Editors' Suggestion | Featured in Physics

Tuning Colloidal Reactions

Ryan K. Krueger,[1,*] Ella M. King,[2,3,*] and Michael P. Brenner[1,4,†]
[1]School of Engineering and Applied Sciences, Harvard University, 29 Oxford Street, Cambridge, Massachusetts 02138, USA
[2]Department of Physics, New York University, 726 Broadway, New York, New York 10003, USA
[3]Simons Foundation, 160 5th Avenue, New York, New York 10010, USA
[4]Department of Physics, Harvard University, 17 Oxford Street, Cambridge, Massachusetts 02138, USA

(Received 15 December 2023; revised 21 August 2024; accepted 7 October 2024; published 27 November 2024)

The precise control of complex reactions is critical for biological processes, yet our inability to design for specific outcomes limits the development of synthetic analogs. Here, we leverage differentiable simulators to design nontrivial reaction pathways in colloidal assemblies. By optimizing over external structures, we achieve controlled disassembly and particle release from colloidal shells. Lastly, we characterize the role of configurational entropy in the structure via both forward calculations and optimization, inspiring new parameterizations of designed colloidal reactions.

DOI: 10.1103/PhysRevLett.133.228201

Both living and nonliving physical systems exhibit complex dynamical behavior, ranging from repair to locomotion to catalysis. Fundamentally, such behaviors arise from sequences of reactions, in which a set of substances (i.e., the reactants) are transformed into a set of different substances (i.e., the products). A rich body of work aims to characterize and tune systems of interacting agents spanning a range of system descriptions, both theoretically [1–3] and experimentally [4–9]. However, for many critical processes in biological systems (e.g., DNA synthesis, protein folding), the components themselves cannot be changed. Instead, to modify these processes, researchers often introduce an external structure (e.g., competitive inhibitors for enzyme inhibition, protein folding chaperones). Thus far, the design of such structures has been bespoke and application specific, necessitating entirely new research programs for each new reaction. For example, while some general theoretical models have provided deep insights into catalysis [10,11], they are largely too abstract to inform experimental design.

To overcome current limitations and tune reactions through the design of external agents, we carry out *inverse design* whereby we optimize the geometry and interactions of such components to achieve a target reaction. While inverse design has been successfully applied to self-assembly [12–16], inverse designing reaction pathways remains a challenge because design parameters must be chosen to favor particular dynamical trajectories. The advent of differentiable simulators [16,17], powered by software libraries developed for machine learning [18], has opened up the possibility of directly designing reactions as the gradient of numerical procedures with respect to control parameters can be computed efficiently.

Here, we design complex reactions using differentiable molecular dynamics (MD) and gradient-based optimization. As an example of a nontrivial reaction, we consider the controlled disassembly of colloidal structures, whereby a particle is extracted from an otherwise complete shell of colloidal particles. Disassembly is central to the dynamic functions of living systems, such as defect repair, self-replication, and catalysis. Existing examples of controlled disassembly in synthetic systems often rely on external forcing to drive the disassembly process [19–22], which provides a direct pathway to tuning behavior. However, for many engineering applications, including those inherent to living systems, the use of external fields is limiting. On the other hand, controlled disassembly in living systems typically relies on local energy consumption [e.g., biological enzymes consuming adenosine triphosphate (ATP)] rather than global fields, but the synthetic design of these systems is significantly more complex.

Inspired by the symmetry of many viral capsids [23,24], we design for the controlled disassembly of icosahedral shells. We consider a fixed shell and only parametrize an external structure that acts upon it, enabling control over disassembly without modifying critical components of the reaction. Importantly, our disassembly mechanism is entirely passive and does not rely on external forcing. As a model for potential engineering applications, we apply our mechanism to provoke the release of a target small particle initially trapped inside the shell. Controlled disassembly serves here as a striking example of a complex reaction because the reaction requires a finely-tuned interaction energy to keep the remaining shell stabilized while still performing the desired particle extraction. We start from a

*These authors contributed equally to this work.
†Contact author: brenner@seas.harvard.edu

rigid structure and thereafter proceed to quantify the role of flexibility in the structure by computing free energy landscapes both for predefined extrema and for structures optimized via a chosen parameterization of configurational entropy, opening the door to novel designed reactions.

Results—We implement controlled disassembly in a colloidal patchy particle system. Patchy particles have long been used to emulate interactions in soft materials [25,26] and offer tremendous tunability in designed interactions. Optimizing said systems to achieve specific behaviors has been made possible by the recent development of patchy particle simulations within a differentiable library [16]. In particular, we aim to remove a single particle from a shell composed of patchy particles in a controlled manner without disrupting the remaining shell structure. To that end, we tune disassembly without changing any properties of the shell itself. Instead, we introduce an external structure that interacts with the shell to disassemble it in the desired manner. We term the external structure a "spider" due to its geometry.

The shell is modeled as a collection of patchy particles forming an icosahedron where each patch corresponds to a contact with a neighboring particle. Each patchy particle consists of a central sphere and a set of rigidly attached patches. Patches interact via a Morse potential ($\epsilon_V = 10.0$, $\alpha_V = 5.0$) and central spheres interact via soft-sphere repulsion ($\epsilon_{ss} = 10^4$). Importantly, the geometry and interaction energy of shell-comprising particles are fixed throughout the optimization. Although we focus on the disassembly of the icosahedron, our framework can be easily adapted for other shell geometries (see Supplemental Material (SM) for octahedral shells [27]).

For the spider, we consider several different models with a set of core similarities. All models contain a ring of "base" particles and a "head" particle that sits above the ring along its symmetry axis. The head is connected to the base particles by repulsive bars, making the entire structure a cage-like object that is open on one end. An attractive particle type (either the head or a third particle species) interacts with the shell-comprising particles via a Morse potential, whereas base particles and connecting bars interact with shell particles via soft sphere repulsion. Unlike the shell, the geometry and interaction energy of the spider are parameters of the optimization. See Fig. 1(c) for an overview of this parameterization. All interaction energies in our system are parameterized with simple, physics-based potentials.

Given a specified parameterization for the spider geometry and interactions, we run an ensemble of differentiable molecular dynamics simulations (see Fig. 1(d) and SM [27]). To focus our optimization procedure on the challenges specific to disassembly, we initialize the spider bound to the shell and therefore ignore the period in which the spider is freely diffusing. We optimize over the parameters that characterize the geometry of the spider and its interaction with the shell (8 parameters for the optimizations in Fig. 2). To optimize our system, we perform gradient descent to minimize a loss function. The loss is constructed from two competing terms: one that rewards a final state in which the target particle is extracted, and one that penalizes a strong interaction between the spider and non-target particles. The second term, which we refer to as the "remaining energy" term, tends to reward pathways in which the remaining shell holds its shape after the spider extracts the target particle.

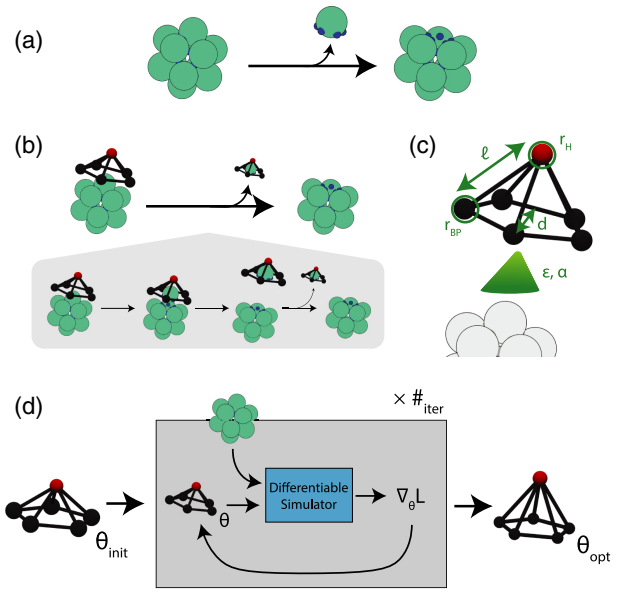

FIG. 1. Tuning the interaction potential of an external structure, the spider, to achieve a desired reaction of disassembly. (a) A single particle is removed from an icosahedron. (b) A candidate mechanism: the spider extracts the target particle via an attractive potential and detaches from the remaining shell. (c) Parametrization of the spider geometry and interaction potential with the shell. The red particle is the "head" particle, situated above the four black "base" particles that constitute the ring. The interaction energy between spider and shell is depicted as a green triangle. We optimize over all labeled parameters, as well as the cutoff of the interaction energy (not depicted). (d) High-level depiction of our optimization pipeline: analytic gradients are computed via a differentiable molecular dynamics simulator and parameters of the spider are updated via gradient descent.

We formalize the loss function as follows. Consider an icosahedral shell comprised of a collection of particles $V = \{\vec{v}_1, \vec{v}_2, \ldots, \vec{v}_n\}$ where $n = 12$. We seek to extract a target particle $\vec{v}_j$ from the shell while leaving the remaining shell $V \backslash \vec{v}_j$ intact. We can measure the degree to which $\vec{v}_j$ is successfully extracted via the following expression:

$$\mathcal{L}_{\text{extract}}(V) = -\sum_{i \neq j} d(\vec{v}_i, \vec{v}_j), \tag{1}$$

where $d(\vec{v}_i, \vec{v}_j)$ denotes the Euclidean distance between particles $\vec{v}_i$ and $\vec{v}_j$. Note the negative sign as we formulate our optimization problem to minimize the loss. Next, we

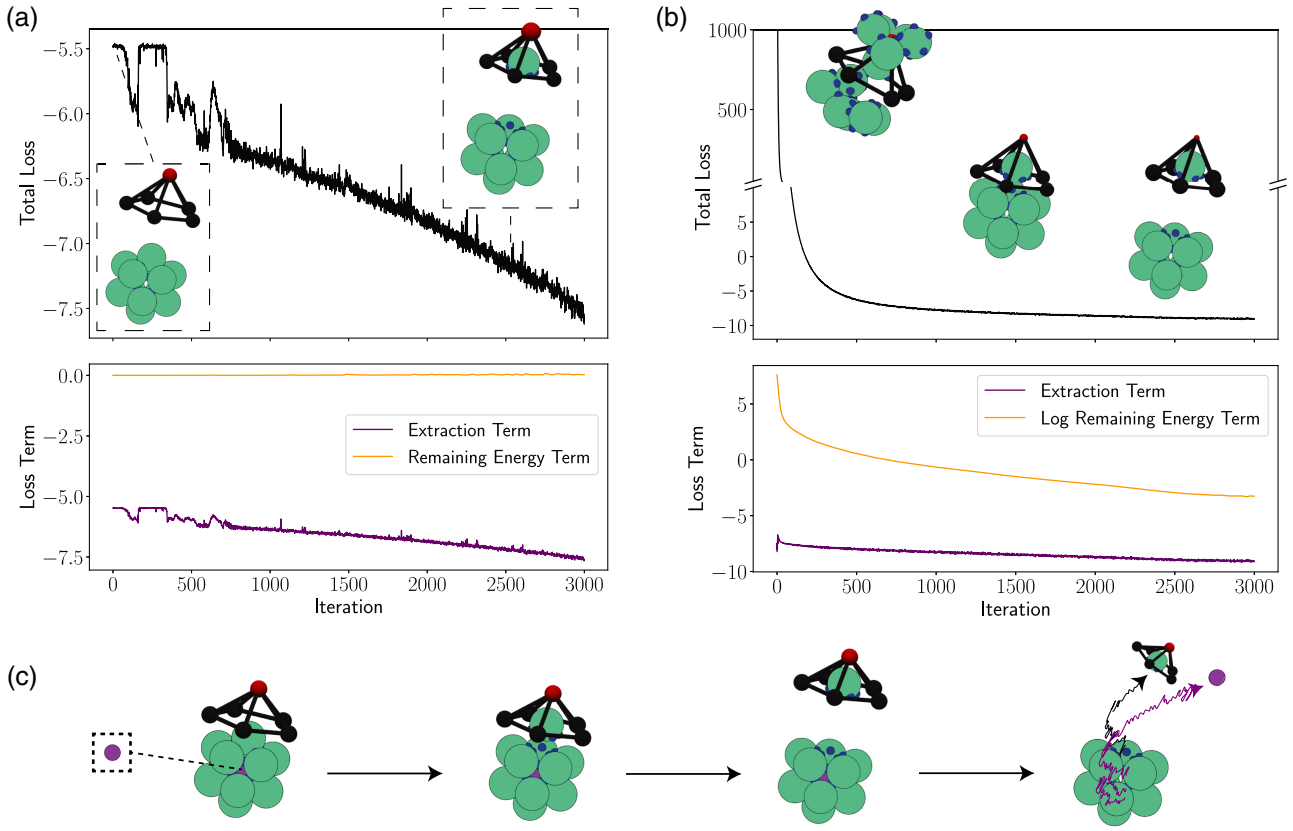

FIG. 2. Optimizing the geometry and interaction potential of a rigid spider. (a) In the limit of a weak initial spider-shell interaction, the initial spider simply diffuses away from the shell. By the 3000th iteration, the spider geometry and interaction energy are optimized to extract the target particle while still diffusing away from the remaining particles, leaving them intact. (b) In the limit of a strong initial spider-shell interaction, the initial spider extracts the target particle but does not diffuse away, disturbing the integrity of the remaining shell. As optimization progresses, the interaction is tuned to only extract the particle without disrupting the remaining shell. Upon convergence, the spider geometry and interaction energy are tuned to maintain extraction while diffusing away. Insets depict representative states after 10 000 MD steps at the corresponding iteration. (c) Schematic of a simulation of an optimized spider provoking the release of a particle from an icosahedral shell.

minimize the interaction energy between attractive site(s) and nontarget shell particles:

$$\mathcal{L}_{\text{remain}}(V,A) = \sum_{\vec{a}\in A}\left(\sum\nolimits_{i\neq j} U_{\text{m}}(\vec{a},\vec{v}_i)\right)^2, \qquad (2)$$

where A denotes the set of attractive sites and $U_{\text{m}}(\vec{a},\vec{v}_i)$ represents the interaction energy between the attractive site $\vec{a}$ and a shell particle $\vec{v}_i \in V$. For the spider depicted in Fig. 1, the head particle is the only attractive site. We calculate the "remaining energy" term, $\mathcal{L}_{\text{remain}}$, of the total loss with respect to the initial configuration, i.e. the first timestep of the simulation. All other terms depend on the dynamics of the system, so we evaluate them on the final state. In all simulations, the spider is initially bound to the target particle and we integrate the system for 1000 time steps (see SM [27]).

Rigid spider—We begin with a minimal conception of the spider: a rigid body consisting of only a head particle and five base particles which reflect the five-fold symmetry of the icosahedron. The base particles are attached to each other by rigid bars, forming a cage-like structure that is open on one end. The attractive interactions between the spider and the icosahedron are restricted to interactions between the head particle in the spider and the patches on the icosahedral vertices.

We explore two limits of our optimization procedure (Fig. 2). First, we perform an optimization where the spider is initialized to interact weakly with the shell particles ($\log(\epsilon_H) = 3.0$, $\alpha_H = 1.5$). In this limit, the spider simply diffuses away from the shell at long timescales without extracting the target particle. Initially, we observe variable changes consistent with increasing the interaction between spider head and shell to achieve particle extraction: ϵ_H increases, the head height decreases, and the head particle radius increases. In the following iterations, we observe parameter changes focused on maintaining extraction while reducing the interaction strength between the spider and the rest of the shell. The head height increases, consistent with minimizing the remaining energy, but to maintain particle

extraction, α_H decreases (increasing the range of the Morse potential). This suggests that tightly coupled, nontrivial parameter changes drive extraction while maintaining minimal interaction with the remaining shell.

Next, we perform an optimization in the opposite limit in which the spider is initialized to interact strongly with the shell ($\log(\epsilon_H) = 10.5$, $\alpha_H = 1.5$). Initially, this interaction is so strong that the spider not only extracts the target particle but it also disrupts the remaining shell. This can be seen in the large value of the remaining energy loss term, which penalizes the energy between the spider head and non-target particles. Throughout the optimization, we observe variable changes consistent with tuning the interaction strength to maintain extraction while minimizing off-target interactions: ϵ_H decreases, α_H increases, the head radius decreases, the head height increases, and the base particle radius increases. When evaluated on longer simulations, the converged parameter set also achieves spontaneous diffusion of the spider-particle complex away from the remaining undisturbed shell. Note that neither changes in the random seed nor perturbations to the initial parameters consistently yield similar optimized parameters (see SM [27]).

Contrasting the high and low energy optimization regimes reveals the inherent delicacy in tuning the spider to achieve extraction and subsequent diffusion from the shell. The spider-shell interaction must be sufficiently strong to extract the target particle, but simultaneously weak enough to not disturb non-target particles and to diffuse away from the shell within the timescale of our simulations. This tension is reflected in the behavior of the loss terms in each optimization. Overall, in the weak-interaction limit, the term penalizing interactions with non-target particles remains negligible while the extraction term drives optimization; in the strong-interaction limit it is the same energy-penalizing term that dominates the loss. Our optimized reactions represent a notion of balance that is necessary for biologically relevant functions, e.g. the controlled release of a particle from a closed shell [see simulation in Fig. 2(c)]. This serves as a toy example of a potential target for engineering applications, such as drug delivery via a viral shell.

Flexible spider—The configurational entropy of the spider can serve as a control knob for tuning reactions. While the optimized results in Fig. 2 highlight that spider geometry dramatically impacts its performance, the rigid formulation cannot access configurational entropy. Here, we define a modified form of our spider permitting varying degrees of flexibility. Rather than the head serving as the sole attractive site, we introduce a ring of attractive sites consisting of one site per spider leg positioned between the base and head particles [Fig. 3(a)]. In this way, bonds at the spider base connecting individual legs can be made flexible or removed entirely. The resulting fluctuations due to leg flexibility directly change the interaction strength between spider and extracted particle and thus modulate the entropic contribution to the interaction. In this scheme, the head particle only interacts repulsively with the icosahedron.

Instead of considering the probability of extraction, we focus on the release of an already extracted particle since this process is likely to be strongly influenced by configurational entropy. We reason that increased entropy in the extracted state (i.e., the extracted particle bound to the spider) would favor particle release compared to the fully rigid spider because fluctuations in the spider configuration would reduce the effective attraction. To test this hypothesis, we quantitatively measure free energy differences corresponding to particle release between three versions of the modified spider with varying flexibility: (i) a fully rigid spider with fixed bonds between all adjacent base particles, (ii) a partially flexible spider resulting from the removal of two base bonds, and (iii) a fully flexible spider via the removal of all base bonds.

We compute the free energy of release using each of the three models. The distance from the extracted particle to the head directly relates to its release. We therefore use this metric as the order parameter for free energy calculations [Fig. 3(b)]. We compute the free energy diagrams for each spider using the weighted histogram analysis method (WHAM; see SM [27] and Refs. [28,29]) and use a fixed set of parameters for the spider geometry and interaction determined via a single optimization with the fully rigid spider (see SM [27]). As expected, the more flexible the geometry, the more favorable the released state: there is a smaller change in free energy between the attached and released states for more flexible geometries [Fig. 3(c)].

Next, we optimize configurational entropy directly. We define a spider in which all pairs of (i) adjacent and (ii) next-nearest neighbor base particles are connected with springs whose spring constants are free parameters [Fig. 3(d)]. To bias the optimization procedure toward a spider with an increased likelihood of release, we define an additional loss term representing the total interaction energy between attractive sites and extracted particle:

$$\mathcal{L}_{\text{release}}(\vec{v}, A) = \sum_{\vec{a} \in A} [U_{\text{m}}(\vec{a}, \vec{v})]^2. \tag{3}$$

The weaker the interaction, the easier the release. Here, we optimize over a longer (1500 step) simulation than in the rigid case. We average the new loss term over states sampled from the final 500 simulation steps, while maintaining extraction within the first 1000 steps (see SM [27]). To give the optimization algorithm more freedom to promote release, we rescale the loss describing extraction such that it changes minimally beyond a specified maximum value. The optimization algorithm can then reduce extraction efficiency without penalty.

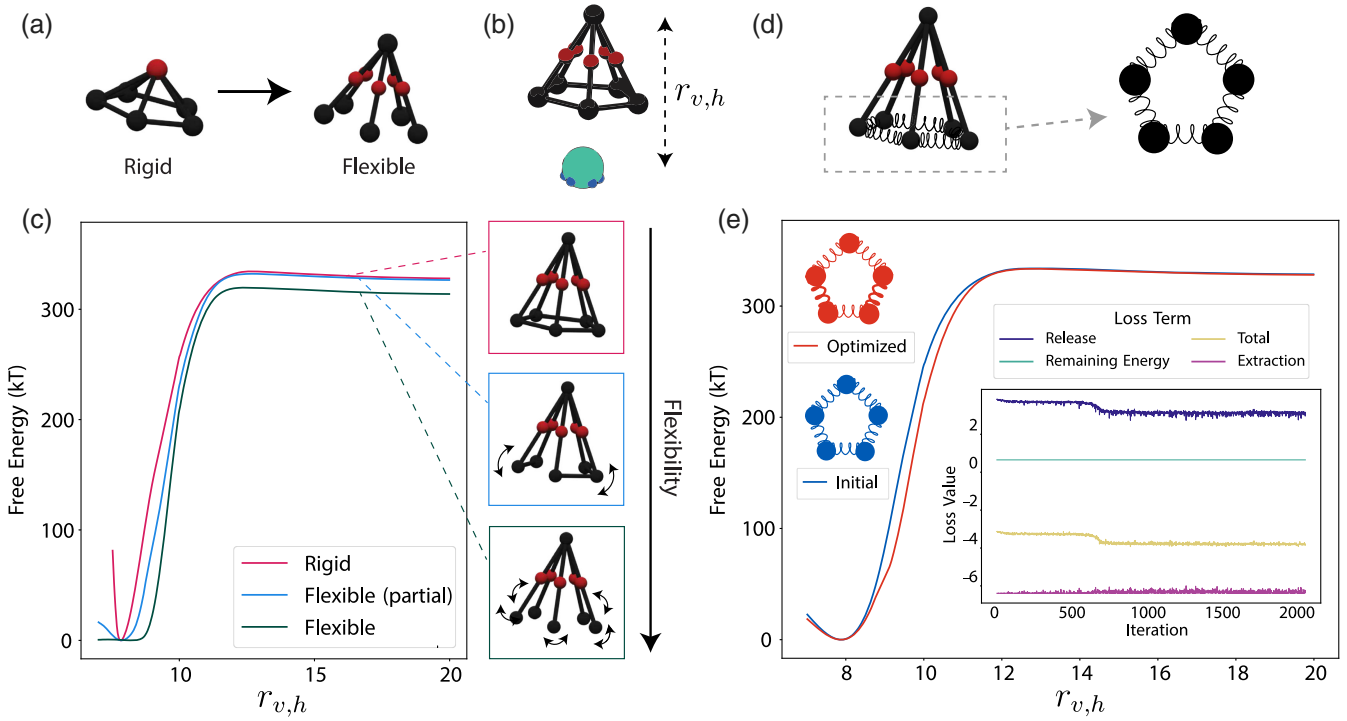

FIG. 3. Role of configurational entropy in the release of a target particle. (a) A modified version of the original spider in which a ring of attractive sites are positioned between the base and head particles on each spider leg. (b) We use the distance from the extracted particle to the spider head as the order parameter approximating particle release. (c) Free energy diagrams for the order parameter depicted in (b) for three variants of the modified spider: (i) fully rigid, (ii) partially flexible, defining two rigid substructures which are free to rotate about the head, and (iii) fully flexible, in which all base bonds are removed. The parameters for the spider geometry and interaction are determined via an optimization of the rigid configuration. Free energy diagrams are computed via the Weighted Histogram Analysis Method (WHAM). (d) A parameterization of the configurational entropy of the modified spider in which the base bonds are represented as springs whose spring constants are free parameters. Identical springs are also placed between all next-nearest neighbors to parameterize the bond angles (not shown). (e) The free energies of the initial spider with uniform spring constant and the optimized spider as computed via WHAM. Inset: an optimization over the spring constants depicted in (d) to maintain extraction and minimize remaining energy while also minimizing the average energy between extracted particle and attractive sites over an additional 500 simulation steps. The parameters defining the spider geometry and interaction are set to those used to compute free energies in (c).

We optimize these spring constants using the modified loss function while keeping all other parameters fixed. We initialize all spring constants to be the same value (i.e., $\log(k) = 2.0$) and fix the spider geometry and interaction parameters to those used in Fig. 3(c). The optimized solution has a wider well but maintains the same well depth. As a result, the free energy difference between the extracted and released states is lower in the optimized configuration than in the initial one at intermediate distances. Interestingly, the optimization procedure naturally converges to a solution with an asymmetric distribution of spring constants. Directly tuning this asymmetry is a promising avenue for future work.

Discussion—In this Letter, we achieve nontrivial reactions via designed external structures. We consider the case of controlled disassembly of an icosahedral shell composed of patchy particles, in which there is an inherent tension between initiating disassembly and maintaining the integrity of the remaining substructure. We show how the parameters governing a rigid external structure can be finely tuned to minimize a loss function representing this tension. We find that the optimized spider provokes particle release. We then add configurational entropy by introducing a flexible spider geometry, and quantify the influence of flexibility by comparing free energy landscapes for varying degrees of flexibility. Our framework naturally accommodates parameterizations of configurational entropy. Upon optimization, a spider with asymmetrically flexible base legs favors release over the initialized uniform configuration.

Since we optimize directly with respect to the numerically integrated dynamics, our method is general enough to study a wide range of systems. Foremost, it may enable experimental realizations of theoretical models that were otherwise limited by an inability to finely tune interaction energies. For example, Ref. [30] introduces a model of self-replicating colloidal clusters in which kinetic traps can be avoided by tuning the interaction energies, but dissociation of a new cluster from its parent (a necessary step for replication) required an artificial trigger event in numerical

simulations. In contrast, our designed parameters lead to spontaneous dissociation of the spider-particle complex away from the remaining shell. The computational flexibility of the method could also easily enable users to restrict the parameter regime to experimentally realizable interactions. This could be done for DNA coated colloids, e.g., by optimizing the DNA sequences that define the interaction strength [31–34].

Numerical instabilities can arise when optimizing over dynamical simulations. The primary limitation we observe is that gradients become unstable and very large for long simulations. There are several possible approaches to reducing instability in gradients in such cases. One standard method to mitigate instabilities in the context of differentiable programming is gradient clipping [35,36]. One could also decrease the total number of time steps by training an emulator to resolve the dynamics with a larger time step than is possible with standard integrators, following similar work for deterministic systems [37–39]. An alternative approach would be to integrate differentiable simulations with enhanced sampling methods to sample low probability events without the need for long simulation times.

We rely on gradient-based optimization due to its scalability and performance. Our method naturally scales to larger and more complex systems since (i) gradient calculation via automatic differentiation only requires a single simulation, (ii) reverse-mode automatic differentiation scales efficiently with the number of parameters [40], and (iii) the gradient explicitly captures interdependencies which is essential to efficiently tuning complex behavior. We anticipate that our approach and proposed design rules will be applicable to physical reactions beyond the colloidal regime.

Acknowledgments—We thank Sam Schoenholz for his work developing JAX-MD, the members of the Brenner Group for helpful discussions, and Stefania Ketzetzi for helpful feedback on the text. This work is supported by the Office of Naval Research (ONR N00014-17-1-3029, ONR N00014-23-1-2654), the NSF (Grant No. DMR-1921619), the NSF AI Institute of Dynamic Systems (Grant No. 2112085), and the Simons Foundation (Grant No. 1141499). E. M. K. acknowledges support from a Simons Foundation Junior Fellowship under Grant No. 1141499.

[1] G. van Anders, D. Klotsa, A. S. Karas, P. M. Dodd, and S. C. Glotzer, Digital alchemy for materials design: Colloids and beyond, ACS Nano **9**, 9542 (2015).

[2] Z. Zeravcic, V. N. Manoharan, and M. P. Brenner, Colloquium: Toward living matter with colloidal particles, Rev. Mod. Phys. **89**, 031001 (2017).

[3] Z. M. Sherman, M. P. Howard, B. A. Lindquist, R. B. Jadrich, and T. M. Truskett, Inverse methods for design of soft materials, J. Chem. Phys. **152**, 140902 (2020).

[4] V. X. Truong, L. L. Rodrigues, and C. Barner-Kowollik, Light-and mechanic field controlled dynamic soft matter materials, Polym. Prepr. **13**, 4915 (2022).

[5] Q. Li, A. P. Schenning, and T. J. Bunning, Light-responsive smart soft matter technologies, Adv. Opt. Mater. **7**, 1901160 (2019).

[6] Y. Wang, Y. Wang, D. R. Breed, V. N. Manoharan, L. Feng, A. D. Hollingsworth, M. Weck, and D. J. Pine, Colloids with valence and specific directional bonding, Nature (London) **491**, 51 (2012).

[7] A. McMullen, M. Muñoz Basagoiti, Z. Zeravcic, and J. Brujic, Self-assembly of emulsion droplets through programmable folding, Nature (London) **610**, 502 (2022).

[8] R. Niu, C. X. Du, E. Esposito, J. Ng, M. P. Brenner, P. L. McEuen, and I. Cohen, Magnetic handshake materials as a scale-invariant platform for programmed self-assembly, Proc. Natl. Acad. Sci. U.S.A. **116**, 24402 (2019).

[9] N. B. Schade, M. C. Holmes-Cerfon, E. R. Chen, D. Aronzon, J. W. Collins, J. A. Fan, F. Capasso, and V. N. Manoharan, Tetrahedral colloidal clusters from random parking of bidisperse spheres, Phys. Rev. Lett. **110**, 148303 (2013).

[10] Z. Zeravcic and M. P. Brenner, Spontaneous emergence of catalytic cycles with colloidal spheres, Proc. Natl. Acad. Sci. U.S.A. **114**, 4342 (2017).

[11] M. Muñoz-Basagoiti, O. Rivoire, and Z. Zeravcic, Computational design of a minimal catalyst using colloidal particles with programmable interactions, Soft Matter **19**, 3933 (2023).

[12] S. Torquato, Inverse optimization techniques for targeted self-assembly, Soft Matter **5**, 1157 (2009).

[13] D. Chen, G. Zhang, and S. Torquato, Inverse design of colloidal crystals via optimized patchy interactions, J. Phys. Chem. B **122**, 8462 (2018).

[14] M. Rechtsman, F. Stillinger, and S. Torquato, Designed interaction potentials via inverse methods for self-assembly, Phys. Rev. E **73**, 011406 (2006).

[15] Y. Ma and A. L. Ferguson, Inverse design of self-assembling colloidal crystals with omnidirectional photonic bandgaps, Soft Matter **15**, 8808 (2019).

[16] E. M. King, C. X. Du, Q.-Z. Zhu, S. S. Schoenholz, and M. P. Brenner, Programmable patchy particles for materials design, Proc. Natl. Acad. Sci. U.S.A. **121**, 1748e2311891121 (2024).

[17] S. Schoenholz and E. D. Cubuk, JAX M.D.: A framework for differentiable physics, Adv. Neural Inf. Process. Syst. **33**, 11428 (2020).

[18] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: Composable transformations of Python + NumPy programs (2018).

[19] Y.-j. Jung, H. Kim, H.-K. Cheong, and Y.-b. Lim, Magnetic control of self-assembly and disassembly in organic materials, Nat. Commun. **14**, 3081 (2023).

[20] S. Tottori, L. Zhang, K. E. Peyer, and B. J. Nelson, Assembly, disassembly, and anomalous propulsion of microscopic helices, Nano Lett. **13**, 4263 (2013).

[21] F. Martínez-Pedrero, F. Ortega, J. Codina, C. Calero, and R. G. Rubio, Controlled disassembly of colloidal aggregates confined at fluid interfaces using magnetic dipolar interactions, J. Colloid Interface Sci. **560**, 388 (2020).

[22] M. A. Kostiainen, P. Ceci, M. Fornara, P. Hiekkataipale, O. Kasyutich, R. J. Nolte, J. J. Cornelissen, R. D. Desautels, and J. van Lierop, Hierarchical self-assembly and optical disassembly for controlled switching of magnetoferritin nanoparticle magnetism, ACS Nano **5**, 6394 (2011).

[23] D. L. Caspar and A. Klug, Physical principles in the construction of regular viruses, in *Cold Spring Harbor Symposia on Quantitative Biology* (Cold Spring Harbor Laboratory Press, New York, 1962), Vol. 27, pp. 1–24, 10.1101/sqb.1962.027.001.005.

[24] R. Twarock and A. Luque, Structural puzzles in virology solved with an overarching icosahedral design principle, Nat. Commun. **10**, 4414 (2019).

[25] Z. Zhang and S. C. Glotzer, Self-assembly of patchy particles, Nano Lett. **4**, 1407 (2004).

[26] S. C. Glotzer and M. J. Solomon, Anisotropy of building blocks and their assembly into complex structures, Nat. Mater. **6**, 557 (2007).

[27] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.133.228201 for simulation and optimization details, system definitions, optimized parameter values, and octahedral optimizations.

[28] S. Kumar, J. M. Rosenberg, D. Bouzida, R. H. Swendsen, and P. A. Kollman, The weighted histogram analysis method for free-energy calculations on biomolecules. I. The method, J. Comput. Chem. **13**, 1011 (1992).

[29] A. Grossfield, Wham: The weighted histogram analysis method, Available at http://membrane.urmc.rochester.edu/?page_id=126, version 2.0.11.

[30] Z. Zeravcic and M. P. Brenner, Self-replicating colloidal clusters, Proc. Natl. Acad. Sci. U.S.A. **111**, 1748 (2014).

[31] W. B. Rogers, A mean-field model of linker-mediated colloidal interactions, J. Chem. Phys. **153**, 124901 (2020).

[32] W. B. Rogers and J. C. Crocker, Direct measurements of DNA-mediated colloidal interactions and their quantitative modeling, Proc. Natl. Acad. Sci. U.S.A. **108**, 15687 (2011).

[33] M. Matthies, R. Krueger, A. Torda, and M. Ward, Differentiable partition function calculation for RNA, Nucleic Acids Research **52**, e14 (2023).

[34] R. Krueger, M. P. Brenner, and K. Shrinivas, Generalized design of sequence-ensemble-function relationships for intrinsically disordered proteins, bioRxiv 2024 (2024).

[35] J. G. Greener, Differentiable simulation to develop molecular dynamics force fields for disordered proteins, Chem. Sci. **15**, 4897 (2024).

[36] M. C. Engel, J. A. Smith, and M. P. Brenner, Optimal Control of Nonequilibrium Systems Through Automatic Differentiation, Phys. Rev. X **13**, 041032 (2024).

[37] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. Battaglia, Learning to simulate complex physics with graph networks, in *International Conference on Machine Learning* (PMLR, 2020), pp. 8459–8468, https://proceedings.mlr.press/v119/.

[38] K. R. Allen, T. Lopez-Guevara, K. Stachenfeld, A. Sanchez-Gonzalez, P. Battaglia, J. Hamrick, and T. Pfaff, Physical design using differentiable learned simulators, arXiv:2202.00728.

[39] Y. Bar-Sinai, S. Hoyer, J. Hickey, and M. P. Brenner, Learning data-driven discretizations for partial differential equations, Proc. Natl. Acad. Sci. U.S.A. **116**, 15344 (2019).

[40] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, Automatic differentiation in machine learning: A survey, J. March. Learn. Res. **18**, 1 (2018).