

Lawrence Berkeley National Laboratory

LBL Publications

Title

Learning protocols for the fast and efficient control of active matter

Permalink

<https://escholarship.org/uc/item/4354g4hb>

Journal

Nature Communications, 15(1)

ISSN

2041-1723

Authors

Casert, Corneel

Whitelam, Stephen

Publication Date

2024-10-01

DOI

10.1038/s41467-024-52878-2

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>



Peer reviewed

Learning protocols for the fast and efficient control of active matter

Received: 7 May 2024

Accepted: 19 September 2024

Published online: 23 October 2024

 Check for updatesCorneel Casert ^{1,2} & Stephen Whitlam ¹ ✉

Exact analytic calculation shows that optimal control protocols for passive molecular systems often involve rapid variations and discontinuities. However, similar analytic baselines are not generally available for active-matter systems, because it is more difficult to treat active systems exactly. Here we use machine learning to derive efficient control protocols for active-matter systems, and find that they are characterized by sharp features similar to those seen in passive systems. We show that it is possible to learn protocols that effect fast and efficient state-to-state transformations in simulation models of active particles by encoding the protocol in the form of a neural network. We use evolutionary methods to identify protocols that take active particles from one steady state to another, as quickly as possible or with as little energy expended as possible. Our results show that protocols identified by a flexible neural-network ansatz, which allows the optimization of multiple control parameters and the emergence of sharp features, are more efficient than protocols derived recently by constrained analytical methods. Our learning scheme is straightforward to use in experiment, suggesting a way of designing protocols for the efficient manipulation of active matter in the laboratory.

Active particles extract energy from their surroundings to produce directed motion^{1–4}. Natural active particles include groups of animals and assemblies of cells and bacteria^{5–7}; synthetic active particles include active colloids and Janus particles^{8,9}. Active matter, collections of active particles, displays emergent behavior that includes motility-induced phase separation^{10,11}, flocking^{12,13}, swarming¹⁴, pattern formation^{15,16}, and the formation of living crystals¹⁷.

Recent work has focused on controlling such behavior by creating active engines^{18–27}, controllably clogging and unclogging microchannels²⁸, doing drug delivery in a targeted way^{29,30}, controlling active fluids through topological defects^{31–33}, and creating micro-robotic swarms with controllable collective behavior^{34–36}. For such applications, efficient time-dependent protocols are important^{37–40}. Methods for identifying efficient protocols, such as reinforcement learning, have been used to optimize the navigation of active particles in complex environments^{41–43} and induce transport in self-propelled disks using a controllable spotlight⁴⁴.

For purely diffusive (passive) molecular systems, analytic methods allow the identification of optimal time-dependent protocols for a range of model systems^{45–48}. These results establish that rapidly-varying and discontinuous features are common components of optimal protocols, and are useful for benchmarking numerical approaches^{49,50}. However, active-matter systems are more complicated to treat analytically than passive systems, requiring the imposition of protocol constraints in order to make optimization calculations feasible for even the simplest model systems. Two recent papers derive control protocols for confined active overdamped particles by assuming that protocols are slowly varying and smooth⁵¹ or have a specific functional form⁵². In this paper we show numerically that relaxing these assumptions leads to more efficient control protocols for those systems. In particular, we demonstrate the importance of allowing jump discontinuities and rapid variations in control protocols, similar to those seen for overdamped passive systems.

To learn protocols to control active matter we use the neuroevolutionary method described in refs. 49,53–55, which we adapted

¹Molecular Foundry, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA. ²Department of Physics and Astronomy, Ghent University, 9000 Ghent, Belgium. ✉e-mail: ccasert@lbl.gov; swhitelam@lbl.gov

from the computer science literature^{56–58}. Briefly, we encode a system’s time-dependent protocol in the form $\mathbf{g}_\theta(t/t_f)$. Here \mathbf{g} is the output vector of a deep neural network, corresponding to the control parameters of the system (which in this paper consist of the activity of the particles and the spring constant of their confining potential), θ is the set of neural-network weights, t is the elapsed time of the protocol, and t_f is the total protocol time. It is also straightforward within this scheme to consider a feedback-control protocol, by considering a neural network $\mathbf{g}_\theta(t/t_f, \mathbf{v})$, where \mathbf{v} is a vector of state-dependent information⁴⁹. We apply the protocol to the system in question, and compute an order parameter ϕ that is minimized when it achieves our desired objective (such as inducing a state-to-state transformation while emitting as little heat as possible). Neural networks are flexible function approximators, and they can be used to represent protocols that are free of the constraints imposed in recent analytical work: they do not have to follow a specific functional form, and they can be used to represent protocols that possess discontinuities and rapidly-varying features. The neural-network weights θ are iteratively adjusted by a genetic algorithm in order to identify the protocol whose associated value of ϕ is as small as possible.

This approach is a form of deep learning – in the limit of small mutations and a genetic population of size 2 it is equivalent to noisy gradient descent on the objective ϕ ⁵⁹ – and so comes with the benefits and drawbacks of deep learning generally. Neural networks are very expressive, and if trained well can identify “good” solutions to a problem, but these solutions are not guaranteed to be optimal^{60,61}. We must therefore be pragmatic, and (as with other forms of sampling) verify that protocols obtained from different starting conditions and from independent runs of the learning algorithm are consistent. Consequently, we call the protocols identified by the algorithm “learned” rather than “optimal”. In general, we have found the method to be easy to apply and to solve the problems we have set it: we have benchmarked the method – see refs. 49,55 and Fig. S1 in the Supplementary Information (SI) – against exact solutions⁴⁵ and other numerical methods^{48,50,62,63}. In this paper we use it to produce protocols that are closer to optimal than the protocols obtained by other methods^{51,52}.

Importantly, the neuroevolutionary learning algorithm uses information that is accessible in a typical experiment. While in this paper we have learned protocols for the control of simulation models (these protocols could then be applied to experiment if the simulation model is a good enough representation of the experiment⁶⁴), the same learning algorithm can also be applied directly to experiment. The success of this method as discussed in the following sections therefore

demonstrates the potential of neural-network protocols for the control of active matter in the laboratory.

Results

Active particle in a trap of variable stiffness

In this section we consider the problem of Section IIIA of ref. 51, a single active Ornstein-Uhlenbeck particle^{65–67} in a one-dimensional harmonic trap of stiffness $\alpha(t)$. A schematic of this model is shown in Fig. 1a. The particle has position r and self-propulsion velocity v . It experiences overdamped Brownian motion with diffusion constant D and mobility μ , such that

$$\dot{r}(t) = v(t) - \mu\alpha r(t) + \sqrt{2D}\eta(t). \tag{1}$$

Here η is a Gaussian white noise term with zero mean and unit variance. The self-propulsion velocity v follows an Ornstein-Uhlenbeck process with persistence time τ and amplitude D_1 , such that $\langle v \rangle = 0$ and $\langle v(t)v(t') \rangle = D_1\tau^{-1}e^{-|t-t'|/\tau}$. The parameter D_1 is zero in the passive limit.

The trajectory-averaged heat associated with varying $\alpha(t)$ from α_i to α_f in time t_f is^{19,51,66}

$$\langle Q \rangle = \frac{1}{2}(\alpha_i x_i - \alpha_f x_f) + \frac{1}{2} \int_0^{t_f} dt \dot{\alpha}(t)x(t) + \frac{D_1 t_f}{\tau\mu} - \int_0^{t_f} dt \alpha(t)y(t). \tag{2}$$

Here $\langle \cdot \rangle$ denotes an average over dynamical trajectories, and we have defined $x \equiv \langle r^2 \rangle$ and $y \equiv \langle rv \rangle$. For time-dependent quantities $q(t)$ we use the notation $q_i \equiv q(0)$ and $q_f \equiv q(t_f)$ to denote initial and final values. The first line of Eq. (2) is the passive heat (minus the change in energy plus the work done by changing the trap stiffness), and the second line is the active contribution to the heat. The first term on the second line is constant for fixed t_f (describing the heat dissipated to sustain the self-propelled motion), and plays no role in selecting the protocol.

For a given protocol $\alpha(t)$, the time evolution of x and y is given by the equations⁵¹

$$\frac{1}{2} \dot{x}(t) + \mu\alpha(t)x(t) = y(t) + D, \text{ and} \tag{3}$$

$$\tau\dot{y}(t) + \gamma(t)y(t) = D_1,$$




Fig. 1 | Protocols $\alpha(t)$ controlling the stiffness of a trap confining an active Ornstein-Uhlenbeck particle. Model parameters: $D = 1, \mu = 1, \tau = 1, D_1 = 2$. **a** Schematic of the model. **b** Neural-network protocols $\alpha(t)$ that minimize $\langle Q \rangle$, Eq. (2), for different protocol lengths t_f . **c** Heat associated with the protocols. The dotted and

dashed lines correspond to the values Q_1 , Eq. (6), and Q_2 , Eq. (7), respectively. Inset: now for $D_1 = 0.01$. The orange line (“DPF”) is the result extracted from ref. 51. **d** as (b), but requiring minimum heat while enacting a state-to-state transformation (SST); see Eq. (11).

where $\gamma(t) \equiv 1 + \mu\tau\alpha(t)$. The system starts in the steady state associated with the trap stiffness α_i , and so its initial coordinates are

$$x_i = \frac{1}{\alpha_i\mu} \left(\frac{D_1}{\gamma_i} + D \right) \text{ and } y_i = \frac{D_1}{\gamma_i}. \quad (4)$$

Ref. 51 sought protocols that carry out the change of trap stiffness $\alpha_i = 1 \rightarrow \alpha_f = 5$ with minimum mean heat, Eq. (2). The theoretical framework used in that work assumes that protocols $\alpha(t)$ are smooth and are not rapidly varying (see Section S2 for a discussion of this point). Here we revisit this problem using neuroevolution. We find that heat-minimizing protocols are not in general slowly varying or smooth, but can vary rapidly and can display jump discontinuities. The protocols we identify produce considerably less heat than do the protocols identified in ref. 51 (see Fig. 1 and Fig. S2).

To learn a protocol $\alpha(t)$ that minimizes heat, we encode a general time-dependent protocol using a deep neural network. We choose the parameterization

$$\alpha_\theta(t) = \alpha_i + (\alpha_f - \alpha_i)(t/t_f) + g_\theta(t/t_f), \quad (5)$$

where g is the output of a neural network whose input is t/t_f (restricting the scale of inputs to a range $[0, 1]$ typically allows training to proceed faster than when inputs can be numerically large). We constrain the neural network so that $\alpha_i \leq \alpha_\theta(t) \leq \alpha_f$, meaning that it cannot access values of α outside the range studied in ref. 51. When we relax this constraint we find protocols that produce less heat, in general, than the protocols that observe the constraint. We impose the constraint to allow us to make contact with ref. 51, and because experimental systems have constraints on the maximum values of their control parameters. Initially the weights and output of the neural network are zero, and so we start by assuming a protocol that interpolates linearly with time between the initial and final values of α . We train the neural network by genetic algorithm to minimize the order parameter $\phi = \langle Q \rangle$, given by Eq. (2), which we calculate for a given protocol by propagating (3) for time t_f , using a forward Euler discretization with step $\Delta t = 10^{-3}$. An example of the learning process is shown in Fig. S2a.

In Fig. 1b we show, for the choice $D_1 = 2$, that heat-minimizing protocols learned by the neural network vary between a step-like jump at the final time, for small values of t_f , and a step-like jump at the initial time, for large values of t_f (all protocols shown in this work are provided in the Supplementary Data 1 file). For intermediate values of t_f we observe a range of protocol types. These protocols include non-monotonic and rapidly-varying forms, and show jump discontinuities at initial and final times. In Sec. S3A, we discuss the effect of the model parameters on the range of t_f for which these non-trivial protocols result in a lower value of $\langle Q \rangle$ than the step protocols.

The heat associated with the final-time step protocol is just that associated with the initial steady state, and is

$$Q_1 = \frac{D_1 t_f}{\mu\tau(1 + \alpha_i\mu\tau)}. \quad (6)$$

The heat associated with the initial-time jump protocol can be calculated from Eqs. (2) and (3), and is

$$Q_2 = \frac{\alpha_f}{2} (x_i - x_2(t_f)) + \frac{D_1 t_f}{\mu\tau(1 + \alpha_f\mu\tau)} - \frac{D_1 \tau \alpha_f}{\gamma_f} \left(\frac{1}{\gamma_i} - \frac{1}{\gamma_f} \right) \left(1 - e^{-\gamma_i t_f / \tau} \right), \quad (7)$$

where

$$x_2(t) \equiv (x_i - x_f) e^{-2\mu\alpha_f t} + x_{ss} + 2D_1 \left(\frac{1}{\gamma_i} - \frac{1}{\gamma_f} \right) \left(2\mu\alpha_f - \frac{\gamma_f}{\tau} \right)^{-1} \times \left(e^{-\gamma_i t / \mu} - e^{-2\mu\alpha_f t} \right). \quad (8)$$

Note that x_{ss} is given in Eq. (10). For large t_f we have

$$Q_2 \approx \frac{D_1 t_f}{\mu\tau(1 + \alpha_f\mu\tau)}, \quad (9)$$

which is the heat associated with the final steady state.

In Fig. 1c we show that the heat values associated with the trained neural-network protocols interpolate, as a function of t_f , between the values Q_1 and Q_2 . Our conclusion is that this optimization problem is solved by protocols that are rapidly varying, have a variety of functional forms, and display jump discontinuities. As shown in the inset of Fig. 1c and in Fig. S2, these protocols produce values of heat considerably smaller than those associated with the protocols derived in ref. 51. (In the latter figure we also show that it is possible to construct smooth but rapidly-varying protocols that can produce values of heat arbitrarily close to the discontinuous protocols identified by the learning procedure).

The protocols just described are valid solutions to the heat-minimization problem defined in ref. 51. However, some of them are not meaningful in experimental terms. For instance, for small values of t_f , the heat-minimizing protocol is a step function at the final time. This protocol is a solution to the stated problem, but effects no change of the system's microscopic coordinates. All the heat associated with the subsequent transformation of the system is ignored, simply because we have stopped the clock.

We therefore argue that it is more meaningful to search for protocols that minimize heat subject to the requirement of a state-to-state transformation. That is, we require that a specified change in the system's state has occurred. We modify the problem studied in ref. 51 to search for protocols that minimize the mean heat (2) caused by a change of trap stiffness $\alpha_i = 1 \rightarrow \alpha_f = 5$, subject to the completion of a state-to-state transformation (SST) between the initial steady state (4) and that associated with the final-time value of α_f ,

$$x_{ss} = \frac{1}{\alpha_f\mu} \left(\frac{D_1}{\gamma_f} + D \right) \text{ and } y_{ss} = \frac{D_1}{\gamma_f}. \quad (10)$$

As before, we impose the experimentally-motivated constraint $\alpha_i \leq \alpha_\theta(t) \leq \alpha_f$.

To solve this dual-objective problem we choose the evolutionary order parameter

$$\phi = \Delta + c \text{ if } \Delta \geq \Delta_0 \text{ and } \phi = \langle Q \rangle \text{ otherwise.} \quad (11)$$

Here $\Delta^2 \equiv (x_f - x_{ss})^2 + (y_f - y_{ss})^2$ measures the difference between the final-time system coordinates and their values (10) in the final steady state; $\Delta_0 = 10^{-3}$ is the tolerance with which we wish to achieve this steady state; and $c = 100$ is an arbitrary constant whose only role is to make the first clause of (11) always larger than the second (protocols and heat values depend weakly on the value of the threshold Δ , but not in a way that affects our general conclusions). Minimizing (11) will minimize heat emission for a protocol $\alpha(t)$ that in time t_f effects a state-to-state transformation within the precision Δ_0 .

In Fig. 1d we show protocols that minimize heat while achieving SST (see also Sec. S4). These protocols have a variety of forms, which involve rapidly-varying portions and jump discontinuities, and that tend, for large t_f , to the initial-time jump form. For times $t_f \lesssim 1.3$ the

learning algorithm could not identify a protocol that could achieve SST. We discuss how this minimal time is affected by the model parameters in Sec. S3B.

The heat emission associated with these protocols is shown in panel (c). The time t_f for which least heat is emitted is about $t_f = 1.74$, for this choice of D_1 . (For heat optimization alone, the minimum heat is $\langle Q \rangle = 0$, and is shown by Eq. (6) to occur at time $t_f = 0$, a conclusion different to that drawn in Fig. 3 of ref. 51. This strange result follows from the fact that the instruction to minimize heat comes with no requirement that the system change state.)

For comparison, we show the heat emission associated with the linear protocol $\alpha_{\text{lin}}(t) = \alpha_i + (\alpha_f - \alpha_i)(t/t_f)$ (square symbols). The linear protocol emits considerably more heat than learned protocols (note the log scale of the figure), and fails to achieve SST for times $t_f \lesssim 60$.

We conclude that the model of the confined active particle studied in ref. 51 is best controlled by protocols $\alpha(t)$ that are in general rapidly varying and exhibit jump discontinuities – similar to protocols for overdamped passive systems – whether the goal is to minimize heat or to do so while also inducing SST. We note that while the evolutionary training of the neural network is a numerical procedure, the protocols it identified allowed us to derive analytic results for the minimum heat produced for sufficiently small and large trajectory lengths, Eq. (6) and Eq. (7) respectively.

Active particle of variable activity in a trap of variable stiffness State-to-state transformation in least time.

In this section we consider the problem of ref. 52, an active Brownian particle confined by a two-dimensional harmonic potential $U(\boldsymbol{\rho}) = \frac{1}{2}k\rho^2$ with stiffness k . The particle is described by the position vector $\boldsymbol{\rho} = (\rho \cos \phi, \rho \sin \phi)$ and orientation θ , and moves in the direction $\hat{\boldsymbol{e}}(\theta) = (\cos \theta, \sin \theta)$ with constant speed u_0 . Its dynamics is described by the Langevin equation

$$\begin{aligned} \frac{d\boldsymbol{\rho}}{d\tau} &= u_0 \hat{\boldsymbol{e}}(\theta) - \mu k \boldsymbol{\rho} + \sqrt{2D_t} \boldsymbol{\xi}_r(\tau) \\ \frac{d\theta}{d\tau} &= \sqrt{2D_\theta} \xi_\theta(\tau), \end{aligned} \tag{12}$$

where τ is the time; μ is the mobility; D_t and D_θ are translational and rotational diffusion coefficients, respectively; and $\boldsymbol{\xi}_r(\tau)$ and $\xi_\theta(\tau)$ are Gaussian white noise terms with zero mean and unit variance. Upon introducing the dimensionless variables $\boldsymbol{r} \equiv \boldsymbol{\rho} \sqrt{D_\theta/D_t}$ and $t \equiv \tau D_\theta$, Eq. (12) reads

$$\begin{aligned} \frac{d\boldsymbol{r}}{dt} &= \lambda \hat{\boldsymbol{e}}(\theta) - \kappa \boldsymbol{r} + \sqrt{2} \boldsymbol{\xi}_r(t) \\ \frac{d\theta}{dt} &= \sqrt{2} \xi_\theta(t), \end{aligned} \tag{13}$$

where $\kappa \equiv \mu k/D_\theta$ and $\lambda \equiv u_0/\sqrt{D_\theta D_t}$ are dimensionless versions of the spring constant and the self-propulsion speed (λ is the Péclet number). These dimensionless variables are the control parameters of the problem.

The steady-state probability distribution function $\mathcal{P}_{\text{ss}}(r, \chi)$ of the system depends only on $r \equiv |\boldsymbol{r}|$ and $\chi \equiv \theta - \phi$, and is known exactly⁶⁸. The steady state associated with the control-parameter choices κ and λ can be classified as passive or active (Fig. 2a): in the passive phase, the radial probability distribution $P(r)$ is peaked at the trap center, while in the active phase it is peaked at $r > 0$.

This model system is motivated by experiments involving spherical Janus particles, whose self-propulsion speed can be tuned through light intensity⁹, confined in a trap constructed by acoustic waves⁶⁹. For a typical experimental setup the control parameters are bounded as $0 \leq \lambda \leq 11$ and $1 \leq \kappa \leq 7^{9,52,69}$.

The problem described in ref. 52 is to find a time-dependent protocol $(\lambda(t), \kappa(t))$ that obeys the bounds of the previous paragraph

and that minimizes the time t_f required to transform the distribution $\mathcal{P}(r, \chi)$ from a passive steady state at $(\lambda_i, \kappa_i) = (2.5, 4)$ to an active one at $(\lambda_f, \kappa_f) = (5, 4)$. Using an ansatz constrained so that the distribution $\mathcal{P}(r(t), \chi(t))$ has at all times the form of the steady-state distribution $\mathcal{P}_{\text{ss}}(r, \chi)$ with effective values for the control parameters, the authors of that paper found a protocol that completed the state-to-state transformation in time $t_f \approx 0.44$. This was achieved by inserting this ansatz into the Fokker-Planck equation for the time evolution of the probability distribution. In order for analytical calculations to be feasible, one also has to assume that the effective value of the stiffness is constant throughout time. Under these assumptions, the protocol for the activity λ is determined solely by the protocol for the stiffness κ , see Eq. (6) in ref. 52. The resulting protocol is shown in Fig. 2a, b.

A neural-network ansatz for the protocol is free of the restrictions required for the analytical calculations: it does not assume a functional form for the probability distribution at intermediate times, and the protocols for $\lambda(t)$ and $\kappa(t)$ are independent. With a trained neural-network ansatz for the protocol $(\lambda(t), \kappa(t))$, we find that the state-to-state transformation can be achieved about three times as rapidly as it is under the analytic protocol of ref. 52. For a simulation of fixed time t_f we use a genetic algorithm to train the neural network to minimize the order parameter $\phi = \Delta$, the mean-squared error between the target distribution $\mathcal{P}_{\text{ss}}^*(r, \chi)$ associated with the control-parameter values (λ_f, κ_f) and the distribution $\mathcal{P}(r(t_f), \chi(t_f))$ obtained at the end of the simulation. The latter was calculated from 10^5 independent trajectories of (13) under a given neural-network protocol.

The protocol learned by the neural network for time $t_f = 0.16$ is shown in Fig. 2a, b, together with the protocol of ref. 52. Both show sharp jumps in trap stiffness, decreasing it abruptly to its smallest possible value (we discuss the effect of the control parameter ranges on the learned protocols in Sec. S3C). The neural-network protocol achieves the transformation more quickly because it also enacts a sharp jump in activity, setting it to the maximum possible value (the constraints imposed in ref. 52 mean that if one control parameter achieves its maximum value in an abrupt way, the other is not free to do so). Near the end of the learned protocol both parameters are abruptly changed to their final values.

In Fig. 2c, we show the temporal evolution of $P(r)$ for the learned protocol. Starting from an initial distribution peaked at the origin, the peak of $P(r)$ overshoots the peak of the target distribution (they are not at that time of the same shape). The peak of $P(r)$ is later brought back toward the target when stiffness and activity are set to their maximal and minimal values, respectively. Subsequently, both are set to their final values.

In Fig. 2d, we show the final-time distribution of χ for the learned protocol, which matches the target distribution.

In Fig. 2e we show the value of Δ obtained by protocols trained at various fixed simulation times t_f . For times $t_f \geq 0.15$, the learned protocol produces a small constant value of Δ consistent with the value produced by the protocol of ref. 52 (horizontal line). For times $t_f \lesssim 0.15$ the value of Δ increases sharply with decreasing t_f , indicating that the state-to-state transformation cannot be achieved with the same precision.

State-to-state transformation with work extraction. It is possible to extract work during the state-to-state transformation. Setting $t_f = 0.44$, the transformation time of the protocol of ref. 52, we used a genetic algorithm to train a neural network to minimize the objective

$$\phi = \Delta + c \text{ if } \Delta \geq \Delta_0 \text{ and } \phi = \langle W \rangle \text{ otherwise.} \tag{14}$$

Here Δ_0 is the mean-squared error associated with the protocol of ref. 52 (calculated using 10^5 trajectories), and $c = 100$ is an arbitrary

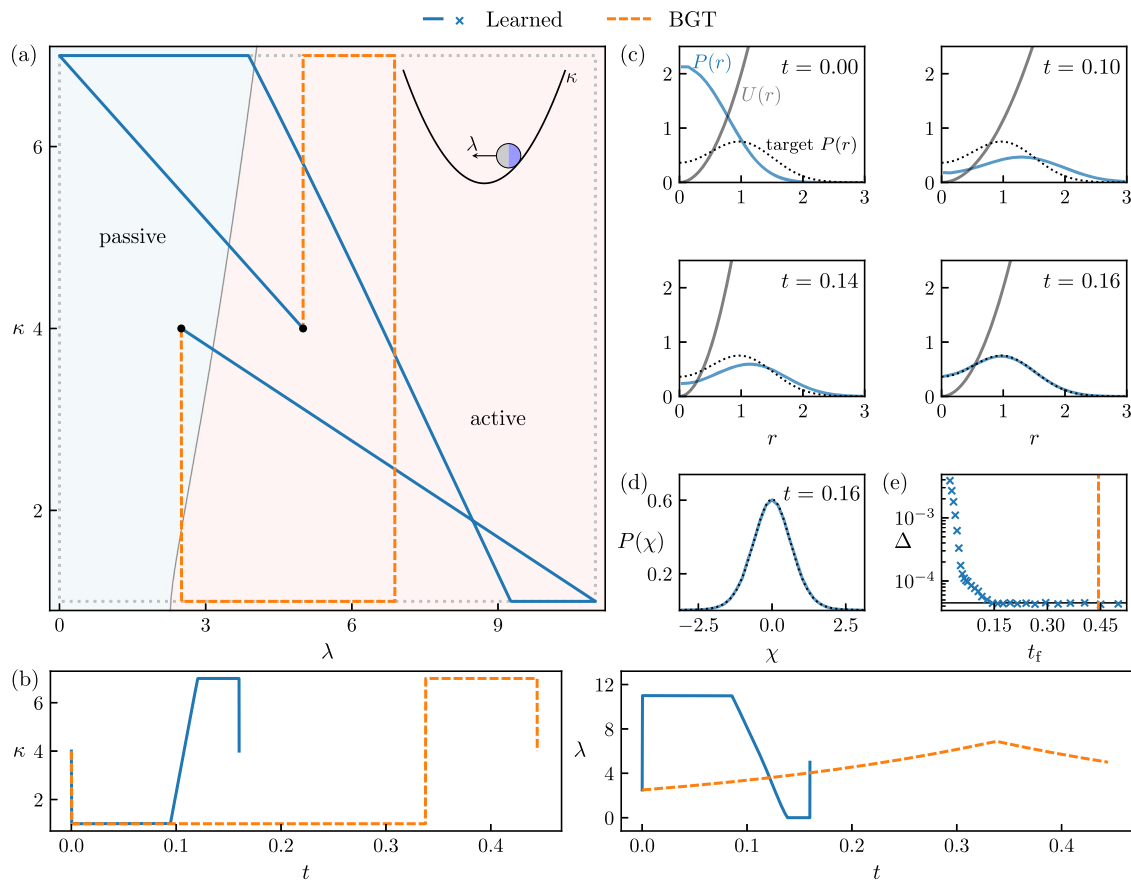


Fig. 2 | State-to-state transformation for a confined active Brownian particle with control protocol $(\lambda(t), \kappa(t))$. **a** Parametric protocols from ref. 52 (orange, “BGT”) and this work (blue). The latter achieves the transformation about three times more rapidly than the former. The black dots denote initial and final points; the dotted line denotes the bounds for the control parameter values. The schematic in the top-right corner is a one-dimensional schematic of this two-dimensional system. **b** The protocols of panel (a) as a function of time **c** Temporal evolution of

the radial distribution function $P(r)$ together with the target distribution (dotted line) and the potential $U(r)$, for the learned protocols shown in (a, b). **d** Final-time distribution for χ . The dotted line is the target distribution. **e** Mean-squared error Δ between the final distribution and the exact solution, averaged over 10^5 trajectories. The dashed vertical line is the transformation time for the protocol of ref. 52, and the horizontal line is the value of Δ associated with that protocol.

constant whose only role is to make the first clause of (14) always larger than the second. The quantity $\langle W \rangle$ is the mean work, in units of μ/D_0 , given by

$$\langle W \rangle = \int_0^{t_f} dt \dot{\kappa} \left\langle \frac{\partial U}{\partial \kappa} \right\rangle = \frac{1}{2} \int_0^{t_f} dt \dot{\kappa} \langle r^2 \rangle. \tag{15}$$

Minimizing (14) will minimize the mean work associated with a protocol $(\lambda(t), \kappa(t))$ that in time t_f effects the state-to-state transformation to a precision Δ_0 .

The protocol learned in this way is shown in Fig. 3a, b, together with the protocol of ref. 52. Panels (c) show the effect of the learned protocol on the radial probability distribution. The neural-network protocol increases κ to its maximum value at the beginning of the protocol. Doing so costs work, but only small amounts because the system is initially in a passive phase and so $\langle r^2 \rangle$ is small. The protocol also increases λ to a large (but sub-maximal) value, which begins to drive the distribution into the active phase, so increasing $\langle r^2 \rangle$. Subsequently, κ is decreased to its target value, causing a decrease of energy and allowing net extraction of work.

Figure 3d shows the work distributions $P(W)$ associated with the learned protocol and that of ref. 52. The latter results in a broad distribution of work values, and on average requires a large input of work

to enact the transformation. By contrast, the work distribution obtained using the learned protocol is sharply peaked at a negative value, and the mean work is negative.

In Fig. 3e we show mean work as a function of time for the two protocols. The learned protocol requires an input of work at early times in order to extract net work at later times. This solution was identified by a genetic algorithm using an order parameter (14) that depends only on quantities evaluated at the final time point. As a result, the protocol is not biased toward any particular functional form. By contrast, greedy reinforcement-learning algorithms, which at all times attempt to reduce the objective function, would (without special shaping of the reward function) be unlikely to find the solution shown here.

Work extraction from confined, interacting active particles

We now consider the case of N interacting active Brownian particles placed within the two-dimensional harmonic trap of the previous section. Particle i evolves according to the Langevin equation

$$\begin{aligned} \frac{d\mathbf{r}_i}{dt} &= \lambda \hat{\mathbf{e}}_i(\theta) - \kappa \mathbf{r}_i - \partial_{\mathbf{r}_i} \sum_{j \neq i} V(\mathbf{r}_{ij}) + \sqrt{2} \boldsymbol{\xi}_r(t) \\ \frac{d\theta_i}{dt} &= \sqrt{2} \xi_\theta(t), \end{aligned} \tag{16}$$

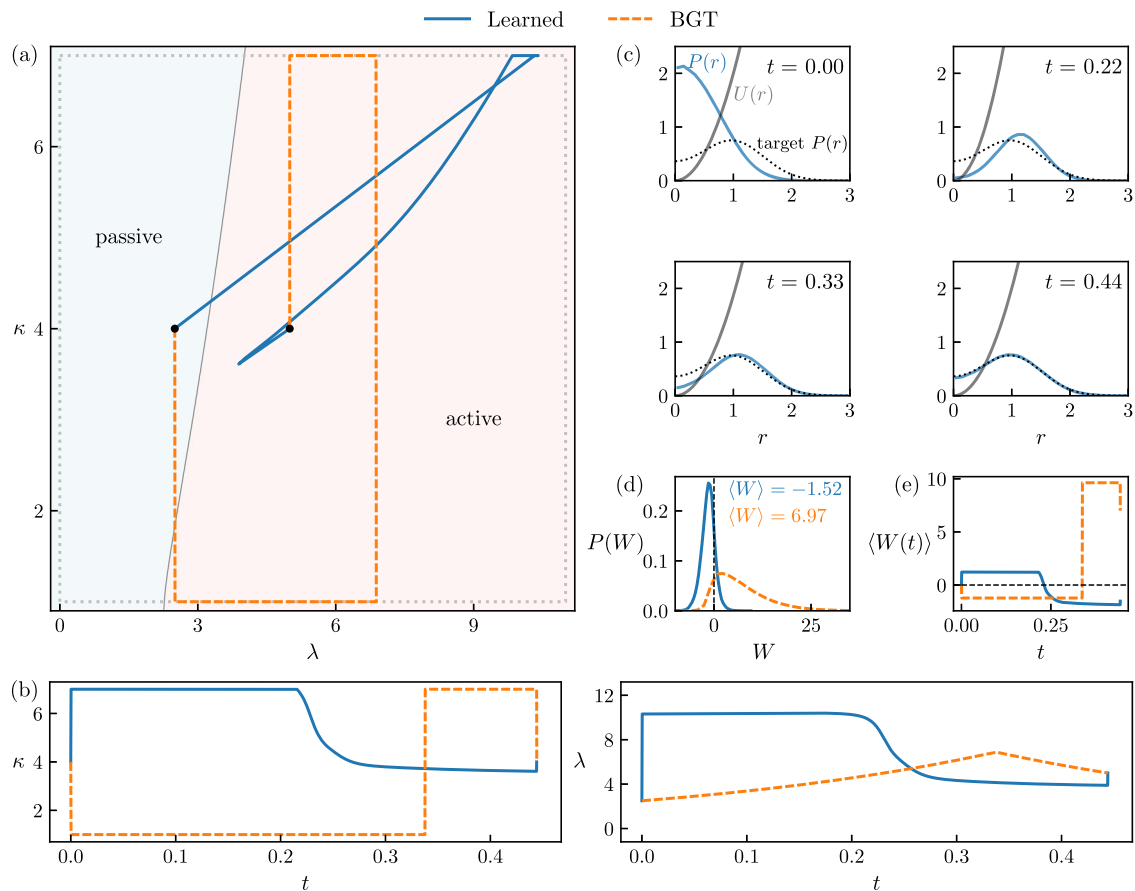


Fig. 3 | State-to-state transformation with extracted work for a confined active Brownian particle with control protocol $(\lambda(t), \kappa(t))$. Similar to Fig. 2, but now the learning algorithm is told to enact the state-to-state transformation of Fig. 2, in time

$t_f = 0.44$, while minimizing work done; see Eq. (14). Panels (a–c) are analogous to those of Fig. 2. **d** Distribution of work for the two protocols. **e** Mean work as a function of time for the two protocols.

whose terms are similar to those of (12) with the addition of the Weeks-Chandler-Andersen interaction

$$V(x) = \begin{cases} 4\epsilon \left[(\sigma/x)^{12} - (\sigma/x)^6 \right] + \epsilon & (x < 2^{1/6}\sigma) \\ 0 & (\text{otherwise}), \end{cases} \quad (17)$$

which takes as its argument the inter-particle separation $r_{ij} \equiv |r_j - r_i|$. We set σ and ϵ to 1.

We wish to learn protocols that minimize the mean work done upon reducing the trap stiffness from $\kappa_i = 5$ to $\kappa_f = 2$, in time t_f , observing the bounds on the control parameter values as in the previous section. Here work is

$$\langle W \rangle = \frac{N}{2} \int_0^{t_f} dt \dot{\kappa} R^2, \quad (18)$$

where $R^2 \equiv N^{-1} \sum_{i=1}^N \langle r_i^2 \rangle$. The angle brackets indicate an average over dynamical trajectories. We start from a steady state at $\lambda_i = 0$, but place no constraints (beyond those of the control-parameter bounds) on the value of λ_f . Such a transformation could be used as part of a cycle for an active engine^{18–20}.

No analytical solutions are known for this many-body system, but a protocol can be learned in exactly the same way as for the single-particle problems considered previously, using a genetic algorithm to train a neural network to minimize $\phi = \langle W \rangle$. The latter was calculated from 10^3 independent trajectories.

In Fig. 4a we show the result of this learning procedure for trajectory time $t_f = 1$ and a number of particles between $N = 1$ and $N = 40$. In

Figs. S7 and S8 we provide additional details of learned protocols for the cases $N = 12$ and $N = 40$. In all cases work can be extracted, $\langle W \rangle < 0$. However, the extracted work per particle is a non-monotonic function of N , attaining a minimum value for $N = 12$. For this particular problem, the many-body system becomes more efficient than the one-body system for $N > 25$. This finding suggests that particular cycles of many-body active engines may function more efficiently with certain particle numbers.

The learned protocols that produce the work values in Fig. 4a initially increase λ to its maximum value. For small N they initially increase κ to its maximum value, while for large N they initially increase κ to close to its maximum value (see Figs. S7 and S8). This initial increase of κ costs work (as with the protocols in the previous section), but the significant increase in R^2 as the activity is increased allows for net work extraction upon the subsequent reduction of κ .

As N is increased from 1, the amount of work that can be extracted per particle initially goes down. This decrease results from the fact that particles repel each other, and so R^2 in the passive initial state is significantly larger for $N > 1$ than for $N = 1$; see Fig. 4b. Increasing κ (at early times) therefore costs more work per particle than for the case $N = 1$. Work can still be extracted from this system, but less efficiently than for the single-body system.

For N sufficiently large, however, the situation changes: R^2 for large λ and κ becomes much larger than it is for a single particle (Fig. 4b, c) (for small N , the value of R^2 for large λ and κ is not much larger than it is for $N = 1$, because particles spread out to form a ring and can adopt a mean radial position similar to that preferred by a single particle). This change allows for greater work extraction per particle when κ is decreased later on in the protocol. For $N > 25$, this

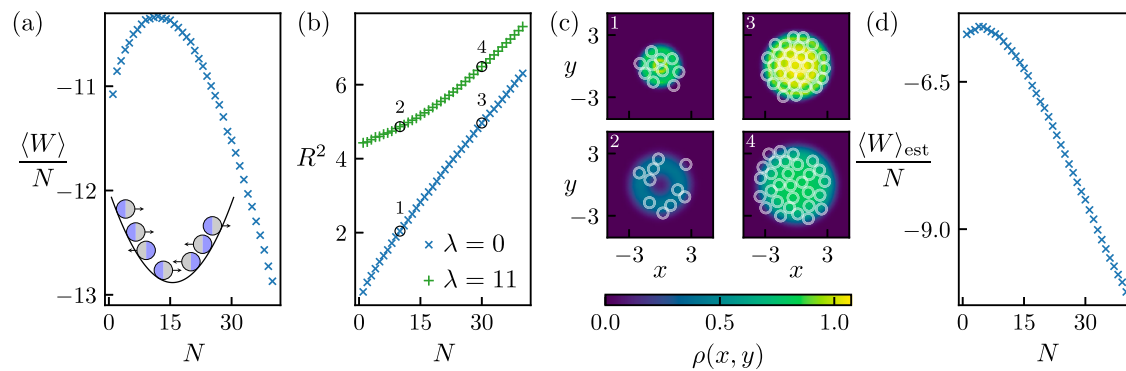


Fig. 4 | Many-body active engine segment. **a** Average work $\langle W \rangle$ per particle, as a function of the number of particles N , for neural-network protocols trained to minimize work by controlling the activity and confining potential of a set of interacting active Brownian particles. The many-body system is more efficient than the one-body system for $N > 25$. **b** Steady-state values of $R^2 \equiv N^{-1} \sum_{i=1}^N \langle r_i^2 \rangle$ as a

function of N . Here $\kappa = 5$, and λ is set to its minimum (blue line) and maximum (green line) values. **c** Steady-state particle density at the control parameters indicated in **(b)**. Instantaneous particle positions for one realization of the system at steady state are shown in white. **d** A simple model of the learned protocols, Eq. (19), captures the non-monotonicity of panel **(a)**.

effect exceeds that described in the previous paragraph, and the many-body system provides more work per particle than a one-body system.

To illustrate the origin of the non-monotonicity seen in Fig. 4, we consider a simplified protocol that instantaneously sets κ and λ to their maximum values, waits until the system reaches a steady state, and then sets $\kappa = \kappa_F$. The work per particle associated with this protocol is given by

$$\frac{\langle W \rangle_{\text{est}}}{N} = \frac{1}{2}(\kappa_{\text{max}} - \kappa_i)R_{\lambda_i, \kappa_i}^2 - \frac{1}{2}(\kappa_F - \kappa_{\text{max}})R_{\lambda_{\text{max}}, \kappa_{\text{max}}}^2, \quad (19)$$

where $R_{\lambda, \kappa}^2$ denotes the steady-state value of R^2 measured at (λ, κ) . In Fig. 4c we show that Eq. (19) is a non-monotonic function of N . It is not a quantitatively accurate model of the learned protocols, but captures one important feature of their behavior.

Extending the simulation time to $t_f = 10$ allows for even greater work extraction. This improvement is achieved by a learned protocol that substantially changes the system's activity twice, from passive to active to passive again (see Fig. S6).

Discussion

We have shown that the efficient control of active-matter systems generally requires protocols containing discontinuities and rapidly-varying features. We discuss how these can be obtained efficiently using neural networks trained with evolutionary methods. We found protocols that achieve particular tasks – enacting state-to-state transformations or changing control parameters with minimal energy input – that were more efficient than those derived recently by constrained analytical methods. Beyond these prototypical one-body problems, we showed that neural-network methods can aid in the design of protocols that achieve extraction of work from many-body active systems. The approach we use is general and can be applied with relatively little modification to a wide variety of active-matter systems. For instance, the optimization of active engines is a quickly growing field of research, garnering interest from both a numerical and experimental perspective^{18–27}, and our results demonstrate that neural networks can lead to optimal protocols for their control, for both interacting and non-interacting active particles. The learning scheme used here can be applied to experiment the way it is applied to simulations, suggesting a way of designing protocols for the efficient manipulation of active matter in the laboratory.

Data availability

The protocols shown in this work are provided in the Supplementary Data 1 file.

Code availability

An example of code for optimizing neural-network protocols for control problems in statistical physics can be found in ref. 70.

References

- Ramaswamy, S. The mechanics and statistics of active matter. *Annu. Rev. Condens. Matter Phys.* **1**, 323–45 (2010).
- Romanczuk, P., Bär, M., Ebeling, W., Lindner, B. & Schimansky-Geier, L. Active brownian particles. *Eur. Phys. J. Spec. Top.* **202**, 1–162 (2012).
- Marchetti, M. C. et al. Hydrodynamics of soft active matter. *Rev. Mod. Phys.* **85**, 1143 (2013).
- Bechinger, C. et al. Active particles in complex and crowded environments. *Rev. Mod. Phys.* **88**, 045006 (2016).
- Cavagna, A. & Giardina, I. Bird flocks as condensed matter. *Annu. Rev. Condens. Matter Phys.* **5**, 183–207 (2014).
- Elgeti, J., Winkler, R. G. & Gompper, G. Physics of microswimmers—single particle motion and collective behavior: a review. *Rep. Prog. Phys.* **78**, 056601 (2015).
- Needleman, D. & Dogic, Z. Active matter at the interface between materials science and cell biology. *Nat. Rev. Mater.* **2**, 1–14 (2017).
- Zöttl, A. & Stark, H. Emergent behavior in active colloids. *J. Phys.: Condens. Matter* **28**, 253001 (2016).
- Buttinoni, I., Volpe, G., Kümmel, F., Volpe, G. & Bechinger, C. Active brownian motion tunable by light. *J. Phys.: Condens. Matter* **24**, 284129 (2012).
- Cates, M. E. & Tailleur, J. Motility-induced phase separation. *Annu. Rev. Condens. Matter Phys.* **6**, 219 (2015).
- O’Byrne, J., Solon, A., Tailleur, J. & Zhao, Y. An Introduction to Motility-induced Phase Separation, in *Out-of-equilibrium Soft Matter*. <https://doi.org/10.1039/9781839169465-00107> (The Royal Society of Chemistry, 2023).
- Toner, J., Tu, Y. & Ramaswamy, S. Hydrodynamics and phases of flocks. *Ann. Phys.* **318**, 170–244 (2005).
- Chaté, H. Dry aligning dilute active matter. *Annu. Rev. Condens. Matter Phys.* **11**, 189–212 (2020).
- Be’er, A. & Ariel, G. A statistical physics view of swarming bacteria. *Mov. Ecol.* **7**, 1–17 (2019).
- Liebchen, B. & Levis, D. Collective behavior of chiral active matter: Pattern formation and enhanced flocking. *Phys. Rev. Lett.* **119**, 058002 (2017).
- Liebchen, B. & Löwen, H. Synthetic chemotaxis and collective behavior in active matter. *Acc. Chem. Res.* **51**, 2982–2990 (2018).

17. Palacci, J., Sacanna, S., Steinberg, A. P., Pine, D. J. & Chaikin, P. M. Living crystals of light-activated colloidal surfers. *Science* **339**, 936–940 (2013).
18. Ekeh, T., Cates, M. E. & Fodor, É. Thermodynamic cycles with active matter. *Phys. Rev. E* **102**, 010101 (2020).
19. Fodor, É. & Cates, M. E. Active engines: Thermodynamics moves forward. *Europhys. Lett.* **134**, 10003 (2021).
20. Kumari, A., Pal, P. S., Saha, A. & Lahiri, S. Stochastic heat engine using an active particle. *Phys. Rev. E* **101**, 032109 (2020).
21. Cocconi, L., Knight, J. & Roberts, C. Optimal power extraction from active particles with hidden states. *Phys. Rev. Lett.* **131**, 188301 (2023).
22. Saha, T. K. et al. Information engine in a nonequilibrium bath. *Phys. Rev. Lett.* **131**, 057101 (2023).
23. Holubec, V., Steffenoni, S., Falasco, G. & Kroy, K. Active brownian heat engines. *Phys. Rev. Res.* **2**, 043262 (2020).
24. Holubec, V. & Marathe, R. Underdamped active brownian heat engine. *Phys. Rev. E* **102**, 060101 (2020).
25. Datta, A., Pietzonka, P. & Barato, A. C. Second law for active heat engines. *Phys. Rev. X* **12**, 031034 (2022).
26. Gronchi, G. & Puglisi, A. Optimization of an active heat engine. *Phys. Rev. E* **103**, 052134 (2021).
27. Pietzonka, P., Fodor, É., Lohrmann, C., Cates, M. E. & Seifert, U. Autonomous engines driven by active matter: Energetics and design principles. *Phys. Rev. X* **9**, 041032 (2019).
28. Caprini, L., Cecconi, F., Maggi, C. & Marconi, U. M. B. Activity-controlled clogging and unclogging of microchannels. *Phys. Rev. Res.* **2**, 043359 (2020).
29. Ghosh, A., Xu, W., Gupta, N. & Gracias, D. H. Active matter therapeutics. *Nano Today* **31**, 100836 (2020).
30. Luo, M., Feng, Y., Wang, T. & Guan, J. Micro-/nanorobots at work in active drug delivery. *Adv. Funct. Mater.* **28**, 1706100 (2018).
31. Norton, M. M., Grover, P., Hagan, M. F. & Fraden, S. Optimal control of active nematics. *Phys. Rev. Lett.* **125**, 178005 (2020).
32. Shankar, S., Scharrer, L. V. D., Bowick, M. J. & Marchetti, M. C. Design rules for controlling active topological defects. *Proc. Natl Acad. Sci.* **121**, e2400933121 (2024).
33. Floyd, C., Dinner, A. R. & Vaikuntanathan, S. Learning to control nonequilibrium dynamics using local imperfect gradients. arXiv preprint arXiv:2404.03798 (2024).
34. Rubenstein, M., Ahler, C., Hoff, N., Cabrera, A. & Nagpal, R. Kilobot: A low cost robot with scalable operations designed for collective behaviors. *Robot. Autonomous Syst.* **62**, 966–975 (2014).
35. Yigit, B., Alapan, Y. & Sitti, M. Programmable collective behavior in dynamically self-assembled mobile microrobotic swarms. *Adv. Sci.* **6**, 1801837 (2019).
36. Balda, A. B., Argun, A., Callegari, A. & Volpe, G. Playing with active matter. arXiv preprint arXiv:2209.04168 (2022).
37. Gupta, D., Klapp, S. H. L. & Sivak, D. A. Efficient control protocols for an active Ornstein-Uhlenbeck particle. *Phys. Rev. E* **108**, 024117 (2023).
38. Guéry-Odelin, D., Jarzynski, C., Plata, C. A., Prados, A. & Trizac, E. Driving rapidly while remaining in control: Classical shortcuts from Hamiltonian to stochastic dynamics. *Rep. Prog. Phys.* **86**, 035902 (2023).
39. Chennakesavalu, S. & Rotskoff, G. M. Probing the theoretical and computational limits of dissipative design. *J. Chem. Phys.* **155**, 194114 (2021).
40. Shankar, S., Raju, V. & Mahadevan, L. Optimal transport and control of active drops. *Proc. Natl Acad. Sci.* **119**, e2121985119 (2022).
41. Monderkamp, P. A., Schwarzendahl, F. J., Klatt, M. A. & Löwen, H. Active particles using reinforcement learning to navigate in complex motility landscapes. *Mach. Learn.: Sci. Technol.* **3**, 045024 (2022).
42. Nasiri, M. & Liebchen, B. Reinforcement learning of optimal active particle navigation. *N. J. Phys.* **24**, 073042 (2022).
43. Nasiri, M., Löwen, H. & Liebchen, B. Optimal active particle navigation meets machine learning. *Europhys. Lett.* **142**, 17001 (2023).
44. Falk, M. J., Alizadehyazdi, V., Jaeger, H. & Murugan, A. Learning to control active matter. *Phys. Rev. Res.* **3**, 033291 (2021).
45. Schmiedl, T. & Seifert, U. Optimal finite-time processes in stochastic thermodynamics. *Phys. Rev. Lett.* **98**, 108301 (2007).
46. Gomez-Marin, A., Schmiedl, T. & Seifert, U. Optimal protocols for minimal work processes in underdamped stochastic thermodynamics. *J. Chem. Phys.* **129**, 024114 (2008).
47. Blaber, S., Louwerse, M. D. & Sivak, D. A. Steps minimize dissipation in rapidly driven stochastic systems. *Phys. Rev. E* **104**, L022101 (2021).
48. Zhong, A. & DeWeese, M. R. Limited-control optimal protocols arbitrarily far from equilibrium. *Phys. Rev. E* **106**, 044135 (2022).
49. Whitelam, S. Demon in the Machine: Learning to Extract Work and Absorb Entropy from Fluctuating Nanosystems. *Phys. Rev. X* **13**, 021005 (2023).
50. Engel, M. C., Smith, J. A. & Brenner, M. P. Optimal Control of Nonequilibrium Systems through Automatic Differentiation. *Phys. Rev. X* **13**, 041032 (2023).
51. Davis, L. K., Proesmans, K. & Fodor, É. Active matter under control: Insights from response theory. *Phys. Rev. X* **14**, 011012 (2024).
52. Baldovin, M., Guéry-Odelin, D. & Trizac, E. Control of Active Brownian Particles: An Exact Solution. *Phys. Rev. Lett.* **131**, 118302 (2023).
53. Whitelam, S. & Tamblyn, I. Learning to grow: Control of material self-assembly using evolutionary reinforcement learning. *Phys. Rev. E* **101**, 052604 (2020).
54. Whitelam, S. & Tamblyn, I. Neuroevolutionary learning of particles and protocols for self-assembly. *Phys. Rev. Lett.* **127**, 018003 (2021).
55. Whitelam, S. How to train your demon to do fast information erasure without heat production. *Phys. Rev. E* **108**, 044138 (2023).
56. Holland, J. H. Genetic algorithms. *Sci. Am.* **267**, 66–73 (1992).
57. Mitchell, M. *An introduction to genetic algorithms* (MIT press, 1998).
58. Such, F. P. et al. Deep neuroevolution: genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. arXiv preprint arXiv:1712.06567 (2017).
59. Whitelam, S., Selin, V., Park, S.-W. & Tamblyn, I. Correspondence between neuroevolution and gradient descent. *Nat. Commun.* **12**, 1–10 (2021).
60. Hornik, K., Stinchcombe, M. & White, H. Multilayer feedforward networks are universal approximators. *Neural Netw.* **2**, 359–366 (1989).
61. Bahri, Y. et al. Statistical mechanics of deep learning. *Annu. Rev. Condens. Matter Phys.* **11**, 501–528 (2020).
62. Rotskoff, G. M. & Crooks, G. E. Optimal control in nonequilibrium systems: Dynamic riemannian geometry of the ising model. *Phys. Rev. E* **92**, 060102 (2015).
63. Gingrich, T. R., Rotskoff, G. M., Crooks, G. E. & Geissler, P. L. Near-optimal protocols in complex nonequilibrium transformations. *Proc. Natl Acad. Sci.* **113**, 10263–10268 (2016).
64. Barros, N., Whitelam, S., Ciliberto, S. & Bellon, L. Learning efficient erasure protocols for an underdamped memory. arXiv preprint arXiv:2409.15050 (2024).
65. Martin, D. et al. Statistical mechanics of active ornstein-uhlenbeck particles. *Phys. Rev. E* **103**, 032607 (2021).
66. Fodor, É., Jack, R. L. & Cates, M. E. Irreversibility and biased ensembles in active matter: Insights from stochastic thermodynamics. *Annu. Rev. Condens. Matter Phys.* **13**, 215–238 (2022).
67. Bonilla, L. L. Active ornstein-uhlenbeck particles. *Phys. Rev. E* **100**, 022601 (2019).
68. Malakar, K., Das, A., Kundu, A., Kumar, K. V. & Dhar, A. Steady state of an active Brownian particle in a two-dimensional harmonic trap. *Phys. Rev. E* **101**, 022610 (2020).
69. Takatori, S. C., De Dier, R., Vermant, J. & Brady, J. F. Acoustic trapping of active matter. *Nat. Commun.* **7**, 10694 (2016).

70. Whitelam, S. Demon in the machine: learning to extract work and absorb entropy from fluctuating nanosystems, <https://doi.org/10.5281/zenodo.13777105> switelam/demon, (2022).

Acknowledgements

Work at the Molecular Foundry was supported by the Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231, and the Stevin Supercomputer Infrastructure, provided by the VSC (Flemish Supercomputer Center), funded by Ghent University, FWO and the Flemish Government – department EWI. C.C. was supported through a Francqui Fellowship of the Belgian American Educational Foundation, and by the US Department of Energy, Office of Science, Office of Basic Energy Sciences Data, Artificial Intelligence and Machine Learning at DOE Scientific User Facilities program under Award Number 34532 (a digital twin for in silico spatiotemporally-resolved experiments).

Author contributions

C.C. performed the simulations. S.W. did the analytic work. C.C. and S.W. discussed the results and wrote the paper.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-024-52878-2>.

Correspondence and requests for materials should be addressed to Corneel Casert or Stephen Whitelam.

Peer review information *Nature Communications* thanks the anonymous reviewers for their contribution to the peer review of this work. A peer review file is available.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024