

**Learning to write with the fluid rope trick**

| | |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Journal: | <i>Soft Matter</i> |
| Manuscript ID | SM-ART-02-2023-000177.R1 |
| Article Type: | Paper |
| Date Submitted by the Author: | 04-Sep-2023 |
| Complete List of Authors: | Chaudhary, Gaurav; Harvard University, School of Engineering and Applied Sciences Christ, Stephanie; Harvard University, School of Engineering and Applied Sciences Hart, A.; Massachusetts Institute of Technology, Mechanical Engineering Mahadevan, L.; Harvard University Faculty of Arts and Sciences, Engineering and Applied Science |
| | |

Cite this: DOI: 00.0000/xxxxxxxxxx

Learning to write with the fluid rope trick[†]Gaurav Chaudhary,^{a,b} Stephanie Christ,^a A John Hart,^b and L Mahadevan ^{*a,c}

Received Date

Accepted Date

DOI: 00.0000/xxxxxxxxxx

Direct ink writing, a versatile method of 3D and 4D printing, requires the precise placement of a nozzle just above the print surface to prevent fluid instabilities that cause deviations from the prescribed print path. But what if one could harness the instability associated with the spontaneously folding or coiling of a thin stream of viscous fluid, i.e. use the "fluid rope trick" to write specified patterns on a substrate? Here we use Deep Reinforcement Learning to derive control strategies for the motion of the extruding nozzle and thus the fluid patterns that are deposited on the surface. The method proceeds by having a learner (nozzle) repeatedly interact with the environment (a viscous filament simulator), and improves its strategy using the results of this experience. We demonstrate the outcome of the learned control instructions using experiments to drive a viscous jet and create cursive writing patterns and Pollockian paintings on substrates.

1 Introduction

The rapid evolution of three-dimensional (3D) printing technology enabling new manufacturing capabilities^{1,2} has been driven by the ability to control the additive deposition of complex materials through a computer-controlled nozzle. However, the ultimate limits on print quality are imposed by the rheodynamics of the printed material that is a function of material rheology and the ability to design robust print paths for the nozzle. Defects in printing typically arise in non-uniformly extruded/deposited material and instabilities such as folding and coiling of fluid jets³. A simple protocol to prevent these defects is to force the nozzle trajectory to exactly mimic the target print pattern from a very small height offset.

While this approach increases the accuracy of layer-by-layer direct writing of 3D objects, the result is that printing is typically a very slow process. Moreover, it is challenging to adapt these methods to complex topographies of the substrate on which the deposited material is laid down, or to precisely control sharp turns without distorting the extruded filament. Here, we explore the possibility that 3D printing could be sped up by operating with a nozzle that is elevated above the print surface, and harnessing the dynamic instability of a falling fluid jet^{4–9} to enable rapid ac-

curate printing without requiring the nozzle to exactly mimic the target pattern.

While the coiling instability of a falling jet has been studied in variety of different situations^{3,10–13}, most prior work has focused on explorations of the resulting patterns in response to simple controls of the nozzle. Here we ask if we can actively learn to precisely control the nozzle in space-time to print desired patterns by harnessing folding and coiling instabilities rather than avoiding them. That this can be done, and very well, is evidenced in the striking art of Jackson Pollock, who created paintings by dripping and pouring paint on a canvas from a height while moving his hand¹⁴. More generally, this raises the question: given a final target pattern or shape in a plane, how can we optimize the nozzle trajectory so that a continuous stream of material can be printed accurately and quickly, without having to precisely translate the nozzle along the print path, or bringing the nozzle close to the substrate. Put simply, how can a machine learn and deploy these optimal strategies from experience?

Our approach taps into the recent success of a class of machine learning algorithms known as reinforcement learning methods¹⁵ that are able to harness the expressive power of neural nets (NN) to explore and exploit a large state and action space to find solutions to difficult tasks. The starting point is a physical simulator that characterizes the necessary physics of thin threads of viscous fluid extruded from a height and captures the interactions between agent (nozzle head) and environment (print surface). Then, by building a framework of learning through repeated interactions with the environment, we show that it is possible to find a nozzle trajectory that can print a target pattern. Here, we only consider the case of learning to print a single layer of material onto a flat surface and leave the challenging problem of stacking multiple such layers to build a 3D structure for future

^a E-mail: lmahadev@g.harvard.edu^a School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138^b Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139^c Department of Physics, Department of Organismic and Evolutionary Biology, Harvard University, Cambridge, MA 02138[†] Electronic Supplementary Information (ESI) available: Supplementary Videos 1–4. See DOI: 10.1039/cXsm00000x/‡ Code available: <https://gitlab.com/stchrist/ViscousThreads>

investigation.

2 The dynamics of viscous filaments

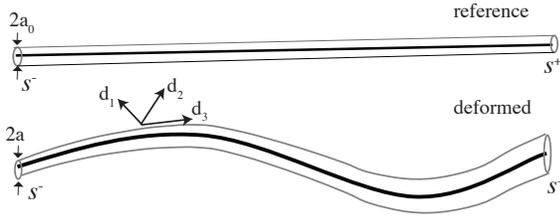


Fig. 1 Representation of the fluid thread and the associated material frame. Taking advantage of the slender aspect ratio of the filament, we can completely define its trajectory in space-time using its center line $x(S, t) \in \mathbb{R}^3$ and an orthonormal material frame $[d_1(S, t), d_2(S, t), d_3(S, t)]$ whose dynamics follow from the balance of linear and angular momentum; Figure adapted from Audoly et al.¹⁶.

The simulator used to model the liquid-rope instability of a viscous fluid follows earlier work on characterizing the coiling of liquid jets¹⁶, which we summarize briefly. A thin stream of viscous liquid is represented using a discrete Lagrangian description that takes advantage of the large aspect ratio (length/radius) of the system to deduce a reduced order description for the fluid thread. Then, the centerline $x(S, t) \in \mathbb{R}^3$ and an orthonormal material frame $[d_1(S, t), d_2(S, t), d_3(S, t)]$ suffices to capture the time evolution of viscous thread as shown in Fig. 1. Here, the material frame is adapted to centerline by requiring d_3 to align with the tangent to the centerline, and d_1, d_2 span the plane normal to centerline's tangent. The dynamical equations of the thread of the thread are given by the balance of linear and angular momentum for the position $x(S, t)$ and the orthonormal frame associated with $d_i(S, t)$, or equivalently, the spin angular velocity $v(S, t)$ in terms of equations¹⁶:

$$\rho A_0 \ddot{x}(S, t) = F_v(S, t) + f(S, t) \quad (1)$$

$$IJ \dot{v}(S, t) = M_v(S, t) + m(S, t) \quad (2)$$

Here, F_v and M_v are internal resultant viscous forces and moments in the thread arising from the rates of stretching, bending and twisting of the thread, while f and m are the external body force and body moment density, respectively. The expressions for the resultant force and moments in the fluid thread are expressed using the kinematics of the centerline representation in terms of the local area A_0 and moments of inertia J at every cross-section S , as well as the viscosity of the fluid (assuming it can be modeled as a simple Newtonian liquid); see^{16, 8} for the detailed expressions of the individual terms in the above equation and their respective discretizations, which form the basis for this work. When supplemented by boundary conditions at the nozzle and free boundary conditions on the substrate, where the filament touches down tangentially and is assumed to adhere instantly, this completes the formulation of the problem.

A critical dynamical variable underlying the generation of complex patterns¹⁷ is the ability to move the nozzle; we assume that it can be moved freely in a plane using horizontal movements and

also allow it to rotate relative to two axes, as shown in Fig. 2 A, with the angles ϕ and θ describing the rotational degrees of freedom (DoF). Assuming that the flow rate (Q) and nozzle extrusion diameter a are constant, controlling the writing/painting pattern on the substrate requires that the agent (nozzle) have time-varying translational and rotational velocities, likely to be bounded within a range of practical relevance. Vertical motion of the nozzle changes the nature of the pattern deposited; at low heights, the jet is rectilinear, then starts to coil as the height is increased past a first threshold, and then the jet coils rapidly as inertial effects start to dominate at even larger falling heights.

In Fig. 2 A (bottom), we show a coiling pattern deposited on the surface ($z = 0$); plotted using the normalized coordinates $x^* = x/a$ and $y^* = y/a$ with the nozzle radius a . For the remainder of the article, we normalize all the relevant lengths with the nozzle radius and velocity with \sqrt{gh} , where g is the gravitational acceleration and h is the nozzle height. The coil radius increases as the height of the nozzle is increased, and when the agent is translated linearly along the y -axis with a constant velocity, following a short transient, we see that an overlaid coiling pattern is deposited on the substrate $z = 0$ similar to what has been observed experimentally¹⁷. These observations suggest that writing a desired pattern requires controlling the various DoFs of the nozzle. But how can we learn the form of the nozzle's actions?

3 A reinforcement learning framework

Reinforcement learning (RL)¹⁵ provides one possible answer by allowing the action of an agent to repeatedly interact with the environment and converge towards an optimal policy without having any prior knowledge of the underlying model. An RL problem is defined in terms of states, actions and rewards. The state $s \in S$ is a quantitative description of the environment at the current time, with S being the set of all possible states. In every state, a set of actions $A(s)$ are available to the agent. By taking an action $a \in A(s)$, the agent transitions from one state to another and receives a numerical scalar reward signal r from the environment. The reward is a measure of how desirable it is to take action a in state s . Overall, the goal of the learner (or agent) is to find an action policy $\pi(a|s)$ which maximizes the cumulative reward over the entire learning episode by exploring different ways of interacting with the environment; this schematic is summarized in Fig. 2 B.

For the printing problem, the current state is defined by the following parameters: the arc length S of the pattern that has been printed so far, the current position of the nozzle (x, y, z, ϕ, θ) with $z \in [h_{\min}, h_{\max}]$, and the last action, corresponding to the velocity of the nozzle ($v_x, v_y, v_z, v_\phi, v_\theta$). If the height of the nozzle z is outside of its prescribed range, the state is considered invalid. Just as the states, the actions (DoFs) that describe the motion of the nozzle are continuous. The agent has up to five degrees of freedom: translation velocities (v_x, v_y, v_z) and rotational velocity (v_ϕ, v_θ). Furthermore, we assume that the range of all actions is bounded within the interval $[v_{\min}, v_{\max}]$. For the position and motion of the nozzle, only the DoF that are part of the action are considered, as the other parameters stay constant. Thus, the dimension of the state is between 3 (one DoF) and 11 (five DoF).

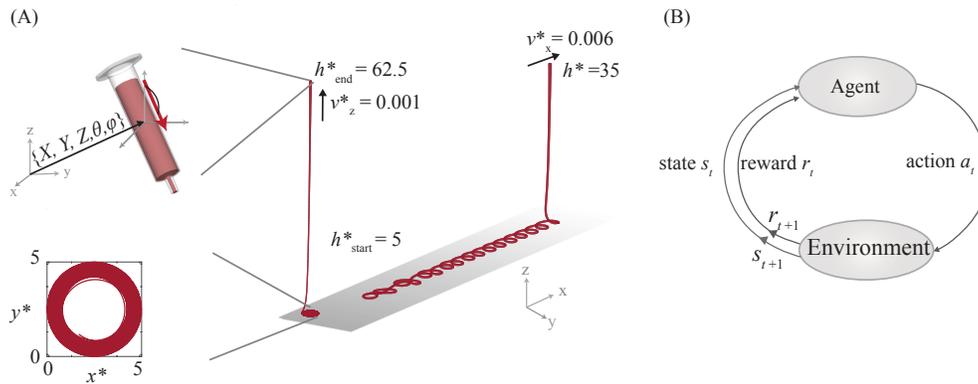


Fig. 2 Printing with the liquid-rope instability. (A) The state of a printing nozzle is represented by its translational coordinates (X, Y, Z) w.r.t. a stationary frame. The azimuthal angle ϕ and polar angle θ further specify the ink extrusion direction (solid red arrow). Simulation of viscous thread ($\eta = 5000$ cP) coiling as the print head is translated vertically at a constant speed. The top view shows coils of increasing radii as the nozzle is raised from $z = 0.5$ to $z = 12.5$ cm. Translating the nozzle at a constant speed of 0.5 cm/s creates periodic coils on the surface as shown for a fixed set of parameters. (B) In the controlled setting, the agent (nozzle) interacts with the environment to improve its action-selection policy π . It receives a state s_t and numerical reward r_t from the environment, based on which it updates its policy and returns the next action a_t . By taking action a_t , the environment transitions to the next state s_{t+1} . It measures its desirability through the reward r_{t+1} and sends the information, along with the new state, back to the agent.

Theoretically, it would be sufficient to use the arc length as the only state variable, however, we find that the rewards converge very slowly in that case.

More specifically, we use an off-policy actor-critic named V-RACER¹⁸ as our RL framework, although many other methods are likely to also work. A summary of this follows: the algorithm trains a neural network to approximate a continuous path $\pi^w(a|s)$ for the nozzle (continuous policy approximation). The policy network is randomly initialized and then iteratively updated through repeated attempts to reach the target following the policy gradient theorem. We employ "Remember and Forget Experience Replay" to reuse past experiences over multiple iterations to update the policy in a stable and data-efficient manner, with hidden network layers with 128 LSTM (long short-term memory) units each. This representational capacity of the network was found to be sufficient for the cases considered in this work. The actuation period (temporal resolution) of our simulations was generally fixed to 0.01 s; varying the number of units and the number of layers did not change the final optimized solutions.

Since the goal of the learning to write is equivalent to depositing a 2D pattern of the same shape as a target input pattern, the reward/penalty is chosen so as to reflect the mismatch between the target and the printed pattern. The target pattern is given as a time-ordered list of 2D coordinates of the path. We implemented the reward by comparing the position of the printed vertices with vertices of the given target shape. Denoting the total arc length of the pattern that has been created so far by S_i – where i is the current learning step, the error is computed by integrating it from the last step to the current step:

$$r = - \int_{S_{i-1}}^{S_i} |\tilde{\mathbf{x}}_{target}(S) - \tilde{\mathbf{x}}(S)| dS \quad (3)$$

Here we have chosen the error at a certain arc length in terms of the absolute difference between the target position and the actual

position, linearly interpolated from their two respective neighbors. In addition to the reward based on the pattern, we give a strong negative reward for invalid states, i.e. for invalid heights z . Using this reward function, the maximum – and optimal – cumulative reward is 0, as this would indicate that the created pattern does not deviate from the target pattern at any point.

To create regular initial conditions for learning, the simulation is first stabilized by using a constant velocity $(v_x, v_y, v_z, v_\phi, v_\theta) = (v_0, 0, 0, 0, 0)$ for a fixed number of time steps until the pattern follows a simple straight line. From there, the agent is allowed to start learning the appropriate set of actions to replicate the target pattern. In practice, in addition to the simulation parameters (nozzle flow/radius and fluid properties) and RL parameters (learning rate, discount factor, size of NN, activation function), we also specify the maximum number of steps per episode n_s , i.e. the number of actions that the agent can take before the episode is over.

4 Using reinforcement to learn writing

To validate that the RL agent takes optimal actions, we first ask if we can reproduce the simplified scenario of printing a straight line from a nozzle. As shown earlier in Fig. 2 A, for an arbitrary unidirectional velocity, an uncontrolled agent, translating with an arbitrary velocity produces a pattern consisting of overlaid coils¹¹. If the nozzle is allowed to vary its planar velocity (v_x and v_y), we see that over time associated with an increase in the number of episodes, the agent converges to create a straight line as shown in Fig. 3 A i (green curves). Consistent with this, the cumulative reward plateaus with increasing episode number (Fig. 3 A ii); the value of the plateau depends on the height of fluid extrusion and bounds on the nozzle actions. To draw a straight line, the agent speed must counter the effective velocity at which the coils are naturally laid on the surface, i.e. the optimal solution must satisfy the relations $v_x \geq \Omega R$, and $v_y \sim 0$.

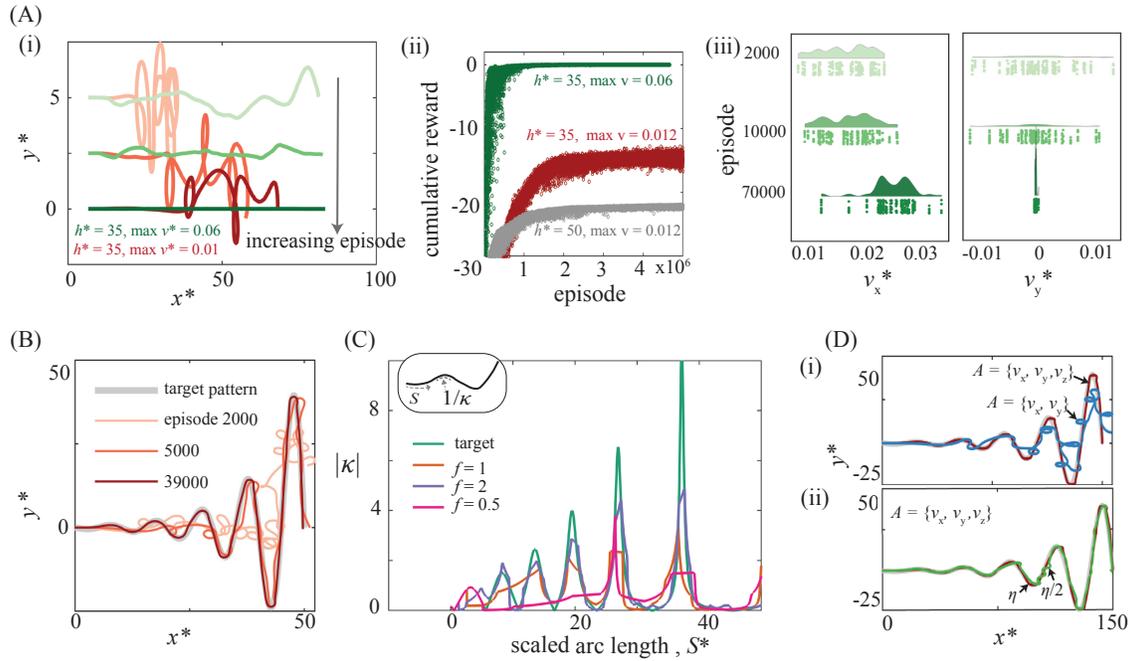


Fig. 3 Using reinforcement learning to control the liquid coiling instability. (A-i) The agent takes actions, $a_t = \{v_x, v_y\}$ while extruding liquid from a fixed height of 7 cm and learns to draw a straight line over multiple iterations, seeking to maximize the total reward. The red and green curves correspond of the maximum limit of 1 cm/s and 5 cm/s in the velocity components. (A-ii) Maximum reward depends on the limits on translation velocity, and the height from which the liquid is extruded. A better control can be learned when extruding from a smaller height for the same limits of velocity components. (A-iii) A density plot of the actions taken at various episodes. To create a straight line, the agent learns to suppress coiling by moving solely in the x-direction. (B) A more complex task of drawing a varying amplitude wave using a larger set of actions, $a_t = \{v_x, v_y, v_z\}$. Unlike the straight line in (A), the agent harnesses the liquid coils to draw the curved pattern. (C) The performance of the agent also depends on the relative scale of liquid coils and target curvature. A comparison of the local curvature ($|\kappa|$) of the drawn pattern at different scales show the challenge in learning when the scale of target curvature is significantly smaller than the nature curvature of liquid coils. (D-i) The constrained set of actions limits the agent's performance in drawing larger curvature. (D-ii) The fluid material properties, such as the kinematic viscosity η , also limit the accuracy of the learned pattern.

The action density plots in Fig. 3 A iii show that the agent does indeed converge to this state. We also test the performance by restricting the action to a lower maximum velocity (shown in red in Fig. 3 A i) and find that the agent is unable to converge to a straight line in such a case, consistent with our expectation. Another experiment we perform to understand the ability to learn is to increase the height of fluid extrusion, keeping the flow rate same as previous cases. A higher height allows for fluid jet to accelerate further, leading to a higher the coiling frequency. In such a case the appropriate actions that are required to prevent depositing coil would require a higher velocity bound, in the absence of which the cumulative reward will decrease further, as shown in Fig. 3 A ii.

To learn to write well, the agent must learn to negotiate curves with complex curvature profiles. Therefore, we need a target pattern with a wide range of curvatures and rate of variation of curvature; an exponentially decaying sinusoidal wave $y = 10e^{-1000x} \sin 2000\pi x$ serves both purposes. The agent is now allowed an additional action $v_z \neq 0$ in addition to the ability to vary (v_x, v_y) , and in Fig. 3 B we show the learned patterns as a function of episode number, and see that after about 5 000 episodes, the correct shape is learned, and the visible accuracy of the task does not improve much beyond that. A closer look at the evolution of the learned pattern shows that the agent first learns to match the pattern on larger length scales, followed by further improvements to closely match the features with sharp changes in curvature, i.e. on smaller length scales. This is consistent with the cumulative reward first showing a sharp increase followed by a very weak increase corresponding to the plateau-like regime in Fig. 3 A ii) associated with later episodes.

An optimal strategy for the agent aims to utilize or avoid the coiling instability depending on the nature of the curvature profiles in the target pattern. Thus, for curvatures comparable to the coiling radius, one should use the natural coiling instability, and otherwise to avoid it by moving quickly. However, when the (position and velocity) actions of the nozzle are bounded, writing patterns with curvatures that are significantly larger than the coil radius can be a challenge. To challenge the agent with such patterns, we change the absolute scale of the target pattern, thus changing the curvature profiles overall, and ask how well the RL agent, with the same bounds on action, learns to draw patterns at different scales. In Fig. 3 C we compare the absolute curvature of the target pattern with that of the learned patterns as a function of scale, using the same exponential form used in Fig. 3 B. For each case, the overall scale of the target pattern is halved or doubled, keeping all other parameters fixed. We see that the magnitude of curvature of the learned patterns is generally different from the target curves at the locations of extreme curvature; the case with smallest scale shows significant deviations from the target, implying that decreasing the scale of the target pattern (hence increasing the curvature), results in a reduction of the agent's ability to print the target. The bounds on velocity and position of the nozzle determine the largest curvature that can be deposited on the surface. A slow moving agent deposits overlaid coils whose radius linearly scales with the height of fluid extrusion (as shown in Fig. 2 A). In such a case, curvature of the deposited

pattern will approach the natural coiling curvature. A fast moving agent very close to the surface ideally can realize any curvature. Restricting agent to a minimum height above the surface, thus requires finding an action sequence to realize curvatures through reinforcement. In the case of smallest scale in Fig. 3 B, the agent is unable to find any action sequence that will allow it to deposit the large curvatures present in the target pattern.

The substrate patterns associated with the jet coiling instability are affected by both the limits on the dynamics of the nozzle as well as the properties of the fluid. To explore their respective roles, in Fig. 3 D i we show the effect of limiting the agent's actions to $v_x, v_y \neq 0, v_z = 0$, i.e. planar motion and compare it to the case when $v_z \neq 0$ (Fig. 3 B); we see that extruding fluid from a constant height introduces an unwanted coiling response in the learned pattern, especially near the regions of large curvature in the target pattern, demonstrating the influence of limiting the action space in the learning problem. In Fig. 3 D ii we show the effect of changing the magnitude of the viscous forces by simply changing the kinematic viscosity (η) of the fluid; reducing the viscosity leads to an increase in the coiling frequency in a predictable way and we see that after a similar number of learning episodes, the learned pattern is not quite as accurate. These tests of deploying RL point to a simple but important lesson; choosing the right range of actions is critical for good performance in any task, and becomes particularly clear in the context of interacting with and learning using physical systems.

To further explore the ability of RL in more realistic and useful tasks, we now demonstrate how to guide a nozzle to mimic curvilinear writing or paint like Pollock. We first prescribe the target pattern: a continuous curve associated with the word "Cambridge" along with an action space (v_x, v_y) , but no vertical velocity, i.e. $v_z = 0$. Since the coiling frequency ($\Omega \sim 1/h$) is higher when fluid is extruded from a larger height (h), the nozzle must translate faster to avoid unstable/unwanted coiling. The natural coiling frequency for the range of extrusion heights explored in Fig. 4 ranges between 0.1- 0.6 1/s. The ratio of agent's translational speed to the natural coiling speed ($\Omega R \sim \Omega h$) varies widely in the range 0.05-5 for each height. This large variation is associated with the use of the coiling mode (< 1) and its avoidance (> 1). In Fig. 4 A, we show how the agent learns to write from different fixed heights in the range $z \in [2, 10]$ cm, after the same number of training episodes. While the overall patterns look similar, a closer look shows noticeable differences in the regions of the curve with large curvature e.g. letters "r", "d" and "g". In Fig. 4 B we show the trajectory of the nozzle at a fixed height as it moves to deposit the pattern on the substrate; in regions of the pattern that are almost straight (or weakly curved), the agent avoids the coiling instability by generally moving rapidly, while it uses the coiling instability to achieve higher curvature in regions where the target pattern demands it. Indeed, the deployment of coils becomes more apparent when printing from higher heights, when the nozzle makes sharper turns in its path (e.g. for the letter "m") to be able to print smooth patterns on the surface.

Continuing to characterize the properties of printing using a nozzle operating from a height, in Fig. 4 C we show the differences between the local direction of motion of the nozzle and the

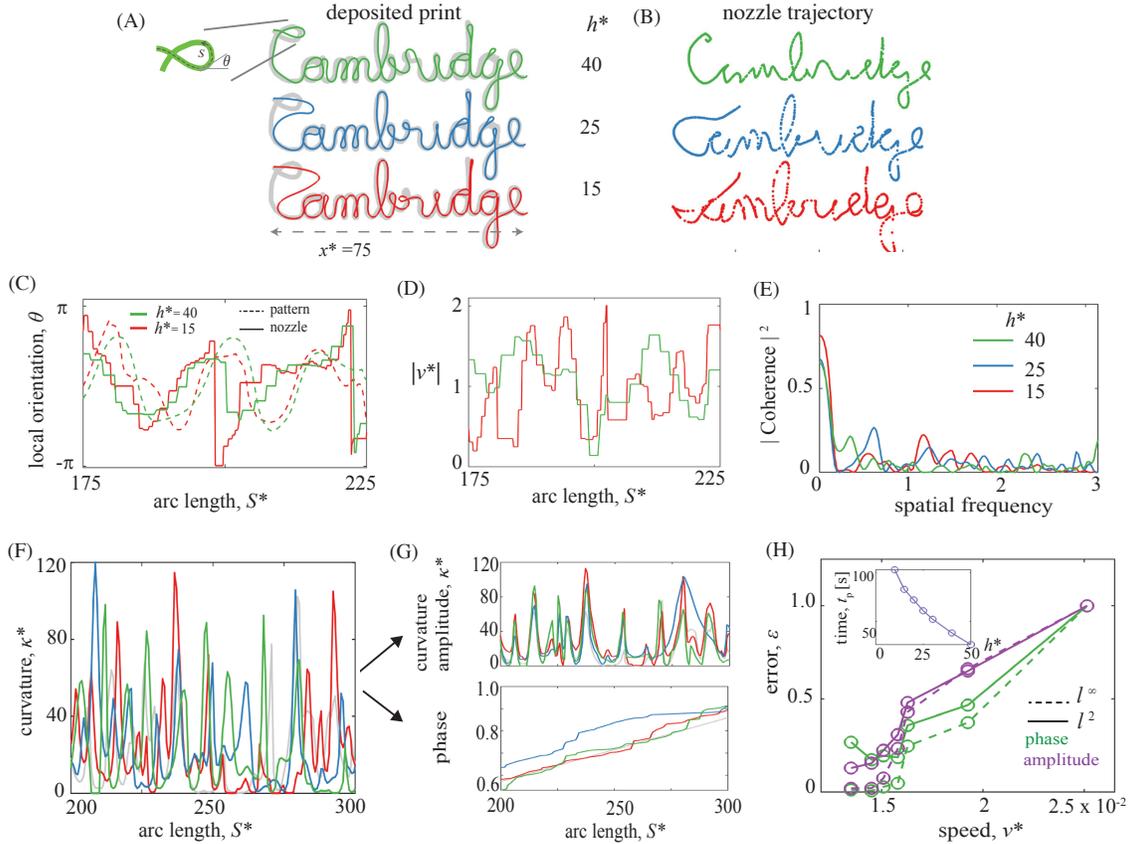


Fig. 4 Using RL to reproduce cursive writing. The agent, with actions $a_t = \{v_x, v_y\}$, can harness liquid coiling to write cursive text. (A) The predicted deposited patterns from fixed heights as shown. In each case the z-component of velocity is set to zero. (B) The nozzle trajectories at corresponding heights are shown. The mismatch between the deposited pattern and the nozzle trajectory is apparent at higher nozzle location. (C) A comparison of local direction (local orientation w.r.t the horizontal θ) of translation of the nozzle and the printed pattern show that non-trivial actions are performed to control the target pattern, and the agent does not just follow the trajectory of target pattern. (D) The normalized magnitude of actions, $|v^*| = \sqrt{(v_x^2 + v_y^2)}/\sqrt{gh}$, that were used from printing are shown. (E) The magnitude squared coherence between the printed patterns and the target pattern is shown. A higher coherence is observed at lower spatial frequencies. (F) The curvature of the cursive text plotted as a function of arc length shows misaligned curvature peaks for patterns deposited from different heights. The error in κ vs S curves can be decomposed into vertical and a horizontal components (in G) using elastic functional data analysis. (H) Changing the height of extrusion allows for directly controlling the total printing time (inset). A speed-error trade-off results as a consequence. Within a given range of actions, a lower error, defined as $\varepsilon = (\int (1 - \kappa(S)/\kappa_0(S))^p dS)^{1/p}$, is achieved when nozzle is held closer to the surface as a result of low-coiling frequency and smaller liquid coil radii. Here we define non-dimensional printing speed $v^* = (S_{total}/t_p)/\sqrt{gh}$

local direction of the pattern that gets laid on the surface. We see that there are two types of errors: at some places, the local orientation of curves appears phase-shifted, while in other regions the amplitudes of the curves are poorly correlated. The phase-shift corresponds to the lag between the nozzle's position and the point of contact of fluid jet with the surface; locations where the two are poorly correlated correspond to the case when the nozzle deploys fluid coils to print. In contrast, the normalized velocity magnitude, $|v^*| = |v|/\sqrt{gh}$, shows frequent changes when the nozzle is held at a small but fixed height as shown in Fig. 4 D, and thus leads to differences in the curvature of the nozzle and the pattern being printed. We see that the nozzle is likely to more closely follow the target pattern while printing the pattern from a lower height.

To quantify the accuracy of learning the task, we use two approaches based on local measures to derive global metrics. Since we work with planar curves, up to rigid motions (that we do not worry about), the target and learned patterns are completely characterized by their scalar curvature as a function of the arc-length. For one measure of the error on the scale of the whole pattern, we define a coherence metric in terms of the spatial Fourier domain as follows :

$$C_{\kappa_0\kappa}(f) = \frac{|P_{\kappa_0\kappa}(f)|^2}{(P_{\kappa_0\kappa_0}(f)P_{\kappa\kappa}(f))} \quad (4)$$

where $P_{\kappa\kappa}(f)$ and $P_{\kappa_0\kappa_0}(f)$ are the power spectral densities of the curvatures $\kappa_0(S)$ (target) and $\kappa(S)$ (learned), respectively, and $P_{\kappa_0\kappa}(f)$ is the cross power spectral density between $\kappa_0(S)$ and $\kappa(S)$.¹⁹, observing that $C_{\kappa_0\kappa}(f) \in [0, 1]$. Plotting the coherence in Fig. 4 C, we note that $\kappa_0(S)$ and $\kappa(S)$ are strongly associated at very low spatial frequencies (or relatively larger scales) while at higher spatial frequencies the association is relatively weak. This is consistent with the observations from other learning experiments as well: the agent can capture the large scale features of the curves either by moving in straighter or slightly curved path or printing coils, but when the target curvature scales are larger than the natural coiling scales, there is a mismatch between the target and learned curvature.

To compare the difference between the target and learned curvature along the trajectory raises a familiar problem of registration - how does one align points along the learned path with points along the target pattern, given that the speed of the nozzle is not necessarily a constant? To solve this problem, we use methods from functional and shape data analysis²⁰ whereby we simultaneously solve the problem of registration and determination of error by insisting on reparametrization invariance of the error metric. A natural solution that presents itself is to align the curvature data for various cases with the curvature of the target curve using the Fisher-Rao metric and the square-root velocity function (SRVF) representation of curvature data²⁰, thus separating out errors in the phase and the amplitude of the curvature. Fig. 4 E shows the aligned curvature of different cases shown in Fig. 4 D and their respective phase shift. Defining the error as the

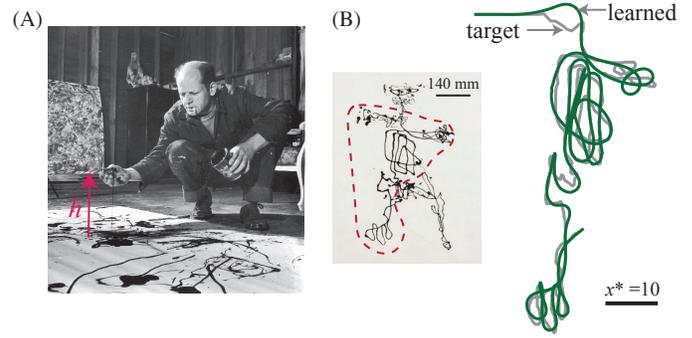


Fig. 5 Using RL to mimic Pollock's paintings. Liquid coiling was extensively utilized by famous American painter Jackson Pollock in his drip paintings (as shown in (A)²¹). (B) With the available set of action (v_x, v_y, v_z) , the agent can learn to draw parts of Jackson Pollock's, *Figure, 1948*, indicating that Pollock's drip painting may owe its complexity to the liquid coiling instability.

$l - p$ norm of the curvature mismatch,

$$\varepsilon = \left[\int \left(1 - \frac{\kappa(S)}{\kappa_0(S)} \right)^p dS \right]^{1/p}, \quad (5)$$

we quantify agent's capability to print from various heights above the surface. Printing with different nozzle heights is associated with a trade-off. Small nozzle heights lead to very accurate learned paths but ones that are very slow, since they do not exploit the instability, while larger nozzle heights lead to faster printing albeit with poorer accuracy. In the inset of Fig. 4 F we show that a higher extrusion height results in a higher coiling frequency and hence faster overall printing. In Fig. 4 F we show the trade-off by comparing the overall curves printed from different heights; faster printing is also less accurate. We note that there is an additional effect, since the filament diameter at the substrate changes with the height of the nozzle, but here we ignore this and only consider matching the target pattern with the centerline of the fluid jet.

Finally we ask if the agent can learn to not just write, but sketch, paint or draw. As an exemplar, we chose part of a painting by Jackson Pollock (*Figure, 1948*²¹), famous for his "painting at a distance" style, wherein he allowed paint to drip and drizzle from a brush or a rod held far above the canvas, thus exploiting both fluid coiling and jetting instabilities¹⁴ as shown in Fig. 5 A. Could an agent learn the complex movements of the artist given the final result? Restricting ourselves to a limited part of the painting which we could easily track as a continuous curve as shown in Fig. 5 B inset, we show that by allowing the action space for the RL agent to be non-zero v_x, v_y, v_z does in fact reproduce a fragment of the Pollockian painting as shown in Fig. 5 B.

5 Physical experiments on learning to write

The real test of the RL protocols described above is of course to deploy them in physical experiments. To do so, we dispense a viscous fluid, silicone oil ($\eta = 500$ cP), through a nozzle at a fixed flow rate, and deposit it along a defined path from above a smooth sheet of paper glued on a flat, stationary platform. We use the RL control learned in the numerical setting as the digital

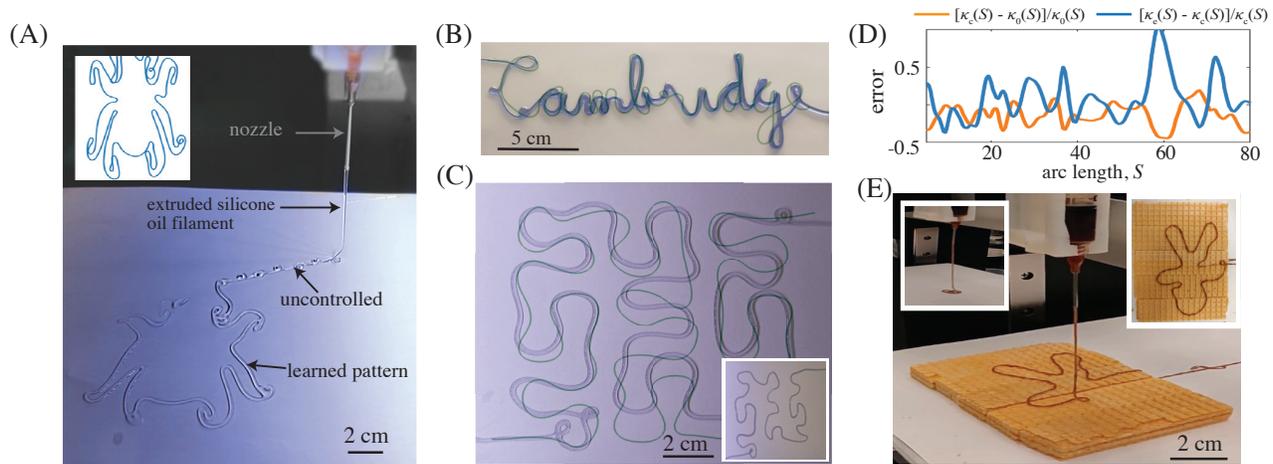


Fig. 6 Testing the predicted RL strategies using physical experiments to write. (A) An experimental print showing a pattern with a first region printed with RL-derived and a second region without control where the nozzle is moved at 0.5 cm/s. (inset) The pattern printed in simulation. (B) RL controlled cursive handwriting using silicone oil, overlaid with the numerically simulated solution. (C) An experimentally printed self-similar structure (solid line), and the numerically learned pattern (dashed line). The mismatch between the experiment and numerical result grows as the arc length grows starting from the bottom left corner. (Inset) The experimental pattern deposited with a noisy set of actions shows the sensitivity of the final pattern to the RL actions. (D) A comparison of error in the curvature along the entire arc length of the pattern in (C) is shown. Here, $\kappa_c(S)$ and $\kappa_e(S)$ are the curvatures of computed and experimentally printed pattern, respectively, from the learned set of actions. (E) In a practical setting, the RL agent can be deployed to automate processes such as printing chocolate syrup on edible surfaces. (insets) a natural coiling of chocolate syrup extruded from a stationary nozzle held at a height of 6.6 cm above the surface and the top view of the final print.

path along which the printer head translates. Since a fixed time step, $\Delta t = 0.01$ s was used in the simulations, we find this digital path by integrating the action set (velocities) at every step. The diameter of the nozzle and flow rate are same as the numerical experiments. In all the experiments the oil is first extruded until the coiling becomes steady, to ensure similar initial conditions as that of the numerical simulations.

We start by showing the difference between the RL protocol and uncontrolled flow, in Fig. 6 A (see S.I. video 1) for an experimental print shown in the inset. Using the learned RL control, the agent is able to smoothly print the pattern whereas when the agent moves at a constant velocity, fluid coiling readily occurs, thus producing a mimic of the given target pattern. To test the ability to write cursively, in Figure 6 B (see S.I. video 2) we show an example parallel to the simulation results in Fig. 4 A using the action set v_x, v_y, v_z , with RL controls based on allowable states for nozzle heights $z \in [4 - 10]$ cm. An overlay of the numerically learned pattern on the experimentally printed pattern shows that the regions of highest errors are generally the regions of maximum curvature, most likely due to velocity mismatch between the RL actions and the experimental implementation. In particular when the nozzle makes sharp changes in direction over a short time corresponding to high curvature regions), the experimentally executed velocity actions are expected to be different from the desired actions. As a consequence, the overall mismatch (error) between the numerical and the experimental pattern builds up as a function of arc length of the pattern. We emphasize this using a self-similar target pattern of a Peano curve in Fig. 6 C (see S.I. video 3) using actions v_x, v_y . We observe that the experimental print and the simulated print show close agreement in the initial part of arc length, but as the printing progresses, the apparent experimental error accumulate over the arc length. A comparison

between the curvatures of computed pattern and the target pattern (computational error), and experimentally printed pattern and computed pattern is shown in Fig. 6 D. In both the cases the error is computed between the aligned curvature vectors using the methodology described in the previous section. It is evident that along the entire arc length, the error in the curvature of the computed pattern deviates consistently from the target curvature. The error between the computed and the experimentally printed pattern has even larger magnitude due the additional experimental errors. In the inset of Fig. 6 C we show the sensitivity of the resulting pattern on the error in optimal set of actions. Here, we add a random noise to the actions ($< 5\%$ of the action value), and use the noisy actions to the print the pattern. The resulting pattern deviates significantly, highlighting the non-trivial nature of the action-reward landscape due to the unsteady and nonlinear effects in the physical problem.

As a test of the RL control approach developed for Newtonian fluids extruded on smooth surfaces, we ask how well the strategy does when printing a thick chocolate syrup on a textured wafer, inspired by a tasty application of our approach. Chocolate syrup is a non-Newtonian fluid with a strain-rate dependent viscosity, but here we assume that it has a constant shear viscosity (~ 100 cP based on ball drop experiments). We find that this approximation results in a reasonable agreement between the printed shape and the target pattern as shown in Figure 6 D (see S.I. video 4). We note that printing from a height naturally handles rough surfaces, unlike the traditional direct-ink write where an irregular surface will result in irregularities in the deposited material.

6 Conclusions

Complementing the substantial literature on understanding and predicting fluid instabilities, here we have focused on controlling

a somewhat non-intuitive low Reynolds number instability associated with the buckling, folding and coiling of a slowly flowing slender viscous filament. Inspired by 3-d and 4-d printing technologies that rely on the movement of a nozzle that dispenses complex fluids onto a substrate from just above, we ask if we can harness the folding and coiling instabilities that arise as soon as the jet falls from a sufficient height above a surface. We answer this in the affirmative by combining a physics-based simulation engine and a variant of reinforcement learning to control the fluid coiling instability and learn to "print at a distance." By varying the action space, material properties and geometric scales that govern the dynamics of viscous coiling, we quantify the performance of an RL agent for a variety of problems, and show that it is possible to learn to write cursively and mimic Pollockian paintings. Deploying the learned policy in physical experiments demonstrates that we can create complex physical patterns leveraging a natural fluid instability. We envision such an approach can be further extended to more challenging scenarios such as printing on non-planar surfaces and using robotic manipulators with greater dexterity and improved motion control, and learn to stack multiple layers out of plane. These extensions would likely require a more sophisticated model to account for the non-flat nature of the substrate.

Acknowledgements

The computations in this paper were run on the FASRC Cannon cluster supported by the FAS Division of Science Research Computing Group at Harvard University. We thank Crystal Owens and the Hatsopoulos Microfluids Laboratory at MIT for help and support with the experiments, and the NSF Harvard MRSEC DMR-2011754, the Simons Foundation and the Henri Seydoux fund for partial financial support.

Notes and references

- 1 A. Sydney Gladman, E. A. Matsumoto, R. G. Nuzzo, L. Mahadevan and J. A. Lewis, *Nature materials*, 2016, **15**, 413–418.
- 2 J. W. Boley, W. M. Van Rees, C. Lissandrello, M. N. Horenstein, R. L. Truby, A. Kotikian, J. A. Lewis and L. Mahadevan, *Proceedings of the National Academy of Sciences*, 2019, **116**, 20856–20862.
- 3 H. Yuk and X. Zhao, *Advanced Materials*, 2018, **30**, 1704028.
- 4 G. Barnes and R. Woodcock, *American Journal of Physics*, 1958, **26**, 205–209.
- 5 L. Mahadevan, W. S. Ryu and A. D. Samuel, *Nature*, 1998, **392**, 140.
- 6 L. Mahadevan, W. S. Ryu and A. D. Samuel, *Nature*, 2000, **403**, 502.
- 7 H.-Y. Kim, M. Lee, K. J. Park, S. Kim and L. Mahadevan, *Nano letters*, 2010, **10**, 2138–2140.
- 8 N. M. Ribe, *Journal of Fluid Mechanics*, 2017, **812**, 9745–9754.
- 9 A. Chakrabarti, S. Al-Mosleh and L. Mahadevan, *Soft Matter*, 2021, **17**, 9745–9754.
- 10 F. Gallaire and P.-T. Brun, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 2017, **375**, 20160155.
- 11 P.-T. Brun, B. Audoly, N. M. Ribe, T. S. Eaves and J. R. Lister, *Physical review letters*, 2015, **114**, 174501.
- 12 P.-T. Brun, C. Inamura, D. Lizardo, G. Franchin, M. Stern, P. Houk and N. Oxman, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 2017, **375**, 20160156.
- 13 R. Passieux, L. Guthrie, S. H. Rad, M. Lévesque, D. Theriault and F. P. Gosselin, *Advanced Materials*, 2015, **27**, 3676–3680.
- 14 A. Herczynski, C. Cernuschi and L. Mahadevan, *Physics Today*, 2011.
- 15 R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction (2nd Edition, Draft)*, MIT Press, Cambridge, MA, 2nd edn, 2017, p. 455.
- 16 B. Audoly, N. Clauvelin, P.-T. Brun, M. Bergou, E. Grinspun and M. Wardetzky, *Journal of Computational Physics*, 2013, **253**, 18–49.
- 17 S. W. Morris, J. H. Dawes, N. M. Ribe and J. R. Lister, *Physical Review E*, 2008, **77**, 066218.
- 18 G. Novati and P. Koumoutsakos, *International Conference on Machine Learning*, 2019, **97**, 4851–4860.
- 19 S. M. Kay, *Modern Spectral Estimation (Englewood Cliffs)*, 1988.
- 20 A. Srivastava and E. P. Klassen, *Functional and shape data analysis*, Springer, 2016.
- 21 J. Pollock, *Figure*, 1970, <https://sammlung.staedelmuseum.de/en/work/figure>.