

Large Scale Automated Investigation of Free-Falling Paper Shapes via Iterative Physical Experimentation

Toby Howison*, Josie Hughes, and Fumiya Iida

Department of Engineering, University of Cambridge, UK

Abstract

Free-falling paper shapes exhibit rich, complex and varied behaviours which are extremely challenging to model analytically. Physical experimentation aids in system understanding but is time-consuming, sensitive to initial conditions and reliant on subjective visual behavioural classification. In this study robotics, computer vision and machine learning are used to autonomously fabricate, drop, analyse and classify the behaviours of hundreds of shapes. The system is validated by reproducing results for falling disks, 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. Steady or periodic behaviours, however, 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 generalizes 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.

1 Introduction

Many natural phenomena such as autorotating seeds,^{1,2} Rayleigh-Taylor flow instabilities³ or chemical oscillations⁴ manifest spontaneous behaviours of considerable complexity. Driven entirely by their inherent non-linear dynamics,⁵⁻⁷ such systems often appear random in their evolution, with a multiplicity of states and possible behaviours. However, there is often underlying structure, for example different behaviour types in different, well-delineated, regions of a suitable non-dimensional parameter space. As well as the intrinsic scientific interest, understanding this structure can aid in the design of systems that exploit complex dynamics, such as for locomotion,⁸⁻¹³ information processing¹⁴ or computation.¹⁵

One example of interest is falling-paper systems. First proposed by James Clerk Maxwell in the 1800s,¹⁶ 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. 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.¹⁷⁻²² These behaviours can be characterised in a behavioural ‘phase diagram’ defined by the non-dimensional moment of inertia I^* and Reynolds number Re . This offers an easily interpretable insight into the relative physical importance of the fluid and body forces in different behavioural patterns. Other simple shapes have been studied, including rectangles,²³⁻²⁵ annuli²⁶ and parallelograms;²⁷ behavioural transitions have also been studied.²⁸⁻³⁰ More complex shapes have so far not been investigated. A number of heavily simplified analytical³¹ and numerical³² solutions have also been presented; however, the problem has evaded a comprehensive solution. Figure 1a summarizes this type of system, in which morphology maps to behaviours via the forward design problem, and similarly, behaviours map to morphology via the inverse design problem.

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 study we utilise 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,³³ with the added benefit of having much greater control over experimental initial conditions.³⁴ Meanwhile, machine learning offers continual,

*th533@cam.ac.uk

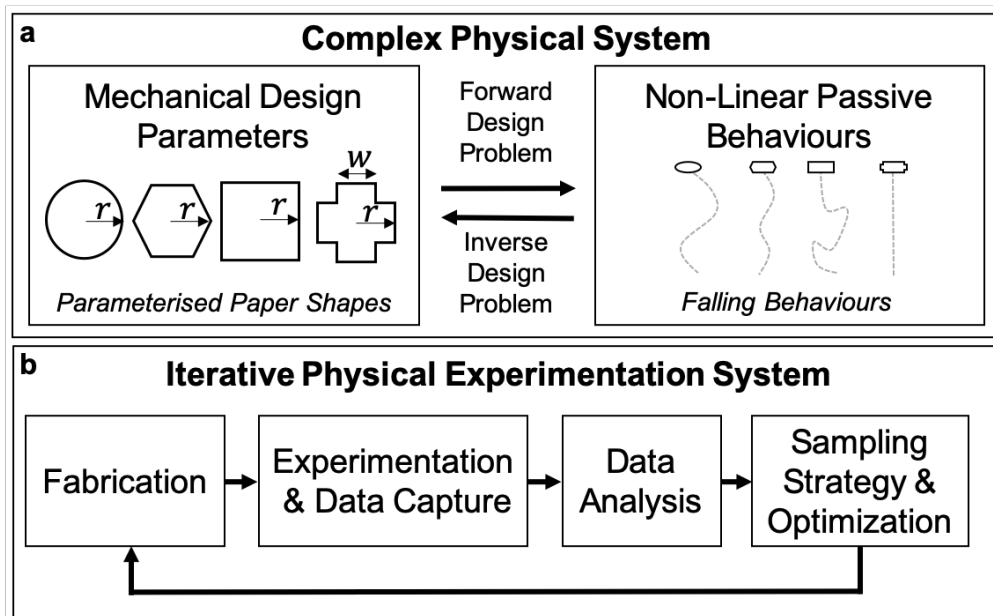


Figure 1: **Schematic of the falling-paper system and proposed iterative physical experimentation approach.** (a) Overview of complex physical systems such as the falling-paper problem. Mechanical design parameters map to behaviours via the forward design problem. Behaviours map to design parameters via the inverse design problem. (b) Flow chart of the iterative physical experimentation system, where paper shapes are fabricated, tested, analysed and sampled automatically.

online data analysis³⁵ to interpret data and suggest efficient sampling strategies.³⁶ The use of robotics in science research, the so-called ‘robot scientists’,^{37–39} is well documented. In biology, robots have been used to automate lab processes⁴⁰ and carry out 1000’s of experiments to efficiently identify gene functionality.^{41,42} When combined with artificial intelligence these systems can even automate hypothesis creation⁴³ and drive future discoveries.⁴⁴ Other methods for exploring high-dimensional spaces for scientific purposes range from reinforcement learning to guiding exploration using intrinsic motivation or entropy.^{45,46} Meanwhile, investigative methods for understanding non-linear dynamics in complex systems using experimental data⁴⁷ and extracting physically meaningful expressions have been presented.^{48–51} However, there has been limited application of ‘robot scientists’ to the broader physical sciences.

We propose an iterative physical experimentation system (IPES; Figure 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,¹⁷ 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 generalizes for the more complex cross, square and hexagon shapes, but that the transitional boundaries shift.

2 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. The experimental setup consists of a laser cutter to manufacture shapes, a UR5 robotic arm to pick and drop them and two high-speed cameras to record the falling behaviour (Figure 2, Supplementary Movie 1). This enables the 3D falling trajectory and observable to be calculated. The observable area is the amount of the shape visible to the camera at any time; it is smallest when the paper is perpendicular to the camera and a largest when the paper is parallel. Hence, it indicates the frequency

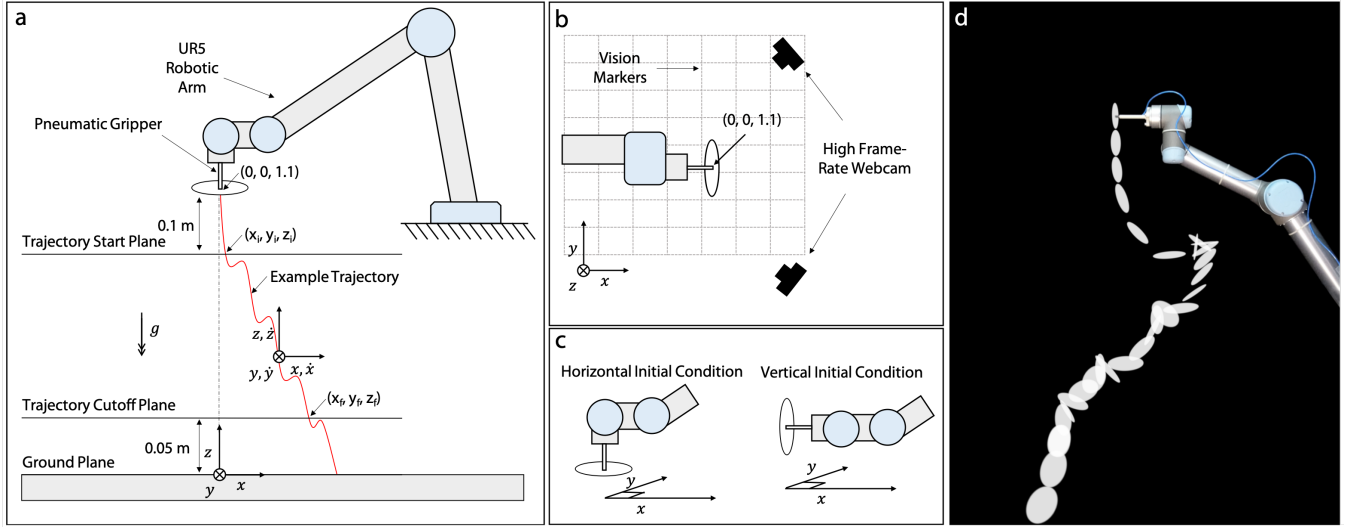


Figure 2: **Diagram of the experimental setup.** (a) Overview of the experimental system showing a UR5 robot arm dropping a paper shape, and the co-ordinate system (x, y, z) used to record the trajectory. (b) Plan view showing the camera setup used to capture 3D falling trajectories. (c) Initial conditions 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.

of shape rotation.

Four parameterised shapes were investigated; circles, to validate the set-up against previous studies,¹⁷ 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 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 1a (see also Supplementary Table 1, Supplementary Figure 1).

We used an adapted version of the behavioural groupings presented in for falling disks,¹⁷ namely:

1. Steady and periodic falling – the disk falls steadily or oscillates back and forth with a horizontal orientation.
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 3 and Supplementary Movie 2 show representative examples of these behaviours for each shape.

Behaviour-characterizing features such as vertical velocity, horizontal velocity and oscillation frequency were extracted from the 3D trajectory, for a full list see Supplementary Table 2. The automated behavioural classification system utilized 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. For classification accuracy and confusion matrices for each shape see Supplementary Table 3 and Supplementary Figure 2. The same classification scheme was used for all four shapes. However, the classification agreement for the cross shapes was significantly lower than for the other shapes, suggesting that alternative parameters may better describe the behaviours of cross shapes.

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 3a) where there is transition from steady falling to chaotic falling.

Figure 4 shows representative trajectories for each shape and behaviour. Chaotic motions were characterized by abrupt changes in vertical velocity in the 3D trajectory and apparently random rotation. Tumbling motions

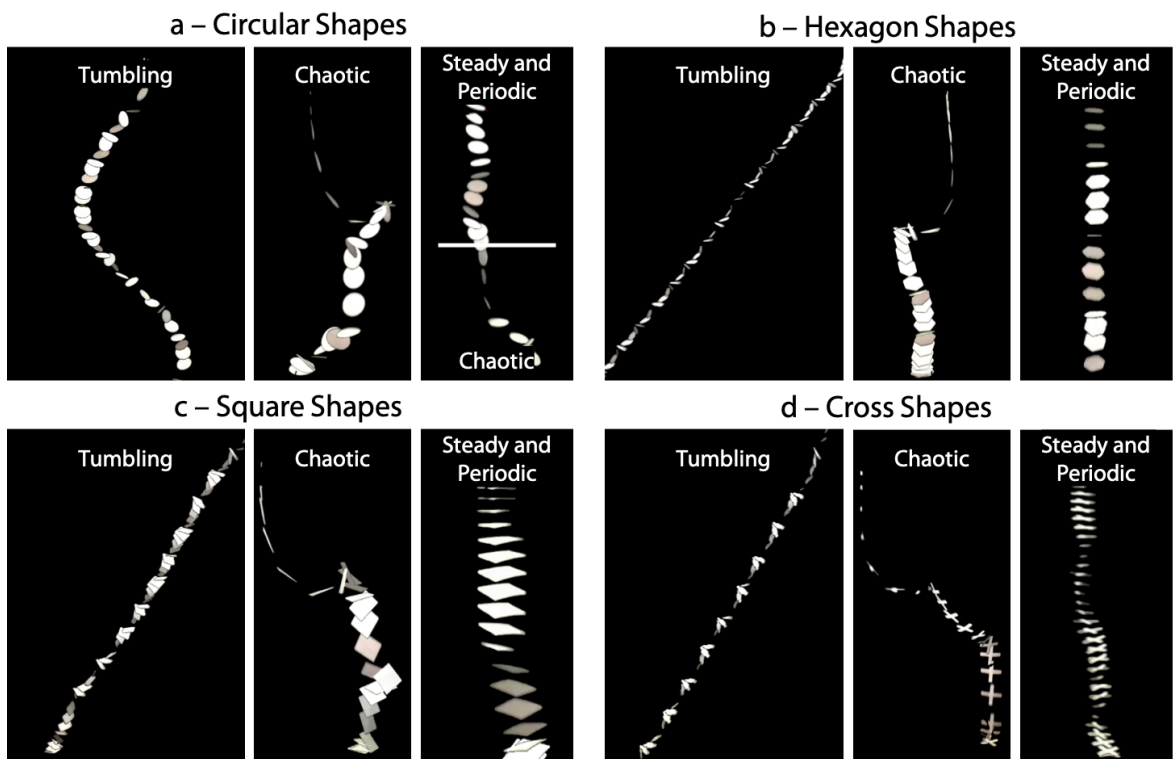


Figure 3: **Time-lapse images of the three falling behaviours observed in each shape** (a) Falling circular shapes, showing a behavioural transition from steady and periodic to chaotic (b) Falling hexagonal shapes (c) Falling square shapes (d) Falling cross shapes.

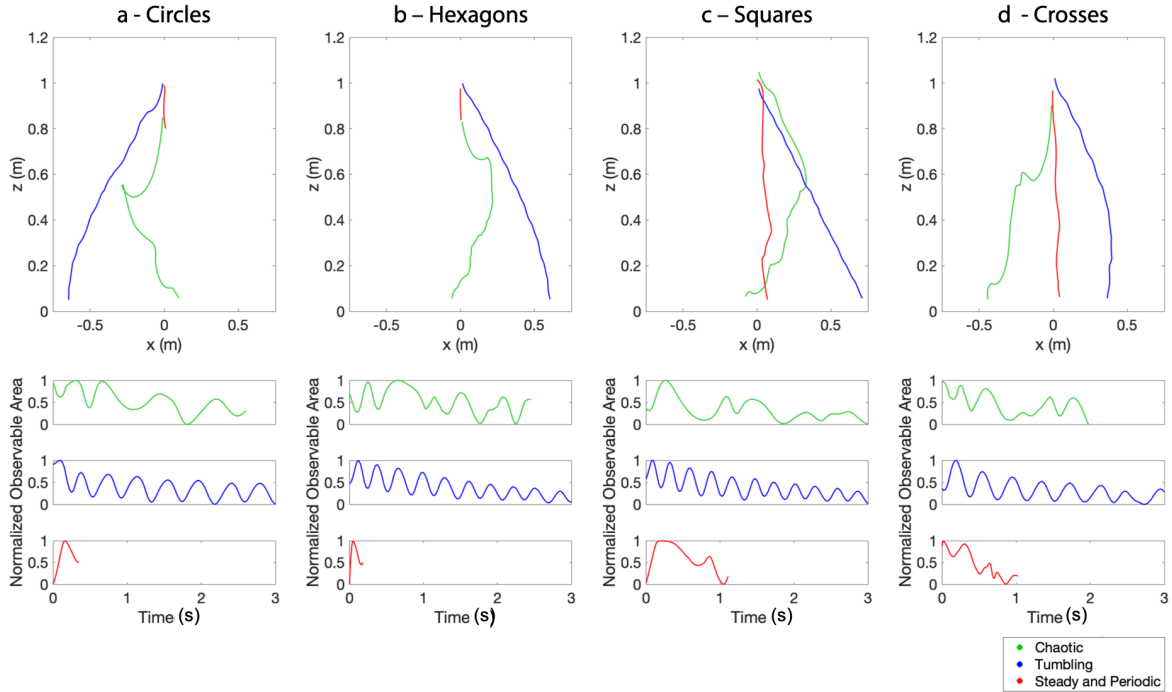


Figure 4: **Trajectories and observable area profiles of falling shapes for different automatically classified behavioural groups.** (a) Top panel shows typical trajectories for the three different behaviours observed for falling circles. The steady trajectory is seen for only the initial period of falling. 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.

were characterized by smooth trajectories and periodic rotation. Steady behaviours were characterized 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 4a,b.

3 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 5a, 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 optimizing 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 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 5b. 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.

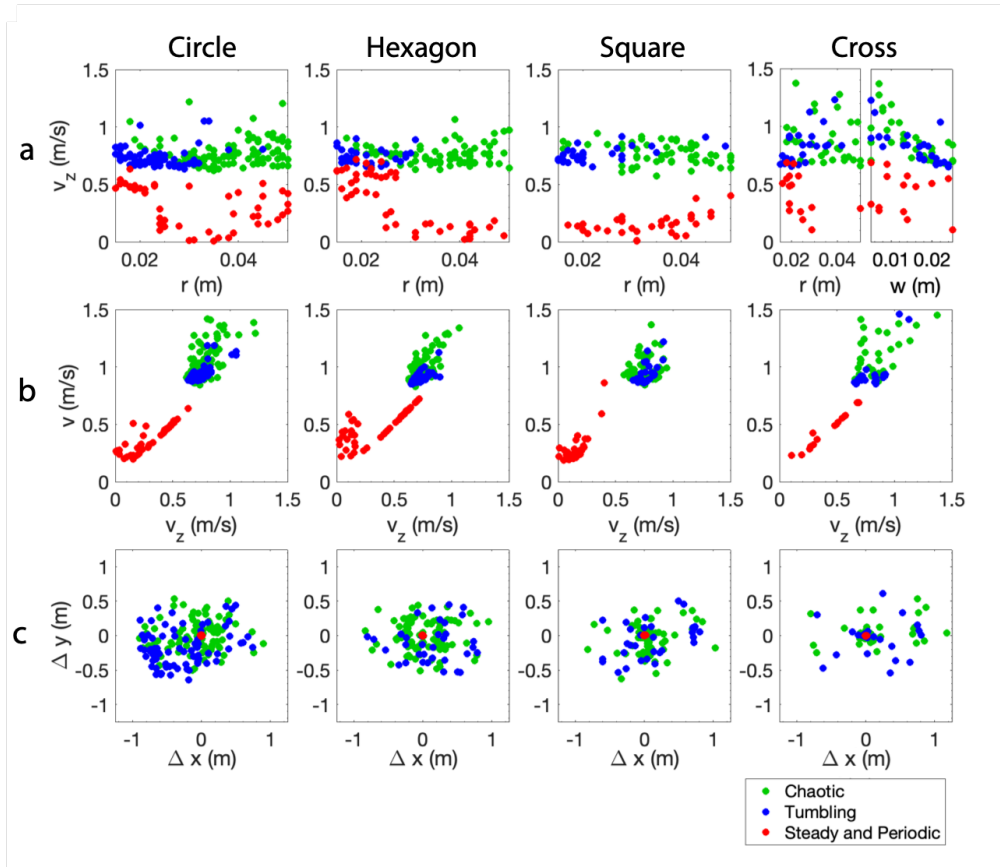


Figure 5: **Variation in measured falling characteristics with respect to design parameters and automatically classified behaviour groups.** (a) Vertical falling speed v_z vs. mechanical design parameters r and w (b) Path velocity v vs. vertical velocity v_z (c) Horizontal displacement Δx and Δy between start and finish of trajectory.

Figure 5c 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. For additional analysis of Δx and Δy relative to initial drop orientation see Supplementary Figure 6.

Although each of the features shown in Figure 5 show some ability to separate or uniquely define behaviours, there is no single output which allows full classification.

4 Behavioural Grouping

To understand how morphological and environmental parameters can influence falling behaviours, non-dimensional parameters have been used to construct behavioural ‘phase-diagrams’.¹⁷ Most commonly among these non-dimensional parameters are the Reynolds number (Re) and dimensionless moment of inertia (I^*), which characterize 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 of¹⁷ 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 6a–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




Figure 6: **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.

observed at the lower end of the Re and I^* range.^{17,21,22} 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.

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 6c–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 circle results increases as the shapes deviate from circular. Hence, the hexagon behaves most similarly, and the cross is least similar. The generalization of the $Re-I^*$ phase space to the hexagon and cross is a previously unreported finding. There may be more appropriate dimensionless quantities to characterize the behaviour of the falling cross.

5 Discussion & Conclusion

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 dynamics³⁹ 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.

In this work we present an approach for iterative physical experimentation to facilitate the scientific investigation of the relationship between mechanical design parameters and non-linear passive dynamics. We demonstrate this

approach by autonomously investigating the falling-paper problem. Understanding this problem requires large quantities of experimental data coupled with intelligent data analysis. In this work 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, the flexibility in morphology, experimental quantity and analysis speed were vastly increased, with experiments and analysis taking an average of just 90 seconds.

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

Previous studies¹⁷ 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 generalizes 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.⁴⁸

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 1.1m. Extending this would allow observation of longer trajectories or potential behavioural switching. In addition to the limitations of the physical system, the exploration strategy as outlined in Fig 1 is basic in this study. 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.

There is scope to investigate far more complex shapes and the role of deformation or even folding during falling. The system could also be used to optimise morphologies based on certain fitness functions, providing a practical and applied usage in design optimization. Indeed, the falling-paper problem is an ideal platform to test optimization strategies for noisy, real-world functions. Finally, further work could be done investigating the transitory behaviour of falling paper. There have been studies on exploring this,²⁹ but the IPES facilitates investigation of a much larger scale.

Finally, it is important to clarify the intended purposes of this work, and where it should be placed in the larger research context. 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.

6 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.

6.1 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.

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 pre-formed 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. For a more detailed set of system diagrams and processes see Supplementary Figures 3-4 .

The shape parameter ranges were chosen to explore a large design space without introducing manufacturing errors or very large deformations. Supplementary Table 1 summarises the parameterised paper shapes.

6.2 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 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. Supplementary Figure 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 5mm 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 $(90^\circ, 0, 0)$.

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 $(90^\circ, 0, 0)$ and $(0, 0, 0)$ 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 if we felt the laboratory conditions could affect this still-air, e.g. during construction when the curtain was being moved.

6.3 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 pre-defined 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 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. The parameters that were extracted are given in Supplementary Table 2. The extracted parameters facilitated automated behavioural classification.

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,²⁹ 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,⁵² 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. The authors have demonstrated this concept on other falling-paper systems.⁴⁸ 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 number of experiments for each shape is given in Supplementary Table 4. The K-Means clustering algorithm was applied to X ,

$$\mathbf{b} = \text{KMC}_N(\mathbf{X}) \quad (1)$$

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 categorized 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 minimized.⁴⁸

Previous work has identified that dimensionless quantities can be used to characterize 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

$$Re = \frac{v_z L}{\nu}, \quad (2)$$

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

$$I_{\text{disk}}^* = \frac{I_{\text{disk}}}{\rho_f d^5} \quad (3)$$

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

$$I_{\text{Hexagon}}^* = \frac{5mr^2}{12\rho_f d^5}, \quad (4)$$

$$I_{\text{Square}}^* = \frac{mr^2}{3\rho_f d^5}, \quad (5)$$

$$I_{\text{Cross}}^* = \frac{m(4r^2 + w^2)}{24\rho_f d^5}. \quad (6)$$

6.4 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. There is scope for implementing smart search and exploration methodologies including genetic algorithms⁵³ or MAP elites.⁵⁴ In particular, by combining automated behavioural classification with such search algorithms, there could be more focused exploration on behavioural boundaries. Supplementary Figure 7 shows the distribution of randomly sampled design parameters for each shape.

Author Correspondence All correspondence and requests for materials should be addressed to author Toby Howison.

Acknowledgements We acknowledge funding from the EPSRC RG92738 and The Mathworks Ltd.

Author Contributions TH, JH and FI conceived the study. TH designed the experimental set-up. TH and JH managed the experiments. TH and JH analysed the data. TH, JH and FI wrote the manuscript.

Competing Interests There are no competing interests.

Data availability Example data is available here <https://github.com/th533/Falling-Paper>.

Code Availability Example code is available here <https://github.com/th533/Falling-Paper>.

References

- ¹ Varshney, K., Chang, S., & Wang, Z.J. The kinematics of falling maple seeds and the initial transition to a helical motion. *Nonlinearity* **25**(1), C1 (2012).
- ² Norberg, R.A. Autorotation, self stability, and structure of single winged fruits and seeds (Samaras) with comparative remarks on animal flight. *Biological Reviews* **48**, 561–596 (1973).
- ³ Mikaelian, K.O. Analytic approach to nonlinear Rayleigh-Taylor and Richtmyer-Meshkov instabilities. *Physical Review Letters* **80**, 508–511 (1998).
- ⁴ Epstein, I.R. & Showalter, K. Nonlinear chemical dynamics: oscillations, patterns, and chaos. *Journal of Physical Chemistry* **100**, 13132–13147 (1996).
- ⁵ Nicolis, G. *Introduction to nonlinear science*. (Cambridge University Press, Cambridge, 1995).
- ⁶ May R.M. Simple mathematical models with very complicated dynamics. *Nature* **261**, 459–467 (1976).
- ⁷ Mustapha, H. & Dimitrakopoulos, R. High-order stochastic simulation of complex spatially distributed natural phenomena. *Mathematical Geosciences* **42**, 457–485 (2010).
- ⁸ Brodbeck, L., Hauser, S. & Iida, F. Morphological evolution of physical robots through model-free phenotype development. *PLOS ONE* **10**, e0128444 (2015).
- ⁹ Vujovic, V., Rosendo, A., Brodbeck, L., & Iida, F. Evolutionary developmental robotics: improving morphology and control of physical robots. *Artificial Life* **23**, 169–185 (2017).
- ¹⁰ Rieffel, J., Knox, D., Smith, S., & Trimmer, B. Growing and evolving soft robots. *Artificial Life* **20**, 143–162 (2014).
- ¹¹ Cheney, N., MacCurdy, R., Clune, J. & Lipson, H. Unshackling evolution. *ACM SIGEVOlution* **7**, 11–23 (2014).
- ¹² Rosendo, A., von Atzigen, M. & Iida, F. The trade-off between morphology and control in the co-optimized design of robots. *PLOS ONE* **12**, e0186107 (2017).
- ¹³ Saar, K.A., Giardina, F. & Iida, F. Model-free design optimization of a hopping robot and its comparison with a human designer. *IEEE Robotics and Automation Letters* **3**, 1245–1251 (2018).
- ¹⁴ Nakajima, K., Hauser, H., Li, T. & Pfeifer, R. Information processing via physical soft body. *Scientific Reports* **5**, 10487 (2015).
- ¹⁵ Hauser, H., Ijspeert, A.J., Fuchslin, R.M., Pfeifer, R. & Maass, W. Towards a theoretical foundation for morphological computation with compliant bodies. *Biological Cybernetics* **105**, 355–370 (2011).
- ¹⁶ Maxwell, J.C. On a particular case of the descent of a heavy body in a resisting medium. *Camb. Dublin Math.* **9**, 145–148 (1854).
- ¹⁷ Field, S.B., Klaus, M., Moore, M.G. & Nori, F. Chaotic dynamics of falling disks. *Nature* **388**, 252–254 (1997).
- ¹⁸ Zhong, H., Chen, S. & Lee, C. Experimental study of freely falling thin disks: Transition from planar zigzag to spiral. *Physics of Fluids* **23**, 011702 (2011).
- ¹⁹ Lee, C. et al. 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 (2013).
- ²⁰ Heisinger, L., Newton, P. & Kanso, E. Coins falling in water. *Journal of Fluid Mechanics* **742**, 243–253 (2014).
- ²¹ Stringham, G., Simons, D., & Guy, H. The behavior of large particles falling in quiescent liquids. *US Government Printing Office* (1969).
- ²² Willmarth, W., Hawk, N. & Harvey, R. Steady and unsteady motions and wakes of freely falling disks. *Physics of Fluids* **7**, 197–208 (1964).
- ²³ Mahadevan, L., Ryu, W.S. & Samuel, A.D. Tumbling cards. *Physics of Fluids* **11**, 1–3 (1999).

- ²⁴ Skews, B.W. Autorotation of rectangular plates. *Journal of Fluid Mechanics* **217**, 33–40 (1990).
- ²⁵ Wang, W.B., Hu, R.F., Xu, S.J. & Wu, Z.N. Influence of aspect ratio on tumbling plates. *Journal of Fluid Mechanics* **733**, 650–679 (2013).
- ²⁶ Vincent, L., Shambaugh, W.S. & Kanso, E. Holes stabilize freely falling coins. *Journal of Fluid Mechanics* **801**, 250–259 (2016).
- ²⁷ Varshney, K., Chang, S. & Wang, Z.J. Unsteady aerodynamic forces and torques on falling parallelograms in coupled tumbling-helical motions. *Physical Review E* **87**, 053021 (2013).
- ²⁸ Belmonte, A., Eisenberg, H. & Moses, E. From flutter to tumble: inertial drag and Froude similarity in falling paper. *Physical Review Letters* **81**, 345–348 (1998).
- ²⁹ Andersen, A., Pesavento, U. & Wang, Z.J. Analysis of transitions between fluttering, tumbling and steady descent of falling cards. *Journal of Fluid Mechanics* **541**, 91–104 (2005).
- ³⁰ Fernandes, P.C., Ern, P., Risso, F. & Magnaudet, J. On the zigzag dynamics of freely moving axisymmetric bodies. *Physics of Fluids* **17**, 098107 (2005).
- ³¹ Pesavento, U & Wang, Z.J. Falling paper: Navier-Stokes solutions, model of fluid forces, and center of mass elevation. *Physical Review Letters* **93**, 144501 (2004).
- ³² Jin, C. & Xu, K. Numerical study of the unsteady aerodynamics of freely falling plates. *Communications in Computational Physics* **3**, 834–851 (2008).
- ³³ Waltz, B. & Buchanan, B.G. Automating science. *Science* **324**, 43–44 (2009).
- ³⁴ Peplow, M. Organic synthesis: the robot-chemist. *Nature* **512**, 20–22 (2014).
- ³⁵ Mjolsness, E. & DeCoste, D. Machine learning for science: state of the art and future prospects. *Science* **293**, 2051–2055 (2001).
- ³⁶ Bottou, L., Curtis, F.E. & Nocedal, J. Optimization methods for large-scale machine learning. *SIAM Review* **60**, 223–31 (2018).
- ³⁷ Soldatova, L.N., Clare, A., Sparkes, A. & King, R.D. An ontology for a robot scientist. *Bioinformatics* **22**, e464–e471 (2006).
- ³⁸ Sparkes, A. et al. Towards robot scientists for autonomous scientific discovery. *Automated Experimentation* **2**, 1 (2010).
- ³⁹ Fan D., et al. A robotic intelligent towing tank for learning complex fluid-structure dynamics. *Science Robotics* **4** (2019).
- ⁴⁰ Chapman, T. Lab automation and robotics: automation on the move. *Nature* **421**, 661–663 (2003).
- ⁴¹ Kachel, V., Sindelar, G. & Grimm, S. High-throughput isolation of ultra-pure plasmid DNA by a robotic system. *BMC Biotechnology* **6**, 9 (2006).
- ⁴² Sparkes, A. et al. An integrated laboratory robotic system for autonomous discovery of gene function. *Journal of the Association for Laboratory Automation*, **15**, 33–40 (2010).
- ⁴³ King, R.D. et al. Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature* **427**, 247–252 (2004).
- ⁴⁴ Vasilevich, A. & de Boer, J. Robot-scientists will lead tomorrow’s biomaterials discovery. *Current Opinion in Biomedical Engineering* **6**, 74–80 (2018).
- ⁴⁵ Bellemare, M. et al. Unifying count-based exploration and intrinsic motivation. *NIPS 29 Proceedings* (2016).
- ⁴⁶ Tang, H. et al. #Exploration: a study of count-based exploration for deep reinforcement learning. *NIPS 30 Proceedings* (2017).
- ⁴⁷ Frankel, F. & Reid, R. Big data: distilling meaning from data. *Nature* **455**, 30–30 (2008).

- ⁴⁸ Howison, T., Hughes, J., Giardina, F. & Iida, F. Physics driven behavioural clustering of free-falling paper shapes. *PLOS ONE* **14**, e0217997 (2019).
- ⁴⁹ Schmidt, M. & Lipson, H. Distilling free-form natural laws from experimental data. *Science* **309**, 1236–1239 (2005).
- ⁵⁰ Bongard, J. & Lipson, H. Automated reverse engineering of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences of the United States of America* **104**, 9943–9948 (2007).
- ⁵¹ Brunton, S.L., Proctor, J.L. & Kutz, J.N. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences of the United States of America* **113**, 3932–3937 (2016).
- ⁵² Hartigan, J.A. & Wong, M.A. Algorithm AS 136: a k-means clustering algorithm. *Applied Statistics* **28**, 100–108 (1979).
- ⁵³ Goldberg, D.E. & Holland, J.H. Genetic algorithms and machine learning. *Machine Learning* **3**, 95–99 (1988).
- ⁵⁴ Mouret, J.B. & Clune, J. Illuminating search spaces by mapping elites. *arXiv preprint*, arXiv:1504.04909 (2015).