Nicholas M. Boffi: Research Statement

I develop algorithms for modeling, simulation, and control of complex dynamical systems in the natural sciences and engineering. My work integrates numerical analysis with machine learning to design principled methods for high-dimensional mathematical problems that are currently out of reach via the standard techniques of computational mathematics.

Research in computational mathematics has flourished through an emphasis on problems defined over a low-dimensional physical space, such as those appearing in continuum mechanics. Yet, high-dimensional mathematical problems are pervasive, and arise when the system under study must be described by degrees of freedom beyond the ambient spatial coordinates alone. Examples include many-body problems in quantum mechanics and statistical physics, whose solutions have wide applications ranging from improving the energy efficiency of batteries to the design of nanoscale devices. However, these problems cannot be addressed with the standard methods of numerical analysis, because their computational expense typically scales poorly with the system dimension.

Methods from machine learning offer a promising solution, as recent advances have demonstrated the remarkable ability of nonlinear parametric representations like neural networks to identify hidden low-dimensional structure (1), thereby grappling with high-dimensionality. There stands immense potential to bring this progress to computational mathematics (2). Yet to do so, it is necessary to develop novel methods that combine machine learning with the traditional techniques of numerical analysis, which come with powerful guarantees necessary for scientific applications.

My research philosophy centers on revisiting classically hard applied mathematical problems from the lens of modern machine learning. I develop principled methods with rigorous guarantees of convergence, accuracy, and stability that are suitable for real-world science and engineering problems. As such, my research profile spans from ***algorithmic and theoretical advances*** in machine learning and control theory to ***empirical advances*** in modeling and scientific simulation.

As a professor, I will lead an interdisciplinary research group that blends ideas from machine learning with dynamical systems and numerical computation. I plan to develop an ***integrated applied mathematics curriculum that fuses modern machine learning with numerical analysis***.

As a primary research focus, I plan to develop methods for high-dimensional mathematical problems built on generative models. In particular, I aim to design algorithms for high-dimensional partial differential equations, for forecasting and coarse-graining of dynamical systems, and for optimal control of nonequilibrium thermodynamic systems. These research thrusts build upon my earlier work in high-dimensional scientific computing and generative modeling (3–7), machine learning and control theory (8–15), scientific computing (16–23), and open source software development (24–28).

Learning flows for many-body Fokker-Planck equations

A canonical example of high-dimensionality arises in the statistical physics of interacting particle systems, where the problem dimension increases linearly with the number of particles. The many-body probability density function describes the system at microscopic granularity, and gives access to several key quantities that cannot typically be computed without its knowledge, such as the system's entropy production rate (EPR). This density solves a high-dimensional Fokker-Planck equation (FPE) (29), which is intractable when approached with the standard techniques of numerical analysis. Leveraging neural networks and flow-based algorithms from generative modeling to design mesh-free methods for the FPE that scale to realistic systems has been a central thrust of my recent research.

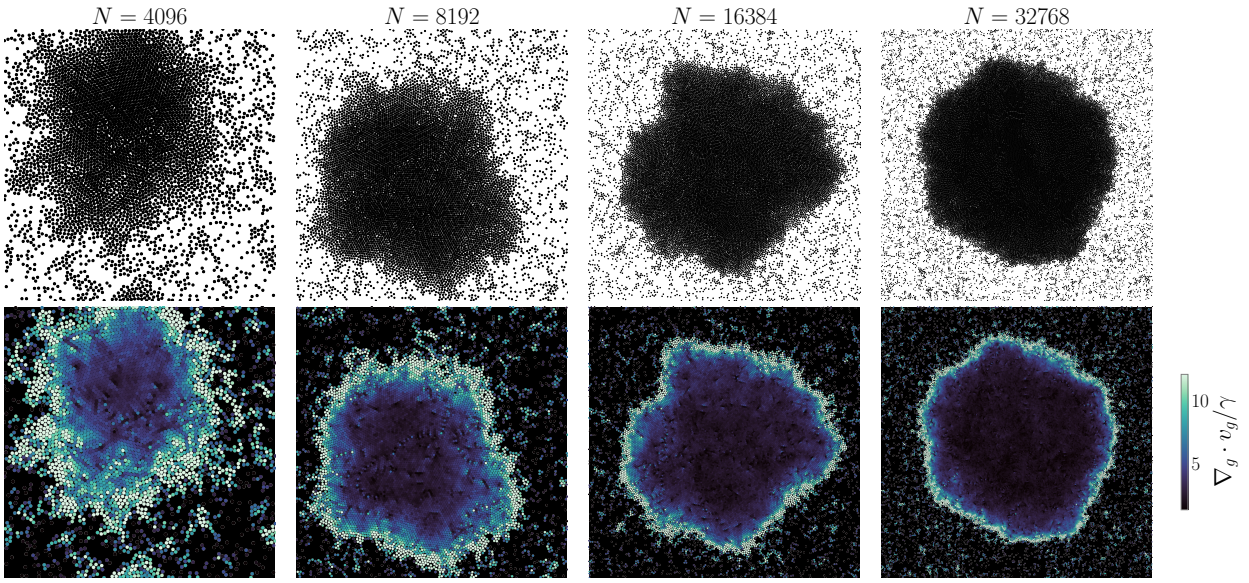

Figure 2: **MIPS.** (Top) MIPS cluster as a function of particle number N. (Bottom) Calculation of the EPR (per particle). The EPR concentrates on the interface, confirming theoretical predictions that nonequilibrium effects are most relevant at phase boundaries (37). The network was trained with $N = 4096$, but due to spatial locality, can transfer to systems of much higher dimension without additional learning. See movie, download recommended.

Stationary FPE. In collaboration with Professor Eric Vanden-Eijnden (Courant Institute), I developed a method to approximate solutions of the stationary FPE in high dimension (4). The method is applicable to active matter systems at statistical steady state, such as a system of particles undergoing motility-induced phase separation (MIPS), a nonequilibrium phenomenon defined by spontaneous coexistence of a liquid and a gas in the absence of attractive interactions (30). Understanding MIPS is important for controlling and preventing biofilm formation, which underlies many forms of bacterial contamination (31, 32). A central goal of the active matter community has therefore been to classify when and where nonequilibrium effects are at play, which can be quantified by the EPR (33–37). ***My method gives the first way to compute the EPR from its first-principles definition with quantifiable accuracy for general systems and realistic sizes.*** The approach leverages recent advances in generative modeling (38), and is built upon estimation of the *score* $\nabla \log \rho$ where ρ solves the FPE. The score leads to a description of the system known as the *probability flow* (Fig. 1), a deterministic dynamical system from which the EPR can be reconstructed. I designed a novel neural network architecture suitable for particle systems that is permutation equivariant and spatially-local, which enables scaling to systems of arbitrary size after training. Using it to study a system of particles undergoing MIPS, ***I found that the EPR concentrates on the interface between phases, confirming predictions from the***

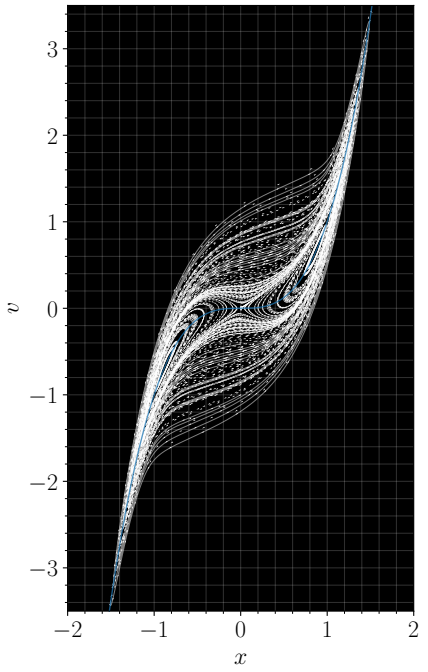

Figure 1: **Probability flow.** Phase portrait of the probability flow for a single active particle in one dimension, with position x and velocity v. This *deterministic* dynamics preserves stationarity of the system, and gives access to the EPR.

theoretical physics community that could not be verified from first-principles (37, 39). Computationally, I showed that the method scales to as many as ∼32,000 particles, corresponding to an FPE in ∼128,000 dimensions (Fig. 2).

Time-dependent FPE. To design nanoscale engines and devices, it is often necessary to drive the system out of equilibrium with an external control, which leads to a time-dependent FPE (40, 41). In collaboration with Professor Eric Vanden-Eijnden, I developed a neural network-based method to solve this equation (3), and I released an open-source implementation (25). The approach, known as *score-based transport modeling*, estimates the score along self-consistent data produced by the probability flow equation itself. I proved that the method controls the Kullback-Leibler divergence from the model to the target, a common distance metric between probability distributions. ***This guarantees that minimizing the loss systematically improves accuracy, thereby giving a principled algorithm for high-dimensional time-dependent FPEs.*** I applied the method to several driven systems, where I showed that it obtains solutions with up to three digits of relative accuracy in as high as 100 dimensions, an intractable dimensionality for standard techniques. In addition, I showed that the method gives stable predictions for the EPR over long time intervals.

Learning flows and diffusions for generative modeling

The algorithms I've developed for solving the FPE rest upon the notion of *dynamical transport of measure*, which studies how maps between probability distributions can be built from motion at the level of the samples. Rooted in the theory of optimal transport (42, 43), this picture has emerged as a central tool in the field of generative modeling, which studies how to sample from a probability distribution given a set of examples. Generative modeling for images has recently risen to fame in popular culture (44–46), and the approach holds remarkable promise for scientific computing, with applications in solving inverse problems (47), sampling molecular configurations for drug discovery (48), and probabilistic forecasting of dynamical systems (49). Yet, a principled mathematical understanding of generative models is still in its infancy, which has inhibited their widespread adoption. ***My work highlights a powerful connection between modern generative models and methods for high-dimensional Fokker-Planck equations, which I have used to enrich both fields.***

Figure 3: **Image super-resolution.** (Top) Low-resolution model input. (Middle) High-resolution model output. (Bottom) Ground truth.

Stochastic interpolants. In collaboration with Professor Eric Vanden-Eijnden and PhD student Michael Albergo, I developed a theory known as the stochastic interpolant method (6), and I released an open-source implementation (24). Stochastic interpolants build a path in the space of measures between an arbitrary "base" distribution and the target distribution by interpolating between samples. I showed that stochastic interpolants generalize *diffusion models* – the current state of the art (38) – by allowing for greater flexibility in the base and the connecting path, and that they can be sampled either deterministically or stochastically with an arbitrary, tunable level of noise. I used this observation to categorize the tradeoffs between transport equations (deterministic) and FPEs (stochastic) for sampling, where I showed that FPEs enjoy stronger theoretical guarantees and often perform better empirically. With PhD student Mark Goldstein (6), we recently extended the method to solve high-dimensional inverse problems such as inpainting and super-resolution (Fig. 3).

Dynamical systems, control, and scientific computing

My current research centers on high-dimensional FPEs and generative modeling, but I maintain a strong interest in dynamical systems, control, and scientific computing, as illustrated below.

Control and prediction. Learning a dynamical system from data can be framed in the language of function approximation, whereby an unknown dynamics must be approximated to a desired degree of accuracy over the state space (50). For many nonlinear dynamical systems, including engineering systems such as self-driving cars and robots, it is critical to learn a dynamical model *online* to counteract modeling errors, a problem known in control theory as *adaptive control*. Classical approaches build a grid in the Fourier domain and estimate parameters in each grid cell, but the size of the resulting model scales exponentially in the state space dimension, which prevents their application to nearly all modern systems. With Professor Jean-Jacques Slotine (MIT) and Dr. Stephen Tu (Google DeepMind), ***I developed a nonparametric algorithm with theoretical guarantees that depend polynomially on dimension rather than exponentially*** (10). I showed empirically that the method scales to control and prediction problems with dimensionality at least as high as 60, surpassing existing approaches by an order of magnitude. To incorporate physics into the method, I introduced a *symplectic kernel* that is adapted to the symmetries of Hamiltonian dynamics.

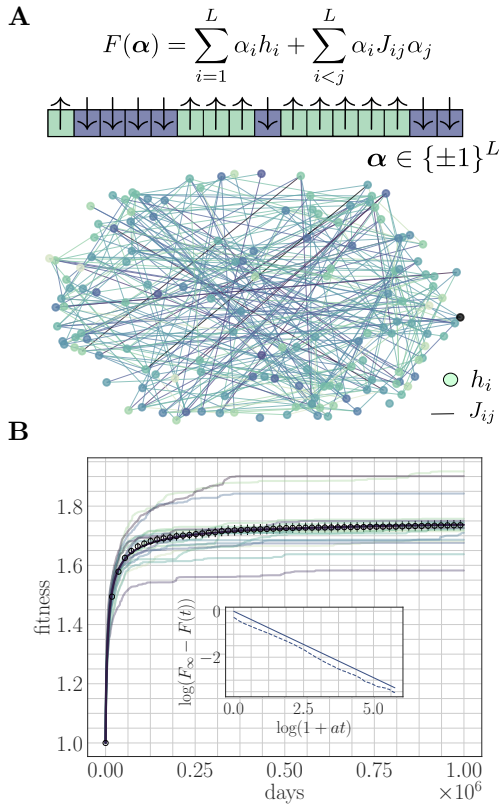

Figure 4: **Spin glass evolutionary dynamics.** (A) The genome is modeled as a sequence of sites, and the fitness F of a genotype α is given by the energy of the Sherrington-Kirkpatrick spin glass (51). (B) Random evolution under experimental conditions reproduces and explains key features of evolving populations, such as a slow power-law fitness trajectory.

Scientific computing. I have diverse interests across the sciences and engineering, ranging from mathematical biology (16), to continuum mechanics (17, 18, 21), to quantum mechanics (19, 20); my approach is unified by the use of mathematical modeling and numerical simulation to tackle questions that are challenging to answer either theoretically or experimentally. In recent work with Ariel Amir (Weizmann Institute), Chris Rycroft (University of Wisconsin), and Yipei Guo (Janelia), I developed a mathematical and computational model of microbial evolutionary dynamics based on spin glass physics (16). Understanding evolutionary processes in microbes is important for preventing the development of antibiotic resistance (52, 53), but many basic aspects are still poorly understood. For example, laboratory experiments have demonstrated the widespread prevalence of microscopic epistasis (the dependence of a mutation's effect on the genotype) and clonal interference (competition between strains) (54, 55). Yet, these aspects are neglected in most studies due to their complexity, so their effect on real-world evolving populations is unclear. To resolve this gap, I developed ***a microscopic simulation environment capable of matching laboratory conditions*** (Fig. 4), inspired by the seminal experiment of Richard Lenski (56, 57). I showed that epistasis leads to a slow power-law fitness trajectory, while clonal interference does not significantly affect the dynamics. This is an important result, because ***previous experimental studies attributed slow trajectories to both phenomena*** (56). I released the simulation as an open source package (26).

Vision and future works

A primary focus of my research group will be the development of flow-based methods for high-dimensional problems in science and engineering; I describe several intended thrusts below.

Flow-based solutions to PDEs. Learning flows is a promising way to solve high-dimensional PDEs, because their Lagrangian frame automatically adapts to low-dimensional structure. I plan to extend my flow-based algorithms for the FPE (3, 4) in several key ways, fusing my work in generative modeling (5–7) with my expertise in scientific computing (17–21). I plan to investigate time-dependent FPEs with nonlocal dynamics to design sampling algorithms that mix faster than those built on local moves. I also plan to develop methods for time-dependent quantum mechanics (58), combining my experience in quantum chemistry (19, 20, 22, 23) with my method for the time-dependent FPE (3). To extend my work on the stationary FPE, I am interested in methods that directly estimate the probability flow instead of constructing it through the score. The flow is often simpler than the score, and can be learned from experimental data when the microscopic dynamics is unknown, which is not possible with the score alone. I also plan to investigate self-consistent algorithms that do not require data from the stationary density, which could avoid slow mixing present in data generation.

Forecasting and coarse-graining. Forecasting and coarse-graining of dynamical systems are fundamental mathematical problems with wide applications across the sciences and engineering, ranging from weather prediction, to motion planning, to model order reduction. I intend to design algorithms for these tasks that combine my work in generative modeling (5–7) with my work in dynamical systems and control (8–13).

To forecast, we can sample the *conditional distribution* of a future state given the current state using a generative model. Some forecasting problems require real-time inference, such as in engineering, while others may not need such rapid inference. This problem is therefore ripe to explore the computational tradeoffs between deterministic and stochastic models, an investigation I have already begun (6). For many forecasting problems, there are physical constraints and symmetries present in the state, and I intend to design models that build these properties into their transport explicitly. One possible approach is to define a path in the space of measures via the minimum action method (59, 60).

Experimental systems modeled by stochastic differential equations are regularly *partially observed*, meaning that some internal state variables cannot be measured. A lower-dimensional equation for the microscopic density of the observed variables can often be defined by marginalization, but an unknown velocity appears in its governing equation (61, 62). I plan to develop methods for coarse-graining that estimate this velocity directly by analogy to the stochastic interpolant method.

Control of nanoscale systems. Efficient control of nanoscale systems is a central goal of modern nonequilibrium thermodynamics, as it promises hope for both the design of efficient synthetic molecular machines and the improved understanding of living systems (34, 63). This is an optimal control problem where one seeks to minimize the *dissipation of energy*, so as to maximize the amount of useful work produced by the machine (64). The dissipation can be related to the total entropy production (41), which could not be computed for generic nonequilibrium systems before the development of my score-based algorithms (3, 4). Building on my work in control theory (9–12), I plan to develop methods that minimize the total entropy production in an outer loop wrapped around its computation. To render these methods efficient, I plan to leverage recent advances in reinforcement learning, which has demonstrated great practical success in the solution of high-dimensional optimal control problems but has not yet been significantly explored in nonequilibrium thermodynamics.

References

[1] Francis Bach. Breaking the curse of dimensionality with convex neural networks. *Journal of Machine Learning Research*, 18(19):1–53, 2017.

[2] Weinan E. Machine Learning and Computational Mathematics. *Communications in Computational Physics*, 28(5):1639–1670, 2020.

[3] Nicholas M. Boffi and Eric Vanden-Eijnden. Probability flow solution of the Fokker–Planck equation. *Machine Learning: Science and Technology*, 4(3):035012, 2023.

[4] Nicholas M. Boffi and Eric Vanden-Eijnden. Deep learning probability flows and entropy production rates in active matter. *arXiv:2309.12991*, 2023.

[5] Michael S. Albergo*, Nicholas M. Boffi*, and Eric Vanden-Eijnden. Stochastic Interpolants: A Unifying Framework for Flows and Diffusions. *arXiv:2303.08797*, 2023.

[6] Michael S. Albergo, Mark Goldstein, Nicholas M. Boffi, Rajesh Ranganath, and Eric Vanden-Eijnden. Stochastic interpolants with data-dependent couplings. *arXiv:2310.03725*, 2023.

[7] Michael S. Albergo, Nicholas M. Boffi, Michael Lindsey, and Eric Vanden-Eijnden. Multimarginal generative modeling with stochastic interpolants. *arXiv:2310.03695*, 2023.

[8] Nicholas M. Boffi and Jean-Jacques E. Slotine. A continuous-time analysis of distributed stochastic gradient. *Neural Computation*, 32(1):36–96, 2020.

[9] Nicholas M. Boffi and Jean-Jacques E. Slotine. Implicit regularization and momentum algorithms in nonlinearly parameterized adaptive control and prediction. *Neural Computation*, 33(3):590–673, 2021.

[10] Nicholas M. Boffi*, Stephen Tu*, and Jean-Jacques E. Slotine. Nonparametric adaptive control and prediction: theory and randomized algorithms. *Journal of Machine Learning Research*, 23(281):1–46, 2022.

[11] Nicholas M. Boffi*, Stephen Tu*, Nikolai Matni, Jean-Jacques Slotine, and Vikas Sindhwani. Learning stability certificates from data. In *Proceedings of the 2020 Conference on Robot Learning*, volume 155 of *Proceedings of Machine Learning Research*, pages 1341–1350, 2021.

[12] Nicholas M. Boffi*, Stephen Tu*, and Jean-Jacques E. Slotine. Regret bounds for adaptive nonlinear control. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, volume 144 of *Proceedings of Machine Learning Research*, pages 471–483, 2021.

[13] Saminda Abeyruwan, Alex Bewley, Nicholas M. Boffi, Krzysztof Marcin Choromanski, David B D'Ambrosio, Deepali Jain, Pannag R Sanketi, Anish Shankar, Vikas Sindhwani, Sumeet Singh, Jean-Jacques Slotine, and Stephen Tu. Agile catching with whole-body mpc and blackbox policy learning. In *Proceedings of The 5th Annual Learning for Dynamics and Control Conference*, volume 211 of *Proceedings of Machine Learning Research*, pages 851–863, 2023.

[14] Nicholas M. Boffi*, Stephen Tu*, and Jean-Jacques Slotine. The role of optimization geometry in single neuron learning. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 11528–11549, 2022.

[15] Thomas Zhang, Stephen Tu, Nicholas M. Boffi, Jean-Jacques Slotine, and Nikolai Matni. Adversarially robust stability certificates can be sample-efficient. In *Proceedings of The 4th Annual Learning for Dynamics and Control Conference*, volume 168 of *Proceedings of Machine Learning Research*, pages 532–545, 2022.

[16] Nicholas M. Boffi, Yipei Guo, Chris H. Rycroft, and Ariel Amir. How microscopic epistasis and clonal interference shape the fitness trajectory in a spin glass model of microbial long-term evolution. *eLife*, 12, 2023.

[17] Nicholas M. Boffi and Chris H. Rycroft. Parallel three-dimensional simulations of quasi-static elastoplastic solids. *Computer Physics Communications*, 257:107254, 2020.

[18] Nicholas M. Boffi and Chris H. Rycroft. Coordinate transformation methodology for simulating quasistatic elastoplastic solids. *Physical Review E*, 101:053304, 2020.

[19] Nicholas M. Boffi, Manish Jain, and Amir Natan. Efficient computation of the Hartree–Fock exchange in real-space with projection operators. *Journal of Chemical Theory and Computation*, 12(8):3614–3622, 2016.

[20] Nicholas M. Boffi, Manish Jain, and Amir Natan. Asymptotic behavior and interpretation of virtual states: The effects of confinement and of basis sets. *The Journal of Chemical Physics*, 144(8):084104, 2016.

[21] Katiana Kontolati, Darius Alix-Williams, Nicholas M. Boffi, Michael L. Falk, Chris H. Rycroft, and Michael D. Shields. Manifold learning for coarse-graining atomistic simulations: Application to amorphous solids. *Acta Materialia*, 215:117008, 2021.

[22] Nicholas M. Boffi, Judith C. Hill, and Matthew G. Reuter. Characterizing the inverses of block tridiagonal, block Toeplitz matrices. *Computational Science & Discovery*, 8(1):015001, 2014.

[23] Matthew G. Reuter, Nicholas M. Boffi, Mark A. Ratner, and Tamar Seideman. The role of dimensionality in the decay of surface effects. *The Journal of Chemical Physics*, 138(8):084707, 2013.

[24] https://github.com/malbergo/stochastic-interpolants.

[25] https://github.com/nmboffi/sbtm.

[26] https://github.com/nmboffi/spin_glass_evodyn.

[27] https://github.com/nmboffi/stzpp.

[28] http://real-space.org/.

[29] Richard Jordan, David Kinderlehrer, and Felix Otto. The Variational Formulation of the Fokker–Planck Equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, 1998.

[30] Michael E. Cates and Julien Tailleur. Motility-Induced Phase Separation. *Annual Review of Condensed Matter Physics*, 6(1):219–244, 2015.

[31] J. W. Costerton, L. Montanaro, and C. R. Arciola. Biofilm in implant infections: its production and regulation. *The International Journal of Artificial Organs*, 28(11):1062–1068, 2005.

[32] H.-C. Flemming. Biofouling in water systems–cases, causes and countermeasures. *Applied Microbiology and Biotechnology*, 59(6):629–640, 2002.

[33] Udo Seifert. Entropy Production along a Stochastic Trajectory and an Integral Fluctuation Theorem. *Physical Review Letters*, 95(4):040602, 2005.

[34] Udo Seifert. Stochastic thermodynamics, fluctuation theorems, and molecular machines. *Reports on Progress in Physics*, 75(12):126001, 2012.

[35] Joel L. Lebowitz and Herbert Spohn. A Gallavotti–Cohen-Type Symmetry in the Large Deviation Functional for Stochastic Dynamics. *Journal of Statistical Physics*, 95(1):333–365, 1999.

[36] Gavin E. Crooks. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Physical Review E*, 60(3):2721–2726, 1999.

[37] Cesare Nardini, Étienne Fodor, Elsen Tjhung, Frédéric van Wijland, Julien Tailleur, and Michael E. Cates. Entropy Production in Field Theories without Time-Reversal Symmetry: Quantifying the Non-Equilibrium Character of Active Matter. *Physical Review X*, 7(2):021007, 2017.

[38] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations. In *Proceedings of the 9th International Conference on Learning Representations*, 2021.

[39] Sunghan Ro, Buming Guo, Aaron Shih, Trung V. Phan, Robert H. Austin, Dov Levine, Paul M. Chaikin, and Stefano Martiniani. Model-Free Measurement of Local Entropy Production and Extractable Work in Active Matter. *Physical Review Letters*, 129(22):220601, 2022.

[40] Shriram Chennakesavalu and Grant M. Rotskoff. Unified, Geometric Framework for Nonequilibrium Protocol Optimization. *Physical Review Letters*, 130(10):107101, 2023.

[41] Erik Aurell, Krzysztof Gawedzki, Carlos Mejia-Monasterio, Roya Mohayaee, and Paolo Muratore-Ginanneschi. Refined Second Law of Thermodynamics for Fast Random Processes. *Journal of Statistical Physics*, 147(3):487–505, 2012.

[42] Gabriel Peyré and Marco Cuturi. Computational Optimal Transport. *arXiv:1803.00567*, 2020.

[43] Cédric Villani. *Optimal Transport*, volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer, Berlin, Heidelberg, 2009.

[44] https://openai.com/dall-e-3.

[45] https://midjourney.com.

[46] https://stability.ai.

[47] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. In *Proceedings of the 10th International Conference on Learning Representations*, 2022.

[48] Marylou Gabrie, Grant M. Rotskoff, and Eric Vanden-Eijnden. Adaptive Monte Carlo augmented with normalizing flows. *Proceedings of the National Academy of Sciences*, 119(10):e2109420119, 2022.

[49] François Rozet and Gilles Louppe. Score-based data assimilation. *arXiv:2306.10574*, 2023.

[50] Robert M. Sanner and Jean-Jacques E. Slotine. Stable adaptive control of robot manipulators using "neural" networks. *Neural computation*, 7(4):753–790, 1995.

[51] David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Phys. Rev. Lett.*, 35, 1975.

[52] Julian Davies and Dorothy Davies. Origins and Evolution of Antibiotic Resistance. *Microbiology and Molecular Biology Reviews*, 74(3):417–433, 2010.

[53] R. Craig MacLean and Alvaro San Millan. The evolution of antibiotic resistance. *Science*, 365(6458):1082–1083, 2019.

[54] Benjamin H. Good and Michael M. Desai. The Impact of Macroscopic Epistasis on Long-Term Evolutionary Dynamics. *Genetics*, 199(1):177–190, 2015.

[55] Craig A Fogle, James L Nagle, and Michael M Desai. Clonal Interference, Multiple Mutations and Adaptation in Large Asexual Populations. *Genetics*, 180(4):2163–2173, 2008.

[56] Michael J. Wiser, Noah Ribeck, and Richard E. Lenski. Long-Term Dynamics of Adaptation in Asexual Populations. *Science*, 342(6164):1364–1367, 2013.

[57] Richard E. Lenski. Experimental evolution and the dynamics of adaptation and genome evolution in microbial populations. *The ISME Journal*, 11(10):2181–2194, 2017.

[58] Elena Orlova, Aleksei Ustimenko, Ruoxi Jiang, Peter Y. Lu, and Rebecca Willett. Deep Stochastic Mechanics. *arXiv:2305.19685*, 2023.

[59] Weinan E, Weiqing Ren, and Eric Vanden-Eijnden. Minimum action method for the study of rare events. *Communications on Pure and Applied Mathematics*, 57(5):637–656, 2004.

[60] Matthias Heymann and Eric Vanden-Eijnden. The geometric minimum action method: A least action principle on the space of curves. *Communications on Pure and Applied Mathematics*, 61(8):1052–1117, 2008.

[61] Alexandre J. Chorin, Ole H. Hald, and Raz Kupferman. Optimal prediction and the Mori–Zwanzig representation of irreversible processes. *Proceedings of the National Academy of Sciences*, 97(7):2968–2973, 2000.

[62] Alexandre J. Chorin, Ole H. Hald, and Raz Kupferman. Optimal prediction with memory. *Physica D: Nonlinear Phenomena*, 166(3-4):239–257, 2002.

[63] David A. Sivak and Gavin E. Crooks. Thermodynamic metrics and optimal paths. *Physical Review Letters*, 108(19):190602, 2012.

[64] Grant M. Rotskoff, Gavin E. Crooks, and Eric Vanden-Eijnden. Geometric approach to optimal nonequilibrium control: Minimizing dissipation in nanomagnetic spin systems. *Physical Review E*, 95(1):012148, 2017.