
Cooperative multi-agent reinforcement learning for high-dimensional nonequilibrium control

Shriram Chennakesavalu
Department of Chemistry
Stanford University
shriramc@stanford.edu

Grant M. Rotskoff
Department of Chemistry
Stanford University
rotskoff@stanford.edu

Abstract

Experimental advances enabling high-resolution external control create new opportunities to produce materials with exotic properties. In this work, we investigate how a multi-agent reinforcement learning approach can be used to design external control protocols for self-assembly. We find that a fully decentralized approach performs remarkably well even with a "coarse" level of external control. More importantly, we see that a partially decentralized approach, where we include information about the local environment allows us to better control our system towards some target distribution. We explain this by analyzing our approach as a partially-observed Markov decision process. With a partially decentralized approach, the agent is able to act more presciently, both by preventing the formation of undesirable structures and by better stabilizing target structures as compared to a fully decentralized approach.

1 Context and related work

Designing nanoscale structures that are tuned to have specific material properties or dynamical behavior is a longstanding goal in the molecular sciences [1–4]. While advances in nanofabrication allow for increasingly intricate design, these approaches are often limited to specific materials and require homogeneity of the molecular components. Many nanoscale systems, however, “self-assemble” in suitable conditions or in the presence of appropriate external driving forces [5–7]. Exploiting self-assembly to manufacture nanoscale components presents its own challenges, as the system must be designed to spontaneously and rapidly form a given target structure. The starting materials can become kinetically trapped, forming unfavorable structures if left to self-assemble without external control [7, 8]. Here we examine the problem of guiding self-assembly to target high yield of the target structure via external nonequilibrium driving forces, which we refer to as a “protocol”. Reinforcement learning (RL) offers a promising toolkit for computing optimal protocols, but there has been little systematic investigation of approaches based on RL in this context.

Here, we consider two minimal models of molecular self-assembly in which particles evolve according to a stochastic, nonequilibrium dynamics. In both models, clusters of particles form spontaneously under appropriate external conditions, and we seek to design protocols that optimize the size of particle clusters to a given target. The external driving forces (e.g., temperature, light intensity) that we consider can be modulated as a function of both space and time (Fig. 1). In the systems considered here, the protocol is consequently high-dimensional because the external field is modulated on a grid with high spatial resolution. Because of the exponentially large action space, it is natural to employ multi-agent reinforcement learning (MARL) in this context. Briefly, MARL extends the RL paradigm of an agent interacting with an environment to one in which multiple agents interact with the same environment concurrently (and possibly with each other).

Several works have investigated reinforcement learning to control condensed phase dynamics, mostly in equilibrium settings [9–12]. These approaches modulate the inter-particle interaction potentials between constituents—fundamentally altering the constituent material—rather than externally modulating them as we do here. Recently, Ref. [13] investigated controlling dynamics of a nonequilibrium active matter system with a low-dimensional protocol using standard RL techniques. Multi-agent reinforcement learning has a substantial literature [14, 15]; strategies that treat each cooperative agent independently are among the most relevant to the current study, including [16–18].

Main contributions:

- We argue that within the framework of a partially observed Markov Decision Process, multi-agent reinforcement learning with improved state estimates lead to improved value.
- We demonstrate the effectiveness of incorporating local state and/or reward information in a MARL approach for controlling self-assembly in a system of active, nonequilibrium colloidal particles and equilibrium Lennard-Jones particles ungoing thermal annealing.

2 Multi-Agent Reinforcement Learning

We consider a partially observed Markov decision process (POMDP), cf. Ref. [19], specified as a tuple consisting of the state space, the action space, the space of observations, the observable function, the reward, the state-action transition operator, the discount factor, and the belief update: $(\mathcal{S}, \mathcal{A}, \mathcal{Z}, \mathcal{O}, \mathcal{R}, T, \gamma, b_0)$. As depicted in Fig. 1, we consider an external control function that can independently tune the strength of an external field on a specific region. We want to dynamically assemble clusters with a user-specified cluster size distribution ρ_* and we take

$$\mathcal{C}(x) = D_{\text{KL}}(\hat{\rho}(h(x))\|\rho_*). \tag{1}$$

as our cost (or inverse reward) function. We discuss in detail how we compute this reward in [Computational approach](#).

The function $h : \mathbb{R}^d \rightarrow \mathbb{N}^n$ takes a local particle configuration and simply counts the number of clusters with k members for each $k < n$; this constitutes the observable \mathcal{O} . The empirical distribution of cluster sizes is denoted $\hat{\rho}$. Due to this mapping, the state information is not perfectly resolved by the agent, but rather the input information is *partially observed*. We decompose the system into a fixed number of independent regions, each of which is controlled by an individual agent. Because each agent shares the same goal, the optimization is “fully cooperative” [14]. As a result, all agents share a common cost function (1) and Q function, allowing for high data efficiency as the experience of each agent is shared with all other agents.

This algorithm follows the MARL paradigm of centralized training but *fully* decentralized execution. However, fully decentralized execution poses a limitation as spatially proximal information is crucial for the self-assembly of clusters. Because of the high spatial resolution of control, the presence or absence of particles and clusters in neighboring grids will no doubt influence the cluster formation for a grid. Incorporating more local information into the observation of the current state leads to improved belief about the current state, which we summarize in the following proposition:

Proposition 1 (Low entropy belief improves expected reward). *Consider a POMDP given by $(\mathcal{S}, \mathcal{A}, \mathcal{Z}, \mathcal{O}, \mathcal{R}, T, \gamma, b_0)$ and the POMDP augmented with local state information in which the observation function is extended to $\tilde{\mathcal{O}} : \otimes_{i=1}^K \mathcal{S}_i \rightarrow \mathcal{Z}$ leading to a distinct evolution of the belief probabilities, denoted \tilde{b} . Let \bar{s} denote a state of the system with a unique optimal action \bar{a} . If $\tilde{\mathcal{O}}$ is such that $b(\bar{s}) = \delta_{s,\bar{s}}$ for all states \bar{s} , i.e. it perfectly resolves the state, then $V^*(\tilde{b}) \geq V^*(b)$ for all $\mathcal{O} \neq \tilde{\mathcal{O}}$.*

To prove this proposition, we simply write the Bellman optimality equation assuming the current state is \bar{s} , for which we know \bar{a} attains the maximal value, and letting τ denote the transition operator on the belief vector induced by T ,

$$\begin{aligned} V^*(\tilde{b}) &= \max_a \left[\sum_s R(s, a) \tilde{b}(s) + \gamma \sum_{s'} p(\tilde{z}|\tilde{b}, a) V^*(\tau) \right] \\ &= \max_a \left[R(\bar{s}, a) + \gamma \sum_{s'} T(s'|\bar{s}, a) V^*(s') \right], \end{aligned} \tag{2}$$

which is equivalent to the fully-observed MDP value function and $\tilde{z} = \tilde{O}(s')$. If the observation increases the uncertainty of the state relative to the perfectly resolved belief vector \tilde{b} , then there exists at least one state $s \neq \bar{s}$ that has belief $p_s > 0$. In other words,

$$V^*(b) = \max_a \left[p_s R(s, a) + (1 - p_s) R(\bar{s}, a) + \gamma \sum_{s'} p(\tilde{z}|\tilde{b}, a) V^*(\tau) \right] \leq V^*(\tilde{b}), \quad (3)$$

because the expected reward for any state $s \neq \bar{s}$ is suboptimal by assumption.

Computational approach : We utilize a fully cooperative multi-agent Clipped Deep Q Learning approach [20, 21] to learn optimized protocols that minimize \mathcal{C} . For a given region consisting of particle configurations x , we define the state $s = h(x)$ and the cost as $\mathcal{C}(x)$. Additionally, we consider a discrete action space \mathcal{A} , representing the strength of the externally modulating field on that region. Proposition 1 suggests that incorporating additional local information should benefit the optimization, so long as that information improves the accuracy of the belief about the current state. To do so, we consider a hierarchy of local region sizes: a plaquette consisting of the current region and its nearest neighbors, a 3×3 grid, a 5×5 grid and the global system. Given particle configurations in the local regions, x' , we define the state of the local region $s' = h(x')$ and the cost of the local region as $\mathcal{C}(x')$.

With this in mind, we extend our approach to include centralized training but *partially* decentralized execution. This corresponds to sharing information among the agents, as we include information about the state and/or cost of local regions. When including state information about the local region, we consider the joint state between a region s and its surrounding local region s' . Additionally, when including cost information about the surrounding region, we consider the average cost between $\mathcal{C}(x)$ and $\mathcal{C}(x')$. We indicate whether or not we included surrounding state and/or cost information with a subscript S for cooperative state and a subscript C for cooperative cost.

3 Active Colloids

Fig. 1 summarizes numerical results for a 2D colloidal particle system driven by externally controlled light [22]. In this system, when the Péclet number is sufficiently large (i.e. when the light-induced activity is high relative to the temperature), the system exhibits a nonequilibrium phase separation, called motility induced phase separation (MIPS) [23].

To explore the limits of an activity-inducing external control, we sought to maintain a steady state distribution consisting of clusters of particles much smaller than the macroscopic aggregate that forms when there is constant activity. We specified a target distribution ρ_* of cluster sizes using a Gaussian target distribution with $\mu = 20$, $\sigma^2 = 9$ and considered a 48×48 grid of control. Consistent with Prop. 1 including information of surrounding grids improves the performance of the optimal protocol. The best performing protocol, $3 \times 3_{CS}$, includes cluster size information and cost information in a local 3×3 region. Surprisingly, including just the cost information of this local region, $3 \times 3_C$, actually hindered the performance compared to the fully decentralized approach. On the other hand, for the Global_C case, when we included information about only the cost of the entire global system, we see that the performance does remarkably well. This underscores the importance of carefully selecting relevant centralized or semi-centralized information to improve performance. Acting in a partially decentralized manner can not only incur an additional computation cost but can also result in a decrease in performance.

Including surrounding information in the best performing approach was important in two cases: when the current region was empty or when it contained a cluster that was at or near the mean of the target distribution. In both these cases, in a fully decentralized setting, the trained agent always selects an action with the highest activity. However, in the partially decentralized setting, the agent enforces a lower activity when the surrounding region contains clusters that are above the mean of the target distribution, preventing the formation of a cluster far larger than the mean target size.

4 Thermal Annealing

Thermal annealing is widely used to improve the yield of nanoscopic materials [24, 25]. Annealing is limited as a mechanism for control because there is essentially only one parameter that can be tuned,

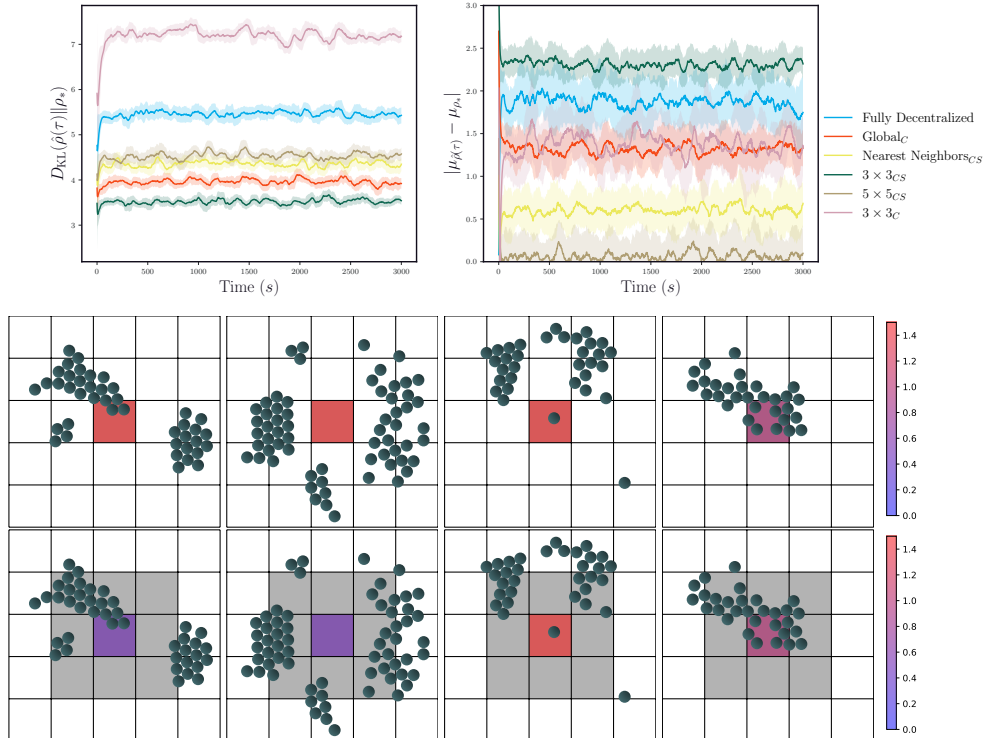


Figure 1: Active Colloids. In the plots above, both the KL divergence from the instantaneous cluster size distribution to the target distribution and the absolute deviation of the mean instantaneous cluster size distribution $\mu_{\hat{\rho}(\tau)}$ from the mean target size μ_{ρ_*} are shown as a function of time for the *optimized* protocols of various partially decentralized approaches and a fully decentralized approach (the protocols are fully trained in these curves). In the illustrations below, we compare the predicted action in the fully decentralized case (no gray shading) and the best performing partially decentralized case, $3 \times 3_{RS}$ (with gray shading). Higher light-induced activities (red) promote cluster formation.

the rate at which the temperature is decreased. We examine an alternative paradigm that exercises more localized control with measurement-guided feedback to design an annealing schedule. Rather than globally tuning a temperature, we locally update the temperature on a grid, see Fig. 2. As above, we fixed a target cluster size distribution, chosen to be a Gamma distribution with variance $\sigma^2 = 1$ and a mean $\mu = 4$ and considered a 15×15 grid of control. Here, the size and density of the system is far lower and we sought to form clusters of a much smaller size compared to the active system. This allowed us to define individual grids that were only slightly larger than an individual particle, providing much finer levels of control. Note that the interaction potential for these particles is distinct from the one that we use to model the active system and there is no nonequilibrium dynamics. The fully decentralized and top performing partially decentralized approaches differed in their behavior when the cluster size was around the mean. In the fully decentralized case, the agent was far more conservative when a cluster size was of size 4 or greater, and would impose a high temperature breaking apart the cluster. On the other hand, in one of the top performing partially decentralized approaches, Nearest Neighbors_{CS}, the agent allowed cluster sizes between 3 and 5 to persist by ensuring a lower temperature, given that there were no clusters in the vicinity.

Broader Impact

In this work, we investigated how a *partially* decentralized reinforcement learning could be used to learn protocols that externally drive a system towards some non-equilibrium steady-state. This approach can be used as a framework for learning protocols for self-assembly in more complex systems. For example, in the design of suprastructures, each agent—or a local grouping of agents—can be used to design individual substructures, which then self-assemble to form larger structures.

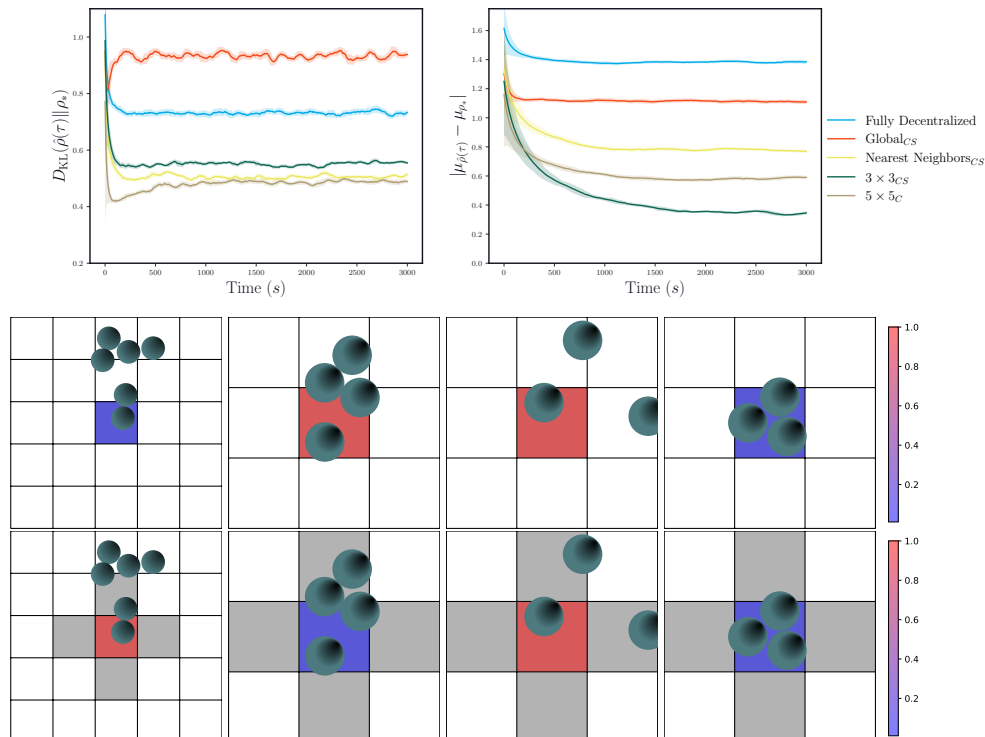


Figure 2: Thermal Annealing. See Fig. 1 caption for descriptions of plots. In the cartoons below, we compare the predicted action in the fully decentralized case (no gray shading) and Nearest Neighbors_{CS} (with gray shading). Lower temperatures (more blue) promote cluster formation.

At present, the key limitation is that the paradigm is most naturally suited to control with spatial resolution. Ultimately, the ability to design protocols for self-assembly of nanoscopic structures, can have wide-ranging impacts across a number of disciplines ranging from medicine to materials science. Potential negative impacts could, in principle, arise from the design of harmful or toxic materials.

Data and Code Availability: The data that support the findings of this study are available from the corresponding author upon reasonable request. Our code is available at <https://github.com/rotskoff-group/marl-design>.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]

- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

References

- ¹Y. Yin and A. P. Alivisatos, "Colloidal nanocrystal synthesis and the organic–inorganic interface", *Nature* **437**, 664–670 (2005).
- ²Y. Ma and A. L. Ferguson, "Inverse design of self-assembling colloidal crystals with omnidirectional photonic bandgaps", *Soft Matter* **15**, 8808–8826 (2019).
- ³K. R. Gadelrab, A. F. Hannon, C. A. Ross, and A. Alexander-Katz, "Inverting the design path for self-assembled block copolymers", *Molecular Systems Design & Engineering* **2**, 539–548 (2017).
- ⁴H. Ronellenfitsch, N. Stoop, J. Yu, A. Forrow, and J. Dunkel, "Inverse design of discrete mechanical metamaterials", *Physical Review Materials* **3**, 095201 (2019).
- ⁵J. C. Cameron, S. C. Wilson, S. L. Bernstein, and C. A. Kerfeld, "Biogenesis of a Bacterial Organelle: The Carboxysome Assembly Pathway", *Cell* **155**, 1131–1140 (2013).
- ⁶C. Sigl et al., "Programmable icosahedral shell system for virus trapping", *Nature Materials*, 1–9 (2021).
- ⁷G. M. Rotskoff and P. L. Geissler, "Robust nonequilibrium pathways to microcompartment assembly", *Proceedings of the National Academy of Sciences* **112**, 201802499 (2018).
- ⁸S. Whitlam and R. L. Jack, "The Statistical Mechanics of Dynamic Pathways to Self-Assembly", *Annual Review of Physical Chemistry* **66**, 143–163 (2015).
- ⁹M. C. Rechtsman, F. H. Stillinger, and S. Torquato, "Optimized Interactions for Targeted Self-Assembly: Application to a Honeycomb Lattice", *Physical Review Letters* **95**, 228301 (2005).
- ¹⁰D. Chen, G. Zhang, and S. Torquato, "Inverse design of colloidal crystals via optimized patchy interactions", *The Journal of Physical Chemistry B* **122**, 8462–8468 (2018).
- ¹¹F. Romano, J. Russo, L. Kroc, and P. Šulc, "Designing Patchy Interactions to Self-Assemble Arbitrary Structures", *Physical Review Letters* **125**, 118003 (2020).
- ¹²Z. Ma, E. Lomba, and S. Torquato, "Optimized large hyperuniform binary colloidal suspensions in two dimensions", *Physical Review Letters* **125**, 068002 (2020).
- ¹³M. J. Falk, V. Alizadehyazdi, H. Jaeger, and A. Murugan, "Learning to Control Active Matter", [arXiv:2105.04641 \[cond-mat\]](https://arxiv.org/abs/2105.04641) (2021).
- ¹⁴M. Tan, "Multi-agent reinforcement learning: Independent versus cooperative agents", in *Machine learning, proceedings of the tenth international conference, university of massachusetts, amherst, MA, USA, june 27-29, 1993*, edited by P. E. Utgoff (1993), pp. 330–337.

- ¹⁵L. Busoniu, R. Babuska, and B. De Schutter, “A Comprehensive Survey of Multiagent Reinforcement Learning”, *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* **38**, 156–172 (2008).
- ¹⁶T. Rashid, M. Samvelyan, C. S. de Witt, G. Farquhar, J. N. Foerster, and S. Whiteson, “QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning”, in *Proceedings of the 35th international conference on machine learning, ICML 2018, stockholm, sweden, july 10-15, 2018*, Vol. 80, edited by J. G. Dy and A. Krause, Proceedings of machine learning research (2018), pp. 4292–4301.
- ¹⁷K. Son, D. Kim, W. J. Kang, D. Hostallero, and Y. Yi, “QTRAN: Learning to factorize with transformation for cooperative multi-agent reinforcement learning”, in *Proceedings of the 36th international conference on machine learning, ICML 2019, 9-15 june 2019, long beach, california, USA*, Vol. 97, edited by K. Chaudhuri and R. Salakhutdinov, Proceedings of machine learning research (2019), pp. 5887–5896.
- ¹⁸P. Sunehag et al., “Value-decomposition networks for cooperative multi-agent learning based on team reward”, in *Proceedings of the 17th international conference on autonomous agents and MultiAgent systems, AAMAS 2018, stockholm, sweden, july 10-15, 2018*, edited by E. André, S. Koenig, M. Dastani, and G. Sukthankar (2018), pp. 2085–2087.
- ¹⁹N. Roy, G. Gordon, and S. Thrun, “Finding Approximate POMDP solutions Through Belief Compression”, *Journal of Artificial Intelligence Research* **23**, 1–40 (2005).
- ²⁰V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, “Playing atari with deep reinforcement learning”, in *NIPS deep learning workshop* (2013).
- ²¹S. Fujimoto, H. van Hoof, and D. Meger, “Addressing function approximation error in actor-critic methods”, in *Proceedings of the 35th international conference on machine learning*, Vol. 80, edited by J. Dy and A. Krause, Proceedings of machine learning research (July 2018), pp. 1587–1596.
- ²²J. Palacci, S. Sacanna, A. P. Steinberg, D. J. Pine, and P. M. Chaikin, “Living Crystals of Light-Activated Colloidal Surfers”, *Science* **339**, 936–940 (2013).
- ²³M. E. Cates and J. Tailleur, “Motility-Induced Phase Separation”, *Annual Review of Condensed Matter Physics* **6**, 219–244 (2015).
- ²⁴G. Makrides, B. Zinsser, A. Phinikarides, M. Schubert, and G. E. Georghiou, “Temperature and thermal annealing effects on different photovoltaic technologies”, *Renewable Energy* **43**, 407–417 (2012).
- ²⁵S. Dey et al., “DNA origami”, *Nature Reviews Methods Primers* **1**, 1–24 (2021).