

Lawrence Berkeley National Laboratory

Molecular Foundry

Title

Low-dissipation self-assembly protocols of active sticky particles

Permalink

<https://escholarship.org/uc/item/1gz9m64b>

Authors

Whitelam, Stephen
Schmit, Jeremy D

Publication Date

2022-12-01

DOI

10.1016/j.jcrysgro.2022.126912

Peer reviewed



Published in final edited form as:

J Cryst Growth. 2022 December 15; 600: . doi:10.1016/j.jcrysgro.2022.126912.

Low-dissipation self-assembly protocols of active sticky particles

Stephen Whitelam^{1,*}, Jeremy D. Schmit^{2,†}

¹Molecular Foundry, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA

²Department of Physics, Kansas State University, Manhattan, KS, 66506, USA

Abstract

We use neuroevolutionary learning to identify time-dependent protocols for low-dissipation self-assembly in a model of generic active particles with interactions. When the time allotted for assembly is sufficiently long, low-dissipation protocols use only interparticle attractions, producing an amount of entropy that scales as the number of particles. When time is too short to allow assembly to proceed via diffusive motion, low-dissipation assembly protocols instead require particle self-propulsion, producing an amount of entropy that scales with the number of particles and the swim length required to cause assembly. Self-propulsion therefore provides an expensive but necessary mechanism for inducing assembly when time is of the essence.

I. INTRODUCTION

Biological processes depend on the assembly of molecular complexes in the face of many constraints [1]. Molecular self-assembly is most effective when it occurs in a short time, with reproducible high yield, and exerts minimal metabolic demands on the cell, but often these constraints are not mutually reconcilable. For instance, the most energy-efficient mode of self-assembly is to allow molecular complexes to find each other through diffusion, but diffusion is slow, requiring timescales that grow as the square of characteristic interparticle separations. Thus while a nanometer protein can diffuse across a micronsized prokaryotic cell in milliseconds, eukaryotic cells, which have sizes on the order of 10 microns, must be subdivided by membrane-bound organelles so that diffusion within them can be efficient. When such subdivision is not possible, such as for transport of neurotransmitters along neuronal axons, cells use active transport by molecular motors to obtain acceptable transit times. However motors require ATP as fuel and so consume energy [1].

In this paper we explore the competing requirements of speed and energy efficiency within a statistical mechanical model of self-assembly. Motivated by the example of intracellular self-assembly we consider a model that represents in a schematic way some of the key features of intracellular complexes, but is not intended to represent a specific biological example. In the model, complexes are represented as particles on a lattice. Particles can

* switelam@lbl.gov . † schmit@phys.ksu.edu .

diffuse and bind to each other. They can also be active, able to move persistently in a particular direction [2–5]. We set them the task of self-assembling in a fixed time, and use evolutionary methods to identify the time-dependent protocols associated with directional motion and binding that achieve this goal with as little entropy production as possible. A corresponding physical scenario relates to the control of activity via the modulation of ATP, combined with the modulation of binding affinities. Examples of the latter include the GTP-dependent stabilization of microtubules [6], turnover of active chaperones [7], and the cyclical progression of molecular motors [8]. We choose entropy production as a convenient measure of the distance from equilibrium of a dynamical trajectory [9, 10], without concern for whether or not its minimization is an organizing principle of biology.

We find that the nature of low-entropy production (or low-dissipation) self-assembly protocols within the model depends on the time allowed for assembly. When there is sufficient time for diffusion to achieve mass transport on the required scales, assembly relies only on energetic interactions. However, when the time allowed for assembly drops below a certain threshold, the combination of diffusion and energetic binding cannot achieve assembly. In this case, particles must undergo self-propulsion, which comes with a large energetic cost. These results, interpreted in the context of cellular assembly, suggest a strong incentive to maintain the system size small enough for diffusion-driven assembly to occur on the required timescales.

In more detail, we encode the time-dependent protocol of the model using a neural network, and use neuroevolutionary methods [11–17] to learn the protocols that achieve a specified degree of self-assembly with the least entropy production (in the medium) possible. There is no guarantee that a given learning scheme will identify the optimum protocol for a given objective, which even in simple systems can involve abrupt changes of control parameters [18–20] (although neural networks of sufficient size can express arbitrary smooth functions, including rapidly-changing ones, and so constitute a good starting point for such a search). To test the learning algorithm we make contact with previous work by learning protocols for low-dissipation magnetization reversal in the Ising model; the protocols learned by the present method are consistent with those learned by path-sampling methods [21, 22], and are essentially equivalent to the least-dissipation pathways for that model. For the active model on which we focus, we find low-dissipation pathways of two kinds, involving self-propulsion or diffusion alone, according to how much time is allotted for assembly. The two mechanisms produce amounts of entropy that differ considerably. We argue that this competition should be present generically, suggesting a principle that natural systems must observe in order to effect a change of phase with least energy expenditure.

In Section II we test the evolutionary framework by making contact with previous work. In Section III we apply it to a model of active particles with interactions. We conclude in Section IV.

II. LOW-DISSIPATION MAGNETIZATION REVERSAL PROTOCOLS IN THE ISING MODEL

We start with the 2D Ising model on a square lattice [23, 24]. The lattice has $N = 50^2$ sites, with periodic boundary conditions in both directions. On each site i is a binary spin $S_i = \pm 1$, and the lattice possesses an energy function

$$E = -J \sum_{\langle ij \rangle} S_i S_j - h \sum_{i=1}^N S_i. \quad (1)$$

Here J (which we set to 1) is the Ising coupling and h is the magnetic field. The first sum in (1) runs over all nearest-neighbor bonds, while the second runs over all lattice sites. We begin with all spins up, giving magnetization $m = N^{-1} \sum_{i=1}^N S_i = 1$, and carry out a Glauber Monte Carlo simulation. At each step of the algorithm a lattice site i is chosen at random, and a proposed change $S_i \rightarrow -S_i$ made. The change is accepted with probability

$$p(\Delta E) = (1 + \exp(\beta \Delta E))^{-1}, \quad (2)$$

where ΔE is the energy change under the proposed move, and is β the reciprocal temperature in units such that $k_B = 1$. If the move is rejected, the original spin state is adopted.

We start with $h = 1$ and $\beta = 1$, and so are below the Ising model critical temperature, $\beta > \beta_c = (2J)^{-1} \ln(1 + \sqrt{2}) \approx 0.44$ [23], in the two-phase region of the phase diagram.

Our goal is to find a time-dependent protocol $\beta(t), h(t)$ that will reverse the magnetization of the model with as little entropy produced as possible (i.e. as can be found by the search algorithm). To do so we encode the protocol using a fully connected single-layer neural network, as shown in Fig. 1(a). The network has one input neuron, which takes the value $t/t_0 \leq 1$, where t is the current number of Monte Carlo steps performed and $t_0 = 100N$ the length of the simulation. The neural network has $K = 1000$ hidden neurons, with hyperbolic tangent activation functions, and two output neurons, which return the values $(\Delta\beta_\theta(t), \Delta h_\theta(t))$, where

$$\Delta\beta_\theta(t) = K^{-1/2} \sum_{i=1}^{K-1} \theta_{3i} \tanh(\theta_{3i+1} t/t_0 + \theta_{3i+2}), \quad (3)$$

and

$$\Delta h_\theta(t) = K^{-1/2} \sum_{i=K}^{2K-1} \theta_{3i} \tanh(\theta_{3i+1} t/t_0 + \theta_{3i+2}). \quad (4)$$

Eqns. (3) and (4) specify the neural network. Here $\theta = (\theta_1, \dots, \theta_1, \dots, \theta_{6K})$ is a $6K$ -dimensional vector of parameters (weights and biases), where K is the number of hidden neurons. Eqns. (3) and (4) can, for large enough K , approximate any smooth functions $\Delta\beta(t)$ and $\Delta h(t)$ [25]. Our goal will be to adjust the parameters θ until the functions $\Delta\beta_\theta(t)$ and $\Delta h_\theta(t)$ expressed by

the neural network achieve our goal of magnetization reversal with as little entropy produced as possible.

The neural network acts at 1000 evenly-spaced time increments, as sketched in Fig. 1(b). Each time it acts, the control parameters are adjusted as

$$\beta(t) \rightarrow \max(0, \beta(t) + \Delta\beta_\theta(t)), \text{ and} \quad (5)$$

$$h(t) \rightarrow h(t) + \Delta h_\theta(t). \quad (6)$$

We quantify the outcome of the simulation using the order parameter

$$\phi \equiv \begin{cases} -m(t_0) - c_1 & (m(t_0) > -1) \\ -c_2\sigma(t_0) & (m(t_0) = -1) \end{cases}. \quad (7)$$

Here $m(t_0)$ is magnetization at the end of the simulation, $c_1 = 2$ and $c_2 = 1/t_0$ are constants, and $\sigma(t_0)$ is the entropy produced during the simulation. There are a number of ways of defining entropy production for a stochastic trajectory [10]; here we consider the entropy change of the medium [9, 26], a sum of terms $\ln(w/w')$ for each process that occurs in the trajectory, where w is the rate for the process and w' the rate for its reverse. In equilibrium, the mean value of this sum is zero.

For the Ising model, the entropy production of the medium for each move is

$$\Delta\sigma = \ln \frac{p(\Delta E)}{p(-\Delta E)} = -\beta\Delta E, \quad (8)$$

using (2), and we define $\sigma(t_0)$ as the sum of terms (8) for each move made (note that both β and the value of h appearing in ΔE are time dependent). The order parameter (7) imposes two objectives: it is maximal for a simulation whose final value of magnetization is -1 , and for which as little entropy as possible was produced. The constants c_1 and c_2 ensure that the value of ϕ for any simulation for which $m(t_0) = -1$ exceeds that for which $m(t_0) \neq -1$, and the former simulations are then distinguished by their entropy produced.

To determine a protocol that maximizes (7), i.e. that results in magnetization reversal with as little entropy produced as possible, we use evolutionary learning on the parameters θ of the protocol-encoding neural networks (sometimes called neuroevolution) [11–17]. This protocol is sketched in Fig. 1(c). We begin with Generation 0, a population of $P = 100$ independent simulations each controlled by a distinct neural network. Each neural network's parameters θ_i are independently and randomly initialized as Gaussian random numbers with zero mean and variance δ_i^2 , i.e. $\theta_i \sim \mathcal{N}(0, \delta_i^2)$. The variance δ_i^2 is also a random number, with $\delta_i \sim |\mathcal{N}(0, \delta_0^2)|$, drawn once for each neural network. The parameter $\delta_0 = 10^{-3}$. Thus each network enacts a different, random protocol. Generation 0 is run for Monte Carlo time t_0 , and the 10 neural networks resulting in the 10 simulations with the largest values of ϕ , Eq. (7), are chosen to be the parents of Generation 1. To construct the 100 members of

Generation 1, we pick 100 times randomly with replacement from the set of 10 parents, and mutate the chosen neural networks by the addition of Gaussian random numbers

$$\theta_i \rightarrow \theta_i + \mathcal{N}(0, \delta_i^2) \quad (9)$$

to each of their parameters. The 100 simulations of Generation 1 are run for Monte Carlo time t_0 , ranked by their values of ϕ , and the 10 highest-ranking neural networks become the parents of Generation 2. As this process continues over evolutionary time, the features of the time-dependent protocols that result in increasingly large values of ϕ are passed to subsequent generations and refined.

In Fig. 2 we show the outcome of this evolutionary learning procedure. Panel (a) shows $-\phi$ for the highest-ranking protocol of each generation as a function of the number n_{gen} of generations of evolutionary learning. After about 10 generations the best protocol has succeeded in reversing the magnetization of the Ising model, and ϕ jumps between the first and second clauses in (7). Thereafter, protocols evolve so as to minimize entropy production while achieving magnetization reversal.

In Fig. 2(b) we show parametric plots of protocols $(\beta(t), h(t))$ for the best protocols from 5 different generations. Protocols start at the point (1,1). The Ising model (infinite system size) critical point and first-order phase transition line are shown as a black dot and black dashed line, respectively. Comparison with panel (a) shows that the first protocols that achieve magnetization reversal pass through the phase transition line. As their entropy production is progressively reduced, protocols move toward the critical point and then pass around it, avoiding the phase transition line and the vicinity of the critical point. To provide context for the scale of Fig. 2(a), note that the entropy produced by flipping all spins following an instantaneous change from $(\beta, h) = (1, 1)$ to $(1, -1)$ (near where the evolved protocols end) is $2N$, corresponding to a value of $\phi = 2N/t_0 = 2N/(100N) = 0.02$, or about the value of the plateau immediately following the jump. The final values of entropy production identified by evolutionary learning are about an order of magnitude smaller.

The evolved low-dissipation magnetization-reversal protocols agree with those obtain by path-sampling techniques [21, 22], showing that the way to reverse magnetization with least dissipation is to heat, reverse the field, and then cool, avoiding the large energy changes that result from the presence of finite surface tension or the large fluctuations near the critical point. The inset to panel (a) emphasizes an additional result of Ref. [22], that several slightly different protocols result in similar entropy production. There may be one optimal protocol, but there are many low-lying protocols that are essentially as good (the number of nature of such protocols can be analyzed to provide insight about the structure of the trajectory ensemble [22, 27]).

In panel (c) we show magnetization and protocol as a function of time, together with time-ordered snapshots, from the simulation with the lowest dissipation obtained by the learning procedure. Below that we show a protocol from a second set of evolutionary learning simulations that were instructed to find magnetization-reversal trajectories with the *greatest* possible entropy production. In this case the strategy is to drive the system through

the phase-transition line and reverse the field only when the temperature is as low as can be achieved on the allotted timescale. The resulting dynamics is a nucleation-like mechanism in the presence of large surface tension and large thermodynamic driving force for phase change.

Results in this section used the Glauber acceptance probability, Eq. (2). Replacing this with the Metropolis acceptance probability $p(\Delta E) = \min(1, \exp[-\beta\Delta E])$ changes the precise dynamics of evolution of the model but not the underlying thermodynamic landscape. In Fig. 3 we show that evolutionary learning of low-dissipation protocols for Metropolis dynamics is qualitatively similar to that for Glauber dynamics.

III. LOW-DISSIPATION SELF-ASSEMBLY PROTOCOLS FOR INTERACTING ACTIVE PARTICLES

Having established that the neuroevolution framework can identify protocols similar to those obtained by other methods, we now consider self-assembly in a model of active particles with interactions. The model is that of Ref. [28] with pairwise nearest-neighbor interactions added, and is summarized in Fig. 4. On a square lattice of size $N = 50^2$ with periodic boundaries we consider $M = 750$ particles, a packing fraction of 0.3. Particles may not overlap. Each particle bears an orientation vector (shown white in the figure) that points toward one of the 4 neighboring sites, and rotates right or left with rate $D = 0.1$. Isolated particles move forward (in the direction of the orientation vector) with rate v_+ , and in each of the other 3 directions with rate $v_0 = 1$. For these parameters the Péclet number, the dimensionless number quantifying the relative scale of self-propulsion and diffusion, is $Pe = (v_+ - v_0)/(2D) = 5(v_+ - 1)$. Two nearest-neighbor particles possess an interaction energy $-\epsilon$, where $\epsilon > 0$ indicates an attraction and $\epsilon < 0$ a repulsion. Interactions modify the rates for motion via a multiplicative term $\min(1, e^{-\Delta E})$, ΔE being the energy change for the move. This model can self-assemble via motility-induced phase separation [29], for sufficiently large Pe , or via conventional phase separation, for sufficiently large ϵ . Our goal is to determine which combinations of these mechanisms are required to promote low-dissipation assembly.

The system advances via a continuous-time Monte Carlo algorithm (also known as kinetic Monte Carlo or the Gillespie algorithm). If r_i is the rate for process i then at each step of the algorithm the process i is chosen with probability $r_i/\sum_i r_i$, the sum running over all possible processes. There are $6M$ possible processes (translations or rotations of each particle in the simulation box), with the rate for translations that result in an overlap being zero. Upon enacting the chosen process, time is advanced by an amount $-\ln \eta/\sum_i r_i$, where η is a random number uniformly distributed on $(0,1]$. Simulations are run for time t_0 , which we varied from of order 100 to of order 105; for the chosen parameters, the characteristic time taken for an isolated particle to diffuse its own length is 1.

At the end of the simulation we evaluate the order parameter

$$\phi \equiv \begin{cases} f_4(t_0) & (f_4(t_0) \leq 0.7) \\ c_1 - c_2 \sigma(t_0) & (f_4(t_0) > 0.7) \end{cases} \quad (10)$$

Here f_4 is the fraction of particles with 4 particles as neighbors, our chosen measure of assembly. $\sigma(t_0)$ is the entropy produced over the course of the simulation. We consider the entropy production of the medium [26], where the entropy produced by each move is $\Delta\sigma = \ln(r/r')$. Here r is the rate for the enacted process and r' the rate for its reverse, which for the various move types is

$$\Delta\sigma = \begin{cases} 0 & (\text{rotation}) \\ -\beta\Delta E & (\text{sideways}) \\ \ln v_+ - \beta\Delta E & (\text{forward}) \\ -\ln v_+ - \beta\Delta E & (\text{backward}) \end{cases} \quad (11)$$

We define $\sigma(t_0)$ as the sum of values of (11) for each move made during the simulation. Finally, the constants $c_1 = \phi_0 = 10^3$ and $c_2 = 1/(100N)$ ensure a separation of scales [30] between the two clauses of (10), and enforce the following dual objective: the instruction to maximize Eq. (10) is the instruction to have the system self-assemble so that 70% of its particles are in a solid-like environment, and, if so, have assembly happen with the least dissipation possible.

We use neuroevolutionary learning to find a time-dependent protocol to maximize ϕ , as in Section II. We start each simulation with parameters $(\text{Pe}, \epsilon) = (0, 0)$, and encode the time-dependent protocol $(\text{Pe}(t), \epsilon(t))$ using a neural network. Similar to before, the input to the neural network is the scaled time t/t_0 , it acts 1000 times within the simulation, and its output each time it acts is $(\Delta\text{Pe}_\theta(t), \Delta\epsilon_\theta(t))$. After every action the simulation protocol is updated as

$$\epsilon(t + \Delta t) \rightarrow \epsilon(t) + \Delta\epsilon_\theta(t), \quad \text{and} \quad (12)$$

$$\text{Pe}(t + \Delta t) \rightarrow \max(0, \text{Pe}(t) + \Delta\text{Pe}_\theta(t)). \quad (13)$$

The evolutionary search on the parameter set θ proceeds as in the previous section, except that we use 50 individuals in each generation and choose the best 5 to propagate to the next generation.

In Fig. 5 we show the results of 5 sets of evolutionary learning simulations, each of which used a different physical time t_0 for the simulations (not all results are shown in each panel). In panel (a) we show the order parameter (10) as a function of evolutionary time for simulations using three values of t_0 ; in each panel, the arrow denotes the direction of increasing t_0 . Each ϕ attains a value consistent with the second clause in (10), showing that at least 70% of particles have self-assembled into a solid-like environment. (This jump happens after fewer evolutionary steps for larger values of t_0 , because for longer simulation times the numerical values of the control parameters (Pe, ϵ) need to be less large in order to induce assembly, and so the neural network needs fewer evolutionary generations in order to attain

those values.) The largest value of ϕ for the largest value of t_0 is larger than those for the other two values of t_0 , showing that assembly in the former case produces less entropy. We will discuss soon why this is.

In Fig. 5(b) we show parametric plots of the protocols $(\text{Pe}(t), \epsilon(t))$ learned after several evolutionary generations (250 in the case of the two smaller times shown, 100 in the other case). For the longer time, the Péclet number remains zero, and only the interparticle attraction is used. In the other two cases the Péclet number is increased to values that promote motility-induced phase separation (a steady-state value of $\text{Pe} \gtrsim 100$ is needed to produce a solid fraction of 0.7 for packing fraction 0.3 [28]), and then the interparticle attraction is made large. Panels (c) and (d) show the outcome of these and similar protocols, showing the fraction f_4 of solid-like particles and the entropy σ produced within a trajectory. The latter panel shows two sets of distinctly different entropy production rates. As we shall describe, these result from the presence or absence of self-propulsion.

Fig. 6 confirms this behavior, showing the smallest values of entropy produced at the end of the simulation, $\sigma(t_0)$, after several generations of evolutionary learning using a range of values of t_0 . A crossover between two different types of behavior is evident.

These results reveal a change of low-dissipation assembly mechanism upon changing t_0 . When the simulation time t_0 is sufficiently large, assembly of the required amount of material can proceed in the presence of diffusive motion and interparticle attractions. But diffusion alone cannot achieve assembly of the required amount of material when t_0 is made too small, in which case self-propulsion is required. However, propulsion is a costly mechanism, and results in considerable entropy production.

To see this, and to understand the scale of Fig. 6, consider that interaction-driven diffusive assembly requires the gain of about $fNz/2$ interparticle bonds, where $f = 0.7$ is the required solid-like fraction, $N = 750$ is the number of particles, and $z = 4$ is the coordination number. Assuming that these bonds are formed in the presence of energy scales of order 1 (in units such that $k_B T = 1$), the total entropy produced (the sum of bond-energy changes) is about 10^3 , consistent with the location of the lower plateau in Fig. 6. By contrast, self-propelled particles produce entropy of order $\ln \text{Pe}$ per step. If $N = 750$ particles each swim of order the box length $L = 50$ at $\text{Pe} = 100$ they would produce an amount of entropy of order 105, consistent with the scale of the higher plateau in Fig. 6.

We expect such a tradeoff to be present in general. The entropy production for energy-driven diffusive assembly scales as N , the number of particles, while that for motility-induced assembly scales as $N \times \ell$ the latter factor being the characteristic particle swim length. Combined with the fact that the second mechanism is faster than the first, protocols that achieve low-dissipation assembly will tend to favor diffusion when time is plentiful, and self-propulsion when it is not. Note that we have not considered the cost to change the protocol, which, for specific physical systems, may enhance or oppose this crossover.

The time at which the crossover occurs will be a function of particle concentration, which sets the characteristic interparticle separation. The timescale for diffusive encounters scales

as the square of this separation, while that for encounters mediated by self-propulsion scales linearly. As a result, we expect the crossover shown in Fig. 6 to shift to larger values of t_0 as particle concentration is reduced.

Time-ordered snapshots of examples of low-dissipation assembly are shown in Fig. 7. There is no strong visual indication of the large difference in entropy production between these examples (self-propulsion would only be evident in a movie), but the example involving self-propulsion (top) assembles considerably faster than that involving only diffusion (bottom).

In Fig. 8 we show the outcome of three sets of evolutionary learning simulations under different constraints. All are done for simulation time $t_0 = 2000$. In panel (a) the neural network can control both $Pe(t)$ and $\epsilon(t)$, but in panels (b) and (c) it can control only $Pe(t)$ or $\epsilon(t)$, respectively. The strategy in (a) is to drive Pe large in order to cause motility-induced phase separation, and then drive ϵ large in order to freeze the system and prevent unnecessary motion. In panel (b) the freezing mechanism is unavailable, and so the increase of Pe is delayed as long as possible. Freezing is an efficient strategy: in the inset to (b) we show that case (b) produces more entropy than case (a). In case (c), diffusive motion cannot achieve the required amount of assembly.

Some protocols, particularly near the crossover in Fig. 6, achieve assembly using moderate values of and small values of Pe , of order 10, not large enough to induce motility-induced phase separation but enough to speed mass transport and enable assembly that could not have happened via diffusion alone. However, because entropy production scales linearly with time but only logarithmically with Pe , we find that the more usual strategy is to quickly drive Pe large, above 100, in order to cause motility-induced phase separation, then drive ϵ large in order to freeze the system and stop excess entropy being produced.

IV. CONCLUSIONS

Motivated by the competition between speed and energy efficiency in the context of biological self-assembly, we have explored self-assembly in a generic model of active particles with interactions. We have used neuroevolutionary learning to find protocols that achieve assembly within a particular time limit, and that produce as little entropy as possible. Enhanced-sampling methods such as transition-path sampling [31] and other forms of protocol learning such as gradient-based reinforcement learning approaches could be applied to this problem [32, 33]. In the context of reinforcement learning, evolutionary methods are relatively simple to implement (they do not require gradient computation through time) and are natural when asked to maximize an order parameter that is specified only at the final time point of a trajectory, a ‘sparse-reward problem’ in the language of reinforcement learning (evolutionary methods can also be used with larger numbers of parameters than we have considered here [15]).

We have found that when time is plentiful, protocols use only attractive interactions and diffusive motion. When time is restricted, self-propulsion is required in order to allow the

required degree of assembly to happen, but propulsion is expensive and results in more entropy being produced than in the diffusive case.

This competition is likely to exist generically, because it depends only on simple scaling arguments and not on specific molecular details. These results, interpreted in the context of cellular assembly, suggest a strong incentive to maintain the system size small enough for diffusion-driven assembly to occur on the required timescales. In specific systems the question of how costly it is to impose the required protocols is also relevant, and these costs may enhance or oppose the competition identified.

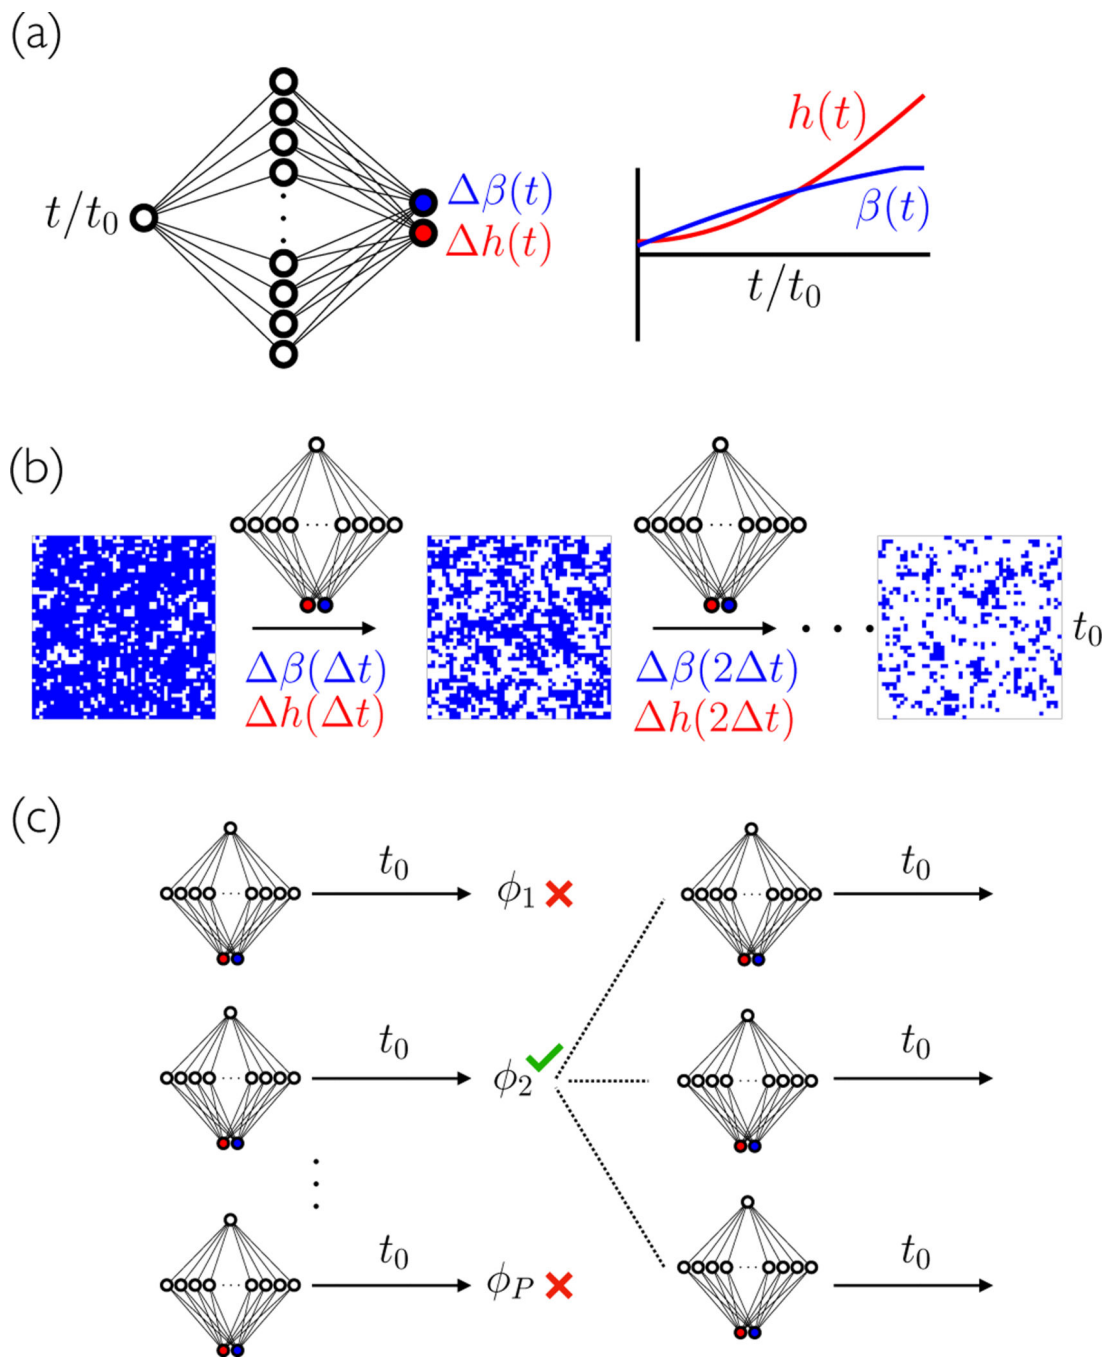
ACKNOWLEDGMENTS

This work was done as part of a User project at the Molecular Foundry at Lawrence Berkeley National Laboratory, supported by the Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. JDS acknowledges support from NIH grant R01GM141235.

References

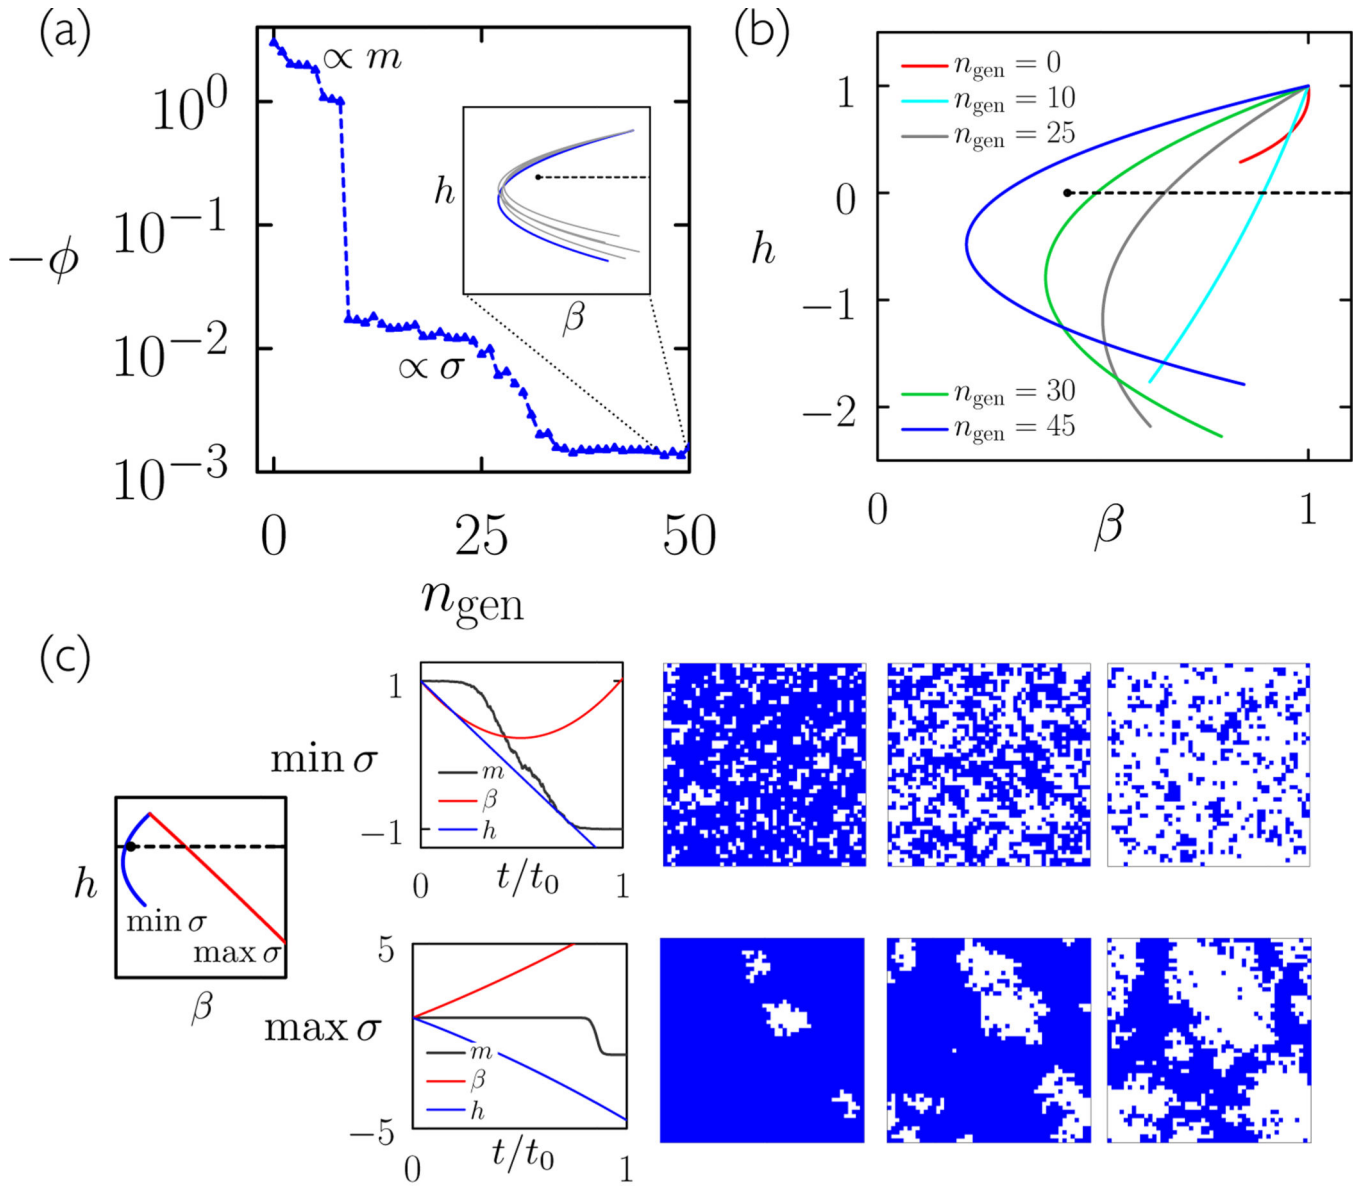
- [1]. Phillips Rob, Kondev Jane, and Theriot Julie, *Physical Biology of the Cell* (Garland Science, 2008) p. 800.
- [2]. Needleman Daniel and Dogic Zvonimir, “Active matter at the interface between materials science and cell biology,” *Nature reviews materials* 2, 1–14 (2017).
- [3]. Ramaswamy Sriram, “The mechanics and statistics of active matter,” *Annu. Rev. Condens. Matter Phys* 1, 323–345 (2010).
- [4]. Fodor Étienne and Marchetti M Cristina, “The statistical physics of active matter: From self-catalytic colloids to living cells,” *Physica A: Statistical Mechanics and its Applications* 504, 106–120 (2018).
- [5]. Hagan Michael F and Baskaran Aparna, “Emergent self-organization in active materials,” *Current opinion in cell biology* 38, 74–80 (2016). [PubMed: 26971116]
- [6]. Desai Arshad and Mitchison Timothy J., “Microtubule polymerization dynamics,” *Annu. Rev. Cell Dev. Biol* 13, 83–117 (1997). [PubMed: 9442869]
- [7]. Thirumalai D, George H. Lorimer, and Hyeon Changbong, “Iterative annealing mechanism explains the functions of the GroEL and RNA chaperones,” *Protein Sci.* 29, 360–377 (2020). [PubMed: 31800116]
- [8]. Gennerich Arne and Vale Ronald D., “Walking the walk: how kinesin and dynein coordinate their steps,” *Curr. Opin. Cell Biol* 21, 59–67 (2009). [PubMed: 19179063]
- [9]. Schnakenberg Jürgen, “Network theory of microscopic and macroscopic behavior of master equation systems,” *Reviews of Modern physics* 48, 571 (1976).
- [10]. Ruelle David, “Entropy production in nonequilibrium statistical mechanics,” *Communications in Mathematical Physics* 189, 365–371 (1997).
- [11]. Holland John H, “Genetic algorithms,” *Scientific american* 267, 66–73 (1992). [PubMed: 1411454]
- [12]. Fogel David B and Stayton Lauren C, “On the effectiveness of crossover in simulated evolutionary optimization,” *BioSystems* 32, 171–182 (1994). [PubMed: 7919114]
- [13]. Floreano Dario, Dürr Peter, and Mattiussi Claudio, “Neuroevolution: from architectures to learning,” *Evolutionary intelligence* 1, 47–62 (2008).
- [14]. Salimans Tim, Ho Jonathan, Chen Xi, Sidor Szymon, and Sutskever Ilya, “Evolution strategies as a scalable alternative to reinforcement learning,” *arXiv preprint arXiv:1703.03864* (2017).
- [15]. Felipe Petroski Such Vashisht Madhavan, Conti Edoardo, Lehman Joel, Stanley Kenneth O, and Clune Jeff, “Deep neuroevolution: genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning,” *arXiv preprint arXiv:1712.06567* (2017).

- [16]. Whitelam Stephen and Tamblyn Isaac, "Learning to grow: Control of material self-assembly using evolutionary reinforcement learning," *Physical Review E* 101, 052604 (2020).
- [17]. Whitelam Stephen and Tamblyn Isaac, "Neuroevolutionary learning of particles and protocols for self-assembly," *Physical Review Letters* 127, 018003 (2021).
- [18]. Schmiedl Tim and Seifert Udo, "Optimal finite-time processes in stochastic thermodynamics," *Physical Review Letters* 98, 108301 (2007).
- [19]. Solon Alexandre Pand Horowitz Jordan M, "Phase transition in protocols minimizing work fluctuations," *Physical review letters* 120, 180605 (2018).
- [20]. Ye Zhuolin, Cerisola Federico, Abiuso Paolo, Anders Janet, Perarnau-Llobet Martí, and Holubec Viktor, "Optimal finite-time heat engines under constrained control," arXiv preprint arXiv:2202.12953 (2022).
- [21]. Rotskoff Grant M and Crooks Gavin E, "Optimal control in nonequilibrium systems: Dynamic riemannian geometry of the ising model," *Physical Review E* 92, 060102 (2015).
- [22]. Gingrich Todd R, Rotskoff Grant M, Crooks Gavin E, and Geissler Phillip L, "Near-optimal protocols in complex nonequilibrium transformations," *Proceedings of the National Academy of Sciences* 113, 10263–10268 (2016).
- [23]. Onsager Lars, "Crystal statistics. i. a two-dimensional model with an order-disorder transition," *Physical Review* 65, 117 (1944).
- [24]. Binney James J, Dowrick NJ, Fisher AJ, and Newman M, *The theory of critical phenomena: an introduction to the renormalization group* (Oxford University Press, Inc., 1992).
- [25]. Cybenko George, "Approximation by superpositions of a sigmoidal function," *Mathematics of control, signals and systems* 2, 303–314 (1989).
- [26]. Seifert Udo, "Entropy production along a stochastic trajectory and an integral fluctuation theorem," *Physical Review Letters* 95, 040602 (2005).
- [27]. Sriraman Saravanapriyan, Kevrekidis Ioannis G, and Hummer Gerhard, "Coarse master equation from bayesian analysis of replica molecular dynamics simulations," *The Journal of Physical Chemistry B* 109, 6479–6484 (2005). [PubMed: 16851726]
- [28]. Whitelam Stephen, Klymko Katherine, and Mandal Dibyendu, "Phase separation and large deviations of lattice active matter," *The Journal of Chemical Physics* 148, 154902 (2018).
- [29]. Cates Michael E and Tailleur Julien, "Motility-induced phase separation," *Annu. Rev. Condens. Matter Phys* 6, 219–244 (2015).
- [30]. The exact numerical values of c_1 and c_2 have no effect on the outcome of learning, provided that the smallest possible value of the second clause of (10) is always greater than the largest possible value of the first clause.
- [31]. Bolhuis Peter G., Chandler David, Dellago Christoph, and Geissler Phillip L., "Transition path sampling: Throwing ropes," *Annu. Rev. Phys. Chem* 53, 291–318 (2002). [PubMed: 11972010]
- [32]. Kaelbling Leslie Pack, Littman Michael L, and Moore Andrew W, "Reinforcement learning: A survey," *Journal of artificial intelligence research* 4, 237–285 (1996).
- [33]. Sutton Richard S and Barto Andrew G, *Reinforcement learning: An introduction* (2018).

**FIG. 1.**

Schematic of the neuroevolutionary learning method. (a) A time-dependent protocol, which in Section II comprises $(\beta(t), h(t))$, is encoded by a neural network whose input is scaled time t/t_0 and whose output is the vector $(\beta(t), h(t))$. (b) The protocol is imposed within a molecular simulation by the neural network, which acts at 1000 evenly-spaced time increments $\Delta t = t_0/1000$. (c) An evolutionary algorithm searches for the neural network (and hence protocol) that maximizes a particular order parameter, ϕ , Eq. (7). A population of P randomly-initialized networks, each controlling one molecular simulation, results in P

outcomes ϕ_i , $i \in [1, P]$. In this schematic the second outcome (“individual”) is the best, $\phi_2 > \phi_{j \neq 2}$, and so that neural network is cloned and mutated in order to build the second generation of the evolutionary scheme. In Section II we set $P = 100$, and the top 10 individuals are retained from one generation to the next.

**FIG. 2.**

Evolutionary learning identifies low-dissipation magnetization-reversal protocols in the Ising model. (a) Order parameter (7) as a function of evolutionary time (inset: parametric protocols for generations 46–50). (b) Parametric protocols $(\beta(t), h(t))$ for 5 different generations; the black dot and dashed line are the Ising model critical point and first-order phase transition line, respectively. (c) Contrasting low- and high-dissipation magnetization-reversal protocols (from left to right: parametric protocols; time-dependent protocols; time-ordered snapshots)

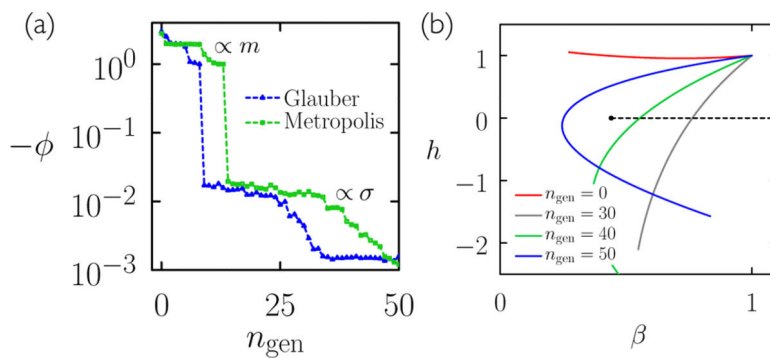


FIG. 3. Protocol learning for the Ising model is qualitatively unchanged upon going from Glauber to Metropolis dynamics. (a) Reproduction of Fig. 2(a) (blue) together with similar results using Metropolis Monte Carlo dynamics (green). (b) Analog of Fig. 2(b) for Metropolis dynamics. Again, the learning algorithm identifies low-dissipation trajectories as those encircling the critical point (black dot) and avoiding the first-order transition line (black dashed line).


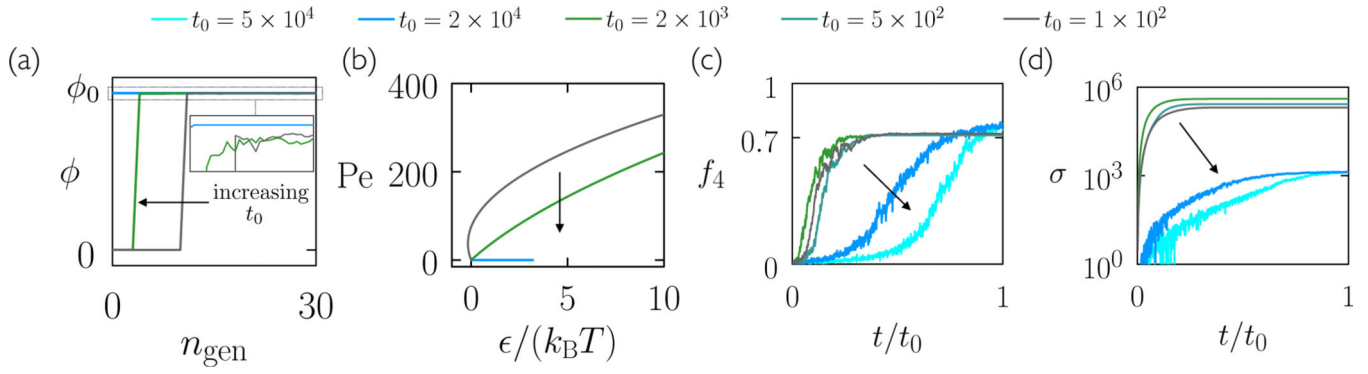


FIG. 4. Schematic of the model of active particles with interactions. (a) Rates for motion of an isolated particle. (b) Example of a rate influenced by the pairwise interparticle interaction; ΔE is the energy change of the move.

**FIG. 5.**

Neuroevolutionary learning simulations instructed to identify a protocol $(Pe(t), \epsilon(t))$ promoting assembly with least dissipation in the model of active particles with interactions. Various values of t_0 , the physical simulation time, are considered. (a) Order parameter (10) as a function of evolutionary time (here and in other panels the arrow points in the direction of increasing t_0). (b) Parametric protocols $(Pe(t), \epsilon(t))$ learned after several evolutionary generations. (c) Solid fraction f_4 and (d) entropy produced as a function of simulation time, using protocols learned after several evolutionary generations.




FIG. 6. Least entropy required to induce assembly, as a function of physical time t_0 , using protocols learned after several evolutionary generations. Two distinct mechanisms are evident.


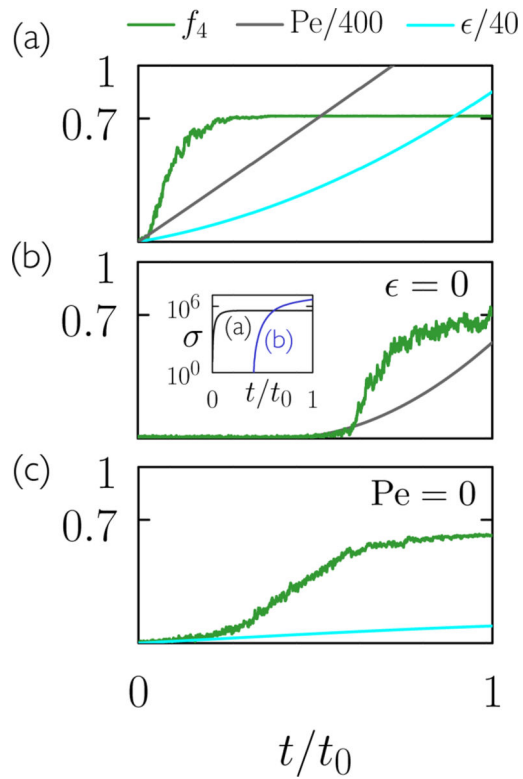


FIG. 7. Time-ordered snapshots of low-dissipation assembly involving self-propulsion (top) and diffusion only (bottom).

**FIG. 8.**

Constrained assembly protocols identified by neuroevolutionary learning. (a) A low-dissipation protocol in which a neural network can control both $Pe(t)$ and $\epsilon(t)$. (b) A low-dissipation protocol in which the neural network controls only $Pe(t)$; this produces more entropy than case (a) (inset). (c) When the neural network can control only $\epsilon(t)$, the required amount of material cannot assemble in the allotted time.