Morphologically programming the interactions of V-shaped falling papers

Toby Howison,¹ Josie Hughes¹ and Fumiya Iida¹

¹University of Cambridge, UK [th533, jaeh2, fi224]@cam.ac.uk

Abstract

The behavioural diversity seen in biological systems is, at the most basic level, driven by interactions between physical materials and their environment. In this context, we investigate the a-life properties of falling paper systems, in which different paper shapes are dropped into free fall and their behaviours observed. These systems have a simple embodiment but highly complex interactions with the environment. Using a synthetic methodology, i.e. understanding by building, we explore how morphology can be used to program certain interactions into the dynamics of a free-falling V-shaped paper. We demonstrate that morphology can encode a stochastic hierarchy of possible behaviours into the system. This hierarchy can be described by a set of conditional switching probabilities and represented in a morphological 'state machine'. We draw a parallel with developmental processes, showing how these can emerge from interaction with the environment. Next, we demonstrate how Bayesian optimisation can be used to optimise morphology in response to a fitness function, in this case minimizing falling speed. Bayesian optimisation allows us to capture the system stochasticity with minimal sampling. By manipulating non-living raw materials such as paper, we are able to analyse how morphology can be used to control and program interactions with the environment. With this bottom-up approach we ultimately aim to demonstrate principles that turn materials into agents that show non-trivial behaviours comparable to those of living organisms.

Introduction

Living systems exhibit extraordinary levels of autonomy and behavioural diversity. This is driven by the embodied interaction between the physical and informational worlds (Pfeifer et al., 2007). At the most basic level, however, behavioural diversity emerges from the interaction of physical materials and components. Hence, investigating how nontrivial behaviours emerge from the interaction between morphology and the environment is a key focus for the design of embodied artificial life-forms.

A particular question is how to design such interactions in a way that gives rise to useful behaviours. There is a growing interest in designing such life-forms over multiple levels of abstraction, from materials to components to robots (Howard et al., 2019). This parallel evolution could facilitate the wide-scale adaptation to environmental niches, with behavioural emergence driven by physical interaction with the environment forming one of the most basic building blocks of this process. There is also increasing interest in the use of data-driven modelling to resolve the reality-gap problem (Mouret and Chatzilygeroudis, 2017). Synthetic methodologies, i.e. understanding by building, have seen significant up-scaling in recent years (Howison et al., 2020; Vujovic et al., 2017; Saar et al., 2018), providing large amounts of ground-truth data on which to train models.

In this context, we are particularly interested in systems with a simple embodiment but complex interactions. Such systems represent the emergence of interaction driven behavioural diversity at a fundamental level. We focus specifically on 'falling paper systems' (Pesavento and Wang, 2004; Tanabe and Kaneko, 1994; Zhong et al., 2011). By manipulating a simple raw material like paper and releasing it into an environment, we can create systems that exhibit properties more complex than might be anticipated from their individual elements. The dynamics of falling paper shapes are driven by highly non-linear interactions with the environment. Yet, at a high level these dynamics manifest as coherent, non-trivial behaviours.

In this paper we explore the properties of falling paper from an a-life perspective. We revisit the V-shaped falling paper system (Howison et al., 2019), in which a 'V' paper shape with an affixed mass is dropped (Fig.1). We investigate this system at the organism level by analysing how falling behaviours can transition between states. By formalising these transitions in a behavioural 'state' diagram we can analyse how morphology encodes developmental-like behaviours into the system. Next, we analyse the system at the population level and use Bayesian optimisation to find a morphology that minimizes falling speed (similar to flying seeds). Bayesian optimisation is highly interesting from an a-life perspective as it offers a structured way to design embodied artificial lifeforms in the real-world.



Figure 1: Schematic of V-shaped falling paper (VSFP) system, adapted from Howison et al. (2019). (a) A paper V-shape is defined by four morphological parameters. m and w are fixed at 10mm and 5g, while l and θ may vary. (b) When dropped from a height of 3m there are four distinct falling behaviours: (1) plummeting; (2) undulation; (3) helicopter rotation; (4) asymmetric rotation. (c) The parameter space can be segmented into areas of behavioural dominance.

Falling paper

'Falling paper systems' are the focus of our research. When different paper shapes (and those of other materials) are released into free fall, a wide range of distinct falling behaviours is observed. For example, circular shapes manifest steady, chaotic or tumbling behaviours depending on morphological and environmental factors (Field et al., 1997). Understanding falling paper systems has been a goal since the time of Maxwell, who first proposed the problem in the 1800s (Maxwell, 1854; Finn, 2007). Usually researched by fluid dynamicists, physicists and mathematicians, the topic can offer a lot to the a-life community.

Physical embodiment Falling behaviours are grounded in the physical embodied world. Paper shapes can be defined by a range of morphological parameters such as density, elasticity and shape. Meanwhile, the environment is characterised by various properties including viscosity, temperature and the presence of air flows. When paper shapes are released into the environment the combination of these

properties induces highly complex interactions, from which emerge stable and distinct behaviours. As discussed, the most basic emergence of behavioural diversity is driven by embodied physical interactions. Indeed, the dynamics of falling papers are very similar to those of stable insect flapping behaviours (Bergou et al., 2007).

Behavioural stochasticity Falling paper systems are nondeterministic. The same shape dropped with the same initial conditions can exhibit completely different dynamics. However, there is structure in the system in the form of behavioural attractors. Certain shapes are more likely to converge to certain behaviours, and there is a hierarchy of likely behaviours a shape could exhibit. Further still, this hierarchy is intrinsically coupled with both the morphology and the environment. This raises interesting questions about behavioural programmability and morphological computation (Müller and Hoffmann, 2017; Nakajima et al., 2013).

Energetics The system is energetically non-conservative, converging to a 'dead state' in the absence of an external energy source. This energetic transience is analogous to biological systems, and in the context of a-life offers a framework to systematically study the interaction driven emergence of dynamic developmental processes including growth, adaptation, resource utilisation and evolution (Taylor et al., 2016).

Complexity The complexity of falling paper systems can be arbitrarily scaled up. We can introduce almost infinite richness into the environment by simply adding various air flows. Simultaneously, we can increase the morphological complexity by altering properties such as shape and porosity, introducing smart materials, or by combining different materials to create hybrid structures. By increasing the degrees-of-freedom within the system, we can increase the complex-ity of interactions and ultimately the system's behavioural diversity. Modulating system complexity may allow us to improve the performance and scope of evolutionary processes (Mitchell et al., 1993; Walker, 2017).

Synthetic methodologies There is a growing interest in using a synthetic methodology for designing embodied a-life agents, for example using evolution (Vujovic et al., 2017; Nygaard et al., 2018) or Bayesian optimisation (Saar et al., 2018; Rieffel and Mouret, 2018). This approach is well suited for falling paper, where the system behaviours are intrinsically linked to the complex interaction with the environment. Conventional modelling is unable to capture this interaction. Luckily, experiments can be carried out quickly and with minimal equipment: paper, cutting tool and a camera. When combined with modern robotic automation, we

have already demonstrated how many hundreds of experiments can be carried out (Howison et al., 2020).

The V-shaped falling paper system

As our case study we revisit the V-shape falling paper (VSFP) system. A V-shaped paper shape with an affixed mass is dropped from a height. The system morphology is fully defined by the four parameters shown in Fig 1a: 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 10mm and 5g. The VSFP system is inspired by plant dispersal mechanisms, where seeds have evolved to fly as far away from the parent tree as possible (Lee and Choi, 2017).

In our previously published study (Howison et al., 2019) we reported on the behavioural diversity seen within the VSFP parameter space. There are four distinct falling behaviours: plummeting, undulating, helicopter rotation and asymmetric rotation (Fig. 1b). As the morphological parameters l and θ vary, so the particular behaviour on which the systems is likely to settle also changes. Figure 1c shows how the design parameter space can be segmented into areas of different dominant behaviours. This representation reveals the population-level structure within the parameter space. However, as mentioned falling paper systems are stochastic so these boundaries are approximations, not absolute rules. The VSFP system is a conceptually simple system, with a design space of only 2 dimensions.

Experimental procedure

A detailed experimental procedure can be found in Howison et al. (2019). Here, we reproduce the relevant procedures for this paper. An Endurance MakeBlock XY engraving/cutting machine was used to cut shapes out of Silvine A4 Graph Refill paper. The paper has a weight of 80 grams per square metre. The affixed 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. Each shape, examples was manually dropped from a height of 3m into still air and using a tip up initial condition. Shapes fell against a black backdrop, and were recorded using a Logitech BRIO camera recording at 120 fps.

Behavioural analysis

The video data was viewed by a human observer to assign behavioural labels to each experiment. This includes the final behavioural state and the transient behaviours observed prior to convergence to a final state. Experiments were limited by the drop height, in this case 3m. Assigning labels to transient behaviours is challenging as the system can pass through multiple behaviours very fast. This human perception and labelling of behaviours is, in itself, an interesting aspect of falling paper systems.

The video data can also be used to extract the falling speed. For a particular experiment with morphological pa-

rameters $x = [l \ \theta]$, the falling time t(x) was manually extracted from the video data, and used to calculate the average falling speed v(x), e.g.

$$v(x) = \frac{h}{t(x)} \tag{1}$$

where h = 3 m is the drop height. Note that this falling speed is the average of transient and converged behaviours, rather than the falling speed of a particular behaviour.

Behavioural switching of falling paper

A key area of research in the future design of embodied agents is developmental processes, e.g. growth and learning Doursat and Sánchez (2014); Kriegman et al. (2018). These form a vital part of a biological systems life, and are a driving force behind the emergence of complex behavioural repertoires. In this context the transient behaviours of falling-paper systems are highly interesting. As papers fall their behaviours switch. Sometimes the behaviours converge, whereas other times switching dominates the fall. There seem to be behavioural 'attractors' to which morphologies are likely, but not guaranteed to converge, and these attractor states are intrinsically encoded into the morphology and environment of the system.

In the VSFP system we have observed that during its lifetime a falling paper can switch between multiple behaviours before reaching a final state (final being dictated by the available drop height). Figure 2a shows an example of the system switching from helicopter rotation to plummeting, while Figure 2b shows switching from undulating to helicopter rotation.

We draw a 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. 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 the context of a-life this allows the investigation of how physical interaction with the environment can induce very basic developmental processes.

Behavioural complexity

We can quantify the complexity of the transient processes encoded into different morphologies by counting the number of behaviour switches in each experiment. For example, some areas of the design space quickly converge to the final state while others may cycle through one or more alternative behaviour states before convergence. The parameter space was discretized into a 7×7 grid. The morphologies at each point were fabricated and then tested five times. We analysed the data-set of 245 experiments to calculated the mean



Figure 2: Examples of behavioural transitions in the VSFP system, showing (a) switching from helicopter rotation to plummeting and (b) undulating to helicopter rotation.

number of switches at each point in the 7×7 discretized parameter space. Using linear interpolation we extended this to estimate the mean number of behavioural switches anywhere in the parameter space. Figure 3 shows this, along with the behavioural boundary lines from Figure 1c.

We see that the mean number of behaviours does not correlate with the boundaries of the dominant behavioural regions. Plummeting and undulating behaviours seem to manifest simpler developmental processes, whereas the helicopter rotation behaviour exhibits the most complex switching process, with a maximum mean of 2.6 switches. This analysis offers an alternative perspective on the VSFP system. Rather than segmenting the design parameter space by dominant behaviours, we can also segment it by the complexity of the transient processes. Returning to the parallel with developmental processes, we see that morphology directly influences the system's capacity to exhibit basic growth and adaptation.

Conditional behavioural switching

The number of behavioural switches across the parameter space provides a high-level understanding of the transient behaviours in the system. We can extend our analysis by looking at the specific structure of transient processes that lead to a particular behaviour. The switching we see between behaviours is conditional, e.g. each behaviour is dependent on the previous states. Furthermore, certain states can only be reached dependent on previous conditions or states. We term this transitory behaviour 'conditional switching'. This framework allows us to identify global rules that control be-



Figure 3: The mean number of behavioural switches across the parameter space along with dominant behavioural boundaries.

havioural switching and to identify attractor states that have a high probability of emergence, and may therefore be more energetically favourable. This 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).

To formalise this we use the idea of state machines, an abstract concept whereby a machine can have different states, but at a given time fulfills only one of them. We can define the conditional probability of switching between different states, and can represent these as a state diagram. We analysed the transient behaviours of morphologies that tend to converge to helicopter or asymmetric rotation behaviours in the VSFP system. By calculating the frequency and direction of switching, we constructed a conditional behavioural switching diagram indicating the possible switching events in the system (Fig. 4). The diagram indicates the conditional probability of any state switching to any other.

This representation allows us to visualise the paths within the behaviour space that can be taken between the initial condition and end state. It summarises the behaviours of a range of morphologies with a common behavioural end state attractor. We can see a hierarchy of attractor states, and these are completely determined by the mechanical design and interaction with the environment. It also provides a structure for the design of certain behaviours and developmental sequences. For the helicopter diagram (Fig.4a), we see that it is highly unlikely to sustain a plummeting behaviour, but that it is relatively likely to observe a helicopter or asymmetric rotation behaviour. Meanwhile, the asymmetric rotation (Fig.4b) shows a similar probability of sustaining a plum-



Figure 4: Stochastic behavioural 'state machine' showing the conditional likelihood for behavioural transitions for (a) helicopter and (b) asymmetric rotation behaviours.

meting behaviour, but different likelihoods of switching between other behaviours. Generally, the there seems to be more switching in paths in the helicopter than asymmetric rotation diagrams.

Behavioural programmability

The simplicity of falling paper systems allows us to explore the design of real-world behaviours. As we have discussed, the VSFP system is highly complex, with a range of conditional behavioural hierarchies across the design parameter space. This stochastic nature makes design challenging, as similar designs may appear to perform very differently. In simulation we could overcome this by using evolutionary algorithms with many thousands of iterations, e.g. Cheney et al. (2014). However when we rely on synthetic methodologies this isn't feasible.

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 blackbox 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 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 (Saar et al., 2018; Rosendo et al., 2017).

We aim to optimise the falling speed v in the VSFP system by using Bayesian optimisation to search for a morphology that minimizes it. This problem is directly bio-inspired. Certain seeds in nature have evolved to fall slowly, allowing them to travel away from the parent tree. 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 same behaviours.

Bayesian optimisation

We start by defining the minimisation problem

$$\min_{x \in A} f(x) \tag{2}$$

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 v 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.



Figure 5: 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.

At any given stage in the optimisation, assume the objective function has been evaluated n times at the points $x_{1:n} = [x_1, \ldots, x_n]^T$ and function values $f_{1:n} = [f(x_1), \ldots, f(x_n)]^T$. 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 **K**, for which we use the Matern kernel with a smoothness parameter of 5/2,

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

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

$$\mathbf{K} = \begin{bmatrix} k(x_1, x_1) & k(x_1, x_2) & \dots & k(x_1, x_n) \\ k(x_2, x_1) & k(x_2, x_2) & \dots & k(x_2, x_n) \\ \vdots & \vdots & \ddots & \vdots \\ k(x_n, x_1) & k(x_n, x_2) & \dots & k(x_n, x_n) \end{bmatrix}$$
(4)

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:

$$\mu(x_k) = \mathbf{k}_k^T \mathbf{K}^{-1} f_{1:n} \tag{5}$$

$$\sigma^2(x_k) = -\mathbf{k}_k^T \mathbf{K}^{-1} \mathbf{k}_k \tag{6}$$

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

$$EI(x_k) = E[\max(0, \mu(x_{\text{best}}) - f(x_k))$$

$$|x_k \sim N(\mu(x_k), \sigma^2(x_k))]$$
(7)

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

$$x_{n+1} = \operatorname{argmax} EI(x) \tag{8}$$

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.

Optimisation results

We implemented the BOA using the MATLAB bayesopt function. We operated the algorithm 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 vwas extracted from the video data and supplied to the BOA.

The optimal solution has morphological parameters l = 82.9mm and $\theta = 51.1^{\circ}$, a falling speed of v = 0.933m/s

and corresponded to an asymmetric falling behaviour. Figure 5 shows various results from the optimisation process. The best observed falling speed after each iterations (the blue line in Fig.5) falls quickly from around 4m/s to around 1.5 m/s, after which it slowly decreases to below 1m/s. 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 (Fig.1c) corresponding to each iteration are primarily helicopter and asymmetric rotation.

Figure 6a shows the fitness landscape as estimated by the GPR model, along with the sampled points. The behavioural boundaries from (Fig.1c) are also shown. We see that the majority of sample evaluations occurred in the rotation behavioral regions of the parameter space. We would expect this as these clearly have a lowed 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 3 we see that this area corresponds to an average number of behavioural switches of around 1.5, generally lower than the helicopter rotation behaviours. As mentioned earlier, the behaviour state diagrams (Fig. 4) indicate there are more switching events in the helicopter rotation system than the asymmetric system. This adds further explanation for the results of the optimisation process.

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, 2011).

Discussion and conclusion

This paper has presented the various a-life like properties of falling paper systems, specifically the VSFP system. We explored how morphology can be used to program certain interactions into the behaviours of falling-paper. The context of this study lies in the apparent reliance on the interactions between physical components and the environment, on all levels, for the emergence embodied behavioural diversity seen in nature. Falling paper systems represent the fundamental emergence of behaviours at the organism level and structure at the population level.

A key finding is demonstrating the stochastic behavioural patterns based on embodied interactions. The number of behavioural switches a falling paper is likely to exhibit depends strongly on its morphological parameters. Furthermore, for a particular dominant behaviour we can construct a probabilistic state transition diagram to visualise this structure. We draw a parallel with developmental processes. The 'de-



Figure 6: 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.

velopment' of behaviours in the falling paper system can be programmed within the morphology. The relationship between and embodied interaction and emergent developmental processes is an interesting field of further research.

We also demonstrated the use of Bayesian optimisation for designing falling behaviours. Bayesian optimisation offers a structured framework for design optimisation in realworld systems, and is well suited to handle the stochasticity of falling paper systems. In the context of designing embodied artificial lifeforms, Bayesian optimisation demonstrates the power of data-driven modelling. The GPR model can capture complex embodied interactions via the function approximation that maps design inputs to fitness outputs. Despite the success of this approach, it is clear that to fully harness the power of embodied interactions in this systems requires a search process that can take account of the behavioural diversity and transience. Bayesian optimisation, for example, is not adaptive to changes in fitness function. Growing research into novelty search algorithms is another key area to research (Mouret and Clune, 2015).

Acknowledgements

This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grant RG92738 for the University of Cambridge Centre for Doctoral Training and The Mathworks.

References

- Bergou, A. J., Xu, S., and Wang, Z. J. (2007). Passive wing pitch reversal in insect flight. *Journal of Fluid Mechanics*, 591:321–337.
- 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.

- Cheney, N., MacCurdy, R., Clune, J., and Lipson, H. (2014). Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. ACM SIGEVOlution, 7(1):11–23.
- Doursat, R. and Sánchez, C. (2014). Growing fine-grained multicellular robots. *Soft Robotics*, 1(2):110–121.
- Field, S. B., Klaus, M., Moore, M., and Nori, F. (1997). Chaotic dynamics of falling disks. *Nature*, 388(6639):252–254.
- Finn, D. L. (2007). Falling paper and flying business cards. SIAM News, 40(4):3.
- Frazier, P. I. (2018). A tutorial on bayesian optimization. arXiv preprint arXiv:1807.02811.
- Howard, D., Eiben, A. E., Kennedy, Danielle Frances fand 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., Giardina, F., and Iida, F. (2019). Physics driven behavioural clustering of free-falling paper shapes. *PloS one*, 14(6).
- Howison, T., Hughes, J., and Iida, F. (2020). Large-scale automated investigation of free-falling paper shapes via iterative physical experimentation. *Nature Machine Intelligence*, pages 1–8.
- Hughes, J., Maiolino, P., and Iida, F. (2018). An anthropomorphic soft skeleton hand exploiting conditional models for piano playing. *Science Robotics*, 3(25):eaau3098.
- Kriegman, S., Cheney, N., and Bongard, J. (2018). How morphological development can guide evolution. *Scientific reports*, 8(1):1–10.
- Lee, I. and Choi, H. (2017). Flight of a falling maple seed. *Physical Review Fluids*, 2(9):090511.
- Lehman, J. and Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223.
- Lizotte, D. J. (2008). *Practical bayesian optimization*. University of Alberta.
- Maxwell, J. C. (1854). On a particular case of the descent of a heavy body in a resisting medium. *Camb. Dublin Math. J*, 9:145–148.
- Mitchell, M., Hraber, P., and Crutchfield, J. P. (1993). Revisiting the edge of chaos: Evolving cellular automata to perform computations. *arXiv preprint adap-org/9303003*.
- 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*.
- 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.
- Nygaard, T. F., Martin, C. P., Samuelsen, 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.
- 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., Lungarella, M., and Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. *science*, 318(5853):1088–1093.
- Rasmussen, C. E. (2003). Gaussian processes in machine learning. In Summer School on Machine Learning, pages 63–71. Springer.
- Rieffel, J. and Mouret, J.-B. (2018). Adaptive and resilient soft tensegrity robots. *Soft robotics*, 5(3):318–329.
- 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).
- 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.
- Tanabe, Y. and Kaneko, K. (1994). Behavior of a falling paper. *Physical Review Letters*, 73(10):1372.
- Taylor, T., Bedau, M., Channon, A., Ackley, D., Banzhaf, W., Beslon, G., Dolson, E., Froese, T., Hickinbotham, S., Ikegami, T., et al. (2016). Open-ended evolution: Perspectives from the oee workshop in york. *Artificial life*, 22(3):408–423.
- 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.
- Walker, S. I. (2017). Origins of life: a problem for physics, a key issues review. *Reports on Progress in Physics*, 80(9):092601.
- 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.