b and (c), the highly ranked solutions show a clear structures in the output space, the best of which is shown in Figure 3.9d. The highest ranked solution has a predictive error $\varepsilon = 0.0204$ – corresponding to one miss-classified behaviour – and a clustering strength $\bar{s} = 0.8581$. The optimal exponent vector was

$$\boldsymbol{\beta}^* = (1.5, 0.5, 0.5, 3, 1, 0) \quad (3.17)$$

which corresponds to optimal PDQs of the form

$$f_1^* = l^{1.5} \theta^{0.5} \bar{\gamma}^{0.5} \quad (3.18a)$$

$$f_2^* = l^3 \theta \quad (3.18b)$$

The plummeting and undulating behaviours both have $f_1^* = 0$. Hence, they can be distinguished using only f_2^* , with the region $f_2^* < 0.24$ being characterised by plummeting behaviours and $f_1^* > 0.24$ characterised by undulating behaviours. The single misclassified behaviour is a plummeting behaviour that has been incorrectly clustered with the undulating behaviours at $f_2^* = 0.275$. The asymmetric rotation behaviours are tightly clustered together with $0.4 < f_1^* < 0.575$ and $f_2^* < 0.075$. The helicopter rotation behaviours are more scattered, with $0.79 < f_1^* < 1.25$ and $f_2^* < 0.35$. f_1^* can be used to distinguish between the helicopter rotation, asymmetric rotation and plummeting/undulating behaviours combined. Only when combined with f_2^* can all four behaviours be distinguished in the PDQ parameter space.

3.5.2 Solution landscape

As well as the optimal solution, we examined the landscape of all K exponent vectors. Figure 3.10 shows the r^k , ε^k , \bar{s}^k and $\beta_{1,\dots,6}$ for all tested solutions, sorted in decreasing rank order. The main characteristics of the solution landscape are presented here.

r^k , ε^k and \bar{s}^k : The top-ten highest ranked solutions all have $\varepsilon^k \leq 0.0612$, with six solutions having the minimum $\varepsilon^k = 0.0204$. Correspondingly, for these ten solutions $\bar{s}^k \approx 0.75$. Beyond this, the predictive error increases to a maximum of $\varepsilon^k = 0.5510$, while the clustering strength decreases to a minimum of $\bar{s}^k = 0.2959$. Across this trend, there are many solutions with a high \bar{s}^k , indicating strong clustering. However, they correspond to low ε values, so are not ranked highly. This shows that strong clustering can be achieved regardless of ε , reinforcing the need to consider both ε and \bar{s}^k . As ε increases and \bar{s} decreases, the r^k decreases to zero, as these solutions are neither distinguish between behaviours or exhibit strong clustering.

β_1 , β_2 and β_3 : These are the exponents corresponding to the first PDQ (3.15a). Over the top 1000 solutions, 55% of β_1 , 52% of β_2 and 98% of β_3 values remain within ± 0.5 of the optimum values of 1.5, 0.5 and 0.5 respectively. After this point, they begin to vary more. β_2 is limited to the range $[0, 2]$, since negative values were unable to be computed as they resulted in a division by zero.

β_4 , β_5 and β_6 : These are the exponents corresponding to the second PDQ (3.15b). Over the top 1000 solutions, 5% of β_4 , 7% of β_5 and 5% of β_6 values remain within ± 0.5 of the optimum values of 1.5, 0.5 and 0.5 respectively.

Figure 3.9b-c shows representative clustering solutions for the 100th and 1000th highest ranked exponent vectors. We can see that as the solution rank increases, the grouping of behavioural groups increases, while the separation between groups decreases. The exponent vector $\beta^k = (0, 0, 1, 0, 0, 1)$ corresponds to PDQs using just the raw measured outputs $\dots \gamma$

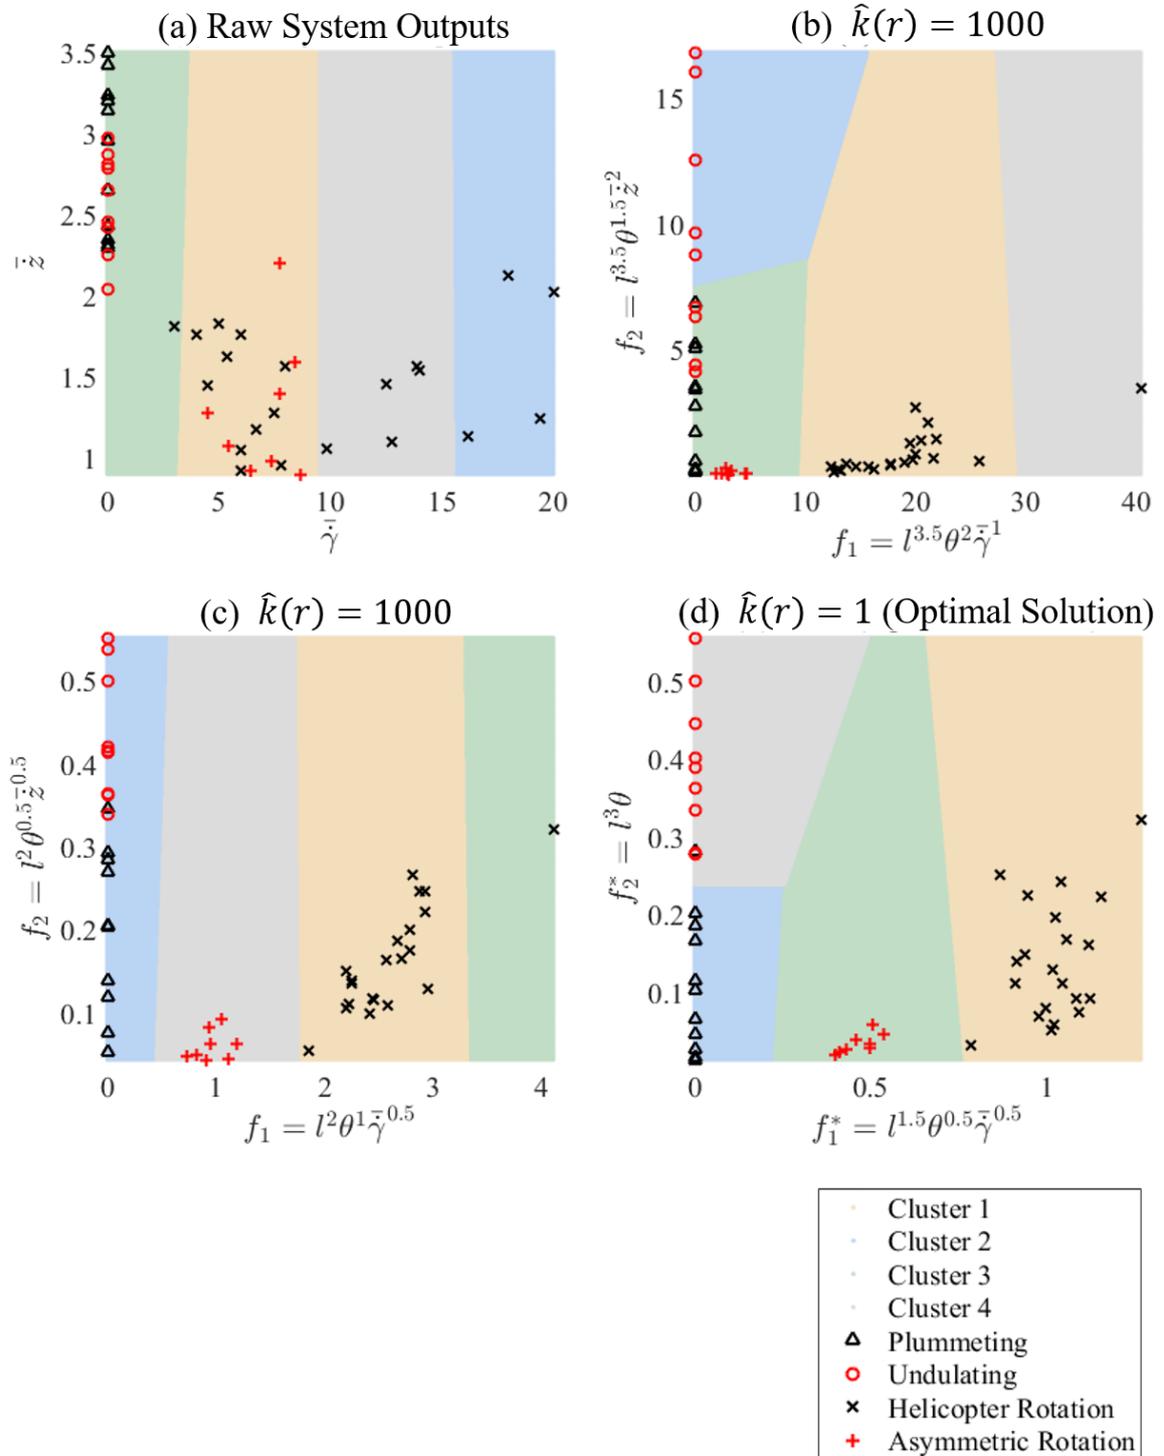


Figure 3.9: Clustering solutions for (a) raw system outputs, (b) 1000th ranked, (c) 100th ranked and (d) optimal exponent vector.

and \bar{z} . This is an interesting solution to examine as the PDBC method is predicated on the notion that the raw outputs alone are not enough to distinguish between system behaviours. This was confirmed, as the solution had a ranking number of $\hat{k}(r) = 35568$ with $\epsilon^k = 0.5102$ and $\bar{s}^k = 0.8550$. Figure 3.9a shows this clustering solution.

3.5.3 Physical significance

Inferring physical significance from the PDBC results is challenging, but some general statements can be made. We consider the optimal PDQs f_1^* and f_2^* separately.

$f_1^* = l^{1.5} \theta^{0.5} \bar{\gamma}^{0.5}$: The analysis of the solution landscape showed that the performance ranking was highly sensitive to exponent changes in this PDQ. This is particularly the case for $\bar{\gamma}$, which is strongly dependent on an exponent of 0.5. Hence, we can infer that this term is key in understanding each behaviour. Furthermore, we can show that f_1^* may represent some form of the Reynolds number Re . First, observe that the term $l\dot{\gamma}$ can represent the wing tip velocity of the rotative structures. Defining this as $V_{\text{tip}} = l\dot{\gamma}$, we can recast f_1^* as

$$f_1^* = (l\theta)^{0.5} (lV_{\text{tip}})^{0.5} \quad (3.19)$$

The Reynolds number is the ratio of a velocity and length term to the kinematic viscosity ν of the liquid under study. In the VSFP system ν is the kinematic viscosity of the air in which the structures fall, and remains unchanged between all experimental observations. Hence, $lV_{\text{tip}} = \nu Re$, and we can recast (3.19) as

$$f_1^* = (l\theta)^{0.5} (\nu Re)^{0.5} \quad (3.20)$$

Hence, following the physical meaning of Re , the rotative behaviours may be characterised by the ratio of inertial to viscous forces at their wing tip, relative to $l\theta$, which is a shape parameter representing the wing length and angle. For the non-rotative behaviours, this analysis doesn't apply as the rotation speed is zero.

$f_2^* = l^3 \theta$: The analysis of the solution landscape showed that the performance ranking is less sensitive to exponent changes in this PDQ. Interestingly, the optimal PDQ doesn't rely on \bar{z} at all. The terms in the PDQ are harder to interpret, but they tell us that the transition between the plummeting and undulating behaviours is governed by the wing length cubed multiplied the wing angle. This term is very similar to a moment of inertia term, indicating that this transition is related to the ease with which the structure can rotate or oscillate relative to the airflow.

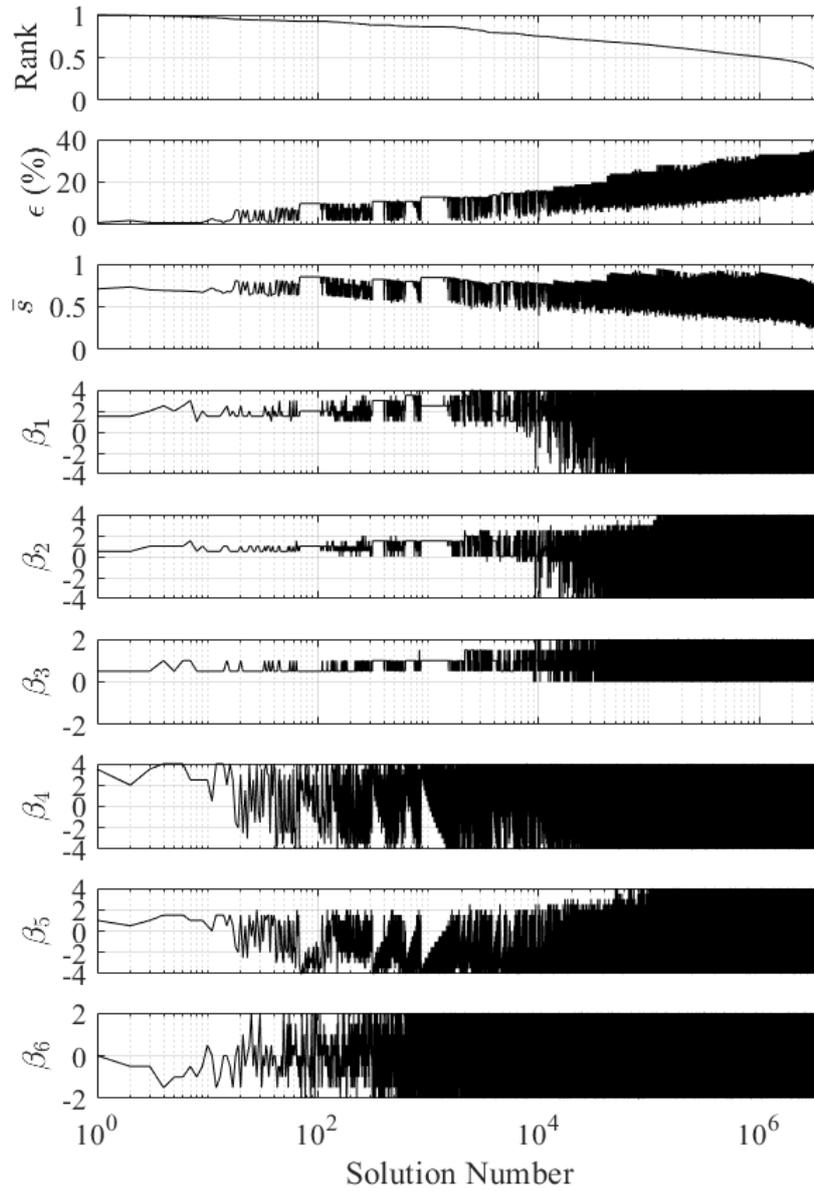


Figure 3.10: Solution landscape for all tested exponent vectors, showing the rank r^k , predictive error ϵ , clustering strength s^k and exponent values $\beta_{1,\dots,6}$.

3.6 Discussion and conclusion

This chapter presented the PDBC method as an algorithm for clustering and aiding understanding of systems with discrete behavioural modes. Furthermore, I presented the VSFP problem, a new category of falling paper systems, and applied the PDBC method to it. The results indicate that the PDBC method is an effective way of finding a parameter space in which behaviours can be clustered together with a high degree of accuracy, with the optimal exponent vector having a predictive error of just $\varepsilon = 0.0204$. In terms of physical significance, the optimal PDQs showed that behaviours can be clustered and categorised using a variant of the rotative Reynolds number on one axis, and a shape factor similar to the moment of inertia on the other. Interestingly, this is a relationship that is common among falling paper problems with (Chrust et al., 2013; Field et al., 1997), for example, reporting a similar behavioural relationship. Hence, this reinforces the hypothesis that choosing the most accurate and strongest clustered solutions reveals physically significant PDQs.

As described in the introduction, there are a range of data-driven algorithms for system understanding. Dynamic Mode Decomposition (DMD) (Kutz et al., 2016) can be used to discover physically meaningful modes and governing equations (Brunton et al., 2016) from high dimensional time series datasets. Meanwhile, the work of Schmidt and Lipson (Schmidt and Lipson, 2009) can distil free-form natural laws directly from time series data. PDBC is conceptually similar to these in that it aids in understanding complex systems, but also fundamentally different in its application.

PDBC is designed to give global insight into systems whose behaviours change across their parameter space. The intended usage is for systems with significantly different behaviours, such as those demonstrated in the VSFP system. In their current forms, the aforementioned alternatives are not well suited to this application. Rather, they would be effective in understanding the dynamics driving a particular behavioural mode. This is highly valuable, but does not provide the same global snapshot as PDBC. Indeed, the comparison is in some ways redundant as the two methods are in fact complimentary; PDBC provides the global picture, while alternative methods provide more specific insight of each behaviour.

More direct comparisons can be made with other approaches from the machine learning community. Support-vector machine (SVM) can be used to classify behaviours in real world systems via feature extraction. Gait analysis is one such example, with (Begg et al., 2005) using SVM for the automated classification of gait in young versus elderly human subjects. The extracted features and decision boundaries are similar to PDQs in that they define a behavioural parameter space. However, the physical relevance of these features is hard to interpret. In this type of application, PDBC would perform the classification while also

outputting physically interpretative PDQs. This would also be the case for other machine learning classifiers such as neural networks.

As previously discussed, in PDBC the user must initially assign behaviours to experimental observations. In the VSFP system this was a relatively simple task, as the observed behaviour were clearly different from each other, allowing a completely unambiguous classification. However, in general the users role in behavioural classification is significant. The user must decide what constitutes a behaviour, then apply this to the system observations. Hence, in the case of behavioural ambiguity, there may be no consensus among users regarding the total number of behaviours in the system.

As a short term solution, there are a few options. Firstly, to use a panel of observers to classify behaviours and reach consensus together. Alternatively, the PDBC algorithm can be run multiple times for each consensus. The solutions can be compared in terms of their solution landscape and physical significance. In the long term, however, the automated interpretation of behaviours presents an interesting challenge. In the case of the VSFP system, motion capture systems could provide a wealth of trajectory data for such a system. The authors hope to implement this into PDBC in the future.

Chapter 4

Transition behaviours in the V-shaped falling paper system¹

Morphological computation is another key area of research for self-structuring embodied agents (Müller and Hoffmann, 2017). Early examples showed the implementation of discrete computation, e.g., XOR gates, via morphology and control (Paul, 2006). Octopus-like tentacles have been shown to function as computation devices when integrated into the framework of reservoir computing (Nakajima et al., 2013). Soft matter computers have also been developed with the purpose of exploiting morphology for useful computation (Garrad et al., 2019). ‘Hot ice’ (sodium acetate) has even been proposed as a mechanism for large-scale parallel computation (Adamatzky, 2009).

In this chapter I present a sensitivity analysis investigating the relationship between morphology and behaviour in the VSFP system. The system can be characterised in terms of its steady behaviours, e.g., the system is in a single attractor state, and behaviour transitions, in which the system switches attractor states (Figure 4.1). Chapter 3 focused on steady behaviours in the system, showing a mapping between morphology and dominant behaviour.

¹This chapter is based on a peer-reviewed publication and a publication currently under review:

- Howison, T., Hughes, J., & Iida, F. (2020). Morphologically programming the interactions of V-shaped falling papers. In *Artificial Life Conference Proceedings*, pages 359–366.
- Howison, T., Hughes, J., & Iida, F. (2021). Morphological sensitivity and falling behaviours of paper V-shapes. *Artificial Life* [Forthcoming: accepted May 2021].

Contributions

T. Howison – devised the main conceptual ideas, carried out experiments, wrote the bulk of the paper and created the figures.

J. Hughes – assisted with carrying out experiments and data analysis, general discussions and proofreading.

F. Iida – contributed to general discussions and proofreading.

I describe the nature of stochastic and deterministic behavioural transitions in the system, and how morphology influences these.

This chapter is structured as follows. In 4.1 I explore how the stochastic nature of behavioural transitions in the system can be represented using switching diagrams. In 4.2 I introduce a new experimental setup that demonstrates deterministic behaviour switching and shows how paper shapes can behave as simple sequential logic circuits. In 4.3 I summarise and conclude how these different analysis approaches relate to each other in terms of the larger VSFP framework.

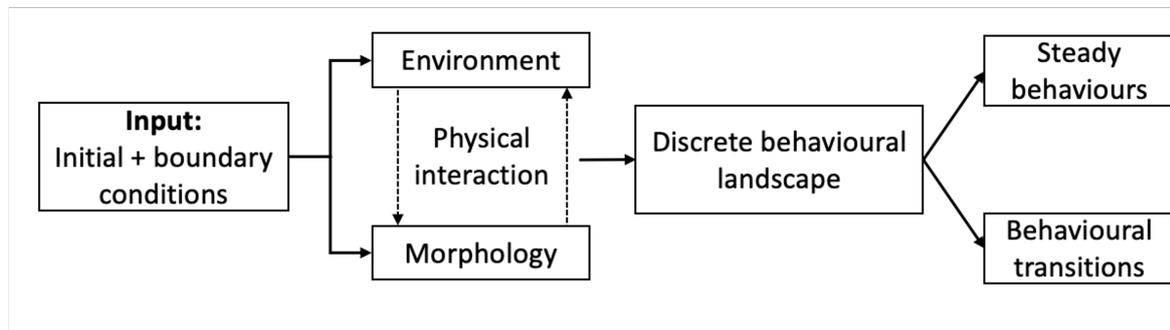


Figure 4.1: The discrete behavioural landscape in the VSFP system can be characterised in terms of steady behaviours and behavioural transitions.

4.1 Stochastic behaviour transitions

The ‘lifetime’ of a paper can be defined as that of when it is released to when it hits the ground, reaching its ‘dead’ state. In the VSFP system we have observed that during its lifetime a falling paper can transition between multiple behaviours. Figure 4.2a shows an example of the a transition from helicopter rotation to plummeting, while Figure 4.2b shows transition from undulating to helicopter rotation.

We draw a loose parallel between the transient behaviours of falling paper and developmental processes seen in nature. The physical interaction with the environment can induce a change in the dynamics of the falling paper, allowing it to switch to a new behavioural attractor state. This switching is completely self-induced since the only energy input is the initial potential energy, and there is only space for a certain number of switches before the shape reaches its dead state. For certain shapes, this switching appears to be highly unstable. For others, there seem to be attractor states from which the behaviour rarely deviates. In this section we investigate the nature of this stochastic behavioural switching and how morphology influences this.

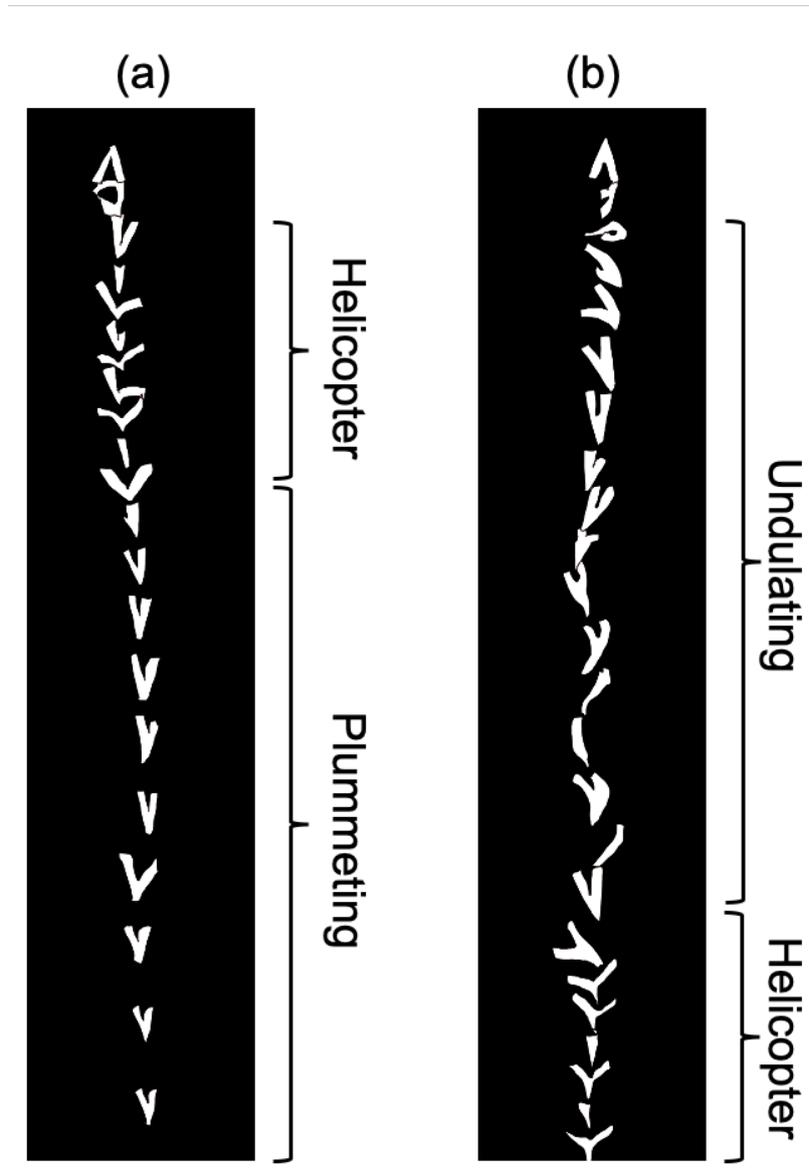


Figure 4.2: Examples of behavioural transitions in the VSPF system, showing (a) switching from helicopter rotation to plummeting and (b) undulating to helicopter rotation.

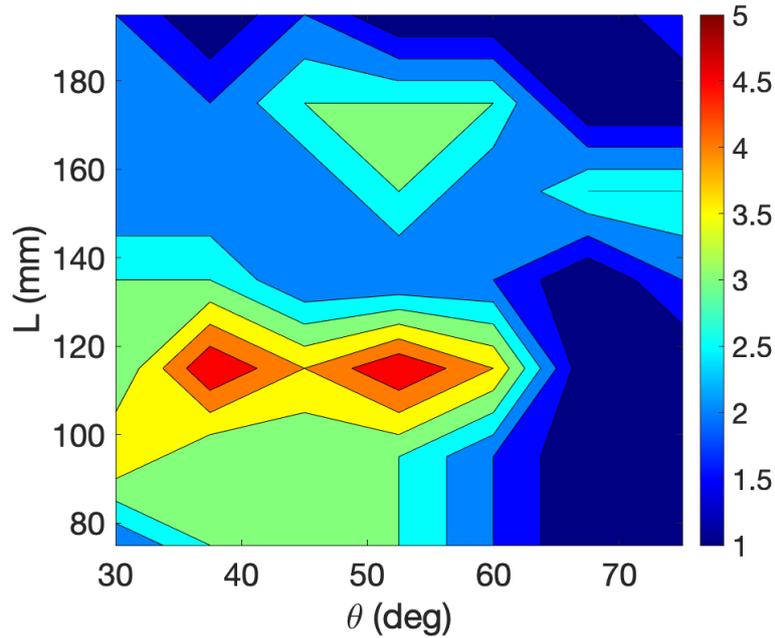
4.1.1 Stochastic transition sequences

We investigated the complexity of behaviour switching processes across the morphological parameter space. Each falling experiment is characterised by a sequence of states, starting from the initial condition and ending with the dead state on the ground. We can assign each possible state a label, e.g., 00 = initial condition, XX = dead state, AS = asymmetric rotation, HE = helicopter rotation, PL = plummeting, UN = undulating. Some experiments have a simple behavioural sequence, e.g., $00 \rightarrow HE \rightarrow XX$, while others can be more com-

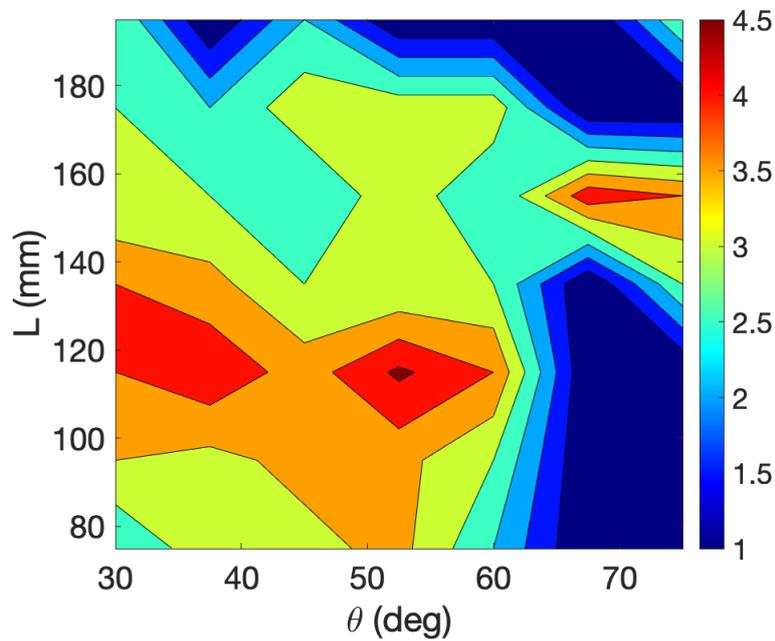
plex, e.g., $00 \rightarrow UL \rightarrow AS \rightarrow UL \rightarrow XX$. When observing the system we see that a particular morphology does not always exhibit a repeatable sequence of behaviours, sometimes instead showing stochasticity in terms of which behaviours emerge in what sequence. For each of the 245 experiments in our dataset of experiments across the morphological parameter space we constructed a behavioural switching sequence, e.g., $00 \rightarrow HE \rightarrow XX$. This was achieved using visual behaviour classification by a human observer. These were grouped in terms of morphology, allowing us to visualise the possible paths of a particular morphology through the behaviour space.

Figure 4.3a shows how the number of unique switching paths observed in the system changes based on morphology. Regions on the upper bounds of both l and θ generally exhibit far fewer switching paths than those elsewhere in the parameter space. These regions of the parameter space correspond to plummeting and undulating behaviours. Figure 4.3b shows the mean number of behaviour switches a particular shape is likely to exhibit based on the observed data. Again we see that morphologies on the extremes, corresponding to plummeting and undulating behaviours, generally exhibit simpler transitions with a low average number of behaviour switches. These diagrams show there is a loose relationship between the dominant behaviour regions identified in Figure 4.1c and the characterisation of transition events, suggesting that dominant behavioural analysis only offers a limited snapshot of the system dynamics.

To extend this analysis we diagrammatically visualised the different switching sequences across the morphological parameter space. The system can be represented as a graph in which nodes represent behaviours and edges represent behaviour transitions. The graphs are characterised by a path between nodes, starting at the initial state 00 and terminating at the dead state XX . Figure 4.4 shows some highlighted examples of these diagrams demonstrating morphologies with one, two three and five switches respectively: each unique transition sequence is indicated with a different colour. The diagrams do not indicate the relative probability of each path occurring: the sample size of five for each morphology is too low to assign confident probabilities. Figure 4.5 shows switching diagrams for each morphology in the tested parameter space. Some morphologies exhibit simple switching, showing only a single behaviour, e.g., helicopter. Other behaviours show far more complex switching sequences with many possible paths between initial condition and dead state. In most cases morphologies corresponding to dominant behaviours of plummeting and undulating exhibit less switching, in many cases showing no behaviours other than their dominant behaviour as indicated in Figure 4.1c. The exception to this trend is plummeting behaviours from morphologies with a low l and θ value, which seem to exhibit many switches.



(a) Number of unique switching paths.



(b) Mean number of switches.

Figure 4.3: (a) Number of unique switching paths across parameter space. (b) Mean number of behaviour switches across parameter space.

By counting all transitions across the parameter space we can assess stochastic nature of each behavioural attractor. Table 4.1 shows the likelihood of of behaviour transitions

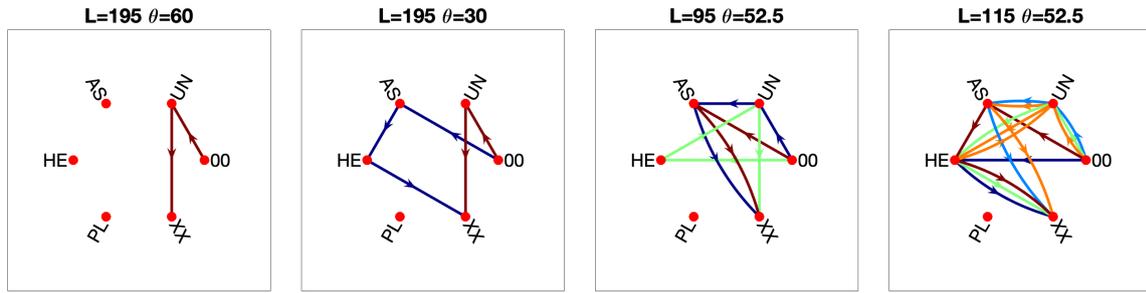


Figure 4.4: Highlighted transition sequence diagrams showing the possible behaviour transitions observe in the VSFP system. 00 = initial condition, XX = dead state, AS = asymmetric rotation, HE = helicopter rotation, PL = plummeting, UN = undulating. Each unique path through the behaviour-space is indicated by a different coloured line. A diagram with one line indicates all experiments for that morphology followed the same behavioural transition sequence. A diagram with more than one line indicates the range of possible paths taken across all experiments for that morphology.

from different initial states. Helicopter and asymmetric behaviours are the most likely to *not* transition (e.g., they go to the terminal state XX), while undulating behaviours are the least likely to reach the terminal state. We can infer from this that rotational behaviours are far stronger attractors than non-rotational. We can also see that helicopter behaviours are the most likely transition behaviours for undulating and asymmetric behaviours, indicating that overall helicopter behaviours are dominant in the system. We can also represent this stochastic switching across the population in the form of a probabilistic switching diagram, e.g., Figure 4.6. Here, we analyse all behaviour switches across the morphological parameter space and indicate the likelihood of moving from one state to any other (e.g., the probabilities from each state sum to 1). The graph representation is similar to Figure 4.4, but differs in that probabilities are explicitly shown and it represents a population of morphologies (rather than a single morphology). Switching probability is indicated by edge colour, visually demonstrating the probabilities outlined in Table 4.1. There are similarities with this representation and Markov chains (Norris and Norris, 1998). Both represent the system as a set of discrete states with associated transition probabilities. However, the switching diagram here does not include a concept of time-horizon that Markov chains rely on, and also represents a population of morphologies rather than a single agent. Markov models also make the assumption that future states are dependent only on current states, a proposition that would need further investigation to prove in our system. While this may be the case with the VSFP system, more work would be needed to prove this. It may be the case that the system could be represented with a higher-order Markov model that takes a number of previous states into account.

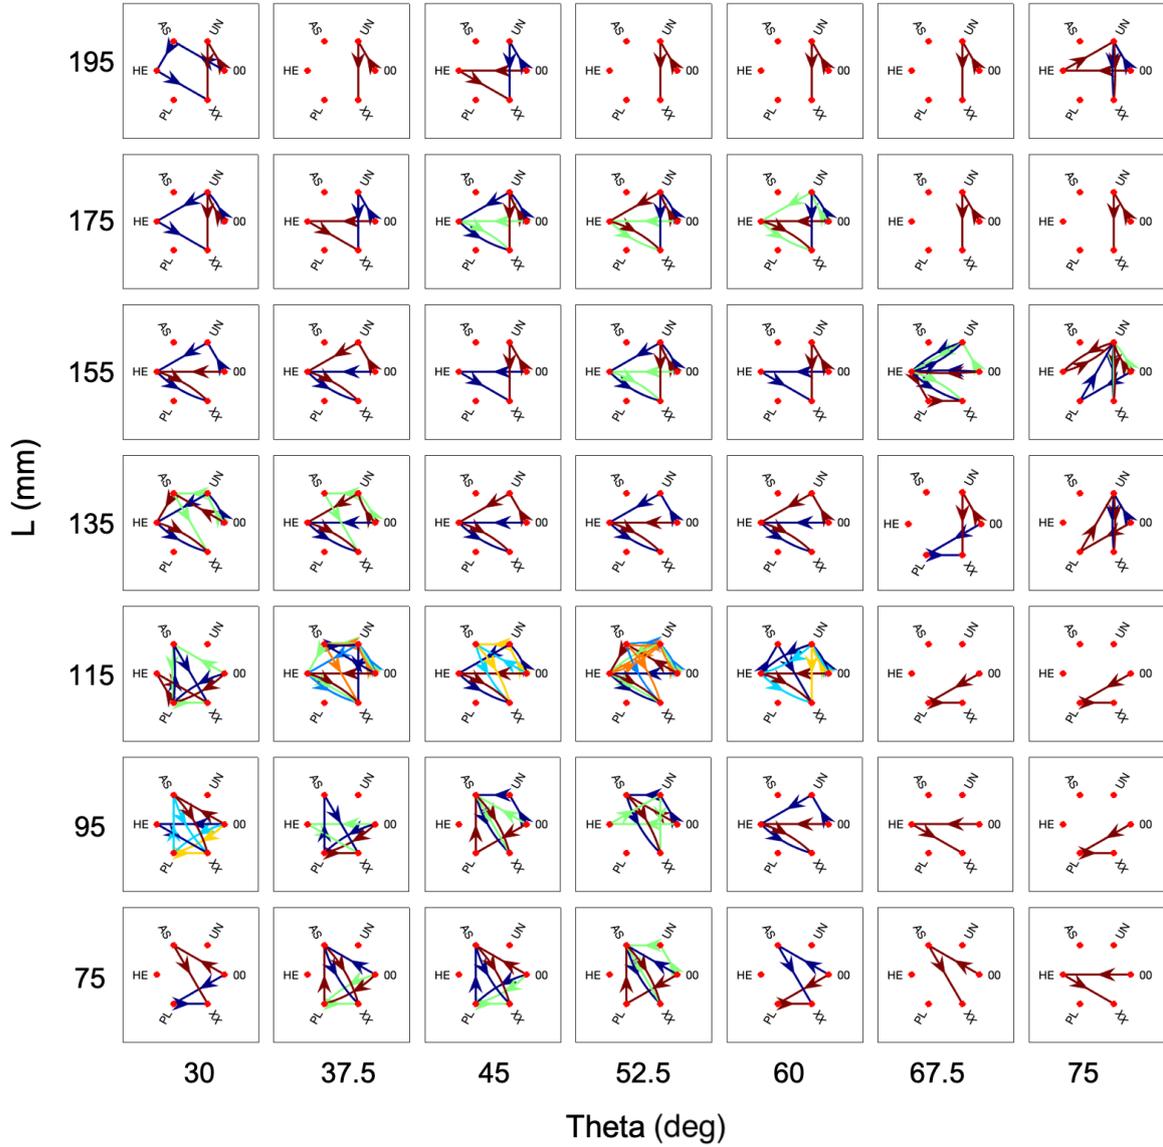


Figure 4.5: Transition diagrams across full morphological parameter space.

Initial state	UN	PL	AS	HE
Transition likelihood	XX (0.48)	XX (0.61)	XX (0.86)	XX (0.92)
	HE (0.38)	AS (0.30)	HE (0.10)	UN (0.06)
	AS (0.14)	UN (0.06)	PL (0.02)	AS (0.01)
	—	HE (0.03)	UN (0.02)	PL (0.01)

Table 4.1: Transition likelihood from each state, averaged across all experiments and morphologies.

In general, the framework here allows us to identify global trends for behaviour transitions and to identify attractor states that have a high probability of emergence, and may therefore be more energetically favourable or stable. The abstraction of passive transient behaviours to a probabilistic interpretation is similar to approaches in other systems. For example, the ‘conditional model’ describing passive hand behaviours that can be achieved only through conditional actions (Hughes et al., 2018).

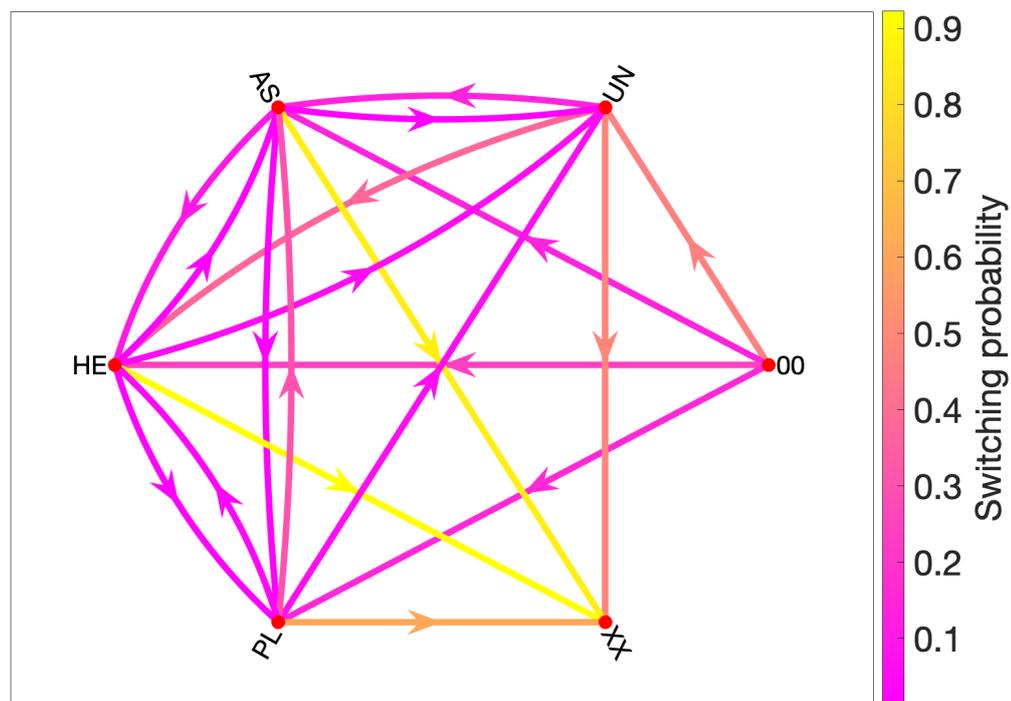


Figure 4.6: Stochastic switching diagram for all designs across the morphological parameter space in the VSFP system. Switching probabilities are indicated by colour. 00 = initial condition, XX = dead state, AS = asymmetric rotation, HE = helicopter rotation, PL = plummeting, UN = undulating.

4.2 Deterministic behaviour switching

Another aspect of transitional behaviours is deterministic switching, e.g., behaviour transitions that are completely predictable. In this section I describe a new experimental platform to investigate these transitions and demonstrate how morphology can perform simplistic sequential logic operations.

4.2.1 Rotational experimental setup

Observation of falling behaviours is limited by the available height in the experimental setup. Given that experiments should ideally be conducted in quiescent air (e.g., not outdoors), this means in many cases it is hard to investigate steady settled behaviours for long periods of time. To overcome this problem we developed a new experimental setup (Figure 4.7) which consists of a platform capable of towing shapes along a circular path in a horizontal plane.

The setup functions as follows. A DC motor is securely attached to a wooden base and rotates a lightweight thin wooden rod. On the end of this rod a paper V-shape can be affixed using thin wire and a rotational coupling that allows the shape to rotate freely with minimal friction. When the motor rotates the shape is towed in a circular path, from which we can calculate the equivalent linear speed, e.g., radius \times angular velocity. Shapes in this approximated falling system exhibit the same fundamental behavioural modes as those in the original system. From the system one can measure the rotational speed ω of the motor (using an optical encoder) and the current behaviour of the shape (by visual classification). Hence, we can explore how behaviours change as the system input is varied.

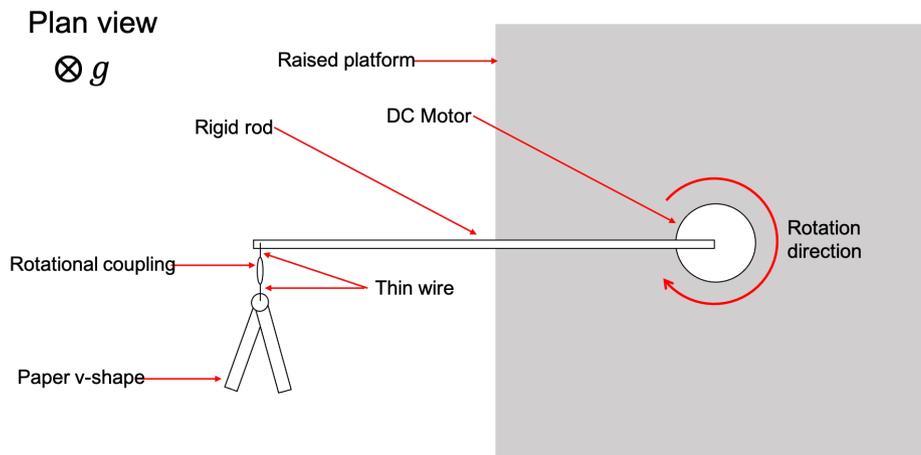
The system can be controlled by altering a PWM signal into the DC motor, mapping to voltage inputs between 0VDC and 6VDC, and is connected to a workbench power supply capable of supplying up to 3A. The equation for the rotational speed ω of an idealised DC motor is

$$\omega = \frac{V}{k} - \frac{T}{k^2}R \quad (4.1)$$

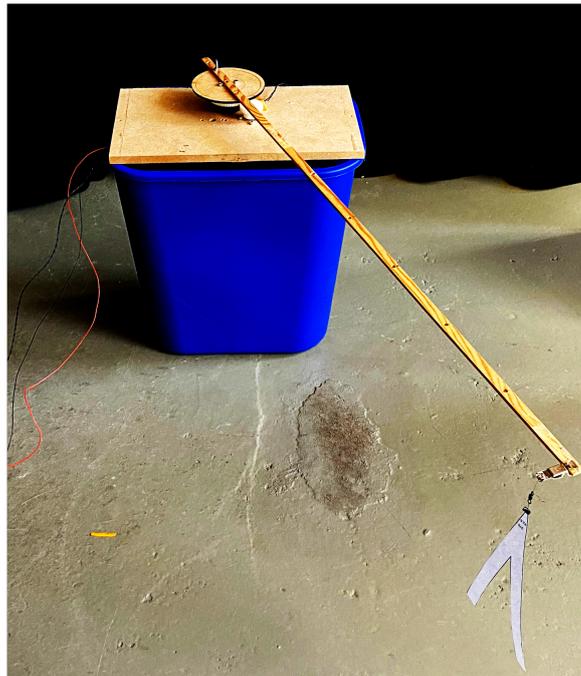
where R is the motor windings resistance, k is the motor constant, V is the input voltage and T is the load torque. R and k are constants. Hence, for a constant voltage V input, increasing load torque T decreases the rotation speed ω . In our system the load torque is dictated by the behaviour of the V-shape. For example, a helicopter behaviour is likely to induce a higher drag, and hence a higher torque than an undulating behaviour. Note, we ignore the torque effect of the rotating rod since this remains constant and the aerodynamic dynamics around the rod itself are likely to be minimal. It follows that ω will be lower for a fixed input voltage with a helicopter behaviour than an undulating behaviour.

This experimental setup allows us to roughly approximate the conditions of a paper shape falling over long periods of time. By changing the voltage, the effective load on the shape is changed, similar to the mass on the shape in true falling experiments. Hence, we can systematically test the behaviours of many shapes over extended time periods. Of course, the system is also subjected to centrifugal forces induced from the rotation. Further still, the loading in the new setup is via a thin wire, which influences how the shape rotates and adds a horizontal dimension to the loading (e.g., it is not purely unidirectional as in the falling

experiments). As such, we cannot claim this setup to be a comprehensive approximation of the falling experiments, and the following analysis is presented independently of the falling experiments.



(a) Plan view of setup.



(b) Photo of setup.

Figure 4.7: (a) Plan view of experimental setup for rotational experiments. Shapes are towed in a circular path, attached to a rigid rod with a rotational coupling. (b) Photo of experimental setup with shaped attached.

4.2.2 Behavioural hysteresis

In the first experiment we sought to induce behavioural transitions by incrementally increasing the input voltage to the system. Recalling (4.1), this increases the speed of rotation, given the torque remains constant. we fabricated ten paper V-shapes with θ fixed at 30° and l varying from 70 mm to 115 mm. For each shape, the motor voltage V was increased from 1.5VDC to 6VDC over 11 equally spaced increments. For each motor voltage the system was allowed to settle for 15 seconds before the rotational (and by extension linear) speed (averaged over 10 seconds) and visually identified behaviour was recorded. we found that 15 seconds was sufficient time for behaviours and speeds to settle. After the highest voltage was tested, the system was subjected to a decreasing voltage input from 6VDC back to 1.5VDC, following the same increments. Again, the rotation speed and behaviour were recorded. As such we were able to observe where behavioural transitions occur, and whether these points were the same for the ‘loading’ (increasing voltage) and ‘unloading’ decreasing voltage phases.

Figure 4.8a shows the results of this experiment. For each morphology we plotted the motor voltage V against the linear speed of the shape v . The behavioural modes have been indicated along with the points at which behaviour transitions occurred. The red lines indicate the loading phase and the blue lines indicate the unloading phase. Three distinct cases were observed:

Case one: purely asymmetric rotation. For morphologies with $l = 70\text{--}85$ mm we see that the only exhibited behaviour is asymmetric rotation. All morphologies exhibit an increase in linear speed with motor voltage, and show similar linear speeds during loading and unloading.

Case two: purely helicopter rotation. For morphologies with $l = 115$ mm we see that the only exhibited behaviour is helicopter rotation. Again there is an increase with voltage, however the speeds are lower than in the purely asymmetric rotation cases, indicating the drag forces induced by helicopter behaviours are larger than those from asymmetric.

Case three: transitional behaviours. For morphologies with $l = 90\text{--}110$ mm we see the most interesting response. In all these cases the behaviours are initially asymmetric rotation. As the voltage increases we see a behavioural transition to helicopter rotation. The point at which this transition occurs is dependent on morphology, with the transition happening earlier at higher l values. After the behaviour transition we see helicopter behaviour for all voltages during the loading phase. During the unloading phase we see that rather than switching back to asymmetric rotation at the transition voltage, the shapes exhibit helicopter rotation all the way to the minimum voltage. Helicopter and asymmetric behaviours have different linear speeds for the corresponding voltage input. We can infer that helicopter

behaviours induce a higher aerodynamic drag than asymmetric behaviours, hence slowing the rotation speed. The system exhibits a form of hysteresis in the speed measurements whereby the system output mapping depends on current and previous input. The strength of helicopter rotation behaviours reinforces our findings in 4.1 in which we showed falling shapes that entered a helicopter rotation attractor were most likely to exhibit this for the full experiment.

All experiments showed a range of linear speeds between 2 and 10 $\text{m}\cdot\text{s}^{-1}$. In our original experiments on the VSFP system (Chapter 3) we observed speeds of between 1 and 2 $\text{m}\cdot\text{s}^{-1}$ for rotative behaviours and 2 and 4 $\text{m}\cdot\text{s}^{-1}$ for non-rotative behaviours. We attribute this difference to the different loading on the shapes compared to the falling experiments, as well as the influence of the other differences summarised earlier in this chapter. Additionally, it was challenging to observe lower speeds in this experimental setup due to the inertial and frictional factors in the motor.

4.2.3 Permanent behaviour changes

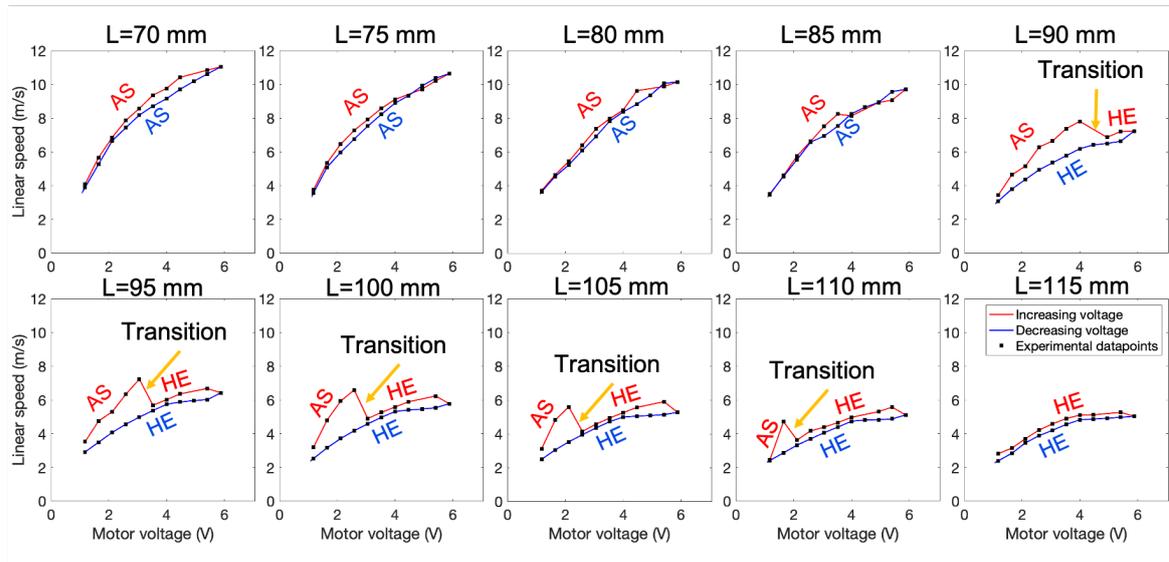
In our second experiment we subjected each morphology to a second cycle of voltage loading-unloading. After the first cycle the system voltage was set to zero and all energy allowed to dissipate. Following this we repeated the loading experiment. Figure 4.8b shows the results of this experiment. For each morphology we plotted the motor voltage V against the linear speed of the shape v . The behavioural modes have been indicated along with the points at which behaviour transitions occurred. The red lines indicate the loading phase and the blue lines indicate the unloading phase. Three distinct cases were observed:

Case one: purely asymmetric rotation. Similar to the first loading cycle, morphologies with $l = 70\text{--}85$ mm only exhibited asymmetric rotation behaviour.

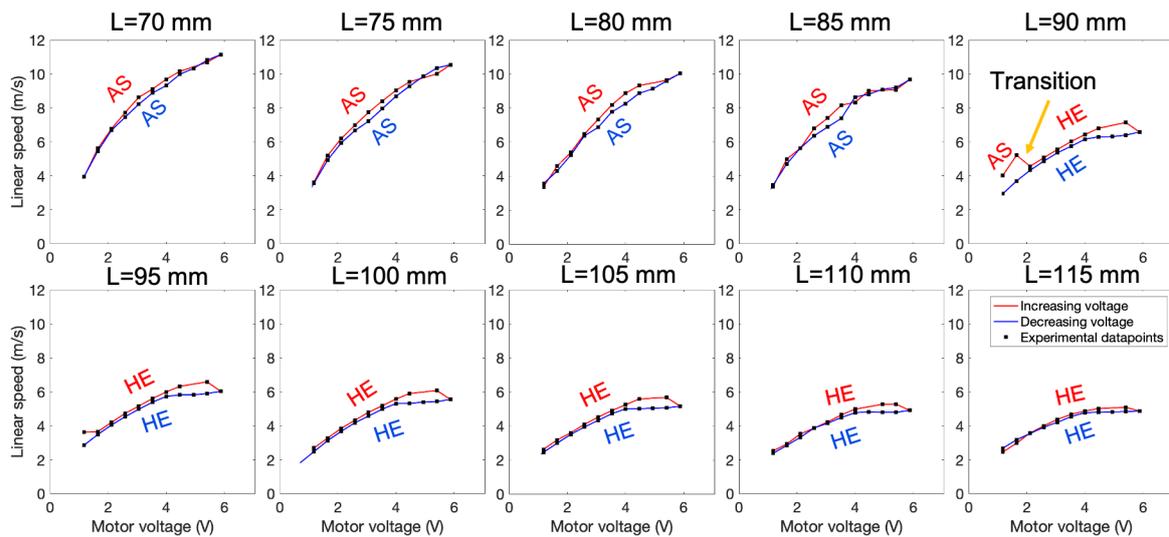
Case two: purely helicopter rotation. Unlike the first experiment, morphologies with $l = 95\text{--}115$ mm exhibit purely helicopter rotation behaviour. In loading cycling one morphologies with $l = 95\text{--}110$ mm exhibited transitional behaviours.

Case three: transitional behaviours Just one morphology, $l = 90$ mm exhibits transitional behaviours. The transition occurs at a far lower voltage than in the first loading cycle.

We can therefore see that during loading cycle one those morphologies that undergo a switch to helicopter rotation exhibit a permanent behaviour transition. Upon inspection of the shapes that exhibit this permanent behaviour change one sees that there is some form of plastic deformation separating the two wings. If we flatten the wings back into place this effect disappears.



(a) Load cycle one.



(b) Load cycle two.

Figure 4.8: Deterministic behaviours transitions. Motor voltage input vs linear speed for morphologies with θ fixed at 30° and l varying from 70 mm to 115 mm. Behavioural modes are annotated, as are transition points (HE = Helicopter behaviour, AS = Asymmetric Behaviour). Red indicates the loading phase in which motor voltage increases, while blue indicates the unloading phase, where motor voltage decreases. (a) shows the first loading cycling (b) shows the second loading cycle.

4.2.4 Behaviour transitions without hysteresis

The transition between helicopter and asymmetric rotation is distinct and apparently irreversible in our experiments. We also investigated the nature of other behavioural transitions

in our setup. Specifically, we investigated the transition between helicopter and undulating behaviours. Figure 4.9 shows the results for these experiments. The first key result is that this behavioural transition does not exhibit hysteresis. Behaviour transitions occur at the same point during loading and unloading, and the voltage to behaviour mapping is one-to-one. We also see unstable behavioural regions, marked as US in Figure 4.9. For $l = 180\text{--}190\text{ mm}$ for example, a range of input motor voltages induce unstable behaviours in which there is random switching between helicopter and undulating behaviours. Recalling the stochastic transition probabilities represented in Table 4.1 we see this random switching is consistent with the most likely transitions of helicopter and undulating behaviours.

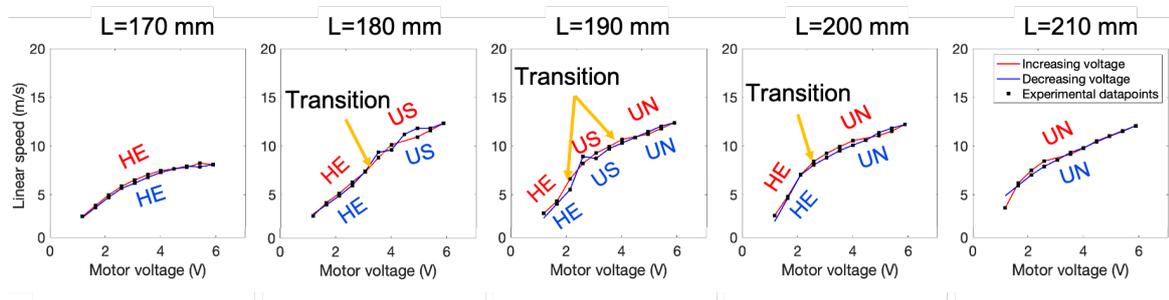


Figure 4.9: Deterministic and stochastic behaviour transitions. Motor voltage input vs linear speed for morphologies with θ fixed at 30° and l varying from 170 mm to 210 mm. Behavioural modes are annotated, as are transition points (HE = Helicopter behaviour, US = Unstable behaviour, UN = Undulating behaviour). Red indicates the loading phase in which motor voltage increases, while blue indicates the unloading phase, where motor voltage decreases.

4.2.5 Sequential logic representation

We now return to the behavioural hysteresis seen in 4.2.2. For some morphologies, a single voltage input can map to more than one behavioural output. Unlike the transition sequences presented in 4.1, these behaviour switches are predictable and deterministic.

The hysteresis effect in the system implies a form of memory. We can further qualify this as rate-independent hysteresis, since the current behaviour is dependent on past states but independent of any transient effects. Memory functionality forms one of the most basic forms of higher-level computation and cognition, and we can demonstrate that the falling V-shapes here can be represented as simplistic sequential logic circuits, aka memory cells. Take, for example the case of $l = 90\text{mm}$, which shows the biggest memory hysteresis effect. We can characterise the system voltage V input as either low (2VDC) or high (6VDC), or alternatively logical TRUE or FALSE. The discrete nature of behavioural outputs allow us to assign these as, for example, helicopter as logical TRUE and asymmetric as logical FALSE.

Given this simple mapping we can tabulate all possible combinations of input power, previous output $B(t-1)$ and current output $B(t)$, e.g., in Table 4.2a. We can see that this tabulation represents the truth table of a looped OR gate (Figure 4.10), see Table 4.2b, which is one of the simplest logic circuits that exhibits memory. Of course, we can arbitrary reassign behavioural outputs, for example setting Asymmetric as high and helicopter as low, leading to a looped NOR gate. Clearly this logical representation is hard to harness as a large-scale computation device. However, it adds to an increasing body of work showing how logic operations can be encoded into physical systems (Adamatzky, 2009; Adamatzky et al., 2017; Safonov, 2018).

$V(t)$	$B(t-1)$	$B(t)$
Low	Asymmetric	Asymmetric
Low	Helicopter	Helicopter
High	Asymmetric	Helicopter
High	Helicopter	Helicopter

(a) Current behaviour $B(t)$ as a function of current voltage input $V(t)$ and previous behaviour $B(t-1)$.

$V(t)$	$B(t-1)$	$B(t)$
FALSE	FALSE	FALSE
FALSE	TRUE	TRUE
TRUE	FALSE	TRUE
TRUE	TRUE	TRUE

(b) Corresponding truth table, showing system behaves as a looped logical OR gate.

Table 4.2: Logical representation of behaviours in rotational experimental system.

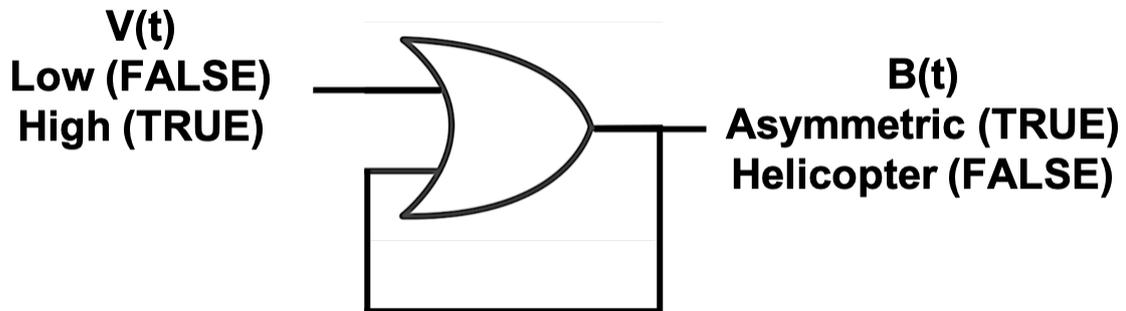


Figure 4.10: Sequential logic circuit. Behaviours transitions in the rotational experiment can be represented as a looped logical OR gate.

4.3 Discussion and conclusion

In this chapter I discussed the properties of behavioural transitions in the VSFP system. The first key contribution is demonstrating the stochastic behavioural patterns based on embodied interactions. Some shapes exhibit simple one-to-one mappings from morphology to behaviour while others exhibit a complex set of possible switching events that can occur

during the shapes lifetime. The number of behavioural switches a falling paper is likely to exhibit depends strongly on its morphological parameters. I demonstrated this transition structure by constructing switching sequence diagrams representing possible paths through the behaviour space, and how these vary with morphology. I showed that across the population of all morphologies we could infer probabilities for different behaviour switching events. From this, the system was found to be less likely to transition out of rotational behaviours. Assigning symbolic labels to system behaviours allows for this systematic analysis. Indeed, it opens up a wide range of possible analysis tools, for example using regular expressions (Thompson, 1968), which could be used to compare the relationship of many different behaviour sequences. A fundamental limitation is the available drop height, in this case 3 m. It is likely that increasing the available drop height for the VSFP system would induce more predictable final behaviours, e.g., there would be more time for transient behaviours to die out. However for multiple morphologies it seems that the behaviours are inherently unstable and would likely continue switching during their descent. Increasing drop height is an area for further study.

The second key contribution is investigating deterministic behavioural transitions. I described a new experimental setup using a rotational platform to investigate approximated falling behaviours over a longer distance. I showed that helicopter to asymmetric transitions can be induced by increasing the energy available to the system, and that these transitions are irreversible for certain morphologies, e.g., they exhibit hysteresis. I demonstrated that these behaviour transitions can be represented using simple sequential logic, implying the system has a memory effect whereby current behaviours are a known function of previous behaviours. Such logical operations in physical systems have been demonstrated before (Adamatzky, 2009; Adamatzky et al., 2017; Safonov, 2018), and while not necessarily useful in a practical sense, this representation demonstrates how such systems can perform rudimentary computation driven by environmental interaction.

The fact that the system can behave in either a stochastic or deterministic way seems to be linked to the environmental conditions. In true free-fall we cannot control the environment, so the paper shapes are unconstrained and small changes in initial conditions can lead to large changes in output behaviours. On the other hand, in the rotational experiment the environment is far more constrained, e.g. the tip of the shape is always connected to the towing line. There is further scope to investigate the influence of these differences in the context of singularity theory, a field of study first proposed in the late 1800's (Maxwell, 1873). Singularity theory deals with, in simple terms, understanding how arbitrarily small changes in a system can give rise to arbitrarily large effects, and is applicable to a wide range scenarios, including dynamical systems. Singularity theory tools may offer ways to

mathematically describe the variations in behaviour we see in the VSFP system (Bruce et al., 1992). Similarly, it could be beneficial to analyse the system in terms of its controllability, for example how easily can we induce switches between different behaviours. The fact we are able to discretize the continuous dynamics into a network of possible behaviours and transitions (e.g. in Figure 4.5) suggests that the system could be analysed from the perspective of controlling complex networks Liu et al. (2011).

Chapter 5

Bayesian optimisation of falling paper¹

The simplicity of falling paper systems allows us to explore the design of real-world behaviours. As we have discussed in Chapters 3 and 4, the VSFP system is highly complex, with a range of behavioural hierarchies across the design parameter space. This stochastic nature makes design challenging, as similar designs may appear to perform very differently with repeated testing. In simulation we could overcome this by using evolutionary algorithms with many thousands of iterations, e.g., (Cheney et al., 2014; Kriegman et al., 2017). However, this isn't feasible when relying on real-world experimentation. Testing designs in the real-world is, however, very important. As we have discussed in previous chapters, a reality-assisted approach is one possible route to designing agents that actively make use of their physical interactions. Furthermore, returning to the work of Brooks (Brooks, 1990, 1991), we should recall the physical grounding hypothesis that focused on the idea that robotic systems should be built in a bottom-up manner based upon a representation based on the physical world. In the case of falling paper, we can only capture the richness of their environmental interactions by testing them in the real-world: simulations are not yet detailed enough. Given that, it's important to explore the practical ways by which we can efficiently explore the design space of the VSFP system.

A promising approach for the structured optimisation of real-world designs is the Bayesian optimisation algorithm (BOA) (Frazier, 2018; Lizotte, 2008). This is a global black-box

¹This chapter is extended from the following peer-reviewed publication:

- Howison, T., Hughes, J., & Iida, F. (2020, July). Morphologically programming the interactions of V-shaped falling papers. In *Artificial Life Conference Proceedings*, (pp. 359–366).

Contributions

T. Howison – devised the main conceptual ideas, carried out experiments, wrote the bulk of the paper and created the figures.

J. Hughes – assisted general discussions and proofreading.

F. Iida – contributed to general discussions and proofreading.

optimisation approach for expensive-to-evaluate functions. BOA uses Gaussian process regression (GPR) (Rasmussen, 2003) to build a data-driven probabilistic system model that updates and improves with increasing function evaluations. This GPR model is in turn used to inform a sampling strategy to efficiently discover high-performing regions of the function parameter space. It has seen usage for controller learning (Calandra et al., 2016; Rieffel and Mouret, 2018), but its use as a morphology optimisation tool has been limited (Rosendo et al., 2017; Saar et al., 2018).

In this chapter I demonstrate the use of Bayesian optimisation to minimise the falling speed \dot{z} in the VSFP system. This problem is directly bio-inspired. Certain seeds in nature have evolved to fall slowly, allowing them to travel away from the parent tree (Lentink et al., 2009). Such behaviours rely heavily on the interaction between the seed and air. It is interesting to see if our system, within this albeit simple design space, converges to the similar behaviours.

This chapter is structured as follows. In 5.1 I briefly introduce the Bayesian optimisation algorithm and how it applies to the VSFP system. In 5.2 I present the results of the optimisation process. In 5.3 I discuss the results with a comparison to the performance of natural seeds seen in nature.

5.1 Bayesian optimisation

We start by defining the minimisation problem

$$\min_{x \in A} f(x) \quad (5.1)$$

where f is our objective function, x is our input and A is the feasible set of input values. In this case, f corresponds to the falling speed \dot{z} of a particular morphology and x to the morphological parameters $x = [l \ \theta]$. There are two key steps to each iteration of Bayesian optimisation; Gaussian process modelling and the acquisition function. We describe these here.

At any given stage in the optimisation, assume the objective function has been evaluated n times at the points $x_{1:n} = [x_1, \dots, x_n]^T$ and function values $f_{1:n} = [f(x_1), \dots, f(x_n)]^T$. These points are used as training data to generate a Gaussian process model of the system, which models the mean and variance of the objective function across the parameter space. This is achieved by first calculating the covariance matrix \mathbf{K} , for which we use the Matern kernel with a smoothness parameter of $5/2$,

$$k(x_i, x_j) = \sigma_f^2 \left(1 + \frac{\sqrt{5}r}{\sigma_l} + \frac{5r^2}{3\sigma_l^2} \right) e^{-\frac{\sqrt{5}r}{\sigma_l}} \quad (5.2)$$

where $i, j = 1, \dots, n$, $r = \sqrt{(x_i - x_j)^T(x_i - x_j)}$, σ_l is the characteristic length scale and σ_f is the signal standard deviation. The point-pair covariance matrix takes the form:

$$\mathbf{K} = \begin{bmatrix} k(x_1, x_1) & k(x_1, x_2) & \dots & k(x_1, x_n) \\ k(x_2, x_1) & k(x_2, x_2) & \dots & k(x_2, x_n) \\ \vdots & \vdots & \ddots & \vdots \\ k(x_n, x_1) & k(x_n, x_2) & \dots & k(x_n, x_n) \end{bmatrix} \quad (5.3)$$

The kernel hyperparameters – σ_l and σ_f – are determined by maximising their marginal likelihood, using local iterative gradient descent (Rasmussen, 2003). Following Bayes' theorem, this is the same as maximising the marginal likelihood of the training data, given the model parameters. Hence, the expected value $\mu(x_k)$ and confidence $\sigma(x_k)$ of any point x_k in the parameter space can be determined:

$$\mu(x_k) = \mathbf{k}_k^T \mathbf{K}^{-1} f_{1:n} \quad (5.4)$$

$$\sigma^2(x_k) = -\mathbf{k}_k^T \mathbf{K}^{-1} \mathbf{k}_k \quad (5.5)$$

where $\mathbf{k}_k = [k(x_1, x_k), \dots, k(x_n, x_k)]$. During each iteration of Bayesian optimisation, the Gaussian process model is updated to incorporate the new training data and better predict the system behaviour.

The Gaussian process model provides information about the expected objective function along with the prediction confidence. The second step of Bayesian optimisation is the acquisition function, which determines which data point to sample next. For this, an acquisition function is used. In this study we use the expected improvement (*EI*) acquisition function, which chooses sampling points based on the expected amount of objective function improvement based on the currently available data points. For any point x_k , the expected improvement is defined as

$$EI(x_k) = E[\max(0, \mu(x_{\text{best}}) - f(x_k)) \mid x_k \sim N(\mu(x_k), \sigma^2(x_k))] \quad (5.6)$$

The next point is therefore chosen by maximising the expected improvement, i.e.

$$x_{n+1} = \operatorname{argmax} EI(x) \quad (5.7)$$

As the Bayesian optimisation algorithm progresses, more data points are gathered and the Gaussian process model predictions become more accurate. Since we are optimising on a

potentially unknown black-box function, we cannot know for sure when an optimal solution has been found; this makes implementing a stopping condition somewhat challenging.

5.2 Optimisation results

We implemented the BOA using the MATLAB `bayesopt` function. The algorithm was operated with a fixed budget of iterations, after which the optimal solution corresponds to the best observed objective function. Following the experimental procedure set out previously, shapes were manufactured using the morphological parameters suggested by the BOA and a single drop test carried out. The falling speed \dot{z} was extracted from the video data and supplied to the BOA.

The optimal solution has morphological parameters $l = 82.9$ mm and $\theta = 51.1^\circ$, a falling speed of $v = 0.933$ m·s⁻¹ and corresponded to an asymmetric falling behaviour. Figure 5.1 shows various results from the optimisation process. The best observed falling speed after each iterations (the blue line in Figure 5.1) falls quickly from around 4 m·s⁻¹ to around 1.5 m·s⁻¹, after which it slowly decreases to below 1 m·s⁻¹. The GPR model estimated falling speed at each iteration follows a similar trend; however, it tends to overestimate the minimum falling speed. The regions in the behavioural landscape (Figure 3.8) corresponding to each iteration are primarily helicopter and asymmetric rotation.

Figure 5.2 shows the fitness landscape as estimated by the GPR model, along with the sampled points. The behavioural boundaries from (Figure 3.8) are also shown. We see that the majority of sample evaluations occurred in the rotation behavioural regions of the parameter space. We would expect this as these clearly have a lowered falling speed. The optimal solution is central in the asymmetric behavioural region, roughly as far as possible from the boundaries with adjacent behaviours. We speculate this is the most stable region of the behavioural landscape. Referring back to Figure 4.4 we see that this area corresponds to an average number of behavioural switches of around 1.5, generally lower than the helicopter rotation behaviours.

5.3 Discussion and conclusion

Bayesian optimisation offers a structured framework in which to optimise morphologies in the real-world. However, even with such a sophisticated search method we are unable to explore the system at a lower level, e.g., behavioural switching. This is one of the limitations of many design optimisation processes. In defining a fitness function and optimisation algorithm we immediately limit the power of the system (Lehman and Stanley, 2011a).

As mentioned, there is a clear parallel between the VSFP and natural seed dispersal systems in nature. Significant research has been carried out into analysing the dynamics of

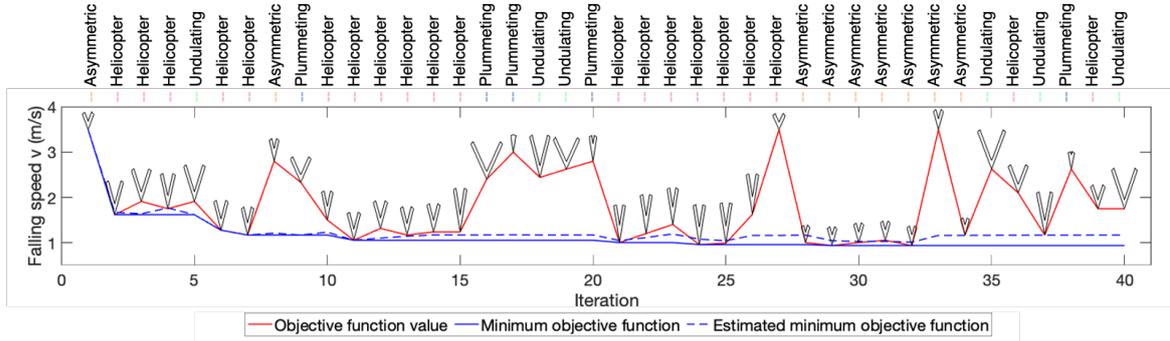


Figure 5.1: Bayesian optimisation results. The minimum observed and predicted falling speed is shown in blue, the observed speed at each iteration is shown in red. The morphology shape at each iteration is shown, as is the corresponding behaviour.

these natural systems (Nave et al., 2021). Here we compare the dynamics of the VSFP to natural systems using a commonly used metric and compare the optimisation results to this.

Following (Lentink et al., 2009) we employ an equation previously used to describe the aerodynamic performance of rotating and gliding seeds in nature:

$$\dot{z} = \sqrt{\frac{W/S}{0.61DF}} \quad (5.8)$$

where \dot{z} is the falling velocity, W/S is the wing loading normalised over the area a , e.g., mg/a and DF is the descent factor, a dimensionless number representing the falling shapes performance. We calculated W/S for each morphology tested in the dataset and combined this with the average velocity measurements and known dominant behaviours to compare how the morphologies tested in this chapter compare to natural seeds. Figure 5.3 shows the results from our experiments alongside a summary of datapoints for natural seeds presented in (Lentink et al., 2009). Also indicated are contours of equal descent factor. Note that in the plot the y axis shows the inverse of \dot{z} , e.g., the time to fall 1 meter. We see that the performance of the VSFP system in terms of falling time is a factor of 2.5 worse than the best seeds seen in nature. However, the majority of natural seeds have a similar falling speed as the rotational behaviours in the VSFP system. Natural seeds are significantly lighter than shapes in the VSFP system, 100–200 mg compared to 5 g. Hence, the descent factor is significantly higher for the slowest falling papers. We see the Bayesian optimisation optimal design has the highest descent factor, as expected as it minimises falling speed for the given mass.

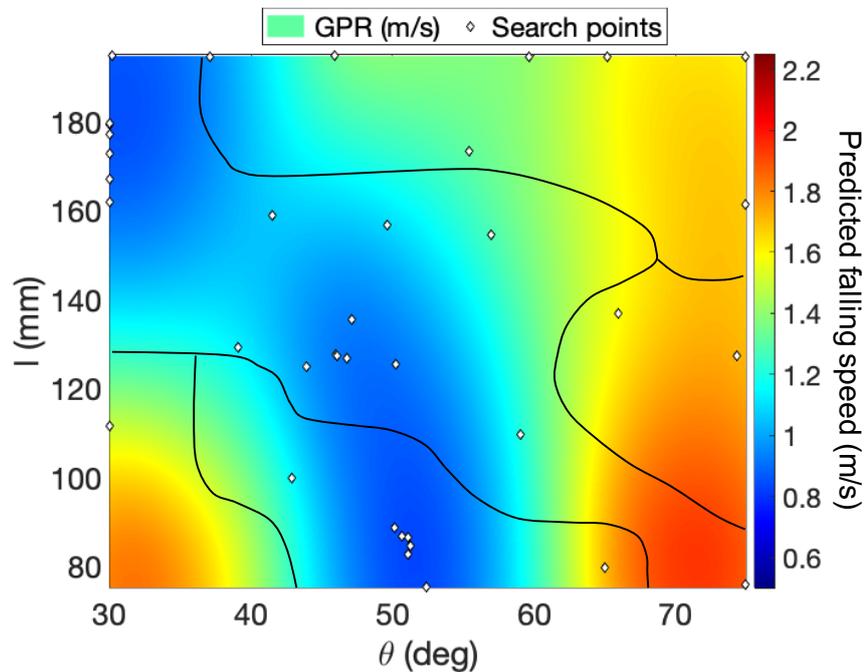


Figure 5.2: Bayesian optimisation sampling strategy and modelling. The Gaussian process regression (GPR) model shows the predicted falling speed at each point in the parameter space. The white markers indicate where the BOA sampled within the parameter space.

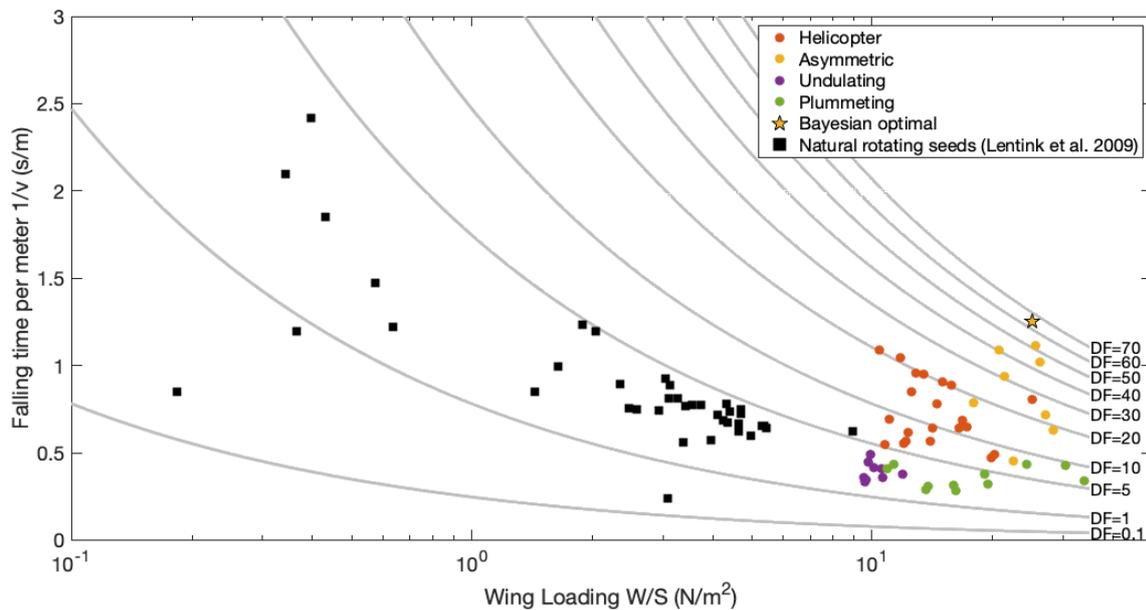


Figure 5.3: Comparison of optimised performance to natural systems, showing the relationship between falling time per meter and wing loading.

Chapter 6

Large-scale automated investigation of free-falling paper shapes via iterative physical experimentation¹

Falling paper systems are very challenging to model. Of course, this is a more general problem for embodied systems in which behaviours emerge from complex environmental interactions. In the absence of any analytically tractable solution one must rely on physical experimentation, which may be labour-intensive, time-consuming and highly sensitive to changes in environmental and initial conditions. Additionally, the search space required to observe an insightful range of behaviours can be extremely large. In this chapter I demonstrate the use of robotic automation, computer vision and machine learning to investigate the falling-paper problem. Robotic automation can be used to cheaply gather large volumes of experimental data (Waltz and Buchanan, 2009), with the added benefit of having much greater control over experimental initial conditions (Peplow, 2014). Meanwhile, machine learning offers continual, online data analysis (Mjolsness and DeCoste, 2001) to interpret data and suggest efficient sampling strategies (Bottou et al., 2018). The use of robotics in science

¹This chapter is based on the following peer-reviewed publication:

- Howison, T., Hughes, J., & Iida, F. (2020). Large-scale automated investigation of free-falling paper shapes via iterative physical experimentation. *Nature Machine Intelligence*, 2(1), 68–75.
- Data and code availability: <https://github.com/th533/Falling-Paper>

Contributions

T. Howison – devised the main concept, constructed the experimental platform, coded the software tools, wrote the bulk of the paper and created the figures.

J. Hughes – assisted with early proof of concept experiments, monitored experimental platform, assisted with carrying out experiments and data analysis, early paper drafts, general discussions and proofreading.

F. Iida – contributed to general discussions and proofreading.

research, the so-called ‘robot scientists’ (Fan et al., 2019; Soldatova et al., 2006; Sparkes et al., 2010a), is well documented. In biology, robots have been used to automate lab processes (Chapman, 2003) and carry out 1000s of experiments to efficiently identify gene functionality (Kachel et al., 2006; Sparkes et al., 2010b). When combined with artificial intelligence these systems can even automate hypothesis creation (King et al., 2004) and drive future discoveries (Vasilevich and de Boer, 2018). Other methods for exploring high-dimensional spaces for scientific purposes range from reinforcement learning to guiding exploration using intrinsic motivation or entropy (Bellemare et al., 2016; Tang et al., 2017). Meanwhile, investigative methods for understanding non-linear dynamics in complex systems using experimental data (Frankel and Reid, 2008) and extracting physically meaningful expressions have been presented (Bongard and Lipson, 2007; Brunton et al., 2016; Schmidt and Lipson, 2009). However, there has been limited application of ‘robot scientists’ to the broader physical sciences.

We propose an iterative physical experimentation system (IPES; Figure 6.1b), in which paper shapes are fabricated, experiments are conducted, data is analysed, and a sampling strategy is implemented, all without human input. In our illustration of this approach, hundreds of different sized circular, square, hexagonal and cross shapes were fabricated, dropped and their falling behaviours classified without human intervention. The number of experiments was an order of magnitude greater than those typically seen in the field, but with a greatly reduced effort. The behavioural classification approach was compared to visual classification from a panel of human observers, showing an average agreement of 86.2%. The system reproduced the original work on falling disks (Field et al., 1997), showing the relationship between Re and I^* holds true in a previously underexplored area of the behavioural phase diagram. Furthermore, the IPES system showed that the $Re-I^*$ relationship generalises for the more complex cross, square and hexagon shapes, but that the transitional boundaries shift.

This chapter is structured as follows. In 6.1 I describe the IPES hardware and software platform. In 6.2 I present the results and discuss various aspects of the relationship between morphological parameters and behaviours, and compare the results to know results from a previous study. In 6.3 I discuss and conclude the chapter.

6.1 Methods

This study presents an automated approach to investigating the falling-paper problem. To achieve near-continuous operation without human intervention, the experimental setup was designed to be robust and autonomous.

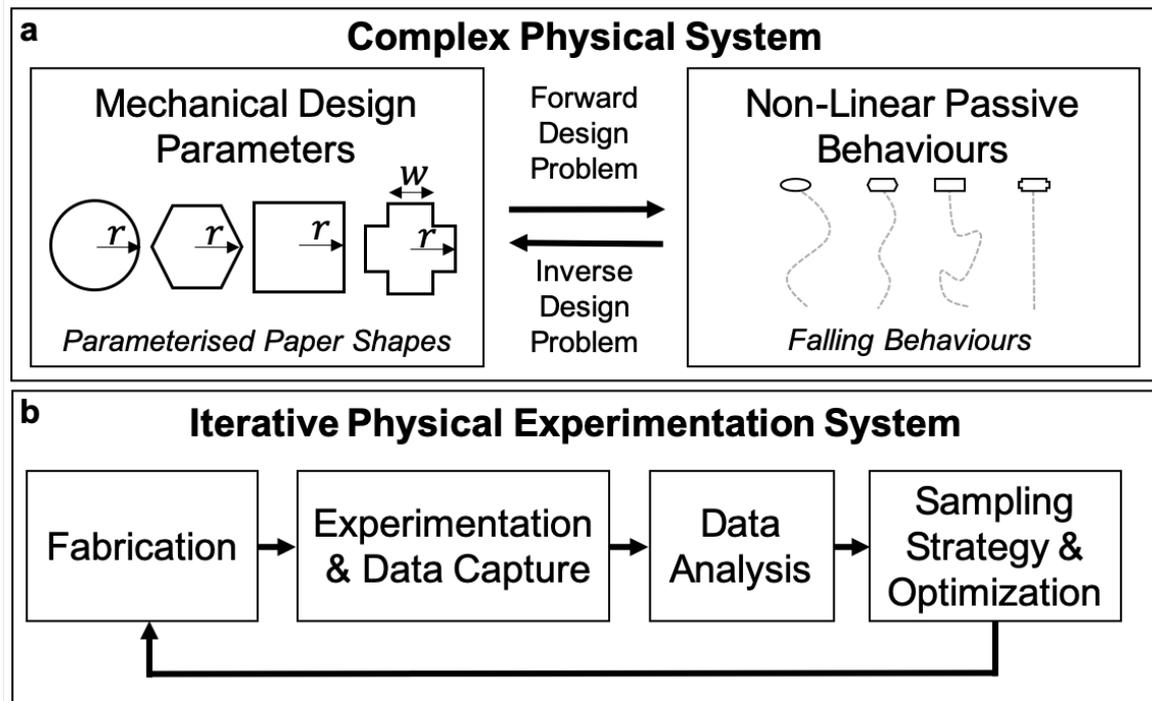


Figure 6.1: Schematic of iterative physical experimentation approach. (a) Overview of complex physical systems such as the falling-paper problem. Mechanical design parameters map to behaviours via the forward design problem. Behaviours map to design parameters via the inverse design problem. (b) Flow chart of the iterative physical experimentation system, where paper shapes are fabricated, tested, analysed and sampled automatically.

6.1.1 Iterative physical experimentation system

Investigating the falling-paper system is challenging due to the complex dynamics, their sensitivity to initial and environmental conditions and the stochasticity in behaviours. The IPES enables us to quickly gather large volumes of data and automatically analyse it to reveal patterns in the underlying dynamics. Figure 6.2 shows the system, which can also be seen in the supplementary video².

6.1.2 Fabrication

The fabrication system was designed to reliably and autonomously cut paper shapes for testing. Shapes were defined by a set of coordinates representing their vertices in two dimensions. Smooth shapes such as circles were discretised along their perimeter; a discretisation of 40 discrete points proved effective.

²<https://doi.org/10.1038/s42256-019-0135-z>

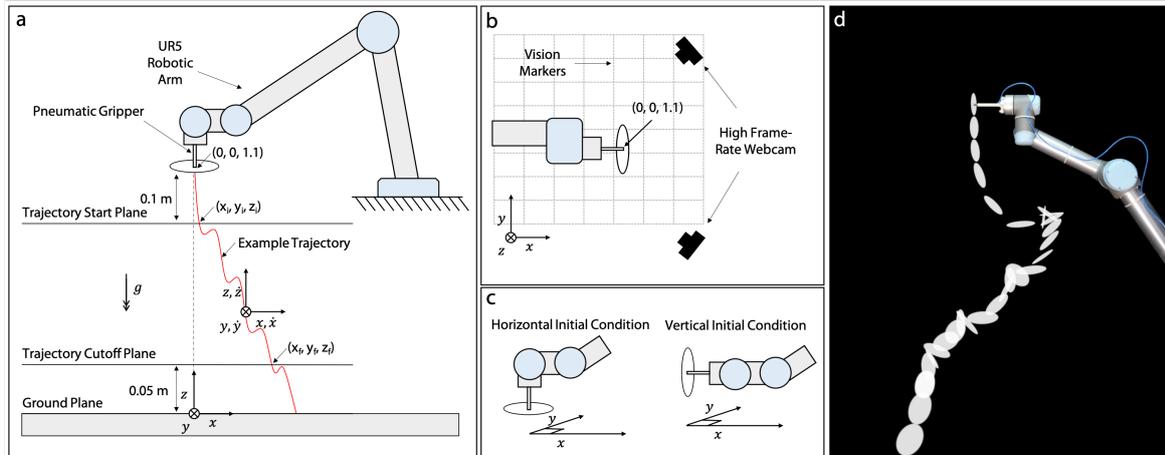


Figure 6.2: Diagram of the experimental setup. (a) Overview of the experimental system showing a UR5 robot arm dropping a paper shape, and the co-ordinate system (x, y, z) used to record the trajectory. (b) Plan view showing the camera setup used to capture 3D falling trajectories. (c) Initial conditions for dropping shapes, chosen randomly in each experiment as a binary choice between the horizontal or vertical case. (d) Time-lapse image of the UR5 robotic arm dropping a paper circle with a vertical initial condition. The shape is exhibiting a chaotic behaviour. The shape outlines have been processed to maximise image quality.

A MakeBlock XY laser engraving and cutting machine was used to cut shapes out of paper. Using laser-cutting technology is advantageous as shapes can be accurately cut to a precision less than one millimetre. Furthermore, there is no degradation in cutting quality over time. The laser cutter uses G-code, a numerical control programming language, to control the laser cutting-head. To allow picking of the fabricated shapes, each shape must be cut at a consistent point within the laser cutter coordinate system, centred at the cut-point $c = (0.1\text{m } 0.1\text{m})$ into the bed of the laser cutter. Each shape was cut using a laser speed rate of 200mm/min; a slower than required speed that minimises cutting failures.

Shapes were cut onto 5-Star branded 70gsm listing paper, which comes as a stream of 2000 continuous sheets. The advantage of using listing paper is that it can be continually fed to the laser cutter and does not have preformed curvature as found in rolled paper. The paper-feeding mechanism was developed such that a continuous stream of listing paper passed under the laser cutter and then through a set of motor-controlled rollers. Once a shape was cut and removed by the pneumatic end-effector, the rollers advanced fresh paper for the next shape. An array of photoresistors was placed under the paper feed path. By detecting changes in resistance as the hole left by the cut shape passed over them, the system could guarantee the removal of used paper. This aided robustness of the setup. Figures 6.3, 6.4 and 6.5 show more details on the experimental set-up and methodology.

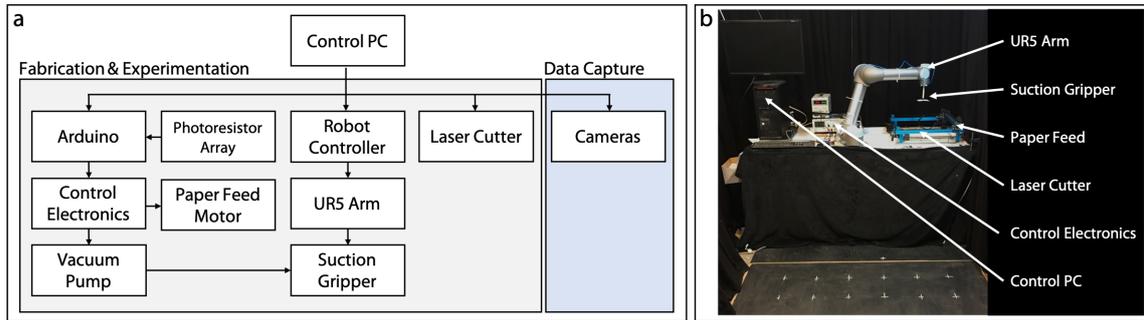


Figure 6.3: Overview of experimental system. (a) Block diagram of the experimental set-up (b) Labelled photograph of the experimental set-up.

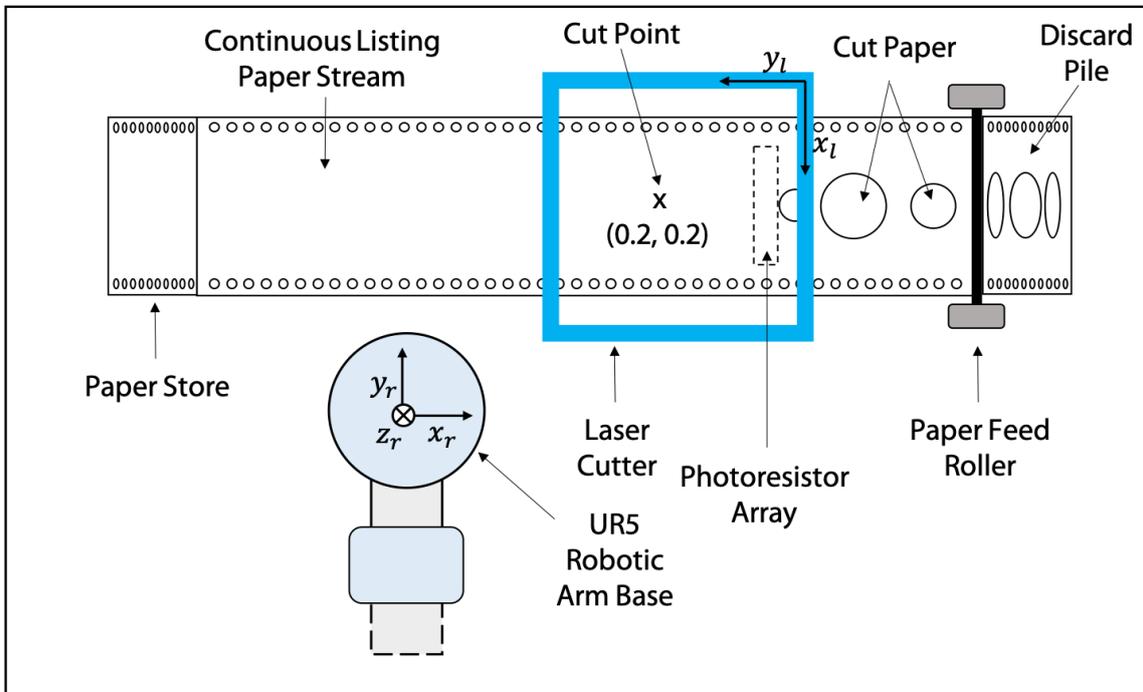


Figure 6.4: Plan view schematic of the laser cutter and paper feed system. The coordinate system of the UR5 robot arm (x_r, y_r, z_r) and laser cutter (x_l, y_l) is shown.

The shape parameter ranges were chosen to explore a large design space without introducing manufacturing errors or very large deformations. Table 6.2 summarises the parameterised paper shapes.

6.1.3 Experimentation and data capture

The experimentation and data capture system were designed to pick up, move and drop fabricated shapes and then capture their falling behaviour. To achieve this a Universal Robots UR5 robotic arm fitted with a custom suction gripper was used. Using a robotic arm allows

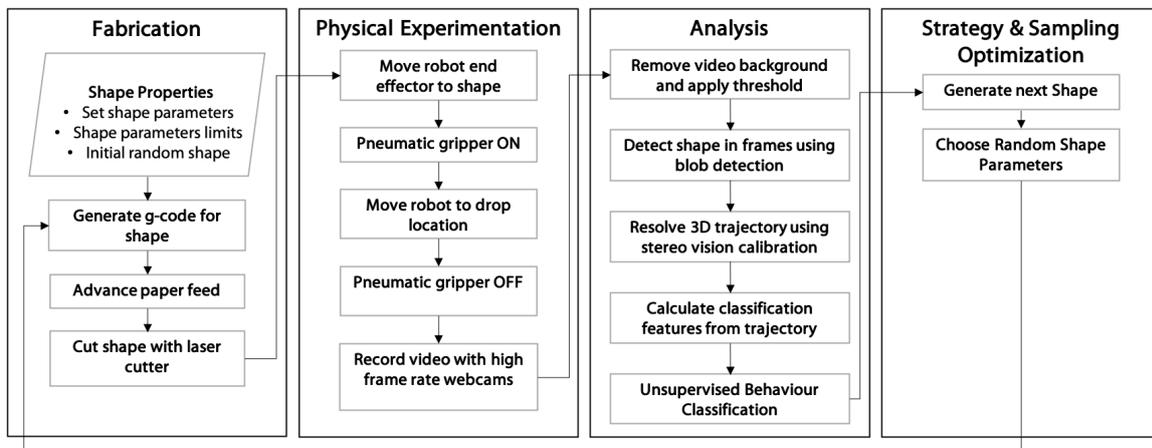


Figure 6.5: Flow chart of system operation, showing the fabrication, experimentation, analysis and sampling strategy steps.

easy automation and precise control over the picking and dropping pose. Additionally, using suction and pneumatic control to pick and drop paper shapes minimises deformation and damage of the paper shapes and allows for rapid experimentation. Figure 6.5 summarises the experimental process.

Once a shape was fabricated, the UR5 arm was programmed to pick and drop the paper using the end-effector, which used a 5 mm rubber suction cup connected to vacuum pump to pick up each shape, and an actuated relief valve to drop each shape. The suction control was operated by a relay system connected to a microcontroller. The pose of the tip of the pneumatic gripper was defined by its position (x, y, z) and rotation (r_x, r_y, r_z) in the robot coordinate system. The picking pose was determined so the gripper was located at the cut point c in the laser coordinate system, with an orientation of $(90^\circ, 0^\circ, 0^\circ)$.

The drop position was 1.1m off the ground and in the centre of the experimental area. When selecting the drop height, a number of factors were considered. A higher drop height is desirable, as it allows longer behavioural observation. However, a higher drop height can lead to more lateral movement. The problem with lateral movement is that shapes can travel out of the camera view, hit the surrounding curtains or land on the experimental apparatus e.g., laser cutter, computer. We found such failures are difficult to detect automatically and required a significant amount of manual intervention to flag and correct. In doing so, behaviours with a high lateral travel distance would be under-represented. 1.1m was therefore chosen as the highest drop height that still prevented excess lateral movement.

Two orientations were used at the drop position, horizontal and vertical, defined by the orientation vectors $(90^\circ, 0^\circ, 0^\circ)$ and $(0^\circ, 0^\circ, 0^\circ)$ respectively. During each experiment, the drop orientation was randomly chosen using a random number generator with two possible

values, each corresponding to an initial condition. Varying the drop angle is necessary as this allows different behaviours to be observed, by randomly choosing the two ‘extreme’ options in angle maximises the range of observed behaviours.

Falling behaviours were recorded using two Logitech BRIO cameras recording at 98 fps. A high frame rate was required to capture the rapidly changing behaviours of falling paper. Each camera was positioned with a side-on view of the experimental platform, on either side of the drop position such that the falling shape was always visible in both camera views. Recording was triggered one second before the paper was dropped and continued for ten seconds. This ensured the whole falling period was captured. To maximise the shape visibility, walls and ground of the experimental area were black; this maximises the contrast with the white paper shape.

The experimental area was situated in a large laboratory and was surrounded by a thick ceiling to floor curtain, designed such that there were no gaps at the floor or ceiling. This was sufficient to create still-air conditions within the curtain, for a range of lab conditions. Experimentation was paused when laboratory conditions could affect this still-air, e.g., during construction when the curtain was being moved.

6.1.4 Data analysis

The data analysis system extracted the three-dimensional trajectory of the falling shape, as well as other physical measurements from the captured video data. Furthermore, it automatically classifies the different falling behaviours in each experiment using a combination of predefined rules and unsupervised clustering.

First, the location of the paper shape in each video frame was determined. Each video frame is compared to a reference background image, taken just before the shape is dropped. Image pixels that differ between images by more than a certain threshold are identified and use to create a binary mask of possible moving objects in the image. The black colour of the ground and walls maximise the contrasts of the white paper shape between the background image and video frame. The MATLAB vision blob analyser is used to detect the largest single connected region in the binary mask, and the centre of mass of this is used to estimate the shape position in the frame. Providing there are no significant changes between frames, this method is effective at locating the shape.

The two webcams were used to create a calibrated stereo-vision pair. Given the image position and timestamps of the shape in each video, the three-dimensional trajectory could be extracted. The first 0.1m of trajectory was discarded (Figure 6.2a), as shapes were sometimes partially occluded by the robot end-effector here. The last 5% of the trajectories was discarded as the trajectory is affected by the ground. Additionally, by identifying the size

of the blob, the observable area of the paper shape visible to the cameras can be determined. The rotational behaviour of the falling shape can be inferred from this parameter. Using the 3D trajectory information and observable area, features which describe the falling can be determined including falling velocity, rotation speed and falling location. A range of parameters were calculated based on the 3D trajectory and observable area measurement. These are shown in Table 6.1. The extracted parameters facilitated automated behavioural classification.

Parameter	Units	Description
Δx	m	Horizontal x displacement
Δy	m	Horizontal y displacement
Δz	m	Vertical z displacement
d	m	Path length
t	s	Falling time
ω	rev/s	Oscillation frequency

Table 6.1: List and description of features extracted from the falling trajectory and observable area.

Steady and periodic behaviours were identified first by using a simple algorithm to segment trajectories into steady or non-steady. These behaviours were defined as those which had a horizontal initial condition and whose position did not deviate horizontally by more than a threshold value (20% of the radius) from the drop position, which was empirically determined. Only horizontal initial conditions give rise to this condition, so only these experiments needed to be considered. As previously documented (Andersen et al., 2005), there are often different behaviours exhibited by the paper during different phases of the falling trajectory. Mostly commonly there is some initial ‘transitory’ phase after initially dropping the paper after which a steady-state phase is found. Hence, all qualifying trajectories were segmented based on the aforementioned identification. In comparison to other work this allows new phases of behaviour to be captured and analysed and allows for a greater understanding of falling behaviours to be obtained.

Tumbling and chaotic behaviours were classified using K-Means clustering (Hartigan and Wong, 1979), an unsupervised machine learning algorithm. The K-Means algorithm is not designed to classify trajectories directly, but can be used to classify trajectory features. Under this approach, a set of scalar features can be extracted from each trajectory and inputted into the K-Means algorithm. Two features were used to classify behaviours. First, the mean rotation frequency, ω , which was extracted from the observable area of each experiment. Peak detection was applied to the observable area time series, with fluctuation in

area corresponding to shape rotation. The assumption is made that each peak corresponds to one full rotation of the paper shape. The frequency ω is then calculated as the number of rotations divided by the total falling time. ω was chosen as it tends to be higher in tumbling behaviours. The second feature used was the standard deviation in vertical velocity $\text{std}(v_z)$, which tends to be higher for chaotic behaviours as the motion is far more varied. These features were chosen empirically by human observation of the different behavioural patterns. Despite this, automating the classification using manually crafted features offers a less subjective and more repeatable classification approach than previous visual approaches.

For each shape, the features for each experiment were stored in the matrix $\mathbf{X} = [\omega_k \text{std}(v_{zk})]$ for $k = 1, \dots, K$, where K was the total number of experiments for that shape. The K-Means clustering algorithm was applied to X ,

$$\mathbf{b} = \text{KMC}_N(\mathbf{X}) \quad (6.1)$$

where N is the number of clusters, two in this case, reflecting the chaotic and tumbling behaviours, and \mathbf{b} is an array of cluster assignments. As is standard practice, the algorithm is run multiple times, three in this case, to avoid clustering anomalies. Labelling behaviours is challenging, even as a human observer. Hence, using automated clustering allows for less subjective classification.

The classification results were compared with those obtained from two human observers. The human observers had access to the same visual data, as well as the extracted 3D trajectories. For each experiment, the observers discussed the features of the experimental falling style and categorised it into one of the three possible behaviours. For experiments where the automated system identified behavioural transitions, the observers classified each behavioural segment independently. To calculate the agreement with the human observers, we uniquely reassigned cluster assignments such that the fraction of misclassified behaviours is minimised.

Previous work has identified that dimensionless quantities can be used to characterise falling disk in fluid behaviours, namely the dimensionless moment of inertia I^* and Reynolds Number Re . Using the shape parameters and the features determined for each experiment these can be calculated for each experiment.

The Reynolds Number is given by

$$Re = \frac{v_z L}{\nu}, \quad (6.2)$$

where v_z is the vertical falling velocity, L is a length scale and ν is the kinematic viscosity of the liquid, in this case air. The diameter parameter $d = 2r$ was used as the length scale. The

non-dimensional moment of Inertia is

$$I_{\text{disk}}^* = \frac{I_{\text{disk}}}{\rho_f d^5} \tag{6.3}$$

where I_{disk} is the mass moment of inertia, ρ_f is the fluid density. $\rho_f d^5$ is an inertial term similar to the mass moment of inertia of an equivalent volume of fluid. For the other shapes we have

$$I_{\text{Hexagon}}^* = \frac{5mr^2}{12\rho_f d^5}, \tag{6.4}$$

$$I_{\text{Square}}^* = \frac{mr^2}{3\rho_f d^5}, \tag{6.5}$$

$$I_{\text{Cross}}^* = \frac{m(4r^2 + w^2)}{24\rho_f d^5}. \tag{6.6}$$

6.1.5 Searching and exploration

In this study, morphological parameters were chosen randomly. Due to the number of experiments undertaken this allowed for exploration of the full parameter-space. Figure 6.6 shows the distribution of sampled design parameters for each shape.

Of course, there is scope for implementing more intelligent and efficient search strategies. These could include genetic algorithms (Goldberg and Holland, 1988) and their variants MAP elites (Mouret and Clune, 2015), as we discussed in Chapter 2, or even Bayesian optimisation as discussed in the Chapter 5. Other machine learning approaches, for example Active Learning (Settles, 2009), could be beneficial. Active learning is an approach that allows algorithms to choose the training data from which they learn, which in our case could mean guiding the parameter search for more focused exploration on behavioural boundaries. Similarly, tools from research into optimal and Bayesian experimental design Chaloner and Verdinelli (1995) offer a structured approach to designing experimental procedures.

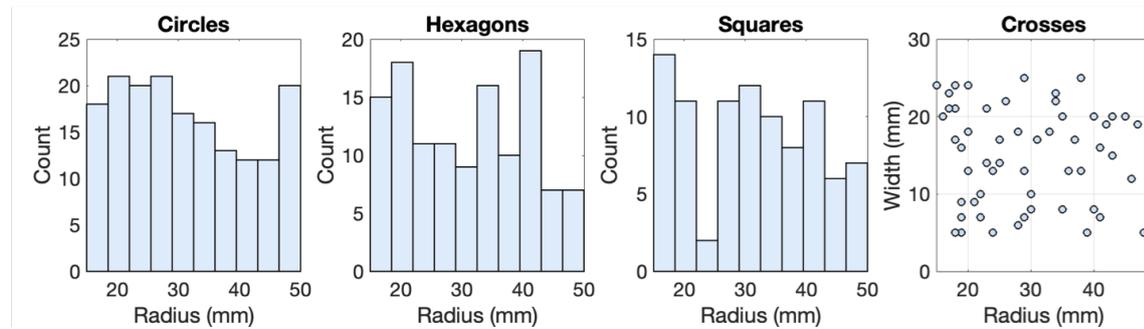


Figure 6.6: Distribution of randomly sampled design parameters for each shape.

6.2 Results

Four parameterised shapes were investigated; circles, to validate the set-up against previous studies (Field et al., 1997), and squares, hexagons and crosses to assess the effect of morphology on falling behaviours for more complex shapes. Shapes were fabricated and dropped from a height of 1.1m and with a randomly chosen initial condition (Figure 6.2c). Experiments took around 90 seconds to complete, with shapes taking 1 to 5 seconds to fall to the ground. Shapes were parameterised by their radius, r and width, w , as shown in Figure 6.1a and Figure 6.7. Table 6.3 shows the total number of experiments conducted for each shape. In the exploratory work for this work, more than 1000 experiments were conducted to validate the system and explore its limitations. Each shape was parameterised so a wide range of morphologies could be observed. The lower and upper limits were set by a range of factors. Smaller shapes were difficult to pick up and track, while larger shapes exhibit very large deformations and took a long time to manufacture. The parameter ranges (Table 6.2) reflects this.

Shape	Parameter	Range (mm)
Circle	r	$25 \leq r \leq 50$
Hexagon	r	$25 \leq r \leq 50$
Square	r	$25 \leq r \leq 50$
Cross	r	$25 \leq r \leq 50$
	w	$5 \leq r \leq 25$

Table 6.2: List of shape parameters and their corresponding ranges.

Shape	Number of Experiments
Circle	170
Hexagon	123
Square	92
Cross	59
Exploratory work (circle, square, rhombus, parallelogram)	1200

Table 6.3: Number of experiments carried out for each shape.

We used an adapted version of the behavioural groupings presented in for falling disks, namely:

1. Steady and periodic falling – the disk falls steadily or oscillates back and forth with a horizontal orientation.

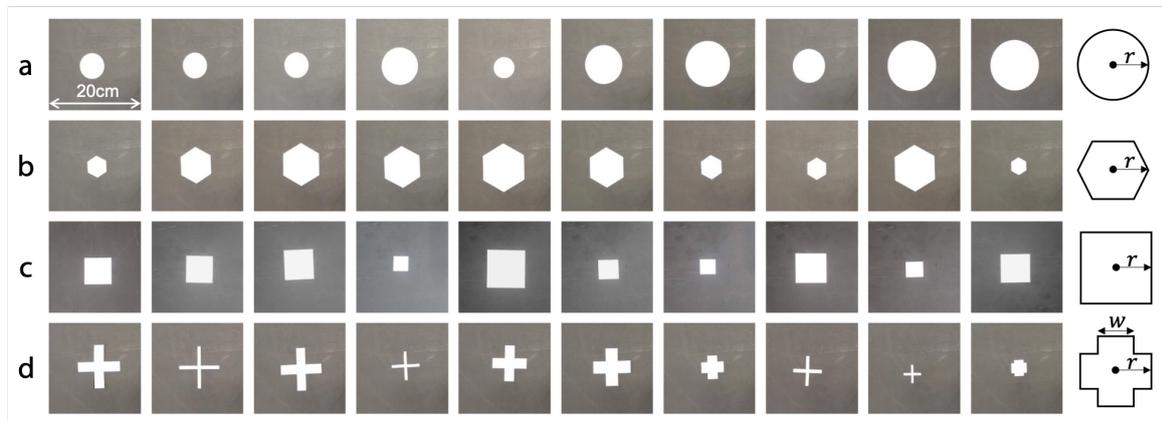


Figure 6.7: A representative sample of manufactured shapes, with design parameters r and w shown. (a) Circles (b) Hexagons (c) Squares (d) Crosses.

2. Tumbling – the disk continuously turns end over end.
3. Chaotic motion – the disk switches between tumbling and swooping motions with no apparent structure.

In other studies, steady and periodic behaviours are considered two separate groups. However, periodic behaviours are typically observed with heavy disks falling in a liquid. In falling-paper systems the fast dynamics make them difficult to differentiate. The main characteristic of both behaviours is that the shape falls vertically down without tumbling or flipping. Figure 6.9 shows representative examples of these behaviours for each shape, which can also be seen in the second supplementary video³.

Behaviour-characterising features such as vertical velocity, horizontal velocity and oscillation frequency were extracted from the 3D trajectory. The automated behavioural classification system utilised these parameters to segment and classify trajectories into one of the three possible behaviours, with an average agreement of 86.2% when compared with human classification. Figure 6.8 shows confusion matrices for each shape and behaviour. These show the percentage of agreement between the automated system and human observer in each case. Generally, the confusion (represented by off-diagonal values) increases with shape complexity. Table 6.4 summarises the agreement between the automated system and human observers for each shape. We see that circles and hexagons show most agreement. We suspect this is because the original behavioural groupings were based on falling disks, and hence the circular shapes fitted these groupings better. The same classification scheme was used for all four shapes. However, the classification agreement for the cross shapes

³<https://doi.org/10.1038/s42256-019-0135-z>

was significantly lower than for the other shapes, suggesting that alternative parameters may better describe the behaviours of cross shapes.

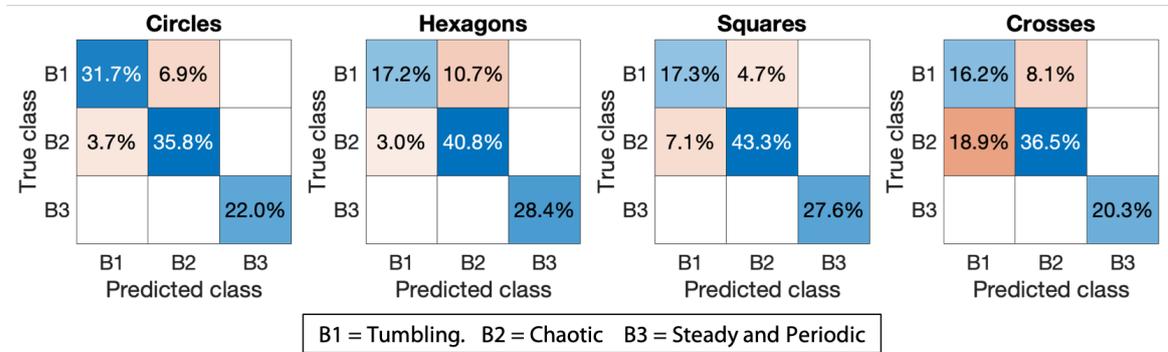


Figure 6.8: Confusion matrices showing the agreement with human observers when classifying behaviours.

Shape	Classification Agreement
Circle	89.5%
Hexagon	86.4%
Square	88.2%
Cross	72.3%
Average	86.2%

Table 6.4: Classification agreement with human observers for each shape.

In some cases, two behaviours were observed in one falling trial, with an initial phase of one behaviour followed by a steady state phase with a secondary behaviour. When this occurred, the trajectories were split into the two different sections and analysed and classified independently. This can be observed for the behaviour of the circular shape (Figure 6.9a) where there is transition from steady falling to chaotic falling.

Figure 6.10 shows representative trajectories for each shape and behaviour. Chaotic motions were characterised by abrupt changes in vertical velocity in the 3D trajectory and apparently random rotation. Tumbling motions were characterised by smooth trajectories and periodic rotation. Steady behaviours were characterised by vertical trajectories with no clear periodic oscillation. In some cases, the steady behaviour was observed only the initial section of the trajectory after which the behaviour transitioned, as can be seen in Figure 6.10a,b.

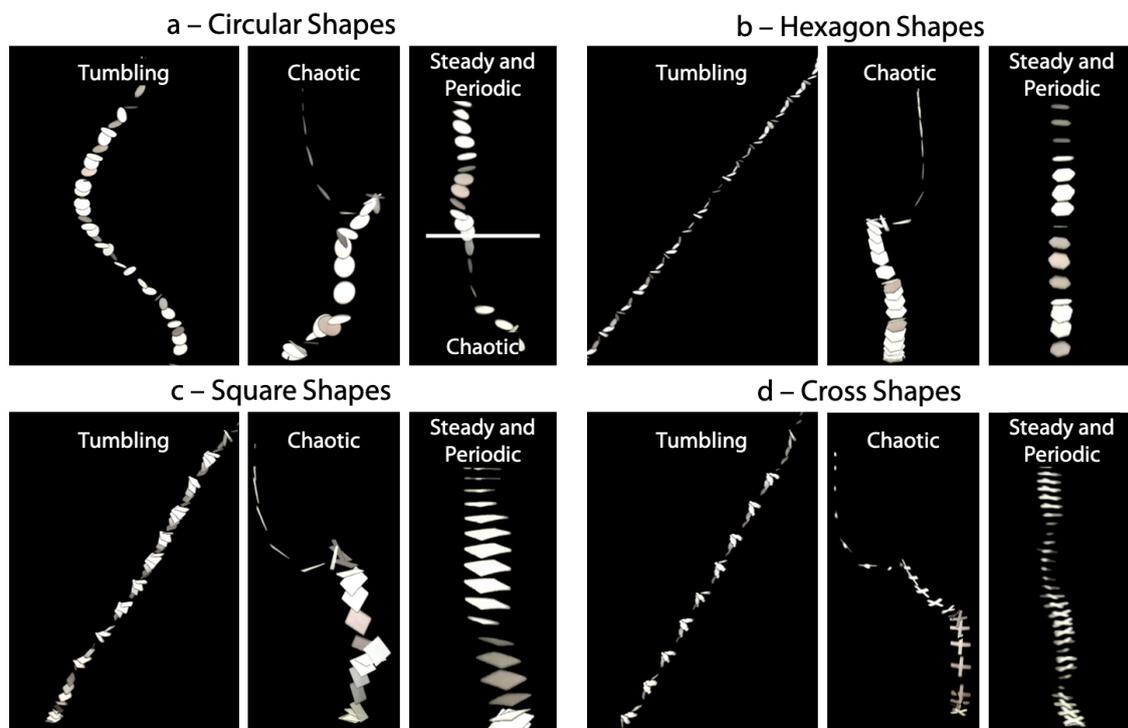


Figure 6.9: Time-lapse images of the three falling behaviours observed in each shape. (a) Falling circular shapes, showing a behavioural transition from steady and periodic to chaotic (b) Falling hexagonal shapes (c) Falling square shapes (d) Falling cross shapes.

6.2.1 Effect of morphological parameters on observable characteristics

The morphological parameters have a significant effect on the characteristics of falling behaviours. Furthermore, a specific morphology can exhibit multiple behaviours depending on initial conditions and environmental interactions. This is shown in Figure 6.11a, where the vertical falling speed v_z (the speed in the z -direction) is plotted against the morphological parameters r and w . Falling speed is an interesting feature to examine, as understanding its relationship with morphology and behaviour is relevant to designing and optimising complex flying structures.

The circle, hexagon and square are all parameterised by one morphological parameter, r . Tumbling motion generally occurs when $r < 0.03\text{m}$, after which chaotic behaviour occurs. Steady and periodic falling behaviours are observed across the full range of r . The cross exhibits a weaker pattern relative to r , however tumbling behaviours tend to occur at lower values. Tumbling patterns generally occur when $w > 0.015\text{m}$, and steady patterns occur across the whole range. This suggests that by introducing a secondary parameter it is possible

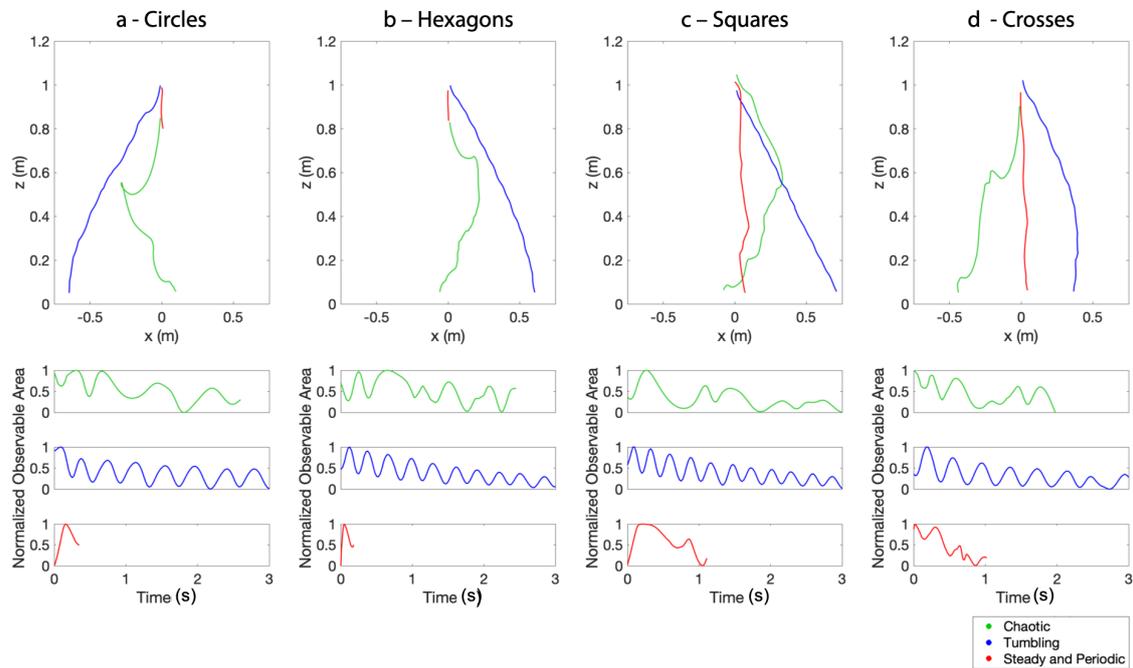


Figure 6.10: Trajectories and observable area profiles of falling shapes for different automatically classified behavioural groups. (a) Top panel shows typical trajectories for the three different behaviours observed for falling circles. The steady trajectory is seen for only the initial period of falling. The bottom figure shows the observable area as the shape falls for the different behaviours, highlighting the rotation of the shape. (b) Trajectory and area plots for hexagonal shapes of different behaviours. (c) Trajectory and area plots for square shapes of different behaviours. (d) Trajectory and area plots for cross shapes.

to modulate falling behaviours. For all shapes, steady falling behaviours generally fall slower than chaotic or tumbling behaviours. However, there is no clear relationship between v_z and tumbling or chaotic behaviours, with small tumbling shapes showing comparable speeds to large chaotic shapes.

The relationship between v_z and path speed v (rate of travel along the path of trajectory) was also examined, as shown in Figure 6.11b. The closer the values of v_z and v are to each other, the more the shape follows a straight vertical trajectory. Hence, we see that steady falling behaviours tend to lie close to the line $v = v_z$. Conversely, chaotic behaviours are spread out, showing that they travel further along their trajectory for the same vertical height loss. Tumbling behaviours tend to lie closely together to the left of the line $v = v_z$, indicating most tumbling behaviours move consistently along their trajectories at constant velocities.

Figure 6.11c shows the horizontal displacement Δx and Δy during the trajectory. Steady and periodic behaviours by definition show minimal horizontal displacement. However, tumbling and chaotic behaviours show greater deviation. In general, chaotic behaviours

appear to have greater horizontal movement from the starting position; however, there is significant variation within the behaviours. The tumbling circle shapes appear to favour travelling in the negative x direction; however, more experiments would be required to fully explore this observed pattern. We speculate this is due to a slight systematic error in the release orientation.

Figure 6.12 shows extended analysis of falling behaviours, specifically the effect of the initial drop condition on distance travelled. Shapes dropped with a vertical initial condition tend to travel further than those dropped with a horizontal initial condition. Even though most shapes eventually transitioned to a chaotic or tumbling behaviour, those dropped horizontally spend more time falling steadily down. Hence, the time available for horizontal displacement was reduced. If we plot the radial distance travelled vs. the design parameter r and initial condition, we see that initial condition dictates displacement more than design.

Although each of the features shown in Figure 6.11 show some ability to separate or uniquely define behaviours, there is no single output which allows full classification.

6.2.2 Behavioural grouping

To understand how morphological and environmental parameters can influence falling behaviours, non-dimensional parameters have been used to construct behavioural ‘phase-diagrams’ (Field et al., 1997). Most commonly among these non-dimensional parameters are the Reynolds number (Re) and dimensionless moment of inertia (I^*), which characterise the effect of the fluid and disk morphology respectively.

When plotted against these quantities, a clear structure emerges that separates the three falling styles of disks. To demonstrate the validity of the IPES in wider context of falling-paper research the results of (Field et al., 1997) were replicated. Re and I^* were calculated for each experiment and overlaid onto the well-known phase diagram of falling disk styles (reproduced manually by the authors), as shown in Figure 6.13a–b.

The automatically classified tumbling, chaotic and steady behaviours matched the original phase diagram with an accuracy of 96%, 71% and 6% respectively. Tumbling behaviours are easiest to classify, as their characteristics are distinct. It is often ambiguous whether a behaviour is chaotic or tumbling. Steady falling has previously been observed at the lower end of the Re and I^* range (Stringham et al., 1969; Willmarth et al., 1964). The relative effects of fluid forces and paper inertia in air in this study meant steady falling in the circle was only observed for short periods of time, before transitioning to another behaviour. The apparent transient existence of steady falling in this area of the phase diagram has not been reported before.

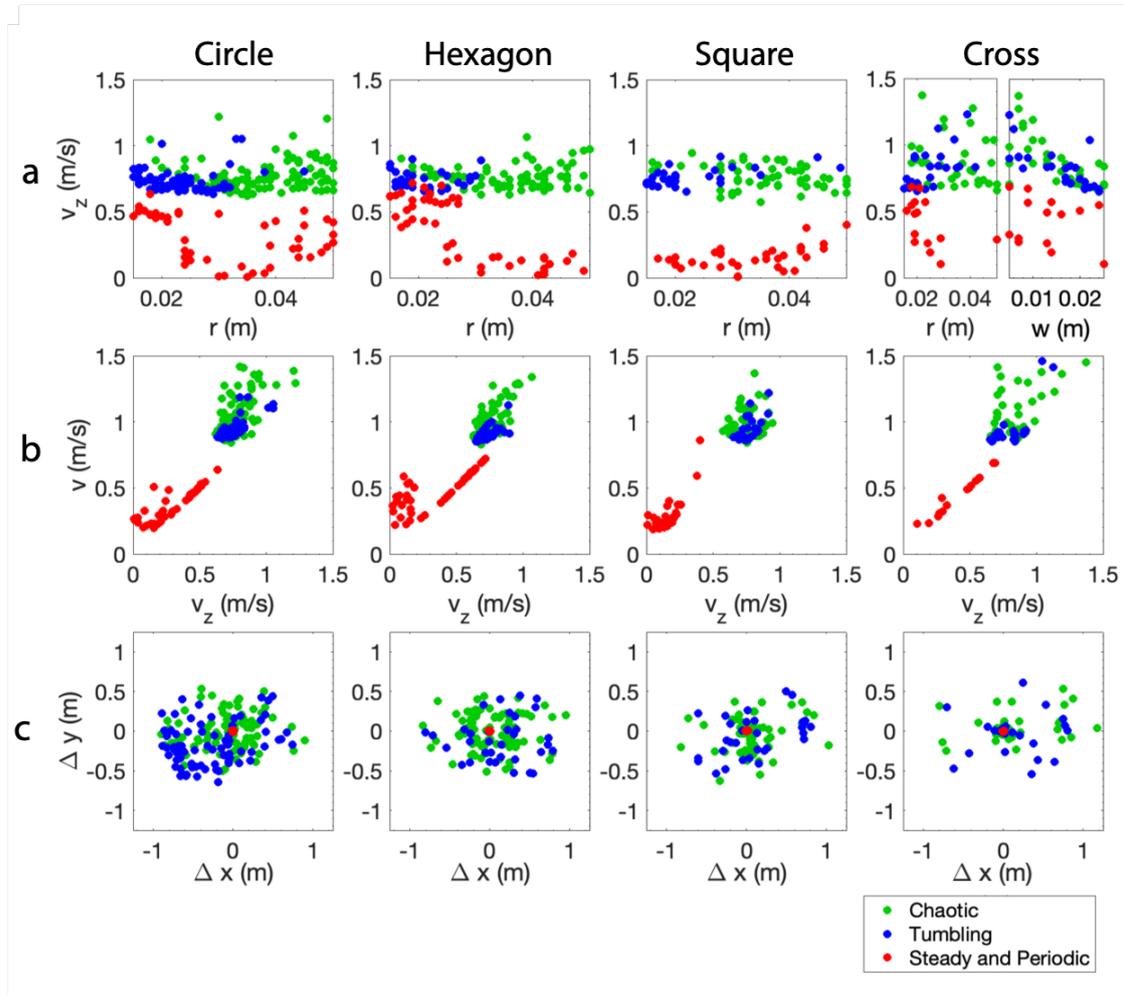


Figure 6.11: Variation in measured falling characteristics with respect to design parameters and automatically classified behaviour groups. (a) Vertical falling speed v_z vs. mechanical design parameters r and w (b) Path velocity v vs. vertical velocity v_z (c) Horizontal displacement Δx and Δy between start and finish of trajectory.

Prior studies have used different materials and fluids, for example metal disks and high viscosity fluids, to extend the range of I^* and Re . Meanwhile, this experimental setup can only cover a limited range of I^* and Re . Previously, this range has been largely underexplored. With the IPES it is possible to observe behavioural boundaries at a much higher resolution.

The $Re-I^*$ phase diagram provides a compact view of behavioural diversity. The applicability of this phase diagram for the hexagonal, square and cross shapes was tested. Figure 6.13c–d shows the automatically classified behaviours plotted against Re and I^* , which were reformulated for each shape. All shapes show a similar overall structure, with high- I^* shapes more likely to exhibit tumbling behaviour, and steady falling shapes exhibiting lower Re . The behavioural boundary line shifts between each shape. The variation relative to the

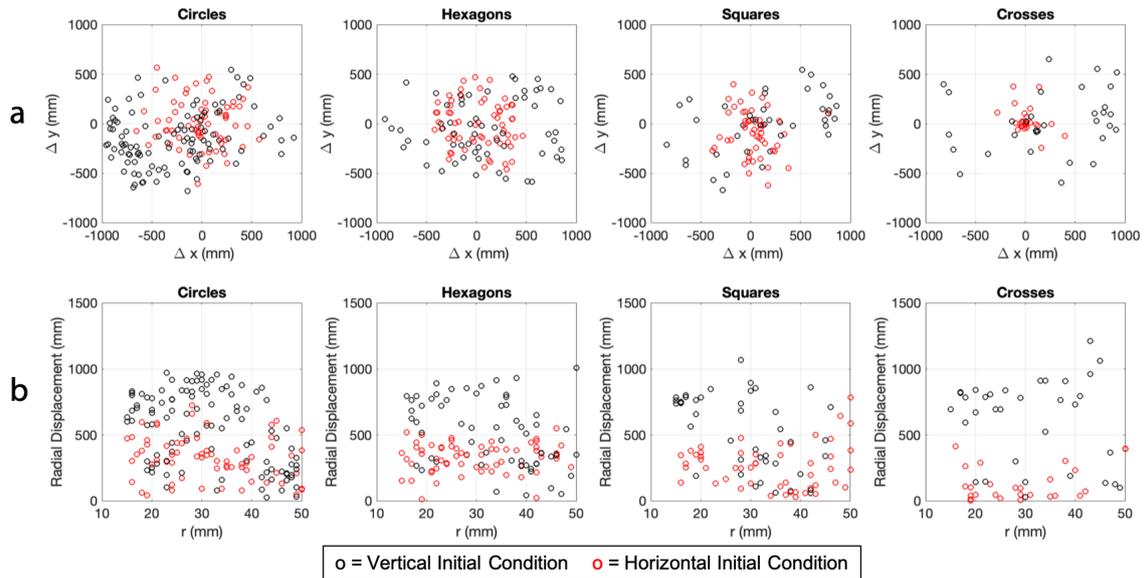


Figure 6.12: Effect of initial condition on horizontal displacement. (a) Horizontal displacement Δx and Δy for each shape and initial condition. We see that shapes with a vertical initial condition tend to travel further. (b) Radial horizontal displacement ($\sqrt{\Delta x^2 + \Delta y^2}$) for the same shapes and initial conditions, again showing how shapes with a vertical initial condition tend to travel further.

circle results increases as the shapes deviate from circular. Hence, the hexagon behaves most similarly, and the cross is least similar. The generalisation of the $Re-I^*$ phase space to the hexagon and cross is a previously unreported finding. There may be more appropriate dimensionless quantities to characterise the behaviour of the falling cross.

6.3 Discussion and conclusion

In this chapter I presented an approach for iterative physical experimentation to facilitate the scientific investigation of the relationship between mechanical design parameters and non-linear passive dynamics. I demonstrated this approach by autonomously investigating the falling-paper problem. Understanding this problem requires large quantities of experimental data coupled with intelligent data analysis.

Using computer vision, it was possible to extract the 3D trajectory and corresponding oscillation. This data was processed and used to automatically segment and classify behaviours, then examine the relationship between morphology, behaviour and output variables such as falling speed (Figure 6.11). The analysis aids and advances understanding of each behaviour, and the corresponding behavioural boundaries. In addition, it provides practical insight into designing structures to exhibit certain behaviours. This could, for example, assist

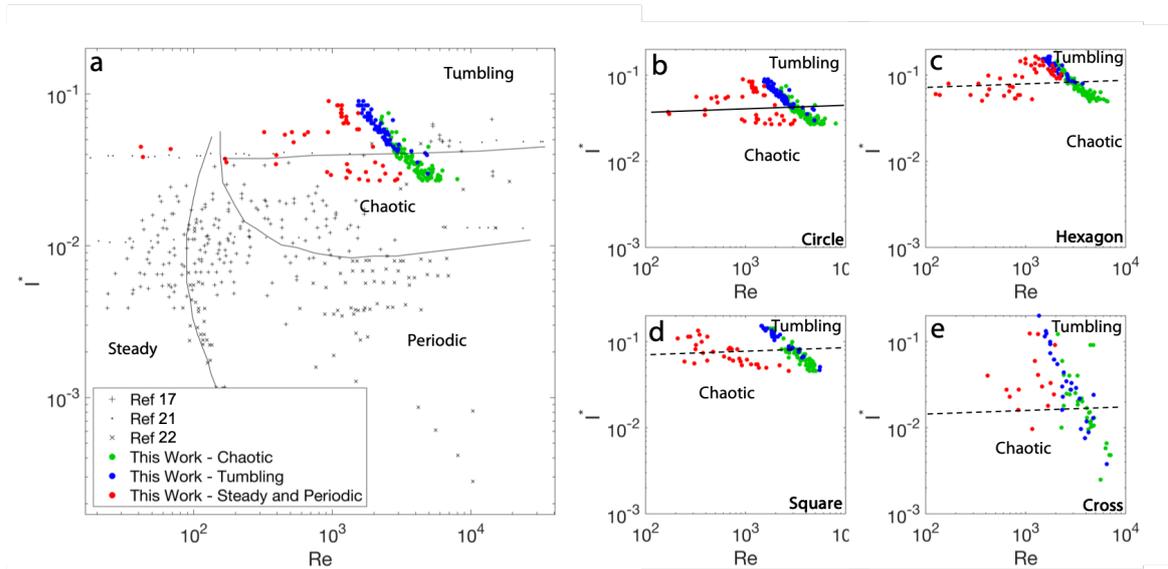


Figure 6.13: Automatically classified falling behaviours in the Reynolds number Re and non-dimensional moment of inertia I^* parameter space. (a) Results for circular disks from this study plotted against results obtained from previous studies, showing the behavioural groups and transition lines. The transition lines indicate the approximate boundaries between behaviours and are added by hand. (b) Close up of the circular disk results from this study (c) Similar plot for Hexagon shapes, the dotted line indicating an approximated behavioural boundary added by hand by the authors (d) Similar plot for square shapes (e) Similar plot for cross shapes.

with optimising the design of flying robots with respect to certain parameters such as stability or falling speed.

Previous studies (Field et al., 1997) have used the parameters Re and I^* to group the behaviours of falling disks, and were used here to validate the IPES. The results showed that the IPES correctly classified tumbling and chaotic behaviours and grouped them in the correct area of the behaviour phase diagram. However, unlike previous studies the system identified steady-falling behaviours in previously unseen areas of the phase diagram. These behaviours manifest for a short period of time in the transient stage of motion, something not analysed in the context of the phase diagram. This suggests the boundaries in the $Re-I^*$ diagram are not as fixed as once thought, and that ‘out-of-group’ behaviours can exist. Investigating previously unexplored hexagons and crosses demonstrated that the $Re-I^*$ phase diagram generalises to other shapes, although the behavioural transitional lines are shifted. However, the behavioural distinctions became less distinct the less circular the shape was, indicating different non-dimensional parameters may be useful.

A key aspect of the falling-paper problem has been the classification of behaviours, which has been previously performed exclusively visually. This is, especially for fast-moving behaviours, somewhat subjective. The human classifiers used in this study were frequently unable to agree on what behaviour a certain shape was exhibiting. While no specific data is available, this occurred in approximately 25% of experiments. Automating this process using unsupervised clustering presents a more repeatable, less subjective classification approach. This is especially true if behaviours are clustered based on well-thought-out parameters, in this case the variation in vertical velocity and frequency of oscillation. The results here showed that one unsupervised classification scheme could differentiate between chaotic and tumbling motion accurately in all four shapes.

The average agreement over all experiments of the human and automated behavioural classification system was 86.2%. As mentioned, in some cases the human observers found visual classification challenging. Hence, we do not refer to this percentage agreement as an accuracy, as this implies the human classified groups are the ground truth. Moving toward semi or fully unsupervised behavioural classification produces behavioural groups based on physically quantifiable differences, rather than visually informed groupings. In doing so, a ground truth can be established based on the natural groupings of behaviours of the parameter space.

As well as the positive results I have shown, the system has some limitations. As discussed, behavioural classification is initially reliant on human supervision to craft features. One of the key potential benefits of an automated system is that it allows us to create system interpretations not based directly on human input. Hence, the system will be fundamentally limited until classification can be decoupled from human input. There are also physical constraints that limit the extent of experimentation. Robotic manipulation of delicate objects, e.g. paper, is a research field in itself, and the range of possible shapes and experiments the system can perform is limited by this. One can imagine that, as robotic technology advances, we will be able to apply automatic experimentation to ever more complex and delicate systems.

Part II

Bernoulli-balls

Chapter 7

Height control of a Bernoulli-ball via bang-bang control¹

There is a long-standing interest in the study of mechanical self-stability – henceforth termed *self-stability* – for the purpose of dynamic motion control (Pfeifer et al., 2007). In biology, it is known that many animals take advantage of mechanical dynamics to maintain stability during walking and running locomotion, for example, in highly unstructured and uncertain terrains (Ijspeert, 2014). Inspired by this, there were a large number of case studies on legged robot locomotion investigating the principles of robust motion control in complex tasks (Iida and Ijspeert, 2016). For example, Passive Dynamic Walkers (PDWs) showed that a purely mechanical system can be designed to take advantage of the pendulum-like dynamics of legs, allowing unaided walking down a shallow slope (McGeer et al., 1990). The exploitation of such mechanical dynamics can be extended by injecting energy, allowing improved operation and walking on flat ground as studied theoretically in (Goswami et al., 1997; Spong et al., 2007) and in robot experiments (Collins et al., 2005). Aside from legged locomotion, there are numerous other studies on self-stabilising systems. The ‘Blind Juggler’ (Reist and D’Andrea, 2012), for example, exploits the passive dynamics of a ball bouncing on

¹This chapter is based on the following peer-reviewed publication:

- Howison, T., Giardina, F., & Iida, F. (2020). Augmenting self-Stability: height control of a Bernoulli-ball via bang-bang control. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3974–3980.

Contributions

T. Howison – devised the conceptual idea, developed models constructed hardware platform, carried out experiments, developed the model and wrote the paper.

F. Giardina – contributed equally to model development and stability analysis, aided with experiments and proofreading.

F. Iida – contributed to general discussions and proofreading.

a parabolic surface to juggle multiple balls without any feedback. Self-stabilising properties can even be observed in a bicycle (Kooijman et al., 2011).

Despite its advantages, self-stabilising systems are generally considered challenging when designing active control mechanisms, because the system behaviours are dictated by their intrinsic mechanical dynamics, which we cannot directly change (Burrige et al., 1999). One problem is the non-linear relationship between the system dynamics and control input making control influence hard to predict (Hunt et al., 2016). Another is latency after active control is applied, because self-stabilisation around a new equilibrium point usually requires long transient periods, making the design of active controllers challenging (Guenther and Iida, 2017).

Various methods exist for the active control of self-stabilising systems. An approach called the energy-shaping method was introduced in order to identify energy inputs for self-stabilising systems (Spong, 1999). Here, the nominal mechanical self-stability is used as a reference trajectory, on top of which the energy input is estimated to handle different environmental conditions. This is an elegant approach to find the minimum energy required to stabilised complex systems, but its applicability is limited to cases where precise system models are available. A more general approach, called the linear quadratic regulator (LQR) tree method, was proposed based on the LQR design methodology (Tedrake et al., 2010). By linearising complex dynamics, the LQR-tree method estimates the energy input to guide systems to target locations. While this approach is generally sound, the requirement of high control bandwidth restrict the applicability in many real-world applications, especially those which have fast dynamics beyond sensing and computational capabilities.

In this chapter, I propose a minimalistic control approach—termed *augmented self-stability*—that utilises bang-bang control to improve the time-to-target of self-stabilising systems without loss of stability. Utilising bang-bang control in self-stabilising systems is advantageous for a number of reasons. First, the control is defined with few discrete parameters that give rise to a certain trajectory. These types of control problems are efficiently optimised using numerical solvers. Hence, systems are often controlled globally using bang-bang before switching to a local feedback controller such as LQR, e.g., in the pendulum swing-up problem (Park et al., 2011). In self-stabilising systems, we can ‘outsource’ the design of a stabilising feedback controller to the intrinsic dynamics of the system. Second, self-stabilising systems reliably follow certain trajectories under certain control inputs. Bang-bang control harnesses these trajectories, which by nature are reliable, hence making the control policy robust to disturbances. Third, the control bandwidth requirement is minimal – there is no requirement for high-frequency sensing and computation of control inputs – meaning control policies can be calculated *a priori*. Finally, It is worth noting that for linear control problems

subject to bounded inputs bang-bang control is time-optimal (Jazar, 2010). Unfortunately, time-optimality of bang-bang control in non-linear systems is not equally guaranteed.

As a case study of our proposed control approach, we study the height control problem of a sphere hovering in a vertical airflow: the so-called Bernoulli-ball. Though this is an interesting control challenge, it has not been studied intensively in the past. The only directly related paper in our search (Nudehi et al., 2017) investigates the feedback control of a similar system, however the self-stabilising properties are not considered or analysed. More broadly, there is a rich literature on magnetic (El Hajjaji and Ouladsine, 2001; Mann and Sims, 2009) and acoustic levitation (Trinh, 1985; Ueha et al., 2000; Xie and Wei, 2001).

This chapter is structured as follows. In 7.1 I introduce the model of the Bernoulli-ball is introduced. In 7.2 I present Lyapunov stability analysis to verify self-stability. In 7.3 I describe how the system self-stability can be augmented with a bang-bang control algorithm. In 7.4 I present experimental results of hardware validation. In 7.5 I demonstrate how further behaviours can be induced in the system with an oscillating airflow. Finally, in 7.6 I discuss and conclude the chapter.

7.1 Model and control

Figure 7.1 shows a schematic of the set-up. A high power fan is used to propel air into the expansion chamber and through a flow straightening nozzle. The airflow is changed by using a pulse width modulation (PWM) signal – which can take an value in the range 0–1 – to alter the fan power. A ping-pong ball is used as the hovered object. A more detailed description of the set-up is provided in 7.4.1.

7.1.1 Dynamics modelling

The Bernoulli-ball exhibits two forms of self-stability. Firstly, it is self-stable in the horizontal plane; when perturbed horizontally the ball returns to its original position. This is commonly explained using Bernoulli's principle, which states that as the speed of a fluid increases, the fluid pressure decreases. Hence, the pressure within an airflow is lower than the surrounding environment, creating a self-stabilising force about the jet centerline (Gerhart et al., 2016). Secondly, it is self-stable in the vertical direction; when released into the airflow, the ball will eventually settle around a nominal height. This is due to the force balance between drag and the ball mass, and the dissipative effect of moving in the airflow. To simplify matters, we first assume that motion in the horizontal plane x - y and the vertical direction z is decoupled. We make this assumption because the horizontal self-stability acts as a feedback loop, keeping the ball centred in the flow.

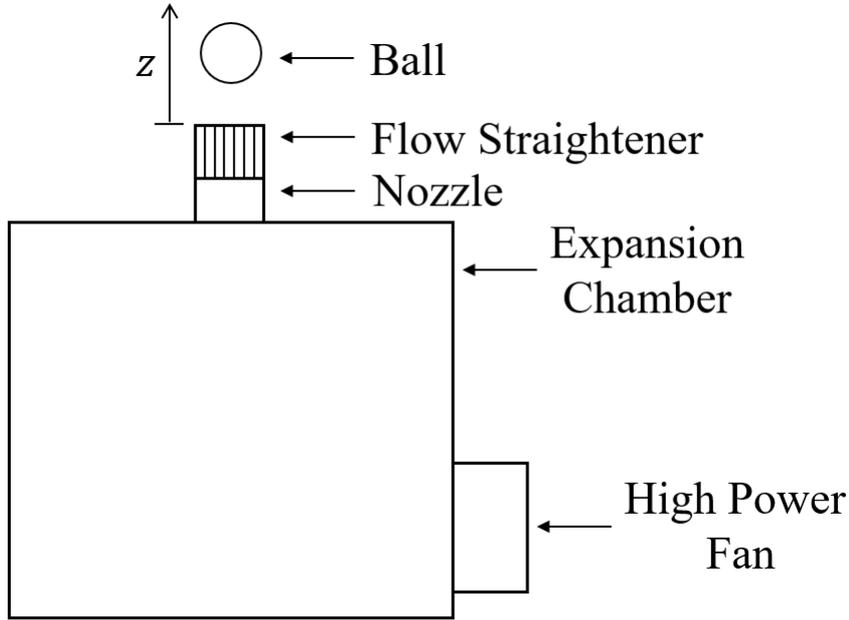


Figure 7.1: Schematic of set-up for hovering a sphere in a vertical airflow. Air is propelled by the fan into the expansion chamber then through the nozzle and flow straightener.

The vertical dynamics are the result of complex interactions between the ball, airflow and environment, meaning the governing equations are highly non-linear. Here, we present a simplified model of the system dynamics. We start by considering the force balance on the ball in the vertical direction,

$$m\ddot{z} = F_d(z, \dot{z}) + F_b - mg \quad (7.1)$$

where z is the vertical distance between the nozzle outlet and ball centre of mass, $F_d(z, \dot{z})$ is the drag force on the ball at a height z and velocity \dot{z} , F_b is the buoyant force due to the weight of displaced air and mg is the ball weight. The buoyant force is minimal, so we have $F_b \approx 0$.

Assume that the air-jet flow expands conically from the nozzle with a cone angle θ and that the velocity profile $V_f(z)$ is approximately parabolic in the radial direction, as shown in Figure 7.2. We define $\bar{V}_f(z)$ as the mean flow velocity at a height z , with the mean velocity at the nozzle outlet being denoted by $V_0 = \bar{V}_f(0)$. Applying flow continuity over z yields an expression for $\bar{V}_f(z)$,

$$\bar{V}_f(z) = \frac{V_0 r_n^2}{(r_n + z \tan \theta)^2} \quad (7.2)$$

where r_n is the nozzle radius.

The mean nozzle outlet velocity changes in response to a change in fan power, which is controlled using a PWM signal denoted by u . We define $f(u)$ as the yet-to-be-determined

mapping between u and V_0 . Therefore,

$$\bar{V}_f(z) = \frac{f(u)r_n^2}{(r_n + z \tan \theta)^2} \quad (7.3)$$

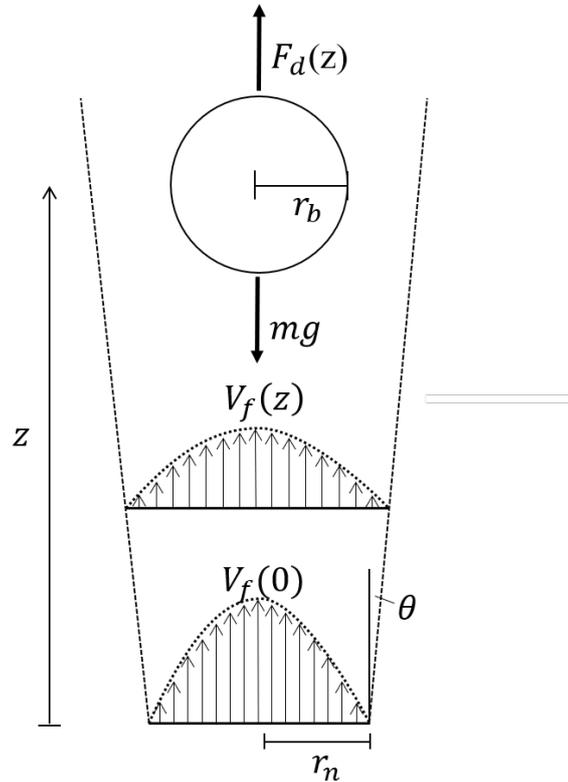


Figure 7.2: Schematic of the airflow model. The flow is assumed to expand conically with a cone angle θ . The flow profile $V_f(z)$ at a height z is modelled as parabolic in the radial direction. The two forces acting on the ball are drag F_d and weight mg .

Next, define the relative airflow around the ball as

$$V(z, \dot{z}) = \bar{V}_f(z) - \dot{z} \quad (7.4)$$

where \dot{z} is the ball velocity. Using the standard approach to drag modelling (e.g., (Gerhart et al., 2016)) gives an expression for $F_d(z)$

$$F_d(z, \dot{z}) = \text{sgn}(V(z, \dot{z})) \frac{1}{2} C_d \pi r_b^2 \rho V^2(z, \dot{z}) \quad (7.5)$$

where C_d is the drag coefficient, r_b is the ball radius and ρ is the density of air. We include the signum function $\text{sgn}(V(z, \dot{z}))$ to ensure that the F_d always acts opposite to $V(z, \dot{z})$, e.g., when $\dot{z} > \bar{V}_f(z)$.

The dynamics of the fan, expansion chamber and airflow, and lags in the control computer and micro-controller introduce latency into the system, which manifests as a lag between a change in control input and a change in the ball force balance. Modelling each of these transient components is challenging, so we model all transience as a lag in u reaching a demanded value, defined as u^* . We set a first-order constraint on the time derivative of u , namely

$$\frac{du}{dt} = k_t(u^* - u) \quad (7.6)$$

where k_t is a time constant to be determined.

Therefore, we can reformulate (7.1) as

$$m\ddot{z} = k_s \frac{1}{2} C_d \pi r_b^2 \rho \left(\frac{f(u) r_n^2}{(r_n + z \tan \theta)^2} - \dot{z} \right)^2 - mg \quad (7.7)$$

where $k_s = \text{sgn}(\bar{V}_f(z) - \dot{z})$. We recast (7.7) in its state-space representation. There are three states

$$\mathbf{x} = \begin{bmatrix} z \\ \dot{z} \\ u \end{bmatrix} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \quad (7.8)$$

Therefore, the state-space representation of the system is

$$\frac{d\mathbf{x}}{dt} = \begin{bmatrix} x_2 \\ k_s \frac{1}{2} C_d \pi r_b^2 \rho \left(\frac{f(x_3) r_n^2}{(r_n + x_1 \tan \theta)^2} - x_2 \right)^2 - mg \\ k_t(u^* - x_3) \end{bmatrix} \quad (7.9)$$

where r_b , ρ , r_n , m and g are known or easily measurable, as shown in Table 7.1, and C_d , $f(\mu)$, θ and k_t are unknown parameters to be identified from the system. Equation (7.9) can be easily integrated using the MATLAB ordinary differential equation solver ODE45, using a non-negative constraint to keep the system in the range $z \geq 0$.

7.1.2 Parameter identification

Table 7.2 shows the identified parameters.

Parameter	Value
r_b	20 mm
r_n	15 mm
ρ	$1.22 \text{ kg}\cdot\text{m}^{-3}$
m	2.7 g
g	$9.81 \text{ m}\cdot\text{s}^{-2}$

Table 7.1: Measured system parameters.

Parameter	Value
θ	1°
C_d	0.185
K_t	3 s^{-1}
$f(u)$	$-30.24u^2 + 47.91u + 2.45$

Table 7.2: Unknown system parameters.

Parameter identification: θ

θ , the velocity profile cone angle, dictates how \bar{V}_f decreases as z increases. A pitot tube was used to measure \bar{V}_f over a range of z and u^* values. A MATLAB nonlinear optimiser (`fminsearch`) was used to minimise the error between the model velocity profile (7.2) and experimental results, finding the optimum cone angle to be $\theta = 1^\circ$. Figure 7.3 shows the experimental results and fitted velocity profile.

Parameter identification: $f(u)$

$f(u)$ defines the mapping between the control parameter u and V_0 . A pitot tube was used to measure V_0 for a range of u values and a second order polynomial was fitted to the data, yielding

$$f(u) = -30.24u^2 + 47.91u + 2.45 \quad (7.10)$$

Figure 7.4 shows the experimental measurements and the fitted model.

Parameter identification: C_d

C_d is an empirically determined parameter that describes the proportion of energy in the flow converted to drag force. The relationship between Reynolds number Re and C_d is well known for a sphere immersed in an infinite uniform flow. The Re regime in this system is of order 10^4 , implying $C_d \approx 0.5$. The mean ball hovering height \bar{z} was recorded for a set of u^* in the range 0.25–0.95 and compared to the model predicted hovering heights. The standard immersed body model with $C_d = 0.5$ overestimated \bar{z} significantly, indicating the actual C_d should be lower. C_d was hand tuned, finding the value that best predicted \bar{z} is $C_d = 0.185$. We

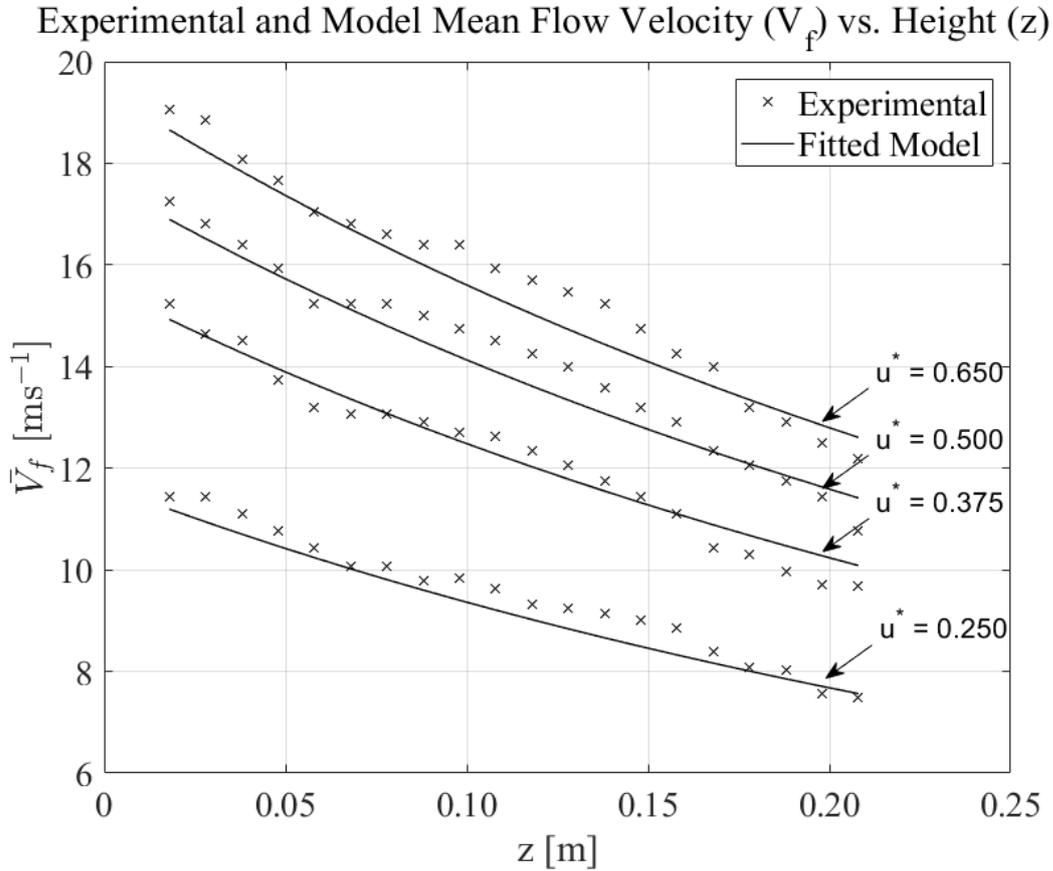


Figure 7.3: Experimentally measured mean flow velocity \bar{V}_f as a function of height z and converged control demand u^* compared with the fitted model.

suggest that this discrepancy between C_d values is because the air-jet diameter is similar to the ball diameter, thus invalidating the assumption of infinite uniform flow. A more rigorous model could investigate how C_d changes as a function of z and $V(z, \dot{z})$; we did not deem this necessary. Figure 7.5 shows the experimental results and model predictions using both C_d values.

Parameter identification: k_t

k_t is used to represent all system latencies. k_t was hand tuned, finding $k_t = 15$, unless the fan is initially off, in which case $k_t = 3$.

7.2 Global asymptotic stability

We have stated that the system exhibits self-stabilising properties in the vertical direction. This is intuitive when observing the system, as the ball seems to always settle around a fixed height for each control input. To be thorough, we present an analysis that proves this self-

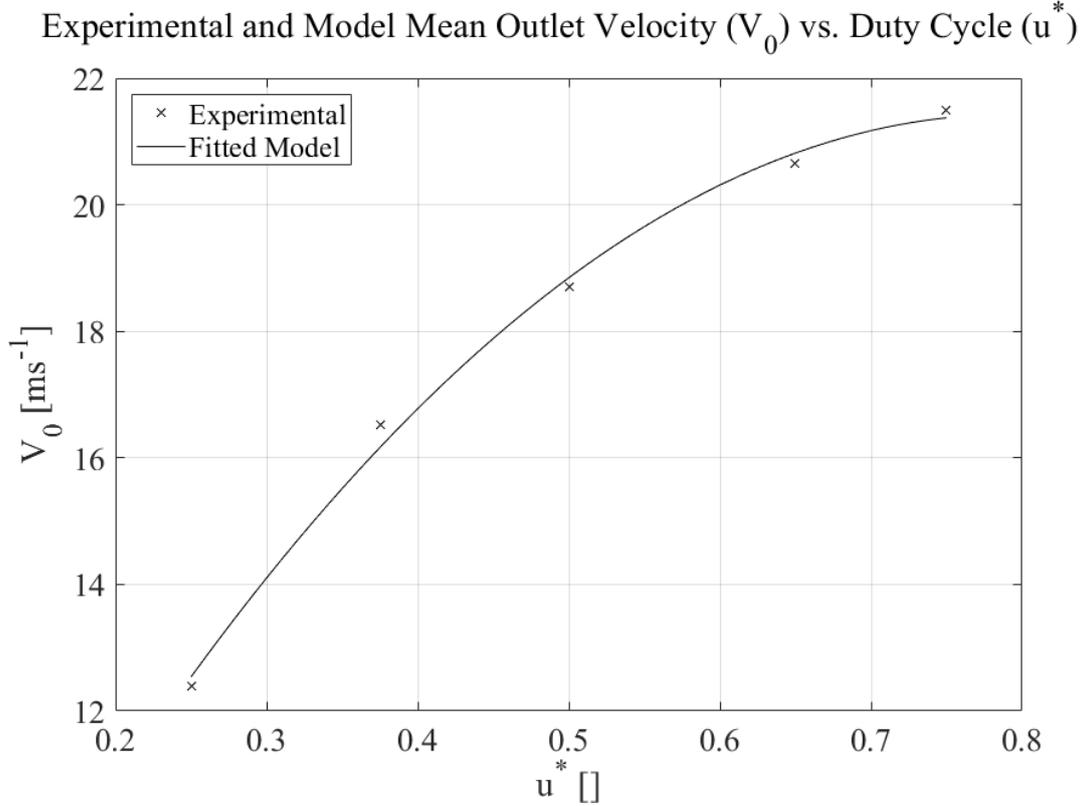


Figure 7.4: Experimentally measured mean outlet nozzle velocity V_0 as a function of converged control demand u^* compared with the fitted model.

stability analytically. Specifically, we show that for any constant and converged control input $u = u^*$, the system will *always* converge to an equilibrium state $\mathbf{x}^* = [z^* \ 0 \ u^*]^T$, providing the initial conditions $\mathbf{x}_{\text{initial}}$ lie on the domain $D_{\mathbf{x}} = \{\mathbf{x} | x_1 > -r_n / \tan \theta\}$; this property is known as global asymptotic stability. Note that $D_{\mathbf{x}}$ is a positive invariant set with respect to the dynamics (7.9) for positive fan speeds due to the singularity at $x_1 = -r_n / \tan \theta$. Of course, in reality $x_1 \geq 0$ as the ball cannot travel below the nozzle, so any physically plausible state \mathbf{x} will by definition lie in $D_{\mathbf{x}}$.

First, define \mathbf{y} as the state variable measured around \mathbf{x}^* , i.e.

$$\mathbf{y} = \begin{bmatrix} z - z^* \\ \dot{z} \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} \quad (7.11)$$

Experimental and Model Mean Hovering Height (z^{mean}) vs. Duty Cycle (u^*)

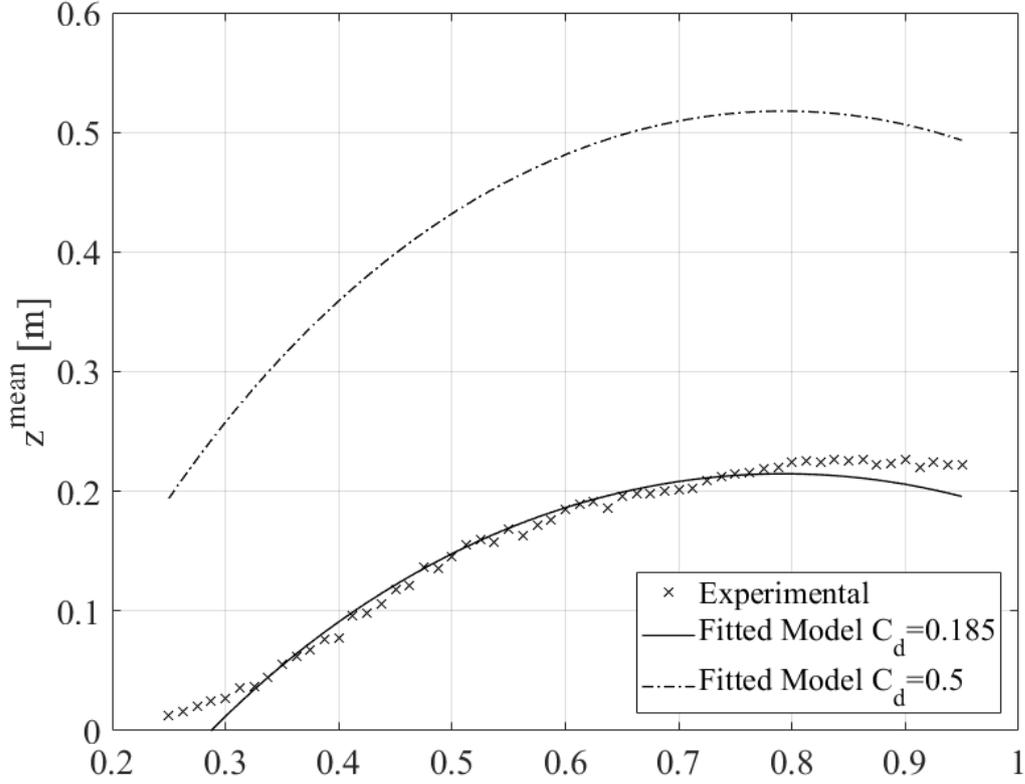


Figure 7.5: Mean steady state hovering height z_{mean} as a function of converged control demand u^* , compared with the model with $C_d = 0.5$ and the fitted model with $C_d = 0.185$.

Since we assume the airflow speed has converged to a constant positive value, we neglect the third system state and assume $u = u^*$. Reformulating (7.9) in this form yields

$$\dot{\mathbf{y}} = \begin{bmatrix} y_2 \\ k_s \frac{C_d \rho r_b^2 \pi}{2m} \left(\frac{f(u^*) r_n^2}{(r_n + (y_1 + z^*) \tan(\theta))^2} - y_2 \right)^2 - g \end{bmatrix} \quad (7.12)$$

valid on the domain $D_{\mathbf{y}} = \{\mathbf{y} | y_1 > -r_n / \tan \theta - z^*\}$. For brevity, denote (7.12) as

$$\dot{\mathbf{y}} = \begin{bmatrix} y_2 \\ F(y_1, y_2) \end{bmatrix} \quad (7.13)$$

We prove the equilibrium $\mathbf{y}^* = [0 \ 0]^T$ is globally asymptotically stable by employing the Lyapunov stability criterion, which requires that a Lyapunov candidate function $L(\mathbf{y})$ exists such that

- $L(\mathbf{0}) = 0$
- $L(\mathbf{y}) > 0 \forall \mathbf{y} \in D_{\mathbf{y}} \setminus \{\mathbf{0}\}$
- $\dot{L}(\mathbf{y}) < 0 \forall \mathbf{y} \in D_{\mathbf{y}} \setminus \{\mathbf{0}\}$, i.e. L decreases along trajectories.

The Lyapunov function can be constructed as

$$L(y_1, y_2) = \frac{1}{2}my_2^2 + W(y_1) \quad (7.14)$$

where $dW/dy_1 = -mF(y_1, 0)$. Hence, $L(y_1, y_2)$ is the sum of kinetic and potential energy, deriving the latter from the sum of gravitational and drag forces but neglecting the effect of the ball velocity y_2 .

To find W , we integrate $-mF(y_1, 0)$, defining the constant of integration such that $W(0) = 0$, to give

$$W(y_1) = -m \left(\frac{K_1}{\tan \theta (K_2 + y_1 \tan \theta)^3} + \frac{K_1}{\tan \theta K_2^3} - gy_1 \right) \quad (7.15)$$

where $K_1 = C_d \rho r_b^2 \pi f^2 (u^*) r_n^4 / 6m$ and $K_2 = r_n + z^* \tan \theta$. Hence, $L(\mathbf{0}) = 0$.

First, we show the state $\mathbf{y} = \mathbf{0}$ is the unique minimum for $L \in D_{\mathbf{y}}$. When in equilibrium, $F(0, 0) = 0$. Since $dW/dy_1 = -mF(y_1, 0)$, W is stationary at this point. Moreover, for $y_1 < 0$, $dW/dy_1 < 0$ and for $y_1 > 0$, $dW/dy_1 > 0$. Hence, the origin is the unique minimum. Similarly, the origin is a global minimum for the term $\frac{1}{2}my_2^2$. Hence, L is the unique minimum for $L \in D_{\mathbf{y}}$.

The time derivative of the Lyapunov function is

$$\dot{L} = my_2 \dot{y}_2 + \frac{dW}{dy_1} \dot{y}_1 \quad (7.16)$$

which, after substituting for y_1 , dW/dy_1 and \dot{y}_2 yields

$$\begin{aligned} \dot{L} &= my_2 (F(y_1, y_2) - F(y_1, 0)) \\ &= y_2 K_3 (\text{sgn}(\bar{V}_f - y_2) (\bar{V}_f - y_2)^2 - \bar{V}_f^2) \end{aligned} \quad (7.17)$$

where $K_3 = \frac{C_d \rho r_b^2 \pi}{2m}$ and $\bar{V}_f = (f(u^*) r_n^2) / ((r_n + (y_1 + z^*) \tan(\theta))^2)$.

Now, we show that L decreases along trajectories, and in fact $\dot{L} < 0$ except instantaneously when $y_2 = 0$. We prove this by considering four cases dictated by the signum function in (7.17).

Case 1: $0 < y_2 \leq \bar{V}_f \rightarrow \text{sgn}(\bar{V}_f - y_2) = 1$ Since $\bar{V}_f^2 > (\bar{V}_f - y_2)^2$, the right hand term in (7.17) dominates so $\dot{L} < 0$.

Case 2: $0 < \bar{V}_f < y_2 \rightarrow \text{sgn}(\bar{V}_f - y_2) = -1$ Both terms in (7.17) are negative so $\dot{L} < 0$.

Case 3: $y_2 < 0 \rightarrow \text{sgn}(\bar{V}_f - y_2) = 1$ The left term is positive and dominates in (7.17), but y_2 is negative so $\dot{L} < 0$.

Case 4: $y_2 = 0 \rightarrow \dot{L} = 0$. This violates the Lyapunov criterion as clearly $\dot{L} \leq 0$, i.e. not strictly negative definite; however, we know from (7.13) that when $y_2 = 0$, $\dot{y}_2 = F(y_1, 0) \neq 0$. Hence, \dot{L} is only zero instantaneously and the dynamics always return to the provably stable regions shown in cases 1–3. (In fact, this point indicates an inflection point for L). These conditions still qualify for LaSalle’s theorem, which is a proof for global asymptotic stability of the equilibrium \mathbf{y}^* in a positive invariant set $D_{\mathbf{y}}$.

Hence, all conditions have been met (with the modification of \dot{L} requiring to be only negative semidefinite on $D_{\mathbf{y}}$ to qualify for LaSalle’s theorem) and the equilibrium \mathbf{y}^* is globally asymptotically stable on $D_{\mathbf{y}}$ which is equivalent to \mathbf{x}^* being globally asymptotically stable on $D_{\mathbf{x}}$ for converged fan speeds. The stability properties of the system can be visualised with a sample trajectory in a phase portrait, as shown in Figure 7.6.

7.3 Augmented self-stability

The Lyapunov analysis proves that the system is self-stable, so will asymptotically converge to a target height when given a constant control demand u_{target}^* , which we can calculate using the dynamics model. However, this self-stability is slow in reaching the target position as the speed of convergence is dictated by the system dynamics, which we cannot change.

Augmenting the self-stable system with a bang-bang control is advantageous as it allows us to alter the dynamics while being guaranteed stability of the system around the target state. Indeed, it is important to note that by using bang-bang control, we are not able to destabilise the system for fan speeds larger than zero. This is possible mainly because the set of possible initial states is positive invariant with respect to the dynamics (7.9), as shown in the Lyapunov analysis in 7.2.

In the next section, I present the bang-bang control policy and analyse it in simulation.

7.3.1 Bang-bang control algorithm

We seek to move the system from an initial state $\mathbf{x}_{\text{initial}}$ to a target state $\mathbf{x}_{\text{target}}$. In our system, these states correspond to target hovering heights; for a target height z_{target} which corresponds to a control input of u_{target}^* , $\mathbf{x}_{\text{target}} = [z_{\text{target}} \ 0 \ u_{\text{target}}^*]^T$. It is important to note here

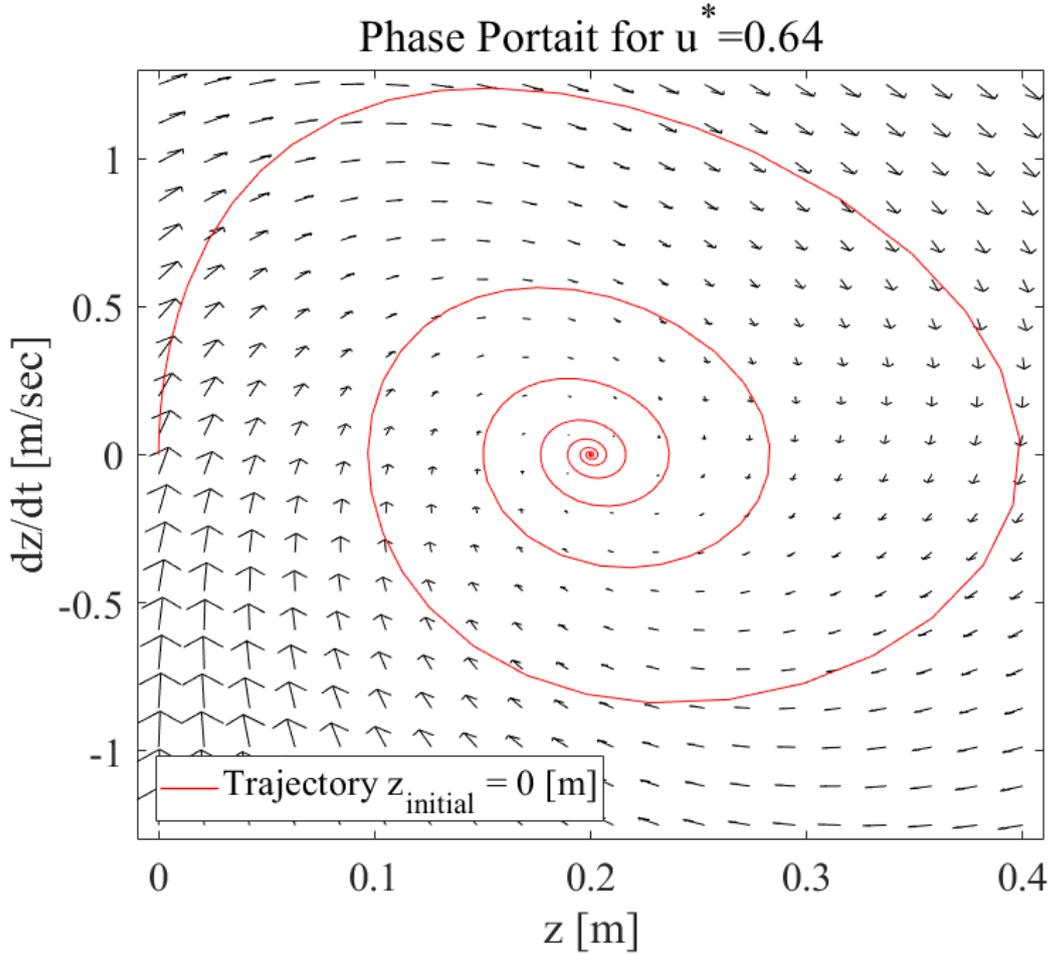


Figure 7.6: Phase portrait and sample trajectory of self-stable vertical dynamics (7.9) around equilibrium $z_{eq} = 0.2$ for converged fan speed $u = u^*$.

that because of the transient system dynamics, control demands are included in the target state.

We implemented a numerical switching time computation method, similar to the method outlined in (Kaya and Noakes, 1996). The proposed bang-bang controller switches between the maximal, minimal and target control inputs u_{high}^* , u_{low}^* and u_{target}^* at a set of three predetermined switching times. We denote the i th switching time as t_i and the initial and final times as t_0 and t_f , such that $0 = t_0 \leq t_1 \leq t_2 \leq t_3 = t_f$. The system trajectory between each switching time, denoted by $\mathbf{x}_i(t)$, is called the i th bang arc, so the total trajectory $\mathbf{x}(t)$ for $0 \leq t \leq t_f$ is the concatenation of $\mathbf{x}_i(t)$ and ends at the final state \mathbf{x}_f . The arc time ξ_i is the time spent on each bang arc, which together are stored in the arc time vector $\xi = [\xi_1 \ \xi_2 \ \xi_3]^T$. Arc times ξ_1 and ξ_2 represent the time spent at u_{high}^* and u_{low}^* , or vice versa, while ξ_3 is a

predefined time used to set a consistent point to measure control performance. Figure 7.7 shows an example of the bang-bang control policy.

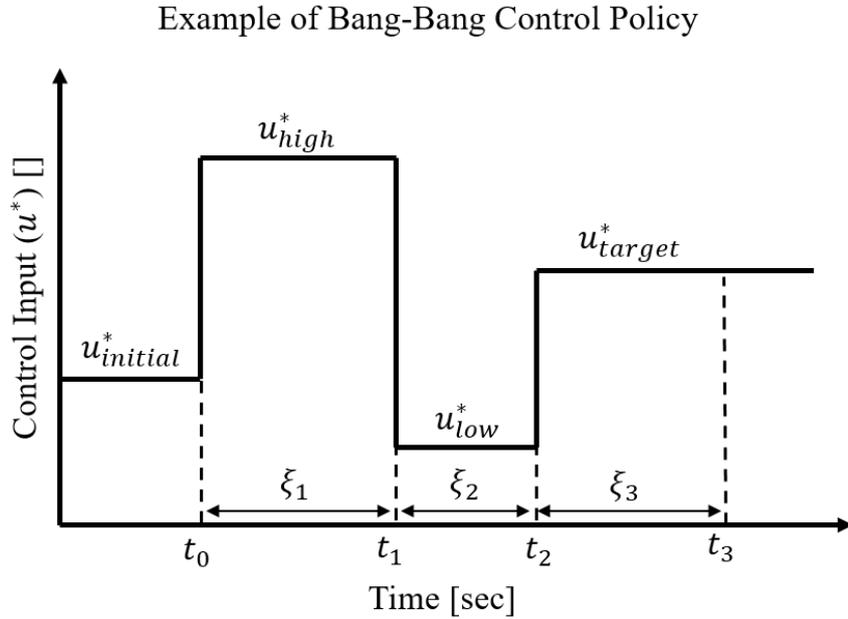


Figure 7.7: Example of a bang-bang control policy showing the switching times $[t_1 \ t_2 \ t_3]$, arc times $[\xi_1 \ \xi_2 \ \xi_3]$, control demands $[u_{high}^* \ u_{low}^* \ u_{target}^*]$ and initial conditions t_0 and $u_{initial}^*$.

We seek to determine ξ such that the error between \mathbf{x}_f and \mathbf{x}_{target} is minimised. We formally define this using a control cost function $g(\xi)$, representing the absolute distance between \mathbf{x}_f and \mathbf{x}_{target} ,

$$g(\xi) = \|\mathbf{x}_f - \mathbf{x}_{target}\| \quad (7.18)$$

Hence, we seek to solve the minimisation problem

$$\min_{\xi} = g(\xi) \quad (7.19)$$

Due to the non-linearity of $g(\xi)$ an analytical solution to (7.19) is not necessarily tractable, so the MATLAB `fminsearch` numerical optimisation function was used. Such numerical solvers require an initial solution guess, which we find using a coarse search of $g(\xi)$ in the range $0 \leq \xi_i \leq 1$.

The bang-bang control method was analysed in simulation for an idealised system with no latency and the full system including latency, allowing us to analyse the effect of latency on the control policy.

7.3.2 Simulation analysis of idealised system

First, consider an idealised system in which we ignore the transient dynamics described in (7.6). This results in a system where control demands are immediately achieved, e.g., $u = u^*$ always.

By using the numerical optimisation method in 7.3.1, we are able to systemically analyse the bang-bang control policy over the whole state-space of initial conditions. When we consider $z_{\text{target}} = 0.1$ m, for example, the `fminsearch` function provides an optimal $[\xi_1, \xi_2]$ for each initial condition z_{initial} .

By scanning through the parameter space of z_{initial} , a single *switching curve* can be obtained, as shown in Figure 7.8a. This curve indicates the point in state-space to switch control inputs in order to achieve time-optimised control.

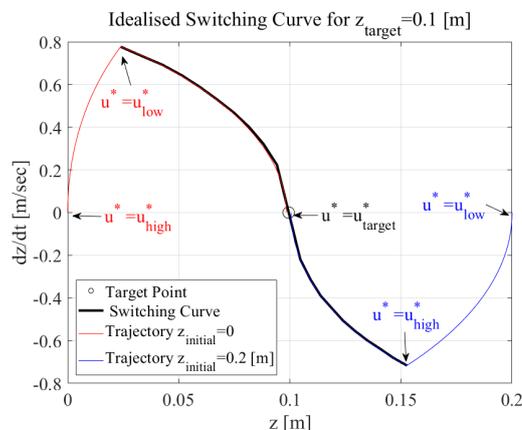
The curve shows that for $z_{\text{initial}} < z_{\text{target}}$, control is achieved by applying u_{high}^* until the trajectory meets the switching curve, at which point the control input switches to u_{low}^* . The controller stays on this trajectory until z_{target} is reached, when it switches to u_{target}^* . For $z_{\text{initial}} > z_{\text{target}}$, control is the same except that u_{low}^* is applied before u_{high}^* , which is intuitive as change in z is negative. The two example trajectories in Figure 7.8a demonstrate the aforementioned behaviour.

7.3.3 Simulation analysis of latency-included system

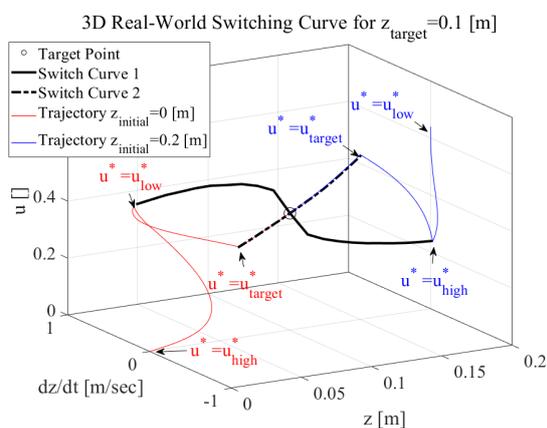
As discussed in 7.1.1, there are physical and computational latencies in the real system, which we model as a lag in the control signal reaching the demanded signal. This complicates the bang-bang control problem, as now the time taken to switch between u_{high}^* , u_{low}^* and u_{target}^* must be considered.

As in 7.3.2, we scanned through the z_{initial} parameter space and generated a switching curve, as shown in Figure 7.8b. We first note that the switching curve exists in three-dimensions, as switching is dependent not only on z and \dot{z} but also the current control signal u . As a result, there are in fact two switching curves. The first indicates at what z , \dot{z} and u we should switch our control demand between u_{high}^* and u_{low}^* , or vice versa. The second indicates when to switch to the target control demand u_{target}^* . Unlike in the idealised system where switching occurs at the $\mathbf{x}_{\text{target}}$, here the final switch happens before the target, so the ball arrives at the target height as the control signal converges to u_{target}^* .

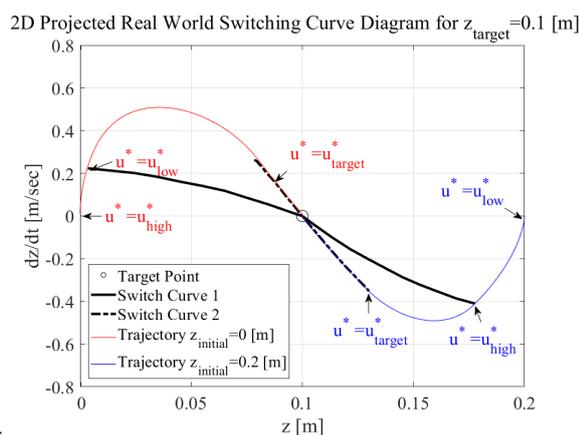
We can better visualise this switching behaviour by projecting the three-dimensional curves in Figure 7.8b onto the z - \dot{z} plane, as shown in Figure 7.8c. Here, we can see that the switching behaviour is generally similar, but that the curve corresponding to switch number one is shallower. This means latency causes switching to happen at lower \dot{z} values, essentially limiting the aggressiveness of the controller and acting as a form of natural system damping.



(a) Switching curve for idealised system with no latency.



(b) Three-dimensional switching curve for real-world system with latency.



(c) Two-dimensional switching curve projection for real-world system with latency.

Figure 7.8: (a) Idealised and (b,c) real-world switching curves for $z_{\text{target}} = 0.1$ m, with example trajectories for $z_{\text{initial}} = 0$ and 0.2 m

7.4 Experimental results

An experimental rig was built to test and validate the dynamics model and control approach described in 7.1.1 and 7.3 in the real-world system. This section describes the experimental set-up, presents a comparison between the model and real-world system and analysis discrepancies between the two.

7.4.1 Experimental set-up

A high-power brushless motor (Diatone M2205 2300KV) fitted with a 3-blade propeller (Diatone 5045) and a constant input voltage of 15 VDC was used to propel air into an

expansion chamber then through a circular nozzle positioned on top of the chamber. To stabilise the outlet jet, a flow straightener was fitted under the nozzle. Control was achieved using an electronic speed control (HobbyKing 20A ESC 3A UBEC) connected to an Arduino UNO interfaced with MATLAB. A pulse width modulation (PWM) signal was used to vary motor speed by changing the available current. The ball height was measured using a high frame rate webcam (Logitech BRIO) capable of delivering 90–120 frames per second. In each frame the ball was located using a simple thresholding algorithm, after which the height was determined using the predetermined camera parameters.

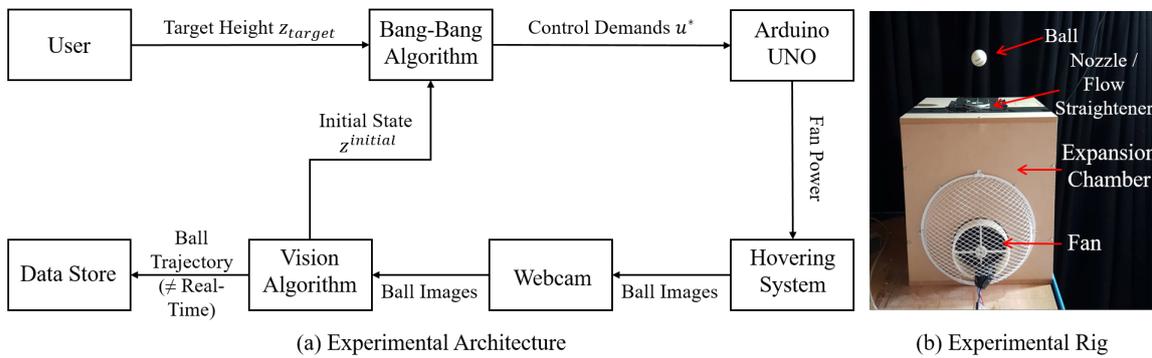


Figure 7.9: Schematic of experimental architecture.

7.4.2 Dynamics model validation

The dynamics model is used to compute bang-bang switching times, described in 7.3. As shown in Figure 7.7, bang-bang control is essentially a series of control step inputs. Each step input leads to oscillatory transient dynamics as the ball moves toward a new equilibrium position. Bang-bang control cancels out these oscillation by changing the control input at an appropriate time. This period of transient dynamics can be characterised by the rise time t_r (the time for the ball to move from 10% to 90% of the steady-state value z_p) and the peak value (the maximum height attained by the ball). The accuracy of the dynamics model during this period will influence the effectiveness of bang-bang control. For example, if the dynamics model overestimates the rise time, then the switching times will be too long, resulting in overshoot.

The parameter identification described in 7.1.2 ensured the dynamics model accurately predicted the settled hovering height and system lags for any valid control input. However, the transient dynamics could not be directly identified. Hence, we experimentally measured the system response to 15 step inputs, from $z = 0$, in the range $0.25 \leq u^* \leq 0.95$ and compared the rise time and peak value characteristics with the corresponding dynamics model trajectory.

Table 7.3 shows the results, while Figure 7.10 shows one representative step response from the experiment. First, we quantitatively discuss the step response characteristics. The experimental time was on average 50 ms faster than the dynamics model prediction, with the error increasing at lower u^* values and decreasing with higher u^* values. The average error in peak value was 0.0262 m, which stayed relatively constant as a function of u^* . The increase in rise time error at lower u^* values is likely because this is where velocity is highest and air-jet diameter is lowest. Under these conditions, our model assumptions about the drag coefficient and flow nature may become invalid, for example, because of increased turbulence around the ball. Anecdotally, we note that horizontal instability was at a maximum here. Conversely, the increase in model accuracy at higher heights is to be expected as the flow is wider and more uniform, as assumed in our model.

u^* m	OL t_r m	MD t_r m	OL z_p m	MD z_p m
0.25	0.2791	0.1576	0.0194	0.0010
0.30	0.5272	0.1594	0.0360	0.0344
0.35	0.4027	0.1728	0.0733	0.0930
0.40	0.1996	0.1839	0.1407	0.1423
0.45	0.2642	0.1921	0.2080	0.1849
0.50	0.2390	0.1976	0.2599	0.2224
0.55	0.2050	0.2016	0.2974	0.2551
0.60	0.2387	0.2025	0.3333	0.2846
0.65	0.2165	0.2018	0.3574	0.3103
0.70	0.2144	0.1992	0.3901	0.3319
0.75	0.1700	0.1944	0.4044	0.3517
0.80	0.1750	0.1881	0.4157	0.3683
0.85	0.1448	0.1810	0.4124	0.3821

Table 7.3: Open-loop (OL) and Model (MD) rise time t_R and peak value z_p .

Qualitatively, we see from Figure 7.10 that the model and experimental trajectories are broadly similar over the whole transient phase, with the period and amplitude of oscillations generally matching. Furthermore, the parameter identification is clearly effective at predicting the mean settled height. However, we note that our model does not include noise, so we cannot predict the oscillation about the mean seen in the experimental results.

7.4.3 Control demonstration

The main purpose of the augmented self-stability approach outlined in 7.3 is to improve the time-to-target as compared to fully open-loop control, while still making use of the self-stabilising properties. We have demonstrated the approach works in simulation, and here seek to demonstrate it also works in the real world system. The possible issues with

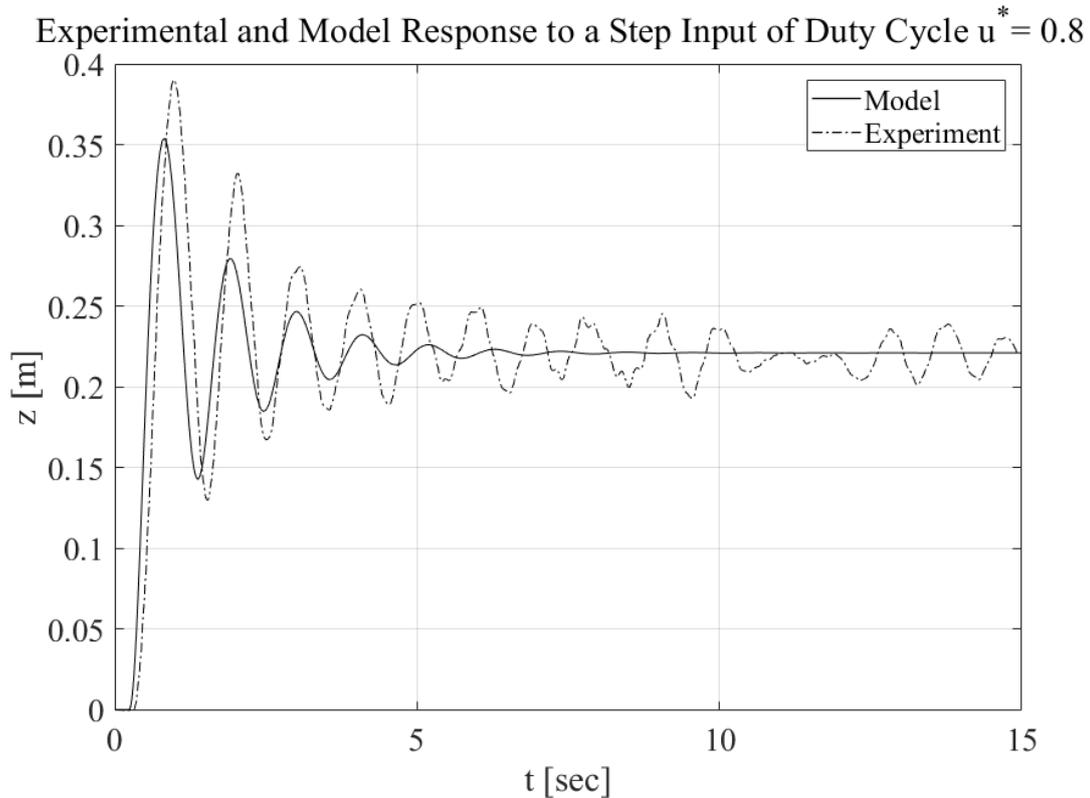


Figure 7.10: Example trajectory of dynamics model and real-world system for control demand $u^* = 0.8$ and $z_{\text{initial}} = 0$ m, measured for 15 seconds.

implementing the algorithm in real life include model inaccuracies, disturbances and system noise. However as described in 7.3, utilising self-stability should provide robustness to such problems.

The control algorithm was applied for three different target heights $z_{\text{target}} = 0.1, 0.15, 0.2$ m, all with the ball initially at rest. The system response was then recorded for a further 5 seconds. For each z_{target} the corresponding open-loop response was also applied for 15 seconds, long enough for the ball to settle to steady state.

Figure 7.11 shows the state-space trajectory for each case. The open-loop response in all cases is characterised by a spiralling trajectory ending in steady-state behaviour, which manifests as an oscillation around the target height. Conversely, the controlled response – additionally shown in snapshot form in Figure 7.12 – is characterised by a curved trajectory directly to steady-state. We measured the time taken for the system to reach its steady state behaviour – the settling time, t_s – in the open-loop and controlled response. For $z_{\text{target}} = 0.1, 0.15, 0.2$ m, the settling time was a factor of 2.4, 3.6 and 4.4 faster in the controlled response, as shown in Table 7.4. This difference in improvement based on z_{target}

can be attributed to a number of factors. As discussed in 7.4.2, the dynamics model is less accurate at duty-cycles, meaning the bang-bang control policy is likely to be less effective when z_{target} is lower. Furthermore, because the system dynamics are highly non-linear we expect the relationship between z_{target} and t_s to also exhibit non-linearity.

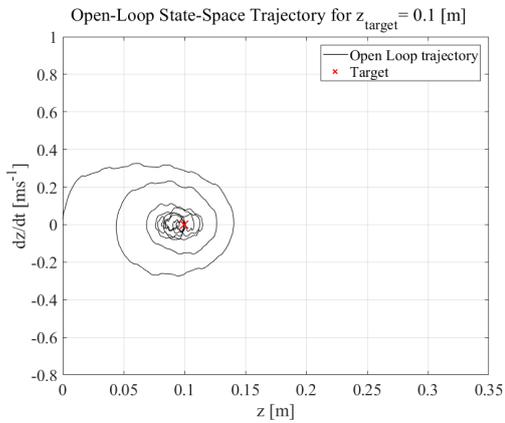
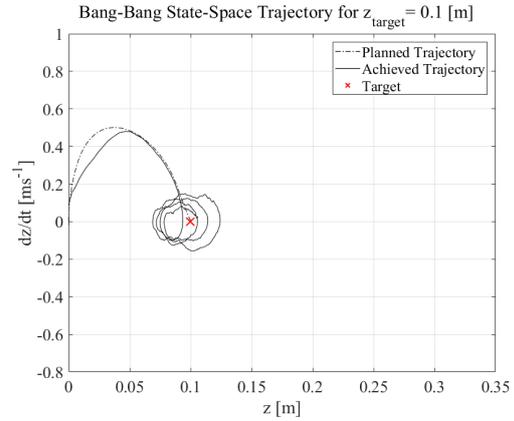
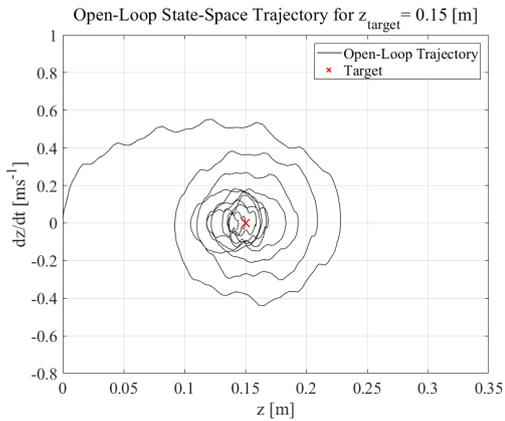
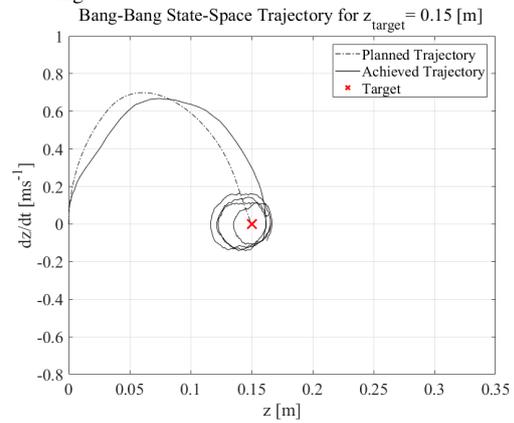
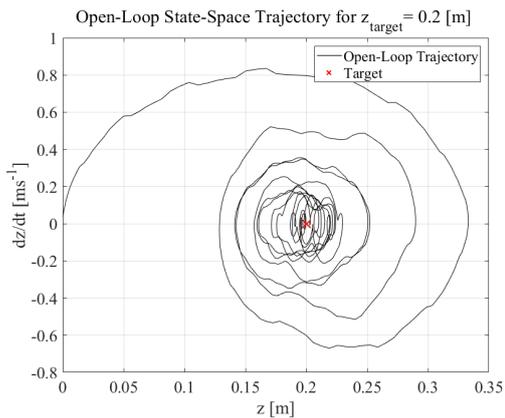
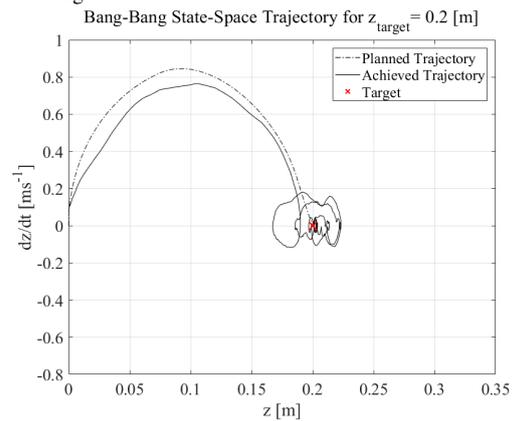
(a) Open-loop response for $z_{\text{target}} = 0.1$ m(b) Planned and achieved bang-bang response for $z_{\text{target}} = 0.1$ m(c) Open-loop response for $z_{\text{target}} = 0.15$ m(d) Planned and achieved bang-bang response for $z_{\text{target}} = 0.15$ m(e) Open-loop response for $z_{\text{target}} = 0.2$ m(f) Planned and achieved bang-bang response for $z_{\text{target}} = 0.2$ m

Figure 7.11: Open-loop response (a,c,e) and bang-bang planned and achieved response (b,d,f) for $z_{\text{target}} = 0, 0.15, 0.2$ m with $z_{\text{initial}} = 0$ m.

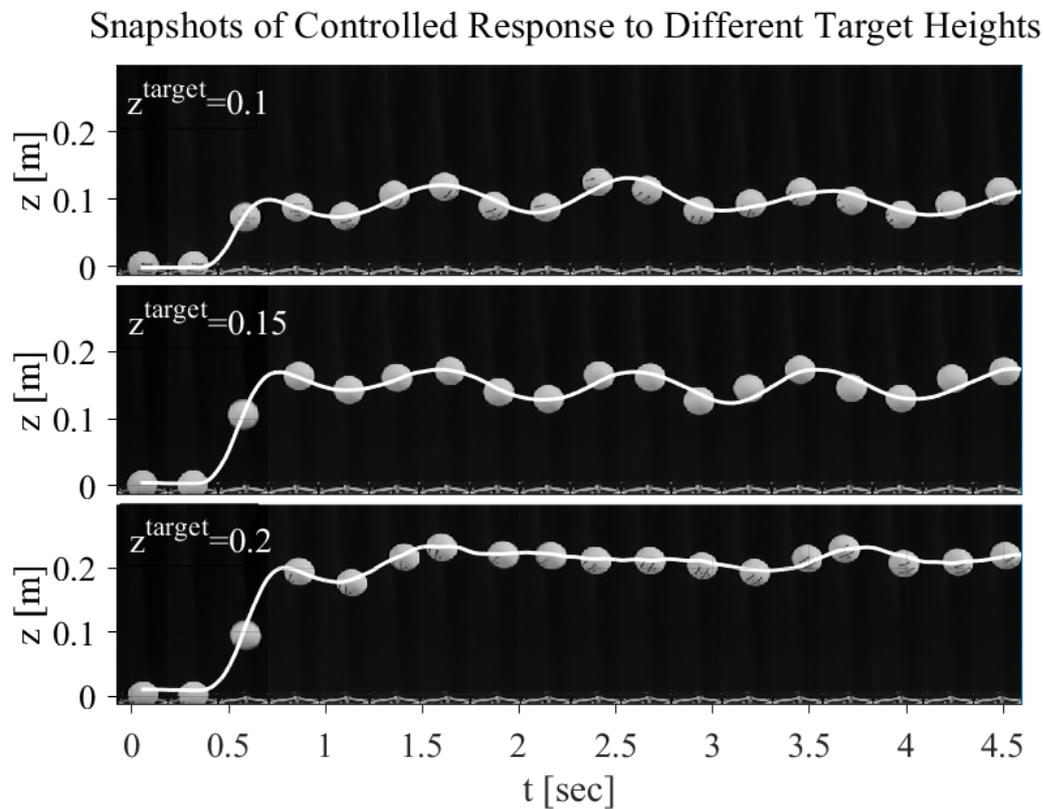


Figure 7.12: Snapshot images recorded over five seconds of controlled response for $z_{\text{target}} = 0.1, 0.15, 0.2$ m and $z_{\text{initial}} = 0$ m

These results show the control approach improves the time-to-target when moving the ball from rest to different target heights. A more challenging task is controlling the ball from a non-zero initial height, because the system noise adds uncertainty in $\mathbf{x}_{\text{initial}}$. We tested the control approach for consecutive trajectory following, where z_{target} switched every 5 seconds and applied open-loop control for comparison. Figure 7.13 shows the time-series trajectories.

Qualitatively, we note that the controlled trajectory is more effective than open-loop at moving the ball between heights. As previously discussed, there also appears to be a relationship between the direction and magnitude of height change and performance difference.

7.4.4 Analysis

The results in 7.4.2 and 7.4.3 demonstrated the effectiveness of the control approach. Here, we investigate the nature of the control cost function $g(\xi)$ and how the time-to-target improvement in simulation translates to an improvement in real-world.

z_{target} m	Open-Loop t_s m	Controlled t_s m	Improv. Factor
0.1	2.3	0.98	2.4
0.15	4.1	1.12	3.6
0.2	5.2	1.20	4.4

Table 7.4: Open-loop and controlled settling time and improvement factor, for $z_{\text{target}} = 0.1, 0.15, 0.2$ m and $z_{\text{initial}} = 0$ m.

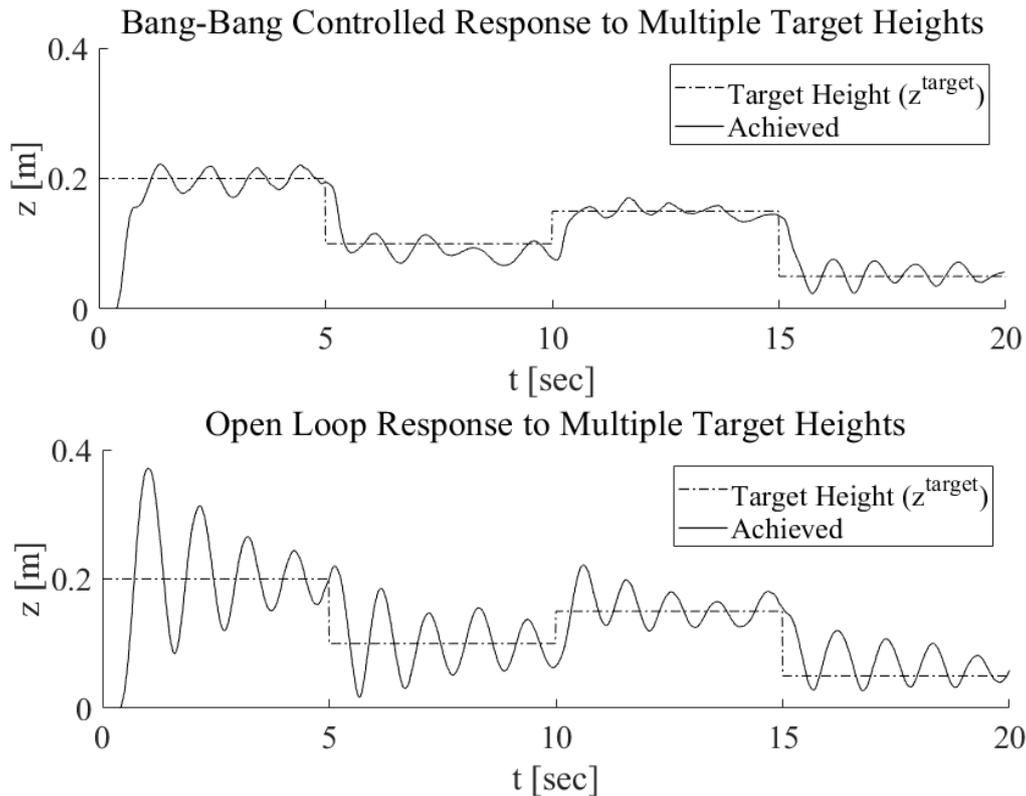


Figure 7.13: Controlled and open-loop response to consecutive trajectory following.

First, we numerically calculated the value of the control cost function $g(\xi)$ in simulation over the parameter space $0 < \xi_1, \xi_2 < 0.5$. The function is represented as a two-dimensional surface, as shown in Figure 7.14a. Assume $g(\xi)$ represents an analogue for time-optimality.

There is a clearly optimal ξ , marked on the figure. The control cost around this is characterised by a curved area of similar control cost. In this area, the relationship between a change in ξ_1 or ξ_2 and $g(\xi)$ is not clear. We suggest this is a manifestation of the highly non-linear system dynamics.

We repeated the same analysis in the real world. The parameter space was discretised into 121 equally spaced points. The resultant bang-bang control policy at each point was

tested and the control cost measured. To account for system noise, $g(\xi)$ was averaged over last 0.25 seconds of motion. A surface was fitted to the points, as shown in Figure 7.14b.

We can see that the general control cost landscape is similar to the simulation, although the absolute $g(\xi)$ appear to be lower. Furthermore, we see that the optimal ξ in the simulation ([0.4 0.2]) and real-world system ([0.4 0.05]) matches in ξ_1 but differs in ξ_2 . We suggest this is due to modelling discrepancies in the system latency. However as shown in 7.4.3, the optimal ξ as calculated from the dynamics model is effective when implemented in the real-world system. One explanation is that since the dynamics model doesn't account for system noise, $g(\xi)$ is measured relative to the absolute target state $\mathbf{x}_{\text{target}}$. However, in reality any bang-bang policy that deposits the ball within the range of system noise is very similar in terms of optimality. Hence, any control policy that is at least close-to-optimal will be effective when implemented in the real system.

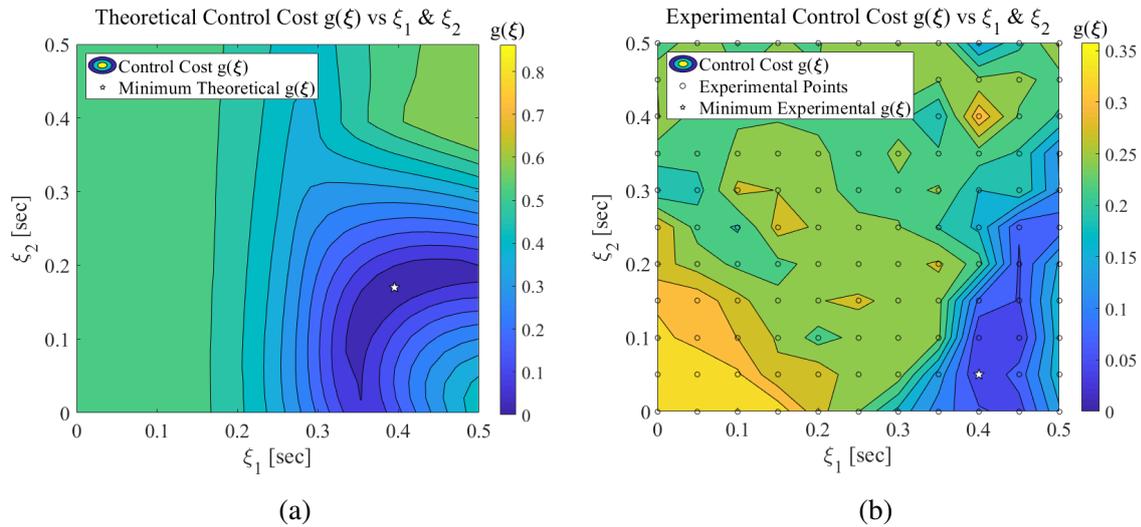


Figure 7.14: (a) Model and (b) experimental control cost $g(\xi)$ as a function of arc times ξ_1 and ξ_2 .

7.5 Additional behaviours: bifurcations via environmental modulation

Changing the mechanical properties, e.g., weight, in a single Bernoulli-ball system does not fundamentally change the types of behaviour the system exhibits. For heavier or lighter balls the two basic behaviour modes are stable hovering or falling. However, we are able to change the system behaviours by modulating the environment, e.g., the airflow. I investigated the effect of an oscillating airflow. Similar to the work on bouncing balls (Holmes, 1982) and the blind juggler (Ronsse et al., 2007), I tested how the system behaves with a periodic input

of varying frequency. The airflow power was set to an oscillating square wave, between a high and low power value, with a frequency ω ($\text{rad}\cdot\text{s}^{-1}$).

By varying ω we can modulate how the trajectory of the ball relative to the environmental conditions. I recorded the ball height trajectory for multiple forcing frequencies over the range $10\text{--}18 \text{ rad}\cdot\text{s}^{-1}$. For each frequency the system was left to settle for 30 seconds, then the trajectory was recorded for 12 periods of the forcing frequency. To measure how the ball trajectory changes relative to the forcing frequency I measured the ball height, z_{cycle} , and velocity, $v_{\text{v-cycle}}$ at the mid-point of each forcing period. Figure 7.15a shows z_{cycle} , measured at the mid-point of the 12 forcing periods, for all tested ω values. Figure 7.15e shows $v_{\text{v-cycle}}$ measured for the same locations. These plots demonstrate how the balls behaviour changes relative to the environment. At lower frequencies, $\omega = 10\text{--}12 \text{ rad}\cdot\text{s}^{-1}$, the ball exhibits a period one type behaviour where the trajectory is synchronised with the environment, moving up and down at the same frequency. The ball height and velocity measured at the midpoint of each forcing period are within a single continuous range, e.g., the ball reaches a similar height and velocity at the mid point of each period, as shown in Figure 7.15b and f. At intermediate frequencies, $\omega = 13\text{--}16.25 \text{ rad}\cdot\text{s}^{-1}$, the ball exhibits period 2 type behaviours. Here, the ball oscillates at approximately half the environmental forcing frequency, moving up and down once for every two forcing periods as shown in Figure 7.15c and g. We see this Figure 7.15a, where for most frequency values the measured height at the mid-point of forcing periods is either a high or low value. This effect is also seen in the measured velocity, how the distinction is less clear since velocity measurements are inherently more noisy than height measurements. We also see that for some frequency values, e.g., $13.75\text{--}14.1 \text{ rad}\cdot\text{s}^{-1}$ the system collapses back to a period one behaviour. For high frequencies, $\omega > 16.5 \text{ rad}\cdot\text{s}^{-1}$ exhibits a range of frequencies, with a weak periodic relationship with the forcing frequency Figure 7.15d and h. These results show the ball exhibits a bifurcation behaviour, similar to that seen in pendulum or bouncing ball systems (Holmes, 1982). We are able to push the system into more complex period two behaviours by manipulating the environment, before it collapses back to period one and steady state behaviours. The frequency ranges in which we can see clear period two behaviours are relatively narrow, implying the system relies on precise environmental conditions to exhibit certain behaviours.

To further analyse the relationship between the environmental forcing frequency and ball trajectory we performed a frequency analysis. By taking the fast Fourier transform (Brigham, 1988) for each trajectory I extracted the dominant frequencies of the ball vertical trajectory and compare them to the know environmental forcing frequency ω . Figure 7.16 shows the frequency power spectrum for the ball trajectory with three airflow forcing frequencies, corresponding to the example trajectories shown in Figure 7.15 b–d and f–h. The higher

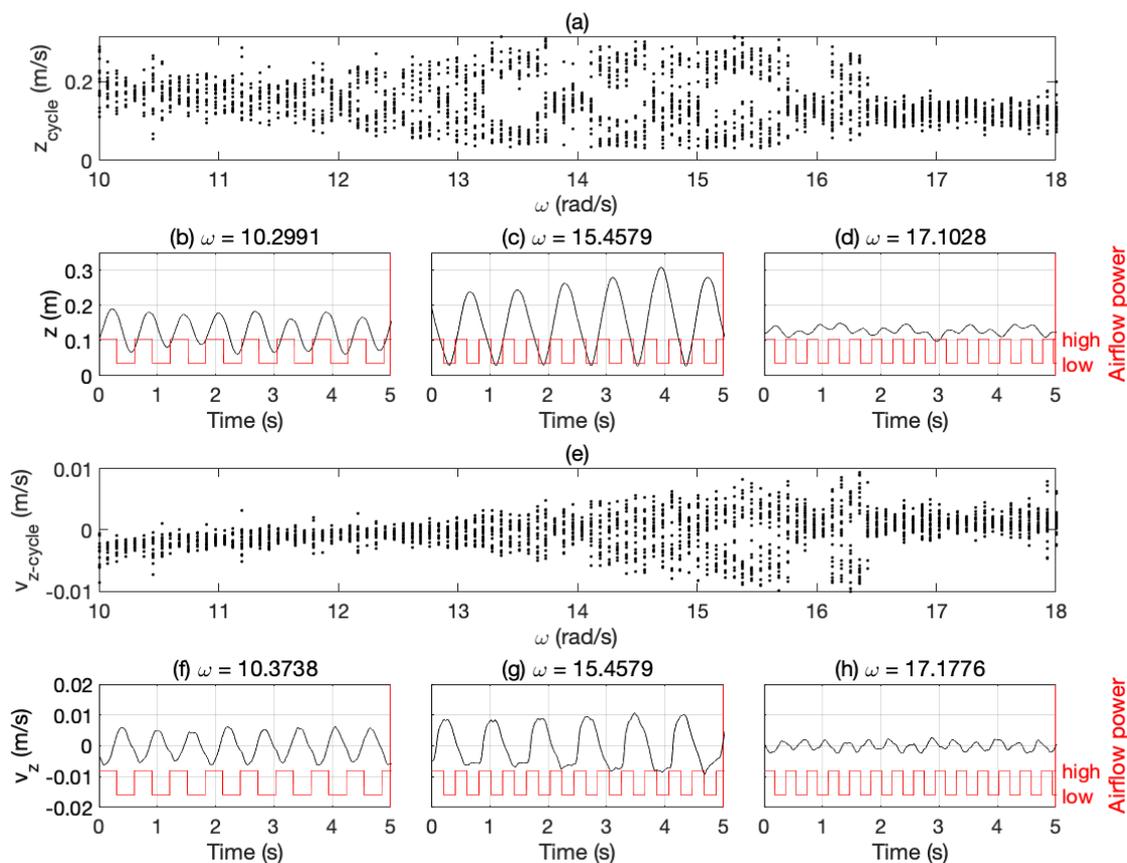


Figure 7.15: Bifurcation behaviours of a hovering ping-pong ball in an airflow oscillating between a high and low power with frequency ω . (a,e) Repeated measurements of vertical height z_{cycle} and velocity $v_{z\text{-cycle}}$ measured at the mid point of the airflow forcing period for range of ω forcing frequencies. (b–d,f–h) vertical height z and velocity v_z of the ping-pong ball at different airflow forcing frequencies, showing period one (b,f), period two (c,g) and period one / random (d,h) behaviours.

the power spectrum value, the higher the dominance of the corresponding frequency is in the ball motion. For period one behaviour (Figure 7.16a) we see that for a forcing frequency of $\omega = 10.30 \text{ rad}\cdot\text{s}^{-1}$ the frequency power spectrum of the ball peaks at a similar value of $10.05 \text{ rad}\cdot\text{s}^{-1}$. For period two behaviour (Figure 7.16b) we see that for a forcing frequency $\omega = 15.46$ the frequency power spectrum of the ball peaks at approximately half this value $7.54 \text{ rad}\cdot\text{s}^{-1}$. For high forcing frequencies (Figure 7.16c) such as $\omega = 17.10 \text{ rad}\cdot\text{s}^{-1}$ we see weaker relationship with the frequency spectrum, which has peaks at values of $16.97 \text{ rad}\cdot\text{s}^{-1}$ and $6.91 \text{ rad}\cdot\text{s}^{-1}$.

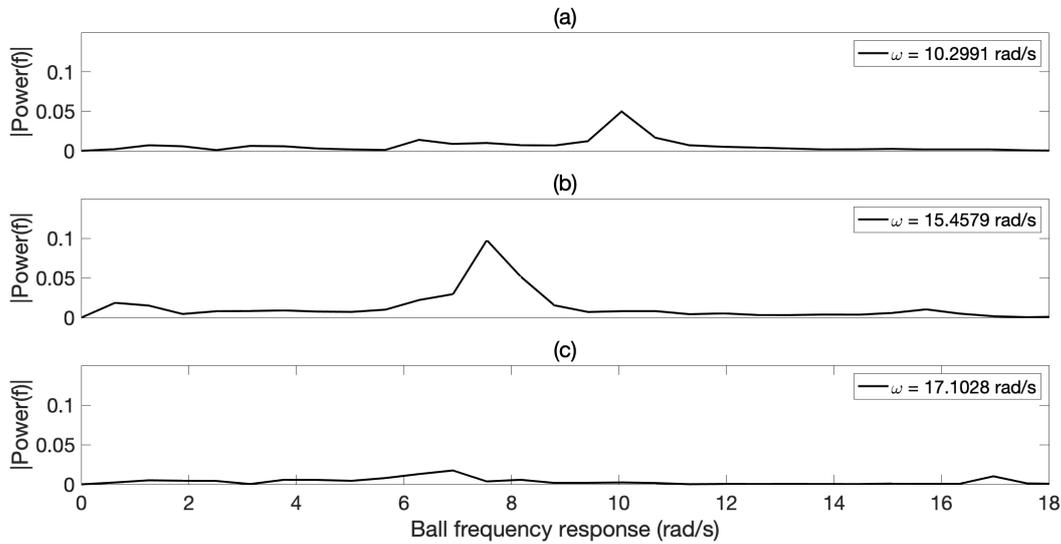


Figure 7.16: Fourier fast transform of hovering ping-pong ball vertical trajectory under three airflow forcing frequencies, leading to period one (a), period two (b) and period one random (c) behaviours.

7.6 Discussion and conclusion

I have shown that the augmented self-stability in the Bernoulli-ball can significantly improve the time-to-target compared to the uncontrolled system while preserving stability. This example demonstrates how fast dynamics can be controlled in a robotic system without the need for high-bandwidth control. We have to be careful in generalising the results to other robotic systems as this approach can technically only guarantee stability in the presence of a globally asymptotically stable equilibrium in the uncontrolled system. However, in the real world experiment we have achieved stable performance in spite of the ground restriction to positive ball heights (the invariant set D_x should include also negative heights as explained in 7.2 to prove stability), which shows that in practice the approach can work for locally self-stable equilibria, allowing its applicability to a larger class of problems including passive dynamic walking (McGeer et al., 1990) and bicycle stability (Kooijman et al., 2011).

An interesting aspect of the single Bernoulli-ball system is that of noise. As show in Figure 7.10, steady-state hovering behaviour is characterised by low-level noise caused by the unstable airflow interacting with the ball. In many ways, this makes the control problem simpler, since all we must do when changing heights is move the ball to within the region of noise that surrounds the target height in the state-space. In the context of embodied intelligence, we see that certain morphology-environment interactions, such as between the airflow and ball, allow an agent to converge to the same behaviour from multiple initial

conditions and in the presence of unpredictable system inputs. We can draw a parallel with the falling paper system, which we have seen can converge to the same behaviour type despite changes in morphology and environment.

In conclusion, augmented self-stability is a simple yet effective approach to improve the performance of self-stable robotic systems. It outsources the difficult tasks of stabilising control to the mechanical feedback loop, leaving the trajectory planning via bang-bang control to a numerical optimisation which can be solved offline. The high performance and robustness of the presented results thus suggests that the focus on the mechanical design to allow for self-stability can be key to overcoming the difficulties arising in control of systems with fast dynamics.

Chapter 8

Stability and behavioural diversity of collective Bernoulli-balls¹

In this chapter I present a novel platform for studying the emergence of individual and collective behaviours in a dynamical system. The platform is based on the so-called Bernoulli-ball, an elegant fluid dynamics phenomena in which spherical objects self-stabilise and hover in an airflow. The process, which actually works for many shaped objects, is driven by fluidic interactions. The vertical dynamics of the ball are governed by a force balance between drag and weight. The horizontal dynamics are commonly attributed to Bernoulli's principle, which describes how pressure at the center of the airflow is lower and hence the ball is drawn in to the airflow. I build on the single Bernoulli-ball system by adding multiple balls to a single airflow, creating a collective Bernoulli-ball system. Adding additional balls to the airflow induces significantly more behavioural diversity since balls can exhibit agent and population behaviours based on their interaction with the environment and other agents.

Collective Bernoulli-ball systems have a number of interesting properties in the context of artificial life and embodied intelligence. The basic setup of the system involves generating an airflow, using a fan for example, and pointing this vertically upward then hovering balls

¹This chapter is based on the following publication, currently under review:

- Howison, T., Crisp, H., Hauser, S. & Iida, F. (2020). On the stability and behavioural diversity of single and collective Bernoulli-balls. [under review as of 01/2022].
- Data availability: <https://github.com/th533/Bernoulli-balls>

Contributions

T. Howison – devised the conceptual idea, carried out experiments, analysed the data, wrote the paper and made the figures.

H. Crisp – aided in carrying out experiments.

S. Hauser – contributed to general discussions and proofreading.

F. Iida – contributed to general discussions and proofreading.

within it. The first key property of the system is the fact that the environmental conditions are controllable and non-linear. The airflow generates a clearly distinct regions in which balls can exhibit hovering behaviours. Since the the properties of the airflow, e.g., speed, width, can vary with height, the dynamics of the interaction between ball and airflow change based on ball height. So, the dynamics governing the behaviour of balls hovering in the airflow are different depending on their location in the environment, with areas of higher and lower stability. Importantly, the environment is also significantly altered by the presence of hovering balls. For example, a hovering ball disrupts the airflow above it, altering the attractor space for other balls in the airflow. Environmental conditions are a key factor in the design and behaviour of artificial agents (Miras et al., 2020). A second key property is the fact we are able to alter the mechanical properties of balls to change how they interact with the airflow and each other. For example, heavier balls will hover at lower heights than lighter balls. Other mechanical properties such as ball size or surface properties can also be altered to change interactions. By altering both morphological and environmental properties we are able to build up the behavioural complexity from a single ball hovering in the airflow to multiple balls with different behaviours. A third key property is the relevance of Bernoulli-ball systems to evolutionary robotics and open ended evolution. One interpretation of the dynamics in the system is that balls are either in an ‘alive’ state, in which energy from the environment is converted into hovering behaviours, or in a ‘dead’ state, in which the ball has fallen out of the airflow. Since hovering balls rely on a structured interaction with the airflow to remain air-born, by adding additional balls we are able to force competition-like behaviour in which the airflow cannot support all balls, and some inevitably fall from the flow. We can interpret this as a the most basic form of evolutionary dynamics, in which the dynamical system self-regulates and removes agents that cannot survive in the given environment, a concept that has been introduced before e.g., (Taylor, 2004).

This chapter is structured as follows. In 8.1 I introduce our experimental setup used to investigate the system. In 8.2 I investigate collective Bernoulli-ball systems with multiple hovering balls. I characterise the basic behaviours different ball configurations exhibit, e.g., stable hovering, interaction or falling. These distinct behaviours within the system offer a natural symbolic representation from which to build probabilistic models of dominant behaviours and behavioural transitions. I investigate the behaviours in isotropic populations, e.g., in which all balls have the same mechanical properties. I demonstrate how morphological and environmental factors influence the nature of behavioural emergence in the system. Next I apply a similar analysis to anisotropic populations, in which the mechanical properties of balls vary across the population. I show how changing morphological properties within the population facilitates further modulation of ball behaviours. In 8.3 I demonstrate how other

more complex properties such as adhesion can induce additional behaviours in the collective system. Finally, in 8.4 I discuss and conclude the chapter.

8.1 Experimental setup

The experiments for collective Bernoulli-balls were conducted by floating balloons in a vertical airflow, as shown in Figure 8.1. The mass of each balloon was altered by affixing small 1 g masses to create arbitrarily heavy or light balls. The airflow is generated by a PWM controllable electric fan with a diameter of 180 mm connected to a constant 12VDC power supply. By varying the control PWM signal we can change the airflow power. The air is passed through a vertical tube of the same diameter and a height of 300 mm. This creates a more steady airflow than is ejected directly from the fan. One or more balloons can be placed in this airflow. A shallow conical catching funnel, constructed of fabric, surrounds the airflow outlet. When balloons fall out of the airflow they are caught by this funnel and are returned via gravity to the airflow outlet. In this way we were able to carry out long experiments with a range of balloon combinations without having to reset the system manually each time a balloon falls from the airflow.

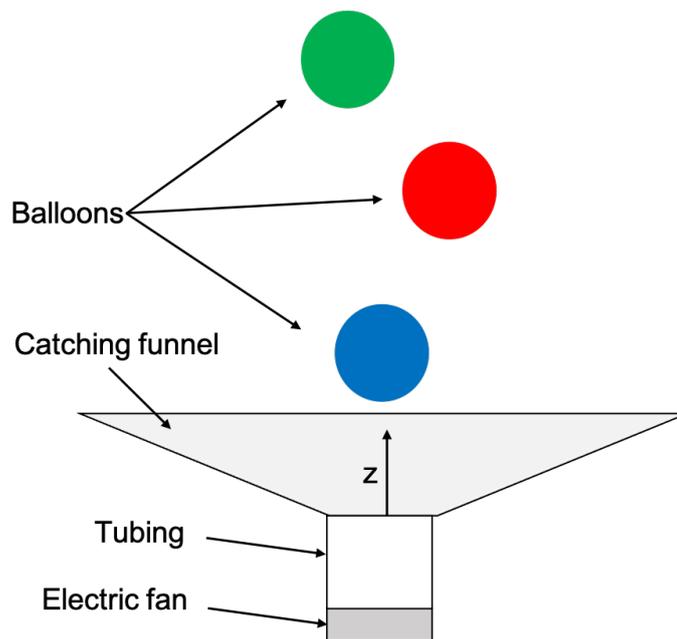


Figure 8.1: Setup for hovering one or more multiple balls in collective Bernoulli-ball system. An electric fan generates a vertical airflow with a controllable power. balls are added to the airflow. When balls fall from the airflow they are caught in the funnel and returned to the airflow base.

Two Logitech Brio webcams were set up to capture the ball trajectories. The position of each ball in pixel coordinates for each camera view was found using a colour thresholding algorithm. We had full control over the ball colour and room lighting, meaning we can in theory track large numbers simultaneously. In some cases balls were partially or fully occluded. If the detected colour region was smaller than a predefined threshold the detection was discarded. Missing trajectories were predicted using Gaussian smoothing and state estimation across the incomplete trajectory, using the MATLAB `fillmissing` function. The two complete pixel trajectories were resolved into a single 3D real-world trajectory using the MATLAB stereo camera tool. The final output of trajectory processing was a full 3D trajectory, with units of mm, in a coordinate frame centered at the airflow outlet.

8.2 Behaviours of collective Bernoulli-balls

Our experiments with a single Bernoulli-ball (Chapter 7) demonstrate the most basic dynamics in the Bernoulli-ball system, and show how more complex behaviours can be induced via environmental manipulation. To increase behavioural complexity we can add additional objects to the airflow, creating a collective Bernoulli-ball system. Hovering multiple ping-pong balls is challenging to stabilise and sensitive to slight environmental changes. Hence, for collective Bernoulli-ball investigations we utilised our experimental setup for hovering balloons (Figure 8.1), which are far lighter, and hence require a less powerful airflow. We conducted a large number of experiments hovering multiple balls in a high and low power airflow. Video examples of these experiments can be seen in the supplementary video². In this section I describe the behavioural diversity observed in these experiments.

8.2.1 Behavioural primitives

The behaviours of balls in the collective-system can be categorised into distinct classes, for example interacting with another ball or falling from the flow. This natural symbolic representation generated by multi-behaviour systems allows us to reduce complex dynamics to a level that captures the diversity in the system but does not need a full analysis of the underlying dynamics, e.g., as shown in (Horibe et al., 2011).

We defined a set of behavioural primitives for the collective Bernoulli-ball system. By observing the system we found that the behaviour of a given ball can be characterised in terms of four basic modalities:

1. **Stable:** a ball exists in a state of *apparent* stability in which it is not mechanically interacting with any other ball in the system.

²<https://youtu.be/lz5HbMwPap8>

2. Interaction 1: a ball is mechanically interacting with *one* other ball, e.g., bouncing off each other.
3. Interaction 2: a ball is mechanically interacting with *one* other ball, which is itself interacting with another ball, e.g., there are three balls involved in simultaneous interaction.
4. Falling: a ball falls out of the airflow.

Figure 8.2 shows examples of these behavioural modalities. Since behaviours are defined at the agent level of individual balls, the system can exhibit multiple behaviours simultaneously. For example, the system can be in a fully stable configuration, a fully interacting configuration or a configuration with two balls interacting and one stable.

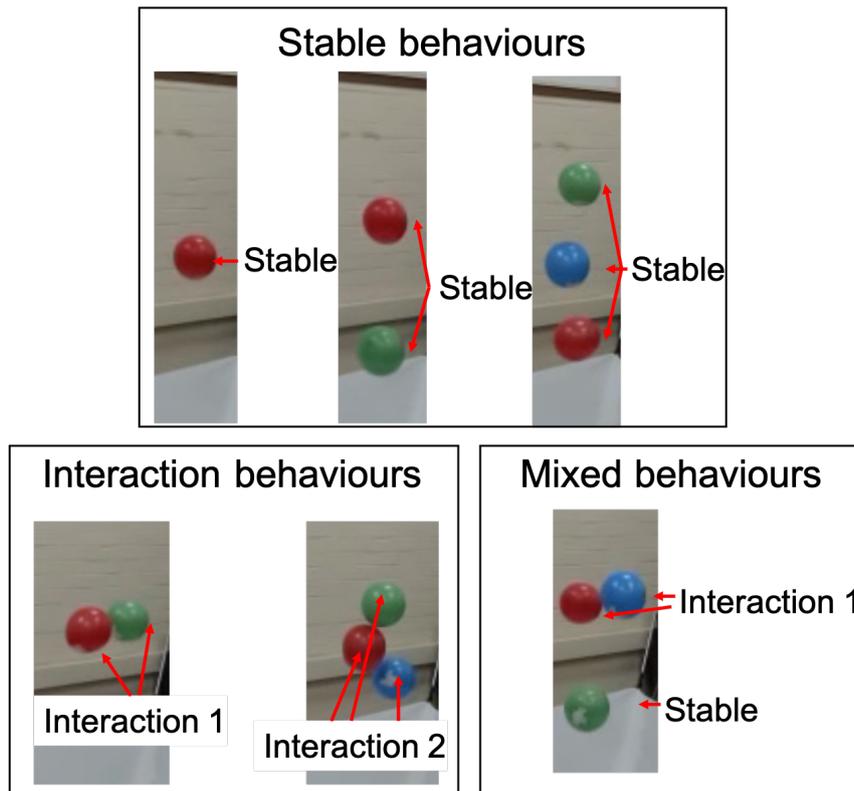


Figure 8.2: Behavioural primitives in one, two and three ball collective Bernoulli-ball systems. Behaviours are defined at the agent level, e.g., each ball as a behaviour. Hence, the collective system can exhibit more than one behaviour (see mixed behaviour panel).

8.2.2 Automatic behavioural classification

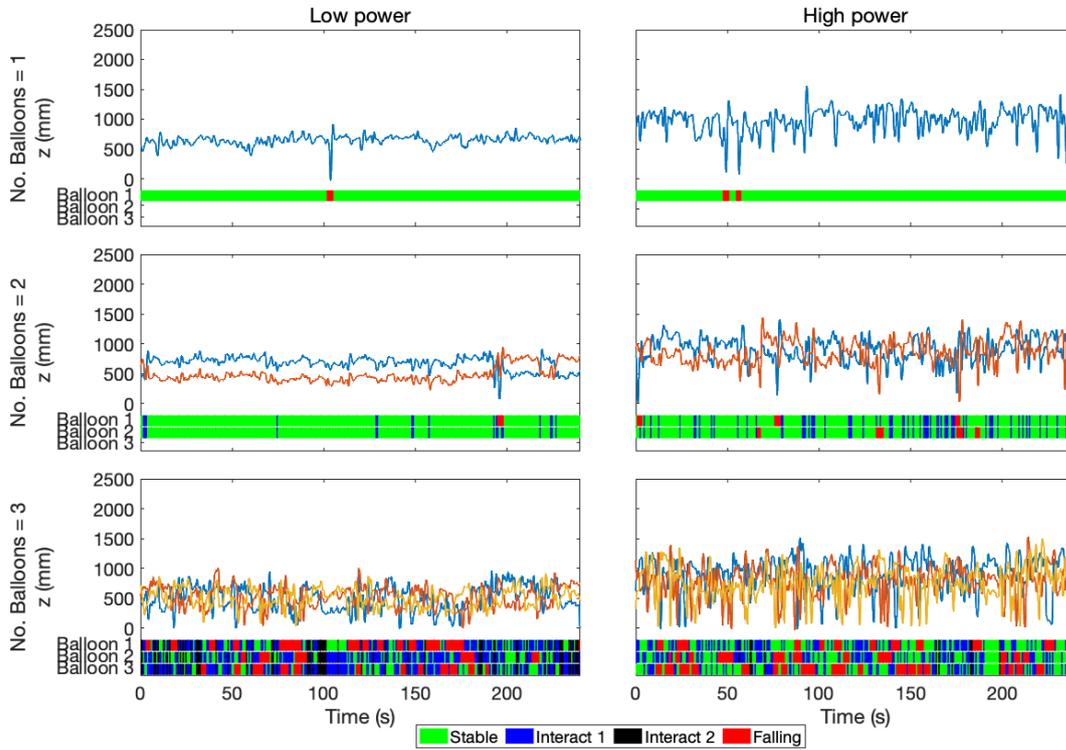
Observer based characterisations are useful to quickly capture the basic modalities of behaviour in the system in an intuitive framework. Systematic analysis and behaviour

categorisation of the data requires more formal behaviour definitions. The 3D trajectory data allows us to achieve this automatically determine the system behaviours at each time-step using the following set of rules:

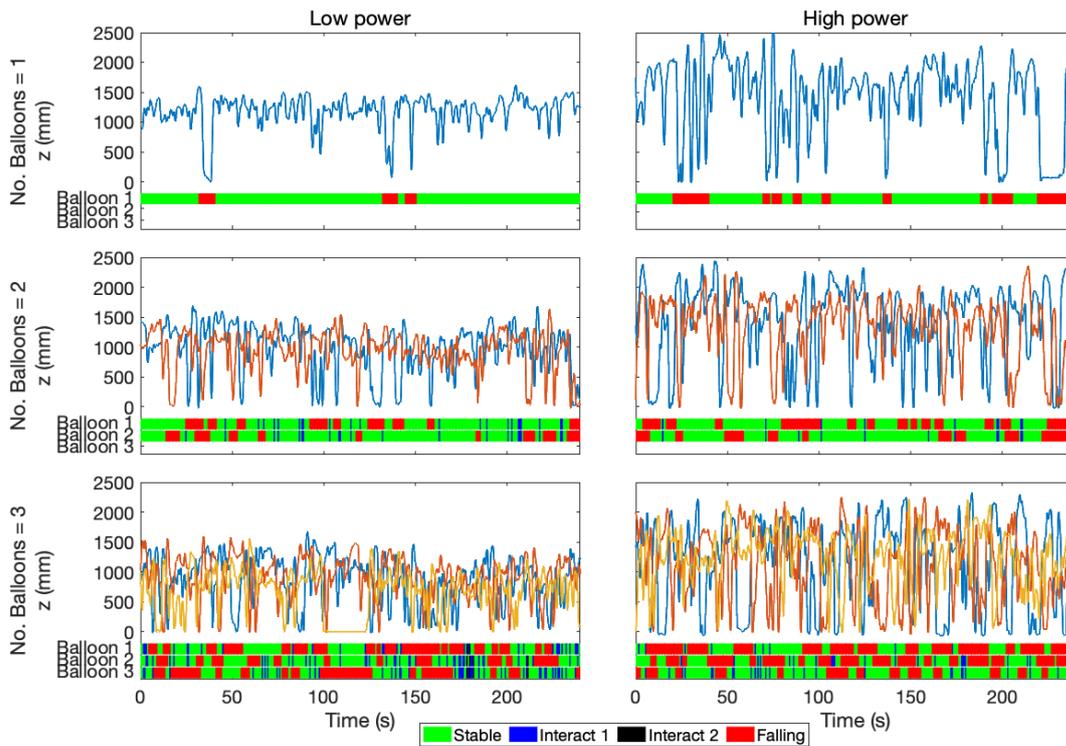
1. Interacting behaviours can be detected by comparing the linear distance between each trajectory in 3D space. If this distance is below 110% of the ball diameters then interaction is defined. If two interacting events are detected within a sliding one second window then interaction is defined for this full time period. This window was chosen as the authors found it to maximise the similarity between the human-observed behaviours and the automatically detected behaviours. In doing so we capture interacting behaviours of balls bouncing off one-another over a prolonged period.
2. A ball falling is detected when the a ball comes within close proximity of the catching funnel. The falling behaviour is defined as the point at which the ball began falling vertically downwards out of the flow, to the point at which the ball has fallen and the re-entered the airflow.
3. Stable behaviour is assigned to all cases where a ball is not interacting or falling.

8.2.3 Behaviours in isotropic populations

We first investigated behaviours of isotropic ball populations, in which all balls are the same size and mass. We tested populations of one, two and three balls for five minutes in a low and high power airflow. The behaviours over the full five minute trajectory were classified. These classifications were used to define the dominant behaviour for each ball in each one second increment of the experiment. We tested two types of population, one comprised of heavier balls weighing 6 g and one with lighter balls weighing 3 g. Figure 8.3a shows a two minute trajectory snapshot of height (z) for populations of one, two and three heavy balls in a low and high powered airflow. Figure 8.3b shows the same data for lighter balls. Figure 8.4 shows a closeup of mean-adjusted trajectory for two heavy balls in a low power airflow, showing the correlation in height. Figure 8.5a and b shows the proportion of time balls in these experiments spent in each behaviour during the full experiment of 5 minutes. The bar charts show the proportions averaged over all balls in the population.



(a) Heavy 6 g ball populations.



(b) Light 3 g ball populations.

Figure 8.3: Time series data and automatically classified behaviours for a two minute snapshot of experimental data, showing ball populations of one, two and three, in a low and high powered airflow, for (a) heavy balls weight 6 g and (b) light balls weighing 3 g.

The time series trajectories and behavioural classifications in Figure 8.3 tell us a number of things. Increasing the number of balls in the system increases the variety in behaviour, and decreases the time spent in a stable behaviour. We expect this since when adding balls we do not change the environment, so balls have to compete for the same airflow to stay in the flow. Heavier balls occupy a smaller vertical range than lighter balls, and generally hover lower than lighter balls. The simple balance between drag force and ball mass explains this behaviour. There is a clear coupling between the heights z of balls in some configurations. For example, a population of two 6 g balls in a low power airflow (Figure 8.3a middle left panel) show a clear correlation between the two ball heights, which move up and down with each other. Importantly we see that the balls are nearly always in a stable behaviour, so this coupling is transmitted through the environment, not by mechanical interaction. Figure 8.4 shows the mean corrected heights of both balls in this configuration for a two minute segment, demonstrating the tight coupling via the environment. The same population in a high power airflow shows a similar, but less obvious, correlation. Lighter balls are not obviously correlated in the same way. We also see that when some balls fall out they can spend a significant time not in the airflow before re-entering, for example 3 g balls in a low power airflow $t=100\text{--}125$ s. During this time balls are in the catching funnel.

Figure 8.5 shows us the average proportion of the five minute experiment balls spent in each behaviour. The figure reinforces our observations from the time series data. In all experiments we see that increasing the fan power, i.e., the energy available in the environment, increases the variance and behavioural diversity of balls. The number of interactions and falls seems to increase, while the time spent in a stable configuration decreases. Heavier balls also appear to be more stable in their interactions, generally spending far longer in the interacting 1 and interacting 2 behaviour modes than the lighter 3 g balls.

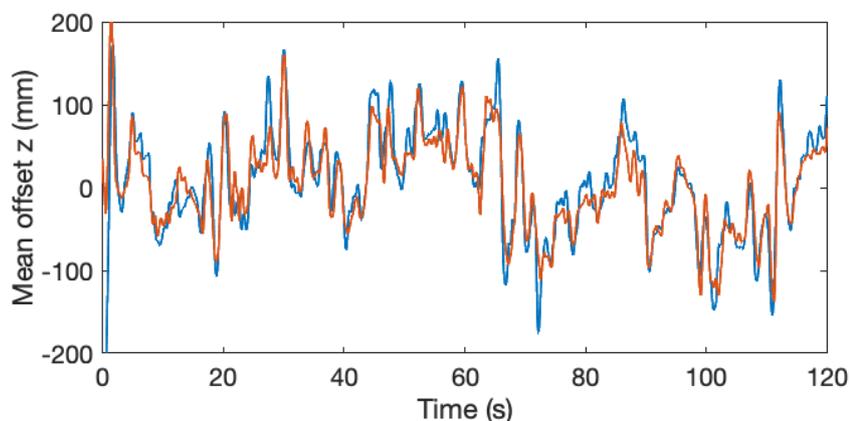


Figure 8.4: Mean corrected height for two heavy balls weighing 6 g in a low power airflow, showing similarity in trajectories through environmental coupling.

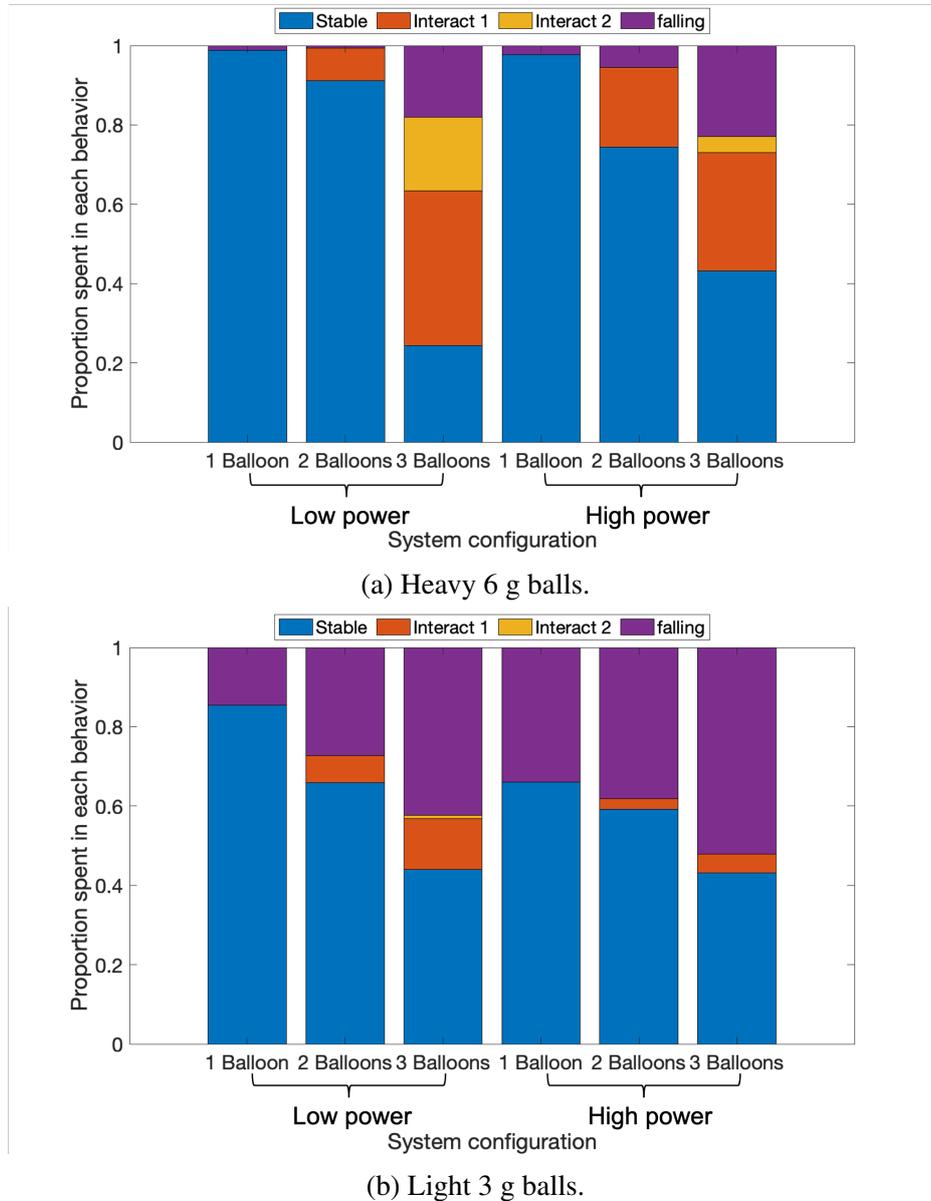


Figure 8.5: Average proportion of time spent in different behavioural modes across different population sizes and airflows for (a) heavy balls weight 6 g and (b) light balls weighing 3 g.

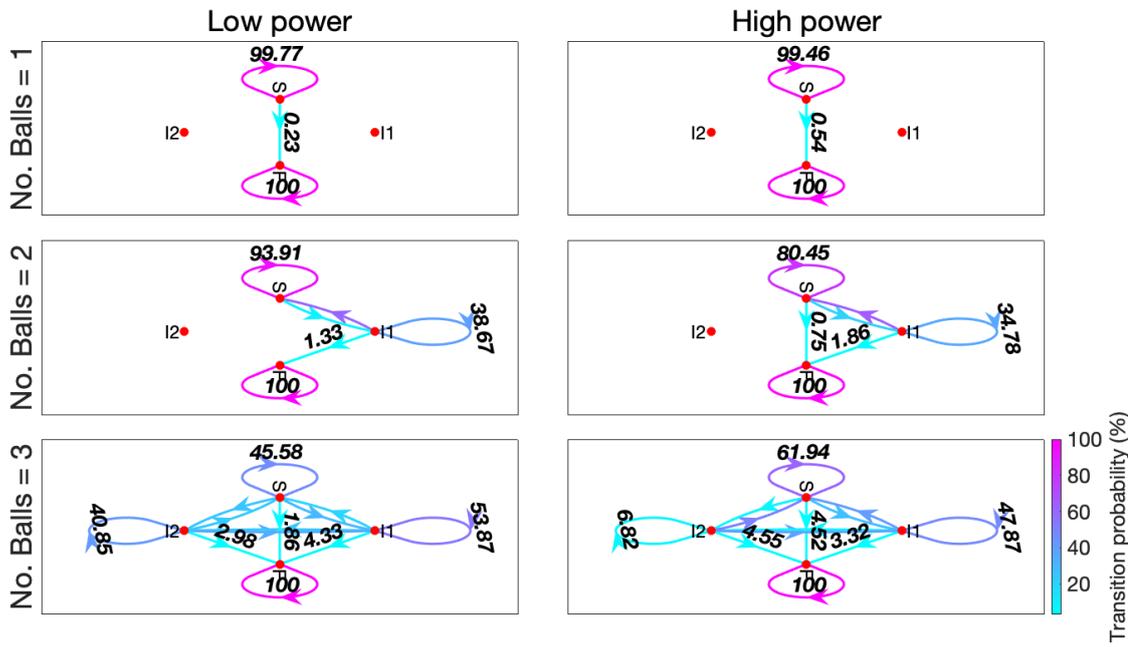
8.2.4 Stochastic behaviour transitions

The basic analysis in Figure 8.3 and 8.5 demonstrates likely behaviours, but cannot tell us about behavioural transitions, another important factor of the system. The discrete representation of hovering behaviours allows us to analyse these transitions. One way of quantifying how behaviours transition in the system is to represent system states in terms of Markov chains, a mathematical system in which state transitions occur based on probabilistic rules (Meyn and Tweedie, 2012). A discrete-time Markov chain (DTMC) represents the

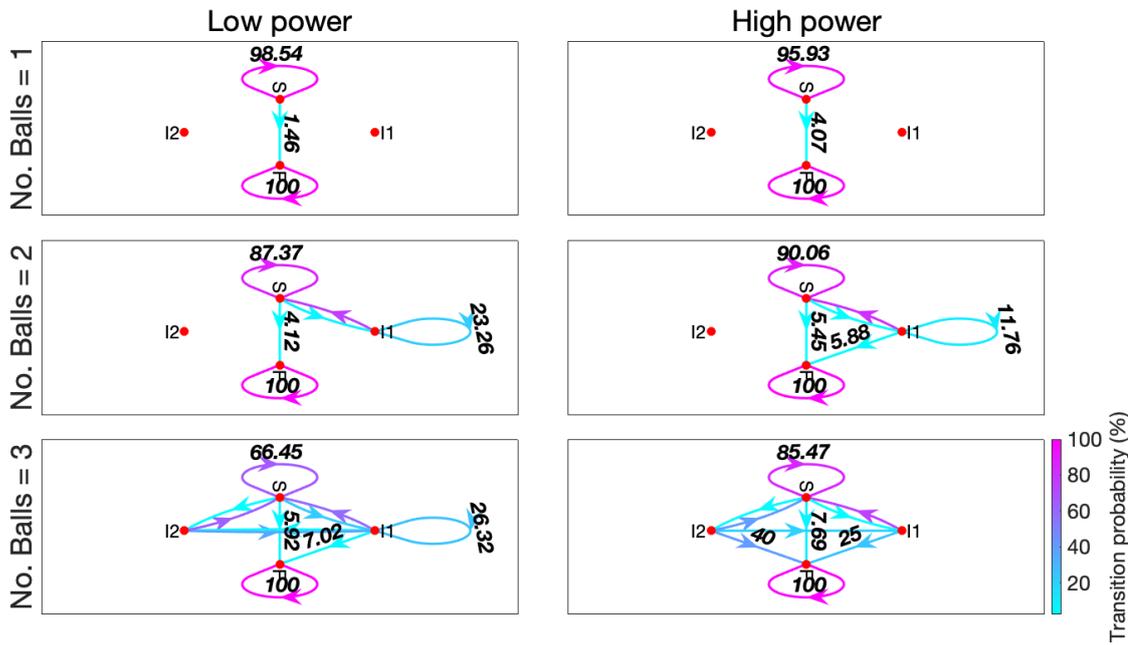
probabilistic switching landscape for moving between states at a given sampling frequency, e.g., every second. This is similar to the mode switching representation shown in (Horibe et al., 2011). Using our segmented trajectory data we are able to record every behaviour transition for each ball in the experimental data. From this we can construct a transition matrix that indicates that probability of different behaviour switches occurring, e.g., stable to interacting 1, interacting 1 to falling.

Figure 8.6 shows a graph representation of the DTMC for each experiment shown in Figure 8.3. The nodes in the graph represent the four possible behaviours a ball can exhibit: S = stable, F = falling, I1 = interacting 1, I2 = interacting 2. The edges between nodes represent the probability of a behaviour transition occurring after the next 1 second increment, with colours showing the likelihood value. The likelihood values have been annotated to show only the likelihood of a ball remaining in the same behaviour, or the ball moving to the falling state (otherwise, the diagram becomes hard to read). The absence of a connection between any two nodes indicates no transition was observed in our experiments, and hence that transition has a likelihood of zero. In our experiments the catching funnel was used to reset the system when balls fell out. In these diagrams we treat falling as an absorbing state from which a system cannot leave, as would be the case in the experimental setup without a catching funnel. For example, when a ball in a three ball population enters the falling state, the remaining two are defined by the behavioural transition diagram for a two ball system. Hence, falling events permanently change the behaviour dynamics for balls remaining in the airflow.

The figure demonstrates some key features of how behaviours change in the collective Bernoulli-ball system. For both a low and high power airflow, increasing the number of balls in the population decreases the likelihood of remaining in a stable behaviour after one second, e.g., 99.77% to 45.58% when moving from one to three heavy 6 g balls in a low power airflow. We expect this since the more balls in the population, the more likely interaction and falling becomes because they are all occupying the same airflow. In general, heavier balls are likely to interact for longer than the light balls. In three ball populations in high power airflow, for example, heavier balls have a 47.87% likelihood of remaining in the interacting 1 behaviour after 1 second, whereas lighter balls have a 0% chance. An obvious conclusion is that the number of available states and transitions increases as the population size increases. A single ball can only exhibit stable or falling, whereas a three ball population exhibits four states and seven possible transitions.



(a) Heavy 6 g balls.



(b) Light 3 g balls.

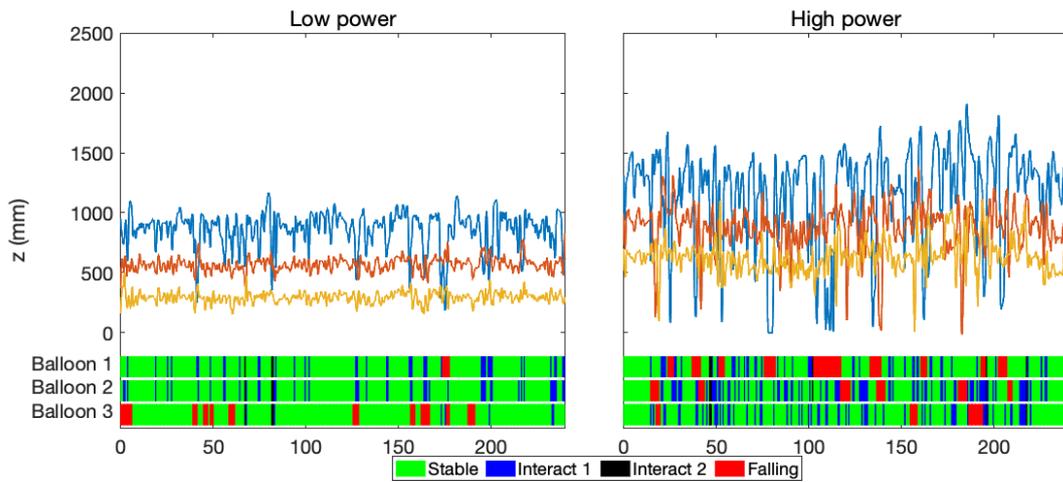
Figure 8.6: DTMC behavioural switching diagrams for ball populations of one, two and three balls, in low and high power airflows, for (a) heavy balls weighing 6 g and (b) light balls weighing 3 g.

8.2.5 Behaviours in anisotropic populations

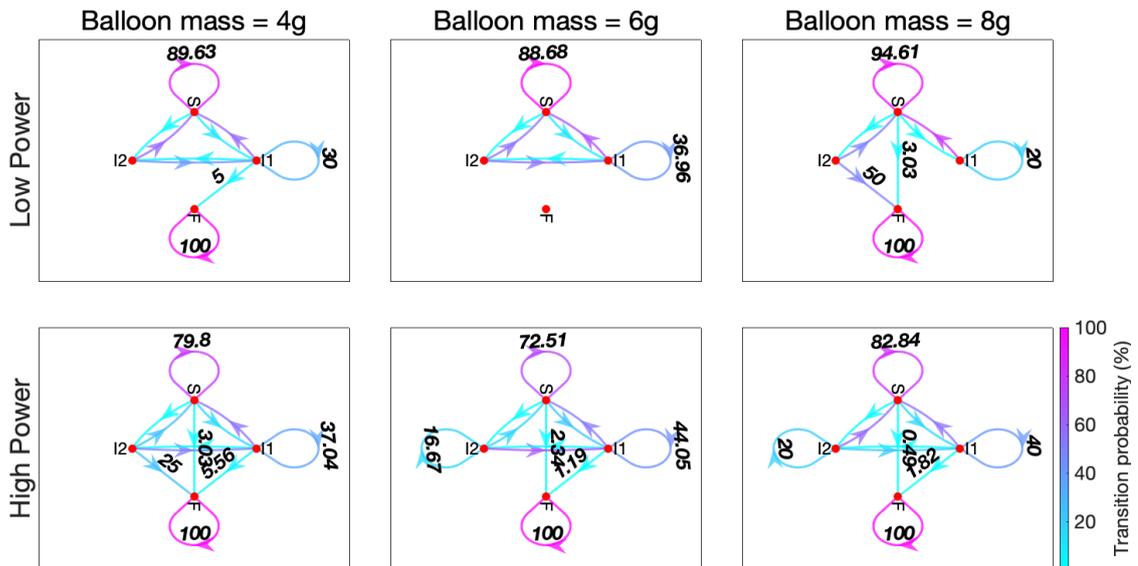
In isotropic populations all balls have the same mechanical properties, and so are governed by the same underlying dynamics. Hence, our primary method to induce behaviour changes is by changing the environment. To introduce more behavioural variation across the population we can also alter the mechanical properties of balls individually to produce an anisotropic population of balls with different properties. We created an anisotropic population of three balls with weighing 4 g, 6 g and 8 g and recorded their behaviours in the same low and high power airflow. The total mass of balls in the system was 18 g, the same as our experiments with three equally weighted 6 g balls (Figure 8.3a). We recorded behaviours for five minutes as with our isotropic population experiments.

Figure 8.7a shows a four minute snapshot of trajectories and segmented behaviours for these experiments, with the associated behaviour classifications. The effect of a distributed mass profile across the population is clear, with a distinct vertical sorting based on the heaviest ball (8 g, yellow line) at the bottom and lightest (4 g, blue line) at the top. This is true for low and high power airflows, but is far more distinct in low power airflows. The spatial structure for both airflows is more distinct than the isotropic population of the same weight (Figure 8.3a) in which there is no vertical sorting. The main similarity between the isotropic and anisotropic populations is the relationship between airflow power and stability. We see that the system is more stable in the low power airflow, showing less interacting and falling, and varying less in height. The time series trajectory show us that, as with our other experiments, increasing the airflow power decreases overall stability. We see that heavier balls are generally more stable. The heaviest balls exhibited fewer interactions and falling behaviours for lower and higher power airflows.

We again constructed DTMC diagrams to represent the stochastic behaviour switches, as shown in Figure 8.7b. In the case of anisotropic populations we constructed a switching diagram for each different ball in the system, since the varying mass changes the intrinsic dynamics, and hence behaviours, that a ball exhibits. The diagrams demonstrate this difference. For the low and high power airflow the heaviest ball is more likely than the lightest ball to remain in a stable behaviour, 94.61% vs 89.63% and 82.84% vs 79.80% for low and high power respectively. For both airflows, the middle ball is the least likely to remain in a stable behaviour, with a probability of 88.68% and 72.51% for low and high power. This is largely because it is frequently interacting with the balls above and below it. For the low power airflow, the middle ball is the least likely to fall, exhibiting a 0% chance of falling in our experiments. Hence, it remains in the airflow through frequent switching between interaction and stable behaviour.



(a) Time series and behaviours for low and high power airflows.



(b) DTMC for individual balls, with different masses, in high and low power airflow.

Figure 8.7: Results for an anisotropic population of balls weighing 4 g (blue line, balloon 1), 6 g (red line, balloon 2) and 8 g (yellow line, balloon 3), showing (a) Time series trajectories and classified behaviours and (b) behaviour switching diagrams.

8.3 Altering other mechanical properties

We can also alter other mechanical and environmental properties to induce more complex behaviours. To investigate possible ways to extend behavioural diversity in the system we ran a set of experiments based around the idea of balls adhering together. The collective nature of

the system naturally suggests this possibility of self-assembling structures, and the inherent interactions that occur within the collective population act as the mechanism on which to allow this process. The setup (Figure 8.8) consists of two vertical fans and four balls. balls 2–4 are standard balls with a mass of 3 g. ball 1 (green, in figure) has an additional layer of adhesive tape wrapped around its circumference. Under the right conditions balls can adhere to this.

Figure 8.9 shows the time series of the vertical z and horizontal x ball positions (as measured from origin shown in Figure 8.8), divided into three phases. In phase 1 of our experiment the system is in a state of dynamic stability, i.e. the balls do not fall out of the airflow but are constantly interacting. At $t = 60$ s balls 1 and 2 (green and purple) adhere together after mechanical interaction. When adhered, their vertical height increases since the effective area under them increases and induces a greater drag force from the two fans. Phase 2 lasts for 20s, after which another mechanical interaction induces adhesion between ball 1 and 3 (green and yellow) and the system enters phase 3. At this stage the interaction within the system has induced the self-assembly of a three-ball structure. This structure has a larger surface area and hovers high up in the airflow. The single remaining ball hovers lower in the airflow, and hence does not interact and adhere to the larger structure.

Unlike in our previous examples, where long-term stability can only be achieved by balls falling out of the airflow, in this experiment long-term stability is achieved via ball adhesion. Via environmental and mechanical interaction the balls adhere together in a formation that allows them to remain stable in the airflow. If the adhesion was not stable, the structure would immediately fall from the airflow. One can conceive that the format of the adhered structure is also influenced by the type of environment. For example if there were many airflows, adhesion may be rare since balls could hover further from each other. Conversely, with only airflow adhesion may happen far quicker. In general, this extension demonstrates the scalability of the collective Bernoulli-ball. We can add arbitrarily many airflows to the environment. In parallel, we can add any number of balls. Finally, we can alter the properties of these balls and airflows to induce different behaviours. We can also frame this adhesion behaviour from the concept of *ultrastability* discussed in Chapter 1 (Ashby, 1961), which describes the behaviour of a system to change its internal structure in response to environmental stimuli. Single Bernoulli-balls are *ultrastable* insofar that they remain stable within an airflow subject to random perturbations (up to a limit). The adhesion experiment shows how they can be *ultrastable* at the population level, since over time the system re-configures to reach a steady-state behaviour. The collective Bernoulli-ball experiments without adhesion are not stable in this way, since in most configurations the stability is reached by ejecting one or more balls from the flow.

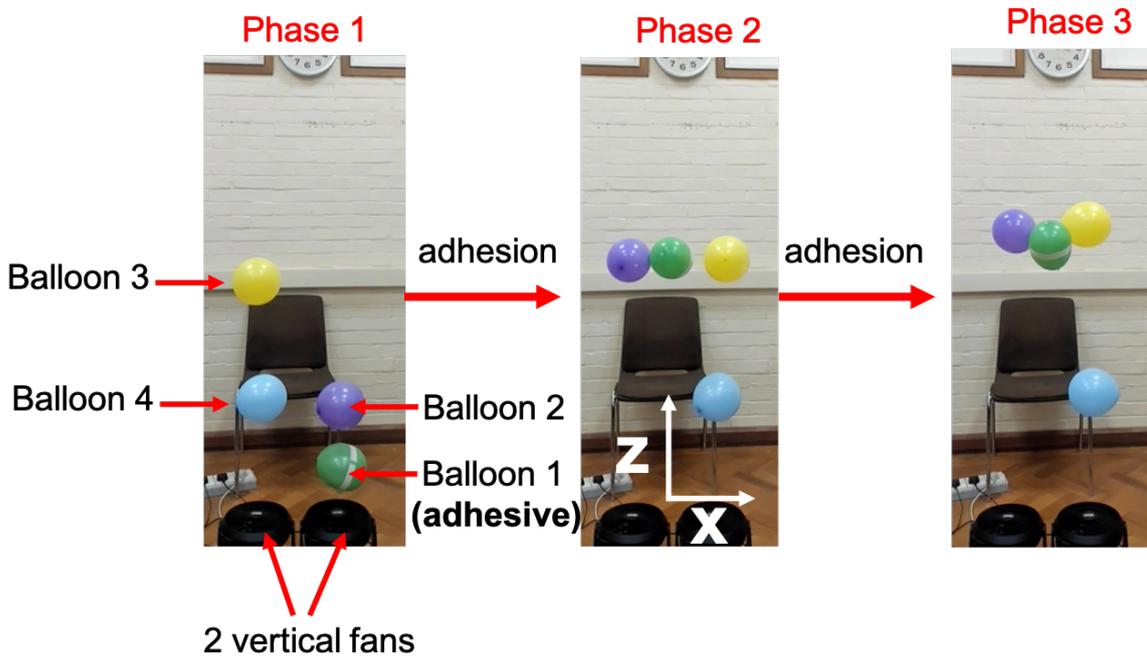


Figure 8.8: Extending behaviours by increasing morphology and environmental complexity using two fans and an adhesive balloon. Images show snapshots of different adhesion phases. Middle panel shows x and z measurement axes.

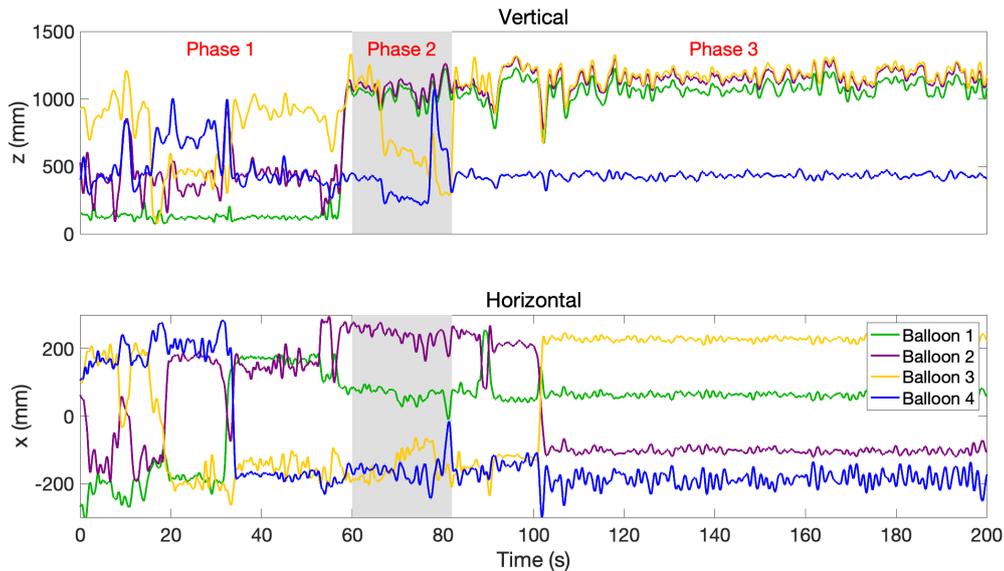


Figure 8.9: Vertical z and horizontal x time series position of balls 1–4 in adhesion experiments.

8.4 Discussion and conclusion

In this chapter I introduced the single and collective Bernoulli-ball systems as a platform on which to study types of behavioural diversity in dynamical systems. The key contribution of the chapter is demonstrating the various ways in which behavioural diversity can be triggered in the system, and how this can be analysed. The different behaviours observed, such as stable, interacting, and falling, are all the result of the complex interaction between agents and a complex environment.

As with much of the work in this thesis, a key feature of the analysis was based around defining and analysing behavioural modalities. Unlike with the falling paper research (Chapters 3-6), behavioural classification in the case of the collective Bernoulli-ball system is far more subjective, and there are many ways of interpreting and labelling behaviours. The rationale behind selecting behavioural modes was to choose a representation that was easily observable at the agent-level, and that could be used to describe population-level behaviours. For example, interacting and stable behaviours between two balloons can be used to describe the population-level behavioural switching in a group of balloons. Other schemes were of course possible, and an interesting avenue of further work would be understanding how to choose the best behavioural representation for a given system.

The results in this chapter indicate a number of key features that make the system useful in the context of self-structuring embodied agents. First is the idea of stochastic cumulative self-organisation. We can describe the system using a set of primitive behaviours as described in Figure 8.2, and this forms an intrinsic symbolic notation by which to describe the system. The emergence of these behaviour states is stochastic insofar that we cannot sense all the states required to predict it. By varying morphological or environmental parameters we can induce different stochastic landscapes. In our experiments we changed the mass of the ball and the power of the airflow. Even with this simple 2DOF design parameter space we can significantly modulate the nature of behavioural emergence in the system. By representing stochastic behaviour transitions in a DTMC (Figure 8.6) we see the likelihood of different behaviour transitions. In doing so we are able to define higher order behaviours from our set of primitives. For example, sequences of stable and interacting behaviour switches are a type of oscillatory behaviour. Similarly, stable configurations followed by interaction followed by stable configurations in a different order are switching behaviours. Hence, the self-organising behaviours are cumulative. Depending on the time-horizon this cumulative property resolves as different higher order behaviours.

Second is the idea of evolutionary dynamics in a dynamical system, an underexplored area in the context of evolving robotic lifeforms (Pfeifer et al., 2005; Taylor, 2004). The system has natural ‘alive’ (ball hovering in airflow) and ‘dead’ (ball fallen out of airflow)

states, which are based on whether or not a ball is utilising the structured environmental landscape to turn convert airflow energy into hovering behaviours. We can view this from the perspective of evolutionary robotics and open ended evolution. The system is initialised with a certain number of balls. If this configuration is stable then agents will remain in the airflow. This stability could be in the form of stable fixed points or a dynamic, interacting equilibrium. If the system is unstable then the system exists in a state of instability until it reaches a critical level, at which time an agent falls from the system (or adheres to another ball). At this stage the collective behaviours reconfigure based on the dynamics of the remaining agents in the airflow. This process repeats until the system reaches a sustainable stable state. The system has a form of intrinsic mortality such that unstable populations cannot survive without ejecting individual agents. Furthermore, we are able to pressure this intrinsic behaviour by structuring the environment and population such that interaction and competition over stable airflow is inevitable

Finally, the system expresses rich complexity in its dynamics while being cheap and easy to experiment with. One of the key features of many complex dynamical systems is the difficulty in capturing their full dynamics in a simulation environment. The Bernoulli-ball system is governed by complex fluid dynamic interactions between the airflow and floating objects. However, all one needs to investigate the system is a minimal set of components, e.g., an electric fan and some balls. More complex setups are of course beneficial, but the fundamental dynamics of the system can be explored with just the simple setup.

Ultimately this system differs from other similar such dynamical systems, e.g., the blind juggler (Reist and D'Andrea, 2012), moving oil droplets (Horibe et al., 2011), passive dynamic walker (McGeer et al., 1990) or floating water elements (Nakajima et al., 2012) due to the scalability and controllability of both morphology and environment. We can easily create arbitrarily complex configurations in which behaviours are governed by the same basic underlying principles that drive the emergence of complex, higher-order behaviour and evolution over multiple time-scales.

Part III

Discussion and conclusion

Chapter 9

Discussion and conclusion

I have presented a body of work relating to the emergence of non-trivial behaviours from low-level embodied physical systems. The framework of iterative and directed exploration of self-structuring embodied agents was used to structure the research. The contributions in this thesis represent a small step toward the goal of creating increasingly complex artificial life-forms. Paradoxically however, the systems studied in this thesis are apparently simplistic. The rationale behind studying these simple systems is that, by investigating how low-level physical interactions give rise to complex behaviours, we can build an understanding of embodied interactions from the bottom up. In this chapter I summarise the key contributions of the thesis and describe any related and future work.

9.1 Discussion of contributions

In this section I discuss the main contributions in terms of the research objectives set out in 1.3.

9.1.1 Thesis objective one: systematic exploration of dynamics in novel self-structuring embodied agents

In Part I I investigated Falling paper systems in the context of self-structuring embodied agents. The VSFP system is a novel addition to this class of system. In Chapter 3, I introduced the VSFP system, and with systematic exploration of the design space showed how morphology dictates dominant behaviours. The population level perspective of the morphology-behaviour mapping offers a compact snapshot of possible behaviours in the system, and what influences them. The key property of these systems is that by varying a small number of morphological parameters, a wide range of complex behaviours can be modulated. As discussed in Chapter 4, some cases the mapping from morphology to behaviour is deterministic, while in other cases it is stochastic. Again, the likelihood of

transitions between different behavioural modes can be influenced with a minimal set of morphological parameters. In Chapter 5, I demonstrated how morphological parameters can be optimally selected based on a fitness function, in this case minimising falling speed. The use of Bayesian optimisation for this task is not a novel innovation. However, demonstrating its use on a system with highly distinct behavioural modes is important since the fitness landscape in this case is non-smooth. Finally, I investigated more general falling paper systems in Chapter 6. More complex morphologies such as hexagons or crosses were investigated and their behaviours categorised. Ultimately, the primary focus of research into falling paper relates to the relationship between morphology and behaviour. Although other studies in fluid dynamics have included environmental factors, e.g., dropping metal plates in oil (Field et al., 1997), this thesis dealt with the static environment of quiescent air.

In Part II I investigated Bernoulli-ball systems. Unlike with falling paper, in this class of system we have far more control over the environment than we do the morphology. In Chapter 7, I explored the case of a single hovering ping-pong ball. This setup used a perfectly spherical ball to hover, and an airflow designed to be as smooth and laminar as possible. Hence, the morphology-environment interactions were as stable as possible, facilitating reduced order modelling and height control. In Chapter 8, on the other hand, I investigated an more unstable configuration. Multiple hovering balloons were placed in an airflow. The balloons are not perfectly spherical and it is challenging to control their exact size and properties. Similarly, the airflow was generated directly from a fan, without any flow smoothing, making the environmental conditions in which the balloons hovered more unstable. The range influence of environmental conditions (e.g., power and number of airflows) and morphological parameters (e.g., ball weight) were systematically investigated and shown to influence the behaviours and likelihood of transitions in the system.

9.1.2 Thesis objective two: development of methods for iterative and directed exploration

This thesis developed two main methods for iterative and directed exploration. I presented the PDBC algorithm in Chapter 3 as a tool for design landscape representation. The results indicated that the PDBC method is an effective way of finding a parameter space in which behaviours can be clustered together with a high degree of accuracy. The PDBC algorithm was designed with the aim providing physical insightful behavioural classification for behaviourally diverse systems. There are many applications in which this is desirable. Automated design optimisation, for example, often focuses on hard to model problems such as the real-world evolution of locomotion (Brodbeck et al., 2015). Behaviourally diverse systems could multiple solutions to such problems. PDBC could be used in conjunction with

quality diversity algorithms such as MAP-Elites (Mouret and Clune, 2015) to optimise such systems and provide a physically insightful snapshot of the solution landscape.

However, there are also systems for which the current PDBC algorithm is inapplicable. Clearly, systems with no clear behavioural diversity are ruled out. More subtly, however, are systems with non-discrete behavioural modes. Here, there may be a clear range of behaviours separated by a continuous transitional zone, in which one behaviour blends into the next; period doubling, for example. Within this transitional zone behavioural classification is ambiguous, making the data acquisition step of PDBC challenging. One approach may be to restrict sampling to areas within the parameter space with very clear behaviours.

In Chapter 6, I presented a platform for large-scale physical experimentation on general falling paper systems. Hundreds of experiments were carried out on a range of paper shapes, with the aim of clarifying the boundary between the various system behaviours. When compared to previous studies (Field et al., 1997), the flexibility in morphology, experimental quantity and analysis speed were vastly increased, with experiments and analysis taking an average of just 90 seconds. One significant way in which robots can contribute to the discovery, exploration and understanding of scientific principles is through large scale automation of highly stochastic experiments, where many data points are needed to uncover otherwise unobservable trends. Additionally, robotic systems offer a consistently high level of accuracy, precision and repeatability in the experimental procedure. Despite this, the application of robotics to exploring physical systems has been limited, especially in the fields of fluid dynamics (Fan et al., 2019) and complex systems. This can be attributed to the conceptual and practical challenges of automation, as well as the prohibitively high cost of some robotic technologies like robotic arms. Nevertheless, the stochastic nature of physical systems represents exactly the type of problem in which robotic automation can be advantageous.

Despite the ability of the system to continuously fabricate, test and analyse falling shapes, it was limited to the Re and I^* values reachable by dropping paper in air. Most notably, in this study only a single weight of paper was considered whereas previous studies have used a range of materials and fluids to expand the search space. Testing alternative materials would allow a greater range of the search space to be considered, and it could be confirmed that the observed patterns hold over these material ranges. Investigating much larger shapes, which would require a stiffer material, could provide interesting insights into the physical scalability of the observed behaviours. Another limitation is the current drop height limitation of 1.1 m. Extending this would allow observation of longer trajectories or potential behavioural switching. An improved exploration strategy could, for example, select morphological

parameters around possible behavioural boundaries or in under-explored areas of the $Re-I^*$ parameter space.

9.2 Outlook and future work

A number of possible directions could be followed to extend the research framework proposed in this thesis.

In Chapter 2, I introduced the framework of reality-assisted evolution for designing embodied soft robots. A continuation of the trend toward a more unified approach to designing soft robots is required. This will be partially driven by the further development of fabrication, material, sensing and actuation technologies. As these become better understood, so will our ability to include them within models and deploy them via large-scale physical experimentation. In parallel, developments in data-driven modelling and transferability approaches will improve the utility of simulation tools, facilitating better predictions of robot behaviours and, crucially, more informed studies into optimal design methods. In the long term, research on *open ended evolution*, e.g., (Huizinga et al., 2018; Lehman and Miikkulainen, 2015), that aims to indefinitely evolve systems without converging to a particular solution could be significant. Developments in these search and optimisation methods are key for structuring the design process toward discovering diverse behaviours. Finally, moving beyond soft robotics I expect to see a trend toward more complex hybrid systems. Biological systems tend to be neither soft nor rigid, so understanding how the combination of different materials improves embodied functionality is a key milestone toward designing complex and intelligent machines that manifest behaviours driven by seemingly conscious action.

In Chapter 3, I presented the PDBC method for generating equations that enable behavioural clustering. For systems which do exhibit discrete behavioural modes, this approach opens up new avenues of analysis and understanding. However, further work is required to apply the method to systems with ambiguous or continuous behavioural phases. Additionally, further work is required in the choice of system parameterisation, output selection and behavioural interpretation. One of the main issues here is the human interpretation of system behaviours. Although relatively clear in the VSFP system, more complex system may exhibit a range of similar behaviours which are hard to distinguish. Hence, there is scope to automate the process of deciding what constitutes a discrete behavioural mode. To fully realise the impact of this approach a more generalised method of approaching and achieving morphological range is required so it is not prescribed or influenced by initial human bias.

In Chapter 4, I investigated the nature of behavioural transitions in the VSFP system, showing a range of empirical results and different representations of the system behaviours.

Future work on this system includes investigating how behaviour switches can be actively induced in falling experiments, either by environmental manipulation or simplistic on-board control. Similarly, larger scale experiments could be carried out to determine transitional probabilities for specific morphologies. In parallel, similar work needs to be extended to other falling paper systems to understand the extent to which their behaviours and environmental interactions can constitute computational processes, e.g., sequential logic.

In Chapter 6, I demonstrated the utility of large-scale physical experimentation on falling paper systems. There is scope to investigate far more complex shapes and the role of deformation or even folding during falling. The large-scale system could also be used to optimise morphologies based on certain fitness functions, providing a practical and applied usage in design optimisation. Indeed, the falling-paper problem is an ideal platform to test optimisation strategies for noisy, real-world functions. Finally, further work could be done investigating the transitory behaviour of falling paper. The work highlights the potential use of robotics in the investigation of highly complex physical systems. Because of their practicability and rich behavioural diversity, falling-paper systems are an ideal case study for robotic intervention such as the IPES presented here. However, the philosophy and approach presented in this study is intended as a template for the wider adoption of robotics and machine learning for investigating complex physical systems, and for designing bio-inspired robotic systems.

In Chapters 7 and 8, I introduced and analysed the single and collective Bernoulli-ball system. The control strategy developed for the height control is a powerful approach for the control of systems in which conventional feedback controllers are hard to implement. To develop this further, alternative control problems should be addressed, for example transferring a single Bernoulli-ball between airflows. This has clear practical applications for non-contact object transportation. Additionally, further study on collective Bernoulli-ball systems is required to investigate further mechanisms to induce behavioural diversity in the system. These studies could help to establish whether the system demonstrates fundamental aspects of how future life-like robotics systems can be built. These aspects include self-organisation and regulation, evolutionary dynamics and collective behaviours. In parallel, they can aid in the discourse relating to how we define life-like properties in man-made systems (Bongard and Levin, 2021).

9.2.1 Final remarks

Ultimately, future work related this thesis should focus on the incremental progression toward designing more complex robotic systems that make use of the fundamental principles outlined. I highlight two key directions here:

1. Robot scientists. One of the fundamental aspects of this thesis was the large-scale experimentation system presented in Chapter 6. Future work should build on the system and design philosophy to push the limits of what we can design and manufacture automatically. As robotic manipulation and fabrication techniques progress, such robot scientists will be able to design and interact with ever more complex systems. Similarly, ever more creative and large machine learning systems should be utilised to improve the robot scientist ‘brain’.
2. Controlling falling paper. A key direction is understanding to what extent we can control and utilise the behaviours of falling paper. There is a clear use case in the fields of flying robotics, for example passive delivery mechanisms such as parachutes. In Chapter 3 I discussed the benefits of certain smart and active materials such as shape memory alloys. Research into how these can be incorporated into falling paper shapes to induce morphology, and hence behaviour, changes is a key direction for further research. Similarly, understanding how falling paper dynamics fit into the broader context of flapping and insect-like robots could prove interesting.

References

- Adamatzky, A. (2009). Hot ice computer. *Physics Letters A*, 374(2):264–271.
- Adamatzky, A., Sirakoulis, G. C., Martinez, G. J., Baluška, F., and Mancuso, S. (2017). On plant roots logical gates. *BioSystems*, 156:40–45.
- Alapan, Y., Yasa, O., Yigit, B., Yasa, I. C., Erkoç, P., and Sitti, M. (2019). Microrobotics and microorganisms: biohybrid autonomous cellular robots. *Annual Review of Control, Robotics, and Autonomous Systems*, 2:205–230.
- Alexander, R. (1989). Optimization and gaits in the locomotion of vertebrates. *Physiological Reviews*, 69(4):1199–1227.
- Andersen, A., Pesavento, U., and Wang, Z. J. (2005). Analysis of transitions between fluttering, tumbling and steady descent of falling cards. *Journal of Fluid Mechanics*, 541:91–104.
- Anderson, D. and Eberhardt, S. (2001). The Newtonian description of lift of a wing. *American Journal of Physics*.
- Angeles, J. (2002). *Fundamentals Of Robotic Mechanical Systems*, volume 2. Springer.
- Ashby, W. R. (1960). The homeostat. In *Design for a Brain*, pages 100–121. Springer.
- Ashby, W. R. (1961). *An introduction to cybernetics*. Chapman & Hall Ltd.
- Auerbach, J. E. and Bongard, J. C. (2012). On the relationship between environmental and morphological complexity in evolved robots. In *Proceedings of the 14th Annual Conference on Genetic and Evolutionary Computation*, pages 521–528.
- Auguste, F., Magnaudet, J., and Fabre, D. (2013). Falling styles of disks. *Journal of Fluid Mechanics*, 719:388–405.
- Begg, R. K., Palaniswami, M., and Owen, B. (2005). Support vector machines for automated gait classification. *IEEE Transactions on Biomedical Engineering*, 52(5):828–838.
- Bellemare, M. G., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. (2016). Unifying count-based exploration and intrinsic motivation. *arXiv preprint arXiv:1606.01868*.
- Belmonte, A., Eisenberg, H., and Moses, E. (1998). From flutter to tumble: inertial drag and Froude similarity in falling paper. *Physical Review Letters*, 81(2):345–348.

- Bergou, A. J., Xu, S., and Wang, Z. J. (2007). Passive wing pitch reversal in insect flight. *Journal of Fluid Mechanics*, 591:321–337.
- Berman, G. J. and Wang, Z. J. (2007). Energy-minimizing kinematics in hovering insect flight. *Journal of Fluid Mechanics*, 582:153–168.
- Bomphrey, R. J., Taylor, G. K., and Thomas, A. L. (2010). Smoke visualization of free-flying bumblebees indicates independent leading-edge vortices on each wing pair. In *Animal Locomotion*, pages 249–259. Springer.
- Bongard, J. and Levin, M. (2021). Living things are not (20th century) machines: updating mechanism metaphors in light of the modern science of machine behavior. *Frontiers in Ecology and Evolution*, 9:147.
- Bongard, J. and Lipson, H. (2007). Automated reverse engineering of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 104(24):9943–9948.
- Bongard, J. C. and Lipson, H. (2005). Nonlinear system identification using coevolution of models and tests. *IEEE Transactions on Evolutionary Computation*, 9(4):361–384.
- Bottou, L., Curtis, F. E., and Nocedal, J. (2018). Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311.
- Bowyer, A. (2014). 3D printing and humanity’s first imperfect replicator. *3D printing and additive manufacturing*, 1(1):4–5.
- Brigham, E. O. (1988). *The Fast Fourier Transform and Its Applications*. Prentice-Hall, Inc., USA.
- Brodbeck, L., Hauser, S., and Iida, F. (2015). Morphological evolution of physical robots through model-free phenotype development. *PLoS One*, 10(6):e0128444.
- Brooks, R. A. (1990). Elephants don’t play chess. *Robotics and autonomous systems*, 6(1-2):3–15.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial intelligence*, 47(1-3):139–159.
- Bruce, J. W., Bruce, J. W., and Giblin, P. (1992). *Curves and Singularities: a geometrical introduction to singularity theory*. Cambridge university press.
- Brunton, S. L., Proctor, J. L., and Kutz, J. N. (2016). Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937.
- Burrige, R. R., Rizzi, A. A., and Koditschek, D. E. (1999). Sequential composition of dynamically dexterous robot behaviours. *The International Journal of Robotics Research*, 18(6):534–555.
- Calandra, R., Seyfarth, A., Peters, J., and Deisenroth, M. P. (2016). Bayesian optimization for learning gaits under uncertainty. *Annals of Mathematics and Artificial Intelligence*, 76(1-2):5–23.

- Caluwaerts, K., Despraz, J., Işçen, A., Sabelhaus, A. P., Bruce, J., Schrauwen, B., and SunSpiral, V. (2014). Design and control of compliant tensegrity robots through simulation and hardware validation. *Journal of the Royal Society Interface*, 11(98):20140520.
- Camarillo, D. B., Milne, C. F., Carlson, C. R., Zinn, M. R., and Salisbury, J. K. (2008). Mechanics modeling of tendon-driven continuum manipulators. *IEEE Transactions on robotics*, 24(6):1262–1273.
- Cellucci, D., MacCurdy, R., Lipson, H., and Risi, S. (2017). 1D printing of recyclable robots. *IEEE Robotics and Automation Letters*, 2(4):1964–1971.
- Chaloner, K. and Verdinelli, I. (1995). Bayesian experimental design: A review. *Statistical Science*, pages 273–304.
- Chapman, T. (2003). Lab automation and robotics: automation on the move. *Nature*, 421(6923):661–663.
- Cheney, N., Bongard, J., and Lipson, H. (2015). Evolving soft robots in tight spaces. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, pages 935–942.
- Cheney, N., Clune, J., and Lipson, H. (2014). Evolved Electrophysiological soft robots. In *Artificial Life Conference Proceedings 14*, pages 222–229.
- Cheney, N., MacCurdy, R., Clune, J., and Lipson, H. (2013). Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proceedings of the 15th Annual Conference on Genetic and Evolutionary Computation*, pages 167–174.
- Choi, C., Schwarting, W., DelPreto, J., and Rus, D. (2018). Learning object grasping for soft robot hands. *IEEE Robotics and Automation Letters*, 3(3):2370–2377.
- Chossat, J.-B., Park, Y.-L., Wood, R. J., and Duchaine, V. (2013). A soft strain sensor based on ionic and metal liquids. *IEEE Sensors Journal*, 13(9):3405–3414.
- Christiansen, E. and Barker, D. H. (1965). The effect of shape and density on the free settling of particles at high reynolds numbers. *AIChE Journal*, 11(1):145–151.
- Chrust, M., Bouchet, G., and Dûek, J. (2013). Numerical simulation of the dynamics of freely falling discs. *Physics of Fluids*, 25(4):044102.
- Cira, N. J., Benusiglio, A., and Prakash, M. (2015). Vapour-mediated sensing and motility in two-component droplets. *Nature*, 519(7544):446–450.
- Clark, A. (2008). *Supersizing The Mind: Embodiment, Action, And Cognitive Extension*. Oxford University Press.
- Clune, J., Mouret, J.-B., and Lipson, H. (2013). The evolutionary origins of modularity. *Proceedings of the Royal Society B: Biological sciences*, 280(1755):20122863.
- Clune, J., Stanley, K. O., Pennock, R. T., and Ofria, C. (2011). On the performance of indirect encoding across the continuum of regularity. *IEEE Transactions on Evolutionary Computation*, 15(3):346–367.

- Collins, S., Ruina, A., Tedrake, R., and Wisse, M. (2005). Efficient bipedal robots based on passive-dynamic walkers. *Science*, 307(5712):1082–1085.
- Connolly, F., Polygerinos, P., Walsh, C. J., and Bertoldi, K. (2015). Mechanical programming of soft actuators by varying fiber angle. *Soft Robotics*, 2(1):26–32.
- Corucci, F., Calisti, M., Hauser, H., and Laschi, C. (2015). Novelty-based evolutionary design of morphing underwater robots. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, pages 145–152.
- Corucci, F., Cheney, N., Giorgio-Serchi, F., Bongard, J., and Laschi, C. (2018). Evolving soft locomotion in aquatic and terrestrial environments: effects of material properties and environmental transitions. *Soft Robotics*, 5(4):475–495.
- Cully, A., Clune, J., Tarapore, D., and Mouret, J.-B. (2015). Robots that can adapt like animals. *Nature*, 521(7553):503–507.
- Davey, J., Kwok, N., and Yim, M. (2012). Emulating self-reconfigurable robots – design of the SMORES system. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4464–4469.
- Della Santina, C. and Rus, D. (2019). Control oriented modeling of soft robots: the polynomial curvature case. *IEEE Robotics and Automation Letters*, 5(2):290–298.
- Dickinson, M. H., Lehmann, F.-O., and Sane, S. P. (1999). Wing rotation and the aerodynamic basis of insect flight. *Science*, 284(5422):1954–1960.
- Dickman, R., Attard, P., and Simonian, V. (1997). Entropic forces in binary hard sphere mixtures: theory and simulation. *The Journal of Chemical Physics*, 107(1):205–213.
- Doncieux, S., Laflaquière, A., and Coninx, A. (2019). Novelty search: a theoretical perspective. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 99–106.
- Doursat, R. and Sánchez, C. (2014). Growing fine-grained multicellular robots. *Soft Robotics*, 1(2):110–121.
- Drotman, D., Jadhav, S., Karimi, M., Dezonias, P., and Tolley, M. T. (2017). 3D printed soft actuators for a legged robot capable of navigating unstructured terrain. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5532–5538.
- Duarte, M., Gomes, J., Oliveira, S. M., and Christensen, A. L. (2017). Evolution of repertoire-based control for robots with complex locomotor systems. *IEEE Transactions on Evolutionary Computation*, 22(2):314–328.
- Duriez, C. (2013). Control of elastic soft robots based on real-time finite element method. In *2013 IEEE International Conference on Robotics and Automation*, pages 3982–3987. IEEE.
- Eiben, A., Bredeche, N., Hoogendoorn, M., Stradner, J., Timmis, J., Tyrrell, A. M., and Winfield, A. (2013). The triangle of life: evolving robots in real-time and real-space. In *Artificial Life Conference Proceedings 13*, pages 1056–1063. MIT Press.

- Eiben, A. E., Kernbach, S., and Haasdijk, E. (2012). Embodied artificial evolution. *Evolutionary Intelligence*, 5(4):261–272.
- El Hajjaji, A. and Ouladsine, M. (2001). Modeling and nonlinear control of magnetic levitation systems. *IEEE Transactions on Industrial Electronics*, 48(4):831–838.
- Ellington, C. P., Van Den Berg, C., Willmott, A. P., and Thomas, A. L. (1996). Leading-edge vortices in insect flight. *Nature*, 384(6610):626–630.
- Fan, D., Jodin, G., Consi, T., Bonfiglio, L., Ma, Y., Keyes, L., Karniadakis, G. E., and Triantafyllou, M. S. (2019). A robotic intelligent towing tank for learning complex fluid-structure dynamics. *Science Robotics*, 4(36):eaay5063.
- Fernandes, P. C., Ern, P., Risso, F., and Magnaudet, J. (2005). On the zigzag dynamics of freely moving axisymmetric bodies. *Physics of Fluids*, 17(9):1–4.
- Fernandes, P. C., Risso, F., Ern, P., and Magnaudet, J. (2007). Oscillatory motion and wake instability of freely rising axisymmetric bodies. *Journal of Fluid Mechanics*, 573:479–502.
- Field, S. B., Klaus, M., Moore, M., and Nori, F. (1997). Chaotic dynamics of falling disks. *Nature*, 388(6639):252.
- Firouzeh, A., Amon-Junior, A. F., and Paik, J. (2015). Soft piezoresistive sensor model and characterization with varying design parameters. *Sensors and Actuators A: Physical*, 233:158–168.
- Fonseca, F. and Herrmann, H. (2005). Simulation of the sedimentation of a falling oblate ellipsoid. *Physica A: Statistical Mechanics and its Applications*, 345(3-4):341–355.
- Frankel, F. and Reid, R. (2008). Big data: distilling meaning from data. *Nature*, 455(7209):30–30.
- Frazier, P. I. (2018). A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*.
- Frutiger, A., Muth, J. T., Vogt, D. M., Mengüç, Y., Campo, A., Valentine, A. D., Walsh, C. J., and Lewis, J. A. (2015). Capacitive soft strain sensors via multicore-shell fiber printing. *Advanced Materials*, 27(15):2440–2446.
- Gandomi, A. and Haider, M. (2015). Beyond the hype: big data concepts, methods, and analytics. *International Journal of Information Management*, 35(2):137–144.
- Garrad, M., Soter, G., Conn, A., Hauser, H., and Rossiter, J. (2019). A soft matter computer for soft robots. *Science Robotics*, 4(33):eaaw6060.
- Gerhart, P. M., Gerhart, A. L., and Hochstein, J. I. (2016). *Munson, Young and Okiishi's Fundamentals of Fluid Mechanics*. John Wiley & Sons.
- Gilday, K., Thuruthel, T. G., and Iida, F. (2020). A vision-based collocated actuation-sensing scheme for a compliant tendon-driven robotic hand. In *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*, pages 760–765. IEEE.

- Gilpin, K., Knaian, A., and Rus, D. (2010). Robot pebbles: one centimeter modules for programmable matter through self-disassembly. In *2010 IEEE International Conference on Robotics and Automation*, pages 2485–2492. IEEE.
- Gilpin, K., Kotay, K., Rus, D., and Vasilescu, I. (2008). Miche: modular shape formation by self-disassembly. *The International Journal of Robotics Research*, 27(3-4):345–372.
- Giorelli, M., Renda, F., Calisti, M., Arienti, A., Ferri, G., and Laschi, C. (2015). Neural network and Jacobian method for solving the inverse statics of a cable-driven soft arm with nonconstant curvature. *IEEE Transactions on Robotics*, 31(4):823–834.
- Goldberg, D. E. and Holland, J. H. (1988). Genetic algorithms and machine learning. *Machine Learning*, 3(2–3):95–99.
- Goswami, A., Espiau, B., and Keramane, A. (1997). Limit cycles in a passive compass gait biped and passivity-mimicking control laws. *Autonomous Robots*, 4(3):273–286.
- Goto, Y. and Tanaka, H. (2015). Purely hydrodynamic ordering of rotating disks at a finite reynolds number. *Nature Communications*, 6:5994.
- Goza, A. and Colonius, T. (2018). Modal decomposition of fluid–structure interaction with application to flag flapping. *Journal of Fluids and Structures*, 81:728–737.
- Gu, G.-Y., Zhu, J., Zhu, L.-M., and Zhu, X. (2017). A survey on dielectric elastomer actuators for soft robots. *Bioinspiration & Biomimetics*, 12(1):011003.
- Guenther, F. and Iida, F. (2017). Energy-efficient monopod running with a large payload based on open-loop parallel elastic actuation. *IEEE Transactions on Robotics*, 33(1):102–113.
- Gul, J. Z., Sajid, M., Rehman, M. M., Siddiqui, G. U., Shah, I., Kim, K.-H., Lee, J.-W., and Choi, K. H. (2018). 3D printing for soft robotics—a review. *Science and Technology of Advanced Materials*, 19(1):243–262.
- Hanczyc, M. M. and Ikegami, T. (2010). Chemical basis for minimal cognition. *Artificial Life*, 16(3):233–243.
- Hans Lugt David W Taylor, A. J. (1983). Autorotation. *Ann. Rev. Fluid Mech*, 15:123–47.
- Hartigan, J. A. and Wong, M. A. (1979). Algorithm as 136: a k-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 28(1):100–108.
- Hauser, S., Mutlu, M., Léziart, P. A., Khodr, H., Bernardino, A., and Ijspeert, A. J. (2020). Roombots extended: challenges in the next generation of self-reconfigurable modular robots and their application in adaptive and assistive furniture. *Robotics and Autonomous Systems*, page 103467.
- Hawkes, E. W., Blumenschein, L. H., Greer, J. D., and Okamura, A. M. (2017). A soft robot that navigates its environment through growth. *Science Robotics*, 2(8):eaan3028.
- Heisinger, L., Newton, P., and Kanso, E. (2014). Coins falling in water. *Journal of fluid mechanics*, 742:243–253.

- Heylighen, F. and Joslyn, C. (2001). Cybernetics and second-order cybernetics. *Encyclopedia of physical science & technology*, 4:155–170.
- Hiller, J. and Lipson, H. (2011). Automatic design and manufacture of soft robots. *IEEE Transactions on Robotics*, 28(2):457–466.
- Hiller, J. and Lipson, H. (2014). Dynamic simulation of soft multimaterial 3D-printed objects. *Soft Robotics*, 1(1):88–101.
- Holmes, P. J. (1982). The dynamics of repeated impacts with a sinusoidally vibrating table. *Journal of Sound and Vibration*, 84(2):173–189.
- Horibe, N., Hanczyc, M. M., and Ikegami, T. (2011). Mode switching and collective behaviour in chemical oil droplets. *Entropy*, 13(3):709–719.
- Hornby, G. S. and Pollack, J. B. (2001). The advantages of generative grammatical encodings for physical design. In *Proceedings of the 2001 Congress on Evolutionary Computation*, volume 1, pages 600–607. IEEE.
- Howard, D., Eiben, A. E., Kennedy, D. F., Mouret, J.-B., Valencia, P., and Winkler, D. (2019). Evolving embodied intelligence from materials to machines. *Nature Machine Intelligence*, 1(1):12–19.
- Howison, T., Hughes, J., and Iida, F. (2020a). Large-scale automated investigation of free-falling paper shapes via iterative physical experimentation. *Nature Machine Intelligence*, 2(1):68–75.
- Howison, T., Hughes, J., and Iida, F. (2020b). Morphologically programming the interactions of v-shaped falling papers. In *Artificial Life Conference Proceedings*, pages 359–366. MIT Press.
- Hu, Y., Liu, J., Spielberg, A., Tenenbaum, J. B., Freeman, W. T., Wu, J., Rus, D., and Matusik, W. (2019). Chainqueen: a real-time differentiable physical simulator for soft robotics. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 6265–6271. IEEE.
- Huang, W., Liu, H., Wang, F., Wu, J., and Zhang, H. (2013). Experimental study of a freely falling plate with an inhomogeneous mass distribution. *Physical Review E*, 88(5):053008.
- Hughes, J. and Iida, F. (2017). 3D printed sensorized soft robotic manipulator design. In *Annual Conference Towards Autonomous Robotic Systems*, pages 627–636. Springer.
- Hughes, J., Maiolino, P., and Iida, F. (2018). An anthropomorphic soft skeleton hand exploiting conditional models for piano playing. *Science Robotics*, 3(25):eaau3098.
- Huizinga, J., Stanley, K. O., and Clune, J. (2018). The emergence of canalization and evolvability in an open-ended, interactive evolutionary system. *Artificial Life*, 24(3):157–181.
- Hunt, J., Giardina, F., Rosendo, A., and Iida, F. (2016). Improving efficiency for an open-loop-controlled locomotion with a pulsed actuation. *IEEE/ASME Transactions on Mechatronics*, 21(3):1581–1591.

- Iida, F. and Ijspeert, A. J. (2016). Biologically inspired robotics. In *Springer Handbook of Robotics*, pages 2015–2034. Springer.
- Iida, F. and Laschi, C. (2011). Soft robotics: challenges and perspectives. *Procedia Computer Science*, 7:99–102.
- Ijspeert, A. J. (2014). Biorobotics: using robots to emulate and investigate agile locomotion. *Science*, 346(6206):196–203.
- Ilievski, F., Mazzeo, A. D., Shepherd, R. F., Chen, X., and Whitesides, G. M. (2011). Soft robotics for chemists. *Angewandte Chemie*, 123(8):1930–1935.
- Iversen, J. D. (1979). Autorotating flat-plate wings: the effect of the moment of inertia, geometry and Reynolds number. *Journal of Fluid Mechanics*, 92(2):327–348.
- Jacob, C. (1994). Genetic l-system programming. In *International Conference on Parallel Problem Solving from Nature*, pages 333–343. Springer.
- Jakobi, N. (1997). Evolutionary robotics and the radical envelope-of-noise hypothesis. *Adaptive Behaviour*, 6(2):325–368.
- Jayaweera, K. and Cottis, R. (1969). Fall velocities of plate-like and columnar ice crystals. *Quarterly Journal of the Royal Meteorological Society*, 95(406):703–709.
- Jazar, R. N. (2010). *Theory Of Applied Robotics: Kinematics, Dynamics, and Control*. Springer Science & Business Media.
- Jin, C. and Xu, K. (2008). Numerical study of the unsteady aerodynamics of freely falling plates. *Communications in Computational Physics*, 3(4):834–851.
- Joachimczak, M., Suzuki, R., and Arita, T. (2015). Improving evolvability of morphologies and controllers of developmental soft-bodied robots with novelty search. *Frontiers in Robotics and AI*, 2:33.
- Kachel, V., Sindelar, G., and Grimm, S. (2006). High-throughput isolation of ultra-pure plasmid dna by a robotic system. *BMC Biotechnology*, 6(1):1–8.
- Kamela, M. (2007). Thinking about Bernoulli. *The Physics Teacher*, 45(6):379–381.
- Katzschmann, R. K., Marchese, A. D., and Rus, D. (2016). Hydraulic autonomous soft robotic fish for 3D swimming. In *Experimental Robotics*, pages 405–420. Springer.
- Kaya, C. Y. and Noakes, J. L. (1996). Computations and time-optimal controls. *Optimal Control Applications and Methods*, 17(3):171–185.
- Khazanov, M., Jocque, J., and Rieffel, J. (2014). Evolution of locomotion on a physical tensegrity robot. In *Artificial Life Conference Proceedings 14*, pages 232–238. MIT Press.
- Kim, S., Laschi, C., and Trimmer, B. (2013). Soft robotics: a bioinspired evolution in robotics. *Trends in Biotechnology*, 31(5):287–294.

- King, R. D., Whelan, K. E., Jones, F. M., Reiser, P. G., Bryant, C. H., Muggleton, S. H., Kell, D. B., and Oliver, S. G. (2004). Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature*, 427(6971):247–252.
- Kooijman, J., Meijaard, J. P., Papadopoulos, J. M., Ruina, A., and Schwab, A. (2011). A bicycle can be self-stable without gyroscopic or caster effects. *Science*, 332(6027):339–342.
- Koos, S., Mouret, J.-B., and Doncieux, S. (2012). The transferability approach: crossing the reality gap in evolutionary robotics. *IEEE Transactions on Evolutionary Computation*, 17(1):122–145.
- Kriegman, S., Blackiston, D., Levin, M., and Bongard, J. (2020a). A scalable pipeline for designing reconfigurable organisms. *Proceedings of the National Academy of Sciences*, 117(4):1853–1859.
- Kriegman, S., Cappelle, C., Corucci, F., Bernatskiy, A., Cheney, N., and Bongard, J. C. (2017). Simulating the evolution of soft and rigid-body robots. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 1117–1120.
- Kriegman, S., Cheney, N., and Bongard, J. (2018). How morphological development can guide evolution. *Scientific Reports*, 8(1):1–10.
- Kriegman, S., Nasab, A. M., Shah, D., Steele, H., Branin, G., Levin, M., Bongard, J., and Kramer-Bottiglio, R. (2020b). Scalable sim-to-real transfer of soft robot designs. In *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*, pages 359–366. IEEE.
- Kutz, J. N. (2017). Deep learning in fluid dynamics. *Journal of Fluid Mechanics*, 814:1–4.
- Kutz, J. N., Brunton, S. L., Brunton, B. W., and Proctor, J. L. (2016). *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*, volume 149. SIAM.
- Kwiatkowski, R. and Lipson, H. (2019). Task-agnostic self-modeling machines. *Science Robotics*, 4(26):eaau9354.
- Laschi, C., Mazzolai, B., and Cianchetti, M. (2016). Soft robotics: technologies and systems pushing the boundaries of robot abilities. *Science Robotics*, 1(1):eaah3690.
- Lau, M., Ohgawara, A., Mitani, J., and Igarashi, T. (2011). Converting 3D furniture models to fabricatable parts and connectors. *ACM Transactions on Graphics (TOG)*, 30(4):1–6.
- Lauga, E. and Nadal, F. (2017). Clustering instability of focused swimmers. *EPL (Europhysics Letters)*, 116(6):64004.
- Lee, C., Su, Z., Zhong, H., Chen, S., Zhou, M., and Wu, J. (2013). Experimental investigation of freely falling thin disks. Part 2. Transition of three-dimensional motion from zigzag to spiral. *Journal of Fluid Mechanics*, 732:77–104.
- Lee, I. and Choi, H. (2018). Scaling law for the lift force of autorotating falling seeds at terminal velocity. *Journal of Fluid Mechanics*, 835:406–420.

- Lehman, J. and Miikkulainen, R. (2015). Enhancing divergent search through extinction events. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, pages 951–958.
- Lehman, J. and Stanley, K. O. (2011a). Abandoning objectives: evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223.
- Lehman, J. and Stanley, K. O. (2011b). Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation*, pages 211–218.
- Lentink, D., Dickson, W. B., Van Leeuwen, J. L., and Dickinson, M. H. (2009). Leading-edge vortices elevate lift of autorotating plant seeds. *Science*, 324(5933):1438–1440.
- Li, J., Liu, L., Liu, Y., and Leng, J. (2019). Dielectric elastomer spring-roll bending actuators: applications in soft robotics and design. *Soft Robotics*, 6(1):69–81.
- Liang, J., Lin, M., and Koltun, V. (2019). Differentiable cloth simulation for inverse problems. In *Advances in Neural Information Processing Systems*, pages 771–780.
- Lipson, H. (2014). Challenges and opportunities for design, simulation, and fabrication of soft robots. *Soft Robotics*, 1(1):21–27.
- Lipton, J. I., MacCurdy, R., Manchester, Z., Chin, L., Cellucci, D., and Rus, D. (2018). Hand- edness in shearing auxetics creates rigid and compliant structures. *Science*, 360(6389):632–635.
- List, R. and Schemenauer, R. S. (1971). Free-fall behavior of planar snow crystals, conical graupel and small hail. *Journal of Atmospheric Sciences*, 28(1):110–115.
- Liu, Y.-Y., Slotine, J.-J., and Barabási, A.-L. (2011). Controllability of complex networks. *nature*, 473(7346):167–173.
- Lizotte, D. J. (2008). *Practical Bayesian Optimization*. University of Alberta.
- Magono, C. and Nakamura, T. (1965). Aerodynamic studies of falling snowflakes. *Journal of the Meteorological Society of Japan. Ser. II*, 43(3):139–147.
- Mahadevan, L., Aref, H., and Jones, S. (1995). Comment on “behavior of a falling paper”. *Physical review letters*, 75(7):1420.
- Mahadevan, L., Ryu, W. S., and Samuel, A. D. (1999). Tumbling cards. *Physics of Fluids*, 11(1):1–3.
- Malley, M., Rubenstein, M., and Nagpal, R. (2017). Flippy: a soft, autonomous climber with simple sensing and control. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6533–6540.
- Mann, B. and Sims, N. (2009). Energy harvesting from the nonlinear oscillations of magnetic levitation. *Journal of Sound and Vibration*, 319(1–2):515–530.
- Maxwell, J. C. (1854). On a particular case of the descent of a heavy body in a resisting medium. *Cambridge and Dublin Mathematical Journal*, 9:145–148.

- Maxwell, J. C. (1873). Does the progress of physical science tend to give any advantage to the opinion of necessity (or determinism) over that of the contingency of events and the freedom of the will. *The Life of James Clerk Maxwell. With a Selection from His Correspondence and Occasional Writings and a Sketch of His Contributions to Science*, 2nd edn. MacMillan, London (1882, 1999). <https://archive.org/details/lifeofjamesclerk00camprich>.
- Maziz, A., Concas, A., Khaldi, A., Stålhand, J., Persson, N.-K., and Jager, E. W. (2017). Knitting and weaving artificial muscles. *Science Advances*, 3(1):e1600327.
- McCormack, J., Dorin, A., Innocent, T., et al. (2004). Generative design: a paradigm for design research. In *Proceedings of Futureground, Design Research Society, Melbourne*.
- McCutchen, C. (1977). The spinning rotation of ash and tulip tree samaras. *Science*, 197(4304):691–692.
- McGeer, T. et al. (1990). Passive dynamic walking. *International Journal of Robotics Research*, 9(2):62–82.
- Mehta, A. M., DelPreto, J., Shaya, B., and Rus, D. (2014). Cogeneration of mechanical, electrical, and software designs for printable robots from structural specifications. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2892–2897. IEEE.
- Mehta, A. M. and Rus, D. (2014). An end-to-end system for designing mechanical structures for print-and-fold robots. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1460–1465. IEEE.
- Meyn, S. P. and Tweedie, R. L. (2012). *Markov Chains And Stochastic Stability*. Springer Science & Business Media.
- Miras, K., Ferrante, E., and Eiben, A. E. (2020). Environmental influences on evolvable robots. *PLoS One*, 15(5):e0233848.
- Mittal, R., Seshadri, V., and Udaykumar, H. S. (2004). Flutter, tumble and vortex induced autorotation. *Theoretical and Computational Fluid Dynamics*, 17(3):165–170.
- Mjolsness, E. and DeCoste, D. (2001). Machine learning for science: state of the art and future prospects. *Science*, 293(5537):2051–2055.
- Moravec, H. (1988). *Mind Children: The Future of Robot and Human Intelligence*. Harvard University Press.
- Morimoto, Y., Onoe, H., and Takeuchi, S. (2018). Biohybrid robot powered by an antagonistic pair of skeletal muscle tissues. *Science Robotics*, 3(18):eaat4440.
- Motzki, P., Khelfa, F., Zimmer, L., Schmidt, M., and Seelecke, S. (2019). Design and validation of a reconfigurable robotic end-effector based on shape memory alloys. *IEEE/ASME Transactions on Mechatronics*, 24(1):293–303.
- Mouret, J.-B. and Chatzilygeroudis, K. (2017). 20 years of reality gap: a few thoughts about simulators in evolutionary robotics. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 1121–1124.

- Mouret, J.-B. and Clune, J. (2015). Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*.
- Mouret, J.-B. and Doncieux, S. (2008). Incremental evolution of animats' behaviours as a multi-objective optimization. In *International Conference on Simulation of Adaptive Behaviour*, pages 210–219. Springer.
- Mouret, J.-B. and Doncieux, S. (2012). Encouraging behavioural diversity in evolutionary robotics: an empirical study. *Evolutionary Computation*, 20(1):91–133.
- Müller, V. C. and Hoffmann, M. (2017). What is morphological computation? on how the body contributes to cognition and control. *Artificial Life*, 23(1):1–24.
- Nakajima, K., Hauser, H., Kang, R., Guglielmino, E., Caldwell, D. G., and Pfeifer, R. (2013). A soft body as a reservoir: case studies in a dynamic model of octopus-inspired soft robotic arm. *Frontiers in Computational Neuroscience*, 7:91.
- Nakajima, K., Hauser, H., Li, T., and Pfeifer, R. (2015). Information processing via physical soft body. *Scientific Reports*, 5:10487.
- Nakajima, K., Nougabeu, A. M. T., Miyashita, S., Göldi, M., Fuchslin, R. M., and Pfeifer, R. (2012). Morphology-induced collective behaviours: dynamic pattern formation in water-floating elements. *PloS One*, 7(6):e37805.
- Nave, G. K., Hall, N., Somers, K., Davis, B., Gruszewski, H., Powers, C., Collver, M., Schmale, D. G., and Ross, S. D. (2021). Wind dispersal of natural and biomimetic maple samaras. *Biomimetics*, 6(2):23.
- Neubert, J. and Lipson, H. (2016). Soldercubes: a self-soldering self-reconfiguring modular robot system. *Autonomous Robots*, 40(1):139–158.
- Norberg, R. Å. (1973). Autorotation, self-stability, and structure of single-winged fruits and seeds (samaras) with comparative remarks on animal flight. *Biological Reviews*, 48(4):561–596.
- Norris, J. R. and Norris, J. R. (1998). *Markov Chains*. Cambridge University Press.
- Nudehi, S. S., Dehmlow, S., and Clark, D. (2017). Position control of a floating ball in a vertical air stream. In *ASME 2017 International Mechanical Engineering Congress and Exposition*, pages V04AT05A002–V04AT05A002. American Society of Mechanical Engineers.
- Nurzaman, S. G., Culha, U., Brodbeck, L., Wang, L., and Iida, F. (2013). Active sensing system with in situ adjustable sensor morphology. *PLoS One*, 8(12):e84090.
- Nurzaman, S. G., Iida, F., Margheri, L., and Laschi, C. (2014). Soft robotics on the move: Scientific networks, activities, and future challenges.
- Nygaard, T. F., Martin, C. P., Samuelson, E., Torresen, J., and Glette, K. (2018). Real-world evolution adapts robot morphology and control to hardware limitations. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 125–132.

- Onal, C. D., Wood, R. J., and Rus, D. (2011). Towards printable robotics: origami-inspired planar fabrication of three-dimensional mechanisms. In *2011 IEEE International Conference on Robotics and Automation*, pages 4608–4613. IEEE.
- Park, M., Kim, Y.-J., and Lee, J.-J. (2011). Swing-up and LQR stabilization of a rotary inverted pendulum. *Artificial Life and Robotics*, 16(1):94–97.
- Paul, C. (2006). Morphological computation: A basis for the analysis of morphology and control requirements. *Robotics and Autonomous Systems*, 54(8):619–630.
- Peng, X. B., Andrychowicz, M., Zaremba, W., and Abbeel, P. (2018). Sim-to-real transfer of robotic control with dynamics randomization. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8. IEEE.
- Peplow, M. (2014). Organic synthesis: the robo-chemist. *Nature News*, 512(7512):20.
- Pesavento, U. and Wang, Z. J. (2004). Falling paper: Navier-stokes solutions, model of fluid forces, and center of mass elevation. *Physical review letters*, 93(14):144501.
- Pfeifer, R. and Bongard, J. (2006). *How The Body Shapes The Way We Think: A New View Of Intelligence*. MIT press.
- Pfeifer, R., Iida, F., and Bongard, J. (2005). New robotics: design principles for intelligent systems. *Artificial Life*, 11(1–2):99–120.
- Pfeifer, R., Lungarella, M., and Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. *Science*, 318(5853):1088–1093.
- Pfeifer, R., Lungarella, M., and Iida, F. (2012). The challenges ahead for bio-inspired soft robotics. *Communications of the ACM*, 55(11):76–87.
- Polygerinos, P., Correll, N., Morin, S. A., Mosadegh, B., Onal, C. D., Petersen, K., Cianchetti, M., Tolley, M. T., and Shepherd, R. F. (2017). Soft robotics: Review of fluid-driven intrinsically soft devices; manufacturing, sensing, control, and applications in human-robot interaction. *Advanced Engineering Materials*, 19(12):1700016.
- Pozzi, M., Miguel, E., Deimel, R., Malvezzi, M., Bickel, B., Brock, O., and Prattichizzo, D. (2018). Efficient fem-based simulation of soft robots modeled as kinematic chains. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8. IEEE.
- Pugh, J. K., Soros, L. B., and Stanley, K. O. (2016). Quality diversity: a new frontier for evolutionary computation. *Frontiers in Robotics and AI*, 3:40.
- Quiroz, M., Kohn, R., Villani, M., and Tran, M.-N. (2018). Speeding up MCMC by efficient data subsampling. *Journal of the American Statistical Association*, pages 1–35.
- Rasmussen, C. E. (2003). Gaussian processes in machine learning. In *Summer School on Machine Learning*, pages 63–71. Springer.
- Reist, P. and D’Andrea, R. (2012). Design and analysis of a blind juggling robot. *IEEE Transactions on Robotics*, 28(6):1228–1243.

- Richards, D. and Amos, M. (2014). Evolving morphologies with cppn-neat and a dynamic substrate. In *Artificial Life Conference Proceedings 14*, pages 255–262. MIT Press.
- Rieffel, J., Knox, D., Smith, S., and Trimmer, B. (2014). Growing and evolving soft robots. *Artificial Life*, 20(1):143–162.
- Rieffel, J. and Mouret, J.-B. (2018). Adaptive and resilient soft tensegrity robots. *Soft Robotics*, 5(3):318–329.
- Rieffel, J. and Pollack, J. (2005). Crossing the fabrication gap: evolving assembly plans to build 3-d objects. In *2005 IEEE Congress on Evolutionary Computation*, volume 1, pages 529–536. IEEE.
- Romanishin, J. W., Gilpin, K., and Rus, D. (2013). M-blocks: Momentum-driven, magnetic modular robots. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4288–4295. IEEE.
- Romano, D., Donati, E., Benelli, G., and Stefanini, C. (2019). A review on animal–robot interaction: from bio-hybrid organisms to mixed societies. *Biological Cybernetics*, 113(3):201–225.
- Ronsse, R., Lefevre, P., and Sepulchre, R. (2007). Rhythmic feedback control of a blind planar juggler. *IEEE Transactions on Robotics*, 23(4):790–802.
- Rosendo, A., Von Atzigen, M., and Iida, F. (2017). The trade-off between morphology and control in the co-optimized design of robots. *PLoS One*, 12(10):e0186107.
- Rosser, K., Kok, J., Chahl, J., and Bongard, J. (2019). Sim2real gap is non-monotonic with robot complexity for morphology-in-the-loop flapping wing design. *arXiv preprint arXiv:1910.13790*.
- Rousseeuw, P. J. (1987). Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65.
- Rubanova, Y., Chen, R. T., and Duvenaud, D. (2019). Latent odes for irregularly-sampled time series. *arXiv preprint arXiv:1907.03907*.
- Rubenstein, M., Cornejo, A., and Nagpal, R. (2014). Programmable self-assembly in a thousand-robot swarm. *Science*, 345(6198):795–799.
- Rudy, S. H., Brunton, S. L., Proctor, J. L., and Kutz, J. N. (2017). Data-driven discovery of partial differential equations. *Science Advances*, 3(4):e1602614.
- Runge, G., Wiese, M., Günther, L., and Raatz, A. (2017). A framework for the kinematic modeling of soft material robots combining finite element analysis and piecewise constant curvature kinematics. In *2017 3rd International Conference on Control, Automation and Robotics (ICCAR)*, pages 7–14. IEEE.
- Rus, D. and Tolley, M. T. (2015). Design, fabrication and control of soft robots. *Nature*, 521(7553):467–475.
- Rus, D. and Tolley, M. T. (2018). Design, fabrication and control of origami robots. *Nature Reviews Materials*, 3(6):101.

- Saar, K. A., Giardina, F., and Iida, F. (2018). Model-free design optimization of a hopping robot and its comparison with a human designer. *IEEE Robotics and Automation Letters*, 3(2):1245–1251.
- Safonov, A. A. (2018). Computing via natural erosion of sandstone. *International Journal of Parallel, Emergent and Distributed Systems*, 33(6):742–751.
- Sakagami, Y., Watanabe, R., Aoyama, C., Matsunaga, S., Higaki, N., and Fujimura, K. (2002). The intelligent ASIMO: system overview and integration. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 3, pages 2478–2483. IEEE.
- Sano, M. and Tamai, K. (2016). A universal transition to turbulence in channel flow. *Nature Physics*, 12(3):249.
- Schmid, P. J. (2010). Dynamic mode decomposition of numerical and experimental data. *Journal of fluid mechanics*, 656:5–28.
- Schmidt, M. and Lipson, H. (2009). Distilling free-form natural laws from experimental data. *Science*, 324(5923):81–85.
- Schmitt, F., Piccin, O., Barbé, L., and Bayle, B. (2018). Soft robots manufacturing: a review. *Frontiers in Robotics and AI*, 5:84.
- Schulz, A., Sung, C., Spielberg, A., Zhao, W., Cheng, R., Grinspun, E., Rus, D., and Matusik, W. (2017). Interactive robogami: an end-to-end system for design of robots with ground locomotion. *The International Journal of Robotics Research*, 36(10):1131–1147.
- Scimeca, L., Hughes, J., Maiolino, P., and Iida, F. (2019). Model-free soft-structure reconstruction for proprioception using tactile arrays. *IEEE Robotics and Automation Letters*, 4(3):2479–2484.
- Settles, B. (2009). Active learning literature survey.
- Shigemi, S., Goswami, A., and Vadakkepat, P. (2018). ASIMO and humanoid robot research at honda. *Humanoid Robotics: A Reference*, pages 55–90.
- Shih, B., Shah, D., Li, J., Thuruthel, T. G., Park, Y.-L., Iida, F., Bao, Z., Kramer-Bottiglio, R., and Tolley, M. T. (2020). Electronic skins and machine learning for intelligent soft robots. *Science Robotics*, 5(41):eaaz9239.
- Silva, D. F. and Maciel, A. (2012). A comparative study of physics engines for modeling soft tissue deformation. In *2012 XXXVIII Conferencia Latinoamericana En Informatica (CLEI)*, pages 1–7. IEEE.
- Sims, K. (1994). Evolving 3D morphology and behaviour by competition. *Artificial Life*, 1(4):353–372.
- Skews, B. W. (1990). Autorotation of rectangular plates. *Journal of Fluid Mechanics*, 217:33–40.
- Smith, E. (1971). Autorotating wings: an experimental investigation. *Journal of Fluid Mechanics*, 50(3):513–534.

- Snoek, J., Larochelle, H., and Adams, R. P. (2012). Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, pages 2951–2959.
- Soldatova, L. N., Clare, A., Sparkes, A., and King, R. D. (2006). An ontology for a robot scientist. *Bioinformatics*, 22(14):e464–e471.
- Sparkes, A., Aubrey, W., Byrne, E., Clare, A., Khan, M. N., Liakata, M., Markham, M., Rowland, J., Soldatova, L. N., Whelan, K. E., et al. (2010a). Towards robot scientists for autonomous scientific discovery. *Automated Experimentation*, 2(1):1–11.
- Sparkes, A., King, R. D., Aubrey, W., Benway, M., Byrne, E., Clare, A., Liakata, M., Markham, M., Whelan, K. E., Young, M., et al. (2010b). An integrated laboratory robotic system for autonomous discovery of gene function. *JALA: Journal of the Association for Laboratory Automation*, 15(1):33–40.
- Spong, M. W. (1999). Passivity based control of the compass gait biped. *IFAC Proceedings Volumes*, 32(2):506–510.
- Spong, M. W., Holm, J. K., and Lee, D. (2007). Passivity-based control of bipedal locomotion. *IEEE Robotics & Automation Magazine*, 14(2):30–40.
- Steels, L. and Brooks, R. (2018). *The Artificial Life Route to Artificial Intelligence: Building Embodied, Situated Agents*. Routledge.
- Stringham, G. E., Simons, D. B., and Guy, H. P. (1969). *The Behaviour of Large Particles Falling in Quiescent Liquids*. US Government Printing Office.
- Strogatz, S. H. (2018). *Nonlinear Dynamics And Chaos: With Applications To Physics, Biology, Chemistry, And Engineering*. CRC Press.
- Sun, Y., Song, Y. S., and Paik, J. (2013). Characterization of silicone rubber based soft pneumatic actuators. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4446–4453. IEEE.
- Suthakorn, J., Cushing, A. B., and Chirikjian, G. S. (2003). An autonomous self-replicating robotic system. In *Proceedings 2003 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM 2003)*, volume 1, pages 137–142. IEEE.
- Tanabe, Y. and Kaneko, K. (1994). Behavior of a falling paper. *Physical Review Letters*, 73(10):1372–1375.
- Tang, H., Houthoof, R., Foote, D., Stooke, A., Chen, X., Duan, Y., Schulman, J., De Turck, F., and Abbeel, P. (2017). # exploration: a study of count-based exploration for deep reinforcement learning. In *31st Conference on Neural Information Processing Systems (NIPS)*, volume 30, pages 1–18.
- Tarapore, D. and Mouret, J.-B. (2014). Comparing the evolvability of generative encoding schemes. In *Artificial Life Conference Proceedings 14*, pages 55–62. MIT Press.
- Taylor, C. and Jefferson, D. (1993). Artificial life as a tool for biological inquiry. *Artificial Life*, 1(1–2):1–13.

- Taylor, T. (2004). Redrawing the boundary between organism and environment. In Pollack, J., Bedau, M. A., Husbands, P., Watson, R. A., and Ikegami, T., editors, *Artificial Life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems*, pages 268–273. MIT Press.
- Tedrake, R., Manchester, I. R., Tobenkin, M., and Roberts, J. W. (2010). LQR-trees: feedback motion planning via sums-of-squares verification. *The International Journal of Robotics Research*, 29(8):1038–1052.
- Thompson, K. (1968). Programming techniques: Regular expression search algorithm. *Communications of the ACM*, 11(6):419–422.
- Toner, J., Tu, Y., and Ramaswamy, S. (2005). Hydrodynamics and phases of flocks. *Annals of Physics*, 318(1):170–244.
- Trianni, V. and López-Ibáñez, M. (2015). Advantages of task-specific multi-objective optimisation in evolutionary robotics. *PLoS One*, 10(8):e0136406.
- Trinh, E. (1985). Compact acoustic levitation device for studies in fluid dynamics and material science in the laboratory and microgravity. *Review of Scientific Instruments*, 56(11):2059–2065.
- Ueha, S., Hashimoto, Y., and Koike, Y. (2000). Non-contact transportation using near-field acoustic levitation. *Ultrasonics*, 38(1–8):26–32.
- Varshney, K., Chang, S., and Wang, Z. J. (2013). Unsteady aerodynamic forces and torques on falling parallelograms in coupled tumbling-helical motions. *Physical Review E*, 87(5):053021.
- Vasilevich, A. and de Boer, J. (2018). Robot-scientists will lead tomorrow’s biomaterials discovery. *Current Opinion in Biomedical Engineering*, 6:74–80.
- Veenstra, F., Jørgensen, J., and Risi, S. (2018). Evolution of fin undulation on a physical knife-fish-inspired soft robot. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 157–164.
- Vergara, A., Lau, Y.-S., Mendoza-Garcia, R.-F., and Zagal, J. C. (2017). Soft modular robotic cubes: toward replicating morphogenetic movements of the embryo. *PLoS One*, 12(1):e0169179.
- Vincent, L., Shambaugh, W. S., and Kanso, E. (2016). Holes stabilize freely falling coins. *Journal of Fluid Mechanics*, 801:250–259.
- von Mammen, S. and Jacob, C. (2007). Genetic swarm grammar programming: ecological breeding like a gardener. In *2007 IEEE Congress on Evolutionary Computation*, pages 851–858. IEEE.
- Vujovic, V., Rosendo, A., Brodbeck, L., and Iida, F. (2017). Evolutionary developmental robotics: improving morphology and control of physical robots. *Artificial Life*, 23(2):169–185.

- Wallin, T., Pikul, J., and Shepherd, R. (2018). 3D printing of soft robotic systems. *Nature Reviews Materials*, 3(6):84–100.
- Waltham, C., Bendall, S., and Kotlicki, A. (2003). Bernoulli levitation. *American Journal of Physics*, 71(2):176–179.
- Waltz, D. and Buchanan, B. G. (2009). Automating science. *Science*, 324(5923):43–44.
- Wang, W. B., Hu, R. F., Xu, S. J., and Wu, Z. N. (2013). Influence of aspect ratio on tumbling plates. *Journal of Fluid Mechanics*, 733:650–679.
- Webster III, R. J. and Jones, B. A. (2010). Design and kinematic modeling of constant curvature continuum robots: a review. *The International Journal of Robotics Research*, 29(13):1661–1683.
- Wehner, M., Truby, R. L., Fitzgerald, D. J., Mosadegh, B., Whitesides, G. M., Lewis, J. A., and Wood, R. J. (2016). An integrated design and fabrication strategy for entirely soft, autonomous robots. *Nature*, 536(7617):451–455.
- Werfel, J., Petersen, K., and Nagpal, R. (2014). Designing collective behaviour in a termite-inspired robot construction team. *Science*, 343(6172):754–758.
- Whitesides, G. M. (2018). Soft robotics. *Angewandte Chemie International Edition*, 57(16):4258–4273.
- Wiener, N. (1965). *Cybernetics, Second Edition: Or the Control and Communication in the Animal and the Machine*. The MIT Press.
- Willmarth, W. W., Hawk, N. E., and Harvey, R. L. (1964). Steady and unsteady motions and wakes of freely falling disks. *Physics of Fluids*, 7(2):197.
- Xie, W. and Wei, B. (2001). Parametric study of single-axis acoustic levitation. *Applied Physics Letters*, 79(6):881–883.
- Yang, Y., Ma, Z., Nie, F., Chang, X., and Hauptmann, A. G. (2015). Multi-class active learning by uncertainty sampling with diversity maximization. *International Journal of Computer Vision*, 113(2):113–127.
- Yarbasi, E. Y. and Samur, E. (2018). Design and evaluation of a continuum robot with extendable balloons. *Mechanical Sciences*, 9(1):51–60.
- Zappetti, D., Mintchev, S., Shintake, J., and Floreano, D. (2017). Bio-inspired tensegrity soft modular robots. In *Conference on Biomimetic and Biohybrid Systems*, pages 497–508. Springer.
- Zhakypov, Z. and Paik, J. (2018). Design methodology for constructing multimaterial origami robots and machines. *IEEE Transactions on Robotics*, 34(1):151–165.
- Zhang, J., Zhou, M., Huang, Y., Ren, P., Wu, Z., Wang, X., and Zhao, S. F. (2017). A smoothed finite element-based elasticity model for soft bodies. *Mathematical Problems in Engineering*, 2017:1467356.

- Zhang, Z., Dequidt, J., Kruszewski, A., Largilliere, F., and Duriez, C. (2016). Kinematic modeling and observer based control of soft robot using real-time finite element method. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5509–5514. IEEE.
- Zheng, G., Goury, O., Thieffry, M., Kruszewski, A., and Duriez, C. (2019). Controllability pre-verification of silicone soft robots based on finite-element method. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 7395–7400. IEEE.
- Zhong, H., Chen, S., and Lee, C. (2011). Experimental study of freely falling thin disks: transition from planar zigzag to spiral. *Physics of Fluids*, 23(1):011702.
- Zhong, H., Lee, C., Su, Z., Chen, S., Zhou, M., and Wu, J. (2013). Experimental investigation of freely falling thin disks. Part 1. The flow structures and reynolds number effects on the zigzag motion. *Journal of Fluid Mechanics*, 716:228.

