MODE: LEARNING COMPOSITIONAL REPRESENTATIONS OF COMPLEX SYSTEMS WITH MIXTURES OF DYNAMICAL EXPERTS

Anonymous authorsPaper under double-blind review

ABSTRACT

Dynamical systems in the life sciences are often composed of complex mixtures of overlapping behavioral regimes. Cellular subpopulations may shift from cycling to equilibrium dynamics or branch towards different developmental fates. The transitions between these regimes can appear noisy and irregular, posing a serious challenge to traditional, flow-based modeling techniques which assume locally smooth dynamics. To address this challenge, we propose MODE (Mixture Of Dynamical Experts), a graphical modeling framework whose neural gating mechanism decomposes complex dynamics into sparse, interpretable components, enabling both the unsupervised discovery of behavioral regimes and accurate long-term forecasting across regime transitions. Crucially, because agents in our framework can jump to different governing laws, MODE is especially tailored to the aforementioned noisy transitions. We evaluate our method on a battery of synthetic and real datasets from computational biology. First, we systematically benchmark MODE on an unsupervised classification task using synthetic dynamical snapshot data, including in noisy, few-sample settings. Next, we show how MODE succeeds on challenging forecasting tasks which simulate key cycling and branching processes in cell biology. Finally, we deploy our method on human, single-cell RNA sequencing data and show that it can not only distinguish proliferation from differentiation dynamics but also predict when cells will commit to their ultimate fate, a key outstanding challenge in computational biology.

1 Introduction

Forecasting the long-term behavior of dynamical systems from sparse, noisy data is a major challenge in computational biology, applied physics and engineering (Bury et al., 2023; Moriel et al., 2024; Romeo et al., 2025). For example, predicting the developmental fates of progenitor cells from single cell RNA sequencing (scRNAseq) data is the subject of intense research in systems biology (Aivazidis et al., 2025). Two factors can make this endeavor especially difficult. First, data may consist of snapshots of unordered states and velocities at a single moment, (x, \dot{x}) , instead of time series. This is often the case in transcriptomic, proteomic and other cellular data, which is acquired by destroying the tissue. Second, the dynamics generating the snapshots may consist of complex mixtures of overlapping behavioral regimes, so that forecasting rules learned for some samples do not generalize to others.

Consequently, a large body of research at the intersection of computational biology and data-driven dynamical systems has been devoted to the modeling of snapshot data. Standard approaches, like neural ordinary/partial differential equations (NODEs/NPDEs) (Lin et al., 2025), flow matching (Haviv et al., 2024; Wang et al., 2025) and dynamical optimal transport (Tong et al., 2020), attempt to fit a continuous-time flow in the form of an ODE or PDE which relates state to velocity. These are powerful, versatile approaches, but they can break down when overlapping dynamical regimes create ambiguous velocity signals. For example, when developing progenitor cells split into two different lineages (Fig. 1, left), cells of different fates (red vs green) can overlap in noisy transition zones (inset). Learning a single flow (e.g. Fig. 1, middle, NODE) tends to average the observed flows in this transition: the estimated developmental process then blurs between the two lineages (Fig. 1, middle, purple cells). This is a fundamental problem with all differential-equation modeling,

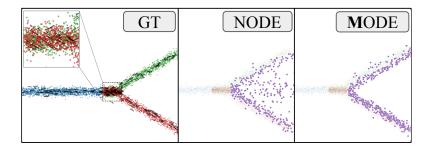


Figure 1: *Branching in in NODE vs MODE.* (*Left*) Ground truth progenitor cells in 2d advance along the blue trunk at a constant rate until they bifurcate into two fates, green and red, passing through an ambiguous switching region (inset). (*Middle*) Models that learn a single flow, like a NODE, can only reconcile switching zones by averaging. (*Right*) Instead, MODE learns a compositional flow with dedicated experts for trunk and branches, helping cells commit to their lineages.

which assumes trajectories change smoothly and cannot cross. With no means to disambiguate the underlying dynamical mixture, flow-based models struggle in these crucial cases.

Here, we tackle the case of branching, overlapping dynamics head-on using a neurally gated mixture of experts, MODE, (Mixture Of Dynamical Experts) which learns compositional representations of complex flows. MODE fits snapshot data using a mixture of dynamical regressors to learn multiple governing laws as well as the transitions between them. As a result, it can tightly model the categorical branching choices which pervade single cell dynamics (Fig. 1, right, MODE). After training, MODE can be rolled out as a stochastic ODE in which evolving particles dynamically shift between different component dynamics. The expert distribution can vary as a function of state, x, helping to model localized shifts in dynamics, or the experts can remain independent of x, in which case MODE can be used to classify heterogeneous mixtures of dynamical agents. Below, we demonstrate MODE's utility in both dynamical classification and forecasting tasks on synthetic and real data taken from across the biological sciences. Concretely, this paper's contributions are:

- Unsupervised classification of heterogeneous dynamical populations. We first show how MODE can be used to classify mixtures of dynamical snapshots in an unsupervised manner. Data represent behaviors of intermingled dynamical agents with different governing equations. We show how MODE far outperforms standard unsupervised baselines and even competes with a supervised multi-layer-perceptron (MLP).
- 2. Systematic noise and sample complexity benchmarking. We run systematic tests on MODE's classification accuracy in increasingly difficult noise and sample complexity conditions, showing its favorable performance in these challenging and realistic cases.
- 3. Forecasting in synthetic biological switches. We then show MODE outperforms existing methods on forecasting the long-term behavior of synthetic genetic switches representing fundamental biological processes, like the cell cycle and branching lineages.
- 4. *Predicting cycle exit in scRNAseq data*. Finally, we demonstrate how MODE can accurately detect, unsupervised, the cell cycle from scRNAseq data and how its mixture distribution can be used to quantitatively forecast the moment of cell differentiation.

2 RELATED WORK

Data-driven dynamical systems. Modeling dynamical systems from data has attracted considerable attention across applied mathematics, machine learning and computational biology (e.g. Schaffer & Kot 1985; Plum & Serra 2025; Khona & Fiete 2022; Costa et al. 2019; Bar et al. 2025). Traditional flow-based approaches, such as Neural Ordinary Differential Equations (NODEs) (Chen et al., 2019), have proven powerful for learning continuous-time dynamics from data by parameterizing the derivative of hidden states through neural networks. These methods excel at capturing smooth trajectories and enable efficient integration for trajectory prediction. However, NODEs inherently

assume a single, globally consistent flow, making them fundamentally unsuitable for systems exhibiting multiple coexisting dynamical regimes or abrupt transitions between behavioral modes. Much the same holds for SINDy (Brunton et al., 2016), which regressors closed form ODEs to data, but also fits single flows. Other approaches attempt to model long term system behavior directly by learning topologically invariants (Moriel et al., 2024), though these frameworks require supervised pre-training and are tailored specifically to limit cycles, as opposed to arbitrary dynamical regimes.

114 115 116

117

118

119

120

121

122

123

124

125

126

127

128

129

108

109

110

111

112

113

Flow learning in computational biology. In computational biology, flow-based methods have found particular success in modeling cellular dynamics through RNA velocity estimation (Bergen et al., 2020) and optimal transport approaches (Tong et al., 2020). Meta Flow Matching (Atanackovic et al., 2024) represents a recent advancement, learning vector fields on the Wasserstein manifold to model population-level dynamics. Recent developments in structured latent velocity modeling (Farrell et al., 2023) have further extended these approaches by incorporating latent variable frameworks to capture complex single-cell transcriptomic dynamics, demonstrating improved performance in scenarios involving temporal gene expression patterns. While these approaches capture important aspects of biological systems, they encounter significant challenges when cellular populations exhibit branching behaviors or discrete fate decisions—scenarios where multiple dynamical laws coexist spatially. The integration of dynamical systems theory as an organizing principle for single-cell biology (Islam & Bhattacharya, 2025) has highlighted the critical need for methods that can simultaneously capture both continuous flows and discrete transitions, particularly in contexts involving cell cycle dynamics (Lugagne et al., 2024) and developmental branching processes. Moreover, the challenge of model discovery from noisy biological data has motivated hybrid approaches that combine neural networks with sparse regression techniques to improve robustness (Wu et al., 2025).

130 131 132

133

134

135

136

137

138

139

140

141

142

143

Switching dynamical systems. The limitations of single-flow models have motivated extensive research into switching dynamical systems, which explicitly model transitions between different dynamical regimes. Classical approaches include piecewise affine models and hybrid systems (Jin et al., 2021), which segment trajectories and fit separate dynamics to each segment. More recent neural approaches (Ojeda et al., 2021; Seifner & Sanchez, 2023) have incorporated deep learning architectures to learn both the underlying dynamics and switching mechanisms simultaneously, often employing attention mechanisms or variational inference frameworks. Continuous-time switching systems (Köhs et al., 2021) extend these ideas by modeling regime transitions as Markov jump processes, providing principled probabilistic frameworks for handling temporal uncertainty. A notable recent advancement is Switched Flow Matching (Zhu & Lin, 2024), which addresses the singularity problem inherent in continuous-time generative models by employing switching ODEs rather than uniform flows, thereby enabling more effective transport between heterogeneous distributions. While theoretically appealing, these methods typically require either prior segmentation of trajectories or complex inference procedures that may not scale effectively to high-dimensional biological data.

148

149

150

151

152

153

154

155

156

157

158

159

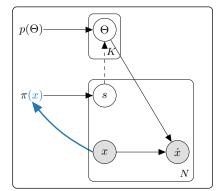
160

161

Mixture of experts (MoE). Mixture methods represent a fundamental approach throughout the data sciences. Early applications of MoE to dynamical systems identification (Lima et al., 2002; Weigend et al., 1995) demonstrated the ability to discover temporal regimes and avoid overfitting through soft partitioning of the input space. Unlike switching systems that make hard assignments, MoE models enable smooth transitions between experts through learned gating functions, making them particularly well-suited for noisy biological data where regime boundaries are inherently ambiguous. However, traditional MoE approaches face several interrelated challenges when applied to dynamical systems. Most critically, they typically lack interpretability in their expert components, as the learned dynamics are often represented by black-box neural networks that provide little insight into the underlying physical or biological mechanisms. Furthermore, these methods very often rely on the availability of time series data, which is not the case of the snapshot data which are so prevalent in cell biology. Recent advances in vector field decomposition (Roy et al., 2006) have shown promise for parametric fields but remain limited to normalized vector fields and cannot distinguish between regions with similar directions but different magnitudes. Similarly, dynamical clustering approaches (Milewski & Tabak) have emphasized interpretable, localized dynamics but remain constrained to specific distance metrics and may not scale effectively to the complex, multi-modal distributions typical of biological systems.

Building on these foundations, our MODE framework unifies interpretable expert dynamics identification with the flexibility of mixture models. Unlike previous approaches that treat clustering and dynamics modeling as separate tasks,

MODE jointly learns spatial organization and governing equations, enabling both accurate classification of heterogeneous populations and precise forecasting of regime transitions. This unified perspective directly addresses key limitations of existing methods: the inability of flow-based models to handle branching dynamics, the computational complexity of switching system inference, and the lack of interpretability in many neural approaches. By combining principled expert modeling with neural gating mechanisms, MODE provides a theoretically grounded yet practically effective solution for modeling complex biological systems with multiple coexisting dynamical behaviors.



3 METHODS

Let $D=\{(x_i,\dot{x}_i)\}_{i=1}^N$, where $x_i\in\mathbb{R}^d$, denote snapshot data sampled from an underlying dynamical system. We assume the data are generated by K latent dynamical laws, each with unknown governing equation f_{Θ_s} parameterized by $\Theta_s\in\mathbb{R}^m$. We place a Laplace prior $\Theta_s\sim \mathrm{Lap}(0,1/\lambda)$ on the parameters, which promotes sparsity. Each data point is assigned to one of the K experts, indexed by $s\in[K]$ which is categorically distributed according to the mixing distribution, $s\sim\pi$. The mixing distribution $\pi(x)$ may depend on the state x (as in the cell cycle, where oscil-

Figure 2: Plate diagram of MODE. Expert parameters Θ (with prior $p(\Theta)$) generate velocities \dot{x} via $f_{\Theta_s}(x)$ under isotropic noise; the expert distribution optionally depends on x (blue), in which case it gates states to specific experts, s.

latory regimes are localized in transcriptomic space) or may be independent of x (as in the case of heterogeneous subpopulations with distinct intrinsic dynamics).

We model the probability of a state's velocity with isotropic Gaussian noise,

$$\dot{x} \mid x, s \sim \mathcal{N}(f_{\Theta_s}(x), \sigma_s^2 I_d)$$
.

The full generative model of the data (Fig. 2) is then

$$p(D \mid \boldsymbol{\Theta}, \sigma, \pi) = \prod_{i=1}^{N} \sum_{s=1}^{K} \pi_s(x_i) \mathcal{N}(\dot{x}_i \mid f_{\Theta_s}(x_i), \ \sigma_s^2 I_d), \tag{1}$$

where $\sigma = (\sigma_s)_{s=1}^K$, giving the conditional negative log likelihood on \dot{x} as

$$-\log p(\dot{x} \mid x) = -\log \sum_{s=1}^{K} \pi_s(x) \exp\left(-\frac{1}{2\sigma_s^2} \|\dot{x} - f_{\Theta_s}(x)\|_2^2 - \frac{d}{2}\log(2\pi\sigma_s^2)\right). \tag{2}$$

Our goal is to estimate π , Θ , σ by minimizing the maximum a posteriori (MAP) objective

$$\mathcal{L}(x, \dot{x}) = -\log \sum_{s=1}^{K} \pi_s(x) \, p(\dot{x} \mid x, s) + \lambda \sum_{s=1}^{K} \|\Theta_s\|_1.$$
 (3)

There are several choices for how to model f_{Θ_s} . In practice, we find that symbolic regression works well: f_{Θ_s} is computed as a linear combination of basis functions, Z, depending on x, so that

$$f_{\Theta_s}(x) = Z(x)\Theta_s,$$

where Θ_s serves as the weights. Throughout, we set Z to be polynomials of order up to c, allowing for explicit control as to the complexity of the fitted approximation. In this form, MODE uses a SINDy-based regressor (Brunton et al., 2016), though other estimators are possible (see Sec. 5).

If π in Eq. 3 does not depend on x, then the MAP can be estimated efficiently using an expectation–maximization (EM) algorithm (see App. C.1). Otherwise, we set $\pi(x)$ to be a multilayer

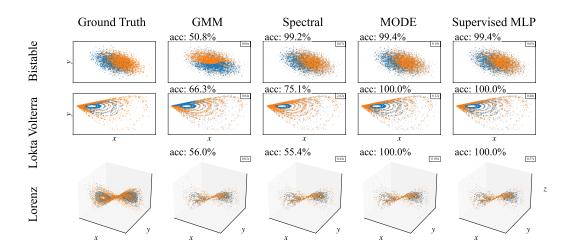


Figure 3: Elementary dynamics discovery on canonical systems. Each system challenges spatial clustering in distinct ways: (**Bistable**) spatially overlapping attractors with opposite rotational dynamics, (**Lotka-Volterra**) nonlinear predator-prey interactions creating complex phase structure, (**Lorenz**) chaotic flow with intricate spatial organization. Ground truth colors indicate true dynamical regimes. Clustering results colors indicate training data (gray) and estimated dunamical regimes (blue, orange). Accuracy of each method is displayed oon the top-left of each subplots.

perceptron (MLP) in the manner of a gating network (Jacobs et al., 1991) and optimize Eq. 3 with stochastic gradient descent. However, when optimizing using gradient descent, the gating function has a tendency to converge to a local minima corresponding to equal probabilities for all points. Thus, we found that performance was improved by regularizing Eq. 3 with the additional terms

$$H_{\text{gate}} = -\frac{1}{B} \sum_{i=1}^{B} \sum_{s=1}^{K} \pi_{i,s} \log \pi_{i,s} \quad ; \quad \text{KL}_{\text{balance}} = \sum_{s=1}^{K} \bar{\pi}_{s} \left(\log \bar{\pi}_{s} - \log \frac{1}{K} \right),$$
 (4)

where H_{gate} encourages confident (low-entropy) gating decisions *per sample* and KL_{balance} prevents expert collapse by balancing usage across experts (see Sec. C).

After training, we can simulate MODE forward in time according to the generative model, with a discrete-time Euler-Maruyama-like update:

$$s_{t+1} \sim \pi(\cdot \mid x_t) \tag{5}$$

$$x_{t+1} = x_t + \Delta t f_{\Theta s_{t+1}}(x_t) + \sigma_B \sqrt{\Delta t} \, \xi_t, \tag{6}$$

where $\xi_t \sim \mathcal{N}(0, I_D)$ is standard D-dimensional Gaussian noise and σ_B controls the stochasticity level. When π is especially peaked, cells can "commit" to a new fate by sampling preferentially from that modal expert.

4 RESULTS

4.1 DISENTANGLING HETEROGENEOUS DYNAMICAL POPULATIONS

Many complex systems in biology comprise heterogeneous mixtures of agents with different dynamical profiles, like tissues with multiple cell types or ecosystems with multiple species. To demonstrate that MODE's compositional dynamics provide essential advantages in disentangling these populations, we evaluate its performance on an unsupervised clustering task using mixture versions of three canonical dynamical systems where "geometry-only" methods systematically fail.

In particular, we generated data from a bistable attracting system (2d), the Lotka-Volterra predatorprey model (2d) and the Lorenz model of chaos (3d). In each case, snapshot samples, (x, \dot{x}) , were randomly drawn from two distinct parameter settings, creating overlapping attracting blobs (bistable), concentric orbits (Lotka volterra) or intertwined paths (Lorenz) (Fig. 3, first column; see Sec.A for data details). Data consisted of snapshots (x, \dot{x}) with additive noise $(\sigma = 0.1)$ split 80-20

into train and test. For each dataset we chose 2 sets of parameters and sampled following numerous initial conditions.

We fit a 2-MODE with x-independent mixture distributions to each of these data sets using an EM algorithm to optimize the expert allocation and dynamical parameters. We also fit a two-class gaussian mixture model (GMM) as well as a spectral clustering algorithm using standard, out-of-the-box parameters from scikit-learn. All models were trained with full snapshot data, (x, \dot{x}) , including velocities. To ensure fair comparison, all methods were tuned via cross-validation, provided with the correct number of clusters K, and evaluated using identical train-test splits (see Sec. C.1 for details). Additionally, we included a supervised MLP baseline using true cluster labels to establish an upper performance bound. Models were evaluated by their normalized mutual information (NMI) and adjusted Rand Index (ARI).

Method		Bistable	Lokta Volterra	Lorenz	Average
GMM	ARI	0.000 ± 0.000	0.069 ± 0.000	-0.000 ± 0.000	0.023 ± 0.000
	NMI	0.000 ± 0.000	0.059 ± 0.000	0.000 ± 0.000	0.020 ± 0.000
Spectral	ARI	0.960 ± 0.000	0.052 ± 0.011	0.000 ± 0.000	0.337 ± 0.004
_	NMI	0.920 ± 0.000	0.113 ± 0.014	0.001 ± 0.000	0.345 ± 0.005
MODE	ARI	$\textbf{0.962} \pm \textbf{0.010}$	$\textbf{0.999} \pm \textbf{0.000}$	$\textbf{0.960} \pm \textbf{0.039}$	$\textbf{0.974} \pm \textbf{0.017}$
	NMI	$\textbf{0.934} \pm \textbf{0.011}$	$\textbf{0.998} \pm \textbf{0.000}$	$\textbf{0.960} \pm \textbf{0.039}$	$\textbf{0.964} \pm \textbf{0.017}$
MLP	ARI	0.977 ± 0.000	1.000 ± 0.000	1.000 ± 0.000	0.992 ± 0.000
	NMI	0.949 ± 0.000	1.000 ± 0.000	1.000 ± 0.000	0.983 ± 0.000

Table 1: ARI and NMI scores in dynamical clustering.

We found that MODE achieved much stronger NMI and ARI scores than both the unsupervised baselines and even competed with the supervised MLP (Table 1). Consistently, the GMM performed the worst, since it could only split ellipsoids in (x, \dot{x}) space, which fails dramatically for these thoroughly mixed systems. Spectral clustering could identify coherent bands from each population, but failed when the bands dissolved into salt-and-pepper noise. MODE, on the other hand, does not rely on phase space geometry, instead clustering systems according to the best division of data into two sparse, dynamical laws. This allows it to cluster populations which look like blobs, bands or true mixtures (Lorenz).

Thus, the improvement is most pronounced in cases purely geometrical clustering fundamentally fails—precisely the scenarios most relevant to biological applications. For noise and sample size benchmarking, see Sec. B.2.

4.2 Forecasting synthetic biological switches

Although Sec. 4.1 demonstrates that MODE can be used to cluster dynamical agents with heterogeneous governing equations, it could be argued that this does not demonstrate MODE's value for modeling dynamics *per se*. After all, a key problem arising in computational biology and elsewhere is not just classifying snapshot data, but also forecasting its behavior into the future.

To that end, we sought to evaluate MODE's utility in forecasting the long-term behavior of dynamical systems with challenging switching behaviors. We examined this capability in two synthetic datasets modeling biological processes fundamental to all multi-cellular organisms: the cell cycle and cell lineage branching. Modeling these processes is made difficult by their inherent stochastic switching behaviors, whereby some cells randomly exit the cycle to continue their development or split off from a current lineage to join a new cell fate. We sought to examine how well MODE's mixture modeling could forecast this behavior compared to standard flow-learning baselines.

For cell cycle data, we used a classic three-variable model of mitotic oscillation from Goldbeter (Goldbeter, 1991) given by fluctuating levels of cyclin, maturation promoting factor (MPF) and protease (Fig. 4, top, first panel). We set the probability of cell cycle exit (differentiation, Fig. 4 top, first panel, orange branch) to be p=0.15 in a small region on the cycle where cyclin was minimal. Cells which randomly committed to leaving the cycle in this region evolved according to a linear dynamics towards a stable node attractor (green X) placed about one cycle radius away. We sampled data uniformly in arc-length around both the cycle and exit path so that slow parts of the

 W_1

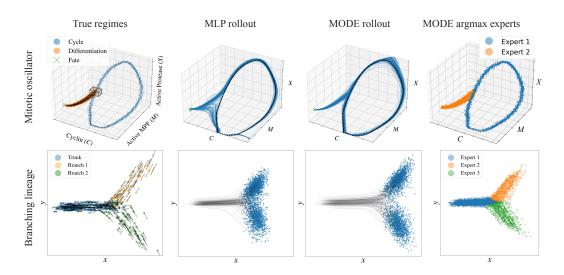


Figure 4: MODE vs MLP baseline on synthetic biological switches. (Top row) Cells evolving according to the Goldbeter oscillator (first panel, blue) could escape (p = 0.15 per time step) to a differentiated state (green X) when they pass through a region of low cyclin (gray box). The MLP (pictured, second panel) and SINDy try to use their single flow to explain the switching system, transiting near the differentiated branch's equilibrium before snaking back wildly towards the cycle. MODE (third panel) simply jettisons cells to the other expert with the appropriate probability since it has located the switching zone (argmax of $\pi(x)$, fourth panel). (Bottom). MODE's' performance advantage on simulated branching lineages (first panel) is similar: the MLP's (pictured, second panel) and SINDy's pushforward distribution blur the early split; MODE (third panel) better separates these branches since its experts locate the branches (fourth panel).

cycle would not be overrepresented. This resulted in 6320 positions and velocities separated into an 80-20 train-validation split.

For branching lineage data, we used a simple two-dimensional switching system (Fig. 4, bottom, first panel) similar to those used in dynamic optimal transport and flow matching studies of transcriptomic dynamics (Farrell et al., 2023; Zhang et al., 2024). A Gaussian blob of cells was advanced with a constant flow along a straight trunk (blue) for a fixed period. After traversing along the trunk, cells randomly branched (50-50) towards two different fates (orange, green) following two linear flows. The resulting 7500 samples were split 80-20 into train and validation.

(44) 11	intotic openiutor		
W_2	$W_{1,C}$	$W_{1,M}$	$W_{1,X}$
0.0015	0.2745 ± 0.0009	0.0896 ± 0.0004	0.1274 ± 0.0006
+0.0025	0.4000 ± 0.0015	0.1500 ± 0.0007	0.2000 ± 0.0010

	,, 1	** 2	77 1,C	* * 1, M	,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,
MLP	0.3288 ± 0.0012	0.3759 ± 0.0015	0.2745 ± 0.0009	0.0896 ± 0.0004	0.1274 ± 0.0006
SINDy	0.5000 ± 0.0020	0.6000 ± 0.0025	0.4000 ± 0.0015	0.1500 ± 0.0007	0.2000 ± 0.0010
MODE	0.0837 ± 0.0005	0.1049 ± 0.0010	0.0590 ± 0.0003	0.0223 ± 0.0002	0.0289 ± 0.0002

(a) Mitotic oscillator

	(b) Branching lineage					
	W_1	W_2	$W_{1,x}$	$W_{1,y}$		
MLP	0.6535 ± 0.0018	0.8371 ± 0.0016	0.1284 ± 0.0006	0.6254 ± 0.0015		
SINDy	0.9000 ± 0.0025	1.1000 ± 0.0030	0.2000 ± 0.0010	0.8000 ± 0.0020		
MODE	0.5713 ± 0.0017	0.7689 ± 0.0015	0.1363 ± 0.0007	0.5452 ± 0.0014		

Table 2: Comparative performance of NODE vs. an MLP and SINDy baseline. Wasserstein scores (lower is better) show distributional discrepancies between final states: W_1 (3D linear cost), W_2 (3D quadratic cost), and the marginal scores (per-coordinate 1D distances).

For the cycle data, we used a 2-expert MODE with up to cubic terms; for branching, a 3-expert MODE with linear terms. In these experiments, the mixture distribution for MODE depended on \boldsymbol{x} and a neural gating function was used, modeled as an MLP with one hidden layer. For baselines, we used a 4-layer MLP and SINDy with cubic (linear) terms for the cycle (branching). All models were trained with the same optimization and early-stopping criteria, and we confirmed both baselines could fit each dynamical component (i.e. Goldbeter cycle, cycle exit, branches) individually. The different models were compared to each other by pushing forward "progenitor cells" (from the cycle or trunk) under the estimated dynamics. The Wasserstein distance between the resulting positions of the cells and their groundtruth counterparts was calculated, using both the full joint and marginal distributions (full training details in Sec. C.2).

We found that MODE outperformed the baseline models on both tasks (Table 2). On the cycle data, the 2-expert MODE correctly identified the cell cycle exit (Fig. 4 top, fourth panel), and the probability assigned to the argmax expert in the exit region was close to the ground truth exit probability ($\pi_0(x) = 0.11 \pm 0.01$) (Fig. 9a). Both baselines tried to learn a single flow accounting for both the stable node and cycle (MLP, Fig. 4, top, second panel), and the resulting circuitous flows significantly increased their Wasserstein error. Note that MODE still provides a very good fit of the Goldbeter oscillator, even though this system uses rational functions outside of our dictionary of basis functions.

On the branching data, MODE's stochastic rollout allowed it to commit (Fig. 4, bottom, third panel) more definitively to each branch, which decreased the error specifically in the *y*-dimensional marginal (i.e. the branching direction) (Tab. 2b). By contrast, both baselines only captured the general quantitative behavior of the switch, spreading noncommittal cells in the branch interior. MODE's experts correctly localized the branching structure (Fig. 4, bottom right) and discovered the underlying governing equations nearly exactly (see Sec. C.2).

4.3 Forecasting cell development fate from SCRNASEQ data

The synthetic cellular processes examined in Sec. 4.2 were challenging to model for regular flow-learning approaches, but they were nevertheless highly idealized compared to their real, biological counterparts. In that spirit, we sought to understand if MODE could be used to forecast the developmental dynamics of the cell cycle from actual scRNAseq data.

To that end, we analyzed single-cell RNA-sequencing data from the U2OS cell line that was integrated with the fluorescent ubiquitination-based cell cycle indicator (FUCCI) in order to investigate the spatiotemporal dynamics of human cell cycle expression (data collected in Mahdessian et al. 2021). The original dataset included 1,152 cells characterized by the expression of 58,884 genes. We follow the preprocessing protocol of Zheng et al. (2023), which involved gene and cell filtering, normalization, log transformation, and the selection of highly variable genes. FUCCI provides a precise indication of which cells are part of a cell-cycle versus those that are not, which gives us a ground truth for two dynamical regimes that mimic our synthetic Goldbeter system of Sec. 4.2: the cycle regime, and the cycle exit regime.

We then used scVelo (Bergen et al., 2020) to estimate the local dynamics of each cell and gene, in the form of an RNA velocity. scVelo computes RNA velocity vectors from spliced and unspliced gene expression counts of a cell population through an analytical model of the transcription, splicing and degradation dynamics of RNA molecules. The gene expression data and the inferred velocities were projected into a 5-dimensional space using principal component analysis (PCA) (Fig. 5a).

We then fit this five-dimensional data with a 2-expert MODE having up to cubic terms, repeating this over ten random initializations to create an ensemble model (further training details in Sec. D.1).

After training, we found that the ensemble averaged expert distribution (Fig. 5b) closely matched the true cell cycle scores (Fig. 5c), indicating that MODE's experts correctly divided gene expression space into a cycling program and a differentiation program. We also observed a characteristic entropic zone between these regimes (Fig. 5b, pink dots), acting as a stochastic escape route for differentiating cells to leave the cycle. We quantitatively confirmed MODE's fit of the true developmental dynamics by calculating the average ROC curve (Fig. 5d) from the ten training initializations, achieving an area-under-the-curve (AUC) score of 0.98, indicating a strong and stable correspondence to the true hidden regimes.

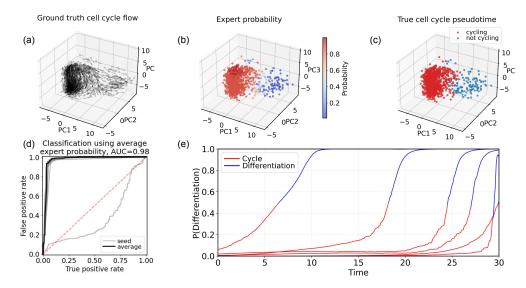


Figure 5: MODE models cell cycle and differentiation dynamics. (a) Single cell RNA sequencing data from the U2OS bone cancer line was preprocessed for RNA velocity with scVelo (first three PCA dimensions shown). (b) We fit a 2-expert MODE to the first five PCA dimensions, revealing a sharp distinction between cycling and differentiating modes (red vs blue) averaged over ten random model initializations. (c) This closely matched ground truth scores computed from pseudotime, which allowed us to (d) accurately, and stably distinguish these two regimes. (e) Furthermore, MODE's stochastic rollout allows us to predict with substantial lead-in time (about a quarter cell cycle, see main text) when cells will commit to their differentiated fate outside of the cycle.

Finally, we examined whether MODE's stochastic rollout could function as a real-time predictor of cell differentiation. We found that cells which eventually differentiated and split off from the cycle exhibited a characteristically smooth, monotonic rise in differentiation-exit probability, in contrast with terminally cycling cells (Fig. 5e). In particular, we found that the differentiation expert's probability for about 90% of terminally cycling cells consistently stayed below 0.1. Setting this as a threshold, we found we could predict whether a cell would differentiate with an average lead-in time of 7.5 units, or about one-quarter of a cell-cycle period.

5 DISCUSSION AND FUTURE DIRECTIONS

Interacting agents in a real-world complex system can exhibit substantial dynamical variety, resulting from intrinsic biases, heterogeneous media, and differing noise levels. While single-flow methods can often explain the long-term behaviors of these heterogeneous systems on average, they can struggle in the case of strong branching, as we have shown. Building on earlier work in switching dynamical systems, MODE aims to discover these differing dynamical regimes from data. We showed the utility of this framework in unsupervised dynamical classification, forecasting in synthetic switches and the modeling of the cell cycle from scRNAseq data.

While these results are promising, MODE is currently limited by its lack of an explicit temporal component. Snapshot data is important and biologically relevant, but including dependence on time could help our framework learn more complex switching and branching schemes. For instance, the current stochastic rollout merely selects the most likely expert at each step. A more sophisticated approach could learn waiting times between switches, in the manner of a continuous-time Markov process. Extensions to higher dimensions, new regressors like neural ODEs, and hierarchical mixtures are also natural future directions.

Overall, we believe that the compositional representation of complex systems is a promising approach which captures the intuition that complicated phenomena should be broken down into simpler parts. MODE shows that this can be done in a straightforward way which can nevertheless outperform baselines on biologically-relevant problems.

REFERENCES

- Alexander Aivazidis, Fani Memi, Vitalii Kleshchevnikov, Sezgin Er, Brian Clarke, Oliver Stegle, and Omer Ali Bayraktar. Cell2fate infers rna velocity modules to improve cell fate prediction. *Nature Methods*, 22(4):698–707, 2025. doi: 10.1038/s41592-025-02608-3.
- Lazar Atanackovic, Xi Zhang, Brandon Amos, Mathieu Blanchette, Leo J. Lee, Yoshua Bengio, Alexander Tong, and Kirill Neklyudov. Meta Flow Matching: Integrating Vector Fields on the Wasserstein Manifold, August 2024. URL http://arxiv.org/abs/2408.14608. arXiv:2408.14608 [cs] version: 1.
- Alon Bar, Ron Moran, Netta Mendelsohn-Cohen, Yael Korem Kohanim, Avi Mayo, Yoel Toledano, and Uri Alon. Pregnancy and postpartum dynamics revealed by millions of lab tests. *Science Advances*, 11(13):eadr7922, 2025.
 - Volker Bergen, Marius Lange, Stefan Peidli, F Alexander Wolf, and Fabian J Theis. Generalizing rna velocity to transient cell states through dynamical modeling. *Nature biotechnology*, 38(12): 1408–1414, 2020.
 - 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, 2016. doi: 10.1073/pnas.1517384113. URL https://www.pnas.org/doi/10.1073/pnas.1517384113.
 - Thomas M. Bury, Daniel Dylewsky, Chris T. Bauch, Madhur Anand, Leon Glass, Tsvi Tlusty, and Bhargav R. Karamched. Predicting discrete-time bifurcations with deep learning. *Nature Communications*, 14(1):6599, 2023. doi: 10.1038/s41467-023-42020-z. URL https://doi.org/10.1038/s41467-023-42020-z.
 - Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural Ordinary Differential Equations. December 2019. doi: 10.48550/arXiv.1806.07366. URL http://arxiv.org/abs/1806.07366. arXiv:1806.07366 [cs].
 - Antonio C Costa, Tosif Ahamed, and Greg J Stephens. Adaptive, locally linear models of complex dynamics. *Proceedings of the National Academy of Sciences*, 116(5):1501–1510, 2019.
 - Spencer Farrell, Christoph Bönisch, Melanie L. Goyal, et al. Inferring single-cell transcriptomic dynamics with structured latent-velo. *Cell Reports Methods*, 3(10):100581, 2023. doi: 10.1016/j.crmeth.2023.100581.
 - Albert Goldbeter. A minimal cascade model for the mitotic oscillator involving cyclin and cdc2 kinase. *Proceedings of the National Academy of Sciences of the United States of America*, 88 (20):9107–9111, 1991. doi: 10.1073/pnas.88.20.9107.
 - Doron Haviv, Aram-Alexandre Pooladian, Dana Pe'er, and Brandon Amos. Wasserstein flow matching: Generative modeling over families of distributions with single-cell genomics applications. 2024. URL https://arxiv.org/abs/2411.00698. arXiv preprint arXiv:2411.00698.
 - Samia Islam and Sudin Bhattacharya. Dynamical systems theory as an organizing principle for single-cell biology. *npj Systems Biology and Applications*, 11(1):85, August 2025. ISSN 2056-7189. doi: 10.1038/s41540-025-00565-3. URL https://www.nature.com/articles/s41540-025-00565-3. Publisher: Nature Publishing Group.
- Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. Adaptive mixtures of local experts. *Neural Computation*, 3(1):79–87, 1991.
- Xiangyu Jin, Jie An, Bohua Zhan, Naijun Zhan, and Miaomiao Zhang. Inferring Switched Nonlinear Dynamical Systems. Formal Aspects of Computing, 33(3):385–406, June 2021. ISSN 0934-5043, 1433-299X. doi: 10.1007/s00165-021-00542-7. URL https://dl.acm.org/doi/10.1007/s00165-021-00542-7.
- Mikail Khona and Ila R Fiete. Attractor and integrator networks in the brain. *Nature Reviews Neuroscience*, 23(12):744–766, 2022.

- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*, 2015. URL https://arxiv.org/abs/1412.6980.
 - Lukas Köhs, Bastian Alt, and Heinz Koeppl. Variational Inference for Continuous-Time Switching Dynamical Systems. In *Advances in Neural Information Processing Systems*, volume 34, pp. 20545–20557. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper/2021/hash/abec16f483abb4f1810ca029aadf8446-Abstract.html.
 - C.Ap.M. Lima, A.L.V. Coelho, and F.J. Von Zuben. Mixture of experts applied to nonlinear dynamic systems identification: a comparative study. In *VII Brazilian Symposium on Neural Networks*, 2002. SBRN 2002. Proceedings., pp. 162–167, November 2002. doi: 10.1109/SBRN. 2002.1181463. URL https://ieeexplore.ieee.org/document/1181463.
 - Zaikang Lin, Sei Chang, Aaron Zweig, Minseo Kang, Elham Azizi, and David A. Knowles. Interpretable neural odes for gene regulatory network discovery under perturbations. *arXiv preprint arXiv:2501.02409*, 2025. doi: 10.48550/arXiv.2501.02409. URL https://arxiv.org/abs/2501.02409. v3.
 - Jean-Baptiste Lugagne, Caroline M. Blassick, and Mary J. Dunlop. Deep model predictive control of gene expression in thousands of single cells. *Nature Communications*, 15(1):2148, March 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-46361-1. URL https://www.nature.com/articles/s41467-024-46361-1. Publisher: Nature Publishing Group.
 - Diana Mahdessian, Anthony J Cesnik, Christian Gnann, Frida Danielsson, Lovisa Stenström, Muhammad Arif, Cheng Zhang, Trang Le, Fredric Johansson, Rutger Schutten, et al. Spatiotemporal dissection of the cell cycle with single-cell proteogenomics. *Nature*, 590(7847):649–654, 2021.
 - Paul A Milewski and Esteban G Tabak. Dynamical Clustering and its Limiting Differential Equations.
 - Noa Moriel, Matthew Ricci, and Mor Nitzan. Let's do the time-warp-attend: Learning topological invariants of dynamical systems. In *International Conference on Representation Learning (ICLR 2024)*, 2024. doi: 10.48550/arXiv.2312.09234. URL https://arxiv.org/abs/2312.09234. arXiv:2312.09234 [cs.LG].
 - Cesar Ojeda, Bogdan Georgiev, Kostadin Cvejoski, Jannis Schucker, Christian Bauckhage, and Ramses J. Sanchez. Switching Dynamical Systems with Deep Neural Networks. In 2020 25th International Conference on Pattern Recognition (ICPR), pp. 6305–6312, Milan, Italy, January 2021. IEEE. ISBN 978-1-7281-8808-9. doi: 10.1109/ICPR48806.2021.9412566. URL https://ieeexplore.ieee.org/document/9412566/.
 - Alex M Plum and Mattia Serra. Dynamical systems of fate and form in development. In *Seminars in Cell & Developmental Biology*, volume 172, pp. 103620. Elsevier, 2025.
 - Nicolas Romeo, Chris Chi, Aaron R. Dinner, and Elizabeth R. Jerison. Characterizing nonlinear dynamics by contrastive cartography. *PRX Life*, 3(3):033022, 2025. doi: 10.1103/jc9p-m3rn. URL https://doi.org/10.1103/jc9p-m3rn.
 - Tristan Roy, Eric Debreuve, Michel Barlaud, and Gilles Aubert. Segmentation of a vector field: Dominant parameter and shape optimization. *Journal of Mathematical Imaging and Vision*, 24 (2):259–276, March 2006. doi: 10.1007/s10851-005-3627-x. URL https://hal.science/hal-00330553. Publisher: Springer Verlag.
 - William M Schaffer and Mark Kot. Do strange attractors govern ecological systems? *BioScience*, 35(6):342–350, 1985.
 - Patrick Seifner and Ramses J. Sanchez. Neural Markov Jump Processes, May 2023. URL http://arxiv.org/abs/2305.19744. arXiv:2305.19744 [cs].

Alex Tong et al. Trajectorynet: A dynamic optimal transport network for single-cell time series. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pp. 7626–7635, 2020.

- Dongyi Wang, Yuanwei Jiang, Zhenyi Zhang, Xiang Gu, Peijie Zhou, and Jian Sun. Joint velocity-growth flow matching for single-cell dynamics modeling. *arXiv preprint arXiv:2505.13413*, 2025. doi: 10.48550/arXiv.2505.13413. URL https://arxiv.org/abs/2505.13413. vl.
- A. S. Weigend, M. Mangeas, and A. N. Srivastava. Nonlinear gated experts for time series: discovering regimes and avoiding overfitting. *International Journal of Neural Systems*, 6(4):373–399, December 1995. ISSN 0129-0657. doi: 10.1142/s0129065795000251.
- Xiaojun Wu, MeiLu McDermott, and Adam L. MacLean. Data-driven model discovery and model selection for noisy biological systems. *PLOS Computational Biology*, 21(1):e1012762, January 2025. ISSN 1553-7358. doi: 10.1371/journal.pcbi. 1012762. URL https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1012762. Publisher: Public Library of Science.
- Zhenyi Zhang, Tiejun Li, and Peijie Zhou. Learning stochastic dynamics from snapshots through regularized unbalanced optimal transport. *arXiv preprint arXiv:2410.00844*, 2024. doi: 10.48550/arXiv.2410.00844. URL https://arxiv.org/abs/2410.00844. Published as a conference paper at ICLR 2025 (oral).
- Shijie C Zheng, Genevieve Stein-O'Brien, Leandros Boukas, Loyal A Goff, and Kasper D Hansen. Pumping the brakes on rna velocity by understanding and interpreting rna velocity estimates. *Genome biology*, 24(1):246, 2023.
- Qunxi Zhu and Wei Lin. Switched Flow Matching: Eliminating Singularities via Switching ODEs. May 2024. doi: 10.48550/arXiv.2405.11605. URL http://arxiv.org/abs/2405.11605. arXiv:2405.11605 [cs].

A SYNTHETIC SYSTEMS

A.1 DISCOVERING ELEMENTARY DYNAMICS IN SPATIALLY OVERLAPPING SYSTEMS

To systematically evaluate MODE's ability to discover elementary dynamics in complex systems, we constructed three canonical test cases that exhibit spatially overlapping regimes with distinct underlying dynamics (Figure 6). These synthetic systems serve as controlled benchmarks where ground truth cluster assignments are known, enabling rigorous quantitative assessment of clustering performance. Each system represents a different class of dynamical complexity commonly encountered in biological and physical applications, specifically designed to challenge spatial clustering methods while providing rich velocity information for dynamics-based approaches.

The **bistable system** (Figure 6, left panels) models two competing attractors that create spatially overlapping but dynamically distinct regimes. Data points are sampled from Gaussian distributions centered at (-0.5,0) and (0.5,0) with spread $\sigma=1.0$, resulting in substantial spatial overlap between the two dynamical modes. The dynamics for each attractor follow:

Mode 0:
$$\dot{x} = -x + 2y$$
, $\dot{y} = -0.5x - y - xy$ (7)

Mode 1:
$$\dot{x} = -x - 2y$$
, $\dot{y} = 0.5x - y + xy$ (8)

where derivatives are computed relative to each attractor's center. As shown in the figure, while the two regimes occupy overlapping spatial regions, their velocity fields exhibit distinctly different flow patterns—Mode 0 displays clockwise rotation while Mode 1 shows counterclockwise dynamics. This configuration exemplifies systems where spatial clustering methods fail due to overlapping support, while velocity-based approaches can successfully distinguish the underlying dynamical regimes.

The **Lotka-Volterra system** (Figure 6, middle panels) implements predator-prey dynamics with two different parameter regimes, each generating distinct oscillatory behaviors. We simulate trajectories

 from 20 uniformly sampled initial conditions over the domain $[10, 50] \times [10, 50]$ and integrate each system for 100 time units. The two dynamical regimes are governed by:

Mode 0:
$$\dot{x} = 0.5x - 0.02xy$$
, $\dot{y} = -0.5y + 0.01xy$ (9)

Mode 1:
$$\dot{x} = 0.5x - 0.04xy$$
, $\dot{y} = -0.6y + 0.01xy$ (10)

The different parameter values create distinct cycle shapes and periods, as visualized by the concentric orbits in the phase portraits. The resulting trajectories exhibit overlapping spatial support where position alone cannot distinguish between regimes, yet the velocity patterns remain characteristic of each underlying dynamics, with Mode 0 generating tighter elliptical cycles and Mode 1 producing broader, more circular orbits.

The **Lorenz system** (Figure 6, right panels) represents chaotic dynamics with two distinct parameter sets, each producing different attractor geometries. We generate trajectories from 20 random initial conditions uniformly distributed over $[-15,15] \times [-15,15] \times [0,40]$ and integrate for 10 time units, discarding the first 1000 time steps to eliminate transients. The two chaotic regimes follow:

Mode 0:
$$\dot{x} = 12(y - x)$$
, $\dot{y} = x(28 - z) - y$, $\dot{z} = xy - 4z$ (11)

Mode 1:
$$\dot{x} = 10(y - x), \quad \dot{y} = x(35.65 - z) - y, \quad \dot{z} = xy - \frac{8z}{3}$$
 (12)

These parameter choices generate attractors with different wing shapes and temporal dynamics. As illustrated in the three-dimensional projections, trajectory segments from different regimes occupy similar spatial regions but exhibit distinct local flow patterns, creating a challenging clustering scenario where velocity information provides the crucial discriminating signal.

For all systems, we add Gaussian noise with standard deviation $\sigma=0.1$ to both position and velocity measurements to simulate realistic experimental conditions. Each dataset contains n=10,000 data points equally distributed between the two dynamical regimes. The resulting datasets exhibit the key challenge that motivates MODE: spatial clustering methods fail due to overlapping support, while velocity information provides the crucial distinguishing signal needed for accurate regime identification. This experimental design directly tests MODE's central hypothesis that incorporating local dynamics can resolve ambiguities that spatial methods cannot handle.

A.2 FORECASTING SYNTHETIC BIOLOGICAL SWITCHES

To evaluate MODE's forecasting capabilities on biologically relevant switching dynamics, we constructed two synthetic datasets that capture fundamental processes in cellular biology: cell cycle dynamics with stochastic exit events and lineage branching processes. These systems represent challenging forecasting scenarios where single-flow models fundamentally fail due to the presence of discrete regime transitions and branching trajectories that violate the assumptions of continuous dynamical systems.

Cell cycle dataset with differentiation exit. The cell cycle dataset implements the classic three-variable Goldbeter oscillator model (Goldbeter, 1991), which describes mitotic oscillations through the temporal dynamics of cyclin concentration, maturation promoting factor (MPF), and protease activity. The Goldbeter model captures essential features of cell cycle regulation through a system of nonlinear ordinary differential equations that generate stable limit cycle behavior representing normal cell division cycles. These are given by:

$$\frac{dC}{dt} = v_i - k_d C - \frac{v_d X C}{K_d + C} \tag{13}$$

$$\frac{dM}{dt} = \frac{V_1(1-M)}{K_1 + (1-M)} - \frac{V_2M}{K_2 + M}, \quad V_1 = \frac{V_{M1}C}{K_c + C}$$
(14)

$$\frac{dX}{dt} = \frac{V_3(1-X)}{K_3 + (1-X)} - \frac{V_4X}{K_4 + X}, \quad V_3 = V_{M3}M$$
 (15)

Parameters were set to the standard values found in Table 3.

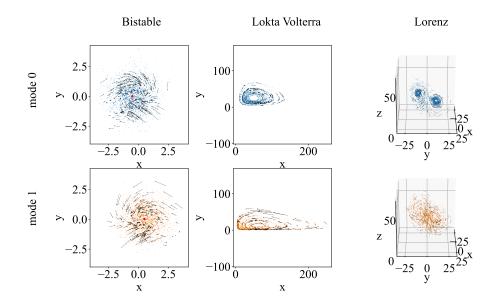


Figure 6: Synthetic datasets for elementary dynamics discovery benchmarks. Phase portraits showing spatially overlapping dynamical regimes for three canonical test systems. Each column represents a different system (bistable, Lotka-Volterra, Lorenz), with Mode 0 (top row, blue) and Mode 1 (bottom row, orange) displaying distinct velocity fields despite spatial overlap. Arrows indicate local flow directions on a subsample of data points. Left: Bistable system with competing attractors showing clockwise vs. counterclockwise rotation patterns. Middle: Lotka-Volterra predator-prey dynamics with different oscillatory regimes producing distinct cycle shapes. Right: Lorenz chaotic system with different parameter sets generating contrasting attractor geometries in 3D projections. These datasets demonstrate the fundamental challenge MODE addresses: spatial clustering fails due to overlapping support, while velocity information provides the discriminating signal for accurate regime identification.

Symbol	Code parameter	Value
$\overline{v_i}$	vi	0.0335
k_d	kd	0.0000
v_d	vd	0.2500
K_d	Kd	0.0200
V_{M1}	VM1	3.2523
K_c	Kc	0.5000
K_1	K1	0.0050
V_2	V2	0.8158
K_2	K2	0.0050
V_{M3}	VM3	1.7020
K_3	K3	0.0050
V_4	V4	1.1580
K_4	K4	0.0050

Table 3: Employed parameters used for Goldbeter's mitotic oscillator.

We augmented the standard Goldbeter oscillator with a stochastic differentiation exit mechanism to simulate the biological phenomenon where cycling cells can irreversibly commit to differentiation. Specifically, we defined a spatial exit zone in the region of minimal cyclin concentration (the G1/S checkpoint region) where cells have a probability p=0.15 per time step of exiting the cell cycle. Cells that stochastically transition in this region evolve according to linear dynamics toward a stable node attractor representing the differentiated state, positioned approximately one cycle radius away from the oscillator trajectory. Specifically, cells evolved according to

765

766

767

768 769

770

771

772

773

774 775

776

777

778

779

780

781 782

783 784

785

786 787 788

789 790

791

792 793 794

795

796

797 798

799

800

801

802

803 804

805

806

807

808

809

where 760

$$\frac{d\mathbf{y}}{dt} = -K\left(\mathbf{y} - \mathbf{x}^*\right),\tag{16}$$

$$\mathbf{x}^* = \begin{bmatrix} -0.3461\\ 0.1481\\ 0.1468 \end{bmatrix}, \qquad K = \begin{bmatrix} 0.6000 & 0 & 0\\ 0 & 0.8000 & 0\\ 0 & 0 & 0.9000 \end{bmatrix}.$$

The resulting dynamics combine three distinct behavioral regimes: (1) stable oscillatory behavior along the limit cycle, (2) stochastic transition events in the exit zone, and (3) deterministic approach to the differentiated attractor. This configuration creates a challenging forecasting scenario where successful prediction requires accurately modeling both the continuous cycle dynamics and the discrete exit decisions that fundamentally alter cell fate trajectories.

Data generation involved uniform sampling in arc-length along both the cycle trajectory and the differentiation path to ensure balanced representation across dynamical regimes and prevent oversampling of slow trajectory segments. This sampling strategy produced n=6,320 position-velocity pairs (x_i, \dot{x}_i) with additive Gaussian noise $(\sigma = 0.1)$, split into 80% training and 20% validation sets. The ground truth exit probability in the transition zone provides quantitative validation of MODE's learned gating function.

Lineage branching dataset. The branching dataset implements a simplified model of cellular lineage specification commonly observed in developmental biology, where a homogeneous progenitor population splits into distinct cell fates through binary decision processes. This system models the essential features of lineage commitment: initial homogeneous dynamics followed by irreversible branching toward distinct terminal states.

Specifically, n = 600 initial cells were drawn from a Gaussian distribution with parameters

$$\mu_0 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \qquad \Sigma_0 = 0.08 I_2.$$

These were evolved according to the system

$$\dot{\mathbf{x}} = A_T \mathbf{x} + \mathbf{c}, \qquad A_T = \begin{bmatrix} 0.15 & -0.05 \\ 0.05 & 0.10 \end{bmatrix}, \quad \mathbf{c} = \begin{bmatrix} 0.6 \\ 0 \end{bmatrix},$$

for 45 steps, at which point 50% of cells were divided between the two lineages.

These new populations evolved according to

$$\dot{\mathbf{x}} = \begin{bmatrix} 0 & 0 \\ \pm s & 0 \end{bmatrix} \mathbf{x} + \mathbf{c}, \qquad s = 0.6,$$

for 30 steps. Throughout, the step size was given by $\Delta t = 0.08$. By sampling all trajectories every step, this process accumulated a total of $600 \times (45 + 30) = 45{,}000$ samples, divided into an 80-20train-validation split.

Experimental design and evaluation metrics. Both datasets challenge forecasting methods through distinct mechanisms: the cell cycle system tests handling of stochastic regime transitions within continuous dynamics, while the branching system evaluates prediction of discrete fate decisions and trajectory divergence. These complementary challenges provide comprehensive assessment of MODE's forecasting capabilities across the spectrum of switching behaviors encountered in biological systems.

Evaluation employed pushforward analysis, where "progenitor cells" sampled from initial conditions (cycle initiation points or trunk origins) were evolved under learned dynamics and compared to ground truth trajectories using Wasserstein distance metrics. This approach directly measures forecasting accuracy for the biological phenomena of interest: predicting cell cycle exit timing and lineage commitment outcomes. Both full joint distributions and marginal distributions were evaluated to assess different aspects of forecasting performance, with particular attention to accuracy in the branching dimension for lineage specification tasks.

B ADDITIONAL EXPERIMENTAL RESULTS

B.1 GOVERNING EQUATION RECOVERY IN MULTI-MODAL DYNAMICAL SYSTEMS

Beyond clustering trajectories into distinct dynamical regimes, MODE simultaneously recovers the governing equations for each discovered mode. This dual capability enables mechanistic interpretation of multi-modal systems by providing sparse polynomial representations of the underlying dynamics. We evaluate MODE's equation discovery performance on three canonical test systems, comparing recovered coefficients against ground truth dynamics.

For each system, we apply MODE with polynomial feature libraries of degree 2, enabling simultaneous trajectory clustering and sparse coefficient identification. The recovered equations provide direct access to the mathematical structure distinguishing different dynamical regimes, offering interpretability that pure clustering methods cannot provide.

Bistable system analysis. The bistable system exhibits spatially overlapping attractors with nonlinear coupling terms that create distinct basins of attraction. Table 4 reveals MODE's capacity for precise coefficient recovery across both modes. The method successfully identifies all structurally important coefficients while maintaining sparsity by correctly setting zero-valued terms to negligible values.

MODE accurately recovers the critical dynamical structure: Mode 0 follows $\dot{x}=-0.5-x+2y$ and $\dot{y}=-0.25-0.5x-1.5y-xy$, while Mode 1 exhibits $\dot{x}=0.5-x-2y$ and $\dot{y}=-0.25+0.5x-1.5y+xy$. The sign reversal in both linear and nonlinear terms creates the bistable character, and MODE's precise identification of these differences (errors $\mathcal{O}(10^{-5})$) demonstrates its sensitivity to dynamical rather than geometric features.

Lotka-Volterra system analysis. The predator-prey dynamics present temporal rather than spatial mode separation, challenging traditional clustering approaches. Table 5 demonstrates MODE's ability to resolve distinct ecological parameter regimes within the same phase space region. The method correctly recovers the prey growth rates (0.5 in both modes) and distinguishes the predation efficiencies: 0.02 for Mode 0 versus 0.04 for Mode 1.

Most critically, MODE identifies the different prey mortality parameters that characterize each regime: -0.5 for Mode 0 and -0.6 for Mode 1. These parameter differences, while seemingly modest, create qualitatively distinct cyclical behaviors that MODE successfully captures through its velocity-informed clustering approach. Coefficient estimation errors remain at the level of $\mathcal{O}(10^{-7})$ for the principal terms.

Lorenz system analysis. The chaotic Lorenz attractor represents the most challenging case for equation discovery, with three-dimensional dynamics and sensitive dependence on initial conditions. Table 6 reveals both the capabilities and limitations of MODE's approach in chaotic regimes. While the method successfully recovers the fundamental structure of the Lorenz equations, estimation accuracy varies significantly across coefficient types.

MODE accurately identifies the critical nonlinear coupling terms: the xy coefficients in the \dot{z} equations (recovered as 1.00 ± 10^{-4} and 1.01 ± 10^{-3}) and maintains the essential xz coupling structure with coefficients of -0.999 ± 10^{-6} and -1.00 ± 10^{-5} . However, the method exhibits larger errors in constant and linear terms, particularly evident in the \dot{y} equations where estimated constant terms $(1.29\pm10^1 \text{ and } -1.59\pm10^1)$ deviate substantially from the true values of 0.

These larger estimation errors in the Lorenz system reflect the inherent challenges of applying sparse regression to chaotic data. The sensitive dependence on initial conditions and the presence of multiple time scales create difficulties in cleanly separating trajectories between modes, leading to coefficient bias when the clustering assignment itself contains uncertainty. Despite these limitations, MODE successfully captures the essential nonlinear structure that governs the chaotic dynamics.

Statistical consistency and limitations. All results represent averages over 10 independent runs with different random initializations, providing estimates of method reliability. For the bistable and Lotka-Volterra systems, coefficient variances remain extremely low $(\mathcal{O}(10^{-7}))$ or smaller), indicat-

Table 4: Equation recovery for the bistable system. Each cell shows the true coefficient, estimated mean \pm variance over 10 runs, and absolute error. MODE achieves near-perfect recovery with errors at the level of $\mathcal{O}(10^{-6})$.

	1	î	x	\hat{x}	y	\hat{y}
mode 0 dx/d	t -0			1 -0.988	$\pm 10^{-5}$ 2	1.98 ± 10^{-5}
mode 0 dy/d	t = -0	.25 -0.248	$\pm 10^{-5}$ -0	0.5 -0.497	$\pm 10^{-5}$ -1 .	$5 -1.49 \pm 10^{-5}$
mode 1 dx/d	t 0.5	$0.497 \pm$	10^{-5} -1	-0.992	$\pm 10^{-6}$ -2	-1.98 ± 10^{-5}
mode 1 dy/d	t -0	.25 -0.244	$\pm 10^{-5}$ 0.5	$0.491 \pm$	-1.	$5 -1.48 \pm 10^{-4}$
	x^2	$\hat{x^2}$	xy	х̂у	y^2	$\hat{y^2}$
mode 0 dx/dt	0	$-0.000152 \pm$	10^{-5} 0	0.00262 ± 10	0^{-5} 0	-0.00307 ± 10^{-6}
mode 0 dy/dt	0	0.000716 ± 10	$^{-5}$ -1	-0.987 ± 10	0^{-5} 0	-0.00317 ± 10^{-6}
mode 1 dx/dt	0	$0.00474 \pm 10^{-}$	0	-0.000486 :	$\pm 10^{-5}$ 0	-0.000675 ± 10^{-5}
mode 1 dy/dt	0	-0.00102 ± 10	0^{-5} 1	0.98 ± 10^{-5}	0	-0.000296 ± 10^{-5}

Table 5: Equation recovery for the Lotka-Volterra predator-prey system. MODE correctly identifies distinct predation rates and mortality parameters that characterize each ecological regime.

	1	î	x	\hat{x}		y	\hat{y}
mode 0 dx/dt/d	t 0	-0.0281 ± 10^{-4}	0.5	0.501	$\pm 10^{-7}$	0	0.000761 ± 10^{-6}
mode 0 dy/dt/d	t 0	-0.006 ± 10^{-4}	0	-2.17	$7e - 05 \pm 10^{-7}$	-0.5	-0.5 ± 10^{-7}
mode 1 dx/dt/d	t 0	0.00661 ± 10^{-4}	0.5	0.499	$\pm 10^{-7}$	0	0.000474 ± 10^{-7}
mode 1 dy/dt/d	lt 0	0.00118 ± 10^{-5}	0	-4.09	$9e - 05 \pm 10^{-8}$	-0.6	-0.6 ± 10^{-7}
	x^2	$\hat{x^2}$		xy	\hat{xy}	y^2	$\hat{y^2}$
mode 0 dx/dt/dt	0	$-3.75e - 07 \pm 10^{-}$	11	-0.02	-0.02 ± 10^{-11}	0	$2.34e - 06 \pm 10^{-10}$
mode 0 dy/dt/dt	0	$-9.74e - 08 \pm 10^{-}$	11	0.01	0.01 ± 10^{-11}	0	$-2.98e - 06 \pm 10^{-10}$
mode 0 dy/dt/dt mode 1 dx/dt/dt	0	$-9.74e - 08 \pm 10^{-1}$ $3.02e - 06 \pm 10^{-11}$		$0.01 \\ -0.04$	$0.01 \pm 10^{-11} \\ -0.04 \pm 10^{-10}$	$0 \\ 0$	$-2.98e - 06 \pm 10^{-10}$ $-3.12e - 05 \pm 10^{-9}$

ing consistent convergence across runs. The Lorenz system shows higher variance, particularly for constant terms, reflecting the added complexity of chaotic trajectory separation.

The equation recovery demonstrates MODE's interpretability advantage: unlike black-box clustering methods, MODE provides explicit mathematical models revealing the mechanisms underlying each dynamical regime. This capability transforms trajectory clustering from pattern recognition into scientific discovery, enabling researchers to understand not just which trajectories belong together, but why they exhibit similar dynamics.

B.2 Noise and training size robustness

To assess the practical applicability of elementary dynamics decomposition, we conduct systematic robustness studies examining how clustering performance degrades under realistic experimental conditions. Real-world dynamical systems inevitably contain measurement noise and are often constrained by limited sample sizes, factors that can severely impact clustering methods that rely on precise velocity estimates.

We evaluated robustness across three canonical test systems (bistable, Lotka-Volterra, Lorenz) by systematically varying two critical experimental parameters: Gaussian noise levels ($\sigma \in \{0.0, 10^{-5}, 10^{-2}, 10^{-1}, 2 \times 10^{-1}\}$) and training dataset sizes ($N \in \{100, 800, 8000\}$). Noise is added to both position x and velocity \dot{x} measurements after normalization ($VAR[x] = VAR[\dot{x}] = 1$) to simulate realistic experimental conditions. For each parameter combination, we evaluate clustering performance using ARI and NMI metrics on held-out test data (20% of each dataset), averaging results over multiple random seeds to ensure statistical reliability.

Table 6: Equation recovery for the Lorenz chaotic system. Despite the system's complexity and three-dimensional dynamics, MODE accurately recovers the key parameters distinguishing the two chaotic attractors.

		1 Î		x	\hat{x}	y	\hat{y}
mode 0 dx/d	t/dt (0.09	93 ± 10^{-2}	-12	-12 ± 10^{-3}	12	11.9 ± 10^{-2}
mode 0 dy/d			$\pm 10^{1}$	28	28.5 ± 10^{0}	-1	-1.1 ± 10^{-2}
mode 0 dz/di		0.0-0.0	0777 ± 10^{-3}	0	-0.653 ± 10^{0}	0	0.266 ± 10^{-1}
mode 1 dx/d			23 ± 10^{-2}	-10	-10.2 ± 10^{-1}	10	10.1 ± 10^{-1}
mode 1 dy/d	t/dt (0 - 1.5	69 ± 10^{1}	35.6		-1	-1.13 ± 10^{-1}
mode 1 dz/da		0.0 - 0.0	0963 ± 10^{-2}	0	0.553 ± 10^{0}	0	-0.212 ± 10^{-1}
	z	â		x^2	$\hat{x^2}$	xy	x̂y
mode 0 dx/dt/dt	0	0.0	386 ± 10^{-2}	0	0.00456 ± 10^{-4}	0	0.00152 ± 10^{-5}
mode 0 dy/dt/dt	0		$.299 \pm 10^{-1}$	0	-0.0516 ± 10^{-2}	0	0.0279 ± 10^{-3}
mode 0 dz/dt/dt	-4	-3	$.92 \pm 10^{-2}$	0	-0.00421 ± 10^{-4}		1 ± 10^{-4}
mode 1 dx/dt/dt	0		$.0406 \pm 10^{-}$	2 0	-0.00311 ± 10^{-4}	0	-0.002 ± 10^{-5}
mode 1 dy/dt/dt	0		65 ± 10^{-1}	0	0.0532 ± 10^{-2}	0	$-0.0362 \pm 10^{-}$
mode 1 dz/dt/dt	-2.6	67 -2	$.76 \pm 10^{-2}$	0	-0.0173 ± 10^{-3}	1	1.01 ± 10^{-3}
	xz	\hat{xz}		y^2	$\hat{\hat{y^2}}$	yz	\hat{yz}
mode 0 dx/dt/dt	0	0.0007	59 ± 10^{-6}	0	-0.000665 ± 10^{-6}	0	-0.000515 ± 10^{-1}
					4		0.00000 1.10-5
node 0 dy/dt/dt	-1	-0.999			-0.0054 ± 10^{-4}	0	0.00202 ± 10^{-5}
		0.0094	$\pm 10^{-4}$	0	-0.000259 ± 10^{-6}	0	-0.00324 ± 10^{-1}
mode 0 dz/dt/dt	-1	0.0094 0.0006	$\pm 10^{-4}$ 64 ± 10^{-6}	0 0	-0.000259 ± 10^{-6} 0.000369 ± 10^{-7}	-	-0.00324 ± 10^{-4} -0.000614 ± 10^{-4}
mode 0 dz/dt/dt mode 1 dx/dt/dt	-1 0	0.0094 0.0006 -1 ± 1	$\pm 10^{-4}$ 64 ± 10^{-6} 0^{-5}	0 0 0	-0.000259 ± 10^{-6} 0.000369 ± 10^{-7} 0.011 ± 10^{-3}	0	-0.00324 ± 10^{-2} -0.000614 ± 10^{-3} 0.00326 ± 10^{-5}
mode 0 dy/dt/dt mode 0 dz/dt/dt mode 1 dx/dt/dt mode 1 dy/dt/dt mode 1 dz/dt/dt	$\begin{array}{c} -1 \\ 0 \\ 0 \end{array}$	0.0094 0.0006 -1 ± 1	$\pm 10^{-4}$ 64 ± 10^{-6}	0 0 0	-0.000259 ± 10^{-6} 0.000369 ± 10^{-7}	0	-0.00324 ± 10^{-4} -0.000614 ± 10^{-4}
mode 0 dz/dt/dt mode 1 dx/dt/dt mode 1 dy/dt/dt	$ \begin{array}{c} -1 \\ 0 \\ 0 \\ -1 \end{array} $	0.0094 0.0006 -1 ± 1	$\pm 10^{-4}$ 64 ± 10^{-6} 0^{-5}	0 0 0	-0.000259 ± 10^{-6} 0.000369 ± 10^{-7} 0.011 ± 10^{-3}	0 0 0	-0.00324 ± 10^{-2} -0.000614 ± 10^{-3} 0.00326 ± 10^{-5}
mode 0 dz/dt/dt mode 1 dx/dt/dt mode 1 dy/dt/dt	$ \begin{array}{c} -1 \\ 0 \\ 0 \\ -1 \end{array} $	0.0094 0.0006 -1 ± 10 -0.008	$\pm 10^{-4}$ 64 ± 10^{-6} 0^{-5}	0 0 0 0	-0.000259 ± 10^{-6} 0.000369 ± 10^{-7} 0.011 ± 10^{-3} -0.00252 ± 10^{-5}	0 0 0	-0.00324 ± 10^{-2} -0.000614 ± 10^{-3} 0.00326 ± 10^{-5}
mode 0 dz/dt/dt mode 1 dx/dt/dt mode 1 dy/dt/dt	$ \begin{array}{c} -1 \\ 0 \\ 0 \\ -1 \end{array} $	0.0094 0.0006 -1 ±1 -0.008	$\pm 10^{-4}$ 64 ± 10^{-6} 0^{-5} 609 ± 10^{-4}	$0\\0\\0\\0\\z^2$	-0.000259 ± 10^{-6} 0.000369 ± 10^{-7} 0.011 ± 10^{-3} -0.00252 ± 10^{-5} $\hat{z^2}$ -0.00187 ± 10^{-5} 0.0076 ± 10^{-4}	0 0 0	-0.00324 ± 10^{-2} -0.000614 ± 10^{-3} 0.00326 ± 10^{-5}
mode 0 dz/dt/dt mode 1 dx/dt/dt mode 1 dy/dt/dt	$ \begin{array}{c} -1 \\ 0 \\ 0 \\ -1 \end{array} $	0.0094 0.0006 -1 ±11 -0.008 mod	$\pm 10^{-4}$ 64 ± 10^{-6} 0^{-5} 609 ± 10^{-4} 609 ± 10^{-4}	$0\\0\\0\\0\\z^2\\0$	-0.000259 ± 10^{-6} 0.000369 ± 10^{-7} 0.011 ± 10^{-3} -0.00252 ± 10^{-5} $\hat{z^2}$ -0.00187 ± 10^{-5}	0 0 0	-0.00324 ± 10^{-2} -0.000614 ± 10^{-3} 0.00326 ± 10^{-5}
mode 0 dz/dt/dt mode 1 dx/dt/dt mode 1 dy/dt/dt	$ \begin{array}{c} -1 \\ 0 \\ 0 \\ -1 \end{array} $	0.0094 0.0006 -1 ±1 -0.008 mod mod	$\pm 10^{-4}$ 64 ± 10^{-6} 0^{-5} 609 ± 10^{-4} The original distribution of the contraction of the co	$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$	-0.000259 ± 10^{-6} 0.000369 ± 10^{-7} 0.011 ± 10^{-3} -0.00252 ± 10^{-5} $\hat{z^2}$ -0.00187 ± 10^{-5} 0.0076 ± 10^{-4}	0 0 0	-0.00324 ± 10^{-2} -0.000614 ± 10^{-3} 0.00326 ± 10^{-5}
mode 0 dz/dt/dt mode 1 dx/dt/dt mode 1 dy/dt/dt	$ \begin{array}{c} -1 \\ 0 \\ 0 \\ -1 \end{array} $	0.0094 0.0006i -1 ±1i -0.008 mod mod mod mod	±10 ⁻⁴ 64 ±10 ⁻⁶ 0 ⁻⁵ 609 ±10 ⁻⁴ de 0 dx/dt/dt e 0 dy/dt/dt e 0 dz/dt/dt	$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$	-0.000259 ± 10^{-6} 0.000369 ± 10^{-7} 0.011 ± 10^{-3} -0.00252 ± 10^{-5} \hat{z}^2 -0.00187 ± 10^{-5} 0.0076 ± 10^{-4} 0.000345 ± 10^{-6}	0 0 0	-0.00324 ± 10^{-2} -0.000614 ± 10^{-3} 0.00326 ± 10^{-5}

Figure 7 reveals that MODE demonstrates competitive performance compared to the MLP base-line while substantially outperforming spatial-only methods across most test conditions. MODE shows robust performance on the bistable and Lotka-Volterra systems, maintaining high clustering accuracy (ARI, NMI > 0.8) even under moderate noise levels where spatial clustering methods fail completely. This advantage stems from MODE's ability to leverage velocity information—even noisy velocity estimates provide directional information that purely spatial clustering cannot access.

The noise robustness patterns reveal system-specific insights: the bistable system shows the most dramatic performance advantages for MODE over spatial methods, with MODE maintaining high accuracy while GMM and Spectral clustering fail completely across most noise levels. This occurs because the bistable system's spatially overlapping attractors become indistinguishable without velocity information as noise increases. The Lorenz system presents the most challenging case, where all methods show performance degradation, though MODE and MLP maintain more robust performance than spatial-only approaches.

The training size analysis reveals differential sample efficiency across systems. On the bistable and Lotka-Volterra systems, MODE maintains reasonable clustering performance even with limited samples (N=100), while spatial methods show severe degradation in small-sample regimes. However, performance variability increases substantially at the smallest sample size (N=100), indicating that while MODE can work with limited data, larger sample sizes ($N\geq800$) provide more reliable results.

The Lotka-Volterra system demonstrates MODE's effectiveness for temporally-separated dynamics, where spatial clustering methods consistently fail regardless of sample size. The underlying advantage is that velocity information provides crucial directional cues about cluster membership that spatial position alone cannot capture—each (x_i, \dot{x}_i) pair encodes both current state and dynamical tendency, enabling more robust clustering even when phase space regions overlap.

These robustness results provide practical guidance for experimental applications. MODE's noise tolerance suggests applicability to real biological data with moderate measurement uncertainties, though performance degrades significantly under high noise conditions ($\sigma > 0.1$). The sample size analysis indicates that while MODE can function with limited data, reliable performance requires adequate sample sizes ($N \ge 800$ recommended) to ensure consistent clustering results.

The comparison with the MLP baseline reveals an important trade-off: while MLP often matches or slightly exceeds MODE's clustering performance, MODE provides the critical advantage of interpretable sparse dynamics. MODE achieves competitive robustness while simultaneously offering mechanistic insight through its recovered governing equations, making it particularly suitable for scientific applications where understanding system structure is as important as clustering accuracy.

B.3 SUPPLEMENTAL FIGURES

C ALGORITHM IMPLEMENTATION DETAILS

We implement two variants of MODE corresponding to different assumptions about the mixing distribution π : MODE Local (state-independent mixing) and MODE Global (state-dependent mixing via neural gating).

C.1 MODE LOCAL

MODE Local implements the case where mixing probabilities $\pi = (\pi_1, \dots, \pi_K)$ are state-independent constants. This variant is particularly suitable for unsupervised clustering of heterogeneous dynamical populations where each subpopulation follows distinct intrinsic dynamics regardless of spatial location.

Algorithm. We optimize the MAP objective (Eq. 3) using an Expectation-Maximization (EM) algorithm with soft assignments. The algorithm alternates between:

E-step: Compute responsibilities $R_{i,k} = p(s_i = k \mid x_i, x_i)$ for each data point i and expert k:

$$R_{i,k} \propto \pi_k \cdot \mathcal{N}(\dot{x}_i \mid Z(x_i)\Theta_k, \sigma_k^2 I_d) \tag{17}$$

where responsibilities are normalized such that $\sum_{k=1}^{K} R_{i,k} = 1$.

M-step: Update expert parameters by solving weighted SINDy regression problems. For each expert k, we fit polynomial coefficients Θ_k by minimizing:

$$\sum_{i=1}^{N} R_{i,k} \|\dot{x}_i - Z(x_i)\Theta_k\|_2^2 + \lambda \|\Theta_k\|_1$$
 (18)

We also update mixing probabilities $\pi_k = \frac{1}{N} \sum_{i=1}^N R_{i,k}$ and noise variances σ_k^2 .

Implementation. The core algorithm is implemented in the MODELocal class with the following key parameters:

- n_clusters: Number of experts K (default: 3)
- degree: Polynomial basis degree for SINDy (default: 2)
- alpha: L1 regularization strength λ (default: 10^{-4})
- max_iter: Maximum EM iterations (default: 150)
- tol: Convergence tolerance on log-likelihood (default: 10^{-5})

The weighted SINDy regression in the M-step uses sample re-weighting with stabilization: for each expert, we scale the design matrix Z and targets \dot{x} by $\sqrt{R_{i,k}}$, apply column-wise standardization without centering, fit Lasso regression, then rescale coefficients back to the original feature scale. This approach improves numerical stability compared to direct weighted least squares.

C.2 MODE GLOBAL

MODE Global implements the case where mixing probabilities depend on state: $\pi(x) = (\pi_1(x), \dots, \pi_K(x))$. This variant uses a neural gating network to model spatially-localized regime transitions, making it suitable for forecasting tasks where dynamical behavior varies smoothly across state space.

Algorithm. We optimize Eq. 3 using stochastic gradient descent with the following neural architecture:

Gating Network: A 3-layer MLP with tanh activations maps normalized states to log-mixing probabilities:

$$\pi(x) = \operatorname{softmax}\left(\operatorname{MLP}\left(\frac{x - \mu_x}{\sigma_x}\right)\right)$$
 (19)

where μ_x , σ_x are empirical mean and standard deviation of the training data.

Expert Networks: Each expert k parameterizes the dynamical law $f_{\Theta_k}(x) = Z(x)\Theta_k$ where Z(x) contains polynomial features up to degree c and Θ_k are learnable weight matrices of shape (P,d) with P polynomial features and d output dimensions.

Regularization. To prevent mode collapse and ensure confident gating decisions, we augment the loss with:

- L1 sparsity: $\lambda \sum_{k=1}^{K} \|\Theta_k\|_1$ promotes sparse expert dynamics
- Gate entropy: $-\frac{1}{B}\sum_{i=1}^{B}\sum_{k=1}^{K}\pi_{i,k}\log\pi_{i,k}$ encourages confident per-sample decisions
- Load balancing: $\sum_{k=1}^{K} \bar{\pi}_k \log(K\bar{\pi}_k)$ prevents expert collapse via KL divergence from uniform distribution

Implementation. The algorithm is implemented in the MODEGlobal class. For all clustering experiments (Sec. 4.1), the key hyperparameters were set to

- K: Number of experts (default: 3)
- degree: Polynomial degree for expert dynamics (default: 2)
- hidden: Hidden units in gating network (default: 64)
- epochs: Training epochs (default: 100)
- 1r: Learning rate (default: 2×10^{-3})
- l1_lambda, ent_lambda, lb_lambda: Regularization weights (defaults: $10^{-4}, 10^{-3}, 5 \times 10^{-4}$)

MODE training and hyperparameters for all experiments are shown in tables Tab. 7, 7, 8 and 10. Wherever parameters are dashed, they were omitted or set to 0 as they had no influence on the result.

All MLP baselines had four hidden layers (except for the two data which had only one layer) with 64 units in each layer and tanh activation functions.

All implementations support multiple random restarts with validation-based model selection to avoid poor local minima. For reproducibility, all experiments use fixed random seeds across method comparisons. GMM and spectral clustering baselines were run in scikit-learn v. 1.7.2. MODE, MLPs and SINDy were implemented in pytorch v. 2.8.0 and were trained with an Adam optimizer (Kingma & Ba, 2015). Experiments were run locally on a cpu cluster with no parallelization. Typical run times range from around a minute for the toy example of Fig. 1 to about ten minutes for the full ten-seed FUCCI ensemble.

Parameter	Description	Value
K	experts	3
LIB_ORDER	poly. degree	0
GATE_HIDDEN	gating hidden sizes	(64,)
ACTIVATION	nonlinearity	tanh
LR	learning rate	2e-3
WEIGHT_DECAY	L2 penalty	1e-5
EPOCHS	training epochs	500
PATIENCE	early stop patience	30
MIN_DELTA	val. loss margin	_
LAM_L1	L1 penalty	1e-4
LAM_ENT	entropy penalty	1e-3
LAM_LB	load balance penalty	5e-4
GRAD_CLIP	gradient clipping	—
BATCH_SIZE	batch size	512

Table 7: Toy Branching (i.e. Fig. 1) hyperparameters.

Parameter	Description	Value
K	experts	2
LIB_ORDER	poly. degree	3
GATE_HIDDEN	gating hidden sizes	(64,)
ACTIVATION	nonlinearity	SILU
LR	learning rate	1e-2
WEIGHT_DECAY	L2 penalty	1e-6
EPOCHS	training epochs	10000
PATIENCE	early stop patience	30
MIN_DELTA	val. loss margin	1e-4
LAM_L1	L1 penalty	1e-3
LAM_ENT	entropy penalty	1e0
LAM_LB	load balance penalty	1e0
GRAD_CLIP	gradient clipping	5.0
BATCH_SIZE	batch size	512

Table 8: Goldbeter oscillator hyperparameters.

D BIOLOGICAL DATASET DETAILS

D.1 FUCCI DATASET

The FUCCI (Fluorescent Ubiquitination-based Cell Cycle Indicator) dataset used in this study originates from the comprehensive single-cell proteogenomics analysis conducted by Mahdessian et al. (2021). This dataset represents a unique integration of single-cell RNA sequencing with fluorescent cell cycle markers in the U2OS human osteosarcoma cell line, providing ground truth annotations for cell cycle phase identification.

Dataset composition and cell cycle labeling. The original dataset comprises 1,152 individual cells characterized by the expression profiles of 58,884 genes. The FUCCI system employs two fluorescently tagged cell cycle markers: CDT1 (tagged with RFP, expressed during G1 phase) and GMNN (tagged with GFP, expressed during S and G2 phases). Cells expressing both markers simultaneously indicate the G1-S transition phase. This dual-marker system provides precise temporal information about cell cycle progression, enabling the identification of cells in cycling versus non-cycling (differentiation) states. The FUCCI markers serve as our ground truth labels for distinguishing between two fundamental dynamical regimes: (1) active cell cycle progression and (2) cell cycle exit leading to differentiation.

Preprocessing pipeline. We followed the preprocessing protocol established by Zheng et al. (2023), which includes several critical steps for robust scRNA-seq analysis:

Parameter	Description	Value
K	experts	3
LIB_ORDER	poly. degree	1
GATE_HIDDEN	gating hidden sizes	(64,)
ACTIVATION	nonlinearity	tanh
LR	learning rate	1e-2
WEIGHT_DECAY	L2 penalty	1e-4
EPOCHS	training epochs	2000
PATIENCE	early stop patience	30
MIN_DELTA	val. loss margin	1e-4
LAM_L1	L1 penalty	1e-3
LAM_ENT	entropy penalty	2e-3
LAM_LB	load balance penalty	1e-1
GRAD_CLIP	gradient clipping	5.0
BATCH_SIZE	batch size	512

Table 9: Lineage branching hyperparameters.

Parameter	Description	Value
K	experts	2
LIB_ORDER	poly. degree	3
GATE_HIDDEN	gating hidden sizes	(64,)
ACTIVATION	nonlinearity	SILU
LR	learning rate	1e-3
WEIGHT_DECAY	L2 penalty	1e-6
EPOCHS	training epochs	5000
PATIENCE	early stop patience	100
MIN_DELTA	val. loss margin	1e-4
LAM_L1	L1 penalty	1e-3
LAM_ENT	entropy penalty	1e0
LAM_LB	load balance penalty	1e0
GRAD_CLIP	gradient clipping	5.0
BATCH_SIZE	batch size	512

Table 10: FUCCI hyperparameters.

- **Gene filtering:** Removal of genes with low expression counts across the cell population to reduce noise and computational burden.
- Cell filtering: Exclusion of cells with abnormally low or high gene counts that might indicate technical artifacts or doublets.
- **Normalization:** Application of library size normalization to account for differences in sequencing depth between cells.
- Log transformation: Log₁₀(count + 1) transformation to stabilize variance and reduce the influence of highly expressed genes.
- Highly variable gene selection: Identification and retention of genes showing significant cell-to-cell variability, which are most informative for capturing biological differences.

RNA velocity estimation. We employed scVelo (Bergen et al., 2020) to estimate RNA velocity vectors from the preprocessed data. scVelo computes velocity by modeling the dynamics of gene expression through an analytical framework that considers transcription, splicing, and degradation rates. Specifically, the method uses the ratio of unspliced to spliced mRNA counts to infer the directional flow of gene expression changes. This provides local velocity estimates \dot{x}_i for each cell i, representing the instantaneous rate of change in gene expression space. The velocity estimation captures the intrinsic directionality of cellular state transitions, making it particularly suitable for identifying developmental trajectories and regime transitions.

Dimensionality reduction and feature space. Given the high-dimensional nature of the gene expression data (58,884 dimensions), we applied Principal Component Analysis (PCA) to project both the gene expression profiles and their corresponding velocity vectors into a 5-dimensional latent space. This dimensionality reduction serves multiple purposes: (1) computational efficiency for MODE training, (2) noise reduction by focusing on the most informative directions of variation, and (3) visualization feasibility for the first three principal components. The choice of 5 dimensions balances the preservation of essential biological signal while maintaining computational tractability for the mixture of experts framework.

MODE training configuration. Full hyperparameters in Tab. 10. An ensemble model created by averaging (e.g. over mixture distributions and expert parameters; experts were aligned by similarity of parameter vectors) over 10 random seeds. All seeds but one (the outlier curve in Fig. 5) divided the data into clear cycle and exit regimes.

Validation methodology. Model performance was evaluated using the FUCCI ground truth labels as the gold standard for cell cycle phase classification. We computed Receiver Operating Characteristic (ROC) curves for each of the 10 model initializations, measuring the ability of the learned expert probabilities to distinguish cycling from differentiating cells. The consistently high Area Under the Curve (AUC) scores (0.98 ± 0.01) across all initializations demonstrate both the reliability of our approach and the biological relevance of the discovered dynamical regimes. Additionally, we validated the biological interpretability of the learned dynamics by examining the correspondence between expert assignments and known cell cycle markers, confirming that the two experts successfully captured the cycling and differentiation programs respectively.

SUPPLEMENTAL VIDEOS

We have also included with this submission four videos showing MLP vs MODE on toy branching data, the stochastic MODE rollout on the Goldbeter oscillator as well as the rollout for our fit of the FUCCI RNA velocity data.

REPRODUCIBILITY STATEMENT

All details necessary to reproduce the results of this work can be found in the appendix and supplementary materials:

- Complete specifications of the MODE architecture, including network architectures, loss functions, and training procedures are provided in Appendix C.
- Hyperparameter choices for all experiments, including learning rates, regularization weights, and model selection criteria can be found in Appendix C.
- Governing equations and parameter values for all synthetic dynamical systems are described in Appendix A for the elementary dynamics benchmarks and Appendix A.2 for the forecasting experiments.
- Detailed preprocessing steps for the FUCCI biological dataset, including gene filtering criteria, normalization procedures, and dimensionality reduction parameters are given in Appendix D.1.
- Source code implementing the MODE framework, including both EM-based and neural-gated variants, is available in the notebooks/ directory with documented examples reproducing all main results.
- All synthetic datasets can be regenerated using the equations and parameters provided in the appendices, while the FUCCI dataset is publicly available from Mahdessian et al. (2021).
- Computational requirements and runtime specifications are documented for all experiments, with typical runtimes.

We have made every effort to ensure that our experimental methodology is transparent and our results are fully reproducible by the research community. Moreover, we provide our complete implementation at https://github.com/anonresearcher22-netizen/MODE.

LARGE LANGUAGE MODEL USAGE STATEMENT

In accordance with ICLR 2026 policies on Large Language Model usage, we disclose that Large Language Models (specifically Claude Sonnet 3.5) were used to assist in the preparation of this manuscript in the following ways:

- Writing assistance: LLMs were used to typeset equations and format tables in LaTeX. It was also used to summarize implementational details from code into sections of the Appendix.
- Literature review support: LLMs assisted in identifying relevant references and helped structure the related work section, though all cited works were independently verified and evaluated by the authors.
- Code documentation and presentation: LLMs were employed to enhance code comments, improve code readability, and assist in the preparation of well-structured, documented implementations in the supplementary materials.

We emphasize that all core research contributions, experimental design, data analysis, and scientific conclusions are entirely the work of the human authors. The authors take full responsibility for the accuracy and validity of all content, including any text that may have been refined with LLM assistance. All LLM-generated content was carefully reviewed, fact-checked, and validated by the authors before inclusion. No LLMs were used for generating experimental results, conducting data analysis, or making scientific claims.

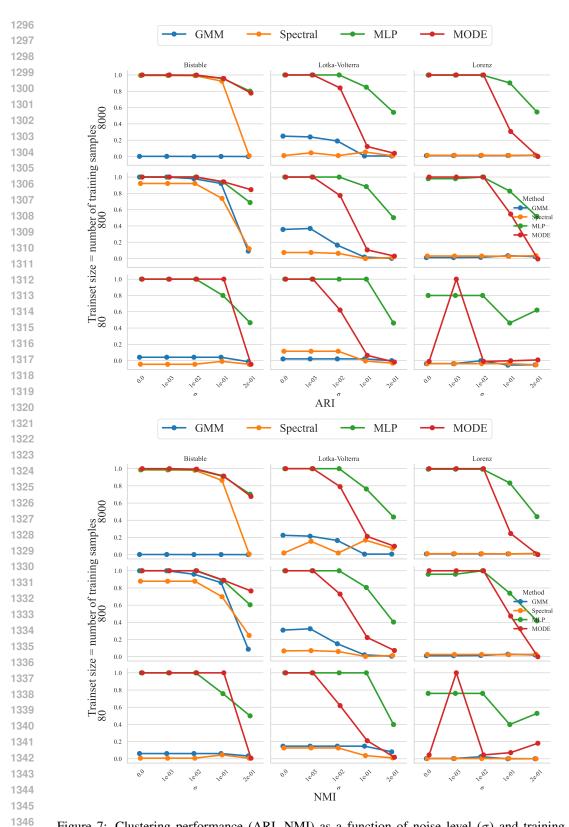
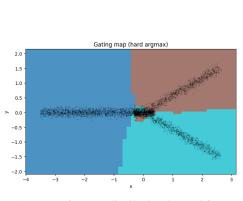
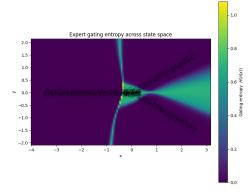


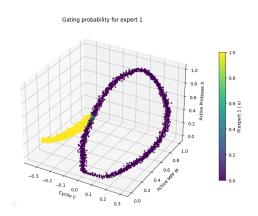
Figure 7: Clustering performance (ARI, NMI) as a function of noise level (σ) and training set size (N) across different dynamical systems. Each subplot corresponds to a specific dataset and metric. MODE demonstrates superior noise tolerance and sample efficiency compared to spatial-only methods, maintaining high performance even under challenging experimental conditions where baseline methods fail.

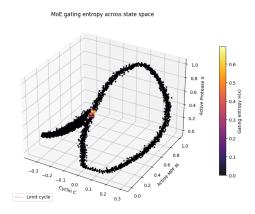




- (a) Argmax of expert distribution learned for toy branching.
- (b) Entropy of expert mixture, highlighting switching region.

Figure 8: Toy branching argmax and entropy figures





- (a) Probability of differentiation expert (expert 1) learned for the Goldbeter system.
- (b) Entropy of expert mixture, highlighting switching region.

Figure 9: Goldbeter probability and entropy figures

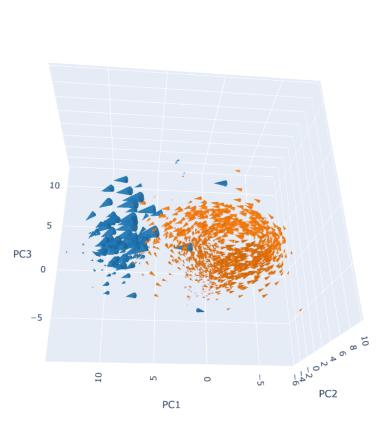


Figure 10: 3D plot of FUCCI cycle data with expert 1 and 2 in orange and blue respectively, highlighting the discovered cycling and exiting modes. Exiting cones doubled in size for visual emphasis.