

000 001 MODE: LEARNING COMPOSITIONAL REPRESENTA- 002 TIONS OF COMPLEX SYSTEMS WITH MIXTURES OF 003 DYNAMICAL EXPERTS 004 005

006 **Anonymous authors**
007 Paper under double-blind review
008
009
010
011
012

ABSTRACT

013 Dynamical systems in the life sciences are often composed of complex mixtures
014 of overlapping behavioral regimes. Cellular subpopulations may shift from cy-
015 cling to equilibrium dynamics or branch towards different developmental fates.
016 The transitions between these regimes can appear noisy and irregular, posing a
017 serious challenge to traditional, flow-based modeling techniques which assume
018 locally smooth dynamics. To address this challenge, we propose MODE (Mix-
019 ture Of Dynamical Experts), a graphical modeling framework whose neural gat-
020 ing mechanism decomposes complex dynamics into sparse, interpretable compo-
021 nents, enabling both the unsupervised discovery of behavioral regimes and accu-
022 rate long-term forecasting across regime transitions. Crucially, because agents in
023 our framework can jump to different governing laws, MODE is especially tailored
024 to the aforementioned noisy transitions. We evaluate our method on a battery of
025 synthetic and real datasets from computational biology. First, we systematically
026 benchmark MODE on an unsupervised classification task using synthetic dynam-
027 ical snapshot data, including in noisy, few-sample settings. Next, we show how
028 MODE succeeds on challenging forecasting tasks which simulate key cycling and
029 branching processes in cell biology. Finally, we deploy our method on human,
030 single-cell RNA sequencing data and show that it can not only distinguish prolif-
031 eration from differentiation dynamics but also predict when cells will commit to
032 their ultimate fate, a key outstanding challenge in computational biology.

033 Forecasting the long-term behavior of dynamical systems from sparse, noisy data is a major chal-
034 lenge in computational biology, applied physics and engineering (Bury et al., 2023; Moriel et al.,
035 2024; Romeo et al., 2025). For example, predicting the developmental fates of progenitor cells from
036 single cell RNA sequencing (scRNAseq) data is the subject of intense research in systems biology
037 (Aivazidis et al., 2025). Two factors can make this endeavor especially difficult. First, data may
038 consist of snapshots of unordered states and velocities at a single moment, (x, \dot{x}) , instead of time
039 series. This is often the case in transcriptomic, proteomic and other cellular data, which is acquired
040 by destroying the tissue. Second, the dynamics generating the snapshots may consist of complex
041 mixtures of overlapping behavioral regimes, so that forecasting rules learned for some samples do
042 not generalize to others.

043 Consequently, a large body of research at the intersection of computational biology and data-driven
044 dynamical systems has been devoted to the modeling of snapshot data. Standard approaches, like
045 neural ordinary/partial differential equations (NODEs/NPDEs) (Lin et al., 2025), flow matching
046 (Haviv et al., 2024; Wang et al., 2025) and dynamical optimal transport (Tong et al., 2020), attempt
047 to fit a continuous-time flow in the form of an ODE or PDE which relates state to velocity. These are
048 powerful, versatile approaches, but they can break down when overlapping dynamical regimes create
049 ambiguous velocity signals. For example, when developing progenitor cells split into two different
050 lineages (Fig. 1, left), cells of different fates (red vs green) can overlap in noisy transition zones
051 where the dynamics are mixed (inset, with green cells having up-slanting velocities; red, down).
052 Learning a single flow (e.g. Fig. 1, middle, NODE) tends to average the observed flows in this
053 transition: the estimated developmental process then blurs between the two lineages (Fig. 1, middle,
purple cells). The problem is all the more acute in cell cycle, where cycling and differentiating
populations intermix (Mahdessian et al., 2021). This is a fundamental problem with all differential-

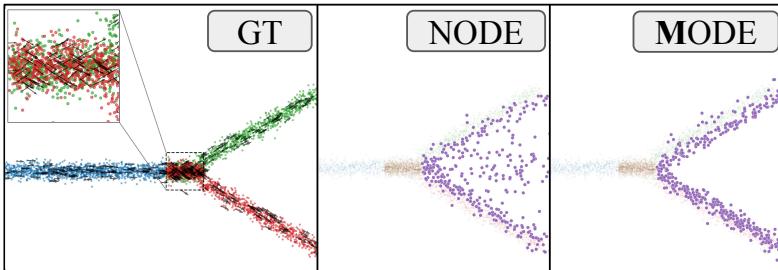
054
055
056
057
058
059
060
061
062
063

Figure 1: *Branching 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; green cells with up-slanting velocities; red, with down-slanting). (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.

equation modeling, 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. Unlike previous switching/hybrid models, MODE jointly learns dynamical regressors and mixture assignments on snapshot data, which allows circuit learning to influence regime discovery. 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. Importantly, we show that the number of experts can also be learned throughout training. 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:

1. *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 how 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 that the mixture distribution can be used to quantitatively forecast the moment of cell differentiation and shed light on the transcriptional mechanisms driving this process.

1 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.,

108 2019) and other regressors like SINDy (Brunton et al., 2016), have proven powerful for learning
 109 continuous-time dynamics from data by parameterizing the derivative of hidden states through neural
 110 networks or symbolic regression. However, they inherently assume a single, smooth flow, making
 111 them fundamentally unsuitable for systems exhibiting multiple coexisting dynamical regimes or
 112 abrupt transitions between behavioral modes.

113

114 **Flow learning in computational biology.** In computational biology, flow-based methods have
 115 found particular success in modeling cellular dynamics through RNA velocity estimation (Bergen
 116 et al., 2020) and optimal transport approaches (Tong et al., 2020). Meta Flow Matching (Atanack-
 117 ovic et al., 2024) represents a recent advancement, learning vector fields on the Wasserstein man-
 118 ifold to model population-level dynamics. Recent developments in structured latent velocity mod-
 119 eling (Farrell et al., 2023) have further extended these approaches by incorporating latent variable
 120 frameworks to capture complex single-cell transcriptomic dynamics, demonstrating improved per-
 121 formance in scenarios involving temporal gene expression patterns. While these approaches capture
 122 important aspects of biological systems, their modeling of branching lineages is completely deter-
 123 ministic. In essence, such modeling of the data assumes that the fate of cells is determined very
 124 early in the differentiation process, much before the actual branching into lineages. As such, these
 125 models do not capture the stochastic nature of biological processes.

126

127 **Switching dynamical systems.** The limitations of single-flow models have motivated exten-
 128 sive research into switching dynamical systems, which explicitly model transitions between dif-
 129 ferent dynamical regimes. Classical approaches include piecewise affine models and hybrid
 130 systems (Jin et al., 2021), which segment trajectories and fit separate dynamics to each seg-
 131 ment. More recent neural approaches (Ojeda et al., 2021; Seifner & Sanchez, 2023) have
 132 incorporated deep learning architectures to learn both the underlying dynamics and switch-
 133 ing mechanisms simultaneously, often employing attention mechanisms or variational infer-
 134 ence frameworks. Continuous-time switching systems (Köhls et al., 2021) extend these ideas
 135 by modeling regime transitions as Markov jump processes, providing principled probabilis-
 136 tic frameworks for handling temporal uncertainty. A related framework is found in Hybrid
 137 SINDy (Mangan et al., 2018), which pre-clusters data and fits cluster-specific dynamical
 138 regressors after the fact, enabling a switching dynamics on top of the basic SINDy framework.

139

140 While theoretically appealing for their ability to model mul-
 141 tiple dynamics simultaneously they suffer from a reliance on
 142 trajectory (not snapshot) data or from their sequential (as op-
 143 posed to joint) fitting of clusters and dynamics, which we will
 144 show below leads to suboptimal results for Hybrid SINDy.

145

146 **Mixture of experts (MoE).** Mixture methods represent a
 147 fundamental approach throughout the data sciences. Early ap-
 148 plications of MoE to dynamical systems identification (Lima
 149 et al., 2002; Weigend et al., 1995) demonstrated the ability
 150 to discover temporal regimes and avoid overfitting through
 151 soft partitioning of the input space. Unlike switching sys-
 152 tems that make hard assignments, MoE models enable smooth
 153 transitions between experts through learned gating functions,
 154 making them particularly well-suited for noisy biological data
 155 where regime boundaries are inherently ambiguous. However,
 156 traditional MoEs typically lack interpretability in their expert
 157 components, as the experts are often represented by black-box
 158 neural networks. Furthermore, these methods very often rely
 159 on the availability of time series data, which is not the case of
 160 the snapshot data which are so prevalent in cell biology.

161

162 Building on these foundations, our MODE framework unifies interpretable expert dynamics iden-
 163 tification with the flexibility of mixture models. Unlike previous approaches that treat clustering and
 164 dynamics modeling as separate tasks, MODE jointly learns spatial organization and governing equa-
 165 tions, enabling both accurate classification of heterogeneous populations and precise forecasting of
 166 regime transitions.

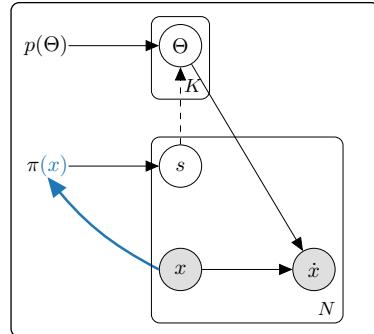


Figure 2: Plate diagram of MODE. Expert parameters Θ (with prior $p(\Theta)$) generate velocities \dot{x} via $f_{\Theta_s}(x)$ under isotropic noise; the expert distribution optionally depends on x (blue), in which case it gates states to specific experts, s .

162

2 METHODS

163

164 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
165 governing equation f_{Θ_s} parameterized by $\Theta_s \in \mathbb{R}^m$. We place a Laplace prior $\Theta_s \sim \text{Lap}(0, 1/\lambda)$
166 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
167 mixing distribution $\pi(x)$ may depend on the state x (as in the cell cycle, where oscillatory regimes
168 are localized in transcriptomic space) or may be independent of x (as in the case of heterogeneous
169 subpopulations with distinct intrinsic dynamics).
170

171 We model the probability of a state's velocity with isotropic Gaussian noise,
172

173
$$\dot{x} \mid x, s \sim \mathcal{N}(f_{\Theta_s}(x), \sigma_s^2 I_d).$$
174

175 The full generative model of the data (Fig. 2) is then
176

177
$$p(D \mid \Theta, \sigma, \pi) = \prod_{i=1}^N \sum_{s=1}^K \pi_s(x_i) \mathcal{N}(\dot{x}_i \mid f_{\Theta_s}(x_i), \sigma_s^2 I_d), \quad (1)$$
178

179 where $\sigma = (\sigma_s)_{s=1}^K$, giving the conditional negative log likelihood on \dot{x} as
180

181
$$-\log p(\dot{x} \mid x) = -\log \sum_{s=1}^K \pi_s(x) \exp\left(-\frac{1}{2\sigma_s^2} \|\dot{x} - f_{\Theta_s}(x)\|_2^2 - \frac{d}{2} \log(2\pi\sigma_s^2)\right). \quad (2)$$
182

183 Our goal is to estimate π, Θ, σ by minimizing the maximum a posteriori (MAP) objective
184

185
$$\mathcal{L}(x, \dot{x}) = -\log \sum_{s=1}^K \pi_s(x) p(\dot{x} \mid x, s) + \lambda \sum_{s=1}^K \|\Theta_s\|_1. \quad (3)$$
186

187 There are several choices for how to model f_{Θ_s} . In practice, we find that symbolic regression works
188 well: f_{Θ_s} is computed as a linear combination of basis functions, Z , depending on x , so that
189

190
$$f_{\Theta_s}(x) = Z(x)\Theta_s,$$
191

192 where Θ_s serves as the weights. Throughout, we set Z to be polynomials of order up to c , allowing
193 for explicit control as to the complexity of the fitted approximation. In this form, MODE uses a
194 SINDy-based regressor (Brunton et al., 2016) (See. App. Sec. C for details), though other estimators
195 are possible (see Sec. 4). Here, we show results in which K is known in advance, but we also
196 provide an algorithm for discovering K from data based on the Akaike (Akaike, 1974) or Bayesian
197 Information Criteria (Schwarz, 1978) (AIC/ BIC; see Sec. Specifically, the algorithm iteratively
198 evaluates whether adding or removing experts improves model fit while penalizing over-complexity.
199 Starting from an initial estimate of K , the gating network's responsibilities are used to track expert
200 usage, and redundant experts are pruned if their assignment drops below a threshold. This adaptive
201 strategy ensures robust identification of the underlying number of dynamical regimes (see Sec. C.4),
202 avoiding manual selection and reducing the risk of overfitting or underfitting.
203

204 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.2). Otherwise, we set $\pi(x)$ to be a multilayer
205 perceptron (MLP) in the manner of a gating network (Jacobs et al., 1991) and optimize Eq. 3 with
206 stochastic gradient descent. However, when optimizing using gradient descent, the gating function
207 has a tendency to converge to local minima corresponding to either a single expert or uniform
208 probabilities everywhere. Thus, we found that performance was improved by regularizing Eq. 3 with the
209 additional terms
210

211
$$H_{\text{gate}} = -\frac{1}{N} \sum_{i=1}^N \sum_{s=1}^K \pi_{i,s} \log \pi_{i,s} \quad ; \quad \text{KL}_{\text{balance}} = \sum_{s=1}^K \bar{\pi}_s \left(\log \bar{\pi}_s - \log \frac{1}{K} \right) \quad (4)$$
212

213 where H_{gate} encourages confident (low-entropy) gating decisions *per sample*, and $\text{KL}_{\text{balance}}$ prevents
214 expert collapse by balancing average usage across experts, $\bar{\pi}_s$, across the dataset (see Sec.C). This
215

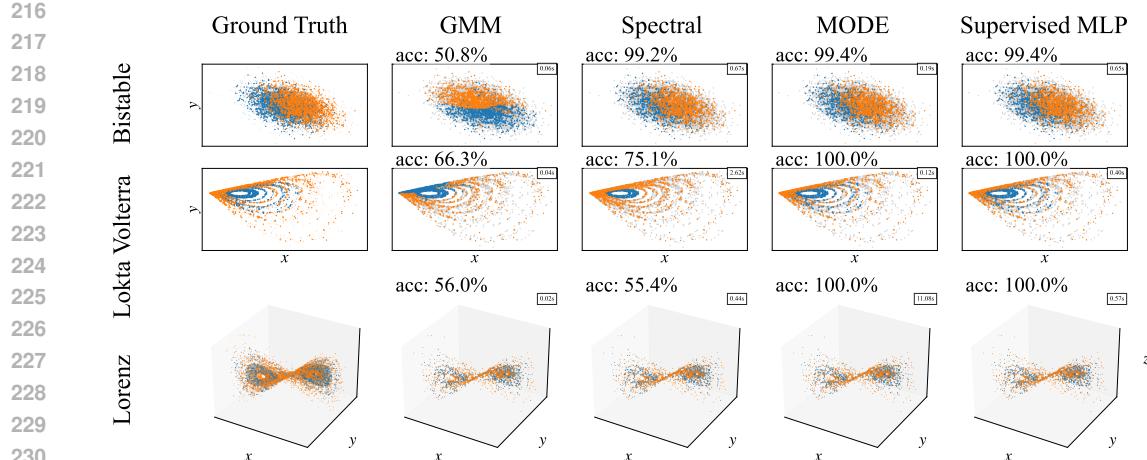


Figure 3: Elementary dynamics discovery on canonical systems. Each system challenges spatial clustering in distinct ways: **(Bistable)** spatially overlapping attractors with opposite rotational dynamics, **(Lotka-Volterra)** nonlinear predator-prey interactions creating complex phase structure, **(Lorenz)**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

combination of terms ensures that the gating network remains both selective and diverse, leveraging the strengths of individual experts while avoiding redundancy.

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

$$s_{t+1} \sim \pi(\cdot | x_t) \quad (5)$$

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

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. Alternative rollout schemes are discussed in Sec. C.

3 RESULTS

3.1 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 predator-prey 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

270 the correct number of clusters K , and evaluated using identical train-test splits (see Sec. C.2 for
 271 details). Additionally, we included a supervised MLP baseline using true cluster labels to establish
 272 an upper performance bound. After optimal cluster-to-class mapping, models were evaluated by
 273 their F1 score, recall, and precision.

Method	Metric	Bistable	Lokta Volterra	Lorenz	Average
GMM	F1	0.527 ± 0.000	1.000 ± 0.000	0.523 ± 0.000	0.683 ± 0.000
	Recall	0.527 ± 0.000	1.000 ± 0.000	0.524 ± 0.000	0.684 ± 0.000
	Precision	0.527 ± 0.000	1.000 ± 0.000	0.524 ± 0.000	0.684 ± 0.000
Spectral	F1	0.988 ± 0.000	0.765 ± 0.032	0.522 ± 0.000	0.758 ± 0.011
	Recall	0.988 ± 0.000	0.793 ± 0.018	0.523 ± 0.000	0.768 ± 0.006
	Precision	0.988 ± 0.000	0.864 ± 0.003	0.523 ± 0.000	0.792 ± 0.001
MODE	F1	0.993 ± 0.000	1.000 ± 0.000	0.927 ± 0.028	0.974 ± 0.009
	Recall	0.993 ± 0.000	1.000 ± 0.000	0.927 ± 0.028	0.974 ± 0.009
	Precision	0.993 ± 0.000	1.000 ± 0.000	0.927 ± 0.028	0.974 ± 0.009
MLP	F1	0.993 ± 0.000	1.000 ± 0.000	1.000 ± 0.000	0.997 ± 0.000
	Recall	0.993 ± 0.000	1.000 ± 0.000	1.000 ± 0.000	0.997 ± 0.000
	Precision	0.993 ± 0.000	1.000 ± 0.000	1.000 ± 0.000	0.997 ± 0.000

Table 1: Summary metrics: F1, Recall and Precision

288 We found that MODE achieved substantially higher F1 scores, recall, and precision than both unsupervised
 289 baselines and even competed with the supervised MLP (Table 1). Consistently, the GMM
 290 performed the worst, since it could only split ellipsoids in (x, \dot{x}) space, which fails dramatically for
 291 these thoroughly mixed systems. Spectral clustering could identify coherent bands from each pop-
 292 ulation, but failed when the bands dissolved into salt-and-pepper noise. MODE, on the other hand,
 293 does not rely on phase space geometry, instead clustering systems according to the best division of
 294 data into two sparse, dynamical laws. This allows it to cluster populations which look like blobs,
 295 bands or true mixtures (Lorenz).

296 Thus, the improvement is most pronounced in cases where purely geometrical clustering fundamen-
 297 tally fails—precisely the scenarios most relevant to biological applications. For noise and sample
 298 size benchmarking, see Sec. B.2.

300 3.2 FORECASTING SYNTHETIC BIOLOGICAL SWITCHES

302 Although Sec. 3.1 demonstrates that MODE can be used to cluster dynamical agents with hetero-
 303 geneous governing equations, it could be argued that this does not demonstrate MODE’s value for
 304 modeling dynamics *per se*. After all, a key problem arising in computational biology and elsewhere
 305 is not just classifying snapshot data, but also forecasting its behavior into the future.

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

313 For cell cycle data, we used a classic three-variable model of mitotic oscillation from Goldbeter
 314 (Goldbeter, 1991) given by fluctuating levels of cyclin, maturation promoting factor (MPF) and
 315 protease (Fig. 4, top, first panel). We set the probability of cell cycle exit (differentiation, Fig. 4
 316 top, first panel, orange branch) to be $p = 0.15$ in a small region on the cycle where cyclin was
 317 minimal. Cells which randomly committed to leaving the cycle in this region evolved according to
 318 a linear dynamics towards a stable node attractor (green X) placed about one cycle radius away. We
 319 sampled data uniformly in arc-length around both the cycle and exit path so that slow parts of the
 320 cycle would not be overrepresented. This resulted in 6320 positions and velocities separated into an
 321 80-20 train-validation split.

322 For branching lineage data, we used a simple two-dimensional switching system (Fig. 4, bottom,
 323 first panel) similar to those used in dynamic optimal transport and flow matching studies of transcrip-
 324 tomic dynamics (Farrell et al., 2023; Zhang et al., 2024). A Gaussian blob of cells was advanced

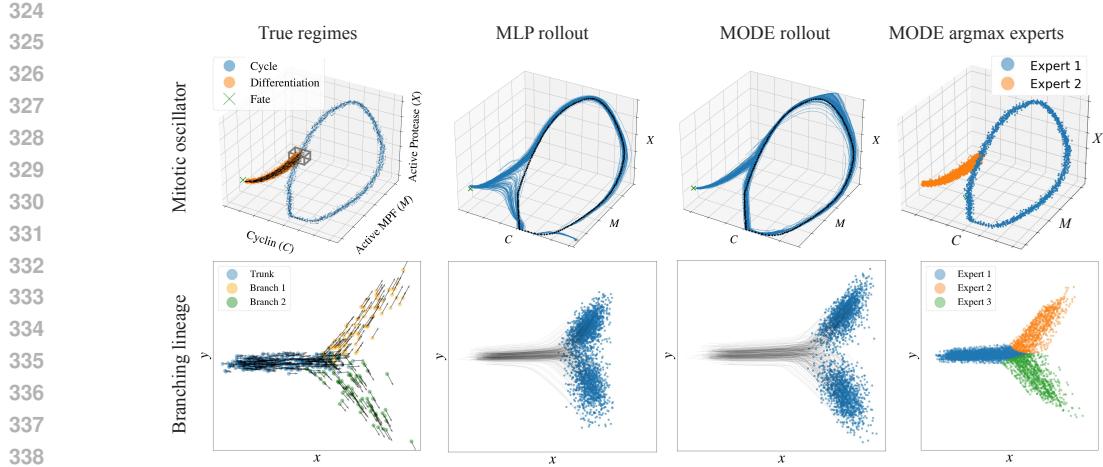


Figure 4: *MODE* vs *MLP* baseline on synthetic biological switches. (Top row) Cells evolving according to the Goldbeter oscillator (first panel, blue) could escape ($p = 0.15$ per time step) to a differentiated state (green X) when they pass through a region of low cyclin (gray box). The MLP (pictured, 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 (fourth panel).

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.

(a) Mitotic oscillator					
	W_1	W_2	$W_{1,C}$	$W_{1,M}$	$W_{1,X}$
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
Hybrid SINDy	0.1097 ± 0.0003	0.1899 ± 0.0005	0.0312 ± 0.0001	0.0593 ± 0.0002	0.0706 ± 0.0003
MODE-NN	0.0913 ± 0.0003	0.1788 ± 0.0005	0.0294 ± 0.0001	0.0392 ± 0.0002	0.0695 ± 0.0003
MODE-Symb	0.0837 ± 0.0005	0.1049 ± 0.0010	0.0590 ± 0.0003	0.0223 ± 0.0002	0.0289 ± 0.0002

(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
Hybrid SINDy	0.6329 ± 0.0017	0.8200 ± 0.0016	0.1304 ± 0.0006	0.6053 ± 0.0015
MODE-NN	0.6124 ± 0.0017	0.8030 ± 0.0016	0.1324 ± 0.0006	0.5853 ± 0.0015
MODE-Symb	0.5713 ± 0.0017	0.7689 ± 0.0015	0.1363 ± 0.0007	0.5452 ± 0.0014

Table 2: Comparative performance of *MODE* vs baselines (see main text) 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 x and a neural gating function was used, modeled as an MLP with one hidden layer. For baselines, we compared NODE with symbolic regressors (MODE-symb) to a (1) 4-layer MLP, (2) SINDy with cubic (linear) terms for the cycle (branching), (3) Hybrid SINDy (Mangan et al., 2018) using these same symbolic hyperparameters and MODE-NN in which the dynamical regressors were neural

ODEs. 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.3).

We found that MODE-Symb outperformed the baseline models on both tasks (Table 2). On the cycle data, the 2-expert MODE-Symb 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. 13a). Single flow baselines (MLP, SINDy) tried to account for both the stable node and cycle (MLP, Fig. 4, top, second panel), which leads to high Wasserstein error. Hybrid SINDy struggled to find the true regimes during its pre-clustering phase and it never let cells exit the cycle. MODE-NN performed closest to MODE-Symb but seemed to suffer from the lack of sparsity constraint, which greatly aids in finding the true dynamical regimes (see App. Sec. C. Note that MODE-Symb 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, baselines only captured the general quantitative behavior of the switch, spreading noncommittal cells in the branch interior. MODE-Symb experts correctly localized the branching structure (Fig. 4, bottom right) and discovered the underlying governing equations nearly exactly (see Sec. C.3).

3.3 FORECASTING CELL DEVELOPMENT FATE FROM scRNASEQ DATA

The synthetic cellular processes examined in Sec. 3.2 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 from Mahdessian et al. in order to investigate the spatiotemporal dynamics of human cell cycle expression, chosen because this data contains a precise labeling for cells that are part of the cell-cycle versus those that are not. These labels give us a ground truth for two dynamical regimes that mimic our synthetic Goldbeter system of Sec. 3.2: the cycle regime, and the cycle exit regime. In Appendix Sec. B.6, we demonstrate similar results on an additional dataset of human fibroblasts.

Following a standard preprocessing protocol (Zheng et al., 2023), we used scVelo (Bergen et al., 2020) to estimate the local dynamics of each cell and gene, in the form of an RNA velocity (full details can be found in Appendix Sec. D.1). The gene expression data and the inferred velocities were projected into a 5-dimensional space using principal component analysis (PCA) for fitting (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. Additionally, Fig. 5d illustrates the stability of this segmentation, showing that all runs consistently converge to the correct number of regimes (cycling vs. differentiation). 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.

We also examined on this data whether the number of experts can be chosen adaptively during training. A metric relying on the Akaike Information Criterion (AIC) achieved by each expert was used in order to change the number of experts used (full details in Appendix Sec. C.4). Fig. 5e shows that our adaptive algorithm consistently converges at 2 experts, and that the AIC flattens out before the selection (Fig. 5f), indicating that it is a stable choice.

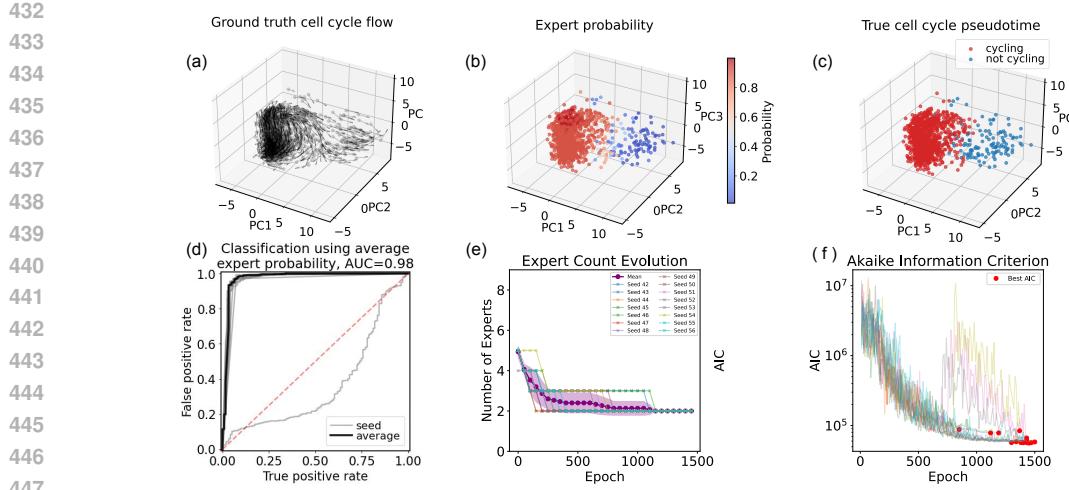


Figure 5: *Modeling cell-cycle and differentiation dynamics using MODE.* (a) Single cell RNA sequencing data from the U2OS bone cancer line was preprocessed for RNA velocity with scVelo (first three PCA dimensions shown). (b) We fit a 2-expert MODE to the first five PCA dimensions, revealing a sharp 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) Evolution of expert count over training epochs for 15 independent runs (colored lines) starting from $K = 5$ experts. The mean trajectory (purple line with markers) and confidence band (shaded region) show convergence to $K = 2$ experts, with most runs stabilizing by epoch 500. Individual runs demonstrate consistent adaptation despite different random initializations. (f) Akaike Information Criterion (AIC) trajectories for all runs. Each curve is truncated at its minimum AIC value (red dots), indicating the optimal model complexity. The log-scale y-axis shows rapid initial decrease followed by stabilization, confirming that the adaptive algorithm successfully identifies the optimal number of experts without manual tuning.

3.4 DISCOVERING CIRCUIT TRANSITIONS DURING CELL-CYCLE EXIT

To validate the biological relevance of MODE’s learned dynamics, we analyzed expert-specific gene expression signatures and regulatory network rewiring during cell cycle exit. Differential expression analysis on experts (based on the analysis of the weighted Jacobian of the experts, see Appendix Sec. D.1) revealed that Expert 0 (cycling cells) upregulated canonical proliferation markers including TOP2A, CDK1, and AURKA (Malumbres & Barbacid, 2009), while Expert 1 (exiting cells) showed enrichment of exit-associated genes such as CAV1, THBS1, and CCN1(Coller et al., 2006) (Table 12 in Appendix Sec. D.1). Known G2/M markers (CCNB1, CDC20, CDK1) exhibited significantly higher expression in Expert 0, while exit markers (CAV1, CDKN1A) were enriched in Expert 1(Coller et al., 2006), confirming that MODE captures the known cycle exit phenotype.

Ordering cells by Expert 0 probability revealed coherent temporal expression patterns (Fig. 6A). Along this pseudotemporal trajectory, proliferation markers gradually decreased while exit markers increased, with the steepest changes occurring at the transition point defined by maximal expert entropy. This sequential pattern suggests coherent regulatory changes during exit that align with known molecular markers for cell cycle exit (Spencer et al., 2013).

To identify mechanistic drivers, we extracted gene regulatory networks (GRNs) from the experts by computing learned Jacobian matrices for data before and after cycle exit (see App. D.1 for methodology). GRN topology underwent substantial rewiring across transition phases (Fig. 6B). Before transition, the network featured prominent interactions among proliferative genes (CDK1–CCNB1–CDC20 module) (Malumbres & Barbacid, 2009). After transition, exit genes (CAV1, THBS1) gained regulatory centrality while proliferative interactions weakened. Notably, CCNB1 interactions shifted from promoting CDK1 activity to engaging exit-associated pathways, consistent with known checkpoint mechanisms(Pines, 1995).

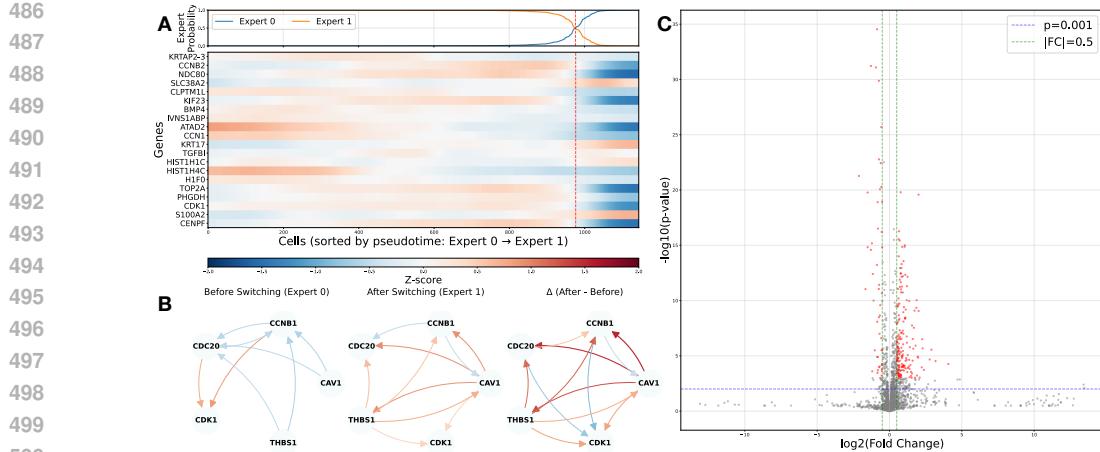


Figure 6: *Biological validation of cell cycle exit dynamics.* **A** Temporal heatmap of top differentially expressed genes along pseudotime (ordered by Expert 0 probability). Upper panel: expert probabilities; lower panel: gene expression z-scores. Red dashed line: transition point (maximum entropy). Proliferation markers decrease while exit markers increase. **B** Gene regulatory network topology before, during, and after transition. Circular graphs show key gene interactions (edge thickness: strength; color: sign). Bottom: differential connectivity. **C** Volcano plot of differential expression in high- vs low-entropy regions. Red: significant genes ($|FC| > 0.5, p < 0.001$). Dashed lines: thresholds.

Analysis of high-entropy cells associated with the moment of transition (top 20% by gating uncertainty) revealed 222 genes with significant expression changes ($|FC| > 0.5, p < 0.001$, Fig. 6C). Top transition-specific genes included *CDC20B* and *PDK4*, which differed from expert-defining markers, suggesting a distinct transcriptional state during fate resolution. These results provide evidence that MODE recovers both cell cycle regulators and mechanistic details of exit dynamics.

4 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. Furthermore, we are intrigued by the incorporation of augmented dynamical variables, like those in augmented neural ODEs (Dupont et al., 2019), which could be used in scRNAseq data to model effects like hidden proteomic influences during cell cycle dynamics.

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.

540 REFERENCES
541

542 Alexander Aivazidis, Fani Memi, Vitalii Kleshchevnikov, Sezgin Er, Brian Clarke, Oliver Stegle,
543 and Omer Ali Bayraktar. Cell2fate infers rna velocity modules to improve cell fate prediction.
544 *Nature Methods*, 22(4):698–707, 2025. doi: 10.1038/s41592-025-02608-3.

545 Hirotugu Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic
546 Control*, 19(6):716–723, 1974. doi: 10.1109/TAC.1974.1100705.
547

548 Lazar Atanackovic, Xi Zhang, Brandon Amos, Mathieu Blanchette, Leo J. Lee, Yoshua Bengio,
549 Alexander Tong, and Kirill Neklyudov. Meta Flow Matching: Integrating Vector Fields on
550 the Wasserstein Manifold, August 2024. URL <http://arxiv.org/abs/2408.14608>.
551 arXiv:2408.14608 [cs] version: 1.

552 Alon Bar, Ron Moran, Netta Mendelsohn-Cohen, Yael Korem Kohanim, Avi Mayo, Yoel Toledano,
553 and Uri Alon. Pregnancy and postpartum dynamics revealed by millions of lab tests. *Science
554 Advances*, 11(13):eadr7922, 2025.

555 Volker Bergen, Marius Lange, Stefan Peidli, F Alexander Wolf, and Fabian J Theis. Generalizing
556 rna velocity to transient cell states through dynamical modeling. *Nature biotechnology*, 38(12):
557 1408–1414, 2020.

558 Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations
559 from data by sparse identification of nonlinear dynamical systems. *Proceedings of the Na-
560 tional Academy of Sciences*, 113(15):3932–3937, 2016. doi: 10.1073/pnas.1517384113. URL
561 <https://www.pnas.org/doi/10.1073/pnas.1517384113>.
562

563 Thomas M. Bury, Daniel Dylewsky, Chris T. Bauch, Madhur Anand, Leon Glass, Tsvi Tlusty,
564 and Bhargav R. Karamched. Predicting discrete-time bifurcations with deep learning. *Na-
565 ture Communications*, 14(1):6599, 2023. doi: 10.1038/s41467-023-42020-z. URL <https://doi.org/10.1038/s41467-023-42020-z>.
566

567 Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural Ordinary
568 Differential Equations. December 2019. doi: 10.48550/arXiv.1806.07366. URL <http://arxiv.org/abs/1806.07366>. arXiv:1806.07366 [cs].
569

570 Hilary A Coller, Liyun Sang, and James M Roberts. A New Description of Cellular Quiescence.
571 *PLoS Biology*, 4(3):e83, March 2006. ISSN 1545-7885. doi: 10.1371/journal.pbio.0040083.
572 URL <https://dx.plos.org/10.1371/journal.pbio.0040083>.
573

574 Antonio C Costa, Tosif Ahamed, and Greg J Stephens. Adaptive, locally linear models of complex
575 dynamics. *Proceedings of the National Academy of Sciences*, 116(5):1501–1510, 2019.

576 Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented neural odes. *Advances in Neural
577 Information Processing Systems*, 32, 2019. URL <https://arxiv.org/abs/1904.01681>.
578 NeurIPS 2019.

579 Spencer Farrell, Christoph Bönisch, Melanie L. Goyal, et al. Inferring single-cell transcriptomic
580 dynamics with structured latent-velo. *Cell Reports Methods*, 3(10):100581, 2023. doi: 10.1016/
581 j.crmeth.2023.100581.

582 Albert Goldbeter. A minimal cascade model for the mitotic oscillator involving cyclin and cdc2
583 kinase. *Proceedings of the National Academy of Sciences of the United States of America*, 88
584 (20):9107–9111, 1991. doi: 10.1073/pnas.88.20.9107.

585 Doron Haviv, Aram-Alexandre Pooladian, Dana Pe’er, and Brandon Amos. Wasserstein flow match-
586 ing: Generative modeling over families of distributions with single-cell genomics applications.
587 2024. URL <https://arxiv.org/abs/2411.00698>. arXiv preprint arXiv:2411.00698.
588

589 Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. Adaptive mixtures
590 of local experts. *Neural Computation*, 3(1):79–87, 1991.
591

594 Xiangyu Jin, Jie An, Bohua Zhan, Naijun Zhan, and Miaomiao Zhang. Inferring Switched Nonlinear
 595 DynamicalSystems. *Formal Aspects of Computing*, 33(3):385–406, June 2021. ISSN 0934-5043,
 596 1433-299X. doi: 10.1007/s00165-021-00542-7. URL <https://dl.acm.org/doi/10.1007/s00165-021-00542-7>.
 597

598 Mikail Khona and Ila R Fiete. Attractor and integrator networks in the brain. *Nature Reviews
 599 Neuroscience*, 23(12):744–766, 2022.
 600

601 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International
 602 Conference on Learning Representations (ICLR)*, 2015. URL <https://arxiv.org/abs/1412.6980>.
 603

604 Lukas Köhs, Bastian Alt, and Heinz Koeppl. Variational Inference for Continuous-Time Switching
 605 Dynamical Systems. In *Advances in Neural Information Processing Systems*, volume 34, pp.
 606 20545–20557. Curran Associates, Inc., 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/abec16f483abb4f1810ca029aadf8446-Abstract.html>.
 607

608

609 C.Ap.M. Lima, A.L.V. Coelho, and F.J. Von Zuben. Mixture of experts applied to nonlinear dy-
 610 namic systems identification: a comparative study. In *VII Brazilian Symposium on Neural Net-
 611 works, 2002. SBRN 2002. Proceedings.*, pp. 162–167, November 2002. doi: 10.1109/SBRN.
 612 2002.1181463. URL <https://ieeexplore.ieee.org/document/1181463>.
 613

614

615 Zaikang Lin, Sei Chang, Aaron Zweig, Minseo Kang, Elham Azizi, and David A. Knowles. Inter-
 616 pretable neural odes for gene regulatory network discovery under perturbations. *arXiv preprint
 617 arXiv:2501.02409*, 2025. doi: 10.48550/arXiv.2501.02409. URL <https://arxiv.org/abs/2501.02409>. v3.
 618

619 Diana Mahdessian, Anthony J Cesnik, Christian Gnann, Frida Danielsson, Lovisa Stenström,
 620 Muhammad Arif, Cheng Zhang, Trang Le, Fredric Johansson, Rutger Schutten, et al. Spatiotem-
 621 poral dissection of the cell cycle with single-cell proteogenomics. *Nature*, 590(7847):649–654,
 622 2021.

623

624 Marcos Malumbres and Mariano Barbacid. Cell cycle, CDKs and cancer: a changing paradigm.
 625 *Nature Reviews Cancer*, 9(3):153–166, March 2009. ISSN 1474-175X, 1474-1768. doi: 10.
 626 1038/nrc2602. URL <https://www.nature.com/articles/nrc2602>.
 627

628 Niall M Mangan, Travis Askham, Steven L Brunton, J Nathan Kutz, and Joshua L Proctor. Model se-
 629 lection for hybrid dynamical systems via sparse regression. *arXiv preprint arXiv:1808.03251*, aug
 630 2018. doi: 10.48550/arXiv.1808.03251. URL <https://arxiv.org/abs/1808.03251>.
 631 22 pages, 5 figures, 1 table. Code repository: https://github.com/niallmangan/hybrid_sindy.
 632

633 Noa Muriel, Matthew Ricci, and Mor Nitzan. Let’s do the time-warp-attend: Learning topological
 634 invariants of dynamical systems. In *International Conference on Representation Learning (ICLR
 635 2024)*, 2024. doi: 10.48550/arXiv.2312.09234. URL <https://arxiv.org/abs/2312.09234>. arXiv:2312.09234 [cs.LG].
 636

637 Cesar Ojeda, Bogdan Georgiev, Kostadin Cvejoski, Jannis Schucker, Christian Bauckhage, and
 638 Ramses J. Sanchez. Switching Dynamical Systems with Deep Neural Networks. In *2020
 639 25th International Conference on Pattern Recognition (ICPR)*, pp. 6305–6312, Milan, Italy, Jan-
 640 uary 2021. IEEE. ISBN 978-1-7281-8808-9. doi: 10.1109/ICPR48806.2021.9412566. URL
 641 <https://ieeexplore.ieee.org/document/9412566/>.
 642

643 J Pines. Cyclins and cyclin-dependent kinases: a biochemical view. *Biochemical
 644 Journal*, 308(3):697–711, June 1995. ISSN 0264-6021, 1470-8728. doi: 10.1042/bj3080697. URL <https://portlandpress.com/biochemj/article/308/3/697/31827/Cyclins-and-cyclin-dependent-kinases-a-biochemical>.
 645

646

647 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.

648 Andrea Riba, Attila Oravecz, Matej Durik, Sara Jiménez, Violaine Alunni, Marie Cerciat, Matthieu
 649 Jung, Céline Keime, William M Keyes, and Nacho Molina. Cell cycle gene regulation dynamics
 650 revealed by rna velocity and deep-learning. *Nature communications*, 13(1):2865, 2022.

651 Nicolas Romeo, Chris Chi, Aaron R. Dinner, and Elizabeth R. Jerison. Characterizing nonlinear
 652 dynamics by contrastive cartography. *PRX Life*, 3(3):033022, 2025. doi: 10.1103/jc9p-m3rn.
 653 URL <https://doi.org/10.1103/jc9p-m3rn>.

654 William M Schaffer and Mark Kot. Do strange attractors govern ecological systems? *BioScience*,
 655 35(6):342–350, 1985.

656 Gideon Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464,
 657 1978. doi: 10.1214/aos/1176344136.

658 Patrick Seifner and Ramses J. Sanchez. Neural Markov Jump Processes, May 2023. URL <http://arxiv.org/abs/2305.19744>. arXiv:2305.19744 [cs].

659 Sabrina L. Spencer, Steven D. Cappell, Feng-Chiao Tsai, K. Wesley Overton, Clifford L. Wang,
 660 and Tobias Meyer. The Proliferation-Quiescence Decision Is Controlled by a Bifurcation in
 661 CDK2 Activity at Mitotic Exit. *Cell*, 155(2):369–383, October 2013. ISSN 00928674. doi:
 662 10.1016/j.cell.2013.08.062. URL <https://linkinghub.elsevier.com/retrieve/pii/S0092867413011434>.

663 Alex Tong et al. Trajectorynet: A dynamic optimal transport network for single-cell time series. In
 664 *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pp. 7626–7635,
 665 2020.

666 Dongyi Wang, Yuanwei Jiang, Zhenyi Zhang, Xiang Gu, Peijie Zhou, and Jian Sun. Joint velocity-
 667 growth flow matching for single-cell dynamics modeling. *arXiv preprint arXiv:2505.13413*, 2025.
 668 doi: 10.48550/arXiv.2505.13413. URL <https://arxiv.org/abs/2505.13413>. v1.

669 A. S. Weigend, M. Mangeas, and A. N. Srivastava. Nonlinear gated experts for time series: discov-
 670 ering regimes and avoiding overfitting. *International Journal of Neural Systems*, 6(4):373–399,
 671 December 1995. ISSN 0129-0657. doi: 10.1142/s0129065795000251.

672 Zhenyi Zhang, Tiejun Li, and Peijie Zhou. Learning stochastic dynamics from snapshots through
 673 regularized unbalanced optimal transport. *arXiv preprint arXiv:2410.00844*, 2024. doi: 10.48550/
 674 arXiv.2410.00844. URL <https://arxiv.org/abs/2410.00844>. Published as a confer-
 675 ence paper at ICLR 2025 (oral).

676 Shijie C Zheng, Genevieve Stein-O’Brien, Leandros Boukas, Loyal A Goff, and Kasper D Hansen.
 677 Pumping the brakes on rna velocity by understanding and interpreting rna velocity estimates.
 678 *Genome biology*, 24(1):246, 2023.

685 A SYNTHETIC SYSTEMS

686 A.1 DISCOVERING ELEMENTARY DYNAMICS IN SPATIALLY OVERLAPPING SYSTEMS

687 To systematically evaluate MODE’s ability to discover elementary dynamics in complex systems,
 688 we constructed three canonical test cases that exhibit spatially overlapping regimes with distinct
 689 underlying dynamics (Figure 7). These synthetic systems serve as controlled benchmarks where
 690 ground truth cluster assignments are known, enabling rigorous quantitative assessment of clustering
 691 performance. Each system represents a different class of dynamical complexity commonly encoun-
 692 tered in biological and physical applications, specifically designed to challenge spatial clustering
 693 methods while providing rich velocity information for dynamics-based approaches.

694 The **bistable system** (Figure 7, left panels) models two competing attractors that create spatially
 695 overlapping but dynamically distinct regimes. Data points are sampled from Gaussian distributions
 696 centered at $(-0.5, 0)$ and $(0.5, 0)$ with spread $\sigma = 1.0$, resulting in substantial spatial overlap be-
 697 tween the two dynamical modes. The dynamics for each attractor follow:

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

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

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

708 **The Lotka-Volterra system** (Figure 7, middle panels) implements predator-prey dynamics with two
 709 different parameter regimes, each generating distinct oscillatory behaviors. We simulate trajectories
 710 from 20 uniformly sampled initial conditions over the domain $[10, 50] \times [10, 50]$ and integrate each
 711 system for 100 time units. The two dynamical regimes are governed by:
 712

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

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

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

718 The **Lorenz system** (Figure 7, right panels) represents chaotic dynamics with two distinct parameter
 719 sets, each producing different attractor geometries. We generate trajectories from 20 random initial
 720 conditions uniformly distributed over $[-15, 15] \times [-15, 15] \times [0, 40]$ and integrate for 10 time units,
 721 discarding the first 1000 time steps to eliminate transients. The two chaotic regimes follow:
 722

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

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

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

727 For all systems, we add Gaussian noise with standard deviation $\sigma = 0.1$ to both position and velocity
 728 measurements to simulate realistic experimental conditions. Each dataset contains $n = 10,000$ data
 729 points equally distributed between the two dynamical regimes. The resulting datasets exhibit the
 730 key challenge that motivates MODE: spatial clustering methods fail due to overlapping support,
 731 while velocity information provides the crucial distinguishing signal needed for accurate regime
 732 identification. This experimental design directly tests MODE’s central hypothesis that incorporating
 733 local dynamics can resolve ambiguities that spatial methods cannot handle.
 734

741 A.2 FORECASTING SYNTHETIC BIOLOGICAL SWITCHES

742 To evaluate MODE’s forecasting capabilities on biologically relevant switching dynamics, we con-
 743 structed two synthetic datasets that capture fundamental processes in cellular biology: cell cycle
 744 dynamics with stochastic exit events and lineage branching processes. These systems represent
 745 challenging forecasting scenarios where single-flow models fundamentally fail due to the presence
 746 of discrete regime transitions and branching trajectories that violate the assumptions of continuous
 747 dynamical systems.
 748

749 **Cell cycle dataset with differentiation exit.** The cell cycle dataset implements the classic three-
 750 variable Goldbeter oscillator model (Goldbeter, 1991), which describes mitotic oscillations through
 751 the temporal dynamics of cyclin concentration, maturation promoting factor (MPF), and protease
 752 activity. The Goldbeter model captures essential features of cell cycle regulation through a system
 753 of nonlinear ordinary differential equations that generate stable limit cycle behavior representing
 754 normal cell division cycles. These are given by:
 755

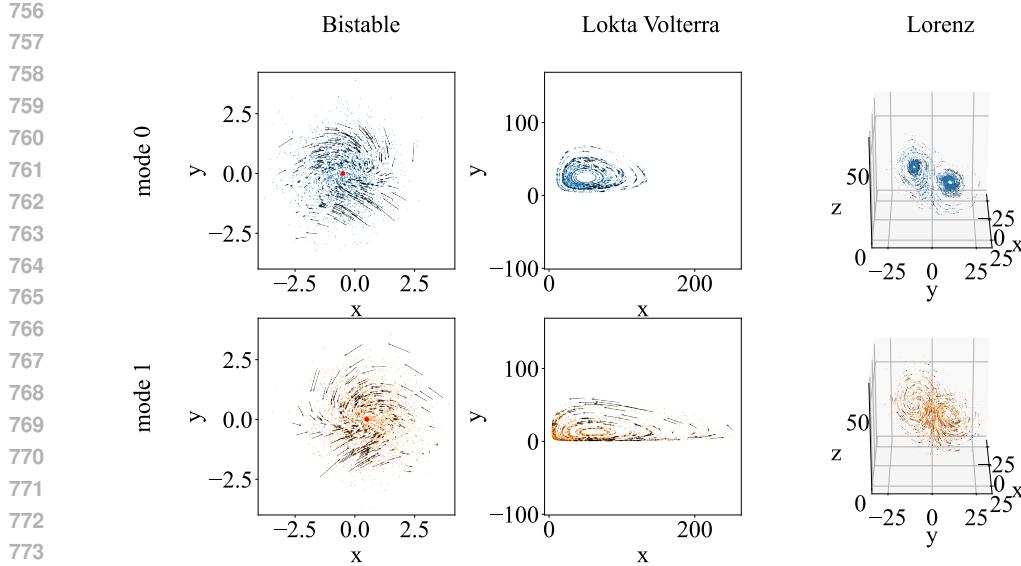


Figure 7: **Synthetic datasets for elementary dynamics discovery benchmarks.** Phase portraits showing spatially overlapping dynamical regimes for three canonical test systems. Each column represents a different system (bistable, Lotka-Volterra, Lorenz), with Mode 0 (top row, blue) and Mode 1 (bottom 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.

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

$$\frac{dM}{dt} = \frac{V_1(1 - M)}{K_1 + (1 - M)} - \frac{V_2 M}{K_2 + M}, \quad V_1 = \frac{V_{M1} C}{K_c + C} \quad (14)$$

$$\frac{dX}{dt} = \frac{V_3(1 - X)}{K_3 + (1 - X)} - \frac{V_4 X}{K_4 + X}, \quad V_3 = V_{M3} M \quad (15)$$

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

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

$$\frac{dy}{dt} = -K (y - x^*), \quad (16)$$

where

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

Symbol	Code parameter	Value
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.

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 over-sampling 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

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

These were evolved according to the system

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

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

These new populations evolved according to

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

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 80–20 train-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 non-linear 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 ± 10^1 and -1.59 ± 10^1) deviate substantially from the true values of 0.

918
 919 Table 4: Equation recovery for the bistable system. Each cell shows the true coefficient, estimated
 920 mean \pm variance over 10 runs, and absolute error. MODE achieves near-perfect recovery with errors
 921 at the level of $\mathcal{O}(10^{-6})$.

	1	$\hat{1}$	x	\hat{x}	y	\hat{y}
mode 0 dx/dt	-0.5	-0.493 ± 10^{-5}	-1	-0.988 ± 10^{-5}	2	1.98 ± 10^{-5}
mode 0 dy/dt	-0.25	-0.248 ± 10^{-5}	-0.5	-0.497 ± 10^{-5}	-1.5	-1.49 ± 10^{-5}
mode 1 dx/dt	0.5	0.497 ± 10^{-5}	-1	-0.992 ± 10^{-6}	-2	-1.98 ± 10^{-5}
mode 1 dy/dt	-0.25	-0.244 ± 10^{-5}	0.5	0.491 ± 10^{-5}	-1.5	-1.48 ± 10^{-4}
	x^2	\hat{x}^2	xy	$\hat{x}y$	y^2	\hat{y}^2
mode 0 dx/dt	0	-0.000152 ± 10^{-5}	0	0.00262 ± 10^{-5}	0	-0.00307 ± 10^{-6}
mode 0 dy/dt	0	0.000716 ± 10^{-5}	-1	-0.987 ± 10^{-5}	0	-0.00317 ± 10^{-6}
mode 1 dx/dt	0	0.00474 ± 10^{-6}	0	-0.000486 ± 10^{-5}	0	-0.000675 ± 10^{-5}
mode 1 dy/dt	0	-0.00102 ± 10^{-5}	1	0.98 ± 10^{-5}	0	-0.000296 ± 10^{-5}

933
 934 Table 5: Equation recovery for the Lotka-Volterra predator-prey system. MODE correctly identifies
 935 distinct predation rates and mortality parameters that characterize each ecological regime.

	1	$\hat{1}$	x	\hat{x}	y	\hat{y}
mode 0 $dx/dt/dt$	0	-0.0281 ± 10^{-4}	0.5	0.501 ± 10^{-7}	0	0.000761 ± 10^{-6}
mode 0 $dy/dt/dt$	0	-0.006 ± 10^{-4}	0	$-2.17e-05 \pm 10^{-7}$	-0.5	-0.5 ± 10^{-7}
mode 1 $dx/dt/dt$	0	0.00661 ± 10^{-4}	0.5	0.499 ± 10^{-7}	0	0.000474 ± 10^{-7}
mode 1 $dy/dt/dt$	0	0.00118 ± 10^{-5}	0	$-4.09e-05 \pm 10^{-8}$	-0.6	-0.6 ± 10^{-7}
	x^2	\hat{x}^2	xy	$\hat{x}y$	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 1 $dx/dt/dt$	0	$3.02e-06 \pm 10^{-11}$	-0.04	-0.04 ± 10^{-10}	0	$-3.12e-05 \pm 10^{-9}$
mode 1 $dy/dt/dt$	0	$6.79e-07 \pm 10^{-12}$	0.01	0.00999 ± 10^{-11}	0	$4.6e-06 \pm 10^{-10}$

949
 950 These larger estimation errors in the Lorenz system reflect the inherent challenges of applying sparse
 951 regression to chaotic data. The sensitive dependence on initial conditions and the presence of mul-
 952 tiple time scales create difficulties in cleanly separating trajectories between modes, leading to co-
 953 efficient bias when the clustering assignment itself contains uncertainty. Despite these limitations,
 954 MODE successfully captures the essential nonlinear structure that governs the chaotic dynamics.

955
 956 **Statistical consistency and limitations.** All results represent averages over 10 independent runs
 957 with different random initializations, providing estimates of method reliability. For the bistable and
 958 Lotka-Volterra systems, coefficient variances remain extremely low ($\mathcal{O}(10^{-7})$ or smaller), indicat-
 959 ing consistent convergence across runs. The Lorenz system shows higher variance, particularly for
 960 constant terms, reflecting the added complexity of chaotic trajectory separation.

961
 962 The equation recovery demonstrates MODE’s interpretability advantage: unlike black-box cluster-
 963 ing methods, MODE provides explicit mathematical models revealing the mechanisms underlying
 964 each dynamical regime. This capability transforms trajectory clustering from pattern recogni-
 965 tion into scientific discovery, enabling researchers to understand not just which trajectories belong to-
 966 gether, but why they exhibit similar dynamics.

967 B.2 NOISE AND TRAINING SIZE ROBUSTNESS

968
 969 To assess the practical applicability of elementary dynamics decomposition, we conduct system-
 970 atic robustness studies examining how clustering performance degrades under realistic experimental
 971 conditions. Real-world dynamical systems inevitably contain measurement noise and are often con-
 972 strained by limited sample sizes, factors that can severely impact clustering methods that rely on
 973 precise velocity estimates.

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

	1	$\hat{1}$	x	\hat{x}	y	\hat{y}
mode 0 $dx/dt/dt$	0	0.0993 ± 10^{-2}	-12	-12 ± 10^{-3}	12	11.9 ± 10^{-2}
mode 0 $dy/dt/dt$	0	1.29 ± 10^1	28	28.5 ± 10^0	-1	-1.1 ± 10^{-2}
mode 0 $dz/dt/dt$	0	-0.0777 ± 10^{-3}	0	-0.653 ± 10^0	0	0.266 ± 10^{-1}
mode 1 $dx/dt/dt$	0	0.0523 ± 10^{-2}	-10	-10.2 ± 10^{-1}	10	10.1 ± 10^{-1}
mode 1 $dy/dt/dt$	0	-1.59 ± 10^1	35.6	35.4 ± 10^{-1}	-1	-1.13 ± 10^{-1}
mode 1 $dz/dt/dt$	0	-0.0963 ± 10^{-2}	0	0.553 ± 10^0	0	-0.212 ± 10^{-1}
	z	\hat{z}	x^2	\hat{x}^2	xy	\hat{xy}
mode 0 $dx/dt/dt$	0	0.0386 ± 10^{-2}	0	0.00456 ± 10^{-4}	0	0.00152 ± 10^{-5}
mode 0 $dy/dt/dt$	0	-0.299 ± 10^{-1}	0	-0.0516 ± 10^{-2}	0	0.0279 ± 10^{-3}
mode 0 $dz/dt/dt$	-4	-3.92 ± 10^{-2}	0	-0.00421 ± 10^{-4}	1	1 ± 10^{-4}
mode 1 $dx/dt/dt$	0	-0.0406 ± 10^{-2}	0	-0.00311 ± 10^{-4}	0	-0.002 ± 10^{-5}
mode 1 $dy/dt/dt$	0	0.265 ± 10^{-1}	0	0.0532 ± 10^{-2}	0	-0.0362 ± 10^{-2}
mode 1 $dz/dt/dt$	-2.67	-2.76 ± 10^{-2}	0	-0.0173 ± 10^{-3}	1	1.01 ± 10^{-3}
	xz	$\hat{x}z$	y^2	\hat{y}^2	yz	\hat{yz}
mode 0 $dx/dt/dt$	0	0.000759 ± 10^{-6}	0	-0.000665 ± 10^{-6}	0	-0.000515 ± 10^{-6}
mode 0 $dy/dt/dt$	-1	-0.999 ± 10^{-6}	0	-0.0054 ± 10^{-4}	0	0.00202 ± 10^{-5}
mode 0 $dz/dt/dt$	0	0.0094 ± 10^{-4}	0	-0.000259 ± 10^{-6}	0	-0.00324 ± 10^{-5}
mode 1 $dx/dt/dt$	0	0.000664 ± 10^{-6}	0	0.000369 ± 10^{-7}	0	-0.000614 ± 10^{-6}
mode 1 $dy/dt/dt$	-1	-1 ± 10^{-5}	0	0.011 ± 10^{-3}	0	0.00326 ± 10^{-5}
mode 1 $dz/dt/dt$	0	-0.00809 ± 10^{-4}	0	-0.00252 ± 10^{-5}	0	0.00249 ± 10^{-5}
	z^2	\hat{z}^2				
mode 0 $dx/dt/dt$	0	-0.00187 ± 10^{-5}				
mode 0 $dy/dt/dt$	0	0.0076 ± 10^{-4}				
mode 0 $dz/dt/dt$	0	0.000345 ± 10^{-6}				
mode 1 $dx/dt/dt$	0	0.00174 ± 10^{-5}				
mode 1 $dy/dt/dt$	0	-0.00652 ± 10^{-4}				
mode 1 $dz/dt/dt$	0	0.00218 ± 10^{-5}				

1007 We evaluated robustness across three canonical test systems (bistable, Lotka-Volterra, Lorenz)
 1008 by systematically varying two critical experimental parameters: Gaussian noise levels ($\sigma \in$
 1009 $\{0.0, 10^{-5}, 10^{-2}, 10^{-1}, 2 \times 10^{-1}\}$) and training dataset sizes ($N \in \{100, 800, 8000\}$). Noise is
 1010 added to both position x and velocity \dot{x} measurements after normalization ($VAR[x] = VAR[\dot{x}] =$
 1011 1) to simulate realistic experimental conditions. For each parameter combination, we evaluate clus-
 1012 tering performance using ARI and NMI metrics on held-out test data (20% of each dataset), averag-
 1013 ing results over multiple random seeds to ensure statistical reliability.

1014 Figure 8 reveals that MODE demonstrates competitive performance compared to the MLP base-
 1015 line while substantially outperforming spatial-only methods across most test conditions. MODE
 1016 shows robust performance on the bistable and Lotka-Volterra systems, maintaining high clustering
 1017 accuracy ($ARI, NMI > 0.8$) even under moderate noise levels where spatial clustering methods
 1018 fail completely. This advantage stems from MODE’s ability to leverage velocity information—even
 1019 noisy velocity estimates provide directional information that purely spatial clustering cannot access.

1020 The noise robustness patterns reveal system-specific insights: the bistable system shows the most
 1021 dramatic performance advantages for MODE over spatial methods, with MODE maintaining high
 1022 accuracy while GMM and Spectral clustering fail completely across most noise levels. This oc-
 1023 curs because the bistable system’s spatially overlapping attractors become indistinguishable without
 1024 velocity information as noise increases. The Lorenz system presents the most challenging case,
 1025 where all methods show performance degradation, though MODE and MLP maintain more robust
 1026 performance than spatial-only approaches.

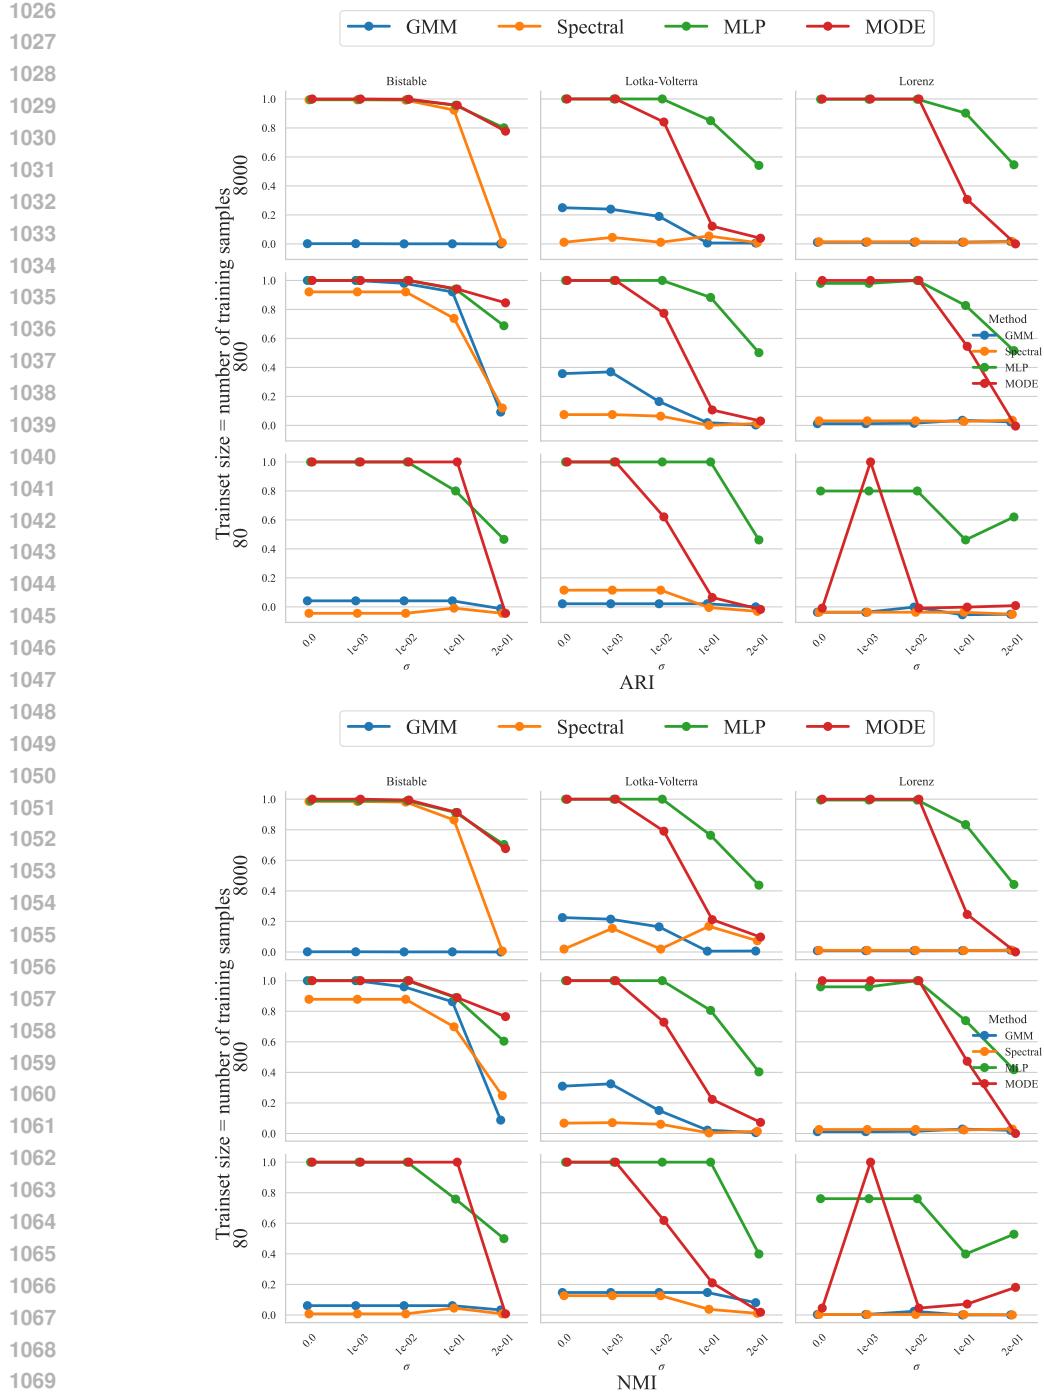


Figure 8: Clustering performance (ARI, NMI) as a function of noise level (σ) and training set size (N) across different dynamical systems. Each subplot corresponds to a specific dataset and metric. MODE demonstrates superior noise tolerance and sample efficiency compared to spatial-only methods, maintaining high performance even under challenging experimental conditions where baseline methods fail.

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,

1080 performance variability increases substantially at the smallest sample size ($N = 100$), indicating
 1081 that while MODE can work with limited data, larger sample sizes ($N \geq 800$) provide more reliable
 1082 results.

1083 The Lotka-Volterra system demonstrates MODE’s effectiveness for temporally-separated dynamics,
 1084 where spatial clustering methods consistently fail regardless of sample size. The underlying ad-
 1085 vantage is that velocity information provides crucial directional cues about cluster membership that
 1086 spatial position alone cannot capture—each (x_i, \dot{x}_i) pair encodes both current state and dynamical
 1087 tendency, enabling more robust clustering even when phase space regions overlap.

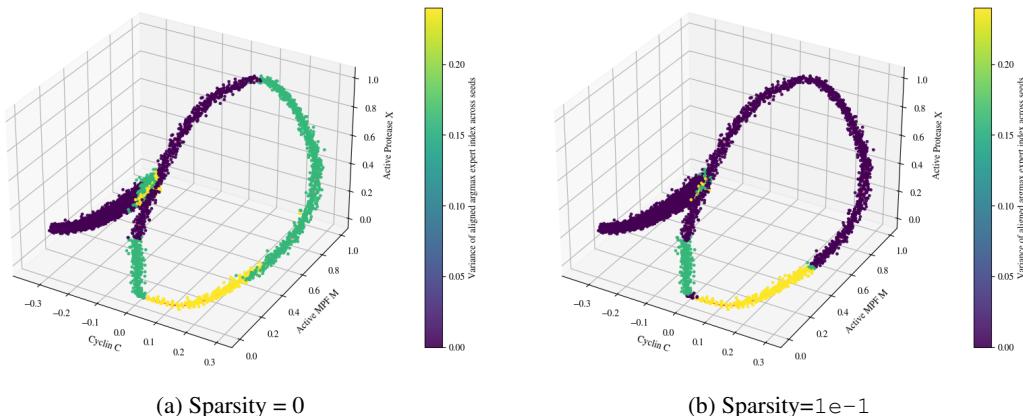
1088 These robustness results provide practical guidance for experimental applications. MODE’s noise
 1089 tolerance suggests applicability to real biological data with moderate measurement uncertainties,
 1090 though performance degrades significantly under high noise conditions ($\sigma > 0.1$). The sample size
 1091 analysis indicates that while MODE can function with limited data, reliable performance requires
 1092 adequate sample sizes ($N \geq 800$ recommended) to ensure consistent clustering results.

1093 The comparison with the MLP baseline reveals an important trade-off: while MLP often matches
 1094 or slightly exceeds MODE’s clustering performance, MODE provides the critical advantage of in-
 1095 terpretable sparse dynamics. MODE achieves competitive robustness while simultaneously offering
 1096 mechanistic insight through its recovered governing equations, making it particularly suitable for
 1097 scientific applications where understanding system structure is as important as clustering accuracy.

1099 B.3 THE ROLE OF SPARSITY AND COMPARISON TO HYBRID SINDY

1100 We found that MODE using symbolic regressors (MODE-Symb) consistently outperformed a ver-
 1101 sion with neural regressors (MODE-NN), suggesting a potential role for sparsity during learning. To
 1102 investigate this phenomenon, we reran the Goldbeter experiment with MODE-Symb while setting
 1103 the sparsity value to $\lambda = 0, 1e-4, 1e-3, 1e-2, 1e-1$. Each model was retrained with 5 random
 1104 seeds.

1105 We found that higher sparsity values encouraged more stable divisions of phase space into dynam-
 1106 ical regimes across retrainings. This is visible if we plot the variance in cluster assignment averaged
 1107 over all cells and across retrainings (Fig. 11) or view this variance across cells having retrained on
 1108 different sparsity values. This complies with our intuition that less sparse (i.e more “complex” sys-
 1109 tems) should be able to fit the dynamics in multiple ways, and more sparse systems should converge
 1110 to the same (in our experience also the true) expert distribution.



1127 Figure 9: Variance in cluster assignment after MODE training with low vs high sparsity on Gold-
 1128 beter.

1129 This highlights an important difference between MODE and comparable switching/hybrid systems
 1130 like Hybrid SINDy, where the regime assignments are learned before training. In our framework,
 1131 model discovery and mixture modeling interact through learning. This means that we can use spar-
 1132 sity, for example, to help discover the true dynamical regimes. See, for example, the comparison
 1133 between the Goldbeter regimes discovered by MODE vs Hybrid SINDy.

1134

1135

1136

1137 Ensemble MoE-SINDy: argmax expert over full dataset

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

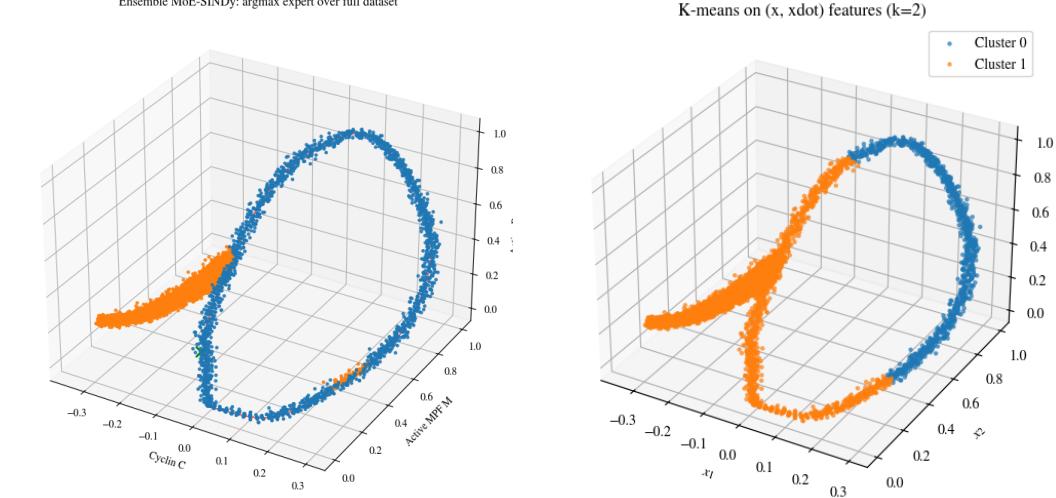
1149

1150

1151

1152

1153



(a) Most likely cluster assignment according to ensemble of 5 MODEs with high sparsity.

(b) Most likely cluster assignment according to ensemble of 5 Hybrid SINDys with high sparsity.

Figure 10: Jointly learning regime assignments with dynamics using MODE correctly identifies true branches (MODE, left). Pre-clustering using Hybrid SINDy systematically leads to poor clustering.

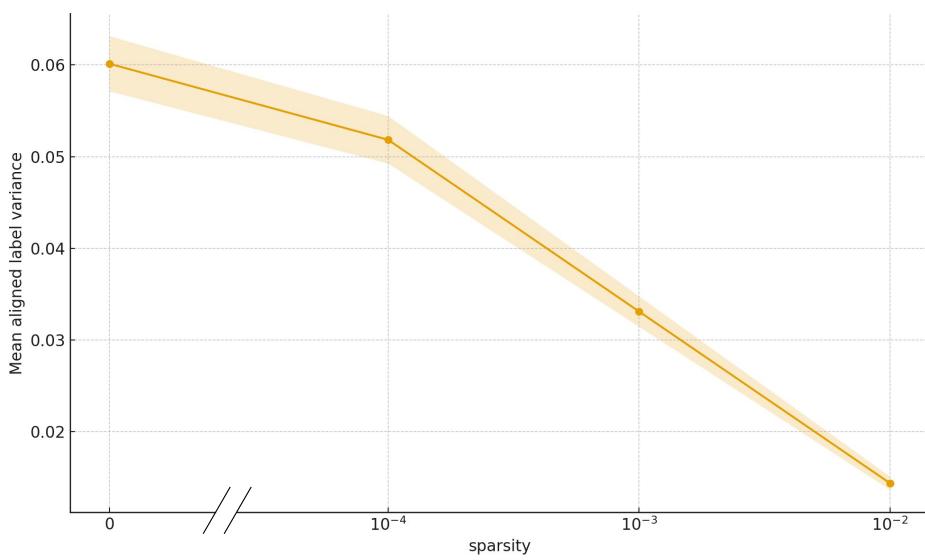


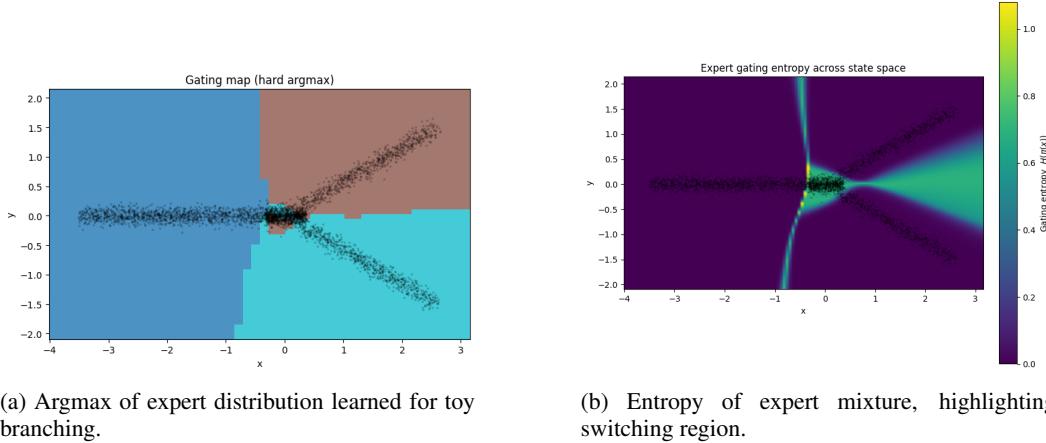
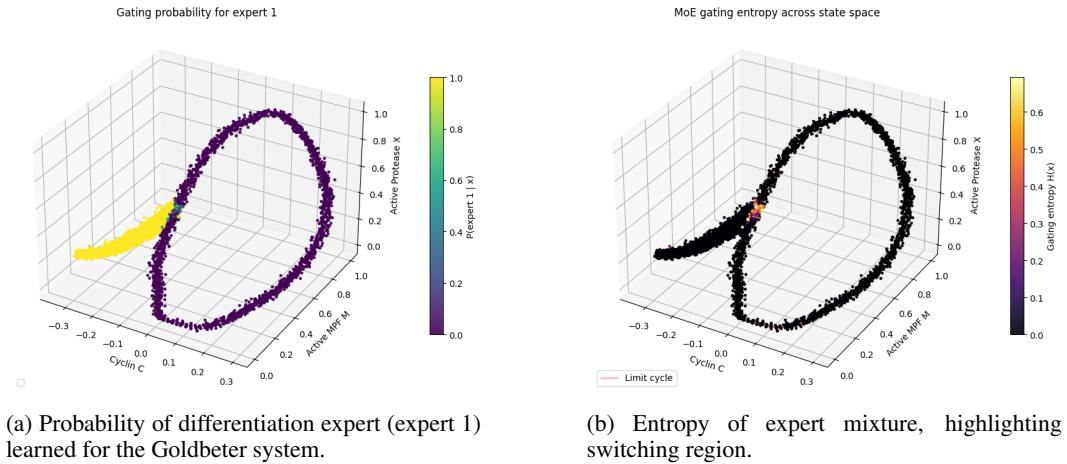
Figure 11: Variance of regime assignment with sparsity.

1184

1185

1186

1187

1188 B.4 SUPPLEMENTAL FIGURES
11891205 Figure 12: Toy branching argmax and entropy figures
12061223 Figure 13: Goldbeter probability and entropy figures
12241225 B.5 OUT-OF-DISTRIBUTION (OOD) GENERALIZATION
1226

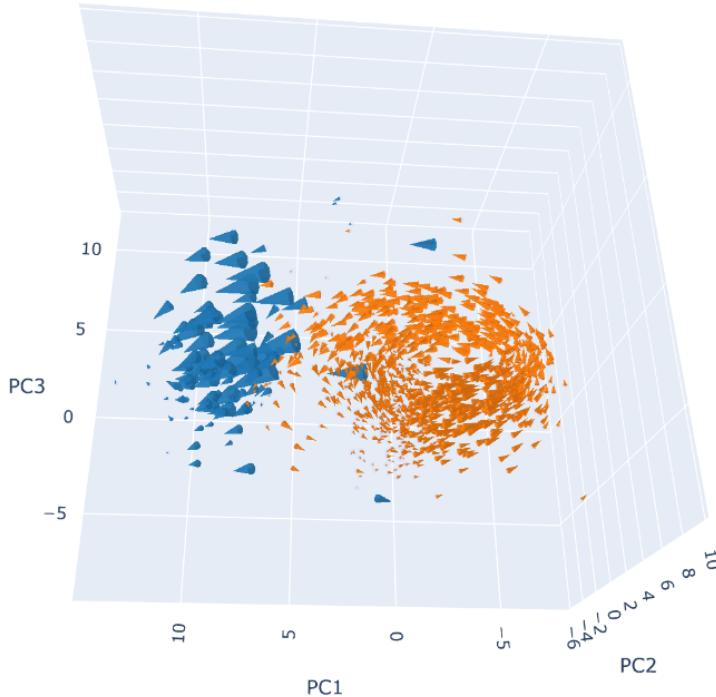
1227 To demonstrate MODE’s ability to generalize to held-out data representing a large, contiguous portion
1228 of phase space, we ran an experiment in which a certain small arc on cell cycle was removed
1229 from the Goldbeter data. The sector was chosen to be away from the switching zone and near where
1230 protease starts being nonzero (red dots, Fig. 15, left). We found no difference (Fig. 15, right) in
1231 MODE’s ability to detect the ground truth dynamical regimes on training (lightweight dots) vs OOD
1232 data (heavy blue dots). This indicates the bias introduced by MODE’s sparse regressors helps ”paint
1233 in” contiguous regions of phase space.

1234 B.6 CELL CYCLING IN HUMAN FIBROBLASTS
1235

1236 In Sec. 3 we used MODE in order to model experimental scRNA-seq data from the U2OS cell-line.

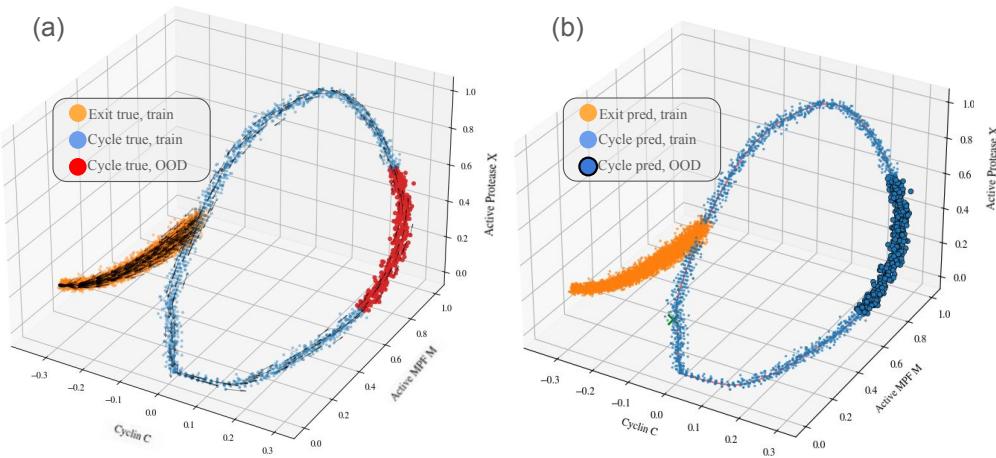
1237 To further demonstrate the applicability of MODE on experimental data, we turned to data from Riba
1238 et al. of human fibroblasts. In this data, the cells can be separated into two subpopulations, only one
1239 of which exhibits genes related to the cell-cycle. A UMAP of this data can be seen in Fig. 16 (a),
1240 where cells are either part of the cell cycle or exiting it. We follow the same preprocessing steps
1241 as for the FUCCI data (Appendix Sec. D.1, estimate velocities using scVelo (Bergen et al., 2020),

1242
 1243
 1244
 1245
 1246
 1247
 1248
 1249
 1250
 1251
 1252
 1253
 1254
 1255
 1256
 1257
 1258
 1259
 1260
 1261
 1262
 1263
 1264
 1265



1266 Figure 14: 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.
 1267

1268
 1269
 1270
 1271
 1272
 1273
 1274
 1275
 1276
 1277
 1278
 1279
 1280
 1281
 1282
 1283
 1284
 1285
 1286
 1287
 1288



1289 Figure 15: *MODE generalizes on Out-of-Distribution data.* (a) a 2-expert MODE was fit to the
 1290 Goldbeter data with a sector of points hidden during training. (b) MODE is able to detect the correct
 1291 dynamical regimes even in the region of space that was held-out during training.

1292
 1293
 1294
 1295

which are then projected onto a 5-dimensional PCA. After the filtration steps, this data contains a total of 5,367 cells. Riba et al. showed it is possible to cluster these into cycling and non-cycling cells.

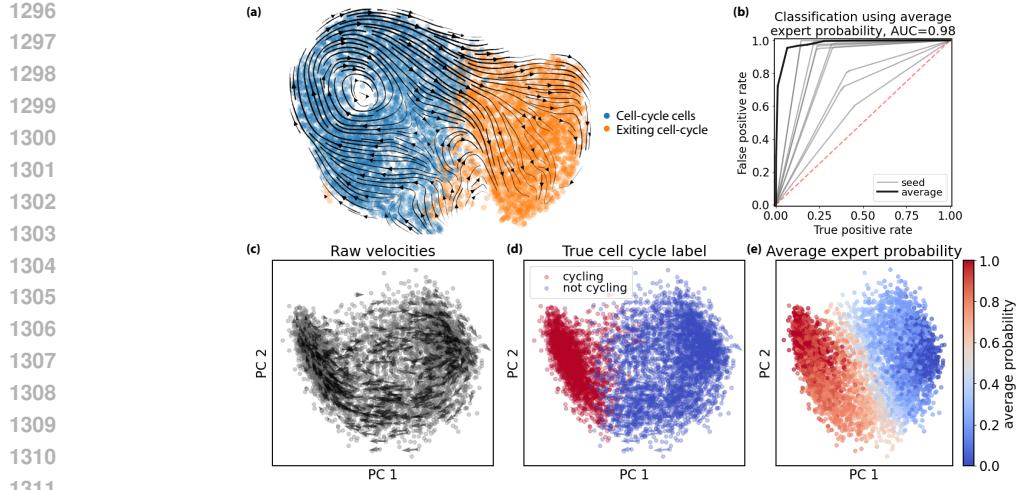


Figure 16: *Using MODE to model cell-cycle and exit dynamics in human fibroblasts.* (a) UMAP of the human fibroblasts data from Riba et al.. The cells can be divided into those that are part of the cell-cycle (in blue) and those that are not part of it (in orange). (b) MODE accurately distinguishes between these two regimes, and is stable between different seeds. (c) PCA projection of the data into two dimensions, with cell velocities. (d) true labeling into cell-cycle (blue) versus not (red). (e) classification according to 10 2-expert MODE, averaged over different seeds. The two different regimens, cycling and not, are captured by the two expert fit using MODE.

Following the results on U2OS, we fit this data with a 2-expert MODE, using up to quadratic terms, and repeated this process over ten random initializations to create an ensemble model (as described in Appendix Sec. D.1). We found that the ensemble averaged expert distribution is able to faithfully distinguish between cells that are cycling versus those that are not (Fig. 16 (b)). As in the U2OS data, the gating function had high entropy for cells between these regions, as shown in Fig. 16 (e).

1350 **C ALGORITHM IMPLEMENTATION DETAILS**

1351
 1352 We implement two variants of MODE corresponding to different assumptions about the mixing dis-
 1353 tribution π : MODE Local (state-independent mixing) and MODE Global (state-dependent mixing
 1354 via neural gating).

1355
 1356 **C.1 SPARSE IDENTIFICATION OF NONLINEAR DYNAMICS (SINDY)**

1357 The SINDy framework (Brunton et al., 2016) forms the foundation of our expert dynamics pa-
 1358 rameterization. Given snapshot data $D = \{(x_i, \dot{x}_i)\}_{i=1}^N$, SINDy assumes the governing equation can be
 1360 expressed as a sparse linear combination of basis functions:

1361
$$\dot{x} = f(x) = Z(x)\Theta,$$

1362 where $Z(x) \in \mathbb{R}^{n_{\text{lib}}}$ is a library of candidate functions and $\Theta \in \mathbb{R}^{n_{\text{lib}} \times d}$ are sparse coefficients
 1363 discovered via regularized regression. For polynomial libraries of degree c , we construct:

1364
$$Z(x) = [1, x_1, \dots, x_d, x_1^2, x_1 x_2, \dots, x_d^c],$$

1365 yielding $n_{\text{lib}} = \binom{d+c}{c}$ terms. The sparsity-promoting Laplace prior $\Theta_s \sim \text{Lap}(0, 1/\lambda)$ in Eq. 3
 1366 induces an ℓ_1 penalty $\lambda \|\Theta\|_1$, implemented via Lasso regression. This enables automatic selection
 1367 of relevant terms, yielding interpretable ordinary differential equations: each row of Θ specifies
 1368 which polynomial terms contribute to each state derivative. For example, with $d = 2$ and $c = 2$, a
 1369 discovered equation might be $\dot{x}_1 = \theta_2 x_2 + \theta_5 x_1 x_2$, directly revealing interaction structure.

1370 In MODE Local (Sec. C), we solve weighted SINDy regressions in the M-step with responsibilities
 1371 $R_{i,k}$ as sample weights. For numerical stability, we scale the design matrix and targets by $\sqrt{R_{i,k}}$,
 1372 apply column-wise standardization without centering, fit Lasso, then rescale coefficients. In MODE
 1373 Global (Sec. C.3), expert weight matrices Θ_k are directly optimized via gradient descent with ℓ_1
 1374 regularization. The degree c controls model complexity: higher degrees capture richer dynamics but
 1375 increase overfitting risk.

1376 We show in the clustering experiments that SINDy can accurately recover the true governing equa-
 1377 tions when they lie within the chosen basis (Tables 4,5). However, SINDy still works well, even
 1378 when the true dynamics are not expressible in closed form in the basis. This is the case with the
 1379 Goldbeter model, which uses rational functions. Polynomials should be an effective basis for this
 1380 data away from singularities, especially when the degree of the basis functions grows larger. We can
 1381 confirm this in a simple experiment in which the polynomial degree, d , is varied and the resulting
 1382 model is evaluated using W_1 and W_2 pushforward losses (Fig. 17).

1383
 1384 **C.2 MODE LOCAL**

1385 MODE Local implements the case where mixing probabilities $\pi = (\pi_1, \dots, \pi_K)$ are state-
 1386 independent constants. This variant is particularly suitable for unsupervised clustering of hetero-
 1387 geneous dynamical populations where each subpopulation follows distinct intrinsic dynamics re-
 1388 gardless of spatial location.

1389
 1390 **Algorithm.** We optimize the MAP objective (Eq. 3) using an Expectation-Maximization (EM)
 1391 algorithm with soft assignments. The algorithm alternates between:

1392 **E-step:** Compute responsibilities $R_{i,k} = p(s_i = k \mid x_i, \dot{x}_i)$ for each data point i and expert k :

1393
$$R_{i,k} \propto \pi_k \cdot \mathcal{N}(\dot{x}_i \mid Z(x_i)\Theta_k, \sigma_k^2 I_d) \tag{17}$$

1394 where responsibilities are normalized such that $\sum_{k=1}^K R_{i,k} = 1$.

1395
 1396 **M-step:** Update expert parameters by solving weighted SINDy regression problems. For each
 1397 expert k , we fit polynomial coefficients Θ_k by minimizing:

1398
$$\sum_{i=1}^N R_{i,k} \|\dot{x}_i - Z(x_i)\Theta_k\|_2^2 + \lambda \|\Theta_k\|_1 \tag{18}$$

1399
 1400 We also update mixing probabilities $\pi_k = \frac{1}{N} \sum_{i=1}^N R_{i,k}$ and noise variances σ_k^2 .

1404

1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

Figure 17: Degree, d , of the SINDy basis was varied while fitting Goldbeter data. We observed monotonically decreasing Wasserstein error as d increased, indicating an increasingly accurate Taylor fit of the rational functions in the true dynamics.

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.

Training Procedure. The EM algorithm for MODE Local is summarized below:

The E-step computes soft assignments (responsibilities) $R_{i,k}$ for each data point based on the current expert parameters (lines 6–10). The M-step updates expert dynamics via weighted SINDy regression (lines 13–18), then updates mixing probabilities and noise variances (lines 21–22). Convergence is monitored via the marginal log-likelihood (line 25). The weighted regression uses sample re-weighting with column standardization (lines 14–17) to improve numerical stability.

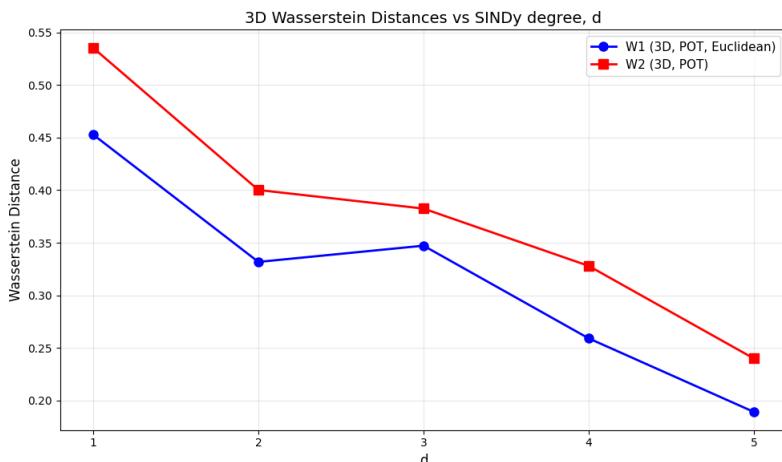
C.3 MODE GLOBAL

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:

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



Algorithm 1 MODE Local Training (EM Algorithm)

Require: Data (X, \dot{X}) , number of experts K , hyperparameters $\{\lambda, \text{degree}\}$

1: **Initialize** Expert parameters $\{\Theta_k\}_{k=1}^K$, mixing probabilities $\pi \in \mathbb{R}^K$, noise scales $\{\sigma_k\}_{k=1}^K$

2: Construct polynomial basis: $Z(X) \leftarrow \text{PolynomialFeatures}(X, \text{degree})$ $\triangleright (N, P)$ features

3:

4: **for** iteration = 1 to N_{iter} **do**

5:

6: *// E-step: Compute responsibilities*

7: **for** $k = 1$ to K **do**

8: $v_k(X) \leftarrow Z(X)\Theta_k$ \triangleright Expert predictions

9: $\ell_{i,k} \leftarrow \log \pi_k + \log \mathcal{N}(\dot{x}_i | v_k(x_i), \sigma_k^2)$ \triangleright Log-likelihood per sample

10: **end for**

11: $R_{i,k} \leftarrow \frac{\exp(\ell_{i,k})}{\sum_{j=1}^K \exp(\ell_{i,j})}$ \triangleright Normalize responsibilities

12:

13: *// M-step: Update parameters*

14: **for** $k = 1$ to K **do**

15: *// Weighted SINDy regression*

16: $w_i \leftarrow \sqrt{R_{i,k}}$ \triangleright Sample weights

17: $Z_w \leftarrow Z \odot w$, $\dot{X}_w \leftarrow \dot{X} \odot w$ \triangleright Weight design matrix and targets

18: Standardize columns of Z_w : $Z_w \leftarrow Z_w / \text{std}(Z_w)$

19: Solve: $\Theta_k \leftarrow \arg \min_{\Theta} \|Z_w \Theta - \dot{X}_w\|_2^2 + \lambda \|\Theta\|_1$ \triangleright Lasso regression

20: Rescale Θ_k to original feature scale

21:

22: *// Update mixing probability and noise*

23: $\pi_k \leftarrow \frac{1}{N} \sum_{i=1}^N R_{i,k}$

24: $\sigma_k^2 \leftarrow \frac{\sum_{i=1}^N R_{i,k} \|\dot{x}_i - v_k(x_i)\|_2^2}{\sum_{i=1}^N R_{i,k}}$

25: **end for**

26:

27: *// Check convergence*

28: $\mathcal{L} \leftarrow \sum_{i=1}^N \log \sum_{k=1}^K \pi_k \mathcal{N}(\dot{x}_i | v_k(x_i), \sigma_k^2)$ \triangleright Marginal log-likelihood

29: **if** $|\mathcal{L} - \mathcal{L}_{\text{prev}}| < \text{tol}$ **then**

30: **break**

31: **end if**

32: $\mathcal{L}_{\text{prev}} \leftarrow \mathcal{L}$

33: **end for**

34:

Ensure: Expert parameters $\{\Theta_k\}_{k=1}^K$, mixing probabilities π , noise scales $\{\sigma_k\}_{k=1}^K$

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:

- **Expert L1 sparsity:** $\lambda_{\text{expert}} \sum_{k=1}^K \|\Theta_k\|_1$ promotes sparse expert dynamics (applied only to linear SINDy experts)
- **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

In practice, we find that a balance of entropy and load balancing terms works fairly well and is not too difficult to tune. Often, setting both to 1 produces fine results which can

1512
 1513
 1514
 1515
 1516
 1517
 1518
 1519
 1520
 1521
 1522
 1523
 1524
 1525
 1526
 1527
 1528
 1529
 1530
 1531
 1532
 1533
 1534
 1535
 1536
 1537
 1538
 1539
 1540
 1541
 1542
 1543
 1544
 1545
 1546
 1547
 1548
 1549
 1550
 1551
 1552
 1553
 1554
 1555
 1556
 1557
 1558
 1559
 1560
 1561
 1562
 1563
 1564
 1565
 be improved if necessary by either simple manual tuning or a grid search. The example of a gridsearch on gating entropy and load balancing in which Wasserstein 1-distance is measured on Goldbeter testing data is given in Fig. 18.

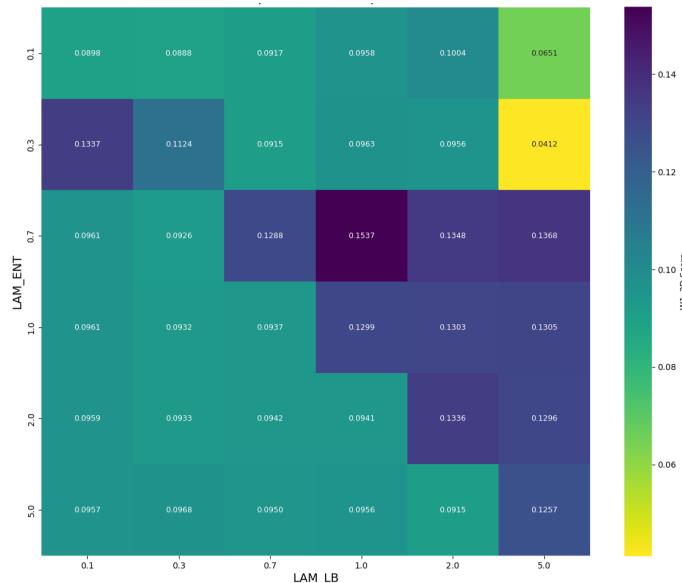


Figure 18: Hyperparameter grid search on the Goldbeter data.

Rollouts All results in the paper use the stochastic rollout described in Eqs. 5. Alternatively, we investigated using "mean" and "max" rollouts in which the average or mode of the expert distribution is used at every step, respectively. This represents a different type of cellular dynamics in which regulatory processes either have a conjoined, simultaneous effect or a winner-takes-all dynamics. Table 7 shows Wasserstein scores on the Goldbeter data for these different dynamics.

	W_1	W_2	$W_{1,C}$	$W_{1,M}$	$W_{1,X}$
Max	0.096034	0.152445	0.022993	0.039216	0.050021
Mean	0.101588	0.160200	0.024213	0.041055	0.055026
Stochastic	0.0837	0.1049	0.0590	0.0223	0.0289

Table 7: Wasserstein distances for different MODE rollout methods. Stochastic rollouts perform best on this Goldbeter data.

Implementation. The algorithm is implemented in the `MODEGlobal` class with two training modes: fixed-K and adaptive-K. For all clustering experiments (Sec. 3.1), the key hyperparameters were set to:

- K: Number of experts (fixed mode: default 2)
- degree: Polynomial degree for expert dynamics (default: 2)
- hidden: Hidden units in gating network (default: 64)
- epochs: Training epochs (default: 1500)
- lr: Learning rate (default: 1×10^{-3})
- $l1_{\text{lambda}}$, ent_{lambda} , lb_{lambda} : Regularization weights (defaults: 10^{-1} , 1, 1)

MODE training and hyperparameters for all experiments are shown in tables Tab. 8, 9 and 11. Wherever parameters are dashed, they were omitted or set to 0 as they had no influence on the result.

1566 All MLP baselines had four hidden layers (except for the two data which had only one layer) with
 1567 64 units in each layer and \tanh activation functions.
 1568

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

1583
1584 Table 8: Toy Branching (i.e. Fig. 1) hyperparameters.
1585

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

1601 Table 9: Goldbeter oscillator hyperparameters.
1602

1603 All implementations support multiple random restarts with validation-based model selection to avoid
 1604 poor local minima. For reproducibility, all experiments use fixed random seeds across method com-
 1605 parisons. GMM and spectral clustering baselines were run in `scikit-learn` v. 1.7.2.
 1606 MODE, MLPs and SINDy were implemented in `pytorch` v. 2.8.0 and were trained with
 1607 an Adam optimizer (Kingma & Ba, 2015). Experiments were run locally on a cpu cluster with no
 1608 parallelization. Typical run times range from around a minute for the toy example of Fig. 1 to about
 1609 ten minutes for the full fifteen-seed FUCCI ensemble shown in Fig. 5 (e), (f).

1610 **Training Procedure.** The complete training algorithm is summarized below:
 1611

1612 The forward pass computes gating probabilities $\pi(x)$ and expert predictions $\{v_k(x)\}$ (lines 7–8),
 1613 which are combined via a Gaussian mixture likelihood (line 11). Regularization terms (lines 14–
 1614 17) prevent mode collapse and encourage sparse, confident expert assignment. The total loss is
 1615 minimized using Adam optimizer (lines 20–21).

1616 C.4 ADAPTIVE EXPERT SELECTION

1617 In many biological systems, the number of distinct dynamical regimes is *a priori* unknown. While
 1618 cell population heterogeneity is ubiquitous in single-cell data, determining how many functionally

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 10: 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 11: FUCCI hyperparameters.

distinct subpopulations exist - and whether these correspond to discrete dynamical regimes or continuous transitions - remains a fundamental challenge. Manual selection of the expert count K is both impractical and prone to overfitting or underfitting. We therefore introduce an AIC/BIC-based strategy for dynamically adapting K during training, allowing the model to automatically discover the intrinsic dimensionality of the system.

Starting with an initial K (default: 5), the algorithm periodically evaluates information criteria to balance model complexity with goodness-of-fit:

$$AIC = 2k + 2N \cdot \text{NLL} \quad (20)$$

$$BIC = k \log(N) + 2N \cdot \text{NLL} \quad (21)$$

where k is the number of model parameters, N is the number of samples, and NLL is the negative log-likelihood. The adaptation rules are:

- **Add expert:** If AIC/BIC plateaus for P_{add} consecutive checks (default: 30 checks \times 50 epochs = 1500 epochs), suggesting the model needs more capacity to improve fit.
- **Remove expert:** If (1) any expert has usage $< \tau_{\text{usage}}$ (default: 5%), indicating redundancy, or (2) AIC/BIC worsens consistently for P_{remove} checks (default: 30), suggesting overfitting. The least-used expert is removed.

To further encourage sparsity in expert assignment, we add the regularization term $\lambda_{\text{gate}} \sum_{p \in \text{gating}} |p|$ to the loss, which penalizes gating network weights and prevents diffuse, non-interpretable expert usage (default: $\lambda_{\text{gate}} = 10^{-3}$).

1674 **Algorithm 2** MODE Global Training

1675

1676 **Require:** Data (X, \dot{X}) , number of experts K , hyperparameters $\{\lambda_{\text{expert}}, \lambda_{\text{ent}}, \lambda_{\text{lb}}\}$

1677 1: **Initialize** Gating network MLP, expert parameters $\{\Theta_k\}_{k=1}^K$, noise scales $\{\sigma_k\}_{k=1}^K$

1678 2: Normalize data: $\mu_x \leftarrow \text{mean}(X)$, $\sigma_x \leftarrow \text{std}(X)$

1679 3: **for** epoch = 1 to N_{epochs} **do**

1680 4: **for** minibatch (x_b, \dot{x}_b) in DataLoader(X, \dot{X}) **do**

1681 5: // Forward pass

1682 6: $\pi(x_b) \leftarrow \text{softmax}(\text{MLP}((x_b - \mu_x)/\sigma_x))$ ▷ Gating probabilities

1683 7: $\{v_k(x_b)\}_{k=1}^K \leftarrow \{Z(x_b)\Theta_k\}_{k=1}^K$ ▷ Expert predictions

1684 8: // Mixture likelihood

1685 9: $\mathcal{L}_{\text{data}} \leftarrow -\frac{1}{B} \sum_{i=1}^B \log \sum_{k=1}^K \pi_k(x_i) \cdot \mathcal{N}(\dot{x}_i | v_k(x_i), \sigma_k^2)$

1686 10: // Regularization

1687 11: $\mathcal{L}_{\text{expert}} \leftarrow \lambda_{\text{expert}} \sum_{k=1}^K \|\Theta_k\|_1$ ▷ Expert sparsity

1688 12: $\mathcal{L}_{\text{ent}} \leftarrow -\lambda_{\text{ent}} \cdot \frac{1}{B} \sum_{i=1}^B \sum_{k=1}^K \pi_{i,k} \log \pi_{i,k}$ ▷ Gate entropy

1689 13: $\bar{\pi} \leftarrow \frac{1}{B} \sum_{i=1}^B \pi_i$ ▷ Average expert usage

1690 14: $\mathcal{L}_{\text{lb}} \leftarrow \lambda_{\text{lb}} \sum_{k=1}^K \bar{\pi}_k \log(K \bar{\pi}_k)$ ▷ Load balancing

1691 15: // Update

1692 16: $\mathcal{L}_{\text{total}} \leftarrow \mathcal{L}_{\text{data}} + \mathcal{L}_{\text{expert}} + \mathcal{L}_{\text{ent}} + \mathcal{L}_{\text{lb}}$

1693 17: Backpropagate $\nabla \mathcal{L}_{\text{total}}$ and update parameters with Adam

1694 18: **end for**

1695 19: **if** early stopping criterion met **then**

1696 20: **break**

1697 21: **end if**

1698 22: **end for**

1699 **Ensure:** Trained gating network and expert parameters

Figure 5(e), (f) demonstrates the robustness of this adaptive strategy across 15 independent training runs with different random initializations. Despite starting from $K = 5$ experts, all runs converge to $K = 2$ experts, with the mean trajectory (purple line) showing rapid initial adaptation followed by stable convergence. The right panel shows that AIC values decrease sharply during adaptation and stabilize once the optimal expert count is reached, with all runs identifying similar final model complexities (red dots mark minimum AIC). The algorithm maintains $K_{\min} \leq K \leq K_{\max}$ (default: $2 \leq K \leq 10$) and resets patience counters after each adaptation to allow the model to stabilize. This consistency across random seeds confirms that the method reliably discovers the intrinsic dimensionality of the dynamical regimes without requiring prior knowledge of the optimal K .

Note on experimental results. While this adaptive selection capability is critical for exploratory analyses where the number of regimes is unknown, *the main results presented in this paper were obtained with fixed K values*, as the number of biologically relevant populations was known *a priori* from experimental annotations. The adaptive strategy is provided as a tool for discovery-driven applications and to demonstrate the model’s robustness to initialization, but was not required for the validation experiments described in Section 3.

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

1728 fluorescently tagged cell cycle markers: CDT1 (tagged with RFP, expressed during G1 phase) and
 1729 GMNN (tagged with GFP, expressed during S and G2 phases). Cells expressing both markers si-
 1730 multaneously indicate the G1-S transition phase. This dual-marker system provides precise tempo-
 1731 ral information about cell cycle progression, enabling the identification of cells in cycling versus
 1732 non-cycling (differentiation) states. The FUCCI markers serve as our ground truth labels for distin-
 1733 guishing between two fundamental dynamical regimes: (1) active cell cycle progression and (2) cell
 1734 cycle exit leading to differentiation.

1735
 1736 **Preprocessing pipeline.** We followed the preprocessing protocol established by Zheng et al.
 1737 (2023), which includes several critical steps for robust scRNA-seq analysis:

- 1738 • **Gene filtering:** Removal of genes with low expression counts across the cell population to
 1739 reduce noise and computational burden.
- 1740 • **Cell filtering:** Exclusion of cells with abnormally low or high gene counts that might
 1741 indicate technical artifacts or doublets.
- 1742 • **Normalization:** Application of library size normalization to account for differences in
 1743 sequencing depth between cells.
- 1744 • **Log transformation:** $\text{Log}_{10}(\text{count} + 1)$ transformation to stabilize variance and reduce the
 1745 influence of highly expressed genes.
- 1746 • **Highly variable gene selection:** Identification and retention of genes showing significant
 1747 cell-to-cell variability, which are most informative for capturing biological differences.

1749
 1750 **RNA velocity estimation.** We employed scVelo (Bergen et al., 2020) to estimate RNA velocity
 1751 vectors from the preprocessed data. scVelo computes velocity by modeling the dynamics of gene
 1752 expression through an analytical framework that considers transcription, splicing, and degradation
 1753 rates. Specifically, the method uses the ratio of unspliced to spliced mRNA counts to infer the
 1754 directional flow of gene expression changes. This provides local velocity estimates \dot{x}_i for each cell
 1755 i , representing the instantaneous rate of change in gene expression space. The velocity estimation
 1756 captures the intrinsic directionality of cellular state transitions, making it particularly suitable for
 1757 identifying developmental trajectories and regime transitions.

1758
 1759 **Dimensionality reduction and feature space.** Given the high-dimensional nature of the gene
 1760 expression data (58,884 dimensions), we applied Principal Component Analysis (PCA) to project
 1761 both the gene expression profiles and their corresponding velocity vectors into a 5-dimensional
 1762 latent space. This dimensionality reduction serves multiple purposes: (1) computational efficiency
 1763 for MODE training, (2) noise reduction by focusing on the most informative directions of variation,
 1764 and (3) visualization feasibility for the first three principal components. The choice of 5 dimensions
 1765 balances the preservation of essential biological signals while maintaining computational tractability
 1766 for the mixture of experts framework.

1766
 1767 **MODE training configuration.** Full hyperparameters in Tab. 11. An ensemble model created by
 1768 averaging (e.g. over mixture distributions and expert parameters; experts were aligned by similarity
 1769 of parameter vectors) over 10 random seeds. All seeds but one (the outlier curve in Fig. 5) divided
 1770 the data into clear cycle and exit regimes.

1771
 1772 **Validation methodology.** Model performance was evaluated using the FUCCI ground truth labels
 1773 as the gold standard for cell cycle phase classification. We computed Receiver Operating Charac-
 1774 teristic (ROC) curves for each of the 10 model initializations, measuring the ability of the learned
 1775 expert probabilities to distinguish cycling from differentiating cells. The consistently high Area Under
 1776 the Curve (AUC) scores (0.98 ± 0.01) across all initializations demonstrate both the reliability
 1777 of our approach and the biological relevance of the discovered dynamical regimes. Additionally,
 1778 we validated the biological interpretability of the learned dynamics by examining the correspon-
 1779 dence between expert assignments and known cell cycle markers, confirming that the two experts
 successfully captured the cycling and differentiation programs respectively.

1780
 1781 **Transcriptional and regulatory validation.** Beyond phase classification accuracy, we validated
 whether MODE’s learned dynamics capture biologically interpretable structure through multilevel

1782 analysis of gene expression signatures, temporal dynamics, and regulatory network topology. Im-
1783 portantly, MODE was trained without any prior knowledge of cell cycle markers or temporal anno-
1784 tations—all biological interpretations were performed post-hoc.
1785

1786 EXPERT-SPECIFIC DIFFERENTIAL EXPRESSION METHODOLOGY. We compared gene expression
1787 between cells assigned to Expert 0 versus Expert 1 using dominant expert assignments as the basis
1788 for grouping. For each gene g , we computed mean expression in each expert group and assessed
1789 differential expression using fold-change $FC_g = \log_2(\bar{x}_g^{(1)}/\bar{x}_g^{(0)})$ with statistical significance via
1790 permutation tests (FDR-corrected $p < 0.001$).

1791 Remarkably, without supervision, Expert 0 spontaneously enriched for canonical proliferation mark-
1792 ers (TOP2A, CDK1, AURKA) and chromatin organization factors (HIST1H1C, HIST1H4C), while
1793 Expert 1 enriched for known cell cycle exit-associated genes (CAV1, THBS1, CCN1, CDKN1A).
1794 Comparison with established marker gene sets from the literature (Malumbres & Barbacid, 2009;
1795 Coller et al., 2006) confirmed that MODE’s unsupervised expert discovery recovered biologically
1796 meaningful cell states (Table 12).
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835

1836
 1837
 1838
 1839
 1840
 1841
 1842
 1843
 1844
 1845
 1846
 1847
 1848
 1849
 1850
 1851
 1852
 1853
 1854
 1855
 1856
 1857
 1858

1859
 1860
 1861
 1862
 1863
 1864
 1865

Upregulated in Expert 1			Downregulated in Expert 1		
Gene	Role	Change	Gene	Role	Change
CAV1	Quiescence	28.0	KRT17	Proliferation	-41.0
CDC20	APC/C activator	26.2	TGFBI	ECM protein	-11.1
THBS1	Anti-angiogenic	23.7	HIST1H1C	Chromatin	-10.2
TPX2	Mitotic spindle	21.2	TOP2A	DNA topology	-7.8
CCNB1	G2/M cyclin	20.5	CDK1	Cell cycle kinase	-5.7

Table 12: Top differentially expressed genes between experts

1866
 1867
 1868
 1869
 1870
 1871
 1872
 1873
 1874
 1875
 1876
 1877
 1878
 1879
 1880
 1881
 1882
 1883
 1884
 1885
 1886
 1887
 1888
 1889

PSEUDOTEMPORAL ORDERING AND TEMPORAL DYNAMICS. We constructed a pseudotime trajectory by sorting cells according to Expert 0 probability $\pi_0(x_i)$, naturally ordering cells from high π_0 to low π_0 states. Gene expression profiles were smoothed using Gaussian kernel smoothing ($\sigma = 50$ cells), and the transition point was identified as the pseudotime value maximizing expert entropy: $t^* = \arg \max_t H(\pi(x_t))$, where $H(\pi) = -\sum_k \pi_k \log \pi_k$.

Along this trajectory, Expert 0-enriched markers (TOP2A, CDK1) decreased while Expert 1-enriched markers (CAV1, THBS1) increased, with the steepest changes occurring near t^* (Fig. 6A). Expression profiles showed gradual downregulation of proliferation-associated genes preceding up-regulation of exit-associated genes, suggesting sequential transcriptional reprogramming consistent with established models of cell cycle exit (Spencer et al., 2013).

GENE REGULATORY NETWORK EXTRACTION AND REWIRING ANALYSIS. For linear SINDy experts with dynamics $\dot{x} = Z(x)\Theta_k$, the local Jacobian in PCA space is computed from the learned coefficient matrices Θ_k and the basis family Z . To map Jacobians to the original gene space, we applied the chain rule through PCA projection $x_{\text{PCA}} = V(x_{\text{gene}} - \mu)$, where $V \in \mathbb{R}^{d_{\text{PCA}} \times d_{\text{gene}}}$ are the PCA components:

$$J_k^{\text{gene}} = V^\top J_k^{\text{PCA}} V$$

We extracted inferred gene regulatory networks (GRNs) for three cell state regions based on expert dominance: before transition (Expert 0 probability > 0.8), during transition (expert entropy in top 20%), and after transition (Expert 1 probability > 0.8). For each region, we computed mean Jacobians over sampled cells ($n = 50$ per region) and applied a sparsity threshold of $|\tilde{J}_{ij}| > 0.15$ after normalization to retain the strongest interactions.

Network topology analysis revealed substantial rewiring across transition phases (Fig. 6B). Before transition, the inferred network featured prominent interactions among proliferative genes including the CDK1–CCNB1–CDC20 module. After transition, exit-associated genes (CAV1, THBS1) gained regulatory prominence while proliferative gene interactions weakened. These patterns of network reorganization are consistent with known regulatory circuitry governing cell cycle progression and exit (Malumbres & Barbacid, 2009; Pines, 1995).

HIGH-ENTROPY TRANSITION ANALYSIS. To identify genes specifically associated with the transition state, we compared gene expression between cells in high-entropy regions (top 20% by gating network entropy: $H(\pi) > \theta_H$, where $\theta_H = \text{percentile}_{80}(H)$) versus low-entropy cells (bottom 20%). This identified 222 genes with significant expression changes ($|FC| > 0.5$, permutation $p < 0.001$, Fig. 6C).

Notably, top transition-associated genes (including CDC20B, PDK4, and RASSF6) differed from the markers defining stable expert states (e.g., TOP2A for Expert 0, CAV1 for Expert 1), suggesting these genes represent transient regulatory factors active during the decision phase rather than stable markers of committed states. This demonstrates MODE’s ability to distinguish stable dynamical regimes from transitional cell states.

VALIDATION SUMMARY. Collectively, these analyses demonstrate that MODE’s unsupervised learning recovers biologically coherent structure: (1) experts spontaneously align with known cell cycle states without prior annotation, (2) pseudotemporal ordering reveals sequential transcriptional dynamics matching established exit mechanisms (Spencer et al., 2013), (3) inferred regulatory networks exhibit topology changes consistent with known cell cycle circuitry (Malumbres & Barbacid, 2009; Pines, 1995), and (4) transition-state cells show distinct gene expression profiles from stable states. Together, this multilevel concordance with domain knowledge validates that MODE captures mechanistically meaningful dynamics rather than arbitrary data partitions.

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.

1944 REPRODUCIBILITY STATEMENT

1945
1946 All details necessary to reproduce the results of this work can be found in the appendix and supple-
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997

- 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.