

THE UNIVERSITY OF CHICAGO

UNCOVERING UNDERLYING DYNAMICS IN NON-EQUILIBRIUM SYSTEMS
THROUGH COMPLEMENTARY MODELING PERSPECTIVES

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
AND
THE FACULTY OF THE DIVISION OF THE BIOLOGICAL SCIENCES
AND THE PRITZKER SCHOOL OF MEDICINE
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

GRADUATE PROGRAM IN BIOPHYSICAL SCIENCES

BY
ELISABETH RENNERT

CHICAGO, ILLINOIS

AUGUST 2026

© 2026 by Elisabeth Rennert

ORCID iD: <https://orcid.org/0009-0004-3375-1844>

for Grandma

TABLE OF CONTENTS

LIST OF FIGURES	vi
LIST OF TABLES	x
ACKNOWLEDGMENTS	xi
ABSTRACT	xii
1 INTRODUCTION	1
1.1 Predicting the behavior of active, self-assembled systems	1
1.2 The actin cytoskeleton and current modeling approaches	4
1.3 Thermodynamic analysis of actin assembly	6
1.4 Dimension reduction and latent space analysis	8
1.5 Inference through generative diffusion	9
1.6 In this work	11
2 A THERMODYNAMIC FRAMEWORK FOR NON-EQUILIBRIUM SELF-ASSEMBLY AND FORCE MORPHOLOGY TRADEOFFS IN BRANCHED ACTIN NETWORKS	14
2.1 Abstract	14
2.2 Introduction	15
2.3 Overview of Branched Actin Networks	19
2.3.1 Brownian Ratchet	19
2.4 A minimal thermodynamic model for branched actin network growth and assembly	20
2.4.1 An analytical steady state solution	23
2.4.2 Steady state and dynamical behavior in various parameter regimes qualitatively resembles some experimental trends	27
2.4.3 Entropy production rate and thermodynamic constraints on force- morphology curves	30
2.4.4 A minimal non-adaptive growth process model identifies dissipative costs of maintaining adaptive response	31
2.4.5 Adaptive properties predicted by our thermodynamic model and driven by the dissipation gap	32
2.5 Discussion	34
2.6 Supplemental Information	38
2.6.1 Effective Model	38
2.6.2 Numerical comparison to experimental values	39
2.6.3 Estimation of adaptation rate	39

3	CHARACTERIZING THE PHASE SPACE OF CHEMICAL REACTION NETWORKS	41
3.1	Introduction	41
3.2	Experimental design and simulated waves	42
3.3	Analysis of the waves using machine learning	43
3.4	Machine learning methods	45
3.5	Conclusion	47
4	DIFFUSION MODELS LEARN UNDERLYING TRENDS IN ACTOMYOSIN NETWORKS AND PREDICT BEHAVIOR AT UNSEEN FILAMENT TURNOVER	49
4.1	Abstract	49
4.2	Introduction	49
4.3	Methods	52
4.3.1	Coarse-grained simulation data captures system scale trends	52
4.3.2	Diffusion models conditioned on physically relevant parameters	53
4.4	Results	56
4.5	Discussion	57
4.5.1	Minimal model of motor binding regimes	60
4.6	Future Work	63
4.7	Supplemental Information	63
4.7.1	Cytosim Simulation	63
4.7.2	Diffusion Model	64
4.7.3	Autoencoder Evaluation	66
4.7.4	Model Variability	66
4.7.5	Motor Binding Model	67
5	CONCLUSION	71
5.1	Discussion	71
5.2	Future Work	72
	REFERENCES	74

LIST OF FIGURES

1.1	Comparison of modeling perspectives	4
1.2	Three strategies for inferring behavior in unsampled regions of parameter space. Filled points are observed. The shaded band is the unseen region. (a) thermodynamic inequalities delimit the accessible region everywhere, without enumerating it. (b) a relationship fit to the observed points is extrapolated into the unseen region. (c) a learned distribution is sampled at unseen conditioning values. . . .	8
2.1	Schematic of the branched actin network physical system.	18
2.2	Depiction of full model. (a) Representation of a state in the full model. The full order of events that have been incorporated into the chain is recorded, emulating a path traced through a full branched network growing. Desaturated components are what is approximated to be the structure of the complete network based on the composition of the filament of interest, resulting in some number of active free ends at the membrane. (b) Transition rates in the full model. $X \in \{N, P, C\}$ Grey squares can represent any type of event, as addition and removal rates depend only on the current number of free ends and most recent event.	21
2.3	The force response of the network impacts both the steady state average and the level of fluctuations around that steady state. (a) The distribution of the number of free ends is symmetrically distributed around the steady state average. As the total network density increases and the polymerization rate slows, the increasing width of the free ends distributions indicates greater fluctuations and a fractionally lower effect of the change of the single discrete nucleation or capping event given the higher overall number of free ends. (b) There is good agreement between the simulated E values, and the self consistent steady state solution. The trends of free ends E and load velocity $v_{load} = \frac{J_p \Delta}{E}$, with increasing load are similar to experimental ones. The divergence of the simulated and analytic solution at high loads reflects the increasingly wide distribution of free ends with increasing load as seen in (a). However this high load state is very close to the stall load as demonstrated by the vanishingly small load velocity in these cases, and might be considered under a different regime.	26
2.4	Force response of a network grown against a constant-stiffness load.	28
2.5	Steady state behavior in various parameter regimes of $c = \sum k_x$ and k_n/k_c . (a) Load velocity profiles normalized by c (top) and steady state polymerization probability (bottom) for various values of c (b) Steady state free ends (top) and polymerization probability (bottom) profiles for various values of k_n/k_c	29
2.6	Thermodynamic bounds on the non-equilibrium structure of the system.	33

2.7	The system adapts dynamically to changes in applied load. (a) The fast response system, v_{load} drops immediately when the load is increased, while the free end density changes more slowly, acting as the mechanism of adaptation. The mimic process fully reconstructs the load velocity, even away from steady states. The effect of the shift in free ends and system composition can be seen in the yellow line, which calculates the mimic process using the unadapted E and P_p values. (b) The adaptation error ($\Delta v_{load}/v_{load_0}$) and rate follow the trends expected in a tradeoff with energy dissipation. Adaptation rate was determined by fitting a logistic function to the free end response curve. Energy dissipation is calculated by the entropy production per monomer addition before the perturbation. . . .	35
3.1	Analysis of chemical waves by ML. (a) Illustration of the reduction of the high dimensional data (sequence of frames for each chemical channel) from image sequences to latent space and the reconstruction of images from the latent space vector. (b) Uniform Manifold Approximation and Projection (UMAP) of phase plot in the latent space. (c) Phase plot in the input parameter space projected from the latent space. Spots represent parameters for sampled data, with the color corresponding to behaviors determined visually from the numerical model. Crosses represent the phase boundary prediction from the latent space analysis. Colors have the same legend as in (b). At present, latent space analysis cannot accurately differentiate between sustained waves and tip waves.	43
3.2	For each space and time point the autoencoder learns a single latent variable from the seven simulated chemical channels. This allows us to capture wave propagation dynamics of the wave pulse as a whole as opposed to the evolution of a single species. Even a single latent variable can reconstruct the original data fairly well.	46
3.3	Learning tasks for phase boundary estimation. (a) The autoencoder learns the minimal space of frames. (b) Linear SVM learns separating hyperplanes between labeled points in latent space. (c) Train recurrent neural network on trajectories in latent space to estimate input parameters. The parameter map is less accurate for more oscillatory results as well as points far from existing parameter combinations.	47
3.4	Phase boundary estimation iteration loop.	48

4.1	(a) A schematic of an actomyosin system containing actin, myosin and crosslinkers. The high myosin region in the center creates a myosin gradient which drives cortical flow towards the center of the simulation box. See SI for simulation details. (b) Trends in average actin curvature and density show a linear density-turnover relationship and a peak in the curvature. Density is calculated as the number density of $0.2\ \mu\text{m}$ filament segments in the simulation box. (c) The average rate of cortical flow of actin towards the center of the box is dependent on the ratio of active stress and viscosity, which itself depends in part on the local curvature and density. (d) Heatmap data used for training. Curvature and density fields are calculated by binning the simulation data and mirroring over the center line (so the top edge is the highest myosin concentration). (e) Joint distributions of curvature and density pixel intensity in log scale. Distributions are across all pixel locations, removing spatial information.	51
4.2	(a) Schematic of trajectories mapping between the turnover dependent data distribution and a Gaussian distributed prior. We train an SGM to learn a score function conditioned on turnover rate that can sample heatmap images from the specified rate. The turnover rate acts as a control parameter rather than a discrete class label, so the conditional data distributions have significant overlap. Distributions shown are of curvature pixel value across all spatial locations. (b) A convolutional autoencoder is trained on the simulated heatmap data. The autoencoder uses 3 convolutional downsampling layers to reduce the heatmap images to a 40 dimensional bottleneck. We use a UMAP of the low dimensional latent space to visualize the data and compare the overlap in latent distributions.	54
4.3	Diffusion models are able to infer nonlinear trends in system level characteristics across varying gaps in training data. (a) Comparison of generated vs simulated (ground truth) average curvature and density with increasing turnover rate. Results are shown for the 6 (top) and 4 (bottom) extreme turnover rates seen in training. Black x's correspond to unseen turnover rates. The model is able to capture the curvature peak with only a few turnover rates at each extreme seen during training. (b) UMAP visualization of autoencoder latent space of heatmap data. Points are averages for the turnover rate indicated by the color, with shape indicating the simulated or generated data. Closer overlap indicates higher similarity between the simulated and generated images.	58

4.4	The joint distribution of curvature and density of the simulation data contains information about the distributions at higher turnover rates. (a) shows the joint curvature density distributions for all simulated turnover rates. The shape of the distribution shifts from a long tailed distribution at low turnover to a more compact distribution after the peak. Circles indicate the turnover rates corresponding to the slices taken in (b). (b) The joint curvature-density distributions at the 2 extreme turnover rates are considered. Slices are taken at the densities corresponding to the average density of the turnover rate indicated by the color. (c) Joint distributions produced by the minimal motor binding model. The model captures the shift in correlation as turnover increases and the trends in average values, though they have substantially less variance than simulated joint density and curvature distributions, leading to less overlap between individual turnover rates.	59
S2.1	State transitions in the effective Markov state model. Green arrows indicate forward rates and red indicates reverse rates.	38
S2.2	The adaptation rate is estimated by fitting a logistic function to the free end recovery curve. The logistic growth rate, k is used as the adaptation rate. For this plot, $L = 142.9, k = 0.18, t_0 = -7.02$	40
S4.1	Illustration of model variability across training epochs	67
S4.2	Comparison of model predictions with naive spline interpolation for 4 and 6 distinct seen turnover rates	67

LIST OF TABLES

S2.1	Values used to estimate experimental conditions	39
S4.1	Cytosim parameter values	64

ACKNOWLEDGMENTS

I have learned and grown so much throughout my PhD, and I would not have been able to do it without all the people who helped me along the way. I want to thank my advisor Suri Vaikuntanathan for all of his insight and guidance. My committee members Aaron Dinner, Margaret Gardel, and Ed Munro provided valuable discussions and feedback throughout my PhD. I also want to thank the Biophysics program, especially Adam Hammond and Michele Wittels, for being constantly supportive, particularly through the first few years. Starting graduate school in the fall of 2019 was not optimal timing. I would also like to thank all my collaborators and coauthors, as well as the other members of the Vaikuntanathan Group, for their help and inspiration.

ABSTRACT

Biological active matter self-assembles and operates far from equilibrium, and its behavior spans many scales. No single modeling perspective captures both the fine-grained behavior of a specific non-equilibrium system and the general principles that govern it. A combination of complementary perspectives, from different regions of the space of modeling approaches, can recover both. This thesis develops that idea across three such systems, using minimal thermodynamic models, dimensionality reduction, and generative machine learning. Together, these approaches characterize behavior in parameter regimes for which detailed data is not available. A minimal thermodynamic model of branched actin self-assembly captures the growth velocity under load, and its entropy production bounds the structures the network can access. The bound reveals a dissipation gap. This energy cost, beyond the work of moving the load, maintains the free-end fluctuations the network uses to adapt. For a chemical reaction-diffusion system producing traveling waves, a learned latent space both predicts the phase boundaries between wave behaviors from a small number of simulations and reduces the dynamics of seven chemical species to a single collective variable. For a simulated actomyosin cortex, a generative diffusion model trained on coarse-grained static images at a few filament turnover rates infers a non-monotonic curvature trend at rates it never saw. A minimal motor-binding model traces the nonlinearity to a diffusion-limited shift in motor binding rates. These perspectives range from data-driven inference to interpretable mechanism, trading fine-grained detail against generality in different ways. The same active matter can be treated as a region to bound, a trend to predict, or a distribution to sample.

CHAPTER 1

INTRODUCTION

1.1 Predicting the behavior of active, self-assembled systems

Active systems are found in a variety of fields [1, 2]. In the simplest description they are characterized by an input of energy that drives them away from equilibrium. Biology offers a rich variety of length scales, from cellular processes such as cytoskeletal assembly [3] and membrane remodeling, to flocking at the organismal level [4], to ecological networks and planetary biospheres. Such systems can also be created synthetically, through actively driven particles [5] or chemical reaction networks [6].

Driving a system away from equilibrium opens up a wide array of behaviors but forfeits the guarantees and assumptions that make equilibrium systems more mathematically tractable. Emergent order and structures inaccessible at equilibrium become possible once a system is actively driven, and this space of systems has been of longstanding interest to many fields. Of particular interest are actively self-assembled systems [7], which typically combine a relatively limited pool of distinct components into composite structures. A characteristic of active self-assembly, particularly in biology, is the interaction of multiple length and time scales: system-scale structures are assembled from a collection of microscopic interactions, with various feedback and control processes able to tune the result. These interactions give rise to rich nonlinear phase behavior, and it is precisely this range of interacting scales that frustrates any attempt to describe such systems from a single vantage point. Even in the narrower example of actin networks, there is no unifying theory predicting network behaviors or material properties.

The number of parameters used to specify a system influences the level of sloppiness [8] in defining the system. Sloppiness describes the level of parameter uncertainty when a given model is fit to data, indicating limited impact of individual parameters on the resulting

behaviors, in contrast to stiff parameters that exhibit a strong response on perturbation. In many highly parametrized models the sensitivity of predictions is concentrated in only a few stiff combinations of parameters, while the remaining sloppy directions can be varied over orders of magnitude with little effect on observable behavior [9, 8], a property found to be near-universal across systems-biology models [10]. Minimal models tend to have both fewer parameters and higher levels of imposed structure. The interpretability of the model comes from the imposed structure that ties parameters to specific meaningful quantities. Essentially using knowledge of the system and existing relationships of components to guess the stiff directions, as opposed to a more highly parameterized model that can learn these key parameters directly from a given dataset. However, in imposing that minimal structure, some of the information in the response is lost, and the detail of the system is locked into the scales considered by the model.

This tradeoff has a deeper origin in the structure of the models themselves, most commonly seen in machine learning in the problems of over and under fitting. The detail encoded in these sloppy directions is precisely what can be discarded in moving toward a more interpretable description, without sacrificing fidelity to system responses that the data can actually resolve. The microscopic detail of a system is largely sloppy with respect to its mesoscale observables, suggesting some collective variable might be best to predict emergent behaviors. This connects back to the problem of systems influenced by multi-scale effects. Including the full picture of microscopic detail can reproduce system behaviors but does not provide insight into which collective variables actually drive observable trends.

How a system is represented therefore sets the scale at which it can be usefully modeled. A representation that retains only coarse-grained observables exposes only the stiff directions of the problem, so the appropriate model is one defined at the intrinsic dimension of that data. A model with very few parameters cannot reproduce an entire microscopic configuration, but it can predict collective variables such as averages. Identifying that re-

duced set of variables is the data-side counterpart of sloppiness: dimension reduction asks how few variables are needed to describe the state of a system, while sloppiness asks how few combinations of parameters a model needs to reproduce it. These are two facets of the same low-dimensional structure [11], and the power of these complementary perspectives is explored more in Chapter 3.

These tradeoffs between model and data scale can define an axis in the space of models. Simply the number of parameters in a model does not capture their variety. The level of structure imposed on the response by the definition of a model plays a significant role in the types of questions it can answer. This defines an axis between the imposed meaning of physics-based and mechanistic models and data driven approaches exemplified by machine learning. This is a fundamental tradeoff between learning emergent features of the data and applying assumed fundamental principles to some initial state. These methods are complementary, and can be bridged through physics-informed learning schemes [12], sparse equation discovery [13], or dimension reduction to identify physically relevant collective variables [14, 15].

This work is concerned with describing and predicting the behavior of active, self-assembled systems from a variety of such perspectives, ranging from data-driven inference that captures mesoscale trends to minimal thermodynamic models that elucidate more general design principles. No single perspective captures both the fine-grained behavior of a specific system and the general principles that govern it. The three contributions that follow each touch on multiple regions of the modeling space (Figure 1.1), and the full picture of the systems emerges through their complementary perspectives.




Figure 1.1: A comparison of modeling perspectives. Various modeling approaches are placed approximately according to the level of structure imposed on the underlying system by the model. At the low end of imposed structure are fully data driven methods that agnostically predict features of the system. Classical machine learning methods fall on this end of the spectrum, using neural networks that are intentionally flexible in structure to learn patterns from the data. In this area are the diffusion model approach discussed in Chapter 4, and the autoencoder dimension reduction. At the mechanistic end are more classical modeling approaches. These impose interactions between components, making predictions by evolving some initial condition over time. Minimal thermodynamic models lie on the far left of the axis, leveraging the imposed thermodynamic structure to predict features of systems with minimal, highly interpretable parameters. Cytosim snapshot from [16].

1.2 The actin cytoskeleton and current modeling approaches

The actin cortex is a self-assembled network of actin and associated proteins localized just beneath the plasma membrane in many cells. This network has highly tunable mechanical properties that allow it to play a crucial role in force generation and transmission within the cell. Branched networks that form in the lamellipodia produce the coordinated pushing forces necessary for cell motility [3], while bundled fibers enable molecular transport and shape modulation during division, and long-range cortical flows establish polarity for a range

of processes including embryonic development [17], patterning, and directed motility [18].

These systems are a complex interplay of assembly, disassembly, modification, feedback, and regulation, mediated through a wide array of actin-associated proteins. These processes span a variety of time and length scales, allowing the cortex to adapt to unpredictable external conditions. This wealth of behavioral modes, however, presents challenges for modeling and prediction: despite the limited number of components, the nonlinear nature of the feedback and control processes thwarts any attempt to describe these systems from the perspective of a single length or time scale.

There is an ongoing effort to model and characterize these networks [19, 16, 18], but the wide range of behaviors arising from network geometry and boundary conditions [20, 21] complicates any overarching hydrodynamic model that can predict network response across the full parameter space. Agent-based simulations calibrated to a specific experimental system, by contrast, have proven predictive of system behavior [22, 23].

Most modeling approaches for actin can be categorized as either detailed mechanistic models or continuum elastic models. Mechanistic models attempt to simulate every reaction and interaction within the network, producing predictions highly dependent on the specific details of the system setup. This usually involves an agent-based representation of mechanical components such as actin, myosin, and crosslinkers, combined with chemical reaction-diffusion fields to account for signaling processes. Cytosim [23] and MEDYAN [24] are both commonly used simulation packages for this approach. The drawback of this approach is its limited generalizability: it requires a very specific enumeration of the boundary geometry, component concentrations, and reaction rates, all of which can affect the network's phase behavior, yet the simulations remain limited to the specified system. The high level of detail also makes these simulations computationally expensive, precluding broad parameter scans that would be necessary for a map from microscopic parameters to resulting mesoscale effects.

Continuum elastic models instead treat the cortex as a coarse-grained viscoelastic gel, balancing active and viscous stresses to evolve the system. With fewer parameters, their predictions are more interpretable in terms of the underlying mechanics driving the trends, and they can predict trends and phase behavior, though their accuracy depends on the specifics of the system and the assumptions built into the fields. These models cannot directly include microscopic detail below the scale of the field coarse-graining. Because both the active and viscous stresses result from the underlying actin structures, this approach needs a way to approximate those effects for the configuration of interest. It also cannot fully capture the fluctuations within the system, which are involved in some of the regulatory processes that modify cortical properties.

1.3 Thermodynamic analysis of actin assembly

Thermodynamic effects are a major part of many biological processes. At the scale of microscopic processes, thermal fluctuations are significant and shape many cellular responses. At equilibrium these fluctuations are described by the fluctuation-dissipation theorem, which relates the fluctuations in an observable to the energy dissipated through them. In the classic example of Brownian motion, this is the balance between the kinetic energy of a particle's motion and the drag that returns that energy to heat. Away from equilibrium these guarantees are lost, but a thermodynamic accounting still constrains what the system can do. In Chapter 2 we use such an approach to model branched actin network assembly.

Inference is what ultimately motivates building a model. At its most general, a model is a structured map from inputs to outputs, and the output follows from the structure of the map and the input alone. It is that structure, rather than prior knowledge of the answer, that supplies the prediction. Across this work we infer in three broad modes, illustrated in Figure 1.2. A model can bound the region a system can occupy, fit a relationship that predicts observables directly and extend it to conditions we have not sampled, or learn the

full distribution of the data and generate new samples from it. The boundaries between these modes are not sharp, but they differ in the kind of answer they produce.

Chapter 2 develops the first of these modes. A thermodynamic perspective lets us tap into the symmetries and conservation laws used in classical equilibrium models. This energetic approach connects the microscale consumption of energy in polymerization to the mesoscale structures it produces. The second law requires the entropy-production rate to be non-negative, which, for a given level of driving, constrains the structures the system can assemble and maintain. Thermodynamic uncertainty relations sharpen the constraint for systems far from equilibrium [25, 26]. The result is a bound on the available structures rather than a prediction of any single one. It marks off an inaccessible region of the space everywhere at once, corresponding to panel (a) of Figure 1.2.

The structures in the system are the resulting self-organization of the non-equilibrium growth process. Organization must be actively maintained against this fluctuating landscape, and that maintenance carries an energetic cost. The connection between energy and structure also means that modifying the driving force modulates the structures and material properties of the network. The force driving the growing network is therefore the primary control parameter for its organization.

Energy is central to the adaptive tradeoffs in the system. Many systems exhibit energy-speed-accuracy tradeoffs in maintaining a particular state. The energy dissipated by fluctuations in assembly and structure at steady state is connected to the system's ability to respond to perturbations. To hold the same state under perturbation, the system must adapt its structures in response. Larger dissipative fluctuations allow this adaptation to occur more quickly and more accurately. In Chapter 2 we find that branched networks pay a dissipative cost beyond what is needed to move the load, and we argue that this excess is what maintains the free-end fluctuations underlying their adaptive response.




Figure 1.2: Three strategies for inferring behavior in unsampled regions of parameter space. Filled points are observed. The shaded band is the unseen region. (a) thermodynamic inequalities delimit the accessible region everywhere, without enumerating it. (b) a relationship fit to the observed points is extrapolated into the unseen region. (c) a learned distribution is sampled at unseen conditioning values.

1.4 Dimension reduction and latent space analysis

In contrast to approaches that impose a mechanism on the system, dimension reduction is a data-driven process. The aim is to identify a lower-dimensional set of collective variables that still describes the system. Many methods exist, from linear ones such as PCA to nonlinear ones such as UMAP [27], t-SNE [28], and autoencoder architectures [29]. Rather than predicting an observable on its own, dimension reduction builds the space in which prediction happens. It is the substrate for inference rather than a mode of it.

These methods are widely used for compression and denoising. Their real power, however, lies in the identification of the key collective variables. Depending on the method, the reduced representation can offer insight into the underlying physics of the system [30]. There are also information based methods developed to identify variables most predictive of the future [14].

Autoencoders are a machine-learning architecture that passes the data through a low-dimensional bottleneck. They are trained with a loss defined by the distance between the reconstructed and original data. The bottleneck can be shrunk until the training loss begins to rise. This reduces the data to the minimal dimension needed to reconstruct it and defines

a latent space. Analysis of this space can provide further insight into the underlying data, since the arrangement of points often aligns with meaningful axes of variation [31, 32].

In Chapter 3 we apply this to a chemical reaction-diffusion network of small organic molecules, a system that produces traveling chemical waves whose form changes across the reaction parameters [6]. We view the system through two complementary reductions. A latent space built from the full simulation data captures the global structure of the parameter space, and fitting boundaries within it predicts the phase behavior at parameters we never simulated. This boundary is a data-driven fit rather than a bound on what is possible as we previously discussed. A more aggressive reduction to a single collective variable instead captures the traveling-wave dynamics. That a single variable sufficiently represents the traveling wave mode is itself evidence that the model the system calls for is low-dimensional. The first reduction tells us where the system changes behavior, and the second tells us how the waves move once it does. Neither alone describes the system, but together they characterize both the topology of its phase space and the dynamics within it.

1.5 Inference through generative diffusion

At the data-driven end of the modeling space are machine-learning approaches, which fit a flexible function to data rather than imposing a mechanism in advance. Machine learning approaches are characterized by three main components: a model, an objective function, and an optimizer. The model is the family of functions available. While there is no formal definition of machine learning specifying the model in question, the way the term is used implies a neural network with many parameters. The objective provides a target for how well the model outputs match the underlying data (alternately described as a loss to minimize). The optimizer is the procedure that adjusts the parameters to minimize the loss. The many parameters let the model capture relationships we could not write down by hand, at the cost of the interpretability that comes with an imposed structure.

Generative diffusion models are one such machine learning approach [33, 34, 35], and the most recent evolution of generative models. They learn the underlying distribution of the data, which lets them sample new points from it. This is the third mode of inference, where unknown conditions can be probed through sampling the learned distribution (Figure 1.2c). The diffusion process at its core is borrowed from non-equilibrium thermodynamics [36]. In a forward process the data is gradually destroyed, and as the diffusion equation evolves each point approaches a Gaussian prior. During training the model learns a score function that predicts the direction to move along the data manifold given the current position and time. To sample, the model follows the reverse process, with each step set by the learned score. For conditioned generation, the model is trained on labeled data so that the score also depends on a conditioning variable, which allows generation to be guided to a chosen point in parameter space.

Like any data-driven model, diffusion models are subject to over- and under-fitting. A model over-parametrized for its training data will memorize the images, learning a set of delta functions rather than the underlying distribution. An underfit model that has not fully learned the map will instead maintain the Gaussian prior through the reverse process.

In Chapter 4 we train a diffusion model on coarse-grained heatmap images of a simulated actomyosin cortex, spanning values of the actin turnover rate. The model is able to generate images from turnover rates never seen in training. From coarse-grained static images, the model extrapolates a system-scale trend across a region of parameter space it never saw in training. The conditioning is what makes this possible – the model learns the relationship between the conditioning variable and the data, and from the fluctuations within that data it extends the trend to new conditions.

The scale at which such a model is useful is set by how the data is represented, because the representation fixes the dimensionality of the relationship the model has to learn. Learning the full microscopic configuration as a function of the conditioning variable is a

high-dimensional problem, and constraining it would require dense sampling across a high-dimensional parameter space. Coarse-graining the data to collective fields, such as the average density and curvature of the network, reduces the task to a low-dimensional relationship between the conditioning variable and those fields. This is the same reduction to collective variables described in the previous section, and it is what lets the relationship be learned from a handful of sampled conditions and extended to ones it was not trained on. In the language of sloppiness, the collective fields are the stiff directions of the problem. The microscopic detail is sloppy, and although a coarse representation cannot reconstruct it, that detail is not what we are trying to predict, nor is it even non-degenerate with respect to the averages. A relatively simple architecture is therefore enough, and the modifications built for richer, hierarchical data, such as noise schedules [37] and classifier-free guidance [38], proved detrimental for our purposes.

The diffusion model can learn that a trend exists and where it goes, but not why. For that we turn to a minimal model of motor binding, which traces the nonlinearity in average curvature to a diffusion-limited shift in motor binding with increasing turnover. This is the second mode of inference, a mechanistic model that reproduces the trend and identifies its cause. The generative model still suffers the lack of interpretability characteristic of heavily parametrized models, but paired with the minimal model it offers a way to extend a detailed, system-specific description toward a more general one. The generative model extrapolates the behavior and the minimal model explains it, and neither alone gives the full account.

1.6 In this work

The three chapters that follow present the major projects of my PhD, each taking up a different non-equilibrium active system from different regions of the modeling space.

In Chapter 2 we find that thermodynamic bounds from a minimal model inform the space of structures a branched actin network can access, and reveal a dissipation gap that

controls its ability to adapt [39]. We build a one-dimensional chain model that captures the loaded growth velocity of the network for a given level of driving, then calculate the entropy production of the growing network and use it to bound the available configurations. The bound exposes a dissipative cost beyond what moving the load requires, which we argue maintains the free-end fluctuations that allow the network to adapt to changing loads.

In Chapter 3 we find that a learned latent space can both predict a system’s behavior and identify its key collective variables [6]. Working with a chemical reaction-diffusion network of small organic molecules, we identify the phase boundaries between wave behaviors from an autoencoder latent space and map them back to the original parameters. Reducing the seven chemical channels to a single latent variable then captures the traveling-wave dynamics, the functionally relevant feature of the system.

In Chapter 4 we find that a generative diffusion model can learn the nonlinear mesoscale trends in actomyosin networks, aided by the presence of a potential single mechanism behind the nonlinearity [40]. We train the model on coarse-grained static images at a few turnover rates and use it to infer the non-monotonic dependence of network average curvature, predictive of cortical flow, for turnover rates never seen in training. A minimal model of motor binding traces that nonlinearity to a diffusion limited shift in motor binding regime, the candidate mechanism that makes the trend coherent enough to learn from sparse data.

Taken together, these projects span the modeling space from data-driven inference to interpretable mechanism. The two machine-learning approaches pull trends and boundaries out of data but pay for it in interpretability, while the thermodynamic model sacrifices fine-grained detail to place global limits on what the system can do. Running through all three is a tradeoff between the system-specific detail needed to predict a particular setup and the ability to generalize across boundary conditions and parameters. No single one of these perspectives captures both the detailed behavior of its system and the principles that govern it. They are complementary instead. The data-driven characterizations supply the system-

specific behavior, the minimal models supply the mechanism and the generalizable limits, and the same non-equilibrium active matter can be viewed through a bound, a predicted trend, or a generated distribution depending on the perspective we take.

CHAPTER 2

A THERMODYNAMIC FRAMEWORK FOR NON-EQUILIBRIUM SELF-ASSEMBLY AND FORCE MORPHOLOGY TRADEOFFS IN BRANCHED ACTIN NETWORKS

2.1 Abstract

Branched actin networks are involved in a variety of cellular processes, most notably the formation of lamellipodia in the leading edge of the cell. These systems adapt to varying loads through force dependent assembly rates that allow the network density and material properties to be modulated. Recent experimental work has described growth and force feedback mechanisms in these systems. Here, we consider the role played by energy dissipation in determining the kind of growth-force-morphology curves obtained in experiments. We construct a minimal model of the branched actin network self-assembly process incorporating some of the established mechanisms. Our minimal analytically tractable model is able to reproduce experimental trends in density and growth rate. Further, we show how these trends depend crucially on entropy dissipation and change quantitatively if the entropy dissipation is parametrically set to values corresponding to a quasistatic state. Finally, we also identify the potential energy costs of adaptive behavior by branched actin networks, using insights from our minimal models. We suggest that the dissipative cost in the system beyond what is necessary to move the load may be necessary to maintain an adaptive steady state. Our results hence show how constraints from stochastic thermodynamics and non-equilibrium thermodynamics may bound or constrain the structures that result in such force generating processes.

2.2 Introduction

The creation and transmission of forces within a cell is a crucial function of the cytoskeleton, allowing the cell to modulate its shape in service of a variety of processes such as division, endocytosis, embryonic development, and motility through complex environments [41, 42, 43, 44]. Branched actin networks are a component of the cytoskeleton that produces a coordinated pushing force localized to the membrane [45, 3]. This is in contrast to the pulling forces generated by kinesins and other motor proteins elsewhere in the cytoskeleton [46].

Branched networks are characterized by their branched structure, with new filaments nucleated off the side of existing ones, through a branching factor such as Arp2/3 [47]. The filaments can then polymerize and push against the load until they are capped, preventing further growth (Fig. 2.1). This structure gives it unique mechanical and elastic properties, distinct from crosslinked or entangled networks [45]. The assembly mechanisms of the network growing against a membrane gives rise to force dependent rates of assembly, and allows branched networks to adapt to varying mechanical forces, allowing cells to navigate through heterogeneous environments [42]. A major function of these relationships is to drive motility in cells both internally through development of lamellipodia or externally as comet tails propelling invading bacteria [48].

There has been extensive experimental work investigating the molecular mechanisms of branched network assembly and load response [49, 45, 50, 51, 52]. However, *in vivo* these networks are difficult to isolate to probe force responses [53, 50], making it preferable to use reconstituted components *in vitro*. Additionally the filament and monomer level network structure can only be resolved with electron microscopy techniques [54, 55], so approximations from bulk network data are frequently used to describe the growing network. A common experimental setup for branched networks is reconstituted networks grown on surfaces coated in nucleation promoting factor (NPF) in an atomic force microscopy (AFM)

cantilever, which provides load and displacement information relating to the network’s elastic properties and growth rate. This information can be combined with fluorescence imaging, which provides densities and incorporation rates of the various species as the network grows, to allow interrogation of the microscopic processes driving branched network growth [45, 49].

When the branched network is growing against a load, the network density and growth rate are sensitive to force. Higher load results in increased filament density, decreased growth velocity, and an increase in efficiency of the network as a molecular motor [45]. This force response is caused by changes in the network structure, both through shifts in the interaction angle of filaments with the membrane [56] and the density of growing filaments [55, 45]. This force feedback response allows the network to respond to and modulate membrane tension to create non-equilibrium shapes as well as adapt to heterogeneous extracellular environments [57, 44, 42].

The complexity generated by this limited set of interactions has motivated a variety of models and simulations. On the most complex end are highly detailed whole cell simulation systems that seek to closely replicate experimental conditions [42]. However more minimal models can still capture network responses. For example, the shift in orientation angle in response to time-varying load was investigated through a Gillespie-like branch addition model [56]. An exactly solvable continuum model of an ensemble of filaments acting as pure ratchets against a fluctuating membrane with imposed inward drift can capture membrane velocity though not network density fluctuations [58]. More recently, a coarse-grained model used insights gleaned from a reinforcement learning approach to show how aspects of the external forces are encoded in the structure of the network [59].

We seek to create a minimal analytically tractable model that provides thermodynamic insight into the feedback mechanisms and force response trends in the system. Branched network assembly is a non-equilibrium process driven by ATP hydrolysis involved in actin polymerization which can serve to break detailed balance and drive the system to a non-

equilibrium state even in the absence of direct force generation as seen in molecular motors. This energy input can be dissipated in maintaining the structure of the network or used to perform work moving a load, usually a membrane.

Here we show how our minimal thermodynamic model is sufficient to capture experimental trends (Fig. 2.6). Next, by analyzing the entropy production rate, we show that these trends are only obtainable at substantial dissipative cost. We suggest that the extra dissipative cost, beyond what might be necessary to simply push a load, might be due to the need to maintain an adaptive steady state. Indeed, while the fluctuations in network free end density dissipate energy, they may also allow the system to maintain the motion of the membrane across shifts in load as part of an Energy Speed Accuracy (ESA) tradeoff [60]. The force response has been experimentally investigated with a variety of forcing conditions (e.g. constant stiffness, oscillating load) [53, 49] which can vary the tradeoffs in the system. We will focus on the quasistatic case where the load is held constant for a given network. We note previous work has investigated thermodynamic relationships between driving forces and resulting configurations in polymer-like self-assembly systems [61, 62]. This previous work explored simpler systems, but we expand the framework to describe the dynamics of branched network assembly.

The outline of the rest of the paper is as follows. We first describe the interactions involved in assembling the physical system of branched actin networks, then introduce our minimal model. We apply approximations to self-consistently solve for the steady state. We show how our minimal thermodynamic model is consistent with experimental trends. We then calculate the entropy production rate of the growing system, and show how it can provide a parametric bound on the available configurations.




Figure 2.1: NPFs concentrate actin monomers at the membrane. From there they transfer to growing filaments or nucleate new ones. Capping protein prevents further polymerization. The Arp2/3 branching factor interacts with a region of the NPF to transfer an actin monomer to activate the branching factor. Free ends can also polymerize with actin monomers from the solution, but grow four times faster with monomers transferred from the NPF [63], ensuring that network growth remains directed toward the membrane. Free ends interact with the membrane at angle θ . The on and off rates describe the association and dissociation rates for each of the binding events.

2.3 Overview of Branched Actin Networks

While there can be many components involved, branched networks can be reconstituted with only Arp2/3 branching factor, a Nucleation Promoting Factor (NPF) from the WASP family proteins, capping protein (CP), monomeric actin-profilin complex, and filamentous actin [49, 45, 64]. The NPF binds the profilin-actin complex, concentrating it at the membrane. The actin can then be transferred to the end of an existing filament and continue its polymerization or bind an Arp2/3 complex and nucleate a new filament, as described in Fig. 2.1. The CP binds to free ends and prevents further polymerization, meaning the filament will fall away from the growing front of the network. The growing ends of actin filaments can also bind to the NPFs which prevents nucleation and tethers the free end to the membrane in a process called barbed end interference [65]. This effect, along with competition between the Arp2/3 and the growing filaments for the pool of NPF-associated actin monomers [63] contributes to the negative feedback of uncapped free ends on the nucleation rates of new free ends [49]. Capping is necessary in the system to reduce the number of free ends in the system, which in turn increases nucleation rates as capped ends cannot bind NPFs or use NPF-associated actin monomers to polymerize [66]. Because Arp2/3 must bind to an existing filament to nucleate a new one, the nucleation process is autocatalytic. However, the negative feedback from high numbers of free ends prevent exponential nucleation [63] and allows the system to fluctuate around some average steady state network density. This steady state density depends on the rates of the various processes in the system, and can shift in response to load as the rates are impacted by Brownian Ratchet effects.

2.3.1 *Brownian Ratchet*

Both polymerization and capping of filaments in contact with the membrane can be described as a Brownian Ratchet process [49, 67]. Brownian ratchets are a theoretical framework to describe force generation through thermal fluctuations and steric hindrance. When a filament

is in contact with the membrane, a new monomer cannot be added unless fluctuations create a large enough gap to intercalate it. For a filament growing at angle θ to the membrane, the gap size would be $\Delta = \delta \sin \theta$, where δ is the monomer size. To move a load f that distance requires energy $C_f = f\Delta$, so the probability it will occur through thermal fluctuations is proportional to $\exp(-\frac{C_f}{k_B T})$ [68]. For an ensemble of filaments polymerizing in contact with the membrane, we assume the load is evenly distributed across all free ends (E), scaling the polymerization rate by $\exp(-\frac{C_f}{E k_B T})$.

The average free end length of a filament in the network is determined by the ratio of polymerization and capping rates. Because both these processes are Brownian ratchets, the rates scale exponentially with the energy required to move the membrane. Both processes occur at the same site on a free end, so the difference in energy cost to add each respective monomer type depends only on the size, δ . Both monomers are similar in size, allowing the rates to scale together with load, and keeping the ratio of the rates and thus the free end lengths constant, as has been observed experimentally [45]. The Brownian Ratchet component of the capping rate is crucial to the force response feedback loop in the network. If the number of free ends is higher than the steady state value, capping is amplified due to lower per-filament load, which in turn reduces E . Conversely, if the density is low, polymerization and capping are both suppressed and nucleation is encouraged, allowing the system to return to the steady state average.

2.4 A minimal thermodynamic model for branched actin network growth and assembly

To investigate the thermodynamic bounds on the structural fluctuations in this type of system, we develop a minimal model that still captures the force response trends observed in experiments. This model emulates a column of branched actin network growing against a load, as is frequently employed in experimental setups [45, 49, 53, 51].




Figure 2.2: Depiction of full model. (a) Representation of a state in the full model. The full order of events that have been incorporated into the chain is recorded, emulating a path traced through a full branched network growing. Desaturated components are what is approximated to be the structure of the complete network based on the composition of the filament of interest, resulting in some number of active free ends at the membrane. (b) Transition rates in the full model. $X \in \{N, P, C\}$ Grey squares can represent any type of event, as addition and removal rates depend only on the current number of free ends and most recent event.

As our model assembles, it is characterized by the number of free ends (E), which is analogous to the network density and impacts the rates at which components are incorporated. We assume spatial homogeneity as well as a constant load, so the load is evenly distributed across all free ends and the per-filament load is determined by E, allowing us to model the system as a one dimensional series of nucleation (N), polymerization (P), and capping events (C), that can be thought of like monomers added to a growing chain. We further assume a large enough reservoir that background monomer concentrations remain constant as the network grows. This allows us to define constant per-filament addition rates for the zero load case, where the Brownian Ratchet component is 1. These rates are based on the background monomer concentrations and the k_x^{on} association rates, where $x \in \{n, p, c\}$, as shown in Fig. 2.1,

$$\begin{aligned} k_n &= k_n^{on} d_{NPF}[M] \\ k_p &= k_p^{on} d_{NPF}[M] \\ k_c &= k_c^{on} [CP] \end{aligned} \tag{2.1}$$

where d_{NPF} is the density of NPFs at the load, $[M]$ is the concentration of soluble actin monomers, and $[CP]$ is the concentration of capping protein. The corresponding removal rates are simply the k_x^{off} rates.

Nucleation and capping events respectively increase or decrease, the number of free ends (E), while polymerization moves the load forward. The addition and removal rates for the events depend on the overall configuration of the system are depicted in Fig. 2.2, with only the most recent event able to be removed, in which case the second most recent is now the ‘tip’ of the chain. This linear picture is meant to approximate tracing a single path or ‘filament of interest’ through an existing network and can be viewed as a mean field model. We then take that sequence of events to be representative of the average structure of the network as a whole.

To make analytical progress we make one further approximation (which is justified subsequently). In the model depicted in Fig. 2.2, the history of the events on the filament of interest is completely captured as the order of events is preserved. However, this model's infinite state space make it difficult to solve the master equation, so we will instead consider an effective Markov state model defined by the most recent event and the number of free ends, E , shown in Fig. S2.1. In this simplified effective model, the absence of a history results in uncertainty about the appropriate state to transition to upon event removal. Prediction of the next state relies on calculating conditional probabilities based on the second most recent event. In the case described in Fig. 2.2b top, the reverse rate for the effective model would become $k_n^{off}(E+1)P(X, E|N, E+1)$. Because the number of free ends is prescribed by the type of the removed event the conditional probability can be expressed simply as $P(X|E)$.

Due to the autocatalytic nature of filament nucleation with the negative feedback of E , the number of free ends in a growing network fluctuates around some steady state $E_{ss} = \langle E \rangle$, representing the most probable value of E . Our objective is to replicate the growth and force response of a network comprising numerous polymerizing filaments in a minimal model that can be analyzed thermodynamically. Given that these are bulk properties resulting from an ensemble of filaments, we prioritize the average behavior over extreme and infrequent states. To capture the bulk steady state behavior of the network we approximate

$$\begin{aligned} E &\approx E_{ss} \\ P(N|E_{ss}) &\approx P(N) = P_n \end{aligned} \tag{2.2}$$

when E is close to E_{ss} , which is where the system spends the most time.

2.4.1 *An analytical steady state solution*

A consequence of our steady state assumption is that the state probabilities are proportional to the flux of that event type into the system J_x . This means that

$$P_x = \frac{J_x}{J_{tot}} \quad (2.3)$$

Where $J_{tot} = J_n + J_p + J_c$. For there to be a steady state E , the flux of nucleation and capping events into the system must be equal. So, by setting $J_n = J_c$ we get

$$k_n - k_n^{off} E_{ss} P_n = k_c E_{ss} e^{-\frac{C_f}{E_{ss}}} - k_c^{off} P_c \quad (2.4)$$

and solve for E_{ss}

$$E_{ss}(\mathbf{P}) = \frac{C_f(k_n + k_c - P_c)}{C_f k_n^{off} P_n + (k_n + k_c^{off} P_c)A} \quad (2.5)$$

$$A = W_0 \left(\frac{C_f k_c \exp \left(-\frac{C_f k_n^{off} P_n}{k_n + k_c^{off} P_c} \right)}{k_n + P_c} \right)$$

where $W_0(x)$ is the Lambert W function.

Eq. 2.5 implies that the steady state structure can be completely described by the state probabilities $\mathbf{P}$. To solve for the steady state where $\frac{d\mathbf{P}}{dt} = 0$, we construct a matrix $\mathbf{M}$ such that $\frac{d\mathbf{P}}{dt} = \mathbf{M}\mathbf{P}$. $\mathbf{M}$ is determined by the master equation for our minimal model described in Eq. S2.1. The eigenvector of $\mathbf{M}$ corresponding to the zero eigenvalue will give us the steady state $\mathbf{P}$.

Now that we have developed and solved our minimal model of branched network assembly, we verify the model through simulations and comparison to experimental findings. Once the association and dissociation rates in the system are set, we can simulate the growing system using a kinetic Monte Carlo algorithm. Events (adding a new monomer through nucleation, polymerization or capping) are successively added and removed to the chain, with rates determined by the current number of free ends, as described in the full model in Fig. 2.2.

The previous event is part of the current state, so conditional probabilities are not needed to identify the new tip of the chain upon a removal event. Nucleation and capping events increment or decrement E , respectively. We evolve the system until it reaches a steady state. Once at the steady state we let the system assemble until the chain of events reaches a large enough end length to estimate the steady state event probabilities. These are calculated simply from the composition of the chain. Regardless of the initial number of free ends, the system will reach and fluctuate around a steady state $E_{ss} = \langle E \rangle$, as well as steady state event probabilities $\mathbf{P}$.




Figure 2.3: The force response of the network impacts both the steady state average and the level of fluctuations around that steady state. (a) The distribution of the number of free ends is symmetrically distributed around the steady state average. As the total network density increases and the polymerization rate slows, the increasing width of the free ends distributions indicates greater fluctuations and a fractionally lower effect of the change of the single discrete nucleation or capping event given the higher overall number of free ends. (b) There is good agreement between the simulated E values, and the self consistent steady state solution. The trends of free ends E and load velocity $v_{load} = \frac{J_p \Delta}{E}$, with increasing load are similar to experimental ones. The divergence of the simulated and analytic solution at high loads reflects the increasingly wide distribution of free ends with increasing load as seen in (a). However this high load state is very close to the stall load as demonstrated by the vanishingly small load velocity in these cases, and might be considered under a different regime.

2.4.2 *Steady state and dynamical behavior in various parameter regimes*

qualitatively resembles some experimental trends

The structures generated by the simulation using the full model allow us to calculate the probability of incorporating each event type. We then use this to self-consistently calculate the steady state solution in the effective 3-state model. The distributions agree between the simulation and the steady state solution as shown in Fig. 2.3b. Our results also agree with the general force response trends observed experimentally of networks growing under step-wise increasing load, with convex decay in growth velocity and increase to saturation of free ends with increasing load [45, 49]. We also investigate the behavior of spring-like loads, capturing an initial plateau followed by decay of the load velocity, as shown in Fig. 2.4. We simulated the full model with a Hookean spring-like load defined as $f = k_{spring}x(t)$, where $x(t)$ is the instantaneous position of the front, and the energy to move the load one monomer length remains $C_f = f\Delta$. The load response of this system demonstrated some of the plateau structure in seen in experimental works [51, 52]. Interestingly, higher spring constants lead to larger plateaus, likely due to slower initial velocity allowing the system to adapt more fully before reaching the high load regime where system composition begins to change.

We examined a range of values of the base forward rates to identify the quantities crucial to system behavior. We refer to the base forward rates generically as k_x where $x \in n, p, c$ and k_{rx} similarly refers to the base reverse rates. Fig. 2.5 (a) [top] shows when the load velocity profiles are normalized by their $c \equiv \sum k_x$ value, they seem to converge asymptotically as the c is increased. Further, these results also predict that as c is increased, the system shifts to a polymerization driven regime (as indicated by a high value of P_p). Hence, for the low load regime where $k_x \gg k_{rx}$, which is where many experimental systems operate, we may expect the force response to scale with the unloaded forward rates. This echoes experimental data in which force velocity curves are normalized by unloaded growth rates [49, 52, 51].

We also considered changing the ratio k_n/k_c Fig. 2.5 (b). For a fixed k_n/k_c , we find that increasing the load shifts slowly the system from a polymerization dominated regime Fig. 2.5 (b) (bottom). This is consistent with Fig. 2.5 (a)(bottom). Further, the number of free ends increases with the load in a manner that is qualitatively consistent across a range of values of k_n/k_c . This consistency in our parameter scans shows how the results from our minimal model can be extended to different experimentally important regimes.




Figure 2.4: (a) The instantaneous force response of a network grown against a load with constant stiffness. In this case the load exerts a force $f(t) = x(t)k_{spring}$, where $x(t)$ is the current displacement of the membrane. The energy required to move the membrane on monomer gap size remains $C_f(t) = f(t)\Delta$. This system cannot be solved for a non equilibrium steady state due to the constant change in C_f until the network reaches stall. However the response with an initial plateau at low C_f followed by decay of load velocity shows trends similar to experimental findings in cases with spring-like loads. (b) After the initial drop in monomer addition rate, it recovers as the system begins to approach stall and shift the system composition. [51, 52]




Figure 2.5: Steady state behavior in various parameter regimes of $c = \sum k_x$ and k_n/k_c . (a) Load velocity profiles normalized by c (top) and steady state polymerization probability (bottom) for various values of c (b) Steady state free ends (top) and polymerization probability (bottom) profiles for various values of k_n/k_c

2.4.3 Entropy production rate and thermodynamic constraints on force-morphology curves

Using the forward and reverse rates in the effective model, we define the total entropy production rate $\dot{S}$, which is the sum over the entropy production of each of the possible transitions. This allows us to calculate

$$\begin{aligned} \frac{\dot{S}}{J_{tot}} = & -\frac{C_f}{E}(P_p + P_c) + P_c \ln \left(\frac{k_n k_c}{k_n^{off} k_c^{off} P_c^2} \right) \\ & + P_p \ln \left(\frac{k_p}{k_p^{off} P_p} \right) \end{aligned} \quad (2.6)$$

the average entropy production per monomer addition.

The steady state assumption from Eq. 2.2 we used to simplify the model introduces some error. If $E \approx E_{ss}$ then $P(E) \approx P(E_{ss}) \approx 1$, which would be represented by a delta function. However, the distribution $P(E)$ in the full model, as shown in Fig. 2.3a, is Gaussian, with broader distributions seen at higher loads. The broader the underlying distribution of E values is, the farther from a delta function and thus the worse our approximation is. The increasing deviation of the distributions correlates with the slight disagreement between steady state solution and simulation seen in Fig. 2.3b at high load. These broader distributions are also caused by a suppression of polymerization at higher loads, so more time is spent adding and removing free ends which drives the fluctuations.

Eq. 2.6 provides a basis to obtain thermodynamic bounds between various quantities. For example we define

$$\epsilon_{load} \equiv \frac{C_f}{E}(1 - P_n) \quad (2.7)$$

as the thermodynamic cost per monomer addition due to the load. The requirement that

$\dot{S} \geq 0$ in Eq. 2.6 can now be used to bound and connect such thermodynamic costs to quantities that parameterize the structure such as E the number of free ends in the system. In Fig. 2.6a, we show how the non-negativity of the entropy production rate can be parametrically expressed as a bound on available states for a given load or monomer addition cost. In Fig. 2.6b, we describe a similar plot but now connecting the number of free ends E to the force on the membrane, C_f . Specifically, the second law bound on ϵ_{load} shown in Fig. 2.6a can be expressed as

$$\epsilon_{load} \leq P_c \ln \left(\frac{k_n k_c}{k_n^{off} k_c^{off} P_c^2} \right) + P_p \ln \left(\frac{k_p}{k_p^{off} P_p} \right) \quad (2.8)$$

with the bound on the load in Fig. 2.6b as

$$C_f \leq \left(P_c \ln \left(\frac{k_n k_c e^{C_f/E}}{k_n^{off} k_c^{off} P_c^2} \right) + P_p \ln \left(\frac{k_p}{k_p^{off} P_p} \right) \right) E \quad (2.9)$$

using the assumptions from Eqs. 2.3 and 2.4 we can express the monomer probabilities as

$$P_c = \frac{k_c E e^{-C_f/E} - k_n}{k_c^{off} - k_n^{off} E} \quad (2.10)$$

$$P_n = P_c, \quad P_p = 1 - 2P_c$$

2.4.4 *A minimal non-adaptive growth process model identifies dissipative costs of maintaining adaptive response*

The motion of the load due to polymerization and the maintenance of the system structure through free end density fluctuations can be viewed as linked but separable processes. The load response behavior can be recreated by a simple 1D Brownian walker. For a given system configuration defined by E and P and a load condition C_f , we define a mimic process that

takes steps of size Δ forward and back with rates

$$\begin{aligned} k_+ &= k_p \exp -C_f/E \\ k_- &= k_p^{off} P_p. \end{aligned} \tag{2.11}$$

We can further define the entropy production of this mimic process as

$$\dot{S}_{mimic} = (k_+ - k_-) \ln \frac{k_+}{k_-} \tag{2.12}$$

Interestingly, this entropy production rate results in bounds that are very close to the bounds discussed in the previous section (Fig. 2.6). Because the entropy production of this process is doing work moving the load, the rest of the entropy production of the system is dissipated to maintain the structure of the system. We call this cost a dissipation gap. In the next section we argue that this extra dissipative cost may in fact be connected to the cost associated with maintaining an adaptive response in Arp2/3 networks.

2.4.5 Adaptive properties predicted by our thermodynamic model and driven by the dissipation gap

Given that a primary function of branched actin networks is in motility through inhomogeneous networks, we next asked if we can consider our model as an adaptive system, that engages in Energy Speed Accuracy (ESA) tradeoffs between the energy dissipated by the system and the accuracy and speed of that adaptation. Applying this framework to a branched system, the adaptation problem is maintaining consistent motion (v_{load}) in response to perturbation in the external environment ($C_f \rightarrow C_f + \Delta C_f$).

The perturbation response exhibited by our model is shown in Fig. 2.7, where v_{load} drops immediately upon increased load, while E increases to the new steady state value in the slow response that adapts the system to its new conditions in an adaptation pro-




Figure 2.6: The entropy production imposes thermodynamic bounds on the non-equilibrium structure of the system, characterized by E . The grey shaded area indicate inaccessible configurations. (a) shows bound on ϵ_{load} from Eq. 2.8, while (b) uses the bound in Eq. 2.9, both parametrically derived from the entropy production in Eq. 2.6. For a given average system density, there is some maximum load or monomer addition cost that the system has been assembling under. The distance between the model prediction and the mimic process bound reflects the amount of energy dissipated by the growing system. Yellow points in both plots represent values calculated from the experimental data of Li *et al.* [49] (Fig 1D bottom), scaled as described in Eq 2.13. Scaling parameters were $s = 19.1$, $\beta = -9.5$, and $\kappa = .093$

cess like those described in [60]. We speculate that this tradeoff is fueled by the energy dissipation of the nucleation/capping processes, which can be defined as $\dot{S} - \dot{S}_{mimic}$. The necessity of the energy overhead needed to maintain the system is demonstrated in Fig. 2.7a, which demonstrates the impact of the system characteristics adapting in response to the perturbation. The mimic process uses the instantaneous values to calculate its load velocity $v_{lm}(E(t), P_p(t), C_f + \Delta C_f)$. However, without adaptation in the system, the mimic process would predict $v_{lm}(E_0, P_{p0}, C_f + \Delta C_f)$, which does not maintain the desired growth, as shown in Fig. 2.7a. For our system, we define the adaptation error as the change in v_{load} relative to the original value. The adaptation rate is extracted from fitting a logistic function to the free end recovery curve. The tradeoffs follow the same trends seen in [60], with higher energy dissipation allowing for greater accuracy, and adaptation speed upon a perturbation.

2.5 Discussion

The entropy production we calculated allows us to place thermodynamic bounds on the available structures in the system. We can also compare this to experimental data from Li *et al.* [49] (Figure 1D bottom), of normalized branched end density vs growth stress. We scaled the experimental data to our steady state solution values. The experimental data was presented in terms of the growth stress Γ and normalized branched end density ρ . We scaled these using fitting parameters s, β, κ to account for unit and normalization differences. as follows

$$\begin{aligned} E_{exp} &= \rho s - \beta \\ C_{f,exp} &= \Gamma \kappa \\ \epsilon_{load,exp} &= \frac{C_{f,exp}}{E_{exp}} (1 - P_n(E_{exp})) \end{aligned} \tag{2.13}$$

where $P_n(E)$ is calculated from Eq. 2.5.




Figure 2.7: The system adapts dynamically to changes in applied load. (a) The fast response system, v_{load} drops immediately when the load is increased, while the free end density changes more slowly, acting as the mechanism of adaptation. The mimic process fully reconstructs the load velocity, even away from steady states. The effect of the shift in free ends and system composition can be seen in the yellow line, which calculates the mimic process using the unadapted E and P_p values. (b) The adaptation error ($\Delta v_{load}/v_{load_0}$) and rate follow the trends expected in a tradeoff with energy dissipation. Adaptation rate was determined by fitting a logistic function to the free end response curve. Energy dissipation is calculated by the entropy production per monomer addition before the perturbation.

The load only needed a single parameter as the zero point aligned with the simulation. This is analogous to the velocity scaling based on c described in Fig. 2.5. The branched end densities, however, were normalized, so we could not scale them directly based off predicted rates. A second parameter was necessary to create an offset accounting for variations in the unloaded network density between our simulation and the experiment. We then optimized the parameters for the best agreement. As additional confirmation, if we estimate the rate constants using the values noted in Table S2.1 and use experimental reagent concentrations from [49], we get an unloaded velocity within about 8% error demonstrating further agreement with experiments.

With the qualitative agreement between the experimental and theoretical trends we can now comment on the importance of energy dissipation to maintain these non-equilibrium structures so that they can adapt to varying extracellular environments. The parametric forms inferred in Eq. 2.8 and 2.9 correspond to the quasistatic limit where no excess entropy is dissipated. These parametric forms substantially differ from those observed in experiments and in our theoretical model. This demonstrates that the energy dissipation, in the form of polymerization-depolymerization, capping, filament turnover, are all crucial to obtain the experimentally observed force morphology relations.

We also interrogate the bound on available structures using thermodynamic uncertainty relations [26]. These are a class of relations that extend the fluctuation-dissipation relation to non-equilibrium systems, describing how entropy production rates can bound the fluctuations in an evolving system. When applied to self-assembly processes, the thermodynamic uncertainty relation can be expressed as

$$\dot{S} \geq \frac{v^2}{D} \quad (2.14)$$

where $v = \langle J_{tot} \rangle = \frac{\Delta N}{\Delta t}$, the total flux of events into the system, and $D = \lim_{\tau \rightarrow \infty} \frac{\langle \Delta N^2 \rangle}{2\tau}$. D captures the fluctuations in growth rate in the system, and is analogous to a diffusion

coefficient [26]. This provides a tighter bound on $\dot{S}$ than the second law bound given by $\dot{S} \geq 0$. However, in this system the inclusion of the fluctuations in assembly rate did not strongly impact the thermodynamic bounds. This is likely because the system growth does not have a linear response to increasing loads. Since the parametric forms obtained in this manner differ substantially from the observed theory profiles, we can again conclude that far from equilibrium fluctuations are most likely responsible for the maintenance of the force-morphology relations of branched Arp2/3 networks.

In summary, the bulk properties of growing branched networks under constant load can be captured in a minimal thermodynamic model. From this limited system we can determine bounds on the available structures expressed in terms of experimentally accessible quantities. Our work quantifies how dissipative processes in the background play a crucial role in allowing the system to adapt to a range of load conditions without fundamentally altering the underlying rates, determining the force-morphology-growth tradeoff curves.

This could be useful in analysis of the system, especially in interactions with heterogeneous environments where the load or resistance conditions are not fully known. This theoretical framework could be also used to analyze data from branched networks grown under unknown conditions, and potentially make inferences about the conditions the network has encountered. The model is also much more sensitive to changes in load than the underlying rates of the system making it robust to changes in background concentrations. In the future we would like to expand the thermodynamic analysis of branched networks to include responses to time-varying and heterogeneous loads such as one would encounter in cellular environments through interactions with membranes. We can also investigate the interactions between different non-equilibrium systems as in the interactions between assembling networks and growing membranes.

2.6 Supplemental Information

2.6.1 Effective Model




Figure S2.1: State transitions in the effective Markov state model. Green arrows indicate forward rates and red indicates reverse rates.

Applying the steady state approximations in Eq. 2.2 to the effective model in Fig. S2.1, we can express the master equation for the minimal model

$$\begin{aligned}
 \frac{dP_n}{dt} &= -(k_n^{off} E_{ss}(1 - P_n) + k_p E_{ss} e^{-C_f/E_{ss}} + k_c E_{ss} e^{-C_f/E_{ss}}) P_n + (k_n + k_p^{off} E_{ss} P_n) P_p \\
 &\quad + (k_n + k_c^{off} E_{ss} P_n) P_c \\
 \frac{dP_p}{dt} &= -(k_p^{off} E_{ss}(1 - P_p) + k_n + k_c E_{ss} e^{-C_f/E_{ss}}) P_p + (k_p E_{ss} e^{-C_f/E_{ss}} + k_p^{off} E_{ss} P_p) P_n \\
 &\quad + (k_p E_{ss} e^{-C_f/E_{ss}} + k_c^{off} P_p) P_c \\
 \frac{dP_c}{dt} &= -(k_c^{off} (1 - P_c) + k_n + k_p E_{ss} e^{-C_f/E_{ss}}) P_c + (k_c E_{ss} e^{-C_f/E_{ss}} + k_n^{off} E_{ss} P_c) P_n \\
 &\quad + (k_c E_{ss} e^{-C_f/E_{ss}} + k_p^{off} E_{ss} P_c) P_p
 \end{aligned} \tag{S2.1}$$

2.6.2 Numerical comparison to experimental values

Parameter	Value	Source
k_{poly}^{on}	$11\mu M^{-1}s^{-1}$	Pollard 1986 [69]
k_{poly}^{off}	$1s^{-1}$	Pollard 1986 [69]
k_{cap}^{on}	$6.3\mu M^{-1}s^{-1}$	Wear <i>et. al.</i> 2003 [70]
k_{cap}^{off}	$5 \times 10^{-4}s^{-1}$	Wear <i>et. al.</i> 2003 [70]
k_{nuc}^{on}	$6.3\mu M^{-1}s^{-1}$	Mullins <i>et. al.</i> 1998 [71]

Table S2.1: Values used to estimate experimental conditions

We could not locate an experimental estimate of k_{nuc}^{off} , but we will assume that $k_{nuc}^{on}/k_{nuc}^{off} \approx 10^4$. Using these values and the conditions from Li *et. al.* ($5\mu M$ actin, $100nM$ CP) we can simulate the model. This resulted in an unloaded v_{load} of $6.7 \mu m/min$, which agrees with the experimental values from [49] of $6.2 \mu m/min$

2.6.3 Estimation of adaptation rate




Figure S2.2: The adaptation rate is estimated by fitting a logistic function to the free end recovery curve. The logistic growth rate, k is used as the adaptation rate. For this plot, $L = 142.9$, $k = 0.18$, $t_0 = -7.02$

CHAPTER 3

CHARACTERIZING THE PHASE SPACE OF CHEMICAL REACTION NETWORKS

3.1 Introduction

The work described in this chapter is from [6], where we used latent space analysis to characterize a chemical reaction network designed by our experimental collaborators. Inspired by the level of control in biological reaction networks, there is significant interest in using reaction networks of small organic molecules to create tunable chemical waves. Among many other applications they can be used in soft robotics as a method of controlling hydrogel actuators. In [72] they design such an actuator using stressed and unstressed hydrogel connected by disulfide crosslinks. Thiol species break these bonds, so a traveling autocatalytic thiol wave can be used to cause a mechanical transformation in the material. The functional application of these waves drives interest in the characterization and control of these autocatalytic fronts. Characterizing this system is difficult because the system characterization is high dimensional with interacting chemical fields of several species, while the behavior it exhibits is set by only a few control parameters and falls into a handful of qualitative classes. This gap is what makes dimension reduction the natural tool.

In this chapter we use a learned latent space to identify phase boundaries between resulting wave behaviors. We also use dimension reduction to identify a single collective variable that characterizes the traveling wave behavior. Together these show that a data-driven latent space can provide insights into the behavior of this far from equilibrium system.

3.2 Experimental design and simulated waves

The full chemical structure of the reaction network can be found in [6], but the crucial elements of the reaction are source species that react to form thiols. This reaction is catalyzed by the thiols, driving the initial formation of the wave front. This is combined with negative feedback from maleimide, which removes the thiols.

The experimental basis of the oscillatory system is an unstirred hydrogel flow reactor. Source species flow along the sides of a hydrogel and diffuse into the center. The source species are split between two channels to prevent reaction at the gel interface, so the center axis of the reactor is where the initial reactions take place. The impact of the diffusivity of maleimide was investigated by attaching varying lengths of polyethylene glycol to the maleimide. The difficulty of synthesizing these altered maleimide species meant only three different diffusivities were experimentally tested. To further explore the space, our collaborators designed a 2D numerical reaction-diffusion simulation. I ran this simulation to produce the data we used for training our machine learning models.

This setup can result in a range of oscillatory behaviors: sustained waves traveling the length of the reactor, sustained fluctuations only at the far end of the reactor (tip waves), damped waves, or no initiation of waves at all. The control parameters considered were the source concentration and diffusion coefficient of maleimide. The phase plot can be seen in Fig. 3.1c. We wished to predict how these control parameters affect the resulting oscillatory behavior.




Figure 3.1: Analysis of chemical waves by ML. (a) Illustration of the reduction of the high dimensional data (sequence of frames for each chemical channel) from image sequences to latent space and the reconstruction of images from the latent space vector. (b) Uniform Manifold Approximation and Projection (UMAP) of phase plot in the latent space. (c) Phase plot in the input parameter space projected from the latent space. Spots represent parameters for sampled data, with the color corresponding to behaviors determined visually from the numerical model. Crosses represent the phase boundary prediction from the latent space analysis. Colors have the same legend as in (b). At present, latent space analysis cannot accurately differentiate between sustained waves and tip waves.

3.3 Analysis of the waves using machine learning

To better understand the chemical waves in our system, we analyzed this system using a convolutional autoencoder structure that passes the data through a low dimensional bottleneck while still preserving enough information to reconstruct the input data [73, 74], ensuring that

essential features are still present in the data (Fig. 3.1a). We trained the autoencoder using images generated by the numerical model because they contain more chemical information and less noise than experimental data while still capturing all the essential features of the experimental system. The low dimensional latent space learned by the autoencoder contains the most important features of large and potentially noisy input data space. Dynamics in this reduced dimensional space can additionally provide insights into the dynamics of the full system [32, 75]. We use these ML techniques to identify phase boundaries on wave behaviors as well as investigate a minimal dynamical model of the system.

We estimated the phase boundaries in the parameter space of the system using machine learning techniques to predict boundary points and behaviors for parameter combinations not sampled by simulations [76]. From a small number of simulations, we used a linear Support Vector Machine (SVM) to calculate the latent space boundary between different labeled resulting behaviors (Fig. 3.1b). The separating hyperplane was then projected back to the space of input parameters, allowing prediction of behavior directly from input parameters (Fig. 3.1c). The initial boundary prediction is noisy, but converges quickly with a small number of resampled points from the identified region of interest. We resample based on points selected in the latent space and then projected back to the parameter space as it provides more efficient coverage of less frequent behaviors [77, 78]. Machine learning guided sampling and boundary prediction of this system allows estimates of behavior transition boundaries with significantly fewer samples than would be needed in the full parameter space.

To propose a minimal model of the system, we first averaged the data from the numerical model along the short axis for each channel, giving us a series of 1D concentration profiles over time for each chemical species (Fig. 3.2). We then trained an autoencoder on the combination of concentrations for each time and space point, allowing us to reduce 7 channels to only 1. This suggests that there is a minimal model that could represent concentration profiles

of reactive species through only a single dynamical variable [32]. We find that the following model qualitatively describes the waves seen in the one dimensional latent variable. The wave velocity as a function of the latent variable, u , is given as

$$v(u) = v_0 + \bar{v} \nabla \ln(u), \quad (3.1)$$

which allows us to connect the wave propagation to the gradient of a chemical potential-like quantity. The full wave can then be modeled as

$$\frac{\partial u}{\partial t} = D_u \nabla^2 u - \nabla v(u) + u(m(x) - u) \quad (3.2)$$

where $m(x)$ is the background source gradient, analogous to the source chemicals diffusing out from the flow channels, and approximated as linear. This wave has the structure of a Fisher-like wave with a changing background concentration. Overall, this analysis helps identify some of the important driving forces and characteristics of the resultant waves. Specifically, in the latent dimension, the waves are convected by the equivalent of a chemical potential-like gradient. Further, source and sink terms reminiscent of those encountered in population dynamics models also drive the wave forms.

3.4 Machine learning methods

First, using a uniform sampling of parameter values in a region of interest, we ran 25 simulations of the numerical model and then used a convolutional autoencoder on the individual frames of the simulation to reduce the 7 spatial concentration profiles to a latent space of 500 variables. The autoencoder was able to accurately reconstruct the original frames, indicating that the latent space is able to fully capture the details of the system. We then trained a recurrent neural net to map from sequences of points in the latent space to the input parameter space varying D_I (diffusion coefficient of maleimide) and I_0 (concentration




Figure 3.2: For each space and time point the autoencoder learns a single latent variable from the seven simulated chemical channels. This allows us to capture wave propagation dynamics of the wave pulse as a whole as opposed to the evolution of a single species. Even a single latent variable can reconstruct the original data fairly well.

of maleimide supplied to the gel). To predict the phase boundary, we used a linear SVM classifier to find the boundaries between different behaviors. We then used the parameter map we had previously trained to project the boundary into the parameter space (Fig. 3.3). These boundaries are hyperplanes in the latent space, and trajectories on the hyperplane need to be sampled to project back to the parameter space. To achieve this we took the closest point on the hyperplane to each existing point from the frames of the original simulations. The parameter map is the most accurate at estimating points close to parameter combinations it was trained on, so if a point is projected to the hyperplane from too far away it can move into regions of the latent space the parameter estimator has not seen, leading us to estimate the phase boundary using only points close enough to the hyperplane. The estimated boundary is very sensitive to the choice of distance cutoff, but there are a few metrics that can be considered based on characteristics of the data and the shape of the boundary. Distance from the separating hyperplane can be used as a predictive metric for




Figure 3.3: Learning tasks for phase boundary estimation. (a) The autoencoder learns the minimal space of frames. (b) Linear SVM learns separating hyperplanes between labeled points in latent space. (c) Train recurrent neural network on trajectories in latent space to estimate input parameters. The parameter map is less accurate for more oscillatory results as well as points far from existing parameter combinations.

behavior [79, 30].

Once the boundary is predicted we can refine it using a small number of resampled points to update the parameter estimator and SVM boundary to converge on an estimated phase boundary quite quickly (Fig. 3.4). The improvement of the boundary is not sensitive to the exact choice of resampled points as long as they cover an area near the boundary, preferably spaced away from existing points to aid the parameter estimator.

3.5 Conclusion

An alternative to the reductionist approach for constructing dynamic chemical systems is the use of machine learning (ML), which captures the behavior of a system as a whole in response to changes in external parameters. The first step in this direction is to train ML on the numerical model of the system. In this work, we demonstrated that ML, using latent space analysis, identifies and extracts the most relevant parts of extremely high dimensional data and captures essential dynamic features of the system. As a result, it can predict the




Figure 3.4: Phase boundary estimation iteration loop.

phase space of the system and suggest the dimensionality of a minimalistic model that would capture some essential properties of the system. For design purposes, the next step would involve extracting major correlations between system behaviors and control parameters within a large parameter space using ML. These correlations can then be used to guide the design of such systems.

CHAPTER 4

DIFFUSION MODELS LEARN UNDERLYING TRENDS IN ACTOMYOSIN NETWORKS AND PREDICT BEHAVIOR AT UNSEEN FILAMENT TURNOVER

4.1 Abstract

Generative diffusion models have demonstrated an ability to produce novel images sampled from the learned underlying data distribution. These models are able to infer system characteristics for parameter combinations that were not seen during training. We investigate the ability of these models to infer trends in biological data from limited samples. Specifically, we consider the response of system scale behaviors such as cortical flow in a simulated actomyosin system as we tune filament turnover rates. We train a diffusion model on coarse-grained actin curvature and density heatmap images, and are able to generate images from conditioning variables not seen during training. These images are predictive of nonlinear trends in the system. We also consider characteristics of the system that allows this level of inference, such as the strong linear relationship between average density and filament turnover in the system, and explore minimal underlying dynamics with a motor binding model.

4.2 Introduction

The actin cortex is crucial to the generation and transmission of forces within a cell. This actively assembled biopolymer network interacts with a variety of proteins that modulate the network structure and growth [80]. The balance between active stresses generated by myosin motors and network viscosity can produce coordinated and sustained flows in the actin cortex [81, 16]. Cortical flows establish polarity within the cell to direct motility

[82, 83], cytokinesis [84] or morphogenesis [17, 85, 86]. This wide range of applications is possible due to the highly tunable properties of actin networks, partly controlled by their regulatory interactions with various polarity establishing proteins.

Cortical flows are implicated in the establishment of polarity in *C. elegans* embryos through asymmetrical contractions establishing a localization of PAR proteins to the developing anterior pole of the cell [17]. The patterning protein CDC-42 creates a gradient of the molecular motor myosin which drives the cortical flow toward the anterior half of the cell [87, 22]. The ability to predict cortical flows and their impacts on cellular processes across parameter perturbations is currently limited by the availability of detailed observations of the actin cortex’s structure and dynamics, whether through experiments or detailed, and thus computationally expensive, simulations. Given the high level of network variability that can be produced even from a limited set of components, an ability to predict mesoscale system behaviors from limited samples would allow more rapid exploration of high dimensional parameter spaces. Generative diffusion models have emerged in recent years as powerful tools to create new samples from complex high dimensional distributions [35, 33, 88, 36, 34, 89]. In this work we investigate the ability of these models to generalize across the parameter space and predict cortical flow in novel regimes.

Experiments in reconstituted actin networks [90] and in *C. elegans* embryos [22] have shown that the plastin, formin, and profilin concentrations can non-monotonically tune the rate of cortical flow, suggesting a mechanism to control flows in embryo development. Across perturbations, the filaments show a consistent rate of nucleation with distinct assembly and disassembly phases. We will describe this actin assembly/disassembly rate as the turnover throughout this paper. This system’s cortical flow behavior in response to varying turnover can be reproduced through carefully tuned Cytosim simulations [23, 22, 16]. In these simulations, there is a non-monotonic response of cortical flow with increasing turnover, shown in Fig. 4.1c. Our main result shows how generative diffusion models are able to predict

non-monotonic trends just from static images provided at points from different mechanical regimes, suggesting a capacity to identify underlying interactions and predict resulting system scale dynamical behaviors like cortical flow.




Figure 4.1: (a) A schematic of an actomyosin system containing actin, myosin and crosslinkers. The high myosin region in the center creates a myosin gradient which drives cortical flow towards the center of the simulation box. See SI for simulation details. (b) Trends in average actin curvature and density show a linear density-turnover relationship and a peak in the curvature. Density is calculated as the number density of $0.2 \mu\text{m}$ filament segments in the simulation box. (c) The average rate of cortical flow of actin towards the center of the box is dependent on the ratio of active stress and viscosity, which itself depends in part on the local curvature and density. (d) Heatmap data used for training. Curvature and density fields are calculated by binning the simulation data and mirroring over the center line (so the top edge is the highest myosin concentration). (e) Joint distributions of curvature and density pixel intensity in log scale. Distributions are across all pixel locations, removing spatial information.

4.3 Methods

4.3.1 Coarse-grained simulation data captures system scale trends

The specific simulation setup we consider is a 2D system of filamentous actin, myosin and crosslinkers with a high myosin region at the center, as shown in Fig. 4.1a. The imposed myosin gradient causes anisotropic contractility which drives flows of actin [86] towards the center of the box. We employ the turnover dynamics described in [16], where the filament nucleation rate is fixed and the filaments grow for a fixed time at the specified turnover rate before disassembling at that same turnover rate. We simulate this system using the Cytosim [23] simulation package at 15 different actin turnover rates. This configuration generates consistent cortical flows towards the center of the simulation box, in agreement with experimentally observed trends during polarity establishment in the *C. elegans* embryo [22].

To prepare the training data for our diffusion model, we reduced the full simulation data to coarse-grained fields that are still predictive of network responses to perturbations. Analysis of autoencoder latent space suggests that both the density and curvature fields are required to predict flows [16]. Additionally, local actin curvature and density can also be measured experimentally through video microscopy, making the ability to infer trends from these fields broadly relevant. We create heatmap images of the actin curvature and density fields by dividing the simulation box into $1\text{ }\mu\text{m} \times 1\text{ }\mu\text{m}$ bins, averaging the simulated network values within those bins, shown in Fig. 4.1d.

Due to the fixed actin growth window, the turnover rate determines the length of the filaments in the system. In combination with the constant nucleation rate, this results in the strong linear relationship between turnover and actin density shown in Fig. 4.1b. The peak in curvature is due to a decrease in network mesh size with increasing actin density. A smaller mesh size means the motors are more likely to bind a filament, exhausting the

available pool of motors. After the motors are mostly bound, increasing actin density results in a drop in density of bound motors per length of assembled actin. Both the cortical flow and average curvature exhibit a peak as the turnover rate increases, with the curvature peak corresponding to the point with the highest motor density on the filaments. This is indicative of an underlying shift in the motor binding regime. This transition can also be seen in the shift in shape of the curvature-density joint distribution (Fig. 4.1e), with a much longer tailed curvature distribution at low turnover, shifting to a more compact distribution at high turnover. The long tail in the curvature distribution at low turnover indicates more extreme fluctuations in the relatively lower amount of F-actin.

If a generative diffusion model is able to infer the nonlinear curvature trend in the system arising from an underlying shift in motor activity, that indicates an ability to generalize across mechanical paradigms that could allow it to predict more complex behaviors like cortical flow.

4.3.2 *Diffusion models conditioned on physically relevant parameters*

We train a score-based generative model (SGM) to learn a score function in the space of actin heatmap images conditioned with the actin turnover rate. During training our model learns a conditional score function given by $\mathbf{S}(\mathbf{x}_t, c) = \nabla \log P_t(\mathbf{x}_t|c)$, where c is the class label, in our case the turnover rate. Our goal is to sample generated heatmaps conditioned with turnover rates not seen during training. We considered a variety of models trained on subsets of the 15 simulated turnover rates. For these models we adjust the number of frames seen per class in training such that all models are trained on 54,000 total images.

A physically relevant continuous conditioning variable allows us to avoid the problem of identifying an embedding that preserves the semantic relationships between class labels as is necessary with natural language descriptions, though it limits us to sampling heatmaps along that axis. As the turnover rate increases the resulting network behaviors, and thus underlying




Figure 4.2: (a) Schematic of trajectories mapping between the turnover dependent data distribution and a Gaussian distributed prior. We train an SGM to learn a score function conditioned on turnover rate that can sample heatmap images from the specified rate. The turnover rate acts as a control parameter rather than a discrete class label, so the conditional data distributions have significant overlap. Distributions shown are of curvature pixel value across all spatial locations. (b) A convolutional autoencoder is trained on the simulated heatmap data. The autoencoder uses 3 convolutional downsampling layers to reduce the heatmap images to a 40 dimensional bottleneck. We use a UMAP of the low dimensional latent space to visualize the data and compare the overlap in latent distributions.

distributions of the heatmap images, shift. The high levels of overlap in distributions across turnover rates indicate that an individual turnover rate might contain information about neighboring rates. The conditional score functions learned by diffusion models capture the relationship between the data and conditional labels to the extent they can even be used as a classifier [91]. This indicates that the conditional score functions contain some core features of how the data varies across classes, and might pick up on underlying relationships and trends.

The structure and complexity of the training data has a strong impact on the types of models that are most effective and the level of prediction that can be drawn from the model. With this work we focus on the coarse-grained heatmap data rather than the full simulation as we are concerned with the ability of the model to infer nonlinearity in the large scale trends, rather than making specific predictions of microscopic structures. The heatmap data has only limited spatial anisotropy, as can be seen in Fig. 4.1d, with areas of high curvature slightly more likely at the top of the images due to the myosin gradient. This means the primary challenge is in capturing both the averages and the extreme fluctuations in the system. The coarse-grained nature of the data combined with the high levels of overlap between different simulated turnover rates indicates there is not much hierarchy in the features the model learns, so the diffusion trajectories do not speciate (Fig. 4.2a) into separate distributions as is seen in more hierarchical data [92] like photographs. This allows us to use the simplest version of an SGM without noise scheduling or high levels of tuning and still learn characteristics of the underlying data distributions. We learn more general relationships within the data that indicate system level trends than details of the microscopic interaction of the individual components.

While we train on static heatmap images without the temporal correlations necessary to calculate cortical flow, there are other machine learning methods that attempt to predict the dynamic behaviors of such networks from static snapshots [93, 94, 95]. Other approaches to

predict cortical flow could include a hydrodynamic model for this specific system configuration that connects the curvature and density fields to cortical flow. Additionally the model could be trained on sequences of heatmap images to also learn the temporal correlations that reflect the cortical flow in the system.

To evaluate the quality of the generated images, we train a convolutional autoencoder on the heatmap data. This model has a 40 dimensional bottleneck which provides a representation of key features necessary to reconstruct the data. We further reduce the dimensionality by taking a UMAP projection of the autoencoder latent space to visualize the heatmap distributions in 3D as shown in Fig. 4.2b. In both the simulated and generated heatmaps, the images are distributed along an underlying manifold corresponding to the turnover rate, with some overlap between neighboring rates, indicating this projection largely captures the high level features of the data. This further allows us to evaluate the overlap of the latent image distributions between simulated and generated heatmap images.

4.4 Results

Our model is able to successfully predict the global trends in average density and curvature, even across the nonlinearity in curvature as the system shifts motor binding regimes. For a model trained on the 6 extreme turnover rates (the 3 highest and 3 lowest turnover values) of our simulation, it was able to generate images that followed the average curvature and density trends from the full simulation as shown in Fig. 4.3a. The model could still predict the nonlinearity with as few as 4 seen turnover rates. While a naive spline interpolation was only slightly worse at predicting average trends with 6 seen rates, at 4 rates the generative model had substantially better predictions, shown in Fig. S4.2. Furthermore, a spline interpolation provides information only about average values without predicting any qualities of the distribution. The latent distributions in the autoencoder bottleneck were also in agreement as can be seen in the UMAP visualizations in Fig. 4.3b, indicating enough identifying

features were created in the generated images to target them to the correct areas in the autoencoder latent space. There can be some tradeoffs in the generated images between matching the averages and fully capturing the long tailed nature of the true distributions at lower turnover rates. The generated images tended to be more Gaussian in distribution, which is not surprising given the Gaussian prior of the diffusion model.

The size of the gap in turnover rate significantly affected the ability of the model to predict behaviors. Training on 4 turnover rates seems to be the minimum number of turnover rates and largest prediction gap size that can give accurate results. However, there can still be a lot of variation in model accuracy both during training and between models trained on the same subsets of data, illustrated in Fig. S4.1. This indicates there could be numerous spurious minima in the high dimensional loss landscape that the model can fall into without fully learning the details of the distribution. The model variability can be overcome by training an ensemble of models and choosing the ones with closest overlap between simulated and generated images at the training turnover rates.


The ability of these models to generate image data that follow underlying physical trends suggests a further ability to infer dynamic behaviors like directly predicting cortical flow with the right model architecture and resolution of data. The complexity of what the model needs to garner from the data is also demonstrated in that the limit in the model’s ability to infer behavior is not completely defined by the amount of data or even the turnover gap size, but also the number of distinct turnover rates seen in training.

4.5 Discussion

Given the ability of the model to predict characteristics of the actomyosin system at unseen turnover rates, we consider what features of the data allow the model to make these predictions from limited data. The lack of tuning required for the model and the simplicity of the underlying SDE indicate the existence of some minimal underlying dynamics driving these

a

curvature, density vs turnover



b

average umap values




Figure 4.3: Diffusion models are able to infer nonlinear trends in system level characteristics across varying gaps in training data. (a) Comparison of generated vs simulated (ground truth) average curvature and density with increasing turnover rate. Results are shown for the 6 (top) and 4 (bottom) extreme turnover rates seen in training. Black x's correspond to unseen turnover rates. The model is able to capture the curvature peak with only a few turnover rates at each extreme seen during training. (b) UMAP visualization of autoencoder latent space of heatmap data. Points are averages for the turnover rate indicated by the color, with shape indicating the simulated or generated data. Closer overlap indicates higher similarity between the simulated and generated images.




Figure 4.4: The joint distribution of curvature and density of the simulation data contains information about the distributions at higher turnover rates. (a) shows the joint curvature density distributions for all simulated turnover rates. The shape of the distribution shifts from a long tailed distribution at low turnover to a more compact distribution after the peak. Circles indicate the turnover rates corresponding to the slices taken in (b). (b) The joint curvature-density distributions at the 2 extreme turnover rates are considered. Slices are taken at the densities corresponding to the average density of the turnover rate indicated by the color. (c) Joint distributions produced by the minimal motor binding model. The model captures the shift in correlation as turnover increases and the trends in average values, though they have substantially less variance than simulated joint density and curvature distributions, leading to less overlap between individual turnover rates.

trends.

The quality of the predictions we have shown indicates that there is information about other unseen turnover rates contained in the training data. We propose that the linear relationship between turnover and density allows the model to map fluctuations away from the mean at seen turnover rates to behaviors at unseen rates. This is aided by the high degree of overlap in the distributions, even across the full range of turnover rates considered. In Fig. 4.4b, we show slices from the joint distribution at the average densities associated with various turnover rates in the simulation data. The shape of these slices changes as the density approaches the point corresponding to the curvature peak, and the mode shift in motor saturation. These slices reflect how the trends across turnover rate are linked to density fluctuations within the network. Even at very low turnover rates, and thus low average densities, there are still some areas of the network experiencing a high enough actin density that a different motor binding regime can be observed. That might allow the diffusion model to learn the densities at which these shifts occur with relation to the conditioning variable and accurately reflect the curvature nonlinearity in the generated images. The slices also reflect the extent to which these fluctuations can be effective at predicting behavior, as the amount of information in the slices drops off with more extreme fluctuations.

4.5.1 *Minimal model of motor binding regimes*

If the diffusion model is picking up on system wide trends in the fluctuations to predict different regimes, then there is likely some characteristic underlying interaction giving rise to these trends. We consider a model for the number of bound motors, M_b , from a pool of M total motors, with the rate equation given by

$$\frac{dM_b}{dt} = -k_{\text{off}}M_b + P_{\text{fil}}(\xi)k_{\text{on}}(M - M_b) \quad (4.1)$$

where k_{on} and k_{off} are motor binding and unbinding rates, and P_{fil} is the probability that the motor will diffuse into a filament during its lifetime.

This probability is given by

$$P_{\text{fil}}(\xi) = \exp\left(-\sqrt{\frac{k_{\text{rem}}}{2D}} \xi\right) \quad (4.2)$$

where k_{rem} is the motor removal rate that determines motor lifetime, and D is the motor diffusion constant. P_{fil} varies based on the mesh size ξ of the actin network. Network mesh size is deeply connected to the details of the crosslinking within the system, which we are not considering in this model. However, in quasi-2D actin networks, the mesh size has been described as scaling with actin concentration generically as $[c_A]^{-1/2}$ [96, 97]. We will use $\xi = C_1/\sqrt{N_a - C_2}$, with $C_{1,2}$ as fitting parameters to tune the mesh size scaling, where C_1 sets the scale of the mesh sizes, corresponding to the mesh size at the minimum turnover considered, and C_2 is an offset related to the minimal actin density needed to create a mesh. However these variables are approximate, so we fit them to get the best agreement between the normalized curves of $\sqrt{M_b/N_a}$ and the simulated average curvature. N_a is the number of assembled actin monomers in the system, proportional to the actin density. N_a fluctuates independently of the number of bound motors, distributed as a compound Poisson process in the number of filaments N_f and filament length l ,

$$\begin{aligned} N_a &= \sum_{i=1}^{N_f} l_i \\ N_f &\sim \text{Poisson}(2k_{\text{nuc}}t_{\text{grow}}) \\ l_i &\sim \text{Uniform}(0, k_{\text{act}}t_{\text{grow}}) \end{aligned} \quad (4.3)$$

The distributions of N_f and l are a result of the specific actin growth dynamics in our system, depending on the filament nucleation rate k_{nuc} , the filament growth time t_{grow} , and

the turnover rate k_{act} . See SI for details.

We can then sample the joint fluctuations in N_a and M_b by sampling a value of N_a and using that to determine the Markov transition rates in motor binding, which we can use to sample $P(M_b|N_a)$. This model describes the fluctuations in assembled actin and bound motors in the system, but to compare to the joint distributions of the heatmap data we need to connect actin curvature to bound motor density. If we assume the filament bending energy can be approximated by a worm-like chain, and each motor makes a unit energy contribution to the bending of a local segment, the curvature is approximately proportional to $\sqrt{M_b/N_a}$.

While the joint distributions the model produces (Fig. 4.4c), do not fully capture the fluctuations of the simulation distributions, they do reflect a shift between 2 distinct motor binding regimes with increasing turnover. As the pool of available motors is exhausted, the fluctuations in the system become dominated by density fluctuations. The model may begin to diverge from the Cytosim simulation more at very high turnover rates as bundling effects, which were not considered in the model, become more impactful. This motor binding model shows that there is a minimal framework that captures the average system response to increasing turnover.

The feedback relationship from the limited pool of motors guides the core trend in average values. These trends can then be easily picked up by the diffusion model due to the strong linear relationship between density and turnover rate allowing the correlations of the joint distribution to reflect the slope of the trend in average curvature vs turnover. The range of fluctuations in the training data then determines the interpolation distance the model is able to extend these trends across. Much of the predictive power of this diffusion model comes from the fact that the conditioning variable for the model is closely related to the average of a fluctuating data channel.

4.6 Future Work

Expanding the diffusion model to predict sequences of images could allow us to more directly predict cortical flow. Applying this procedure to experimental actomyosin data would be the primary application to fully leverage the predictive power of these models and allow for explorations of parameter space with relatively fewer experimental points. This procedure could also be applied to characterize other systems. Given that long tailed distributions are common in biological data, developing a diffusion model with a long tailed prior distribution, as is discussed in [98, 99], could aid in fully capturing these distributions. Another possible method could be to use active or time correlated noise in the diffusion, which also serves to push values more toward the extremes of the distribution.

Additionally, an investigation of how well the model can infer trends across other control parameters could give more insight into the underlying system dynamics. Considering a parameter without the linear relationship to the data could disrupt the ability of the edges of the distributions to predict new parameter values. However, more spatially organized data could introduce new features for the model to learn relationships from.

4.7 Supplemental Information

4.7.1 *Cytosim Simulation*

The actomyosin network was simulated using the Cytosim simulation package. The simulation is implemented as described in [16] section S1A, with an increased number of motors (4800 initial motors instead of 1600). Actin filaments are nucleated in a $20 \times 32 \mu\text{m}$ box at a fixed rate and grow at the specified turnover rate for time $t_{\text{grow}} = 8.5\text{s}$ before disassembling. Motors are added to the simulation box at a low background rate of 30 s^{-1} . Then, to create the motor gradient driving the cortical flow, they are added at a high rate to the center $10 \mu\text{m}$ of the box at an additional rate of 180 s^{-1} . Motors are then removed uniformly through

the box. Crosslinkers are recycled uniformly across the simulation box. The simulations are run for 100 seconds, recording the system state every second. We do not consider the first 20 s of the simulation run to allow it to reach steady state.

To bin the simulation data into heatmap images, the simulation box is split into bins ($1\text{ }\mu\text{m}$ for the main results). The simulation models actin filaments as a sequence of $0.2\text{ }\mu\text{m}$ segments. The density is calculated as the number density of segments within the bin, while the curvature is calculated as the inverse radius of the circle through the centers of a given segment and its two neighboring segments.

	Actin	Myosin	Xlinker
nucleate	121 s^{-1}	—	—
add	$212\text{--}740\text{ s}^{-1}$	$30\text{--}210\text{ s}^{-1}$	50 s^{-1}
remove	$212\text{--}740\text{ s}^{-1}$	210 s^{-1}	50 s^{-1}
bind	—	5 s^{-1}	60 s^{-1}
diffusion	—	10	10
unbind	—	0.15 s^{-1}	0.05 s^{-1}
initial number	steady state	4800	44500

Table S4.1: Cytosim parameter values

4.7.2 Diffusion Model

The generative diffusion model is an SGM that learn a score function for the diffusion process on heatmap images. The equations for forward and reverse processes are shown in Fig 4.2 respectively given by

$$\begin{aligned}
d\mathbf{x}_t &= -\mathbf{x}_t dt + \sqrt{2}d\mathbf{B}_t \\
d\mathbf{x}_\tau &= [\mathbf{x}_\tau + 2\mathbf{S}(\mathbf{x}_\tau, c)]d\tau + \sqrt{2}d\mathbf{B}_\tau
\end{aligned}
\tag{S4.1}$$

where $\mathbf{S}(\mathbf{x}_\tau, c)$ is the class conditional score function and c is the turnover rate. The reverse process is solved during sampling using the Euler-Maruyama sampling scheme.

Architecture

The model uses an NCSN++ style U-Net. We use a base width of 32, with channel multipliers (1,2,4). This results in feature dims of (32, 64, 128) across 3 resolution levels created by 3 residual blocks. Timesteps are embedded into 256 dimensions with a 2 layer MLP, and injected into residual blocks. There are 3 encoder blocks with 3x3 stride-2-convolution, followed by a transformer bottleneck, and 3 decoder blocks with a nearest neighbor interpolation followed by a 3x3 conv. This structure means the bottleneck dimension is image size/4, which for the standard 16x20 image resolution primarily considered in this work results in a bottleneck of 4x5 resolution. Residual blocks use BatchNorm and SiLU activation.

To add conditioning to the model, we considered a few different methods to embed the class label (turnover rate). These included a randomly initialized MLP, mapping turnover rate to a circle, and simply using the class index. Once the label is embedded, it is added as additional channels to the image data, and input into the U-Net without any additional noise. The method of embedding did not result in any significant variation in generated image quality from the model. This is likely because the turnover rate already has an established numerical ordering, so we do not need to learn any semantic relationships between different classes. For all the models discussed in this work we use the circle embedding of $(\sin(c), \cos(c))$, where c is the range of turnover rates normalized to $[0, \pi/2]$, with that range selected to confine the embedding to positive values. We chose this to allow for a null token at (0,0) that is equidistant from all the turnover rates. This feature of the embedding is useful for classifier free guidance methods. The general idea of these methods is to combine scores from a conditional and an unconditional model to use both global and class specific features to more fully sample the space [38, 100]. These methods were detrimental for our

applications, likely because the predictive power of our model comes from the relationship of the data fluctuations to the control parameter, and the lack of class independent structures in the data make a purely conditional model preferable.

The model is trained on heatmap data from a varying number of turnover rates. To compare across models with varying numbers of seen turnover rates, we use 54,000 training frames evenly split between the rates.

4.7.3 Autoencoder Evaluation

To identify key features of the simulated heatmap data and develop a metric to compare the high dimensional distributions of heatmap data, we use a convolutional autoencoder to reduce the 16x20x2 heatmap images to a 40 dimensional latent vector. We use 3 convolutional downsampling layers for the encoder, and similarly 3 upsampling layers in the decoder, all with SiLU activations.

4.7.4 Model Variability

There is some variability in how fully the model learns the underlying distributions, especially for lower number of seen turnover rates. This is illustrated in Fig S4.1. The model does not converge to a single state even after extended training. Instead after some amount of initial training, the results begin fluctuating around some local minima. Different initializations of training show similar behaviors, but can find different local minima in the landscape. The effect is more pronounced with fewer seen turnover rates, suggesting a loss of detail in the distribution the model is able to learn with larger inference gaps in turnover. This variability in results can be counteracted by training ensembles of different models then selecting the best predictions of the seen turnover rates. This can be measured through the mean squared difference in average curvature and density or through an FID-like metric comparing the Gaussian fits of the simulated and generated distribution in a reduced latent space.




Figure S4.1: Illustration of model variability across training epochs




Figure S4.2: Comparison of model predictions with naive spline interpolation for 4 and 6 distinct seen turnover rates

4.7.5 Motor Binding Model

Filament length distribution

A filament grows at rate k_{act} until it reaches length $L \sim \text{Poisson}(k_{\text{act}} t_{\text{grow}})$, then decays at rate k_{act} . Assuming each assembly/disassembly step takes the same expected time, the conditional probability that a filament of maximum length n is in state l is:

$$P(l \mid L = n) = \frac{1}{n}. \quad (\text{S4.2})$$

Marginalizing over L (a filament of length l is only possible if $l \leq L$):

$$p(l) = \sum_{n=l}^{\infty} P(L = n) \frac{1}{n} = \sum_{n=l}^{\infty} \frac{\lambda^n e^{-\lambda}}{n!} \frac{1}{n} \quad (\text{S4.3})$$

High- λ (long filament) limit:

At high $\lambda = k_{\text{act}} t_{\text{grow}}$, end effects are negligible, and the distribution becomes approximately uniform:

$$p(l) \approx \frac{1}{\langle L \rangle} = \frac{1}{\lambda} \implies l \sim \text{Uniform}(0, k_{\text{act}} t_{\text{grow}}). \quad (\text{S4.4})$$

$$\mathbb{E}[l] = \frac{\langle L \rangle}{2} = \frac{k_{\text{act}} t_{\text{grow}}}{2} \quad (\text{S4.5})$$

$$\text{Var}(l) = \frac{1}{12} (k_{\text{act}} t_{\text{grow}})^2. \quad (\text{S4.6})$$

Compound Poisson Process for N_a

$$\lambda_{\text{act}} = k_{\text{act}} t_{\text{grow}} \quad \lambda_{\text{nuc}} = 2 k_{\text{nuc}} t_{\text{grow}}. \quad (\text{S4.7})$$

The number of filaments $N_f \sim \text{Poisson}(\lambda_{\text{nuc}})$

Total assembled actin, N_a

$$N_a = \sum_{i=1}^{N_f} l_i, \quad N_f \sim \text{Poisson}(\lambda_{\text{nuc}}), \quad l_i \sim \text{Uniform}(0, \lambda_{\text{act}}). \quad (\text{S4.8})$$

$$\begin{aligned} \mathbb{E}(N_a) &= \lambda_{\text{nuc}} \mathbb{E}(l) = \lambda_{\text{nuc}} \frac{\lambda_{\text{act}}}{2} = k_{\text{nuc}} k_{\text{act}} t_{\text{grow}}^2 \\ \text{Var}(N_a) &= \lambda_{\text{nuc}} \mathbb{E}(l^2) = \lambda_{\text{nuc}} \frac{\lambda_{\text{act}}^2}{3} = \frac{2}{3} k_{\text{nuc}} k_{\text{act}}^2 t_{\text{grow}}^3 \end{aligned} \quad (\text{S4.9})$$

Mean in actin density scales linearly with k_{act} .

Probability that a motor encounters a filament

Motors are added randomly to the simulation box, and undergo 2d Brownian diffusion. The displacement x at time t is distributed as $P(x, t) \sim \text{Normal}(0, 4Dt)$ The probability that the

motor reaches a filament at distance ξ within lifetime t is:

$$P(x > \xi | t) = 2 \left(1 - \frac{1}{2} \left(1 + \operatorname{erf} \left(\frac{\xi}{\sigma\sqrt{2}} \right) \right) \right) = 1 - \operatorname{erf} \left(\frac{\xi}{\sqrt{8Dt}} \right). \quad (\text{S4.10})$$

The motor lifetime is exponentially distributed with removal rate $k_{\text{rem}} = 210 \text{ s}^{-1}$:

$$P(t) = k_{\text{rem}} e^{-k_{\text{rem}} t}. \quad (\text{S4.11})$$

Marginalizing over the motor lifetime:

$$P(x > \xi) = P_{\text{fil}}(\xi) = \int_0^\infty \left(1 - \operatorname{erf} \left(\frac{\xi}{\sqrt{8Dt}} \right) \right) k_{\text{rem}} e^{-k_{\text{rem}} t} dt \quad (\text{S4.12})$$

$$= \exp \left(-\sqrt{\frac{k_{\text{rem}}}{2D}} \xi \right) \quad (\text{S4.13})$$

The distance to encounter a filament depends on the network mesh size. Average mesh size scales inversely with actin concentration, approximately as $1/\sqrt{c_{\text{actin}}}$

$$\xi = C_1 / \sqrt{N_a - C_2} \quad (\text{S4.14})$$

where C_1 and C_2 are fitting constants to tune the mesh size scaling.

Motor Binding Rate Equation and Markov Chain

Let M_b = number of bound motors out of total M . The mean-field rate equation is:

$$\frac{dM_b}{dt} = -k_{\text{off}} M_b + P_{\text{fil}}(\xi) k_{\text{on}} (M - M_b). \quad (\text{S4.15})$$

At steady state:

$$M_b^{ss} = \frac{M}{\frac{k_{\text{off}}}{P_{\text{fil}} k_{\text{on}}} + 1}. \quad (\text{S4.16})$$

To capture fluctuations (rather than just using the average N_s), a Markov chain for the full distribution $P(M_b)$ is used:

$$\begin{aligned} \frac{dP(M_b)}{dt} = & -(P_{\text{fil}} k_{\text{on}}(M - M_b) + k_{\text{off}} M_b) P(M_b) \\ & + P_{\text{fil}} k_{\text{on}}(M - M_b + 1) P(M_b - 1) + k_{\text{off}} (M_b + 1) P(M_b + 1). \end{aligned} \quad (\text{S4.17})$$

CHAPTER 5

CONCLUSION

5.1 Discussion

In this work we explored active matter systems through a range of complementary modeling perspectives. These approaches range in their levels of complexity and imposed structure, and they exercise distinct modes of inference: bounding a region of state space, predicting observables from new inputs, and sampling from a learned distribution. In each system a combination of modeling approaches predicted the behavior while connecting it to a minimal underlying mechanism.

A thermodynamic bound on branched actin networks showed that their adaptive force response carries a dissipative cost beyond the work of moving the load, while the full Markov model supplied the configurational trends the bound constrains. In the chemical reaction-diffusion system, a learned latent space identified the collective variables and located the boundaries between wave regimes, while a minimal one-dimensional model gave the speed and decay of the wave. In the actomyosin cortex, a generative diffusion model inferred a nonlinear mesoscale trend across filament turnover rates it was never trained on, while a motor binding model gave the mechanism behind the nonlinearity. No single one of these perspectives captures both the specific behavior and the general principle that governs it. The combination does.

The active nature of these systems gives rise to characteristics that recur across the works. Each system is multiscale, and the balance between its microscopic interactions and its mesoscale response sets the length and time scales that a model can predict. The imposed driving that places a system out of equilibrium also takes a concrete and parallel form. It appears as the driving chemical potential from the background actin pool in branched networks, and as the chemical gradient maintained by the flow reactor in the reaction-

diffusion waves. Both are imposed thermodynamic forces of the same kind, expressed through very different chemistry.

The representation of the data is itself part of the model, and it sets the scale of the questions a model can answer. The autoencoder latent space and the coarse-grained density maps are both choices of representation that fix the dimensionality of the learnable problem. A coarser representation gives a lower-dimensional problem that can be learned from sparse data, which is what allowed the diffusion model to extrapolate beyond its training conditions. A finer representation resolves more structure but raises the dimensionality that has to be learned. Where to set this scale is not obvious, and it depends on the level of detail each model can resolve.

5.2 Future Work

The perspectives developed here suggest directions both within each system and across them. In branched network assembly, the most direct extension is to time varying and spatially heterogeneous loads. This would introduce spatial coupling between different regions of the network, and could extend the thermodynamic bounding to more complex systems. The thermodynamic description of branched network growth could be coupled with driven membrane systems, such as ones described in [101]. The inverse problem is also worth considering: inferring the loading conditions a network has experienced from its measured structure, which the model’s greater sensitivity to load than to background rates should favor.

In chemical reaction networks, many potential control variables remain unexplored. The latent space analysis could be leveraged to identify new control parameters and governing principles of the traveling waves, which would be especially useful in predicting changes in wave behavior across reactor geometries. Further development of the 1D wave model could help identify the control parameters governing wave speed and the phase behavior.

For the diffusion work, further development is needed to connect the interpolation of curvature trends to the resulting cortical flows. This would require training the model with sequences of images to pick up dynamics or developing an external predictor of cortical flow from the heatmap data. Another complementary avenue is investigating the effect of the scale of coarse-graining on the predictive ability of the model. At finer scales more of the spatial structure of the network begins to emerge, and analysis of the optimal coarsening could further elucidate the factors driving cortical flow. It would also be relevant to efforts in analyzing experimental data and inform microscopy with the relationship between coarse-graining and the point spread function. Additionally, given the long tailed distribution of the curvature data, a diffusion model with a long tailed prior might capture the full distribution better. That could further probe the relationship between prior distribution and generation. This approach is, however, significantly complicated by the need to analytically calculate the forward process, limiting the possible distributions.

In a broader scope, there are connections between the thermodynamic perspectives and generative approaches we have considered. The thermodynamics of the diffusion process itself has been studied, including bounds relating the accuracy of generation to the entropy production of the dynamics [102] and reformulations of the process in the language of statistical mechanics [103]. These treat the generative dynamics as a thermodynamic object. A less explored variant would instead derive a thermodynamic bound on the accessible structures of the modeled system and use it to guide the reverse process toward feasible configurations. Operationally this resembles guiding generation toward a tilted target distribution, as in classifier-free guidance [38], with the tilt set by a physical bound rather than a learned classifier. The branched network bound on available structures is a natural candidate. We do not expect this to be a departure from existing thermodynamically informed generative models, but a specific case that ties a derived structural bound to a generative model of the same system.

REFERENCES

- [1] M. C. Marchetti, J. F. Joanny, S. Ramaswamy, T. B. Liverpool, J. Prost, Madan Rao, and R. Aditi Simha. Hydrodynamics of soft active matter. *Reviews of Modern Physics*, 85(3):1143–1189, July 2013. ISSN 0034-6861, 1539-0756. doi:10.1103/RevModPhys.85.1143. URL <https://link.aps.org/doi/10.1103/RevModPhys.85.1143>.
- [2] Daniel Needleman and Zvonimir Dogic. Active matter at the interface between materials science and cell biology. *Nature Reviews Materials*, 2(9):17048, July 2017. ISSN 2058-8437. doi:10.1038/natrevmats.2017.48. URL <https://www.nature.com/articles/natrevmats201748>.
- [3] Thomas D. Pollard and Gary G. Borisy. Cellular Motility Driven by Assembly and Disassembly of Actin Filaments. *Cell*, 112(4):453–465, February 2003. ISSN 0092-8674, 1097-4172. doi:10.1016/S0092-8674(03)00120-X. URL [https://www.cell.com/cell/abstract/S0092-8674\(03\)00120-X](https://www.cell.com/cell/abstract/S0092-8674(03)00120-X).
- [4] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the National Academy of Sciences*, 105(4):1232–1237, January 2008. doi:10.1073/pnas.0711437105. URL <https://www.pnas.org/doi/10.1073/pnas.0711437105>.
- [5] Clemens Bechinger, Roberto Di Leonardo, Hartmut Löwen, Charles Reichhardt, Giorgio Volpe, and Giovanni Volpe. Active particles in complex and crowded environments. *Reviews of Modern Physics*, 88(4):045006, November 2016. doi:10.1103/RevModPhys.88.045006. URL <https://link.aps.org/doi/10.1103/RevModPhys.88.045006>.
- [6] Arpita Paikar, Xiuxiu Li, Liat Avram, Barbara S. Smith, István Sütő, Dezső Horváth, Elisabeth Rennert, Yuqing Qiu, Ágota Tóth, Suriyanarayanan Vaikuntanathan, and Sergey N. Semenov. Chemical waves in reaction-diffusion networks of small organic molecules. *Chemical Science*, 16(2):659–669, 2025. doi:10.1039/D4SC06351A. URL <https://pubs.rsc.org/en/content/articlelanding/2025/sc/d4sc06351a>.
- [7] Daniel A. Fletcher and Phillip L. Geissler. Active Biological Materials. *Annual Review of Physical Chemistry*, 60(1):469–486, May 2009. ISSN 0066-426X, 1545-1593. doi:10.1146/annurev.physchem.040808.090304. URL <https://www.annualreviews.org/doi/10.1146/annurev.physchem.040808.090304>.
- [8] Mark K. Transtrum, Benjamin B. Machta, Kevin S. Brown, Bryan C. Daniels, Christopher R. Myers, and James P. Sethna. Perspective: Sloppiness and emergent theories in physics, biology, and beyond. *The Journal of Chemical Physics*, 143(1):010901, July 2015. ISSN 0021-9606, 1089-7690. doi:10.1063/1.4923066. URL

<https://pubs.aip.org/jcp/article/143/1/010901/566995/Perspective-Sloppiness-and-emergent-theories-in>.

- [9] Benjamin B. Machta, Ricky Chachra, Mark K. Transtrum, and James P. Sethna. Parameter Space Compression Underlies Emergent Theories and Predictive Models. *Science*, 342(6158):604–607, November 2013. doi:10.1126/science.1238723. URL <http://www.science.org/doi/full/10.1126/science.1238723>.
- [10] Ryan N. Gutenkunst, Joshua J. Waterfall, Fergal P. Casey, Kevin S. Brown, Christopher R. Myers, and James P. Sethna. Universally Sloppy Parameter Sensitivities in Systems Biology Models. *PLOS Computational Biology*, 3(10):e189, October 2007. ISSN 1553-7358. doi:10.1371/journal.pcbi.0030189. URL <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.0030189>.
- [11] Alexander Holiday, Mahdi Kooshkbaghi, Juan M. Bello-Rivas, C. William Gear, Antonios Zagaris, and Ioannis G. Kevrekidis. Manifold learning for parameter reduction. *Journal of Computational Physics*, 392:419–431, September 2019. ISSN 00219991. doi:10.1016/j.jcp.2019.04.015. URL <https://linkinghub.elsevier.com/retrieve/pii/S0021999119302487>.
- [12] 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. *Journal of Computational Physics*, 378:686–707, February 2019. ISSN 0021-9991. doi:10.1016/j.jcp.2018.10.045. URL <https://www.sciencedirect.com/science/article/pii/S0021999118307125>.
- [13] Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, April 2016. doi:10.1073/pnas.1517384113. URL <https://www.pnas.org/doi/10.1073/pnas.1517384113>.
- [14] Matthew S. Schmitt, Maciej Koch-Janusz, Michel Fruchart, Daniel S. Seara, Michael Rust, and Vincenzo Vitelli. Information theory for data-driven model reduction in physics and biology, April 2024. URL <http://arxiv.org/abs/2312.06608>. arXiv:2312.06608 [cond-mat].
- [15] Rohit Supekar, Boya Song, Alasdair Hastewell, Gary P. T. Choi, Alexander Mietke, and Jörn Dunkel. Learning hydrodynamic equations for active matter from particle simulations and experiments. *Proceedings of the National Academy of Sciences*, 120(7):e2206994120, February 2023. doi:10.1073/pnas.2206994120. URL <https://www.pnas.org/doi/abs/10.1073/pnas.2206994120>.
- [16] Yuqing Qiu, Elizabeth D. White, Edwin M. Munro, Suriyanarayanan Vaikuntanathan, and Aaron R. Dinner. Elucidating the Role of Filament Turnover in Cortical Flow using Simulations and Representation Learning, October 2023. URL <http://arxiv.org/abs/2310.10819>. arXiv:2310.10819 [cond-mat].

- [17] Edwin Munro, Jeremy Nance, and James R. Priess. Cortical Flows Powered by Asymmetrical Contraction Transport PAR Proteins to Establish and Maintain Anterior-Posterior Polarity in the Early *C. elegans* Embryo. *Developmental Cell*, 7(3):413–424, September 2004. ISSN 1534-5807. doi:10.1016/j.devcel.2004.08.001. URL [https://www.cell.com/developmental-cell/abstract/S1534-5807\(04\)00276-X](https://www.cell.com/developmental-cell/abstract/S1534-5807(04)00276-X).
- [18] Zeno Messi, Franck Raynaud, Nathan Goehring, and Alexander B. Verkhovsky. Balanced contractility and adhesion drive polarization in a minimal elastic actomyosin network, August 2025. URL <https://www.biorxiv.org/content/10.1101/2025.08.26.672435v1>. ISSN: 2692-8205 Pages: 2025.08.26.672435 Section: New Results.
- [19] Deb Sankar Banerjee, Akankshi Munjal, Thomas Lecuit, and Madan Rao. Actomyosin pulsation and flows in an active elastomer with turnover and network remodeling. *Nature Communications*, 8(1):1121, October 2017. ISSN 2041-1723. doi:10.1038/s41467-017-01130-1. URL <https://www.nature.com/articles/s41467-017-01130-1>.
- [20] Ashwini Krishna, Mariya Savinov, Niv Ierushalmi, Alex Mogilner, and Kinneret Keren. Size-dependent transition from steady contraction to waves in actomyosin networks with turnover. *Nature Physics*, 20(1):123–134, January 2024. ISSN 1745-2481. doi:10.1038/s41567-023-02271-5. URL <https://www.nature.com/articles/s41567-023-02271-5>.
- [21] Maya Malik-Garbi, Niv Ierushalmi, Silvia Jansen, Enas Abu-Shah, Bruce L. Goode, Alex Mogilner, and Kinneret Keren. Scaling behaviour in steady-state contracting actomyosin networks. *Nature Physics*, 15(5):509–516, May 2019. ISSN 1745-2473, 1745-2481. doi:10.1038/s41567-018-0413-4. URL <https://www.nature.com/articles/s41567-018-0413-4>.
- [22] Elizabeth D. White, Younan Li, Rachel Kadzik, Yuqing Qiu, Aaron R. Dinner, and Edwin M. Munro. Microscopic Control of Cortical Flows in Polarized *C. elegans* Zygotes, January 2025. URL <https://www.biorxiv.org/content/10.1101/2025.01.12.632635v1>. Pages: 2025.01.12.632635 Section: New Results.
- [23] Francois Nedelec and Dietrich Foethke. Collective Langevin dynamics of flexible cytoskeletal fibers. *New Journal of Physics*, 9(11):427, November 2007. ISSN 1367-2630. doi:10.1088/1367-2630/9/11/427. URL <https://dx.doi.org/10.1088/1367-2630/9/11/427>.
- [24] Konstantin Popov, James Komianos, and Garegin A. Papoian. MEDYAN: Mechanochemical Simulations of Contraction and Polarity Alignment in Actomyosin Networks. *PLOS Computational Biology*, 12(4):e1004877, April 2016. ISSN 1553-7358. doi:10.1371/journal.pcbi.1004877. URL <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1004877>.
- [25] Andre C. Barato and Udo Seifert. Thermodynamic Uncertainty Relation for Biomolecular Processes. *Physical Review Letters*, 114(15):158101, April 2015.

- doi:10.1103/PhysRevLett.114.158101. URL <https://link.aps.org/doi/10.1103/PhysRevLett.114.158101>.
- [26] Michael Nguyen, Yuqing Qiu, and Suriyanarayanan Vaikuntanathan. Organization and Self-Assembly Away from Equilibrium: Toward Thermodynamic Design Principles. *Annual Review of Condensed Matter Physics*, 12(1):273–290, March 2021. ISSN 1947-5454, 1947-5462. doi:10.1146/annurev-conmatphys-031218-013309. URL <https://www.annualreviews.org/doi/10.1146/annurev-conmatphys-031218-013309>.
 - [27] Leland McInnes, John Healy, and James Melville. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. September 2020. doi:10.48550/arXiv.1802.03426. URL <http://arxiv.org/abs/1802.03426>. arXiv:1802.03426 [stat.ML].
 - [28] Laurens van der Maaten and Geoffrey Hinton. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008. URL <http://jmlr.org/papers/v9/vandermaaten08a.html>.
 - [29] G. E. Hinton and R. R. Salakhutdinov. Reducing the Dimensionality of Data with Neural Networks. *Science*, 313(5786):504–507, July 2006. doi:10.1126/science.1127647. URL <https://www.science.org/doi/10.1126/science.1127647>.
 - [30] S. S. Schoenholz, E. D. Cubuk, D. M. Sussman, E. Kaxiras, and A. J. Liu. A structural approach to relaxation in glassy liquids. *Nature Physics*, 12(5):469–471, May 2016. ISSN 1745-2481. doi:10.1038/nphys3644. URL <https://www.nature.com/articles/nphys3644>. Number: 5.
 - [31] Yusuf Dalva and Pinar Yanardag. NoiseCLR: A Contrastive Learning Approach for Unsupervised Discovery of Interpretable Directions in Diffusion Models, December 2023. URL <http://arxiv.org/abs/2312.05390>. arXiv:2312.05390 [cs].
 - [32] Rabea Seyboldt, Juliette Lavoie, Adrien Henry, Jules Vanaret, Mariela D. Petkova, Thomas Gregor, and Paul François. Latent space of a small genetic network: Geometry of dynamics and information. *Proceedings of the National Academy of Sciences*, 119(26):e2113651119, June 2022. doi:10.1073/pnas.2113651119. URL <https://www.pnas.org/doi/10.1073/pnas.2113651119>.
 - [33] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. *Advances in Neural Information Processing Systems*, 2020. URL <https://proceedings.neurips.cc/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf>.
 - [34] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations, February 2021. URL <http://arxiv.org/abs/2011.13456>. arXiv:2011.13456 [cs].

- [35] Prafulla Dhariwal and Alex Nichol. Diffusion Models Beat GANs on Image Synthesis.
- [36] Jascha Sohl-Dickstein, Eric A Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. *Proceedings of the 32nd International Conference on Machine Learning*, 2015. URL <https://proceedings.mlr.press/v37/sohl-dickstein15.pdf>.
- [37] Alex Nichol and Prafulla Dhariwal. Improved Denoising Diffusion Probabilistic Models, February 2021. URL <http://arxiv.org/abs/2102.09672>. arXiv:2102.09672 [cs.LG].
- [38] Jonathan Ho and Tim Salimans. Classifier-Free Diffusion Guidance, July 2022. URL <http://arxiv.org/abs/2207.12598>. arXiv:2207.12598 [cs].
- [39] Elisabeth Rennert and Suriyanarayanan Vaikuntanathan. Thermodynamic framework for nonequilibrium self-assembly of branched actin networks. *bioRxiv*, page 2023.11.19.567734, January 2023. doi:10.1101/2023.11.19.567734. URL <http://biorxiv.org/content/early/2023/11/19/2023.11.19.567734.abstract>.
- [40] Elisabeth Rennert, Agnish Kumar Behera, Yuqing Qiu, and Suriyanarayanan Vaikuntanathan. Diffusion models learn underlying trends in actomyosin networks and predict behavior at unseen filament turnover, July 2026. URL <https://www.biorxiv.org/content/10.64898/2026.05.26.727950v2>. ISSN: 2692-8205 Pages: 2026.05.26.727950 Section: New Results.
- [41] Alba Diz-Muñoz, Michael Krieg, Martin Bergert, Itziar Ibarlucea-Benitez, Daniel J. Muller, Ewa Paluch, and Carl-Philipp Heisenberg. Control of Directed Cell Migration In Vivo by Membrane-to-Cortex Attachment. *PLOS Biology*, 8(11):e1000544, November 2010. ISSN 1545-7885. doi:10.1371/journal.pbio.1000544. URL <https://journals.plos.org/plosbiology/article?id=10.1371/journal.pbio.1000544>.
- [42] Xindong Chen, Yuhui Li, Ming Guo, Bowen Xu, Yanhui Ma, Hanxing Zhu, and Xi-Qiao Feng. Polymerization force-regulated actin filament-Arp2/3 complex interaction dominates self-adaptive cell migrations, April 2023. URL <https://www.biorxiv.org/content/10.1101/2023.04.15.536869v2>. Pages: 2023.04.15.536869 Section: New Results.
- [43] Thomas D. Pollard and John A. Cooper. Actin, a Central Player in Cell Shape and Movement. *Science*, 326(5957):1208–1212, November 2009. ISSN 0036-8075, 1095-9203. doi:10.1126/science.1175862. URL <https://www.science.org/doi/10.1126/science.1175862>.
- [44] Jessica Planade, Reda Belbahri, Micaela Boiero Sanders, Audrey Guillotin, Olivia du Roure, Alphée Michelot, and Julien Heuvingh. Mechanical stiffness of reconstituted actin patches correlates tightly with endocytosis efficiency. *PLOS Biology*, 17(10):e3000500, October 2019. ISSN 1545-7885. doi:10.1371/journal.pbio.3000500. URL <https://journals.plos.org/plosbiology/article?id=10.1371/journal.pbio.3000500>.

- [45] Peter Bieling, Tai-De Li, Julian Weichsel, Ryan McGorty, Pamela Jreij, Bo Huang, Daniel A. Fletcher, and R. Dyche Mullins. Force Feedback Controls Motor Activity and Mechanical Properties of Self-Assembling Branched Actin Networks. *Cell*, 164(1-2):115–127, January 2016. ISSN 1097-4172. doi:10.1016/j.cell.2015.11.057.
- [46] Daniel A. Fletcher and R. Dyche Mullins. Cell mechanics and the cytoskeleton. *Nature*, 463(7280):485–492, January 2010. ISSN 1476-4687. doi:10.1038/nature08908. URL <https://www.nature.com/articles/nature08908>. Number: 7280.
- [47] Peter Bieling and Klemens Rottner. From WRC to Arp2/3: Collective molecular mechanisms of branched actin network assembly. *Current Opinion in Cell Biology*, 80:102156, February 2023. ISSN 09550674. doi:10.1016/j.ceb.2023.102156. URL <https://linkinghub.elsevier.com/retrieve/pii/S0955067423000054>.
- [48] Julie A. Theriot, Timothy J. Mitchison, Lewis G. Tilney, and Daniel A. Portnoy. The rate of actin-based motility of intracellular *Listeria monocytogenes* equals the rate of actin polymerization. *Nature*, 357(6375):257–260, May 1992. ISSN 1476-4687. doi:10.1038/357257a0. URL <https://www.nature.com/articles/357257a0>. Number: 6375.
- [49] Tai-De Li, Peter Bieling, Julian Weichsel, R Dyche Mullins, and Daniel A Fletcher. The molecular mechanism of load adaptation by branched actin networks. *eLife*, 11:e73145, June 2022. ISSN 2050-084X. doi:10.7554/eLife.73145. URL <https://doi.org/10.7554/eLife.73145>.
- [50] Thomas Pujol, Olivia du Roure, Marc Fermigier, and Julien Heuvingh. Impact of branching on the elasticity of actin networks. *Proceedings of the National Academy of Sciences of the United States of America*, 109(26):10364–10369, June 2012. ISSN 0027-8424. doi:10.1073/pnas.1121238109. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3387051/>.
- [51] Sapun H. Parekh, Ovijit Chaudhuri, Julie A. Theriot, and Daniel A. Fletcher. Loading history determines the velocity of actin-network growth. *Nature Cell Biology*, 7(12):1219–1223, December 2005. ISSN 1476-4679. doi:10.1038/ncb1336. URL <https://www.nature.com/articles/ncb1336>. Number: 12.
- [52] Marcus Prass, Ken Jacobson, Alex Mogilner, and Manfred Radmacher. Direct measurement of the lamellipodial protrusive force in a migrating cell. *Journal of Cell Biology*, 174(6):767–772, September 2006. ISSN 0021-9525. doi:10.1083/jcb.200601159. URL <https://doi.org/10.1083/jcb.200601159>.
- [53] Ovijit Chaudhuri, Sapun H. Parekh, and Daniel A. Fletcher. Reversible stress softening of actin networks. *Nature*, 445(7125):295–298, January 2007. ISSN 1476-4687. doi:10.1038/nature05459. URL <https://www.nature.com/articles/nature05459>. Number: 7125.

- [54] Alexander B. Verkhovsky, Oleg Y. Chaga, Sébastien Schaub, Tatyana M. Svitkina, Jean-Jacques Meister, and Gary G. Borisy. Orientational Order of the Lamellipodial Actin Network as Demonstrated in Living Motile Cells. *Molecular Biology of the Cell*, 14(11):4667–4675, November 2003. ISSN 1059-1524. doi:10.1091/mbc.e02-10-0630. URL <https://www.molbiolcell.org/doi/full/10.1091/mbc.e02-10-0630>.
- [55] Jan Mueller, Gregory Szep, Maria Nemethova, Ingrid De Vries, Arnon D. Lieber, Christoph Winkler, Karsten Kruse, J. Victor Small, Christian Schmeiser, Kinneret Keren, Robert Hauschild, and Michael Sixt. Load Adaptation of Lamellipodial Actin Networks. *Cell*, 171(1):188–200.e16, September 2017. ISSN 00928674. doi:10.1016/j.cell.2017.07.051. URL <https://linkinghub.elsevier.com/retrieve/pii/S0092867417308887>.
- [56] Julian Weichsel and Ulrich S. Schwarz. Two competing orientation patterns explain experimentally observed anomalies in growing actin networks. *Proceedings of the National Academy of Sciences*, 107(14):6304–6309, April 2010. doi:10.1073/pnas.0913730107. URL <https://www.pnas.org/doi/abs/10.1073/pnas.0913730107>.
- [57] Matthias Krause and Alexis Gautreau. Steering cell migration: lamellipodium dynamics and the regulation of directional persistence. *Nature Reviews Molecular Cell Biology*, 15(9):577–590, September 2014. ISSN 1471-0080. doi:10.1038/nrm3861. URL <https://www.nature.com/articles/nrm3861>. Number: 9.
- [58] Anthony J. Wood, Richard A. Blythe, and Martin R. Evans. Solvable model of a many-filament Brownian ratchet. *Physical Review E*, 100(4):042122, October 2019. ISSN 2470-0045, 2470-0053. doi:10.1103/PhysRevE.100.042122. URL <https://link.aps.org/doi/10.1103/PhysRevE.100.042122>.
- [59] Shriram Chennakesavalu, Sreekanth K. Manikandan, Frank Hu, and Grant M. Rotsekoff. Adaptive nonequilibrium design of actin-based metamaterials: fundamental and practical limits of control, July 2023. URL <http://arxiv.org/abs/2306.10778>. arXiv:2306.10778 [cond-mat].
- [60] Ganhui Lan, Pablo Sartori, Silke Neumann, Victor Sourjik, and Yuhai Tu. The energy–speed–accuracy trade-off in sensory adaptation. *Nature Physics*, 8(5):422–428, May 2012. ISSN 1745-2473, 1745-2481. doi:10.1038/nphys2276. URL <https://www.nature.com/articles/nphys2276>.
- [61] Michael Nguyen and Suriyanarayanan Vaikuntanathan. Design principles for nonequilibrium self-assembly. *Proceedings of the National Academy of Sciences*, 113(50):14231–14236, December 2016. ISSN 0027-8424, 1091-6490. doi:10.1073/pnas.1609983113. URL <https://pnas.org/doi/full/10.1073/pnas.1609983113>.
- [62] Yuqing Qiu, Michael Nguyen, Glen M. Hocky, Aaron R. Dinner, and Suriyanarayanan Vaikuntanathan. A strong nonequilibrium bound for sorting of cross-linkers on growing biopolymers. *Proceedings of the National Academy of Sciences*, 118(38):e2102881118,

- September 2021. ISSN 0027-8424, 1091-6490. doi:10.1073/pnas.2102881118. URL <https://pnas.org/doi/full/10.1073/pnas.2102881118>.
- [63] R. Dyche Mullins, Peter Bieling, and Daniel A. Fletcher. From solution to surface to filament: actin flux into branched networks. *Biophysical Reviews*, 10(6):1537–1551, December 2018. ISSN 1867-2469. doi:10.1007/s12551-018-0469-5. URL <https://doi.org/10.1007/s12551-018-0469-5>.
 - [64] Thomas P. Loisel, Rajaa Boujemaa, Dominique Pantaloni, and Marie-France Carlier. Reconstitution of actin-based motility of *Listeria* and *Shigella* using pure proteins. *Nature*, 401(6753):613–616, October 1999. ISSN 0028-0836, 1476-4687. doi:10.1038/44183. URL <https://www.nature.com/articles/44183>.
 - [65] Johanna Funk, Felipe Merino, Matthias Schaks, Klemens Rottner, Stefan Raunser, and Peter Bieling. A barbed end interference mechanism reveals how capping protein promotes nucleation in branched actin networks. *Nature Communications*, 12(1):5329, September 2021. ISSN 2041-1723. doi:10.1038/s41467-021-25682-5. URL <https://www.nature.com/articles/s41467-021-25682-5>. Number: 1.
 - [66] Orkun Akin and R. Dyche Mullins. Capping Protein Increases the Rate of Actin-Based Motility by Promoting Filament Nucleation by the Arp2/3 Complex. *Cell*, 133(5):841–851, May 2008. ISSN 0092-8674, 1097-4172. doi:10.1016/j.cell.2008.04.011. URL [https://www.cell.com/cell/abstract/S0092-8674\(08\)00498-4](https://www.cell.com/cell/abstract/S0092-8674(08)00498-4).
 - [67] A. Mogilner and G. Oster. Cell motility driven by actin polymerization. *Biophysical Journal*, 71(6):3030–3045, December 1996. ISSN 00063495. doi:10.1016/S0006-3495(96)79496-1. URL <https://linkinghub.elsevier.com/retrieve/pii/S0006349596794961>.
 - [68] C.S. Peskin, G.M. Odell, and G.F. Oster. Cellular motions and thermal fluctuations: the Brownian ratchet. *Biophysical Journal*, 65(1):316–324, July 1993. ISSN 00063495. doi:10.1016/S0006-3495(93)81035-X. URL <https://linkinghub.elsevier.com/retrieve/pii/S000634959381035X>.
 - [69] T D Pollard. Rate constants for the reactions of ATP- and ADP-actin with the ends of actin filaments. *Journal of Cell Biology*, 103(6):2747–2754, December 1986. ISSN 0021-9525. doi:10.1083/jcb.103.6.2747. URL <https://doi.org/10.1083/jcb.103.6.2747>.
 - [70] Martin A. Wear, Atsuko Yamashita, Kyoungtae Kim, Yuichiro Maéda, and John A. Cooper. How Capping Protein Binds the Barbed End of the Actin Filament. *Current Biology*, 13(17):1531–1537, September 2003. ISSN 0960-9822. doi:10.1016/S0960-9822(03)00559-1. URL <https://www.sciencedirect.com/science/article/pii/S0960982203005591>.
 - [71] R. Dyche Mullins, John A. Heuser, and Thomas D. Pollard. The interaction of Arp2/3 complex with actin: Nucleation, high affinity pointed end capping, and formation of

- branching networks of filaments. *Proceedings of the National Academy of Sciences*, 95(11):6181–6186, May 1998. doi:10.1073/pnas.95.11.6181. URL <https://www.pnas.org/doi/abs/10.1073/pnas.95.11.6181>.
- [72] Arpita Paikar, Alexander I. Novichkov, Anton I. Hanopolskyi, Viktoryia A. Smaliak, Xiaomeng Sui, Nir Kampf, Ekaterina V. Skorb, and Sergey N. Semenov. Spatiotemporal Regulation of Hydrogel Actuators by Autocatalytic Reaction Networks. *Advanced Materials*, 34(13):2106816, April 2022. ISSN 0935-9648. doi:10.1002/adma.202106816. URL <https://onlinelibrary.wiley.com/doi/full/10.1002/adma.202106816>.
- [73] Junhai Zhai, Sufang Zhang, Junfen Chen, and Qiang He. Autoencoder and Its Various Variants. In *2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 415–419, October 2018. doi:10.1109/SMC.2018.00080. URL <https://ieeexplore.ieee.org/abstract/document/8616075/references#references>.
- [74] Michael Tschannen, Olivier Bachem, and Mario Lucic. Recent Advances in Autoencoder-Based Representation Learning, December 2018. URL <http://arxiv.org/abs/1812.05069>. arXiv:1812.05069 [cs, stat].
- [75] Jonathan Colen, Alexis Poncet, Denis Bartolo, and Vincenzo Vitelli. Interpreting neural operators: how nonlinear waves propagate in non-reciprocal solids, April 2024. URL <http://arxiv.org/abs/2404.12918>. arXiv:2404.12918 [cond-mat, physics:physics].
- [76] Thomas Gaskin, Grigorios A. Pavliotis, and Mark Girolami. Neural parameter calibration for large-scale multiagent models. *Proceedings of the National Academy of Sciences*, 120(7):e2216415120, February 2023. doi:10.1073/pnas.2216415120. URL <https://www.pnas.org/doi/abs/10.1073/pnas.2216415120>.
- [77] Martin J. Falk, Finnegan D. Roach, William Gilpin, and Arvind Murugan. Curiosity-driven search for novel non-equilibrium behaviors, August 2023. URL <http://arxiv.org/abs/2211.02589>. arXiv:2211.02589 [cond-mat].
- [78] Chengyu Dai and Sharon C. Glotzer. Efficient Phase Diagram Sampling by Active Learning. *The Journal of Physical Chemistry B*, 124(7):1275–1284, February 2020. ISSN 1520-6106. doi:10.1021/acs.jpcb.9b09202. URL <https://doi.org/10.1021/acs.jpcb.9b09202>.
- [79] Ekin D. Cubuk, Samuel S. Schoenholz, Efthimios Kaxiras, and Andrea J. Liu. Structural Properties of Defects in Glassy Liquids. *The Journal of Physical Chemistry B*, 120(26):6139–6146, July 2016. ISSN 1520-6106. doi:10.1021/acs.jpcb.6b02144. URL <https://doi.org/10.1021/acs.jpcb.6b02144>.
- [80] Thomas D. Pollard. Actin and Actin-Binding Proteins. *Cold Spring Harbor Perspectives in Biology*, 8(8):a018226, August 2016. ISSN , 1943-0264. doi:10.1101/cshperspect.a018226. URL <http://cshperspectives.cshlp.org/content/8/8/a018226>.

- [81] Peter J. Foster, Alexandra Zampetaki, Jinghui Liu, Sebastian Fürthauer, and Nikta Fakhri. Active mechanics of sea star oocytes, April 2026. URL <https://www.biorxiv.org/content/10.1101/2022.04.22.489189v2>. ISSN: 2692-8205 Pages: 2022.04.22.489189 Section: New Results.
- [82] Jörg Renkawitz, Kathrin Schumann, Michele Weber, Tim Lämmermann, Holger Pflücke, Matthieu Piel, Julien Polleux, Joachim P. Spatz, and Michael Sixt. Adaptive force transmission in amoeboid cell migration. *Nature Cell Biology*, 11(12): 1438–1443, December 2009. ISSN 1476-4679. doi:10.1038/ncb1992. URL <https://www.nature.com/articles/ncb1992>.
- [83] Paolo Maiuri, Jean-François Rupprecht, Stefan Wieser, Verena Rupprecht, Olivier Bénichou, Nicolas Carpi, Mathieu Coppey, Simon De Beco, Nir Gov, Carl-Philipp Heisenberg, Carolina Lage Crespo, Franziska Lautenschlaeger, Maël Le Berre, Ana-Maria Lennon-Dumenil, Matthew Raab, Hawa-Racine Thiam, Matthieu Piel, Michael Sixt, and Raphaël Voituriez. Actin Flows Mediate a Universal Coupling between Cell Speed and Cell Persistence. *Cell*, 161(2):374–386, April 2015. ISSN 00928674. doi:10.1016/j.cell.2015.01.056. URL <https://linkinghub.elsevier.com/retrieve/pii/S0092867415001804>.
- [84] Younan Li and Edwin Munro. Filament-guided filament assembly provides structural memory of filament alignment during cytokinesis. *Developmental Cell*, 56(17):2486–2500.e6, September 2021. ISSN 1534-5807. doi:10.1016/j.devcel.2021.08.009. URL [https://www.cell.com/developmental-cell/abstract/S1534-5807\(21\)00668-7](https://www.cell.com/developmental-cell/abstract/S1534-5807(21)00668-7).
- [85] Matteo Rauzi, Pascale Verant, Thomas Lecuit, and Pierre-François Lenne. Nature and anisotropy of cortical forces orienting Drosophila tissue morphogenesis. *Nature Cell Biology*, 10(12):1401–1410, December 2008. ISSN 1476-4679. doi:10.1038/ncb1798. URL <https://www.nature.com/articles/ncb1798>.
- [86] Mirjam Mayer, Martin Depken, Justin S. Bois, Frank Jülicher, and Stephan W. Grill. Anisotropies in cortical tension reveal the physical basis of polarizing cortical flows. *Nature*, 467(7315):617–621, September 2010. ISSN 1476-4687. doi:10.1038/nature09376. URL <https://www.nature.com/articles/nature09376>.
- [87] Anne Sailer, Alexander Anneken, Younan Li, Sam Lee, and Edwin Munro. Dynamic Opposition of Clustered Proteins Stabilizes Cortical Polarity in the *C. elegans* Zygote. *Developmental Cell*, 35(1):131–142, October 2015. ISSN 15345807. doi:10.1016/j.devcel.2015.09.006. URL <https://linkinghub.elsevier.com/retrieve/pii/S1534580715005869>.
- [88] Tim Dockhorn, Arash Vahdat, and Karsten Kreis. Score-Based Generative Modeling with Critically-Damped Langevin Diffusion, March 2022. URL <http://arxiv.org/abs/2112.07068>. arXiv:2112.07068 [stat].

- [89] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion Models: A Comprehensive Survey of Methods and Applications. *ACM Computing Surveys*, 56(4):1–39, April 2024. ISSN 0360-0300, 1557-7341. doi:10.1145/3626235. URL <https://dl.acm.org/doi/10.1145/3626235>.
- [90] Patrick M. McCall, Frederick C. MacKintosh, David R. Kovar, and Margaret L. Gardel. Cofilin drives rapid turnover and fluidization of entangled F-actin. *Proceedings of the National Academy of Sciences*, 116(26):12629–12637, June 2019. doi:10.1073/pnas.1818808116. URL <https://www.pnas.org/doi/abs/10.1073/pnas.1818808116>.
- [91] Benjamin Kuznets-Speck, Jaekwon Jung, Pornchanan Pholraksa, Adrianne Zhong, Leon Schwartz, Ekta Prashnani, Suriyanarayanan Vaikuntanathan, and Yogesh Goyal. Interpretable Thermodynamic Score-based Classification of Relaxation Excursions, December 2025. URL <https://www.biorxiv.org/content/10.1101/2025.11.26.690838v1>. ISSN: 2692-8205 Pages: 2025.11.26.690838 Section: New Results.
- [92] Antonio Sclocchi, Alessandro Favero, and Matthieu Wyart. A Phase Transition in Diffusion Models Reveals the Hierarchical Nature of Data, December 2024. URL <http://arxiv.org/abs/2402.16991>. arXiv:2402.16991 [stat].
- [93] Siddhartha Saha, Qixin Yang, Wolfgang Losert, Alexandre V. Morozov, and Anirvan M. Sengupta. Spatiotemporal feature learning for actin dynamics. *PLOS ONE*, 20(3):e0318036, March 2025. ISSN 1932-6203. doi:10.1371/journal.pone.0318036. URL <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0318036>.
- [94] Matthew S. Schmitt, Jonathan Colen, Stefano Sala, John Devany, Shailaja Seetharaman, Alexia Caillier, Margaret L. Gardel, Patrick W. Oakes, and Vincenzo Vitelli. Machine learning interpretable models of cell mechanics from protein images. *Cell*, 187(2):481–494.e24, January 2024. ISSN 0092-8674, 1097-4172. doi:10.1016/j.cell.2023.11.041. URL [https://www.cell.com/cell/abstract/S0092-8674\(23\)01331-4](https://www.cell.com/cell/abstract/S0092-8674(23)01331-4).
- [95] Ming Han, John Devany, Michel Fruchart, Margaret L. Gardel, and Vincenzo Vitelli. Learning noisy tissue dynamics across time scales, October 2025. URL <http://arxiv.org/abs/2510.19090>. arXiv:2510.19090 [cond-mat].
- [96] Francis Cavanna and José Alvarado. Quantification of the mesh structure of bundled actin filaments. *Soft Matter*, 17(19):5034–5043, 2021. doi:10.1039/D1SM00428J. URL <https://pubs.rsc.org/en/content/articlelanding/2021/sm/d1sm00428j>.
- [97] Margaret L. Gardel, Karen E. Kasza, Clifford P. Brangwynne, Jiayu Liu, and David A. Weitz. Chapter 19 Mechanical Response of Cytoskeletal Networks. In *Methods in Cell Biology*, volume 89, pages 487–519. Elsevier, 2008. ISBN 978-0-12-372521-9. doi:10.1016/S0091-679X(08)00619-5. URL <https://linkinghub.elsevier.com/retrieve/pii/S0091679X08006195>.

- [98] Qi Lian, Yu Qi, and Yueming Wang. Cauchy Diffusion: A Heavy-tailed Denoising Diffusion Probabilistic Model for Speech Synthesis. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(23):24549–24557, April 2025. ISSN 2374-3468. doi:10.1609/aaai.v39i23.34634. URL <https://ojs.aaai.org/index.php/AAAI/article/view/34634>.
- [99] Robert Richardson, H. Dennis Tolley, and Kenneth Kuttler. Expanding the Standard Diffusion Process to Specified Non-Gaussian Marginal Distributions, August 2025. URL <http://arxiv.org/abs/2508.03617>. arXiv:2508.03617 [math].
- [100] Krunoslav Lehman Pavasovic, Jakob Verbeek, Giulio Biroli, and Marc Mezard. Classifier-Free Guidance: From High-Dimensional Analysis to Generalized Guidance Forms, May 2025. URL <http://arxiv.org/abs/2502.07849>. arXiv:2502.07849 [cs].
- [101] Jordan L. Shivers, Michael Nguyen, Aaron R. Dinner, Petia M. Vlahovska, and Suriyanarayanan Vaikuntanathan. Renormalized Mechanics and Stochastic Thermodynamics of Growing Vesicles. *PRX Life*, 4(1):013012, January 2026. ISSN 2835-8279. doi:10.1103/qxn3-xx5n. URL <https://link.aps.org/doi/10.1103/qxn3-xx5n>.
- [102] Kotaro Ikeda, Tomoya Uda, Daisuke Okanohara, and Sosuke Ito. Speed-accuracy relations for diffusion models: Wisdom from nonequilibrium thermodynamics and optimal transport, May 2025. URL <http://arxiv.org/abs/2407.04495>. arXiv:2407.04495 [cond-mat].
- [103] Luca Ambrogioni. The Statistical Thermodynamics of Generative Diffusion Models: Phase Transitions, Symmetry Breaking, and Critical Instability. *Entropy*, 27(3):291, March 2025. ISSN 1099-4300. doi:10.3390/e27030291. URL <https://www.mdpi.com/1099-4300/27/3/291>.