

Differentiable learning and control of free-energy-driven pattern dynamics

Alexander E. Cohen ^{1,*} Samuel Degnan-Morgenstern ^{1,†} Simon Daubner ²
 Jörn Dunkel ^{3,‡} and Martin Z. Bazant ^{3,1,§}

¹*Department of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

²*Dyson School of Design Engineering, Imperial College London, London, United Kingdom*

³*Department of Mathematics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*



(Received 20 March 2026; accepted 19 May 2026; published 25 June 2026)

Pattern-forming dynamics govern the behavior of many quantum, electrochemical, and soft-matter systems, yet learning and controlling the corresponding partial differential equation (PDE) models on realistic geometries remains challenging. In this work, we develop a unified, differentiable framework for free-energy-based PDEs that enables end-to-end parameter inference and optimal control directly on image-based domains. Starting from a variational description in terms of an energy or free energy functional, we combine PDE-based smoothing of segmented geometries, the smoothed boundary method, advanced integrators for stiff pattern-forming PDEs, and memory-efficient automatic differentiation implemented in JAX to construct scalable PDE solvers amenable to gradient-based optimization. We demonstrate the capabilities of this framework across four classes of applications: (1) learning Cahn-Hilliard and Allen-Cahn free energies and kinetic laws from noisy spatiotemporal data on complex battery electrode microstructures; (2) designing time-dependent wetting boundary conditions that steer phase boundaries to desired orientations; (3) optimizing spatially varying reaction rates, interpreted as surface coatings, to suppress phase separation in intercalation electrode models; and (4) computing time-dependent trapping potentials that transfer Bose-Einstein condensates between ground states of qualitatively different Gross-Pitaevskii potentials while minimizing excitations. Together, these results show that variational PDE models, when equipped with differentiable solvers on complex domains, provide a versatile substrate for data-driven discovery, design, and control of pattern-forming materials and quantum fluids, and they point toward tighter integration of physics-based PDE modeling with modern optimization and machine learning pipelines.

DOI: [10.1103/b8kc-vpwq](https://doi.org/10.1103/b8kc-vpwq)

I. INTRODUCTION

Many technologically important materials exhibit pattern-forming dynamics that have a critical impact on their performance and functionality [1–7]. In quantum systems, for example, specific density and phase patterns are desired in quantum states for computing applications [8–14]. In energy materials, chemically driven phase separation produces heterogeneous patterns in electrodes that influences reaction kinetics, mechanical stability, and degradation pathways [15–17]. In biological matter, phase separation and pattern formation at the subcellular and macroscopic scales are involved in spatially organizing biochemical reactions and are implicated in numerous diseases [3,18,19]. Across these domains, the ability to model, predict, and ultimately control pattern dynamics is essential for optimizing performance, designing new functionality, and guiding experiments.

Recent advances in experimental imaging and control have created unprecedented opportunities to observe and manipulate pattern-forming systems. High-resolution imaging modalities enable detailed measurements of the structure and dynamics of materials from the quantum scale to the macroscale [2,5,20–25]. Furthermore, these systems can often be driven or perturbed in controlled ways, including laser-based manipulation of quantum gases [26], surface treatments of electrodes [27,28], and boundary conditions in continuum materials [29]. As imaging and control advances, there is a growing demand for accurate, interpretable, and optimizable models that connect experimental observations to underlying physical processes and design parameters.

A broad mathematical framework for describing such pattern dynamics is based on energy-driven partial differential equations (PDEs) [30]. Many physical systems evolve according to equations derived from a functional derivative of an energy or free energy functional, which encodes the local and nonlocal interactions governing pattern formation. Dissipative systems follow gradient-flow dynamics that decrease the free energy, whereas Hamiltonian systems evolve perpendicular to the gradient direction and hence conserve energy. This distinction is separate from conservation of the order parameter itself: In dissipative systems, the order parameter may be conserved or nonconserved, depending on whether the evolution obeys a local continuity law. A wide variety of familiar models fall under this umbrella: the Cahn-Hilliard

*Contact author: aecohen@mit.edu

†Contact author: stm16109@mit.edu

‡Contact author: dunkel@mit.edu

§Contact author: bazant@mit.edu

and Allen-Cahn equations for phase transformations [31,32]; phase-field models for chemomechanical problems [33], fracture [34], and polymeric materials [35]; biological pattern formation [36]; and the Gross-Pitaevskii equation for Bose-Einstein condensates [17,37]. Despite describing systems that span many orders of magnitude in scale and physical origin, these models share a common variational structure: a free energy or energy functional and evolution equations obtained from its variational derivative. This structure ensures thermodynamic consistency, supports physical interpretability, and provides a natural foundation for parametrization, learning, and optimization.

However, learning, optimizing, and controlling free-energy-driven PDEs in complex geometries remains a significant challenge. Real experimental domains are irregular, noisy, and discretized on pixel grids, which can be difficult to convert into a computational form for most PDE solvers. Estimating parameters of a free energy functional from experimental data constitutes a PDE-constrained inverse problem that is usually solved by differentiating through the nonlinear PDE dynamics. In some cases, direct measurements of the phase variable may not be available, and model parameters must instead be inferred from indirect measurements [38]. Similarly, controlling solutions of these PDEs, either through boundary condition terms or through bulk control terms, requires solving PDE-constrained optimal control problems. These tasks involve differentiating through long time-stepping procedures for stiff, nonlinear PDEs, a process that has traditionally required specialized adjoint equations, careful manual derivations, and case-by-case implementations. Machine learning surrogate models aim to speed up these optimization problems, but they still require large upfront training cost and these purely data-driven methods may not accurately represent the true physics [39,40]. As a result, high-level tasks such as model discovery, parameter inference, and optimal control have remained difficult to perform reliably, especially on noisy, irregular, experimentally derived domains.

In this work, we develop a unified computational framework that enables differentiable learning and control of free-energy-driven pattern dynamics in complex geometries. Our approach combines advanced numerical methods for nonlinear pattern-formation PDEs with PDE-based geometry smoothing and the smoothed boundary method to construct differentiable PDE solvers that operate directly on image-based domains. We demonstrate how free energy parametrizations can be learned from data using PDE-constrained optimization, how boundary conditions and spatial heterogeneities can be optimized to manipulate or suppress pattern formation, and how time-dependent external potentials can be controlled to reshape the states of quantum fluids. Using this framework, we learn Cahn-Hilliard and Allen-Cahn models in complex geometries reconstructed from imaging, control wetting boundary conditions to induce rotation of phase boundaries, design spatial heterogeneity fields that prevent phase separation in energy materials, and optimize time-varying trapping potentials in the Gross-Pitaevskii equation to steer Bose-Einstein condensates into experimentally relevant nontrivial quantum states. Together, these results establish a general, extensible methodology for data-driven discovery, optimization, and control of pattern-forming dynamics gov-

erned by free-energy-based PDEs across a broad range of physical systems.

II. THEORY

Many pattern-forming systems can be described within a variational framework, which provides a unifying description for all models considered in this work. We introduce a real- or complex-valued field $\Psi(\mathbf{x}, t)$ defined on a d -dimensional domain $\Omega \subset \mathbb{R}^d$, together with an energy or free energy functional of the form

$$\mathcal{F}[\Psi] = \int_{\Omega} f(\Psi, \nabla\Psi, \mathbf{x}, t) d\mathbf{x}. \quad (1)$$

The energy density f may include local terms, gradient contributions, nonlocal interactions, and external potentials. In dissipative systems, \mathcal{F} plays the role of a Gibbs free energy that the system seeks to minimize, while in Hamiltonian systems it represents a conserved energy.

The connection between the functional \mathcal{F} and dynamics is made through its variational derivative. For complex-valued fields, this derivative takes the form $\delta\mathcal{F}/\delta\Psi^*$, while for real-valued fields it reduces to $\delta\mathcal{F}/\delta\Psi$. The specific time evolution of the field depends on the physical nature of the system, and, in particular, on whether the dynamics are dissipative or Hamiltonian.

A. Dissipative systems

In dissipative systems, the dynamics drive the system toward states that reduce the free energy. Such systems are commonly described by an order parameter representing a physically conserved or nonconserved quantity. In many applications, and particularly in electrochemistry [17,41], the order parameter represents a concentration field, denoted by $\Psi = c(\mathbf{x}, t)$. The thermodynamic driving force for the temporal evolution of c is the chemical potential, which is defined as the variational derivative of the free energy with respect to the concentration,

$$\mu = \frac{\delta\mathcal{F}}{\delta c}. \quad (2)$$

(For an open system at constant temperature and pressure, \mathcal{F} is the Gibbs free energy, usually denoted by \mathcal{G} .)

When the order parameter is conserved as a result of transport processes, its dynamics are governed by a local conservation law

$$\frac{\partial c}{\partial t} + \nabla \cdot \mathbf{J} = 0, \quad (3)$$

where \mathbf{J} is the flux density. Within the framework of linear irreversible thermodynamics, the flux density is assumed to be proportional to the thermodynamic driving force for sufficiently small departures from local equilibrium. For diffusive transport, this driving force is the gradient of the chemical potential, yielding

$$\mathbf{J} = -M(c)c\nabla\mu, \quad (4)$$

where $M(c)$ is the concentration-dependent mobility [17]. Combining Eqs. (3) and (4) yields

$$\frac{\partial c}{\partial t} = \nabla \cdot \left(M(c) c \nabla \frac{\delta \mathcal{F}}{\delta c} \right), \quad (5)$$

which reduces to the Cahn-Hilliard equation for appropriate choices of $M(c)$ and \mathcal{F} [31].

For nonconserved order parameters, the dynamics are governed by sources and sinks, such as chemical reactions, in addition to diffusive transport. In the simplest case, the evolution of the concentration is determined by a local reaction rate rather than a continuity equation. Near equilibrium, linear irreversible thermodynamics assumes a linear relation between the reaction rate and the chemical potential,

$$\frac{\partial c}{\partial t} = -k(c) \mu, \quad (6)$$

where $k(c)$ is a kinetic coefficient that may be state dependent. This form corresponds to Allen-Cahn-type dynamics. Reaction kinetics may depend nonlinearly on the chemical potential. In this case, the evolution is written as

$$\frac{\partial c}{\partial t} = R(c, \mu), \quad (7)$$

where the reaction rate function R is chosen to be thermodynamically consistent [17].

More generally, chemically reactive mixtures are governed by systems of reaction-diffusion equations, which have the thermodynamically consistent form [41],

$$\frac{\partial c_i}{\partial t} = \nabla \cdot \sum_j L_{ij}(\{c_k\}) \nabla \frac{\delta \mathcal{F}}{\delta c_j} + \sum_m s_{i,m} R_m(\{c_k\}, \left\{ \frac{\delta \mathcal{F}}{\delta c_k} \right\}), \quad (8)$$

where L_{ij} is the Onsager linear response matrix and R_m is the rate of the m th reaction with stoichiometric number $s_{i,m}$ for species i . Although we only consider single species that evolve only by reaction or by diffusion, our framework for differentiable learning and control can be readily extended to the general system of reaction-diffusion PDEs as in Eq. (8).

B. Hamiltonian systems

In contrast to dissipative systems, Hamiltonian systems conserve energy. A prominent example is a Bose-Einstein condensate, whose complex-valued order parameter ψ represents the condensate wavefunction. The dynamics of ψ are governed by the Gross-Pitaevskii equation, which can be written in variational form as [37,42–44]

$$i\hbar \frac{\partial \psi}{\partial t} = \frac{\delta E}{\delta \psi^*}. \quad (9)$$

The associated energy functional is

$$E[\psi] = \int_{\Omega \subseteq \mathbb{R}^3} \left[\frac{\hbar^2}{2m} |\nabla \psi|^2 + V(\mathbf{x}) |\psi|^2 + \frac{g}{2} |\psi|^4 \right] d\mathbf{x}, \quad (10)$$

which includes kinetic energy, external potential energy, and interaction energy between bosons in the dilute limit.

C. Parametrization of models

In this framework, all of these systems are unified by dynamics driven by a variational derivative of an energy or free energy functional. Across both dissipative and Hamiltonian systems, the structure of the dynamics is determined by the choice of energy or free energy functional and the physical constraints governing its evolution. In practice, these models contain parameters describing material properties, including free energies \mathcal{F} [Eq. (5)], interaction strengths g [Eq. (10)], transport coefficients such as mobilities $M(c)$ [Eq. (5)], and reaction kinetics $R(c)$ [Eq. (8)]. In this work, we consider parametrizations of both the energy landscape and the associated dynamics, allowing controlled modification of model behavior while preserving the underlying variational structure. The specific form of these parametrizations is problem dependent and informed by the physics of the system under study.

III. METHODS

To perform PDE-constrained optimization on these models, we develop a differentiable PDE solver that supports stiff diffusion operators, complex geometries derived from images, and parametrized free energies and control fields (Fig. 1). This section describes the computational methods used to perform geometry processing, differentiable time integration, and large-scale optimization.

A. Domain smoothing

Experimental domains are frequently obtained as binary masks from segmentation of microscopy data. To convert these pixelated shapes into smooth, differentiable geometric fields, we solve a modified Allen-Cahn equation with the gradient term $f_{\text{grad}} = \varepsilon |\nabla \phi|^2$ and a double-well homogeneous energy $f_h = 9\phi^2(1-\phi)^2/\varepsilon$ such that ε sets the width of the diffuse interface between domains of $\phi = 0$ and $\phi = 1$. This double-well form is the lowest-order Landau polynomial with stable minima at the two binary phases, and the prefactor $9/\varepsilon$ normalizes the diffuse-interface energy to unity following the Allen-Cahn treatment [45,46]. To preserve the original shape of the binary mask, we remove the effect of curvature minimization [47,48], which leads to the evolution equation

$$\frac{\partial \phi}{\partial t} = \nabla^2 \phi - (1-K) |\nabla \phi| \nabla \cdot \frac{\nabla \phi}{|\nabla \phi|} - \frac{9}{\varepsilon^2} \phi(1-\phi)(1-2\phi). \quad (11)$$

Here, $K \in [0, 1]$ specifies the amount of curvature minimization in the system, where a value of $K = 1$ includes all curvature driving forces that naturally arise from the Laplacian, while a value of $K = 0$ fully removes the mean-curvature contribution from the Laplacian term [48,49]. In practice, a value of around $K = 0.01$ is typically sufficient to retain fine details of the shape and remove pixelization artifacts (Appendix A). The smooth transition region between $\phi = 0$ and $\phi = 1$ will take a hyperbolic tangent form, which makes it consistent with phase field modeling for microstructural evolution [50]. This step is also fully differentiable, enabling gradient-based shape optimization. The smoothed boundary method is then used to modify the governing PDE to find

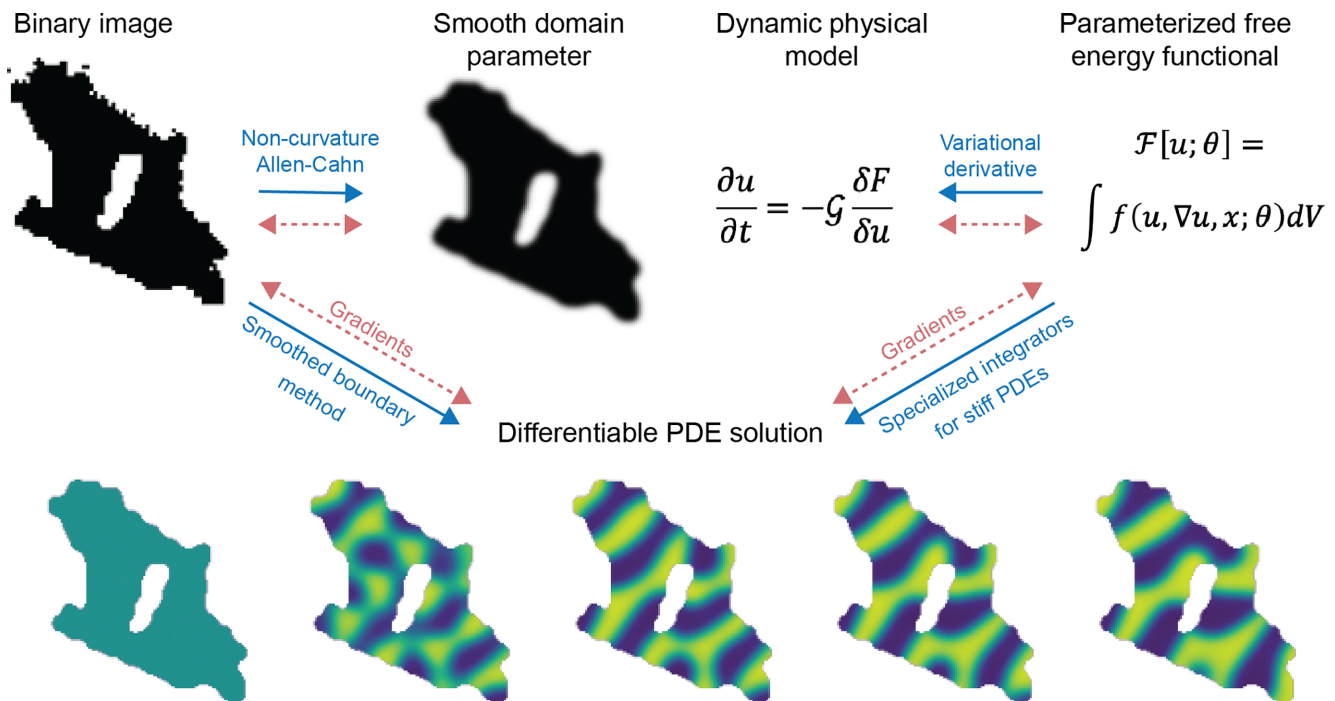


FIG. 1. Overview of methods for producing differentiable simulations of parametrized free energy functional models on smoothed domains from binarized image data. Binary images are converted to smooth domain parameters through the curvature-subtracted Allen-Cahn equation. Variational derivatives of parametrized free energy functionals form PDEs, which can be numerically solved with specialized integrators. Blue arrows indicate the forward simulation process, while red arrows indicate the flow of gradients through the computational pipeline, which can be computed using either forward-mode or reverse-mode automatic differentiation.

a numerical solution on the domain and apply the boundary conditions at the smooth interface [50].

B. Smoothed boundary method

A general framework to solve PDEs on irregular domains without the need for advanced meshing techniques is the smoothed boundary method [50], which has been successfully applied to battery materials [51] and materials science more broadly [50]. The smooth domain mask $\phi \in [0, 1]$ obtained by solving Eq. (11) is used to modify the original PDE such that boundary conditions are imposed on the diffuse transition region where $\nabla\phi \neq 0$. Following the procedure in Ref. [50], we can reformulate the Allen-Cahn equation as

$$\frac{\partial c}{\partial t} = R(c, \mu^*), \quad \text{where } \mu^* = \frac{\partial f_h(c)}{\partial c} - \frac{\kappa}{\phi} \nabla \cdot (\phi \nabla c) + g_\varphi. \quad (12)$$

The term $g_\varphi = -\frac{\kappa}{\phi} |\nabla\phi| |\nabla c| \cos\varphi$ in the modified chemical potential μ^* imposes the boundary angle φ of the concentration field with the smoothed boundary. Similarly, the Cahn-Hilliard equation is modified to

$$\phi \frac{\partial c}{\partial t} = \nabla \cdot \left(\phi \frac{D(c)c}{k_B T} \nabla \left[\frac{\partial f_h(c)}{\partial c} - \frac{\kappa}{\phi} \nabla \cdot (\phi \nabla c) + g_\varphi \right] \right) + \psi R(c) + g_{BC}, \quad (13)$$

where $R(c)$ is a bulk source and $g_{BC} = |\psi| j_N$ represents the boundary flux term. j_N denotes the normal boundary flux that can vary spatially and/or temporally; for a closed system,

$j_N = 0$ holds. For the derivation of other boundary conditions see Ref. [52].

C. Numerical optimization

The PDE solvers are implemented in JAX, enabling just-in-time compilation using the XLA (Accelerated Linear Algebra) compiler, acceleration on graphics processing units (GPUs), vectorization over initial conditions and parameters, and automatic differentiation (AD). Optimization is performed using gradient-based methods, with gradients obtained by differentiating directly through the time-discretized PDE system. Depending on the parametrization and objective function, either forward- or reverse-mode AD is employed. Forward-mode AD has a computational cost that scales linearly with the number of parameters, while reverse-mode AD has a cost that is independent of the number of parameters for scalar-valued objectives. These trade-offs influence both algorithm choice and convergence behavior and are discussed further in Appendix B.

Many pattern-forming PDEs contain stiff diffusion operators such as ∇^2 or ∇^4 terms, which impose severe stability constraints on explicit time-stepping schemes. To address this, we employ semi-implicit Fourier spectral methods and operator splitting techniques that treat stiff linear terms implicitly or exactly, enabling larger stable time steps. We also use a stabilized explicit method, ROCK2 [53], which provides a flexible alternative without requiring implicit solves. These methods are most useful when diffusive or biharmonic terms dominate the stiffness: For an explicit scheme, the stable time

step typically scales as $\Delta t_{\max} = O(\Delta x^h)$, where h is the order of the highest spatial derivative. When such terms are absent, the same differentiable framework can use standard explicit or adaptive integrators, with the stable time step instead limited by reactions, advection, waves, nonlinear source terms, or accuracy requirements. In our numerical tests, semi-implicit and stabilized methods typically allowed stable time steps that were 10–100 times larger than explicit schemes, depending on the PDE and initial condition; this is consistent with reported speedups of two to three orders of magnitude for semi-implicit Fourier spectral methods applied to phase-field equations [54].

We consider optimization problems of the form

$$\min_{\theta} \mathcal{L}(\Psi_{\theta}(\mathbf{x}, t), \theta), \quad (14)$$

subject to the PDE dynamics defined by the free energy functional \mathcal{F} . The parameters θ may specify the free energy density, boundary conditions, time-dependent controls, or domain geometry. When the objective is a least-squares problem and the number of parameters is small or moderate, we employ the Levenberg-Marquardt (LM) algorithm, which benefits from faster convergence through a Gauss-Newton approximation of the Hessian constructed using forward-mode AD of the residuals with respect to the parameters [55] (Appendix B). For problems with large parameter sets or when a Gauss-Newton approximation is not feasible, we use the Broyden-Fletcher-Goldfarb-Shannon (BFGS) algorithm (Appendix B).

IV. RESULTS

Together, these components provide an end-to-end differentiable simulation and optimization framework that enables learning and control for a wide range of variational PDE models. The combination of advanced time integration, domain smoothing, smoothed-boundary operators, and memory-efficient gradient computation allows us to optimize both physical parameters and external controls directly through the PDE dynamics. In the next section, we demonstrate how these tools allow us to (1) infer free energy and kinetic parameters from data on complex geometries, (2) manipulate boundary conditions to achieve targeted interface behavior, (3) optimize spatial heterogeneities to suppress or enhance phase separation, and (4) design time-dependent control protocols for steering quantum fluids.

A. Learning models

To demonstrate the ability of our framework to learn constitutive models directly from data on complex domains, we perform a set of inverse problems in which the objective is to recover an unknown free energy density (or equivalently its variational derivative) from noisy spatiotemporal observations of phase-field evolution. We consider both Cahn-Hilliard [Eq. (5)] and Allen-Cahn [Eq. (6)] dynamics, which, respectively, represent conserved and nonconserved order-parameter evolution, and we perform all inference directly on geometries extracted from image data of real battery graphite electrodes [56] using the smoothed-boundary method [50,57,58].

For each model, we generate synthetic training data by simulating the corresponding PDE in a complex domain and adding varying levels of additive noise. The reported noise level is defined relative to the standard deviation of the simulated concentration field; for example, 30% noise corresponds to additive Gaussian noise with standard deviation 0.3 times the standard deviation of the clean simulation data. The variational derivative of the free energy, the chemical potential, is parametrized as the entropic contribution, an energetic contribution represented by Legendre polynomials, and the gradient term,

$$\frac{\delta \mathcal{F}}{\delta c} = \ln \frac{c}{1-c} + \sum_{n=0}^N P_n(c) - \kappa \nabla^2 c. \quad (15)$$

The mobility in the Cahn-Hilliard equation and the reaction rate's concentration-dependent prefactor in the Allen-Cahn equation are also parametrized by Legendre polynomials, which we exponentiate to ensure positivity. We then optimize the free energy parameters and diffusivity parameters such that the simulated dynamics best match the noisy observations.

The loss function is formulated as the squared error between the simulated field and the observed data,

$$\mathcal{L}(\theta) = \|\Psi_{\theta}(t) - \Psi_{\text{data}}(t)\|^2 + \lambda \theta^{\top} W \theta, \quad (16)$$

where θ is the free energy and kinetic parameters, λ is a regularization parameter, and W is a matrix chosen to regularize derivatives of the functions parametrized by θ (Appendix C). In particular, W is chosen to be the matrix that encodes the Legendre polynomial derivative identities [59]. Optimization is performed with the LM algorithm, which is particularly effective here because the objective is a squared error and the number of learned parameters is moderate due to the one-dimensional Legendre representation. Furthermore, we use a multiple shooting approach in which we sum Eq. (16) for multiple nonoverlapping time intervals. We found this approach to be more robust to noise than a single shooting method and faster due to JAX vectorization over initial conditions (Appendix D).

Our PDE-constrained optimization method successfully recovers the known Cahn-Hilliard model at 30% noise levels [Figs. 2(a) and 2(b)]. Starting from a non-phase-separating free energy initial condition with constant mobility, the models quickly converge to the ground truth. By the fifth optimization step, the LM algorithm has already begun to learn a phase-separating free energy. After ten steps, the homogeneous free energy function closely resembles the ground-truth model and the mobility starts to resemble the ground-truth model. After 90 steps, the learned functional is nearly indistinguishable from the true one, demonstrating rapid convergence and robustness to noise. The only exception is in regions where the order parameter approaches 0 or 1, beyond the equilibrium values of the two domains, which rarely appear in the data.

Similar convergence trends are observed for the Allen-Cahn equation [Figs. 2(c) and 2(d)]. The free energy and reaction prefactor are progressively refined through 5, 10, and 90 iterations [Fig. 2(d)]. Even after only ten steps, the homogeneous free energy and reaction rate closely match the

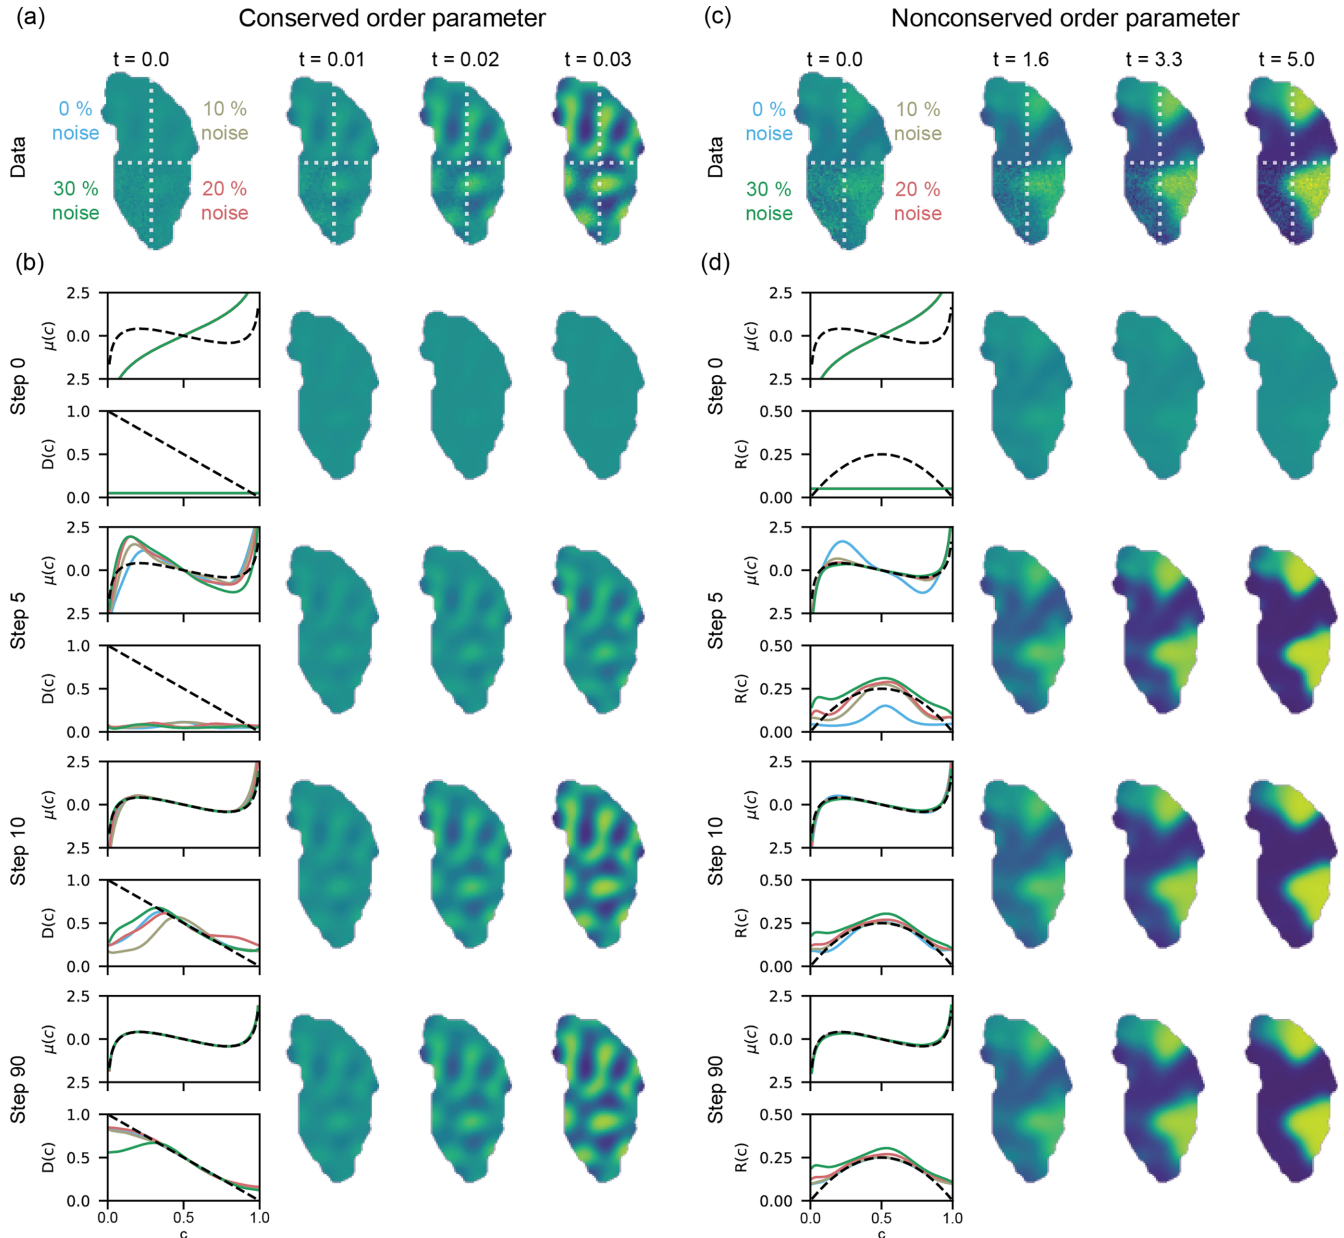


FIG. 2. PDE-constrained optimization recovers conservative (Cahn-Hilliard) and nonconservative (Allen-Cahn) models of phase separation in complex geometries from noisy data. (a) Snapshots from a Cahn-Hilliard simulation used as training data, with four different noise levels added. (b) Convergence of the chemical potential and diffusivity functions during the optimization, along with the simulated model with the corresponding functions. Dashed lines represent ground truth. (c) Snapshots from an Allen-Cahn simulation used as training data, with four different noise levels added. (d) Convergence of the chemical potential and reaction prefactor functions during the optimization, along with the simulated model with the corresponding functions. Dashed lines represent ground truth. Columns correspond to different simulation time points. Line colors indicate the noise level used in the training data: 0% noise is light blue, 10% noise is olive, 20% noise is soft pink, and 30% noise is green.

ground-truth values, again with the exception of regions near $c = 0$ and $c = 1$.

This experiment illustrates the ability of our framework to learn physical models directly from noisy data on arbitrary complex geometries. Furthermore, the combination of performant PDE solvers with the vectorized loss function evaluation and fast convergence of the LM method allows us to accomplish this in minutes on a single GPU.

B. Controlling phase boundary motion

The orientation of a phase boundary in contact with a solid surface is governed by the wetting angle, which is set by the balance of solid-fluid and fluid-fluid surface tensions through Young's equation. In practice, this angle can be modified dynamically, as in the case of electrowetting in which an applied voltage can thermodynamically change the equilibrium wetting angle [60], which has numerous technological applications including in lenses and displays [61]. Here, we

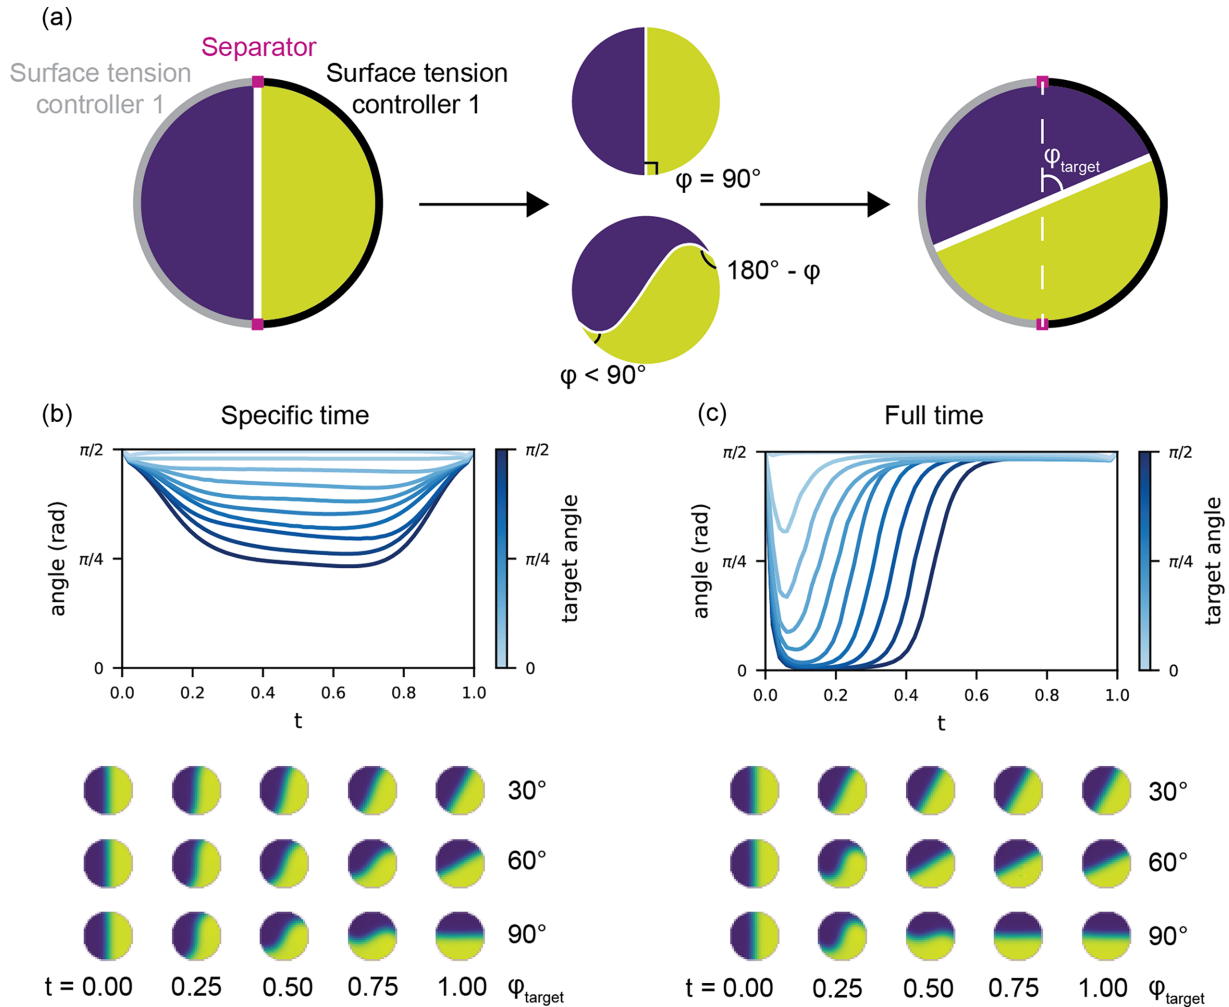


FIG. 3. Successful control of phase boundary motion by controlling boundary conditions using the differentiable simulations. By modulating the surface wetting angles of the two phases, we can rotate the phase boundary to desired angles. (a) Schematic of the system geometry and goal boundary rotation. (b) Optimized wetting angle protocol and resulting simulations to rotate the phase boundary to a desired orientation at the final time. (c) Optimized wetting angle protocol and resulting simulations to rotate the phase boundary to a desired orientation penalizing deviations from the goal orientation throughout the entire trajectory.

demonstrate how our differentiable PDE-constrained optimization framework can be used to determine time-varying wetting angles that rotate a phase boundary inside a circular domain.

The system consists of a circular domain whose boundary is divided into two halves [Fig. 3(a)]. On the left half, the wetting angle is set to φ by the control protocol, where $0 < \varphi < \pi/2$. On the right half, the wetting angle is set to $\pi - \varphi$. Thus, the contact angles will appear opposite on each side of the domain [Fig. 3(a)]. The increased interfacial free energy caused by the longer, curved phase boundary induces its rotation, with the dynamics working to reduce this energy. The surface energy or contact angle term boundary condition enters the equation through the modified chemical potential, as described by Eq. (13).

In the first experiment, we seek to rotate the interface to a specific orientation at the final simulation time [Fig. 3(b)]. The optimization objective is

$$\mathcal{L}(\theta) = \|\varphi_{\text{sim}}(T; \theta) - \varphi_{\text{target}}\|^2, \tag{17}$$

where $\varphi_{\text{sim}}(T; \theta)$ is the simulated phase boundary orientation at the final time T and θ are parameters of the time-dependent contact angle function, which is represented through a piecewise linear interpolation that is constrained to be between 0 and $\pi/2$ (Appendix E). We also add an L2 regularization to these parameters to favor smooth trajectories (Appendix E).

Beginning from an initial guess for θ that remains constant at $\pi/2$ during the course of the simulation, the optimizer rapidly adjusts the wetting condition to drive the interface toward the desired orientation. Despite the nonlinear coupling between the boundary condition and the interface shape, the differentiable PDE solver provides stable gradients, and the optimizer converges to the correct control value in only 10–20 iterations (Appendix E).

In the second experiment, we impose an objective function that drives the interface orientation to match the desired angle at every time point, creating a smooth relaxation of the minimum-time problem [Fig. 3(c)]. The optimization

objective becomes

$$\mathcal{L}(\theta) = \int_0^T \|\varphi_{\text{sim}}(t; \theta) - \varphi_{\text{target}}\|^2 dt. \quad (18)$$

The optimal solution achieves the desired phase-boundary orientation and does so by driving the contact angle closer to zero, producing a faster rotation [Fig. 3(c)]. Together, these experiments illustrate how boundary-energy manipulation through processes like electrowetting can be framed as a PDE-constrained control problem amenable to gradient-based optimization, enabling precise steering of interface orientation.

C. Controlling phase separation

Phase separation is ubiquitous in many intercalation electrode materials, including graphite [62], lithium iron phosphate (LFP) [63,64], lithium titanate [65], and spinel lithium nickel manganese oxide [66]. During operation, these materials often segregate into coexisting high- and low-density phases, producing heterogeneous spatial patterns within electrodes. The onset, morphology, and persistence of such phase-separated domains strongly influence battery performance, state-of-charge estimation, and long-term degradation [17,28,67–69]. Accordingly, understanding and ultimately controlling phase separation offers a promising pathway to improve both the efficiency and lifetime of electrochemical energy storage systems.

One practical mechanism for influencing battery electrode material behavior is the use of surface coatings, which modify the local intercalation kinetics. For example, LFP, a widely used lithium-ion cathode with inherently phase-separating thermodynamics, relies on carbon surface coatings to provide the electronic conductivity needed for practical battery operation [28]. In this section, we demonstrate how our framework can optimize a spatially varying surface-coating profile to minimize phase separation during lithiation.

We model the intercalation dynamics in LFP using an experimentally validated Allen-Cahn reaction model with nonlinear ion-coupled electron transfer (ICET) reaction kinetics [70]:

$$\frac{\partial c}{\partial t} = k(x, y)R(c, \mu), \quad (19)$$

where c is the lithium concentration, $k(x, y)$ is a spatially varying reaction-rate prefactor that is proportional to the carbon coating thickness, $R(c, \mu)$ is the ICET reaction rate, and μ is the chemical potential (Appendix F) [28]. The simulations are performed on real LFP nanoparticle geometries obtained from x-ray imaging [2] under slow-rate galvanostatic conditions.

To minimize the effect of phase separation, we solve the following design problem:

$$\begin{aligned} & \max_{k(x, y)} \text{UC}[c(\mathbf{x}, t)], \\ \text{UC} = 1 - & \frac{\int_0^T \sqrt{\text{Var}[c(\mathbf{x}, t; \mathbf{p})] \cdot \bar{c}(t)(1 - \bar{c}(t))} dt}{\int_0^T \bar{c}(t)(1 - \bar{c}(t)) dt}, \end{aligned} \quad (20)$$

where UC is the uniformity coefficient that quantifies the concentration uniformity throughout the particle, $\bar{c}(t)$ is the spatially averaged concentration, and T is the time interval.

The uniformity coefficient is maximized at $\text{UC} = 1$ when $c(x, y) = \bar{c}$ at all (x, y) . To tractably solve this problem, we represent $k(x, y)$ in the basis of Laplacian eigenfunctions on the particle geometry (Appendix F) and optimize the corresponding coefficients. To ensure that heterogeneity does not alter the overall reaction rate, we impose a linear constraint that preserves the spatial mean of $k(x, y)$ throughout the optimization (Appendix F). Lower-frequency modes correspond to coarse spatial modulation, while higher-frequency modes enable finer spatial control. This representation naturally encodes experimentally realistic constraints on how sharply a coating can vary.

We solve the design problem in Eq. (20) using an increasing number of basis modes and quantify the resulting UC as a function of the spectral length scale associated with each basis [Fig. 4(a)]. The spectral length scale is defined as $l = 1/\sqrt{\lambda_{\text{max}}}$, where λ_{max} is the largest eigenvalue of the Laplacian basis. As l decreases, the optimal UC increases monotonically, reaching a critical length scale near $l \approx 200$ nm, beyond which nearly perfect uniformity is achieved. The corresponding optimal coatings become progressively more finely structured as the spectral length scale decreases [Fig. 4(b)].

To assess the resulting dynamics, we simulate lithiation under both a spatially uniform coating and the optimized coating [Fig. 4(c)]. A uniform reaction rate leads to pronounced phase separation, whereas optimization of $k(x, y)$ produces spatially patterned reaction kinetics that suppress phase separation and maintain near-uniform concentration throughout lithiation. Consistent with this behavior, the optimized coating yields a smooth, solid-solution-like voltage profile, in contrast to the plateaus and voltage jumps characteristic of phase-separating dynamics [Fig. 4(d)].

These results demonstrate that spatially patterned reaction-rate modifications can strongly influence phase separation in intercalation materials. In particular, our findings help explain the experimental observation that nonuniform glucose-derived carbon coatings lead to better rate capability than uniform polydopamine coatings [71]. While improved electrical connectivity in the porous electrode controls interparticle phase separation [72], our work shows that these nonuniform coatings provide the additional benefit of suppressing intraparticle phase separation.

Furthermore, we show that differentiable PDE-constrained optimization can identify and map out effective coating profiles directly on real particle geometries. For example, the optimal carbon coating for realistic LFP platelet particles requires a thicker layer (corresponding to lower reaction rates) in a band around the side crystal facets perpendicular to the b direction [Figs. 4(d) and 8]. This optimized coating can be understood as a strategy to interfere with coherent nucleation [15,73]. The high-density phase is predicted to nucleate on the side facets and propagate to the center only when a sufficiently rich layer has been created to overcome the coherency strain resulting from crystal lattice mismatch of the high- and low-density stable phases. The slower reaction rate from a thicker coating around the rim suppresses the growth of the nucleated surface phase, as the lithium flux is redistributed to the central region of the active b -direction crystal facet with a faster reaction rate from a thinner coating. Such carbon

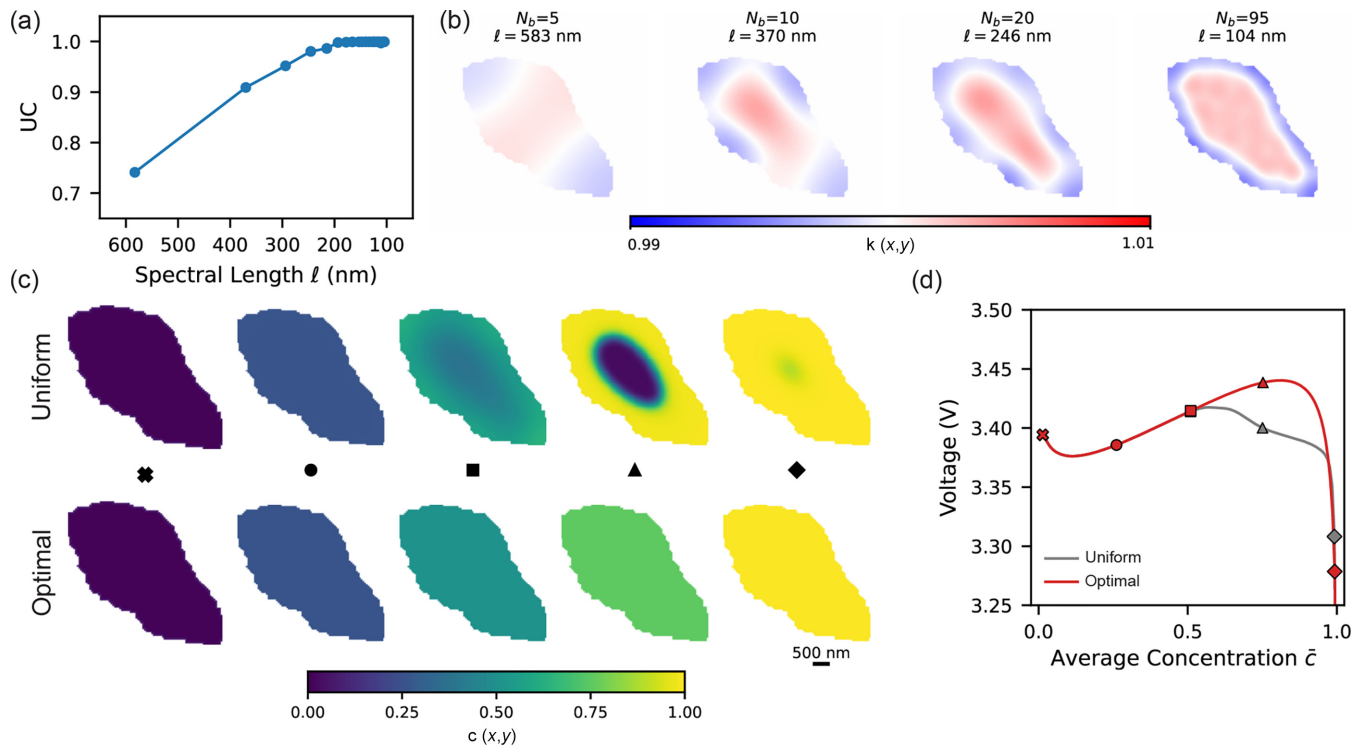


FIG. 4. Suppression of phase separation in LFP via spatially optimized coatings. (a) The maximum spectral length scale of the optimized coating dictates the limit of control; below a critical length scale, concentration uniformity is maximized. (b) Evolution of optimal coating distributions as the spectral length scale decreases, allowing for finer spatial control. (c) Simulated lithiation dynamics comparing uniform vs optimized coatings. The optimized coating maintains internal homogeneity, whereas the uniform coating undergoes strong phase separation. (d) The optimized coating results in a voltage profile exhibiting solid-solution-like dynamics.

coatings profiles could potentially be manufactured by chemical pre patterning of the side facets to favor wetting during synthesis by pyrolysis or sol-gel routes. This is one example of how our approach could enable the design of faster-charging, longer-lasting battery electrodes with suppressed phase separation and improved performance.

D. Controlling quantum states

Bose-Einstein condensates (BECs) play an important role in quantum simulation, quantum sensing, and quantum information processing. Many of these applications require preparing the condensate in nontrivial spatial configurations, and achieving this reliably demands precise dynamical control over the external trapping potential. In principle, such state preparation would ideally proceed adiabatically, ensuring that the condensate remains in the instantaneous ground state of the trap throughout the evolution. However, truly adiabatic transformations require long timescales that are often incompatible with experimental constraints. This challenge has motivated the development of shortcuts to adiabaticity [74], which aim to reproduce the outcome of slow adiabatic processes within finite time. In practice, condensate control is often limited by unwanted excitations, such as sloshing and breathing modes, that arise when traps are changed too quickly or along poorly chosen trajectories [75]. Optimal control theory has therefore become an important tool for designing experimentally feasible protocols that

suppress such excitations while preparing complex quantum states [8,75,76].

Two geometries of particular interest are split condensates, which are used in quantum interferometry and metrology, and toroidal traps, which have been used as quantum analogs of electrical circuits [8–13,75]. Motivated by these applications, we apply our differentiable PDE-constrained optimization framework to design a control protocol that transforms an initial single-well ground state into a combined split-plus-toroid configuration in the shape of a face [Fig. 5(a)].

We parametrize a family of trapping potentials that begins as a single quadratic well and ends in a target potential whose ground state has the desired “face” topology [Fig. 5(b)]. Details of this parametrization are provided in Appendix G. A simple linear interpolation between the initial and final potentials produces significant excitations and density oscillations in the condensate [Fig. 5(a)]. To quantify the mismatch between the evolving state $\psi(t)$ and the desired final state ψ_d , we use [8]

$$\text{infidelity} = 1 - |\langle \psi_d, \psi \rangle|^2. \quad (21)$$

We model the condensate using a two-dimensional reduction of the three-dimensional Gross-Pitaevskii dynamics in Eq. (9), with the energy functional in Eq. (10), appropriate for disk-shaped BECs [37]. The optimization objective minimizes the infidelity from the midpoint of the simulation onward, ensuring that the desired final state is reached at the midpoint and that there are no oscillations after the final state is reached. Our optimized trap trajectory achieves a significantly lower

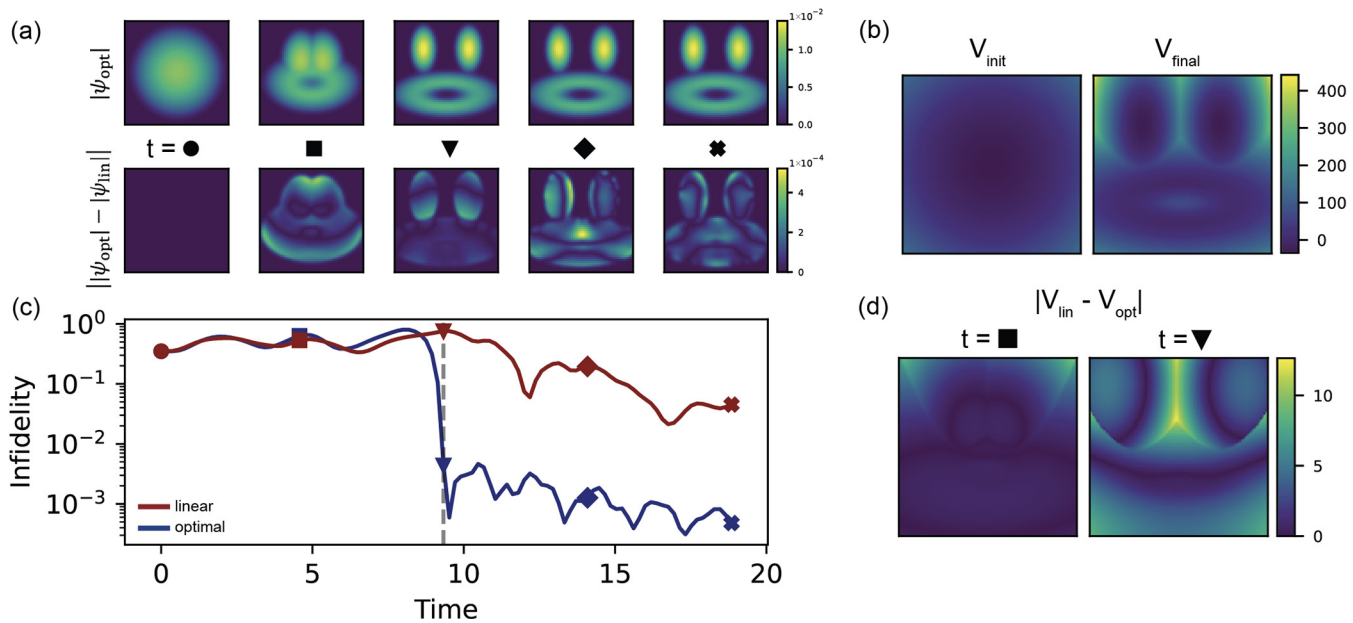


FIG. 5. Optimization of time-dependent applied potentials to dynamically transfer a BEC from an initial single well potential ground state to a split state and toroid. (a) The optimal transition between the two potentials successfully reaches the desired state without oscillations. The difference between the optimal transition and linear transition reveals oscillations due to the linear protocol. Snapshots are shown at $t = 0.0, 4.6, 9.3, 14.1,$ and 18.8 , corresponding, respectively, to the circle, square, triangle, diamond, and x symbols. (b) The initial and final potentials during the state transition. (c) The differences between the linear and optimal protocols are quantified by the infidelity [Eq. (21)]. (d) The difference between the linear and optimal protocols at two time points during the simulation.

infidelity compared to the naive linear ramp and produces a smooth transition from the single-well ground state to the desired split-toroid configuration [Fig. 5(c)]. The difference between the optimized and linearly interpolated potentials is shown in Fig. 5(d), illustrating how the optimizer automatically discovers an effective trap-shaping protocol.

These results demonstrate that differentiable PDE-constrained optimization provides a powerful tool for quantum-state engineering in the Gross-Pitaevskii framework. By optimizing the time-dependent trapping potential, we obtain control protocols that avoid the excitations typically produced by linear trap ramps, enabling precise preparation of complex condensate geometries with direct relevance to quantum technologies.

V. CONCLUSIONS

In this work, we developed and demonstrated a unified framework for learning, optimizing, and controlling pattern-forming dynamics governed by free-energy-based partial differential equations in complex geometries. Starting from a variational description of dissipative and Hamiltonian systems, we show how a broad class of models, from electrochemical phase separation to quantum fluids, can be written in terms of an energy or free energy functional and a kinetic operator applied to its variational derivative. Building on this common structure, we construct differentiable PDE solvers that enable end-to-end optimization over model parameters, boundary conditions, spatial heterogeneities, and time-dependent control fields.

Methodologically, our approach combines several key ingredients: PDE-based smoothing of binarized image domains to obtain smooth boundary fields, the smoothed boundary method to impose physics on irregular shapes, advanced integrators for stiff pattern-forming PDEs, and memory-efficient automatic differentiation through long time-stepping procedures. Implemented in JAX, this pipeline supports GPU-accelerated simulation and gradient computation without hand-derived adjoints, making it straightforward to pose and solve PDE-constrained inverse and control problems for a wide variety of variational models.

We illustrate the capabilities of this framework through four representative applications. The first example shows that Cahn-Hilliard and Allen-Cahn models, including both free energy and kinetic terms, can be learned from noisy data on complex geometries reconstructed from real electrode microstructures. Second, we demonstrate control of phase boundary motion by optimizing the wetting angle or surface energy boundary conditions, using differentiable simulations to design protocols that rotate interfaces to desired orientations. Third, we optimize spatially varying reaction-rate fields, which are realized experimentally as surface coatings, to suppress phase separation in models of intercalation electrodes, highlighting how spatial patterning of kinetics can be used to control phase separation in real battery electrodes, such as carbon-coated LFP. Finally, we apply the same PDE-constrained optimization machinery to the Gross-Pitaevskii equation, designing time-dependent trapping potentials that steer Bose-Einstein condensates from a single-well ground state into a split-plus-toroidal configuration while minimizing excitations.

Taken together, these results show that variational PDE models, when coupled with differentiable solvers, can serve as a powerful substrate for data-driven discovery and control across quantum, electrochemical, and soft-matter systems. Future directions include incorporating more realistic experimental noise and uncertainty models, extending the framework to multiphysics and multiscale couplings, learning not just parameters but functional forms of free energies and kinetic operators, and closing the loop with experiments through real-time or batch experimental design. Conceptually, the framework could also be extended to stochastic free-energy-based PDEs by treating the model output probabilistically, for example, through a Bayesian formulation that infers distributions over parameters or trajectories rather than a single deterministic solution. More broadly, the ideas developed in this work point toward a general strategy for turning physically grounded PDE models into flexible, differentiable components within larger optimization and machine learning pipelines, enabling a tighter integration between theory, computation, and experiment across a wide range of materials.

ACKNOWLEDGMENTS

A.E.C. was funded by the National Defense Science and Engineering Graduate (NDSEG) Fellowship. This work was supported by DARPA APAQuS (J.D.), Schmidt Sciences, LLC (J.D.), and the MathWorks Professorship Fund (J.D.).

A.E.C. and S.D.-M. are co-first authors contributed equally to this work and J.D. and M.Z.B. are co-corresponding authors.

DATA AVAILABILITY

The data that support the findings of this article are openly available [77].

APPENDIX A: DOMAIN SMOOTHING

To assess the effect of the curvature-weighting parameter K and the interface width ε , we perform a systematic regularization study starting from a segmented binary mask [Fig. 6(a)]. The mask is evolved under Eq. (11) for 10 000 time steps with a $dt = 0.01$ using a fifth-order Runge-Kutta method.

Figure 6(b) summarizes the resulting steady-state geometries. When curvature effects are kept (larger K), the interface undergoes noticeable geometric relaxation, with rounding of high-curvature regions and shrinkage of small features. As K decreases, curvature-driven relaxation is progressively suppressed. For small K , the global shape remains effectively unchanged while pixel-scale irregularities are removed.

The parameter ε controls the diffuse interface thickness. Increasing ε broadens the transition region and produces a smoother hyperbolic tangent profile, as shown in the one-dimensional cross sections [Fig. 6(c)]. Importantly, variations in ε modify only the interface width and do not significantly affect the bulk geometry provided that K is sufficiently small.

In all simulations, we select $K = 0.008$, which we find to provide an optimal compromise between geometric fidelity and removal of discretization artifacts. The resulting diffuse domain is fully differentiable and suitable for subsequent adjoint-based optimization.

APPENDIX B: OPTIMIZATION TRADE-OFFS

When differentiating through time-discretized PDE solvers, the computational cost and memory requirements depend strongly on the choice of AD mode. In particular, reverse-mode AD through long time integrations requires storing intermediate states of the forward trajectory, leading

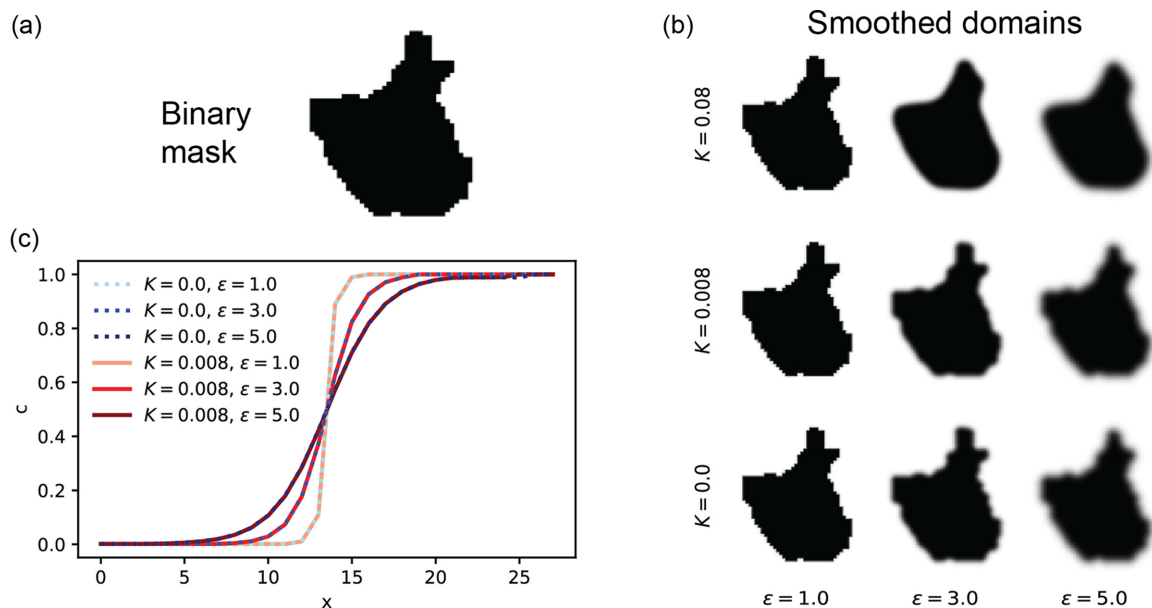


FIG. 6. Diffuse-interface smoothing of a binary mask. (a) Original segmented domain represented as a pixelated binary mask. (b) Domain phase-field profiles obtained from Eq. (11) for different curvature weights K and interface widths ε . Decreasing K suppresses curvature-driven smoothing and preserves geometric features, while increasing ε broadens the diffuse interface. (c) One-dimensional cross sections of the phase field demonstrating the hyperbolic tangent interface profile and its dependence on ε and K .




FIG. 7. Trade-offs between automatic differentiation modes and optimization strategies. (Left) Gradient computation time as a function of the number of parameters. Forward-mode AD scales approximately linearly with parameter dimension, whereas reverse-mode AD remains nearly constant for scalar-valued objectives. (Right) Optimization convergence for different combinations of optimizer and parametrization. Solid lines correspond to neural-network parametrizations and dashed lines to polynomial parametrizations. Blue curves denote LM optimization and red curves denote BFGS. The results illustrate how differentiation mode and parameter dimension influence both computational cost and convergence behavior.

to memory usage that scales linearly with the number of time steps. For high-resolution simulations or long time horizons, this can become prohibitive.

To mitigate this issue, we employ the optimal online recursive checkpointing scheme implemented in *diffrax* [78]. This approach hierarchically partitions the time domain and selectively recomputes intermediate states during the backward pass, reducing memory consumption from $O(T)$ to $O(\log T)$, where T denotes the number of time steps. The reduction in memory comes at the cost of a modest increase in computational effort due to recomputation of segments of the forward trajectory.

The choice between forward- and reverse-mode AD depends on the parameter dimension and optimization strategy. Forward-mode AD scales linearly with the number of parameters, making it efficient for low-dimensional parametrizations. Reverse-mode AD, in contrast, has a computational cost largely independent of parameter dimension for scalar objectives, and is therefore advantageous for high-dimensional representations such as neural networks. For the examples studied in this work, the number of optimized parameters is typically on the order of 100, which lies near the crossover between the forward- and reverse-mode timing curves for the architectures used here (NVIDIA L40S GPU).

These trade-offs are illustrated in Fig. 7. The left panel shows the scaling of gradient computation time with respect to the number of parameters. Forward-mode exhibits approximately linear growth, whereas reverse-mode remains nearly constant. The right panel compares convergence behavior for different combinations of optimizer and parametrization. While reverse mode enables efficient training of high-dimensional neural-network controls, simpler parametrizations may benefit from forward-mode differentiation and enable Gauss-Newton Hessian approximations and thus faster optimization convergence. Together, these considerations guide the selection of differentiation mode and optimization algorithm in our simulations.

APPENDIX C: LEGENDRE POLYNOMIAL REGULARIZATION

Constitutive functions are represented using a truncated Legendre expansion on $c \in [0, 1]$,

$$f(c) = \sum_{n=0}^N \theta_n P_n(2c - 1), \quad (\text{C1})$$

where $P_n(x)$ are the standard Legendre polynomials defined on $[-1, 1]$.

Legendre polynomials are eigenfunctions of the Sturm-Liouville operator

$$L[y] = -\frac{d}{dx} \left[(1-x^2) \frac{dy}{dx} \right], \quad (\text{C2})$$

with eigenvalues $n(n+1)$. Because these eigenvalues grow quadratically with n , higher-order polynomials correspond to increasingly oscillatory (higher-frequency) components of the function.

Motivated by this spectral structure, we regularize the expansion coefficients using the quadratic penalty

$$\theta^\top W \theta, \quad W_{nn} = n(n+1). \quad (\text{C3})$$

This weighting increasingly penalizes high-order modes, thereby suppressing rapid oscillations in the reconstructed constitutive function. The same weighting also admits a direct interpretation in terms of regularization of the derivative of the function. For standard Legendre polynomials,

$$\int_{-1}^1 (P'_n(x))^2 dx = n(n+1). \quad (\text{C4})$$

We therefore choose the regularization term

$$\theta^\top W \theta \quad (\text{C5})$$

so that each coefficient is penalized in proportion to the integral of the squared derivative of the reconstructed function. Together, these two perspectives show that the Legendre

weighting simultaneously suppresses high-frequency spectral components and enforces smoothness by penalizing large derivatives. For $N = 10$, the resulting diagonal weights are

$$\{0, 2, 6, 12, 20, 30, 42, 56, 72, 90, 110\}. \quad (\text{C6})$$

The constant mode ($n = 0$) remains unpenalized, reflecting the fact that uniform shifts do not contribute to the derivative.

APPENDIX D: MULTIPLE SHOOTING OPTIMIZATION

In a conventional single shooting formulation, the PDE is simulated over the entire observation window $[t_0, t_f]$ starting from the initial condition provided by the first snapshot. The loss is then evaluated at specific observation times and gradients are computed by differentiating through the full time integration.

For nonlinear PDEs, however, long time integrations can exhibit sensitive dependence on the initial condition. In the presence of measurement noise, small perturbations in the initial snapshot may propagate and amplify over time, leading to unstable gradients and reduced robustness during optimization. To mitigate these effects, we adopt a multiple shooting strategy. The full time window is decomposed into M nonoverlapping subintervals $[t_i, t_{i+1}]$, and independent forward simulations are performed on each interval using the observed data at time t_i as the initial condition. The total loss is then defined as a sum over intervals,

$$\mathcal{L}(\theta) = \sum_{i=0}^{M-1} \|\Psi_{\theta}^{(i)}(t_{i+1}) - \Psi_{\text{data}}(t_{i+1})\|^2 + \lambda \theta^{\top} W \theta, \quad (\text{D1})$$

where $\Psi_{\theta}^{(i)}$ denotes the solution of the PDE on the interval $[t_i, t_{i+1}]$ initialized from the observed state at t_i .

This formulation provides two principal advantages. First, the effect of noise in any individual snapshot is confined to a single short time interval, preventing error accumulation across the entire trajectory. We have noticed that this significantly improves robustness when the PDE dynamics are sensitive to perturbations in the initial condition and when observation noise is large. Second, the subinterval simulations are independent and can be vectorized or executed in parallel. In our JAX implementation, the solves over different time intervals are batched and compiled together using XLA, yielding substantial speedups compared to a single long simulation. We found that multiple shooting improves convergence stability and reduces wall-clock time relative to single shooting, particularly for long time horizons and noisy data.

APPENDIX E: WETTING ANGLE PARAMETRIZATION

We parametrize the time-dependent boundary contact angle $\varphi(t)$ using a piecewise linear interpolation scheme with N learnable parameters. This approach provides a flexible yet computationally tractable representation of the control function while ensuring physical constraints are satisfied. Let $\{\hat{\varphi}_i\}_{i=1}^N$ denote the learnable parameters. To ensure the contact angle remains within the physically admissible range $[0, \pi/2]$, we apply a smooth transformation:

$$\varphi_i = \frac{\pi}{4}(1 + \tanh(\hat{\varphi}_i)), \quad i = 1, \dots, N. \quad (\text{E1})$$

The hyperbolic tangent transformation maps the unconstrained parameters $\hat{\varphi}_i \in \mathbb{R}$ to the bounded interval $(0, \pi/2)$, which facilitates unconstrained optimization while guaranteeing physically meaningful solutions.

To satisfy the boundary conditions at the initial and final times, we fix the contact angle at $\varphi = \pi/2$ (corresponding to a neutral wetting condition) at $t = 0$ and $t = T$. The full set of control points is thus

$$\boldsymbol{\varphi} = \left(\frac{\pi}{2}, \tilde{\varphi}_1, \tilde{\varphi}_2, \dots, \tilde{\varphi}_N, \frac{\pi}{2} \right)^{\top}, \quad (\text{E2})$$

defined at uniformly spaced time points $t_j = jT/(N+1)$ for $j = 0, 1, \dots, N+1$.

The continuous control function $\varphi(t)$ is obtained via linear interpolation between these control points:

$$\varphi(t) = \varphi_j + \frac{t - t_j}{t_{j+1} - t_j}(\varphi_{j+1} - \varphi_j), \quad t \in [t_j, t_{j+1}]. \quad (\text{E3})$$

To promote smooth control trajectories and avoid rapid oscillations that may be difficult to realize experimentally, we include a regularization term in the objective function that penalizes the discrete approximation of the time derivative. Specifically, we define the regularization functional as

$$\mathcal{R}[\varphi] = \sum_{j=0}^N (\varphi_{j+1} - \varphi_j)^2, \quad (\text{E4})$$

which corresponds to a finite-difference approximation of $\int_0^T |\dot{\varphi}(t)|^2 dt$.

The total objective function for the optimal control problem is then

$$\mathcal{L}_{\text{tot}}[\varphi] = \mathcal{L}[\varphi] + \lambda \mathcal{R}[\varphi], \quad (\text{E5})$$

where $\mathcal{L}[\varphi]$ denotes the data term [Eqs. (17) and (18)] and $\lambda > 0$ is the regularization parameter that controls the trade-off between tracking accuracy and control smoothness.

APPENDIX F: CONTROL OF PHASE SEPARATION

We employ the following mathematical form of the Allen-Cahn model:

$$\frac{\partial c}{\partial t} = k(x, y; \mathbf{p})R(c, \mu, \Delta\phi),$$

$$I = \iint_{\psi} R(c, \mu, \Delta\phi) dA,$$

$$\mu(c) = \log\left(\frac{c}{1-c}\right) + \Omega(1-2c) - \frac{\kappa}{\psi} \nabla \cdot (\psi \nabla c),$$

$$\tilde{\eta} = \mu + \Delta\phi - \Delta\phi^{\ominus},$$

$$R(c, \mu, \Delta\phi) = \sqrt{c}(1-c)(e^{-\alpha\tilde{\eta}} - e^{(1-\alpha)\tilde{\eta}}), \quad (\text{F1})$$

where Ω is the enthalpy of mixing. We employ $\Omega = 4.5$, $\kappa = 0.0053$, and $\Delta\phi^{\ominus} = 3, 42$ V for this study. We will now outline the procedure for generating the spatial prefactor $k(x, y)$.

By converting the mask to a discrete graph \mathcal{G} , we can construct a natural spatial basis for an arbitrary particle geometry by finding the first $N_p + 1$ eigenvectors $\{\mathbf{B}_k\}_{k=1}^{N_p+1}$ of the graph

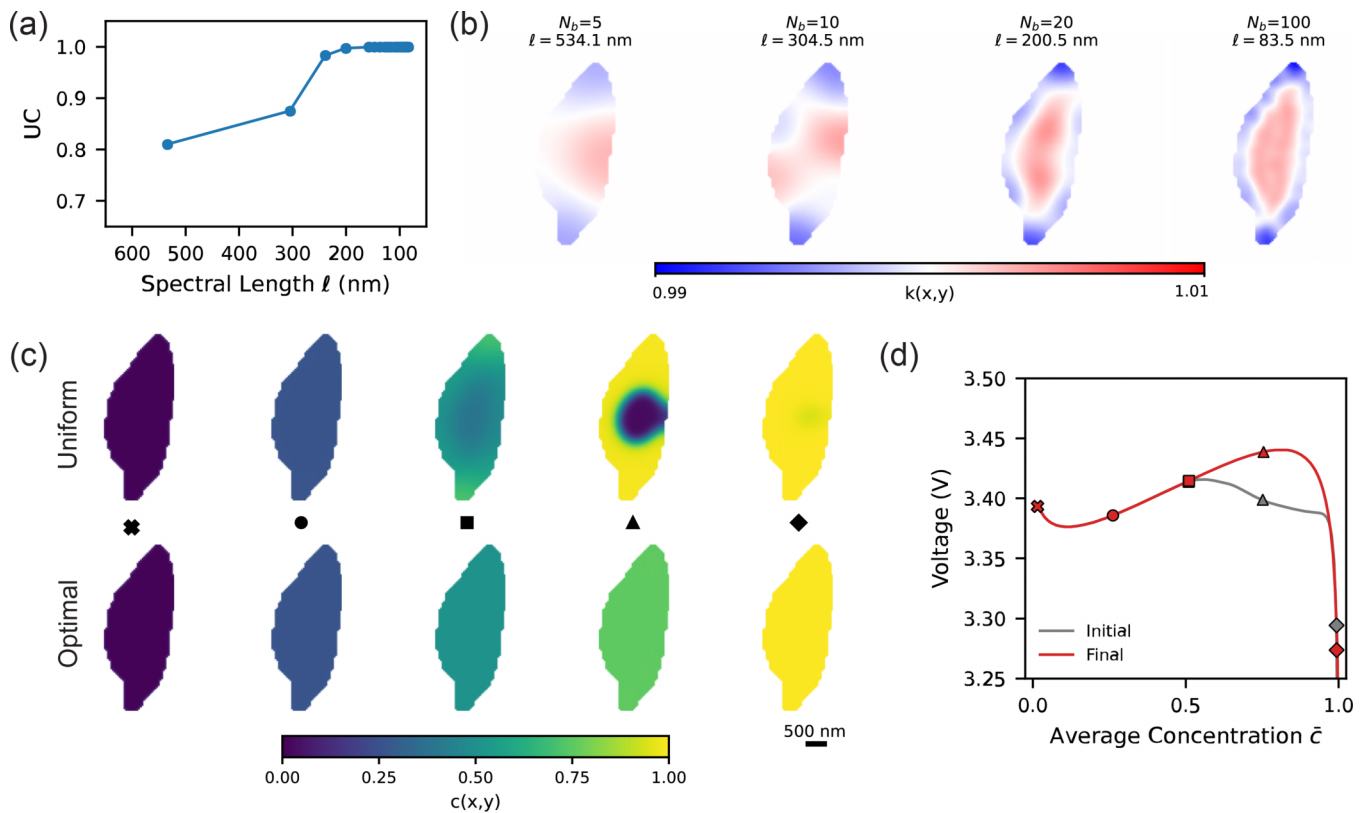


FIG. 8. Suppression of phase separation in an alternative experimental particle geometry. (a) Consistent with the primary geometry, the maximum spectral length scale of the optimized coating determines the limit of control over concentration uniformity. (b) The optimal coating distributions exhibit similar trends to the primary particle; as the spectral length scale decreases, finer spatial patterns emerge, characterized by enhanced kinetics in the center and reduced kinetics along the boundary. (c) Simulated lithiation dynamics comparing uniform and optimized coatings. The optimized coating successfully maintains internal homogeneity. In contrast, the uniform coating undergoes pronounced phase separation, particularly near flatter interior regions, demonstrating that more asymmetric geometries exacerbate asymmetric phase separation. (d) The optimized coating yields a voltage profile characteristic of solid-solution-like dynamics.

Laplacian, where $\mathbf{B} \in \mathbb{R}^{N \times N_p}$. Furthermore, we ensure that for a given set of basis weights, the average value of $k(x, y)$ is the same by defining a matrix operator $\mathbf{M} \in \mathbb{R}^{N \times N}$ (where N is the total spatial nodes) such that

$$\mathbf{M} = \begin{bmatrix} 1 & 1 & \dots & 1 \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix},$$

$$\mathbf{MB} = \begin{bmatrix} \sum_{i=1}^N \mathbf{B}_{1,i} & \sum_{i=1}^N \mathbf{B}_{2,i} & \dots & \sum_{i=1}^N \mathbf{B}_{N_p+1,i} \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{bmatrix}$$

and then parametrizing the null space of \mathbf{MB} denoted by \mathbf{N} to ensure the average is always 0; then we define the discretized form of $k(x, y)$ as follows:

$$\mathbf{k}(\mathbf{p}) = \exp \odot (\mathbf{BNp}), \quad (\text{F2})$$

where $\mathbf{B} \in \mathbb{R}^{N \times (N_p+1)}$, $\mathbf{N} \in \mathbb{R}^{(N_p+1) \times (N_p+1)}$, $\mathbf{p} \in \mathbb{R}^{N_p}$, and the elementwise exponential is applied to ensure positivity.

APPENDIX G: ADDITIONAL PHASE SEPARATION RESULTS

See Fig. 8 for additional results on controlling phase separation.

APPENDIX H: PARAMETRIZATION OF TRAPPING POTENTIALS

We parametrize the time-dependent trapping potential $V(\mathbf{r}, t)$ for the BEC control problem using a flexible functional form that can represent multiple local minima with smooth transitions between configurations. The potential consists of a soft-minimum combination of three anisotropic harmonic wells supplemented by a Gaussian term. The trapping potential is defined as

$$V(\mathbf{r}, t) = V_{\text{wells}}(\mathbf{r}, t) + V_{\text{Gauss}}(\mathbf{r}, t), \quad (\text{H1})$$

where $\mathbf{r} = (x, y)$ denotes the spatial coordinates. This was motivated by a similar form of the potential as in Ref. [8].

The multiwell component employs a log-sum-exp construction to create a smooth approximation to the minimum

of three quadratic potentials:

$$V_{\text{wells}}(\mathbf{r}, t) = -\tau \ln \left[\sum_{k=1}^3 \exp \left(-\frac{Q_k(\mathbf{r}, t)}{\tau^2} \right) \right], \quad (\text{H2})$$

where each quadratic form is given by

$$Q_k(\mathbf{r}, t) = \omega_{x,k}^2(t)[x - x_k(t)]^2 + \omega_{y,k}^2(t)[y - y_k(t)]^2 + b_k(t). \quad (\text{H3})$$

Here, $\omega_{x,k}(t)$ and $\omega_{y,k}(t)$ are the trap frequencies along the x and y directions for the k th well, $(x_k(t), y_k(t))$ is the center position of the k th well, and $b_k(t)$ is an energy bias that controls the relative depth of each well. The parameter $\tau > 0$ controls the smoothness of the soft-minimum operation; in the limit $\tau \rightarrow 0$, Eq. (H2) reduces to $\min_k Q_k(\mathbf{r}, t)$.

The Gaussian component provides additional flexibility for creating the barriers required to form the toroidal wavefunction density pattern:

$$V_{\text{Gauss}}(\mathbf{r}, t) = A(t) \exp \left[-\frac{(x - x_4(t))^2}{2\sigma_x^2(t)} - \frac{(y - y_4(t))^2}{2\sigma_y^2(t)} \right], \quad (\text{H4})$$

where $A(t)$ is the amplitude, $(x_4(t), y_4(t))$ is the center position, and $\sigma_x(t), \sigma_y(t)$ are the widths along each axis.

Time dependence and control parametrization

Each of the 20 time-dependent parameters—comprising the trap frequencies $\{\omega_{x,k}, \omega_{y,k}\}_{k=1}^3$, well positions $\{x_k, y_k\}_{k=1}^3$, biases $\{b_k\}_{k=1}^3$, and Gaussian parameters $\{A, \sigma_x, \sigma_y, x_4, y_4\}$ —is represented using piecewise linear interpolation with fixed boundary conditions. For a generic parameter $\theta^i(t)$, we specify fixed initial and final values θ_0^i and θ_f^i by optimizing for a potential whose ground state matches the desired initial and final density patterns (Fig. 5). Each parameter has N learnable intermediate values $\{\theta_j^i\}_{j=1}^N$. The complete set of control points is

$$\boldsymbol{\theta} = (\theta_0, \theta_1, \theta_2, \dots, \theta_N, \theta_f)^T, \quad (\text{H5})$$

defined at uniformly spaced times $t_j = jT/(N+1)$ for $j = 0, 1, \dots, N+1$. The continuous function $\theta^i(t)$ is obtained via linear interpolation:

$$\theta(t) = \xi_j + \frac{t - t_j}{t_{j+1} - t_j}(\theta_{j+1} - \theta_j), \quad t \in [t_j, t_{j+1}]. \quad (\text{H6})$$

This parametrization ensures that the initial and final trap configurations match the desired boundary conditions while allowing the optimization to determine the optimal intermediate trajectory for each control parameter.

-
- [1] H. Tian, J. Li, and M. Z. Bazant, Multiphase polarization in ion-intercalation nanofilms: General theory including various surface effects and memory applications, *Adv. Funct. Mater.* **33**, 2213621 (2023).
- [2] J. Lim, Y. Li, D. H. Alsem, H. So, S. C. Lee, P. Bai, D. A. Cogswell, X. Liu, N. Jin, Y.-s. Yu, *et al.*, Origin and hysteresis of lithium compositional spatiodynamics within battery primary particles, *Science* **353**, 566 (2016).
- [3] B. Wang, L. Zhang, T. Dai, Z. Qin, H. Lu, L. Zhang, and F. Zhou, Liquid–liquid phase separation in human health and diseases, *Signal Transduct. Target. Ther.* **6**, 290 (2021).
- [4] H. Zhao, A. R. Strom, J. M. Eeftens, M. Haataja, A. Kosmrlj, and C. P. Brangwynne, Condensate driven chromatin organization via elastocapillary interactions, *Biophys. J.* **125**, 4 (2026).
- [5] A. R. Strom, Y. Kim, H. Zhao, Y.-C. Chang, N. D. Orlovsky, A. Košmrlj, C. Storm, and C. P. Brangwynne, Condensate interfacial forces reposition DNA loci and probe chromatin viscoelasticity, *Cell* **187**, 5282 (2024).
- [6] Z. Tan, S. Chen, X. Peng, L. Zhang, and C. Gao, Polyamide membranes with nanoscale Turing structures for water purification, *Science* **360**, 518 (2018).
- [7] S. Ouk Kim, H. H. Solak, M. P. Stoykovich, N. J. Ferrier, J. J. De Pablo, and P. F. Nealey, Epitaxial self-assembly of block copolymers on lithographically defined nanopatterned substrates, *Nature (London)* **424**, 411 (2003).
- [8] J.-F. Mennemann, D. Matthes, R.-M. Weishäupl, and T. Langen, Optimal control of Bose–Einstein condensates in three dimensions, *New J. Phys.* **17**, 113027 (2015).
- [9] C. Ryu, M. F. Andersen, P. Clade, V. Natarajan, K. Helmerson, and W. D. Phillips, Observation of persistent flow of a Bose-Einstein condensate in a toroidal trap, *Phys. Rev. Lett.* **99**, 260401 (2007).
- [10] C. Ryu, P. W. Blackburn, A. A. Blinova, and M. G. Boshier, Experimental realization of Josephson junctions for an atom SQUID, *Phys. Rev. Lett.* **111**, 205301 (2013).
- [11] S. Eckel, J. G. Lee, F. Jendrzejewski, N. Murray, C. W. Clark, C. J. Lobb, W. D. Phillips, M. Edwards, and G. K. Campbell, Hysteresis in a quantized superfluid ‘atomtronic’ circuit, *Nature (London)* **506**, 200 (2014).
- [12] J. Estève, C. Gross, A. Weller, S. Giovanazzi, and M. K. Oberthaler, Squeezing and entanglement in a Bose–Einstein condensate, *Nature (London)* **455**, 1216 (2008).
- [13] L. Pezze, A. Smerzi, M. K. Oberthaler, R. Schmied, and P. Treutlein, Quantum metrology with nonclassical states of atomic ensembles, *Rev. Mod. Phys.* **90**, 035005 (2018).
- [14] J. Hertkorn, J.-N. Schmidt, M. Guo, F. Böttcher, K. S. H. Ng, S. D. Graham, P. Uerlings, T. Langen, M. Zwierlein, and T. Pfau, Pattern formation in quantum ferrofluids: From supersolids to superglasses, *Phys. Rev. Res.* **3**, 033125 (2021).
- [15] D. A. Cogswell and M. Z. Bazant, Coherency strain and the kinetics of phase separation in LiFePO₄ nanoparticles, *ACS Nano* **6**, 2215 (2012).
- [16] D. P. Finegan, A. Quinn, D. S. Wragg, A. M. Colclasure, X. Lu, C. Tan, T. M. Heenan, R. Jervis, D. J. Brett, S. Das, *et al.*, Spatial dynamics of lithiation and lithium plating during

- high-rate operation of graphite electrodes, *Energy Environ. Sci.* **13**, 2570 (2020).
- [17] M. Z. Bazant, Theory of chemical kinetics and charge transfer based on nonequilibrium thermodynamics, *Acc. Chem. Res.* **46**, 1144 (2013).
- [18] S. Boeynaems, A. S. Holehouse, V. Weinhardt, D. Kovacs, J. Van Lindt, C. Larabell, L. Van Den Bosch, R. Das, P. S. Tompa, R. V. Pappu, *et al.*, Spontaneous driving forces give rise to protein- RNA condensates with coexisting phases and complex material properties, *Proc. Natl. Acad. Sci. USA* **116**, 7889 (2019).
- [19] B. R. Sabari, A. Dall’Agnese, A. Boija, I. A. Klein, E. L. Coffey, K. Shrinivas, B. J. Abraham, N. M. Hannett, A. V. Zamudio, J. C. Manteiga, *et al.*, Coactivator condensation at super-enhancers links phase separation and gene control, *Science* **361**, eaar3958 (2018).
- [20] X. Lu, R. E. Owen, W. Du, Z. Zhang, A. Bertei, R. Soni, X. Zhang, F. Iacoviello, D. Li, A. Llewellyn, *et al.*, Unravelling electro-chemo-mechanical processes in graphite/silicon composites for designing nanoporous and microstructured battery electrodes, *Nat. Nanotechnol.* **20**, 1656 (2025).
- [21] Y. Shin, M. W. Zwierlein, C. H. Schunck, A. Schirotzek, and W. Ketterle, Observation of phase separation in a strongly interacting imbalanced Fermi gas, *Phys. Rev. Lett.* **97**, 030401 (2006).
- [22] Z. Lun, A. J. Merryweather, A. Mahadevegowda, S. S. Pandurangi, C. Xu, S. Fairclough, V. S. Deshpande, N. A. Fleck, C. Ducati, C. Schnedermann, *et al.*, Operando single-particle imaging reveals that asymmetric ion flux contributes to capacity degradation in aged Ni-rich layered cathodes, *Energy Environ. Sci.* **18**, 4097 (2025).
- [23] C. Xu, A. J. Merryweather, S. S. Pandurangi, Z. Lun, D. S. Hall, V. S. Deshpande, N. A. Fleck, C. Schnedermann, A. Rao, and C. P. Grey, Operando visualization of kinetically induced lithium heterogeneities in single-particle layered Ni-rich cathodes, *Joule* **6**, 2535 (2022).
- [24] M. A. Ziatdinov, M. Y. Yaman, Y. Liu, D. Ginger, and S. V. Kalinin, Semi-supervised learning of images with strong rotational disorder: Assembling nanoparticle libraries, *Digit. Discov.* **3**, 1213 (2024).
- [25] K. Barakati, H. Yuan, A. Goyal, and S. V. Kalinin, Physics-based reward driven image analysis in microscopy, *Digit. Discov.* **3**, 2061 (2024).
- [26] S. Inouye, Manipulating Bose-Einstein condensates with laser light, Ph.D. thesis, Massachusetts Institute of Technology, 2001.
- [27] Y. Liu, J. Liu, J. Wang, M. N. Bannis, B. Xiao, A. Lushington, W. Xiao, R. Li, T.-K. Sham, G. Liang, *et al.*, Formation of size-dependent and conductive phase on lithium iron phosphate during carbon coating, *Nat. Commun.* **9**, 929 (2018).
- [28] H. Zhao, H. D. Deng, A. E. Cohen, J. Lim, Y. Li, D. Fragedakis, B. Jiang, B. D. Storey, W. C. Chueh, R. D. Braatz, *et al.*, Learning heterogeneous reaction kinetics from X-ray videos pixel by pixel, *Nature (London)* **621**, 289 (2023).
- [29] S. Thutupalli, D. Geyer, R. Singh, R. Adhikari, and H. A. Stone, Flow-induced phase separation of active particles is controlled by boundary conditions, *Proc. Natl. Acad. Sci. USA* **115**, 5403 (2018).
- [30] A. B. Li, L. Miroshnik, B. D. Rummel, G. Balakrishnan, S. M. Han, and T. Sinno, A unified theory of free energy functionals and applications to diffusion, *Proc. Natl. Acad. Sci. USA* **119**, e2203399119 (2022).
- [31] J. W. Cahn and J. E. Hilliard, Free energy of a nonuniform system. I. Interfacial free energy, *J. Chem. Phys.* **28**, 258 (1958).
- [32] J. W. Cahn and S. M. Allen, A microscopic theory for domain wall motion and its experimental verification in Fe-Al alloy domain growth kinetics, *J. Phys. Colloq.* **38**, C7-51 (1977).
- [33] C. V. Di Leo, E. Rejovitzky, and L. Anand, A Cahn–Hilliard-type phase-field theory for species diffusion coupled with large elastic deformations: Application to phase-separating Li-ion electrode materials, *J. Mech. Phys. Solids* **70**, 1 (2014).
- [34] P. K. Kristensen, C. F. Niordson, and E. Martínez-Pañeda, An assessment of phase field fracture: Crack initiation and growth, *Philos. Trans. R. Soc. A* **379**, 20210021 (2021).
- [35] R. Choksi, M. A. Peletier, and J. Williams, On the phase diagram for microphase separation of diblock copolymers: An approach via a nonlocal Cahn–Hilliard functional, *SIAM J. Appl. Math.* **69**, 1712 (2009).
- [36] H. Zhao, A. Košmrlj, and S. S. Datta, Chemotactic motility-induced phase separation, *Phys. Rev. Lett.* **131**, 118301 (2023).
- [37] W. Bao and Y. Cai, Mathematical theory and numerical methods for Bose-Einstein condensation, *Kinet. Relat. Models* **6**, 1 (2013).
- [38] R. C. Ihuaenyi, H. Zhao, R. Fang, R. Bai, M. Z. Bazant, and J. Zhu, Learning reaction-diffusion kinetics from mechanical information, [arXiv:2508.17523](https://arxiv.org/abs/2508.17523).
- [39] M. Raissi, P. Perdikaris, and G. E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* **378**, 686 (2019).
- [40] K. Azizzadenesheli, N. Kovachki, Z. Li, M. Liu-Schiaffini, J. Kossaiifi, and A. Anandkumar, Neural operators for accelerating scientific simulations and design, *Nat. Rev. Phys.* **6**, 320 (2024).
- [41] M. Z. Bazant, Thermodynamic stability of driven open systems and control of phase separation by electro-autocatalysis, *Faraday Discuss.* **199**, 423 (2017).
- [42] L. P. Pitaevskii, Vortex lines in an imperfect Bose gas, *J. Exptl. Theoret. Phys. (U.S.S.R.)* **40**, 646 (1961) [*Sov. Phys. JETP* **13**, 451 (1961)].
- [43] E. P. Gross, Structure of a quantized vortex in boson systems, *Nuovo Cimento* **20**, 454 (1961).
- [44] V. Heinonen, K. J. Burns, and J. Dunkel, Quantum hydrodynamics for supersolid crystals and quasicrystals, *Phys. Rev. A* **99**, 063621 (2019).
- [45] S. M. Allen and J. W. Cahn, A microscopic theory for antiphase boundary motion and its application to antiphase domain coarsening, *Acta Metall.* **27**, 1085 (1979).
- [46] S. Daubner, P. W. Hoffrogge, M. Minar, and B. Nestler, Triple junction benchmark for multiphase-field and multi-order parameter models, *Comput. Mater. Sci.* **219**, 111995 (2023).
- [47] T. Takaki and J. Kato, Phase-field topology optimization model that removes the curvature effects, *Mech. Eng. J.* **4**, 16-00462 (2017).
- [48] S. Daubner, A. E. Cohen, B. Dörich, and S. J. Cooper, voxels: A differentiable physics framework for voxel-based microstructure simulations, *JOSS* **11**, 9733 (2025).
- [49] Y. Sun and C. Beckermann, Sharp interface tracking using the phase-field equation, *J. Comput. Phys.* **220**, 626 (2007).
- [50] H.-C. Yu, H.-Y. Chen, and K. Thornton, Extended smoothed boundary method for solving partial differential equations with general boundary conditions on complex boundaries, *Modell. Simul. Mater. Sci. Eng.* **20**, 075008 (2012).

- [51] S. Daubner, M. Weichel, M. Reder, D. Schneider, Q. Huang, A. E. Cohen, M. Z. Bazant, and B. Nestler, Simulation of intercalation and phase transitions in nano-porous, polycrystalline agglomerates, *npj Comput. Mater.* **11**, 211 (2025).
- [52] X. Li, J. Lowengrub, A. Ratz, and A. Voigt, Solving PDEs in complex geometries, *Commun. Math. Sci.* **7**, 81 (2009).
- [53] A. Abdulle and A. A. Medovikov, Second order Chebyshev methods based on orthogonal polynomials, *Numer. Math.* **90**, 1 (2001).
- [54] L. Q. Chen and J. Shen, Applications of semi-implicit Fourier-spectral method to phase field equations, *Comput. Phys. Commun.* **108**, 147 (1998).
- [55] J. Rader, T. Lyons, and P. Kidger, Optimistix: Modular optimisation in JAX and Equinox, [arXiv:2402.09983](https://arxiv.org/abs/2402.09983).
- [56] S. Agrawal and P. Bai, Dynamic interplay between phase transformation instabilities and reaction heterogeneities in particulate intercalation electrodes, *Cell Rep. Phys. Sci.* **3**, 100854 (2022).
- [57] A. Malik, K. Snyder, M. Liu, and H.-C. Yu, Phase-field electrochemical simulations of reconstructed graphite electrodes, *J. Energy Storage* **77**, 109937 (2024).
- [58] D. Qu, A. Malik, and H.-C. Yu, Physics-based simulation of electrochemical impedance spectroscopy of complex electrode microstructures using smoothed boundary method, *Electrochim. Acta* **432**, 141141 (2022).
- [59] H. Zhao, R. D. Braatz, and M. Z. Bazant, Image inversion and uncertainty quantification for constitutive laws of pattern formation, *J. Comput. Phys.* **436**, 110279 (2021).
- [60] P. Teng, D. Tian, H. Fu, and S. Wang, Recent progress of electrowetting for droplet manipulation: From wetting to super-wetting systems, *Mater. Chem. Front.* **4**, 140 (2020).
- [61] F. Mugele and J.-C. Baret, Electrowetting: From basics to applications, *J. Phys.: Condens. Matter* **17**, R705 (2005).
- [62] T. Ohzuku, Y. Iwakoshi, and K. Sawai, Formation of lithium-graphite intercalation compounds in nonaqueous electrolytes and their application as a negative electrode for a lithium ion (shuttlecock) cell, *J. Electrochem. Soc.* **140**, 2490 (1993).
- [63] A. K. Padhi, K. S. Nanjundaswamy, and J. B. Goodenough, Phospho-olivines as positive-electrode materials for rechargeable lithium batteries, *J. Electrochem. Soc.* **144**, 1188 (1997).
- [64] M. Tang, W. C. Carter, and Y.-M. Chiang, Electrochemically driven phase transitions in insertion electrodes for lithium-ion batteries: Examples in lithium metal phosphate olivines, *Annu. Rev. Mater. Res.* **40**, 501 (2010).
- [65] T. Ohzuku, A. Ueda, and N. Yamamoto, Zero-strain insertion material of $\text{Li}[\text{Li}_{1/3}\text{Ti}_{5/3}]\text{O}_4$ for rechargeable lithium cells, *J. Electrochem. Soc.* **142**, 1431 (1995).
- [66] K. Ariyoshi, Y. Iwakoshi, N. Nakayama, and T. Ohzuku, Topotactic two-phase reactions of $\text{Li}[\text{Ni}_{1/2}\text{Mn}_{3/2}]\text{O}_4$ ($P4_332$) in nonaqueous lithium cells, *J. Electrochem. Soc.* **151**, A296 (2004).
- [67] D. Zhuang and M. Z. Bazant, Scaling analysis of mosaic phase separation in Li-ion batteries, *Phys. Rev. E* **110**, 064142 (2024).
- [68] D. Zhuang, M. L. Li, V. N. Lam, R. D. Braatz, W. C. Chueh, and M. Z. Bazant, Physics-informed design of hybrid pulse power characterization tests for rechargeable batteries, *J. Electrochem. Soc.* **171**, 050510 (2024).
- [69] D. Zhuang, Degradation mechanisms and applications in ion intercalation materials, Ph.D. thesis, Massachusetts Institute of Technology, 2024.
- [70] M. Z. Bazant, Unified quantum theory of electrochemical kinetics by coupled ion–electron transfer, *Faraday Discuss.* **246**, 60 (2023).
- [71] A. R. Iarchuk, V. A. Nikitina, E. A. Karpushkin, V. G. Sergeev, E. V. Antipov, K. J. Stevenson, and A. M. Abakumov, Influence of carbon coating on intercalation kinetics and transport properties of LiFePO_4 , *ChemElectroChem* **6**, 5090 (2019).
- [72] P. Ombrini, S. Pathak, D. Ntagkras, S. K. Pal, P. Karanth, F. M. Mulder, M. Wagemaker, M. Z. Bazant, and A. Vasileiadis, Modeling single-crystal electrodes as a network of primary particles, *Energy Environ. Sci.* **18**, 9468 (2025).
- [73] D. A. Cogswell and M. Z. Bazant, Size-dependent phase morphologies in LiFePO_4 battery particles, *Electrochem. Commun.* **95**, 33 (2018).
- [74] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga, Shortcuts to adiabaticity: Concepts, methods, and applications, *Rev. Mod. Phys.* **91**, 045001 (2019).
- [75] Y. Kuriatnikov, N. Würkner, K. Kumaran, T. Zhang, M. V. Ramana, A. Kugi, J. Schmiedmayer, A. Deutschmann-Olek, and M. Prüfer, Fast coherent splitting of Bose-Einstein condensates, *Phys. Rev. Res.* **7**, 043108 (2025).
- [76] U. Hohenester, P. K. Rekdal, A. Borzì, and J. Schmiedmayer, Optimal quantum control of Bose-Einstein condensates in magnetic microtraps, *Phys. Rev. A* **75**, 023602 (2007).
- [77] A. E. Cohen, S. Degnan-Morgenstern, and S. Daubner, Free-energy-driven-pattern-optimization, <https://github.com/acoh64/free-energy-pattern-optimization>.
- [78] P. Kidger, On neural differential equations, Ph.D. thesis, University of Oxford, 2021.