npj | systems biology and applications

Article

Published in partnership with the Systems Biology Institute

https://doi.org/10.1038/s41540-024-00424-7

Recovering biomolecular network dynamics from single-cell omics data requires three time points

Shu Wang [1,2,3], Muhammad Ali Al-Radhawi[4], Douglas A. Lauffenburger [3] ✉ & Eduardo D. Sontag [4] ✉

Single-cell omics technologies can measure millions of cells for up to thousands of biomolecular features, enabling data-driven studies of complex biological networks. However, these high-throughput experimental techniques often cannot track individual cells over time, thus complicating the understanding of dynamics such as time trajectories of cell states. These "dynamical phenotypes" are key to understanding biological phenomena such as differentiation fates. We show by mathematical analysis that, in spite of high dimensionality and lack of individual cell traces, three time-points of single-cell omics data are theoretically necessary and sufficient to uniquely determine the network interaction matrix and associated dynamics. Moreover, we show through numerical simulations that an interaction matrix can be accurately determined with three or more time-points even in the presence of sampling and measurement noise typical of single-cell omics. Our results can guide the design of single-cell omics time-course experiments, and provide a tool for data-driven phase-space analysis.

In recent decades, experimental single-cell profiling techniques, such as single-cell RNA sequencing[1], multiplexed immunofluorescence[2], or mass[3] and multiparametric flow cytometry[4], have enabled the simultaneous measurement of biomolecular abundances for many ($n = 10^1$–10^5) biomolecules—such as proteins or RNAs, in large numbers ($N = 10^3$–10^6) of cells. These biomolecules constitute highly complex biological networks and their changes over time may in principle be modeled by dynamical systems. The time trajectories of these dynamical systems represent cell states and can be interpreted as "dynamical phenotypes" that are fundamental to understanding biological phenomena at various time-scales, from cell cycles to circadian rhythms to differentiation.

Past works have modeled biomolecular subnetworks by fitting dynamical systems model parameters to other forms of data (e.g., low-dimensional live-cell imaging or bulk population-averaged expression data), enabling predictions about how cell cycle timing can be biochemically manipulated[5], how different cancer signaling pathways respond to drugs[6], or how gene regulatory networks affect cell-type differentiation[7,8]. However, this standard approach of fitting pre-specified models to data is limited by computational feasibility in higher dimensions, prior modeling assumptions, and data availability[9]. Single-cell omics have the potential to alleviate the problem of data availability, but with the caveat that the data often cannot provide single-cell time series due to the destructive sampling typically involved in omics approaches, which also precludes the direct application of powerful data-driven dynamical analysis tools like Dynamic Mode Decomposition[10]. Thus, various computational approaches tailored to single-cell omics data are being developed to aid dynamical analysis of the biomolecular networks inside cells[11], to understand how cell functions might be driven by molecular interactions.

Newly developed methods for inferring single-cell omics dynamics might be crudely categorized as follows: pseudotime[11], which fits data using one-dimensional shapes such as curves, trees, or graphs that are then interpreted as "axes" of time; RNA velocity[12], which estimates the instantaneous rates of change in measured mRNA transcripts based on models of the abundances of different post-transcriptional states of any given gene; potential landscapes[13], which model the single-cell data distribution as being generated by certain classes of dynamical systems such as gradient dynamics; or even extensions of the traditional model-fitting strategies[14]. Intriguingly, these methods often have the peculiar property that they infer time-varying dynamics using as few as one timepoint of single-cell omics data, i.e., without needing to explicitly use time information. Some notable exceptions make explicit use of time information, including Waddington-OT[15], which relies on an optimal transport assumption to trace out dynamics, PRESCIENT[16], which still uses a gradient assumption but incorporates time information when evaluating the goodness-of-fit of

[1]Donnelly Centre, University of Toronto, Toronto, ON, Canada. [2]Molecular Genetics, University of Toronto, Toronto, ON, Canada. [3]Department of Biological Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA. [4]Departments of Bioengineering and Electrical & Computer Engineering, Northeastern University, Boston, MA, USA. ✉e-mail: lauffen@mit.edu; e.sontag@northeastern.edu

candidate potential functions, and Tempora[17], which models dynamics at a coarser level of transitions between cell states while leveraging pathway enrichment information. The Dynamo framework[18] can use different types of time information, e.g. by modeling the metabolic labeling and expression kinetics of time-resolved single-cell RNA sequencing data to infer RNA velocities. Overall, it is not always clear how the various assumptions across various methods may bias the interpretation of single-cell dynamics, especially for methods that can accommodate data with only one timepoint.

Meanwhile, in the standard approach of inferring the equations of an n-dimensional dynamical system $\dot{x} = F(x)$ from finding best-fits to time-series data, one would expect at least $O(n)$ time points to be necessary[19] (depending on the assumed form of $F(x)$) to even have a unique best fit to data. However, single-cell omics experiments are able to survey variation (i.e., single-cell heterogeneity) at a given timepoint, which in general can also provide rich information for identifying a best-fitting model, e.g., different mechanistic models of biochemical processes can be indistinguishable when formulated as deterministic dynamical systems, but become identifiable upon modeling stochastic fluctuations[20]. Thus, it is currently unclear how many time points of single-cell omics data are needed, and how one might feasibly perform data-driven dynamical analysis, while avoiding strong assumptions.

This work provides an initial answer to these questions. As a start, we analyze the problem for linear dynamical systems $\dot{x} = Ax$, which can represent local approximations of nonlinear dynamical systems $\dot{x} = F(x)$ around steady states x_{ss}. The entries of the matrix A represent the network interactions. Through mathematical analysis, we prove that three time points are theoretically necessary and sufficient to uniquely determine A using the statistics of generic single-cell omics time-series data; we use "generic" in the technical sense that all the exceptional instances, in which a particular three-timepoint dataset is insufficient to uniquely determine A, form a set of *measure zero*, i.e., they have a "zero"-probability of occurring amongst all possible datasets. Our results suggest that many existing methods for dynamical analysis of single-cell omics data make, in some sense, stronger assumptions than even linearity, given that they require only one timepoint. Simultaneously, our results also show that an unbiased data-driven analysis of n-dimensional dynamics is quite feasible, as the required number of time points does not scale with the dimension n. Motivated by our mathematical analysis, we use a method of moments to infer dynamics from simulated data. We show that even in the presence of sampling noise and measurement noise that is typical of single-cell omics data, linear dynamics $\dot{x} = A(x - x_{ss})$ can still be estimated accurately using 3 or more time points, up to $n \sim 20$, upon which sample size N (i.e., number of cells) and measurement noise introduce substantial estimation error. Finally, we address the effect of non-linearities or unobserved variables of a dynamical system on the interpretation of dynamics A that can be estimated from data.

Results

Terminology and problem formulation

Consider a single-cell omics experiment. At a time point t_i, a collection of N cells are sampled from a homogeneous cell population. For each of the sampled cells, n features are measured, e.g, the levels of various proteins, RNAs, histone modifications, etc. Thus, the results are recorded in an $N \times n$ measurement matrix.

To model this process, we assume that each of the sampled cells is running an identical dynamical system in an n-dimensional state space, but with possibly different initial conditions at measurement time t_i, i.e., as might ideally be expected of an isogenic cell line cultured in a well. Let the state of each cell be $x_t \in \mathbb{R}^n$ at time $t \geq 0$. The time evolution of each cell is described by the deterministic ODE $\dot{x} = F(x)$ for some unknown function F and an initial condition $x_0 = x(0)$. Since the population is homogeneous, random initial conditions constitute the only source of variation among cells in our model. Therefore, we assume that x_0 is a random variable with a probability density function $f_0(x)$. At time t, the state evolves from x_0 to x_t and the probability density evolves to $f_t(x)$. Therefore, each row of the measurement matrix is a sample of the random variable x_{t_i}, where t_i is the measurement time. Using the measured data, we can estimate the mean and covariance of x_{t_i}. Hence, denote $\mu_t := \langle x_t \rangle = \mathbb{E}[x_t]$, and $\Sigma_t := \langle (x_t - \mu_t)(x(t) - \mu_t)^T \rangle := \mathbb{E}[(x_t - \mu_t)(x_t - \mu_t)^T]$, which we assume to be a positive-definite matrix.

Assume now that F is affine, with a steady state x_{ss} such that $F(x_{ss}) = 0$. Then $\dot{x} = F(x) = A(x - x_{ss})$ for some matrix $A \in \mathbb{R}^{n \times n}$. Therefore, $x_t = e^{tA}(x_0 - x_{ss}) + x_{ss}$. By linearity of expectation, we can describe the time evolution of the mean and the covariance as follows:

$$\mu_t = e^{tA}(\mu_0 - x_{ss}) + x_{ss} \qquad \Sigma_t = e^{tA}\Sigma_0(e^{tA})^T.$$

Note that the time evolution of the mean μ_t depends on the steady-state x_{ss}, while the corresponding equation for the covariance does not. This motivates the following definition:

Definition 1 *Let P, Q be symmetric positive definite matrices. Then (P, Q) is said to be an ordered pair of covariance matrices with respect to a matrix A and time interval $t > 0$ iff $Q = e^{tA}P\left(e^{tA}\right)^T$. Similarly, an ordered pair of vectors (ρ, θ) s.t. $\theta - x_{ss} = e^{tA}(\rho - x_{ss})$ is said to be an ordered pair of means.*

Hence, we are interested in finding the matrix A using the given pairs $(\Sigma_0, \Sigma_t) \equiv : (P, Q)$ and $(\mu_0, \mu_t) \equiv : (\rho, \theta)$ to solve for A. A single pair (P, Q), (ρ, θ) will be insufficient to solve for A, so we consider cases with multiple pairs (P_i, Q_i), (ρ_i, θ_i), in which the P_i's and ρ_i's are distinct over i, i.e., each pair corresponds to a different initial distribution of x_0 as characterized by their first and second moments μ, Σ. This includes the common scenario of measuring a single time series $f_0(x), f_{t_1}(x), f_{t_2}(x), f_{t_3}(x), \ldots$ by taking $P_i = \Sigma_{t_i}$, $Q_i = \Sigma_{t_{i+1}}$, and $\rho_i = \mu_{t_i}$, $\theta_i = \mu_{t_{i+1}}$, but also includes the more general scenario in which multiple time series are measured from different distributions $f_0(x)$, $g_0(x)$, … of initial conditions.

Main theorems

We aim at finding the smallest number of pairs of means and covariance matrices that are needed to find the matrix A of a linear system. We focus our attention on the generic case of a matrix A which is invertible and diagonalizable, possibly with complex eigenvalues. We will show that two generic pairs of means (ρ_i, θ_i) and of covariances (P_i, Q_i) are necessary and sufficient to uniquely determine A and x_{ss}. Also, we will show that three generic pairs of covariances (P_i, Q_i) can be sufficient to uniquely determine A. In either case, the geometric intuition might loosely be characterized as follows: a single pair of covariances can show how the distribution is stretched in various directions by the dynamics of A, but cannot resolve how the distribution is rotated and/or reflected by the dynamics; a second pair of covariances is able to resolve the rotation except for various choices of reflections; finally, the information contained in two pairs of means or a third pair of covariances is able to resolve the ambiguous reflections. In what follows, we denote the group of all orthogonal $n \times n$ matrices as $\mathbb{O}_n$, and the set of positive definite symmetric matrices as $\mathbb{S}_n^+$. Our main results are:

Theorem 1 *For each positive number $t > 0$, let $\mathcal{A}_{n,t}$ be the set of $n \times n$ real matrices A such that, for every eigenvalue $\lambda_j = a_j + ib_j$ of A, it holds that $|b_j| < \pi/t$.*

There is a generic (meaning open, dense, and with complement of measure zero) subset $S \subset \mathbb{S}_n^+ \times \mathbb{S}_n^+ \times \mathbb{R}^n \times \mathbb{R}^n$ with the following property. Suppose given a matrix-vector-tuple $(P_1, P_2, \rho_1, \rho_2) \in S$, two matrices $A, B \in \mathcal{A}_{n,t}$, and two vectors x_A, x_B, such that the following properties hold:

$$e^{tA}P_i(e^{tA})^T = e^{tB}P_i(e^{tB})^T \qquad \text{for } i = 1, 2$$

and

$$e^{tA}(\rho_i - x_A) + x_A = e^{tB}(\rho_i - x_B) + x_B \qquad \text{for } i = 1, 2.$$

Then $A = B$ and $x_A = x_B$. In other words, two matrix-vector pairs uniquely determine A.

Theorem 2 *For each positive number $t > 0$, let $\mathcal{B}_{n,t}$ be the set of $n \times n$ real matrices A such that, for every eigenvalue $\lambda_j = a_j + ib_j$ of A, it holds that $|b_j| < \pi/(2t)$.*

There is a generic (meaning open, dense, and with complement of measure zero) subset $T \subset \mathbb{S}_n^+ \times \mathbb{S}_n^+ \times \mathbb{S}_n^+$ with the following property. Suppose given a matrix-tuple $(P_1, P_2, P_3) \in T$, and two matrices $A, B \in \mathcal{B}_{n,t}$, the following properties hold:

$$e^{tA} P_i (e^{tA})^T = e^{tB} P_i (e^{tB})^T \qquad \text{for } i = 1, 2, 3.$$

Then $A = B$. In other words, three matrix pairs uniquely determine A.

We will show that the pairs of covariance matrices provide us with the majority of the information on the transition matrix e^{tA}, and that subsequently the pairs of means can be used to determine e^{tA} from amongst a finite set of possibilities. To determine A from e^{tA} using matrix logarithms, we will show that sufficiently small t relative to the dynamics of A enables us to determine A uniquely, in close resemblance of the Shannon–Nyquist Theorem, but that time-series data with varying time intervals can effectively remove the small-t constraint.

Remark 1 *It is important that the covariances are assumed to be positive definite as opposed to positive semi-definite, since any subspace along which there is zero variance provides less information about how the dynamics stretch the distribution of cell states. For example, in the extreme limit for which the covariances are simply the zero matrix, one would only have information on how the mean value is shifted by the dynamics, and $O(n)$ time-points would be needed to solve for e^{tA} as opposed to three time-points.*

In what follows, we will show that the first pair of covariances roughly determines e^{tA} up to rotations, that the second pair of covariances determines e^{tA} up to reflections, and that the two means can then determine e^{tA} uniquely.

First pair (P_1, Q_1) gives e^{tA} up to an orthogonal matrix R

We introduce the notations $V \equiv e^{tA}$, and $Q_i = e^{tA} P_i (e^{tA})^T$, giving:

$$Q_1 = V P_1 V^T. \tag{1}$$

As $P_1, Q_1 \in \mathbb{S}_n^+$, they each have a uniquely defined square-root that is also positive definite, $P_i^{1/2}$ and $Q_i^{1/2}$, respectively. Hence, we state the following result.

Theorem 3 *Given a nonsingular $V \in \mathbb{R}^{n \times n}$, and $P_1, Q_1 \in \mathbb{S}_n^+$, the following statements are equivalent:*

1. $Q_1 = V P_1 V^T$
2. *There exists some $R \in \mathbb{O}_n$ such that $V = Q_1^{1/2} R P_1^{-1/2}$.*

Proof To show property 2 implies property 1, we plug $V = Q_1^{1/2} R P_1^{-1/2}$ into the RHS of Eq. (1):

$$\begin{aligned} Q_1 &= (Q_1^{1/2} R P_1^{-1/2}) P_1 (P_1^{-1/2} R^T Q_1^{1/2}) \\ &= (Q_1^{1/2} R P_1^{-1/2}) (P_1^{1/2})^2 (P_1^{-1/2} R^T Q_1^{1/2}) \\ &= Q_1^{1/2} R R^T Q_1^{1/2} \\ &= Q_1^{1/2} Q_1^{1/2} = Q_1. \end{aligned}$$

To show property 1 implies property 2, we rearrange Eq. (1) after taking square roots:

$$\begin{aligned} Q_1^{1/2} Q_1^{1/2} &= V P_1^{1/2} P_1^{1/2} V^T \\ I &= Q_1^{-1/2} V P_1^{1/2} P_1^{1/2} V^T Q_1^{-1/2} \\ I &= (Q_1^{-1/2} V P_1^{1/2}) (Q_1^{-1/2} V P_1^{1/2})^T, \end{aligned}$$

i.e., $R \equiv Q_1^{-1/2} V P_1^{1/2} \in \mathbb{O}_n$, which rearranges to $V = Q_1^{1/2} R P_1^{-1/2}$. □

Thus, a single pair of covariance matrices only specifies $e^{tA} = Q_1^{1/2} R P_1^{-1/2}$ up to rotations and reflections represented by $R \in \mathbb{O}_n$.

Second pair (P_2, Q_2) gives e^{tA} up to a diagonal sign matrix Θ

Assume that we are given two pairs of covariance matrices (P_2, Q_2), (P_1, Q_1) where $P_1 \neq P_2$. Using Theorem 3, the first pair determines the transition matrix V up to a factor $R \in \mathbb{O}_n$. The next lemma shows that the second pair provides additional constraints on R, but also that there are constraints on (P_2, Q_2) if a solution V is to exist:

Lemma 4 *Given a nonsingular $V \in \mathbb{R}^{n \times n}$, and positive definite pairs (P_1, Q_1) and (P_2, Q_2), define*

$$M := Q_1^{-1/2} Q_2 Q_1^{-1/2} \text{ and } N := P_1^{-1/2} P_2 P_1^{-1/2}.$$

Then, then following statements are equivalent:

1. *The following equalities hold:*

$$Q_1 = V P_1 V^T,\ Q_2 = V P_2 V^T. \tag{2}$$

2. *$V = Q_1^{1/2} R P_1^{-1/2}$, for some $R \in \mathbb{O}_n$ such that $MR = RN$.*

Proof Suppose that property 1 is true, so that Eq. (2) holds. It follows from Theorem 3 that there is some $R \in \mathbb{O}_n$ such that $V = Q_1^{1/2} R P_1^{-1/2}$. For this R, $Q_2 = V P_2 V^T$ if and only if $MR = RN$, because of the following equivalences:

$$\begin{aligned} Q_2 = V P_2 V^T &\Leftrightarrow Q_2 = (Q_1^{1/2} R P_1^{-1/2}) P_2 (P_1^{-1/2} R^T Q_1^{1/2}) \\ &\Leftrightarrow (Q_1^{-1/2} Q_2 Q_1^{-1/2}) = R (P_1^{-1/2} P_2 P_1^{-1/2}) R^T \\ &\Leftrightarrow M = R N R^T \\ &\Leftrightarrow MR = RN. \end{aligned}$$

Conversely, suppose that property 2 is true, for some R. Following the above equivalences backwards, we conclude that $Q_2 = V P_2 V^T$, so property 1 holds. □

Observe that, in the previous lemma, the (symmetric and positive-definite) matrices M and N are similar (with an orthogonal similarity transformation R), and thus M and N have the same eigenvalues. The following corollary is clear from the proof and is stated for future reference:

Corollary 5 *Let M, N be defined as in Lemma 4. Then there exist $U_N, U_M \in \mathbb{O}_n$ and a unique diagonal matrix $\Lambda \in \mathbb{R}^n$ with the entries on the diagonal sorted in descending order such that*

$$M = U_M \Lambda U_M^T, N = U_N \Lambda U_N^T. \tag{3}$$

We already pointed out that M and N have the same eigenvalues. We next assume that these eigenvalues are *distinct*. We will show in a later section that this assumption holds generically. We are now ready to state the main result in this subsection. We will say that an $n \times n$ matrix Θ is a *signature* matrix if it is diagonal and $\Theta^2 = I$, that is, its diagonal entries are $+1$ or -1. Every signature matrix is orthogonal, since $\Theta^T \Theta = \Theta^2 = I$.

Theorem 6 *Given a nonsingular $V \in \mathbb{R}^{n \times n}$, and positive definite pairs (P_1, Q_1) and (P_2, Q_2), define*

$$M := Q_1^{-1/2} Q_2 Q_1^{-1/2} \text{ and } N := P_1^{-1/2} P_2 P_1^{-1/2}.$$

Suppose that N has distinct eigenvalues, and M, N are orthogonally similar. Then, then following statements are equivalent:

1. *The following equalities hold:*

$$Q_1 = V P_1 V^T,\ Q_2 = V P_2 V^T. \tag{4}$$

2. *There exist orthogonal matrices U_M and U_N and a signature matrix Θ such that*

$$V = Q_1^{1/2} U_M \Theta U_N^T P_1^{-1/2}. \tag{5}$$

Proof Suppose that property 1 holds. By Lemma 4, we may pick some $R \in \mathbb{O}_n$ such that $MR = RN$, and with this R we have that $V = Q_1^{1/2} R P_1^{-1/2}$. We can then pick U_M, U_N, and Λ as in Corollary 5. Define $\Theta := U_M^T R U_N$, which means that $R = U_M \Theta U_N^T$. Substituting into $V = Q_1^{1/2} R P_1^{-1/2}$, we get the desired equality $V = Q_1^{1/2} U_M \Theta U_N^T P_1^{-1/2}$ Thus, we only need to show that Θ is a signature matrix. The matrix Θ is orthogonal, since orthogonal matrices form a group under matrix multiplication and each of U_M^T, R, and U_N is orthogonal. Observe that

$$\begin{aligned} R = RN \Leftrightarrow (U_M \Lambda U_M^T) R = R(U_N \Lambda U_N^T) \Leftrightarrow \\ \Lambda(U_M^T R U_N) = (U_M^T R U_N)\Lambda \Leftrightarrow \Lambda\Theta = \Theta\Lambda. \end{aligned} \tag{6}$$

Since we know that $MR = RN$, the last equivalence shows that Λ commutes with Θ. Also N has distinct eigenvalues, so it follows that Λ, which is similar to N, has distinct eigenvalues as well. It is an easy exercise in linear algebra to show that an orthogonal matrix that commutes with a diagonal nonsingular matrix must be diagonal. Thus Θ is not only orthogonal but is also diagonal, and therefore $I = \Theta\Theta^T = \Theta^2$, showing that Θ is a signature matrix as claimed.

Conversely, suppose that property 2 holds. Define $R := U_M \Theta U_N^T$, which is an orthogonal matrix. Since M and N are orthogonally similar, they share eigenvalues that can be arranged in a diagonal matrix Λ, which would commute with the signature matrix Θ. Following Eq. (6) in reverse, one finds that $MR = RN$. Thus, by the implication (2) ⇒ (1) in Lemma 4, we have that $Q_1 = VP_1V^T$ and $Q_2 = VP_2V^T$. □

Thus, a second pair of covariance matrices specifies 2^n possible matrices $e^{tA} = Q_1^{1/2} U_M \Theta U_N^T P_1^{-1/2}$, for all the possible signature matrices Θ. Intuitively, this signature matrix compensates for the arbitrary choice of sign on the basis U_M relative to U_N.

Remark 2 *Without assuming Λ has distinct eigenvalues, there are other non-diagonal solutions that satisfy $\Theta \in \mathbb{O}_n$ and $\Lambda\Theta = \Theta\Lambda$. Hence, the number of possible solutions would be much larger than 2^n.*

Two further pairs of means (ρ_1, θ_1) and (ρ_2, θ_2) determine e^{tA} and x_{ss}

Using the first moments of the distributions at different times, we can find additional constraints on e^{tA}, provided the means ρ_1, ρ_2, θ_1, θ_2 are constrained so that a solution for e^{tA} exists.

Lemma 7 *Given a nonsingular $V \in \mathbb{R}^{n\times n}$ and $x_{ss} \in \mathbb{R}^n$, and quadruples $(P_1, Q_1, \rho_1, \theta_1), (P_2, Q_2, \rho_2, \theta_2) \in \mathbb{S}_n^+ \times \mathbb{S}_n^+ \times \mathbb{R}^n \times \mathbb{R}^n$, and for U_M, U_N as defined in Corollary 5, define*:

$$\vec{v} := U_M^T Q_1^{-1/2}(\theta_1 - \theta_2) \qquad \text{and} \qquad \vec{w} := U_N^T P_1^{-1/2}(\rho_1 - \rho_2).$$

Assume that N has distinct eigenvalues, and M, N are orthogonally similar as in Theorem 6. Then, the following statements are equivalent:

1. *The following equalities hold*:

$$Q_1 = VP_1V^T,\ Q_2 = VP_2V^T, \tag{7}$$

$$\theta_1 = V(\rho_1 - x_{ss}) + x_{ss},\ \theta_2 = V(\rho_2 - x_{ss}) + x_{ss}. \tag{8}$$

2. *For some signature matrix Θ, such that $\vec{v} = \Theta\vec{w}$*:

$$V = Q^{1/2} U_M \Theta U_N^T P_1^{-1/2} \tag{9}$$

$$x_{ss} = (V - I)^{-1}(V\rho_1 - \theta_1). \tag{10}$$

Proof Suppose that property 1 holds. We rearrange the equations for θ_1, θ_2 and subtract them from each other:

$$\begin{aligned} \theta_1 - x_{ss} &= V(\rho_1 - x_{ss});\ \theta_2 - x_{ss} = V(\rho_2 - x_{ss}), \\ &\Rightarrow \theta_1 - \theta_2 = V(\rho_1 - \rho_2) \end{aligned}$$

From Theorem 6, we know that $V = Q_1^{1/2} U_M \Theta U_N^T P_1^{-1/2}$, for some signature matrix Θ, so:

$$\begin{aligned} \theta_1 - \theta_2 &= V(\rho_1 - \rho_2) \\ \theta_1 - \theta_2 &= Q_1^{1/2} U_M \Theta U_N^T P_1^{-1/2}(\rho_1 - \rho_2) \\ U_M^T Q_1^{-1/2}(\theta_1 - \theta_2) &= \Theta U_N^T P_1^{-1/2}(\rho_1 - \rho_2). \end{aligned} \tag{11}$$

Since we defined $\vec{v} := U_M^T Q_1^{-1/2}(\theta_1 - \theta_2)$ and $\vec{w} := U_N^T P_1^{-1/2}(\rho_1 - \rho_2) \in \mathbb{R}^n$, we get $\vec{v} = \Theta\vec{w}$.

To solve for x_{ss}, we simply re-arrange the expression for θ_1:

$$\begin{aligned} \theta_1 &= V(\rho_1 - x_{ss}) + x_{ss} \\ V\rho_1 - \theta_1 &= Vx_{ss} - x_{ss} \\ (V - I)^{-1}(V\rho_1 - \theta_1) &= x_{ss} \end{aligned}$$

Conversely, suppose that property 2 holds. By Theorem 6, we know that $Q_1 = VP_1V^T$ and $Q_2 = VP_2V^T$. By simple re-arrangement of the expression for x_{ss}, we arrive back at $\theta_1 = V(\rho_1 - x_{ss}) + x_{ss}$. Furthermore, by following Equations (11) in reverse, we arrive back at $\theta_1 - \theta_2 = V(\rho_1 - \rho_2)$. We may then subtract the expression for $\theta_1 - \theta_2$ from the expression of θ_1 as follows:

$$\begin{aligned} \theta_1 - (\theta_1 - \theta_2) &= V(\rho_1 - x_{ss}) + x_{ss} - (V\rho_1 - V\rho_2) \\ \theta_2 &= V(\rho_2 - x_{ss}) + x_{ss}, \end{aligned}$$

completing statement 1. □

Provided the existence constraint from the above lemma, we may solve for e^{tA} and x_{ss}:

Theorem 8 *Given a nonsingular $V \in \mathbb{R}^{n\times n}$ and $x_{ss} \in \mathbb{R}^n$, and quadruples $(P_1, Q_1, \rho_1, \theta_1), (P_2, Q_2, \rho_2, \theta_2) \in \mathbb{S}_n^+ \times \mathbb{S}_n^+ \times \mathbb{R}^n \times \mathbb{R}^n$, and for U_M, U_N as defined in Corollary 5, define*:

$$\vec{v} := U_M^T Q_1^{-1/2}(\theta_1 - \theta_2) \text{ and } \vec{w} := U_N^T P_1^{-1/2}(\rho_1 - \rho_2).$$

Assume $v_i \neq 0$ and $w_i \neq 0, \forall\, i$, and that $\vec{v} = \Theta\vec{w}$ for some signature matrix Θ. Also assume that M, N are orthogonally similar, and that N has distinct eigenvalues, as in Theorem 6. Then, then following statements are equivalent:

1. *The following equalities hold*:

$$Q_1 = VP_1V^T,\ Q_2 = VP_2V^T, \tag{12}$$

$$\theta_1 = V(\rho_1 - x_{ss}) + x_{ss},\ \theta_2 = V(\rho_2 - x_{ss}) + x_{ss}. \tag{13}$$

2. *There exists a signature matrix $\mathcal{D}$ given by $\mathcal{D}_{ii} := \mathrm{sgn}(v_i/w_i)$, such that*:

$$V = Q_1^{1/2} U_M \mathcal{D} U_N^T P_1^{-1/2}. \tag{14}$$

$$x_{ss} = (V - I)^{-1}(V\rho_1 - \theta_1). \tag{15}$$

Proof Suppose that property 1 holds. By Lemma 7, we know that $V = Q^{1/2} U_M \Theta U_N^T P_1^{-1/2}$, and that $\vec{v} = \Theta\vec{w}$. Since Θ is a signature matrix, then $v_i = \Theta_{ii} w_i$, and so v_i and w_i only differ by a sign. We may find a solution $\mathcal{D}$ for Θ by simply taking $\mathcal{D}_{ii} := \mathrm{sgn}(v_i/w_i)$, choosing the notation sgn() to emphasize that $v_i/w_i = \pm 1$. Also by Lemma 7, we know that $x_{ss} = (V - I)^{-1}(V\rho_1 - \theta_1)$.

Conversely, suppose that property 2 holds. Since $\vec{v} = \Theta\vec{w}$ for some signature matrix Θ, the definition of $\mathcal{D}$ implies that $\Theta = \mathcal{D}$. Thus, by the implication (2) ⇒ (1) in Lemma 7, we recover statement 1. □

There is a unique solution for Θ so long as $v_i, w_i \neq 0\ \forall\, i$, which we show in a later section is generically true. Thus, $V = e^{tA}$ and x_{ss} are uniquely determined from two pairs of means and covariances.

Three pairs of covariances (P_i, Q_i) can determine $\pm e^{tA}$

We have established that two pairs of covariance matrices alone are not enough to determine e^{tA} uniquely, due to ambiguity in "reflections" represented by a signature matrix Θ. A third pair of covariance matrices (P_3, Q_3) is enough to resolve this ambiguity, provided a solution e^{tA} exists:

Lemma 9 *Given a nonsingular $V \in \mathbb{R}^{n\times n}$, and positive definite pairs $(P_1, Q_1), (P_2, Q_2), (P_3, Q_3)$, let U_M, U_N be defined as in Corollary 5, and define*:

$$X := U_N^T P_1^{-1/2} P_3 P_1^{-1/2} U_N \text{ and } Y := U_M^T Q_1^{-1/2} Q_3 Q_1^{-1/2} U_M.$$

Assume that N has distinct eigenvalues, and M, N are orthogonally similar as in Theorem 6. Then, the following statements are equivalent:

1. *The following equalities hold*:

$$Q_1 = VP_1V^T,\ Q_2 = VP_2V^T,\ Q_3 = VP_3V^T \tag{16}$$

2. *For some signature matrix Θ, $V = Q_1^{1/2} U_M \Theta U_N^T P_1^{-1/2}$, such that $Y = \Theta X \Theta$.*

Proof Suppose that property 1 holds. We substitute the previous solution for V from Theorem 6 into the analogous equation to Eq. (1) for the third pair:

$$\begin{aligned} Q_3 = VP_3V^T &\Leftrightarrow Q_3 = Q_1^{1/2} U_M \Theta U_N^T P_1^{-1/2} P_3 (Q_1^{1/2} U_M \Theta U_N^T P_1^{-1/2})^T \\ &\Leftrightarrow U_M^T Q_1^{-1/2} Q_3 Q_1^{-1/2} U_M = \Theta (U_N^T P_1^{-1/2} P_3 P_1^{-1/2} U_N) \Theta, \\ &\Leftrightarrow Y = \Theta X \Theta. \end{aligned} \tag{17}$$

Conversely, suppose that property 2 holds. Since Θ is a signature matrix, and M and N are orthogonally similar, then $Q_1 = VP_1V^T$ and $Q_2 = VP_2V^T$ by Theorem 6. By following Equations (17) in reverse, we get $Q_3 = VP_3V^T$, completing statement 1. □

Provided the existence constraint from the above lemma, we may solve for e^{tA}:

Theorem 10 *Given a nonsingular $V \in \mathbb{R}^{n\times n}$, and positive definite pairs $(P_1, Q_1), (P_2, Q_2), (P_3, Q_3)$, let U_M, U_N be defined as in Corollary 5, and define*

$$X := U_N^T P_1^{-1/2} P_3 P_1^{-1/2} U_N \text{ and } Y := U_M^T Q_1^{-1/2} Q_3 Q_1^{-1/2} U_M.$$

Assume that the node-edge graph $\mathcal{G}_X$ associated to X, treating each non-zero entry $[X]_{ij}$ as an edge weight between the n nodes, has a single connected component, and that $Y = \Theta X \Theta$ for some signature matrix Θ. Also, assume that N has distinct eigenvalues, and M, N are orthogonally similar as in Theorem 6. Then, the following statements are equivalent:

1. *The following equalities hold*:

$$Q_1 = VP_1V^T,\ Q_2 = VP_2V^T,\ Q_3 = VP_3V^T \tag{18}$$

2. *There exist two signature matrices, $\pm\mathcal{S}$, given by $\mathcal{S}_{11} = 1$ and the remaining diagonal entries $\mathcal{S}_{ii} = \text{sgn}([Y]_{ij}/[X]_{ij})$, $i = 2, .., n, i \neq j$, such that either $V = +Q_1^{1/2} U_M \mathcal{S} U_N^T P_1^{-1/2}$ or $V = -Q_1^{1/2} U_M \mathcal{S} U_N^T P_1^{-1/2}$.*

Proof Suppose that property 1 holds. By Lemma 9, we know that for some signature matrix Θ, $V = Q^{1/2} U_M \Theta U_N^T P_1^{-1/2}$ and that $Y = \Theta X \Theta$. The last equation takes the following form entrywise:

$$y_{ij} = \Theta_i \Theta_j x_{ij} \tag{19}$$

for $Y = y_{ij}$, $X = x_{ij}$, and the diagonal entries $\Theta_i = \pm 1$ of Θ. Thus, from the sign of y_{ij}/x_{ij} (provided that $x_{ij} \neq 0$), we may determine the relative sign of Θ_i to Θ_j. Starting by assuming $\Theta_1 = 1$, we can determine Θ_j for all neighboring j'th nodes of node 1 on $\mathcal{G}_X$. Similarly, since $\mathcal{G}_X$ has a single connected component, one may determine the signs of $\Theta_j\ \forall\ j$ by following the edges of $\mathcal{G}_X$. Hence, we define $\mathcal{S}$ to be the diagonal matrix with non-zero entries $\mathcal{S}_{11} = 1$, and $\mathcal{S}_{ii} = \text{sgn}(y_{ij}/x_{ij})$, for $i \geq 2$ and any $j \neq i$, using the notation sgn() to emphasize that $y_{ij}/x_{ij} = \pm 1$. Then $\Theta = \pm\mathcal{S}$ are the two solutions to $Y = \Theta X \Theta$, giving the final solutions $V = \pm Q_1^{1/2} U_M \mathcal{S} U_N^T P_1^{-1/2}$.

Conversely, suppose that property 2 holds. Since we assume that $Y = \Theta X \Theta$ for some signature matrix Θ, the definitions of $\pm\mathcal{S}$ imply that $\Theta = \pm\mathcal{S}$. Thus, by implication (2) ⇒ (1) in Lemma 9, we recover the equalities of statement 1. □

We will show in a later section it is generically true that $\mathcal{G}_X$ has a single connected component.

Remark 3 *Note that determining V up to a sign is the best one can do using covariances $\{P_i, Q_i\}$, since*:

$$P_i = VQ_iV^T = (-V)Q_i(-V^T). \tag{20}$$

Thus, three pairs (P_i, Q_i) determine V up to a sign, and additional pairs provide no further constraints on V.

Genericity of assumptions on (P_i, Q_i) and (ρ_i, θ_i)

We first show three lemmas that will be used in proving genericity, noting that we will use "analytic" or "algebraic" sets to mean sets with positive codimension. The first lemma is, of course, well-known.

Lemma 11 *The set $Z_1 \subset \mathbb{S}_n^+$ of matrices with repeated eigenvalues is an algebraic set.*

Proof Let π_B be the characteristic polynomial of B. The discriminant of π_B (the resultant of π_B and its derivative π_B') for a matrix $B \in Z_1$ with repeated eigenvalues will be zero. The coefficients of π_B are polynomial functions of the entries b_{ij} of B. Thus, Z_1 is an algebraic set. □

Lemma 12 *The uniquely defined positive definite square root $P^{1/2}$ of a positive definite matrix $P \in \mathbb{S}_n^+$ has entries that are analytic functions of the entries of P.*

Proof This follows by applying the tools of holomorphic functional calculus, which insure (for bounded operators in a Banach space) that there is an extension of an analytic function $f(z)$ to a matrix $f(P)$ provided that f is analytic in an open neighborhood of the spectrum of P. The function $\sqrt{z}$ is analytic on the positive real axis $z > 0$, and P has all eigenvalues in this set, so $\sqrt{P}$ is a special case.

An alternative proof of analyticity is by applying the Implicit Function Theorem to the map F: $Q \mapsto Q^2$. The differential of this mapping at a point Q_0 is given by the linear operator L: $Q \mapsto Q_0Q + QQ_0$, and the eigenvalues of this operator are the n^2 pairwise sums of eigenvalues of Q_0. For Q_0 positive definite, these pairwise sums are always positive and hence nonzero. Therefore L is onto and thus F is nonsingular at Q_0, which implies that the solution of $F(Q) = P$ is analytic on P. □

Lemma 13 *Given a positive definite matrix B with distinct eigenvalues, and its eigendecomposition $B = U_B \Lambda_B U_B^T$, the entries of U_B can be chosen as analytic functions of the entries of B.*

Proof For B a positive definite matrix with distinct eigenvalues, let

$$\lambda_{1,B} > \lambda_{2,B} > \ldots > \lambda_{n,B}$$

be its (real) eigenvalues, ordered from largest to smallest. In general, if a polynomial f has a zero of multiplicity one at a point a, then $f'(a) \neq 0$, since $f(x) = (x - a)g(x)$ with $g(a) \neq 0$, and $f'(a) = g(a) + (a - a)g'(a) = g(a)$. This implies that the functions $\lambda_{i,B}$ are analytic on B, as verified by applying the implicit mapping theorem to the implicit equation $\pi_B(z) = 0$ and using that $\pi_B(\lambda_{i,B}) = 0$ and $\pi_B'(\lambda_{i,B}) \neq 0$.

Let Z_2 be the subset of $\mathbb{S}_n^+$ consisting of matrices B for which, for some $i = 1, \ldots, n$, the determinant of the submatrix B_i formed by the last $n - 1$ columns of $B - \lambda_{i,B} I$ has rank $n - 2$. Since B has rank $n - 1$, the rank of B_i can only be $n - 1$ or $n - 2$. Thus asking that B_i have rank $n - 2$ is equivalent to asking that all the $(n - 1) \times (n - 1)$ minors of B_i should be zero. The complement of Z_2 consists of matrices for which all B_i have full column rank $n - 1$. Notice that, since minors are polynomials on the entries of B and the

$\lambda_{i,B}$ are analytic functions on the entries of B, the set Z_2 is an analytic variety. We will let $\mathcal{P}$ denote the complement of the analytic variety $Z_1 \cup Z_2$, for Z_1 as defined in Lemma 11. This is a generic subset of $\mathbb{S}_n^+$, in the sense that it is the complement of an analytic variety, and therefore has measure zero and is also open dense. Any matrix $B \in \mathcal{P}$ admits an orthogonal decomposition $B = U_B \Lambda_B U_B^T$, where

$$\Lambda_B := \text{diag}\,(\lambda_1, B, \lambda_2, B, \ldots, \lambda_n, B).$$

The columns of the matrix U_B are an ordered set $u_{1,B}$, $u_{2,B}$, …, $u_{n,B}$ of orthogonal eigenvectors corresponding to the respective $\lambda_{i,B}$. These vectors $u_{i,B}$ are only unique up to multiplication by ± 1 (because they are multiples of each other and have unit norm). Notwithstanding this non-uniquness, we show next that one can pick the eigenvectors $u_{1,B}$, $u_{2,B}$, …, $u_{n,B}$ as analytic functions of B.

To do so, we will first pick eigenvectors that are orthogonal to each other but not yet of unit norm, in fact with first coordinate equal to 1.

Fix any i and consider the equation $(B - \lambda_{i,B} I)u = 0$. A vector u of the form $(1, r^T)^T$ solves $(B - \lambda_{i,B} I)u = 0$ if and only if $c + B_i r = 0$, and B_i has rank $n - 1$ because $B \in \mathcal{P}$. This means that $r = -B_i^\sharp c$, where $B_i^\sharp$ is the Moore-penrose pseudoinverse of B_i, $B_i^\sharp = (B_i^T B_i)^{-1} B_i^T$, which is an analytic function of the entries of B_i, and hence of B. Finally, we can obtain unit vectors by normalizing u, which preserves analyticity. □

We now show that various assumptions used in the proofs of previous theorems are generic.

Genericity: N has distinct eigenvalues

In Theorem 6, we assume that a matrix $N := P_1^{-1/2} P_2 P_1^{-1/2}$ has distinct eigenvalues, for $P_i \in \mathbb{S}_n^+$.

Lemma 14 *The set $\mathcal{U}$ of matrix-tuples $(P_1, P_2) \in \mathbb{S}_n^+ \times \mathbb{S}_n^+$ for which $N = P_1^{-1/2} P_2 P_1^{-1/2}$ has distinct eigenvalues is open, dense, and has complement of measure zero.*

Proof The set of matrices $N \in \mathbb{S}_n^+$ with distinct eigenvalues constitutes an open subset of $\mathbb{S}_n^+$ that is the complement of a closed algebraic set Z_1 from Lemma 11. Furthermore, the entries n_{ij} of N are polynomials of the entries of $P_1^{-1/2}$ and P_2, and the entries of $P_1^{-1/2}$ are analytic functions of the entries of P_1 by Lemma 12. Thus, since polynomial functions of analytic functions are analytic, the set $\mathcal{U}$ of matrix-tuples $(P_1, P_2) \in \mathbb{S}_n^+ \times \mathbb{S}_n^+$ that produce a matrix N with distinct eigenvalues is generic because it is the complement of an analytic set (which has measure zero). □

Genericity: there is a unique solution to $v_i = \Theta_{ii} w_i$

In Theorem 8, we assume that a vector $\vec{w} := U_N^T P_1^{-1/2} (\rho_1 - \rho_2)$ has no zero entries, for U_N an eigenvector basis of a positive definite matrix $N = P_1^{-1/2} P_2 P^{-1/2}$ with distinct eigenvalues, $P_1 \in \mathbb{S}_n^+$, and $\rho_i \in \mathbb{R}^n$.

Lemma 15 *The set $\mathcal{V}$ of matrix-vector-tuples $(P_1, P_2, \rho_1, \rho_2) \in \mathbb{S}_n^+ \times \mathbb{S}_n^+ \times \mathbb{R}^n \times \mathbb{R}^n$ for which the entries w_i of $\vec{w}$ are nonzero is open, dense, and has complement of measure zero.*

Proof The entries w_i are polynomials in terms of the respective entries of U_N, $P_1^{-1/2}$, ρ_1, and ρ_2. The entries of $P_1^{-1/2}$ are analytic functions of P_1 by Lemma 12. The entries of U_N are analytic functions of $N = P_1^{-1/2} P_2 P_1^{-1/2}$ by Lemma 13, and therefore also analytic functions of the entries of P_1 and P_2. Hence, the entries w_i are analytic functions of $(P_1, P_2, \rho_1, \rho_2)$. The vector w has a zero entry on the union of the n analytic sets defined by $w_i = 0$. Thus, take $\mathcal{V}$ to be the complement of these analytic sets. □

Genericity: $\mathcal{G}_X$ has one connected component

In Theorem 10, we assume that a matrix $X := U_N^T P_1^{-1/2} P_3 P_1^{-1/2} U_N$ has the property that its associated node-edge graph $\mathcal{G}_X$ has a single connected component. We will show that the special case in which G_X is fully connected (i.e., X has no zero entries) is already generic.

Lemma 16 *The set $\mathcal{W}$ of matrix-tuples $(P_1, P_2, P_3) \in \mathbb{S}_n^+ \times \mathbb{S}_n^+ \times \mathbb{S}_n^+$ for which all the entries of X are nonzero is open, dense, and has complement of measure zero.*

Proof The entries x_{ij} of X are polynomials in terms of the respective entries of U_N, $P_1^{-1/2}$, and P_3. The entries of $P_1^{-1/2}$ are analytic functions of P_1 by Lemma 12. The entries of U_N are analytic functions of $N = P_1^{-1/2} P_2 P_1^{-1/2}$ by Lemma 13, and therefore also analytic functions of the entries of P_1 and P_2. Hence, the entries x_{ij} are analytic functions of (P_1, P_2, P_3).

For a given simple node-edge graph $\mathcal{G}$ on n nodes with adjacency matrix $A_\mathcal{G}$, define the analytic set:

$$V_\mathcal{G} := \{(P_1, P_2, P_3) \in \mathbb{S}_n^+ \times \mathbb{S}_n^+ \times \mathbb{S}_n^+ \mid x_{ij} = 0 \,\forall\, i, j \text{ s.t. } A_{\mathcal{G}ij} = 0\}$$

There are a finite number K of graphs $\mathcal{G}_k$ with n nodes and more than one connected component. Thus,

$$Z_\mathcal{G} := \cup_k^K V_{\mathcal{G}_k}$$

is the closed analytic set of all (P_1, P_2, P_3) that result in a graph with more than one connected component. The complement $(\mathbb{S}_n^+ \times \mathbb{S}_n^+ \times \mathbb{S}_n^+) \setminus Z_\mathcal{G}$ of (P_1, P_2, P_3)'s resulting in one connected component is therefore open with complement of measure zero. □

Determining A given $V = e^{tA}$ or ± V

In the previous sections, we have shown how to determine V or $\pm V$ uniquely. Given a unique $V = e^{tA}$, we can then solve for A up to a discrete equivalence class: if some of V's eigenvalues are complex, the matrix logarithm $\log(V)$ would be multi-valued[21]. Concretely, we diagonalize $V = W D_V W^{-1}$ (which is diagonalizable by the assumption that A is diagonalizable), and take the logarithms of its eigenvalues to solve for A, then:

$$A = \frac{1}{t} W \log(D_V) W^{-1} \tag{21}$$

where each eigenvalue $z_j = r_j e^{i\theta}$ along the diagonal D_V has multiple logarithms $\log(r_j) + i(\theta + 2\pi k)$ for $k \in \mathbb{Z}$. Thus, we may solve up to an equivalence class of linear dynamical systems $\{A_{\overrightarrow{k}}\}$, for $\overrightarrow{k} \in \mathbb{Z}^n$ denoting the possible imaginary shifts in the period of the eigenvalues. All members of $\{A_{\overrightarrow{k}}\}$ share common eigenvectors and real-parts of eigenvalues, which can often be represented by the uniquely-defined *principal* logarithm $\text{Log}(V)$ in $A_p = \frac{1}{t}\text{Log}(V)$ with imaginary components restricted to $[-\pi, \pi]$.

If we assume for a given $t > 0$, it is known that A belongs to the set $\mathcal{A}_{n,t}$ of matrices with eigenvalues $\lambda_j = a_j + ib_j$ for which $|b_j| < \pi/t$ (i.e., the oscillatory components of solutions $x(t)$ to $\dot{x} = Ax$ have periods larger than the time interval t), then we can uniquely identify A as $A = \frac{1}{t}\text{Log}(V)$.

Lemma 17 *Suppose $A \in \mathcal{A}_{n,t}$. Given $V = e^{tA}$, $A = \frac{1}{t} Log(V)$.*

Proof The j'th eigenvalue of $\log(V)$ can be any of $l_j = ta_j + i(tb_j + 2\pi k)$ for $k \in \mathbb{Z}$, for a_j, b_j the corresponding real and imaginary components of A's eigenvalue λ_j. Suppose that $|b_j| < \pi/t$, i.e.,

$$-\pi < tb_j < \pi, \ \forall j.$$

Then the principal logarithm $\text{Log}(V)$ corresponds to the choice of $k = 0, \forall j$, and $\text{Log}(V) = tA$. □

Remark 4 *The requirement on the time interval coincides with that of the Shannon–Nyquist Theorem[22], but our result differs since there is no single signal $x(t)$ being measured, and even hypothetical solutions $x(t)$ fulfilling $\dot{x} = Ax$ contain infinitely high frequencies from the real exponential components of $x(t)$.*

In the case that we only identify ± V, if we define the set $\mathcal{B}_{n,t}$ of real $n \times n$ matrices such that $|b_j| < \pi/(2t)$, $\forall\, j$, then only one of $\text{Log}(\pm V)$ will belong to $\mathcal{B}_{n,t}$.

Lemma 18 *Suppose $A \in \mathcal{B}_{n,t}$. Given $V = e^{tA}$ and $-V$, $A = \frac{1}{t} Log(V)$ and $\frac{1}{t} Log(-V) \notin B_{n,t}$.*

Proof Since $\mathcal{B}_{n,t} \subset \mathcal{A}_{n,t}$, then by Lemma 17, $A = \frac{1}{t} Log(V)$.

Given that V can be diagonalized as $V = WD_VW^{-1}$ with eigenvalues $z_j = x_j + iy_j$, then $-V = W(-D_V)W^{-1}$ with corresponding eigenvalues $-(x_i + iy_j)$ that have an additional imaginary phase of π. Thus, for $\lambda_j = ta_j + itb_j$ the eigenvalues of $\mathrm{Log}(V)$, the eigenvalues of $\frac{1}{t}\,\mathrm{Log}\,(-\mathrm{V})$ are:

$$l_j = a_j + i(b_j \pm \pi/t).$$

Supposing that $A \in \mathcal{B}_{n,t}$, i.e., $|b_j| < \pi/(2t)$, then either:

$$\begin{gathered}-\pi/(2t) + \pi/t < b_j + \pi/t < \pi/(2t) + \pi/t \ \text{ or } \ -\pi/(2t) - \pi/t < b_j - \pi/t < \pi/(2t) - \pi/t \\ \Rightarrow \pi/(2t) < b_j + \pi/t < 3\pi/(2t) \ \text{ or } \ -3\pi/(2t) < b_j - \pi/t < \pi/(2t).\end{gathered}$$

which both imply that $|b_j \pm \pi/t| > \pi/(2t)$, and that therefore $\frac{1}{t}\,\mathrm{Log}\,(-\mathrm{V}) \notin \mathcal{B}_{\mathrm{n,t}}$. □

Remark 5 *If even one of V's eigenvalues is real, and positive, and distinct from V's other eigenvalues, then* $\log(-V)$ *cannot be a real matrix because* $-V$ *would have a real, negative eigenvalue distinct from the other eigenvalues of* $-V$[21]. *In such cases, one can identify V and then apply Lemma 17. Thus, Lemma 18 is only used when all eigenvalues of V are complex.*

Finally, we note that one might also consider solving for A, given V, by making use of data for which the time intervals vary. This possibility is motivated by:

Lemma 19 *For a diagonalizable and invertible A, suppose* $e^{tB} = e^{tA}$, *for some* $t > 0$ *and* $A \neq B$. *Given* $s > 0$, $e^{sA} \neq e^{sB}$ *for generic* $s \in \mathbb{R}_{>0}$, *i.e., when s/t is irrational.*

Proof If $e^{tB} = e^{tA}$, then each eigenvalue λ_B of B differs from some eigenvalue λ_A of A by:

$$\lambda_B - \lambda_A = \frac{2\pi k_1}{t}, \qquad \text{for } k_1 \in \mathbb{Z}.$$

In order for $e^{sA} = e^{sB}$, their eigenvalues must be the same. The eigenvalues of e^{sA} and e^{sB} will be $e^{s\lambda_A}$ and $e^{s\lambda_B}$, and can only be equal if:

$$\lambda_B - \lambda_A = \frac{2\pi k_2}{s}, \qquad \text{for } k_2 \in \mathbb{Z}.$$

Thus, $e^{sA} = e^{sB}$ only if:

$$\frac{2\pi k_1}{s} = \frac{2\pi k_2}{t} \ \Leftrightarrow \ \frac{s}{t} = \frac{k_1}{k_2}, \qquad \text{for } k_1, k_2 \in \mathbb{Z}.$$

Since $e^{sA} = e^{sB}$ requires s/t to be rational, and the rationals have measure zero in $\mathbb{R}_{>0}$, generically $e^{sA} \neq e^{sB}$. □

For example, suppose that after collecting data in an experiment with three time points at intervals t, one has determined that the true $A \in \{A \underset{k}{\rightarrow}\}$, for which $e^{tA} = e^{tB}$ for any $B \in \{A \underset{k}{\rightarrow}\}$. If this experiment is repeated, changing only the interval to s so that s/t is irrational, one may determine $V_s = e^{sA}$. By the above lemma, only the true $A \in \{A \underset{k}{\rightarrow}\}$ will have the property that $e^{sA} = V_s$, since $e^{sB} \neq e^{sA} = V_s$ for all other $B \in \{A \underset{k}{\rightarrow}\}$. It is possible that fewer time points (at different intervals) are needed to uniquely determine A, but we do not investigate this further here.

Remark 6 *In practice, data will always contain noise from finite sampling or measurement noise, and so it is likely difficult to distinguish between the statistical differences of data taken at intervals s, t for s/t irrational versus rational, since the rationals are dense. Still, uneven intervals clearly help eliminate possible choices for A: whereas with a single interval t, we must assume, for example, that* $|b_j| < \pi/t$, *if* $s/t = p/q$ *for* $p, q \in \mathbb{Z}$, *then for* $a := \mathrm{lcm}(p, q)$ *the least common multiple of p and q, we can loosen the requirement to* $|b_j| < a\pi/t$.

Proof of main theorems

To prove Theorem 1, we choose the generic set S to be $(\mathcal{U} \times \mathbb{R}^n \times \mathbb{R}^n) \cap \mathcal{V}$ from Lemmas 14 and 15. By Theorem 8, we determine $V = e^{tA}$ and x_{ss} uniquely. By Lemma 17, we then recover A uniquely from V.

To prove Theorem 2, we choose the generic set T to be $(\mathcal{U} \times \mathbb{S}_n^+) \cap \mathcal{W}$ from Lemmas 14 and 16. Then, by Theorem 10, we may determine $V = e^{tA}$ and $-V$. By Lemma 18, we may then recover A uniquely from $\pm V$.

Approximate solutions for e^{tA}

Throughout, we have assumed the existence of a solution V to the simultaneous equations $Q_i = VP_iV^T$ for $i = 1, 2, \ldots, k$, but in general there does not exist a solution for $k > 1$ and arbitrary (P_i, Q_i). This is already strongly suggested by the fact that V represents n^2 indeterminates v_{ij}, and that for each i the equation $Q_i = VP_iV^T$ represents $(n^2 + n)/2$ quadratic equations. For $k = 2$, there are already $n^2 + n$ equations, suggesting an overdetermined system of equations. This is implied, for example, in Corollary 5, in which the two otherwise arbitrary positive-definite matrices M, N must be similar if a solution V is to exist: if P_1, P_2, Q_1, Q_2 were simply four arbitrary matrices in $\mathbb{S}_n^+$, M and N could be any elements of $\mathbb{S}_n^+$ and would therefore not necessarily be similar to each other.

In the absence of exact solutions V to the simultaneous quadratic matrix equations $Q_i = VP_iV^T$, one may still desire approximate solutions V that minimize total error $\min_V \sum_i d(Q_i, VP_iV^T)$ for some metric $d(X, Y)$, e.g., in applications to regression. We are not aware of an analytical solution to this particular minimization problem, although various other quadratic optimization problems have been studied before[23,24]. However, it is possible to reframe the respective solutions of V from the first, second, and third pairs of covariances (P_i, Q_i) as well-known quadratic optimization problems, and sequentially solving these three optimization problems may at least provide informed initial guesses for downstream numerical optimization.

For the first pair $\min_V d(Q_1, VP_1V^T)$, there are infinitely many exact solutions $V = Q_1^{1/2}RP_1^{-1/2}$ for $R \in \mathbb{O}_n$. For the second pair, one seeks solutions $\min_R d(M, RNR^T)$ for symmetric M, N, which is known as a *two-sided orthogonal Procrustes problem*[25] when d is the metric induced by the Frobenius norm. The optimal solutions R happen to still be $R = U_M\Theta U_N^T$ for Θ any signature matrix. Finally, for the third pair, one seeks optimal solutions $\min_\Theta d(Y, \Theta X\Theta)$, which can be determined by *quadratic unconstrained binary optimization* (QUBO)[26] for appropriate choices of d such as that induced by the L_1 vector-norm.

Specifically, to reframe $|Y - \Theta X\Theta|_1 = \sum_{i,j}|Y_{ij} - X_{ij}\Theta_i\Theta_j|$ as a QUBO problem, we first rearrange the error terms into Ising Hamiltonian terms:

$$|Y_{ij} - X_{ij}\Theta_i\Theta_j| = \frac{|Y_{ij} - X_{ij}| - |Y_{ij} + X_{ij}|}{2}\Theta_i\Theta_j + \frac{|Y_{ij} - X_{ij}| + |Y_{ij} + X_{ij}|}{2},$$

which can be collected into the symmetric pairwise interaction matrix J that defines the variable term of the Hamiltonian:

$$J_{ji} = J_{ij} = \frac{|Y_{ij} - X_{ij}| - |Y_{ij} + X_{ij}|}{2} + \frac{|Y_{ji} - X_{ji}| - |Y_{ji} + X_{ji}|}{2}.$$

Since Ising-like Hamiltonians are equivalent to general QUBO problems by a change-of-basis, we may then use existing mixed-integer programming packages to solve for optimal Θ^*'s that minimize $\Theta J\Theta$. Thus, from three pairs $\{P_i, Q_i\}$, one may compute a solution V' that is "stepwise optimal":

$$V' = Q_1^{1/2}U_M\Theta^*U_N^TP_1^{-1/2}.$$

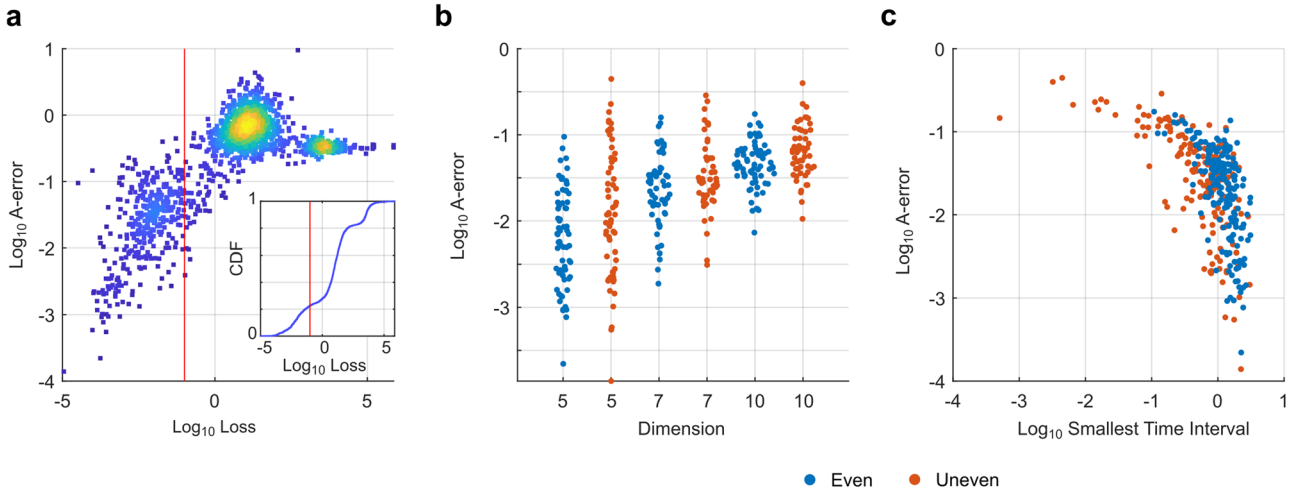

Fig. 1 | Numerical solutions of A for even and uneven time intervals. a Density scatter plot of loss L_{series} at convergence for gradient descent from random initializations, and the resulting average entry-wise error magnitude between ground-truth A and converged solutions A^*. Inset shows cumulative distribution function of loss values. Red line indicates loss cutoff for solutions assumed to have converged numerically to a global optimum. **b** Entry-wise error of solutions A^* to ground-truth for different dimensions, and time-series with either even (blue) or uneven (orange) time intervals. **c** Dependence of A^* error on the smallest simulated time interval in a time-series.

Remark 7 *One may also stop at two pairs of covariances, and determine the diagonal entries of the signature matrix Θ by comparing the signs of $\vec{v}$ and $\vec{w}$ and subsequently finding a solution for x_{ss} as described in Theorem 8.*

Numerical solutions of e^{tA} in the presence of noise and non-linearity

Whereas the previous section shows the possibility of solving for n-dimensional linear dynamics using distributions from only three time points, here we numerically demonstrate the empirical performance for determining A given various practical issues of real-world single-cell omics data. Specifically, we consider sampling noise, measurement noise, hidden variables, and non-linearities in the dynamics. Notably, when dealing with nonlinear dynamical systems $\dot{x} = F(x)$, our aim will be to approximately recover the Jacobian at some or all of the steady state points $\{x_i^*\}$ s.t. $\dot{x} = 0$, by the local approximation $F(x) = A(x - x_i^*) + O\big((x - x_i^*)^2\big)$. Practically, unstable steady state points will be harder, but not impossible, to observe than stable ones when one only has finitely many data points.

We first define a loss function for the purposes of numerical optimization. For a sequence of covariance matrices $\{C_i\}_{i=0,\ldots,k}$ and means $\{m_i\}_{i=0,\ldots,k}$ measured at different times t_i, assumed to originate from a common starting distribution with mean μ_0 and covariance Σ_0 at t_0, we define the loss-function $L_{\text{series}}(A, \Sigma_0, \mu_0, x_{ss})$:

$$L_{series}\left(A, \mu_0, \Sigma_0, x_{ss}\right) = \sum_{i=0}^{k} d^2\left(e^{t_i A}\Sigma_0 e^{t_i A^T}, C_i\right) + \alpha\delta^2\left(e^{t_i A}(\mu_0 - x_{ss}) + x_{ss}, m_i\right) \tag{22}$$

for some metric $d(P, Q)$ for $P, Q \in \mathbb{S}_n^+$, a metric $\delta(x, y)$ for $x, y \in \mathbb{R}^n$, and an unknown steady state x_{ss}. We also include a tuning parameter α to weigh the covariance-loss and mean-loss differently. In the case that data originates from different starting distributions with means $\{m_{0r}\}_{r=1,\ldots,j}$ and covariances $\{\Sigma_{0r}\}_{r=1,\ldots,j}$, one may sum the loss-functions of each time series:

$$L\big(A, \{m_{0r}\}, \{\Sigma_{0r}\}, x_{ss}\big) = \sum_{r=1}^{j} L_{\text{series}}\big(A, \mu_{0r}, \Sigma_{0r}, x_{ss}\big). \tag{23}$$

Depending on availability of data and measurement error, it may also be preferable to instead fix $\Sigma_0 = C_0$ and $\mu_0 = m_0$, and optimize L by only varying A and x_{ss}.

As for the specific choice of $d^2(P, Q)$, the Frobenius norm-induced metric:

$$d^2(P, Q) = \text{tr}\left((P - Q)(P - Q)^T\right)$$

is one possible choice. We instead considered a "Log"-metric tailored to covariance matrices[27]:

$$d^2(P, Q) = \text{tr}\left(\log^2(P^{-1/2}QP^{-1/2})\right). \tag{24}$$

In the following numerical simulations, we chose to only vary A and x_{ss} to optimize L_{series}, fixing $\Sigma_0 = C_0$, $\mu_0 = m_0$; this choice facilitates downstream analysis of our numerical results in relation to our theoretical results, but allows any errors in C_0 and m_0 to propagate to future time points during optimization. Thus, letting Σ_0 and μ_0 be free parameters in practical applications may yield slightly better results. Also, we made use of the Log-metric for d^2 and the Euclidean L^2 metric for δ^2, with a weight factor $\alpha = 0.1$.

Numerically fitting data with unequal time intervals

Before considering noise, we first examined whether exact solutions for A could be easily determined by minimizing L_{series} from the exact covariances Σ_i and means μ_i generated by random A's, for three time points with intervals bounded above by $\pi/|b_j|$. For numerical simulations in dimensions $n = 5, 7, 10$, we found that the majority (~80%) of randomly initialized gradient-descent optimizations of L_{series} did not converge to a solution A^* corresponding to the ground-truth A (Fig. 1a), based on both the value of the loss L_{series} and the actual average entry-wise error $|(A^* - A)_{ij}|$ to ground-truth A. We assumed that $L_{\text{series}} < 10^{-1}$ corresponded to convergence to the global minimum. From these results, we conclude that there are, as one might expect, many local minima in L_{series}, which can be circumnavigated by using the formulae for e^{tA} in Theorems 8 and 10.

Then, we also explored whether three time points was sufficient to solve for A when the time points t_i did not occur at equal intervals, i.e., $t_0 - t_1 \neq t_1 - t_2$, and are bounded above by $\pi/|b_j|$. For $n = 5, 7, 10$, we again generated random matrices A as ground-truth, setting $x_{ss} = 0$, and tried to solve for A by optimizing L_{series} using a simulated time-series. Keeping only optimization solutions for which $L_{\text{series}} < 0.1$, we compared average entry-wise error of the solutions A^* for both even and uneven intervals (Fig. 1b). Uneven time intervals were chosen uniformly at random between $[0,\ \pi/|b_{\max}|]$ for $|b_{\max}|$

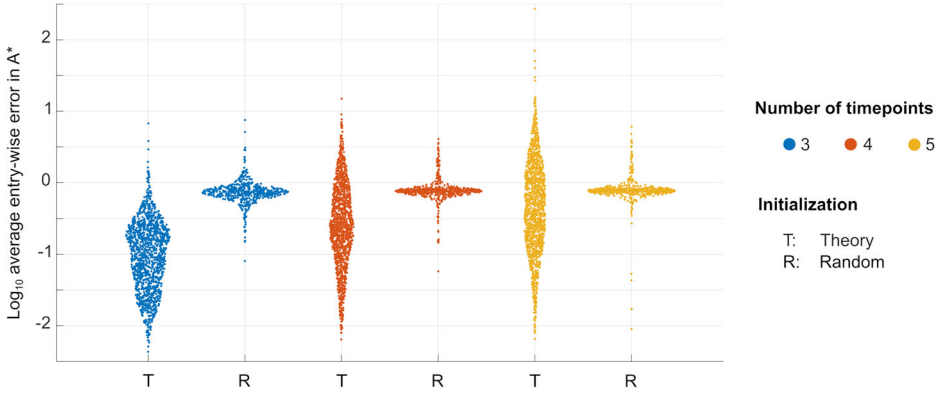

Fig. 2 | Estimation error of initialization strategies in the presence of sampling noise. Average entry-wise error of locally optimal solutions A^*, depending on whether initialization was chosen at random (R) or by theory (T). Results are sorted by number of time points. Simulations span dimensions $n = 5, 7, 10, 15, 20$ and sample sizes $N = 300, 1000, 3000$.

the magnitude of the largest imaginary component of a given A's eigenvalues. We then simulated a corresponding even-interval time series by using three time points spaced evenly between 0 and t_2 for the same A's, μ_0's, and Σ_0's. Since we have analytically shown that three time points are necessary and sufficient to guarantee uniqueness for equal intervals, the range of entry-wise errors from optimizing L_{series} for even intervals (blue) are representative of expected numerical error from optimization as opposed to the lack of a unique solution.

We found that entry-wise errors of A^* were comparable between optimizing with even intervals and uneven intervals. In the few instances for which uneven time intervals led to relatively large errors, at least one interval $t_{i+1} - t_i$ was often two or three orders of magnitude less than in the even case (Fig. 1c); this corresponds to the limit $(t_{i+1} - t_i) \to 0$ that can occur in the uneven case, in which three time points of data effectively approaches two time points worth, and there is no unique solution. Thus, these simulations suggest that three unevenly spaced time points may also be generically sufficient to solve for the linear dynamics A.

Recovering dynamics from data with sampling and measurement noise

Next, we evaluated how well one can recover a ground-truth A using noisy estimates $\{C_i\}$ and $\{m_i\}$ of covariance matrices and means under various practical conditions. First, to investigate the effects of sampling noise, we specified ground-truth matrices A and initial distributions (μ_0, Σ_0) as was done above, for dimensions $n = 5, 7, 10, 15, 20$, computed the ground-truth trajectory of distributions for 3, 4, or 5 time points evenly spaced between $[0, 0.9\pi/|b_{\max}|]$, and then generated $N = 300, 1000, 3000$ samples from each ground-truth distribution at each timepoint to compute estimated covariances and means $\{C_i\}$, $\{m_i\}$. Then, we numerically optimized for minimal-loss solutions A^*, and evaluated the average entry-wise error $|(A^* - A)_{ij}|$.

During numerical optimization, we asked if the theory-motivated solution would still be informative in a noisy setting, by using the approximate solutions described in a previous section to initialize optimization. We compared the error of converged solutions to those resulting from initializing randomly over a range of entries between $[-1.5, 1.5]$ to emulate the practical uncertainty in the range of parameters (i.e., entries) of A. The resulting average entry-wise error of any converged solution A^* to ground truth A are compared in Fig. 2 for these two initialization strategies, regardless of whether the loss L_{series} was low or high. We found that random initialization was seldom able to converge to the better optima often found by theory-based initialization, indicating that the approximate solutions described in a previous section can even be useful for efficiently finding "good fits" when given noisy data.

We then examined how much estimation error results from sampling noise, assuming that a globally optimal solution is found. To do so, for any given ground-truth A, we optimized multiple times using a mix of theory-based and random initialization strategies to improve the chances of finding a global optimum, keeping the lowest-loss solution as the final solution A^*. Examining only the subset of simulations in which the optimized loss $L_{\text{series}} < 0.3n$, we found that the entries of A could often be estimated with an average error often between 10^{-2} and 10^{-1} (Fig. 3), which was small compared to the original entries' range of 1. Errors generally increased with increasing dimension n, and decreased with sample size N and number of time points T, as one would expect, and for example, in the worst case of $n = 20$, $T = 3$, and $N = 300$, we saw that noise made it difficult to recover a good estimate A^* of A.

We also explored the effects of measurement noise on estimation error by defining a new loss function that adds a measurement error covariance matrix E to L_{series}:

$$L_{\text{noise}}(A, \mu_0, \Sigma_0, x_{ss}, E) = \sum_{i=0}^{k} d^2\left(e^{t_i A}\Sigma_0 e^{t_i A^T} + E, C_i\right) + \alpha\delta^2(e^{t_i A(\mu_0 - x_{ss})}, m_i).$$

Depending on the situation, E could also be simplified, e.g., to a diagonal matrix, or even to a matrix $\sigma^2 I$ for a single noise-parameter σ. We computed covariances and means as before, but adding Gaussian noise with covariance matrix $\sigma^2 I$ to each datapoint. Specifically, we chose σ values that resulted in the total variance at time 0 (i.e., $\text{tr}(\Sigma_0 + \sigma^2 I)$) being composed of ~10%, 26%, 55%, or 78% noise. In Fig. 4a, we show the estimation errors for $T = 3$ and $N = 3000$ when $L_{\text{noise}} < 0.5n$, minimizing L_{noise} while keeping σ as a fixed parameter corresponding to the simulated scenario. Even when measurement noise is modeled perfectly as in our simulations, we found that estimation error increases with increasing noise. In the case of $T = 3$, $N = 3000$, only $n = 5$ dynamics could be estimated with average entry-wise error on the order of ~0.1, and typically only when measurement noise was not the majority of the total variance. For $T = 4$ time points, higher-dimensional dynamics could again be partially recovered (Fig. 4b), although again only when measurement noise was low.

Overall, our simulations suggest that even with the practical challenges of sampling and measurement noise that one might encounter in single-cell omics data, one can still use three or more time points of estimated covariances and means to recover an underlying dynamical system A of dimension $n > 5$ with reasonable estimation error.

Remark 8 *Although sufficiently small time intervals $t_{\max}$ are necessary to theoretically guarantee that A can be recovered uniquely, it may sometimes be more favorable in the presence of noise to choose time intervals in which the data distributions have the largest change between time-points, i.e., the most "signal-to-noise". Even if this means choosing a larger interval $t_{\max}$, it is still possible to recover A's eigenvectors and eigenvalue's real-parts.*

Estimating stable and unstable subspaces of nonlinear dynamical systems

While the results thus far have been restricted to linear dynamical systems, we wondered if finding the best-fitting linear dynamical system for a nonlinear system $\dot{x} = F(x)$ could still provide useful information about $F(X)$.

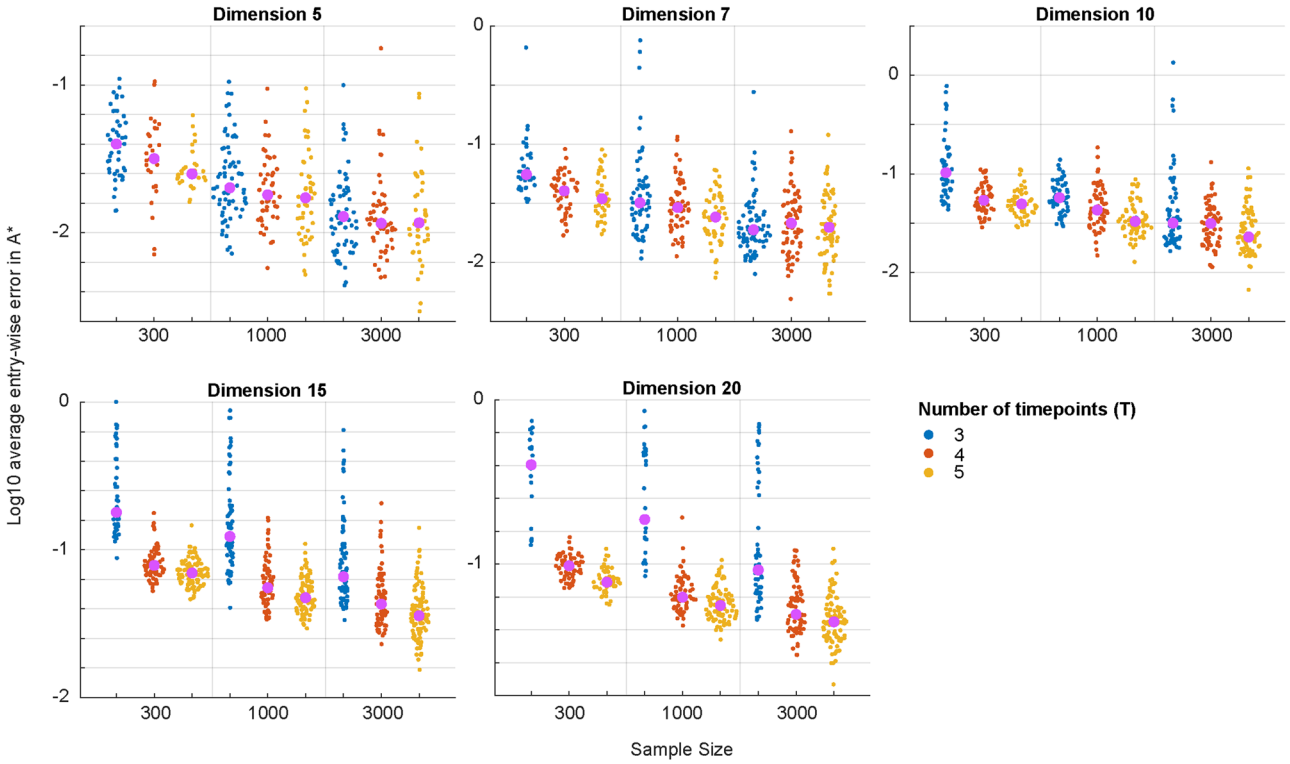

Fig. 3 | Estimation errors from sampling noise. Average entry-wise errors of best multi-initialization solutions A^* in relation to dimension n, sample size N, and number of time points T. Magenta dots show medians.

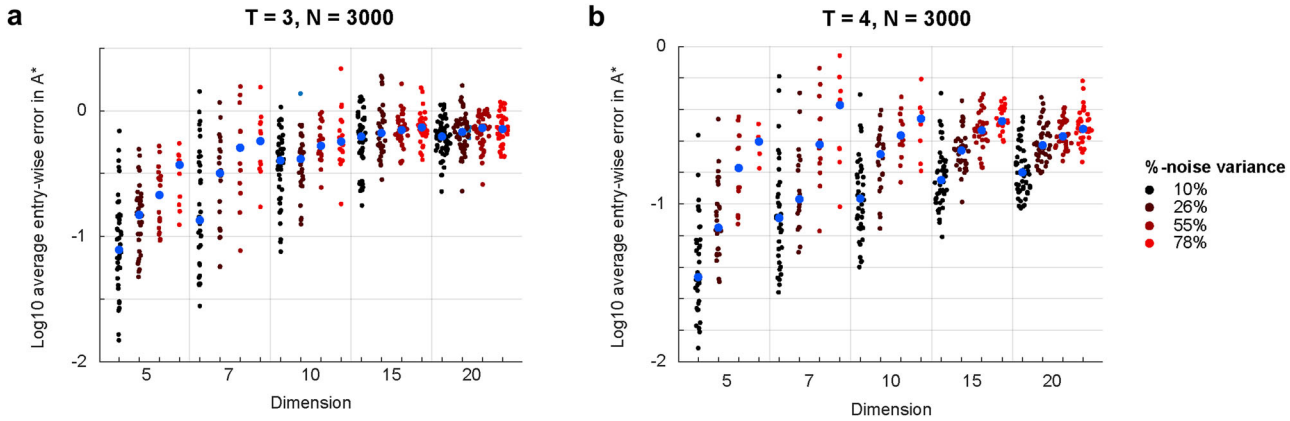

Fig. 4 | Estimation errors from measurement noise. a Average entry-wise errors of best multi-initialization solutions A^* for varying scales of measurement noise and dimension, for $T = 3$, $N = 3000$. **b** Same as (**a**) for $T = 4$. Blue dots show medians.

In particular, dynamical systems theory often analyzes the qualitative dynamics of a system $\dot{x} = F(x)$ by first finding all the steady states $\{x_i^*\}$ s.t. $F(x_i^*) = 0$, and then analyzing their respective associated linear dynamical systems via the Jacobians $J_i = D_x F(x)|_{x=x_i^*}$ to determine local stable or unstable subspaces, as a way to piece together key features of the phase portrait. We sought to infer similar properties of a dynamical system without knowing the equations for nonlinear $F(x)$ beforehand, but instead using time-varying data distributions generated by $F(x)$ and fitting linear dynamics $\dot{x} = A_i(x - x_i^*)$ using the covariance matrices Σ_t and means μ_t.

We evaluated how well the stable and unstable subspaces of J_i can be recovered from data around an isolated steady state x_i^* of randomly generated $F(x)$ by simulating trajectories as follows. First, we generated a trajectory of distributions for a linear dynamical system $\dot{x} = Ax$ as above. To simulate a nonlinear dynamical system, we then randomly generated an autodiffeomorphism $H : \mathbb{R}^n \to \mathbb{R}^n$ by composing three randomly generated autodiffeomorphisms $H = g \circ R \circ h$: both g and h were defined coordinate-wise using random univariate monotonic splines (note the Jacobians of g and h are thus diagonal), and R was a rotation chosen randomly from $\mathbb{O}_n$ uniformly on the Haar measure; we also ensured that $g(0) = h(0) = R(0) = 0$. Finally, we applied H to each x_{t_i} drawn from the distributions corresponding to trajectories of $\dot{x} = Ax$, thereby giving samples y_{t_i} of a nonlinear dynamical system $\dot{y} = F(y) = D_x H|_{y=H^{-1}(x)} A H^{-1}(y)$ with a unique steady state at 0. Examples of three different autodiffeomorphisms $H_i(x)$ for $n = 2$ are shown in Fig. 5a, alongside the image of a trajectory (red). We think of these nonlinear systems $\dot{y} = F(y)$ as representatives for nonlinear dynamics near a steady state x_{ss}.

For various choices of n, we then used the estimated covariances C_t and means m_t from the trajectories of $\dot{y} = F(y)$ to optimize L_{series}.

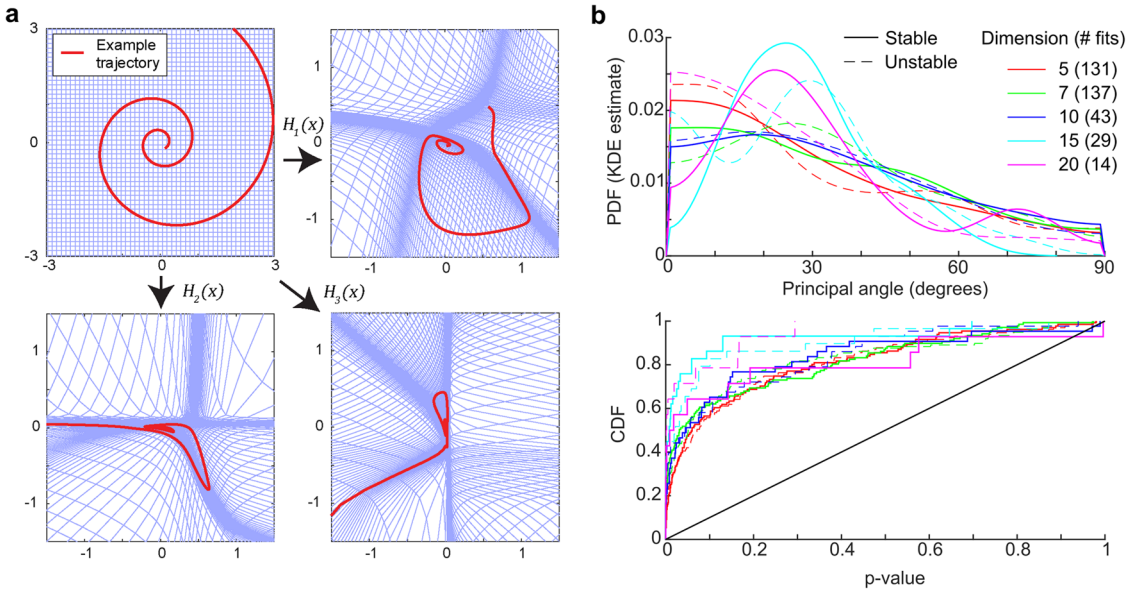

Fig. 5 | Estimating qualitative features of the Jacobian of nonlinear dynamical systems near steady state. a Three examples of diffeomorphisms $H_i(x) : \mathbb{R}^2 \to \mathbb{R}^2$ used to generate nonlinear systems from linear ones. The image of blue gridlines and a red oscillating trajectory are shown. **b** Principal angles and p-values between the stable or unstable subspaces of A^* and J, for various dimensions. Number of fits for which the principal angle was meaningful (i.e., neither A^* nor J had a stable/unstable subspace of full dimension n) are denoted in brackets.

We specifically chose $N = 3 \times 10^4$ and $T = 6$ so that we could eliminate errors due to noise, as our interest was in whether the best fit A^* could even be related to $D_yF(y)|_{y=0}$.

Given a linear fit A^*, we then compared the deviation of the stable and unstable subspaces S_{stab}, S_{unst} (i.e., the respective eigen-subspaces corresponding to eigenvalues with negative or positive real-part) of A^* and $J = D_yF(y)|_{y=0}$ by determining the *principal angle*[28] between them. The results are shown in Fig. 5b for cases where the principal angle is well-defined, i.e., both A^* and $D_yF(y)|_{y=0}$ had nontrivial subspaces. Corresponding p-values for these principal angles are also shown. Based on these results, the stable and unstable subspaces of A^* and J are significantly similar in at least 60−80% of cases. Furthermore, we did not exclude simulations based on the loss L_{series}, since it was unclear whether a large loss corresponded to poor initialization or to the systematic error of fitting a linear system to a nonlinear system, and so the insignificant ~ 30% of $F(y)$ contains at least some cases of poor initialization. Thus, fitting linear dynamics to single-cell omics data generated by a nonlinear dynamical system $F(y)$ is likely informative about stable and unstable subspaces of $F(y)$.

Interpretations of A in the presence of hidden variables

Practically, if a complex system is governed by $\dot{x} = F(x)$ for $x \in \mathbb{R}^n$, experiments often only observe a subset of the entries of x, i.e., data is often only available for $y \in \mathbb{R}^m$ for $m < n$ and $y = \mathcal{P}x$ for $\mathcal{P}$ a linear projection matrix onto the subset of m coordinates. This is likely the case even for a single-cell RNA sequencing experiment, which despite measuring expression from $n \sim 10^5$ different genes in a cell, does not measure the proteins that govern the coupled expression dynamics between genes. In such a case, the dynamics of the observed variables $y(t)$ generally cannot be described by an autonomous dynamical system $\dot{y} = G(y)$, and the methods presented herein for estimating linear dynamics A do not immediately give interpretable results. However, with additional assumptions, we may still infer properties of A.

Assuming that a linear system $\dot{x} = Ax$ has a slow-fast time-scale separation, i.e., the real parts of the eigenvalues $\lambda_k = a_k + ib_k$ have a gap s.t. $a_1 < \cdots < a_s \ll a_{s+1} < \cdots < a_n$, and furthermore that $a_{s+1}, \cdots a_n < 0$, then for times $t \gg a_{s+1}^{-1}$, trajectories $x(t)$ are arbitrarily close to the s-dimensional subspace L_s spanned by the eigenvectors $\{\vec{v}_i\}_{i=1,2,\ldots,s}$. Consequently, any dynamics $x(t) \in \mathbb{R}^n$ may be approximated by the "slow" dynamics of $x(t) \in L_s$, given by the "slow" eigenvalues $\{\lambda_i\}_{i=1,\ldots,s}$ and their corresponding eigenvectors.

If the projection $\mathcal{P}$ restricts to a diffeomorphism between $L_s \leftrightarrow \mathcal{P}(L_s) \in \mathbb{R}^m$, the dynamics of $y(t) := \mathcal{P}x(t) \in \mathcal{P}(L_s)$ are smoothly equivalent to those of $x(t) \in L_s$, and so $y(t) \in \mathcal{P}(L_s)$ may be described by a linear dynamical system $\dot{y} = By$ with the same eigenvalues $\{\lambda_i\}_{i=1,\ldots,s}$ as A. To ensure that $\mathcal{P}|_{L_s}$ is a diffeomorphism, we need for $L_s \cap \ker(\mathcal{P}) = 0$ so that the map is injective; this is generically true if $\dim(\mathcal{P}(L_s)) = s \le m$ since $\ker(\mathcal{P})$ is a $(n-m)$-dimensional subspace, and L_s can be any arbitrary s-dimensional subspace, and an $(n-m)$-dimensional subspace intersects a generic m-dimensional subspace only at 0. Then, surjectivity follows by the definition of a linear projection, and differentiability follows from linearity. Under these conditions, we may recover B using the observable data $y \in \mathbb{R}^m$, which recapitulates the slow eigenvalues $\lambda_{i\le s}$ of the full dynamical system A.

Furthermore, we can recover partial information about the eigenvectors $\vec{v}_{i\le s}$. For the slow dynamics on L_s, which we denote by $A|_{L_s}$, let its eigendecomposition be denoted by $A|_{L_s} = WEW^{-1}$ for W the matrix of A's eigenvectors $\vec{v}_i$, and E the diagonal matrix with the same eigenvalues $\lambda_{i\le s}$ as A, but setting $\lambda_{i>s} = 0$. Due to smooth equivalence via the linear diffeomorphism $\mathcal{P}_{L_s}$, the observable dynamics $\dot{y} = By$ relate to $\dot{x} = A|_{L_s}x$ as:

$$B = \mathcal{P}|_{L_s} A|_{L_s} \mathcal{P}|_{L_s}^{-1} = \mathcal{P}|_{L_s} WEW^{-1} \mathcal{P}|_{L_s}^{-1},$$

which implies that the s columns of $P|_{L_s} W$ corresponding to the eigenvalues $\lambda_{i\le s}$ form the eigenbasis of B. In many practical cases, $\mathcal{P}$ is simply a coordinate projection, and so $\mathcal{P}\vec{v}_i$ simply gives the observable coordinates of $\vec{v}_i$. Thus, from the relative weights of different observed variables y_i in the eigenvectors of B, we may recover the observed variables' relative weights in the slow eigenvectors $\vec{v}_i$ of the full dynamical system A.

A cartoon of these properties is shown in Fig. 6 for $m = 2$ and $n = 3$. In single-cell high-dimensional data, it is often assumed or justified from empirical analysis that the measured variables y lie on a s-dimensional manifold for $s \ll m < n$, often with s <10. In such cases, the relevant dynamics are effectively of dimension s as opposed to m, and fitting linear

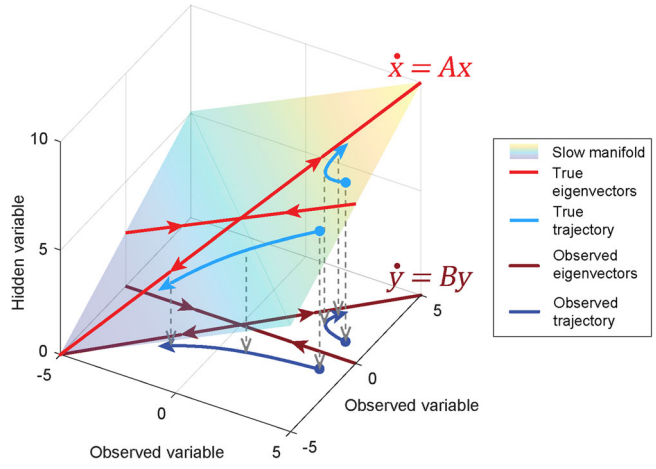

Fig. 6 | Linear dynamical properties preserved under projection. Toy example of a linear dynamical system $\dot{x} = Ax$ restricted to a plane in $\mathbb{R}^3$ representing a slow manifold, alongside the projection of trajectories (blue) from $\mathbb{R}^3$ onto the coordinate subspace $\mathbb{R}^2$ representing the subspace of experimentally observed variables. The observed dynamics in $\mathbb{R}^2$ are given by $\dot{y} = By$, and has eigenvectors (red) and eigenvalues related to $\dot{x} = Ax$.

dynamics to the data can be feasible even with current levels of experimental sampling or measurement noise.

Discussion

We have demonstrated the feasibility of inferring linear dynamics $\dot{x} = A(x - x_{ss})$ from single-cell omics time-series data, which can be used to infer the biomolecular dynamics near steady state x_{ss} that may represent cell states or differentiation fates. We showed that two pairs of data distributions $\{f_0(x), f_t(x)\}$ and $\{g_0(x), g_t(x)\}$ are generically needed to uniquely determine e^{tA} from the first- and second-order moments μ_t, Σ_t (in practice, estimated means and estimated covariance matrices), and that A could be uniquely identified so long as the time interval t between measurements is sufficiently short compared to the oscillatory components of solutions $x(t)$ to A. In other words, it theoretically requires only three sequential time points to determine the $n \times n$ matrix A, regardless of how large n is, i.e., how large of a biochemical interaction network one wishes to investigate.

We then demonstrated by simulation that A can be estimated accurately in the presence of sampling and measurement noise that is encountered in single-cell omics data, using only 3–5 time points, although noise and sample size become limiting factors past dimensions of n ~20. While increasing the number of time points, the sample size, and the signal-to-noise can help push the limit, an alternative strategy may simply be to perform dimensionality reduction as a "denoising" process, as is often done in single-cell data analysis. For example, single-cell omics time-series data may only portray significant changes over time in a k-dimensional subspace (e.g., the first k Principal Components), and so k-dimensional dynamics could be evaluated before translating back to the original n-dimensional coordinates. Also, the timing of sampling time points is crucial in the presence of noise, as one would ideally choose time points that capture the transient dynamics, before they are obscured by noise, along all n dimensions. This may not always be possible in biological systems for which multiple timescales exist simultaneously, and so a few time points would need to be sampled at each timescale one wishes to experimentally investigate. In future work, one might analytically explore the nonlinear optimization landscape that arises for noisy data, the statistics and confidence intervals of estimated dynamical properties, and better ways to choose multiple initializations in a theory-motivated fashion, so that well-designed experiments can efficiently fit higher-dimensional A to data and dynamics can be analyzed rigorously.

We also addressed how information about dynamical eigenvalues and eigenvectors can still be recovered even in the presence of "hidden variables", if one assumes timescale separation with a sufficiently low-dimensional slow subspace relative to the dimension of the observed data. Even when the dynamics $\dot{x} = F(x)$ are nonlinear, we find that fitting a model $\dot{x} = A(x - x^*)$ in the neighborhood of a steady state x^* can still allow for describing stable and unstable subspaces, akin to the traditional analytical approach of qualifying nonlinear dynamics via analysis of linearized systems near steady state. Alternatively, one might also consider the linear fit achieved here for a nonlinear dynamical system to be a coarse-grained model that "smooths out" non-linearities, and interpret the coarse-grained model without an explicit connection to the Jacobian of $F(x)$; such an approach has been fruitful in analyzing protein dynamics and may also be useful in single-cell omics analysis[29]. Either way, we hope that our proposed method can help to perform data-driven phase space analysis, without prior knowledge of the dynamic equations $\dot{x} = F(x)$.

In the setting of single-cell data analysis, our approach for inferring time-dependent behavior contrasts with many existing methods (e.g., "pseudo-time" trajectories[11]) in that we explicitly use experimental information on time, whereas other methods often use simplifying assumptions on the form of $F(x)$ in order to infer time-dependent behavior. Our results highlight how strong their assumptions must be: even a linear autonomous $F(x)$ requires at least three time points to determine $F(x) = A(x - x_{ss})$, by leveraging the rich information in the covariance matrices Σ between single cells; a general nonlinear $F(x)$ would require ≥3 time points to determine $F(x)$ in some "unique" sense. More precisely, from an algebraic geometric perspective, each pair of covariances offer $n(n + 1)/2$ quadratic equations that can be used to solve for the n^2 unknown entries that comprise the matrix e^{tA} in the linear case, and so it is not entirely surprising that two pairs of covariances, which offer $n^2 + n$ quadratic equations in total, could solve up to a finite set of 2^n solutions for e^{tA}, noting that the usual adage of needing as many equations as there are unknowns does not apply exactly for systems of nonlinear equations. A more general nonlinear $F(x)$ would involve additional unknown parameters, and therefore require additional time points to constrain those parameters. Similar themes can be found in statistical physics, in which infinitely many non-gradient dynamical systems can generate the same steady-state distribution[30], and so the assumptions of both gradient dynamics as well as steady-state is what allows for certain single-cell dynamics inference methods to determine dynamics using the data distribution of a single timepoint[29]. While our results show the sufficiency of a few time points for only linear $F(x)$, a careful algebraic accounting of parameters in other biologically relevant, nonlinear $F(x)$ may also reveal that $O(1)$ time points are sufficient to uniquely determine dynamics from data.

Our procedure for inferring properties of $F(x)$ from estimated means and covariances can be integrated into various workflows for inferring dynamics from data. Our analysis shows that the vast majority of experimental information on A comes from covariance data (i.e., single-cell heterogeneity). We did not explicitly consider what information can be further gleaned from higher-order moments of the data, although we note that for linear $F(x)$ and Normal distributions $f_0(x)$, there is no additional information beyond that provided by means and covariances. On a related note, it may also be of benefit to integrate single-cell omics data, which offers the advantage of high-dimensional measurement, with traditional time series data of only a few proteins/variables, which offers high time resolution, to better infer the dynamical system $F(x)$. For example, it may be easier to determine the threshold on the time intervals $t<\pi/|b_{\max}|$ using a time series experiment, or to estimate the neighborhoods in which $F(x)$ is well-approximated by linear dynamics. Furthermore, prior information about $F(X)$ might also be used in tandem with the analysis provided by our procedure: databases on gene or protein interactions[31] might be used to fix certain entries of A to 0 if there are no previously recorded correlations implying network interactions; other data analysis methods for inferring biochemical reaction network properties from single-cell data (e.g., stoichiometry of reactions[32]) could also constrain the form of A.

Finally, to incorporate our data-driven local linear analysis into a more global framework, one next step may be to consider how data analysis may be carried out on multistable dynamical systems. Since multi-modal

distributions $f_t(x)$ are a frequent occurrence in single-cell data, one might hope to derive relations between the effects of nonlinear dynamics $F(x)$ on $f_t(x)$ and data cluster analysis (which might also be regarded implicitly as analyzing higher-order moments). For example, mixture modeling of distributions $f_t(x)$ as linear combinations of components $f_t(x) = \sum_k p_k(x)$ may offer a natural starting point for modeling a dynamical system with multiple hyperbolic steady states and no other complications such as limit cycles or strange attractors. Local linear analysis might then be carried out on each mixture component $p_k(x)$. Ultimately, one might imagine developing methods to use the Liouville/transport equation[33] to solve for $F(x)$ using $f_t(x)$, but it is unclear how much data and time points would be needed to do so in well-defined or practical manner. In general, we hope that the rich information in single-cell omics data may be fully leveraged to understand the dynamical systems and biomolecular networks that govern cell states and fates.

Methods

Numerical optimization by gradient descent

We implemented the loss-functions L_{series} and L_{noise} in MATLAB, and solved numerically using *fminunc*() with the following options: MaxFunctionEvaluations=10^6, OptimalityTolerance=10^{-7}, MaxIterations=10^6, FiniteDifferenceStepSize=10^{-7}, StepTolerance=10^{-10}, and FiniteDifferenceType=central.

Generating ground-truth dynamical systems *A* and distributions

In all cases, ground-truth dynamical systems A were randomly generated by choosing the entries A_{ij} from a uniform distribution on $[-0.5, 0.5]$. For each such A, a corresponding time series was defined by choosing time points t_i evenly spaced between $[0,\ 0.9\pi/|b_{\max}|]$ for $b_{\max}$ the largest imaginary component amongst all the eigenvalues of A. In the case where uneven time intervals were tested, the t_i were chosen uniformly at random in the same range, and then re-ordered.

Then, ground-truth distributions were defined by just the parameters μ_i, Σ_i: first, we generated random μ_0 and Σ_0 by taking the estimated mean and covariance of $n+2$ random n-vectors with entries chosen uniformly at random from $[0, 10]$. Then, using A and the t_i's, we computed μ_i and Σ_i using $e^{t_i A}$. To generate noisy covariances C_i and means m_i, in the case of sampling noise we took the estimates for N samples from the normal distributions given by μ_i and Σ_i. In the case of measurement noise, we did the same but sampling instead from normal distributions given by μ_i and $\Sigma_i + \sigma^2 I$, for $\sigma^2 = 1, 3, 10, 30$. Since the diagonal entries of Σ_0 as generated here have an expectation value of ~ 8, the chosen values of σ^2 correspond to measurement noise percentages $\sigma^2/(8+\sigma^2)$ around ~ 10, 26, 55, and 78%.

Initialization strategies

To generate an initialization based on theory, we took three sequential covariances C_i (at random if there were more than three time points), and determined a guess $V = Q_1^{1/2} R P_1^{-1/2}$ for $P_1 = C_0$ and $Q_1 = C_1$. For R, we used the estimate $R = U_M \Theta U_N^T$ for U_M and U_N the orthogonal eigenvector bases of M, N as defined in Theorem 6, taking $P_2 = C_1$ and $Q_2 = C_2$. For Θ, we defined the signature matrix $\mathcal{D} = \text{sgn}(v_i/w_i)$ from Theorem 8 for $\vec{v}$, $\vec{w}$. We then determined an initialization x_{ss} as in Theorem 8, using the guess for V. An initialization for A was given by $\text{Log}(V)/\delta t$, for δt the first two time points. If $\text{Log}(V)$ was complex, we took only the real-parts of $\text{Log}(V)/\delta t$ as the initialization for A.

To allow for multiple initializations for a given A, additional initializations were generated by choosing the entries of A and x_{ss} uniformly at random from the ranges $[-0.5, 0.5]$ or $[-1.5, 1.5]$, or taking the theory-based initialization of A but the random guess for x_{ss}. Each of these three initializations was used once per given A.

Generating random homeomorphisms *H*(*x*)

To generate a random diffeomorphism $H(x) : \mathbb{R}^n \leftrightarrow \mathbb{R}^n$, we generated three separate diffeomorphisms g, R, h on $\mathbb{R}^n$ and composed them as $g \circ R \circ h$. R was an orthogonal matrix selected randomly from the uniform Haar measure, using the standard technique of applying QR decomposition to normally distributed variables[34]. Both g and h were defined coordinate-wise by a random monotonic spline, e.g., for the i'th coordinate, $g_i(x) = s(x_i)$ for the spline $s(x) : \mathbb{R} \to \mathbb{R}$.

The splines $s(x)$ were randomly generated by first defining a spline $p(x)$ at the points $p(-3) = y_1$, $p(0) = y_2$, and $p(3) = y_3$ for y_i's chosen from the uniform distribution on the interval $[-3, 3]$ and sorted in ascending order, and $p(x) = x$ for $x = -6, -5, -4, 4, 5, 6$. Using these nine points we fit the piecewise cubic hermite interpolating polynomial to define $p(x)$, using the *pchip*() function in MATLAB. Then, we defined $s(x) = p(x) - p(0)$, so that $s(0) = 0$, and thus $h(0) = g(0) = 0$ and $H(0) = 0$.

To ensure that the distribution of trajectories was mostly contained in the nonlinear regions of the $s(x)$'s domain, we generated the parameters μ_0 and Σ_0 as before, but choosing entries for the n-vectors used to generate Σ_0 uniformly in $[0, 3]$, and entries for μ_0 in $[0, 1]$.

During numerical optimization of L_{series} for the nonlinear simulations, we used 5 initializations with entries randomly chosen from $[-0.5, 0.5]$ per ground-truth system $\dot{y} = F(y)$.

Null distribution of subspace principal angles

To compute the p-value for the principal angle between two subspaces (of respective dimensions k and l contained in $\mathbb{R}^n$), we generated 100 pairs of k-dimensional and l-dimensional subspaces and computed the principal angle between each pair to create a simulated null distribution. Each subspace was generated at random by first selecting an orthogonal matrix $U \in \mathbb{O}_n$ with uniform probability on the Haar measure, and then selecting the first k or l columns of U to span a subspace.

Code availability

All original code has been deposited at GitHub (https://github.com/shuwang543/covar_lin_dyn) and is publicly available as of the date of publication (https://doi.org/10.5281/zenodo.12604544).

Received: 10 May 2024; Accepted: 10 August 2024;
Published online: 27 August 2024

References

1. Nawy, T. Single-cell sequencing. *Nat. Methods* **11**, 18–18 (2014).
2. Lin, J.-R., Fallahi-Sichani, M. & Sorger, P. K. Highly multiplexed imaging of single cells using a high-throughput cyclic immunofluorescence method. *Nat. Commun.* https://doi.org/10.1038/ncomms9390 (2015).
3. Spitzer, M. H. & Nolan, G. P. Mass cytometry: Single cells, many features. *Cell* **165**, 780–791 (2016).
4. O'Donnell, E. A., Ernst, D. N. & Hingorani, R. Multiparameter flow cytometry: advances in high resolution analysis. *Immune Netw.* **13**, 43 (2013).
5. Sible, J. C. & Tyson, J. J. Mathematical modeling as a tool for investigating cell cycle control networks. *Methods* **41**, 238–247 (2007).
6. Fröhlich, F. et al. Efficient parameter estimation enables the prediction of drug response using a mechanistic pan-cancer pathway model. *Cell Syst.* **7**, 567–579.e6 (2018).
7. Huang, S. Non-genetic heterogeneity of cells in development: more than just noise. *Development* **136**, 3853–3862 (2009).
8. van Mourik, S. et al. Continuous-time modeling of cell fate determination in arabidopsis flowers. *BMC Syst. Biol.* **4**, 1–13 (2010).
9. Kapfer, E.-M., Stapor, P. & Hasenauer, J. Challenges in the calibration of large-scale ordinary differential equation models. *IFAC-PapersOnLine* **52**, 58–64 (2019).
10. Schmid, P. J. Dynamic mode decomposition and its variants. *Annu. Rev. Fluid Mech.* **54**, 225–254 (2022).
11. Saelens, W., Cannoodt, R., Todorov, H. & Saeys, Y. A comparison of single-cell trajectory inference methods. *Nat. Biotechnol.* **37**, 547–554 (2019).

12. Gorin, G., Fang, M., Chari, T. & Pachter, L. Rna velocity unraveled. *PLoS Comput. Biol.* **18**, e1010492 (2022).
13. Schiebinger, G. Reconstructing developmental landscapes and trajectories from single-cell data. *Curr. Opin. Syst. Biol.* **27**, 100351 (2021).
14. Loos, C. & Hasenauer, J. Mathematical modeling of variability in intracellular signaling. *Curr. Opin. Syst. Biol.* **16**, 17–24 (2019).
15. Schiebinger, G. et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell* **176**, 928–943.e22 (2019).
16. Yeo, G. H. T., Saksena, S. D. & Gifford, D. K. Generative modeling of single-cell time series with prescient enables prediction of cell trajectories with interventions. *Nat. Commun.* **12**, 3222 (2021).
17. Tran, T. N. & Bader, G. D. Tempora: cell trajectory inference using time-series single-cell RNA sequencing data. *PLoS Comput. Biol.* **16**, e1008205 (2020).
18. Qiu, X. et al. Mapping transcriptomic vector fields of single cells. *Cell* **185**, 690–711.e45 (2022).
19. Sontag, E. D. For differential equations with r parameters, 2r+1 experiments are enough for identification. *J. Nonlinear Sci.* **12**, 553–583 (2003).
20. Browning, A. P., Warne, D. J., Burrage, K., Baker, R. E. & Simpson, M. J. Identifiability analysis for stochastic differential equation models in systems biology. *J. R. Soc. Interface* **17**, 20200652 (2020).
21. Culver, W. J. On the existence and uniqueness of the real logarithm of a matrix. *Proc. Am. Math. Soc.* **17**, 1146–1151 (1966).
22. Shannon, C. E. A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 623–656 (1948).
23. Albers, C., Critchley, F. & Gower, J. Applications of quadratic minimisation problems in statistics. *J. Multivar. Anal.* **102**, 714–722 (2011).
24. Meng, F. et al. Procrustes: a Python library to find transformations that maximize the similarity between matrices. *Comput. Phys. Commun.* **276**, 108334 (2022).
25. Schönemann, P. H. On two-sided orthogonal Procrustes problems. *Psychometrika* **33**, 19–33 (1968).
26. Kochenberger, G. et al. The unconstrained binary quadratic programming problem: a survey. *J. Comb. Optim.* **28**, 58–81 (2014).
27. Förstner, W. & Moonen, B. Geodesy-The Challenge of the 3rd Millennium. In *A Metric for Covariance Matrices* (eds Grafarend, E. W. et al.) 299–309 (Springer Berlin Heidelberg, 2003).
28. Absil, P.-A., Edelman, A. & Koev, P. On the largest principal angle between random subspaces. *Linear Algebra Appl.* **414**, 288–294 (2006).
29. Xing, J. Reconstructing data-driven governing equations for cell phenotypic transitions: integration of data science and systems biology. *Phys. Biol.* **19**, 061001 (2022).
30. Xing, J. Mapping between dissipative and Hamiltonian systems. *J. Phys. A: Math. Theor.* **43**, 375003 (2010).
31. Cline, M. S. et al. Integration of biological networks and gene expression data using Cytoscape. *Nat. Protoc.* **2**, 2366–2382 (2007).
32. Wang, S., Lin, J.-R., Sontag, E. D. & Sorger, P. K. Inferring reaction network structure from single-cell, multiplex data, using toric systems theory. *PLoS Comput. Biol.* **15**, e1007311 (2019).
33. Villani, C. *Optimal Transport* (Springer Berlin Heidelberg. Grundlehren der mathematischen Wissenschaften, 2009).
34. Mezzadri, F. How to generate random matrices from the classical compact groups. *Not. Am. Math. Soc.* **54**, 592–604 (2007).

Acknowledgements

S.W. and D.A.L. were supported by NIH IMPAcTB 75N93019C00071 and NIH HIPC grant U19-AI167899. M.A.A and E.D.S. were supported in part by grants AFOSR FA9550-21-1-0289, AFOSR FA9550-22-1-0316, and NSF DMS2052455.

Author contributions

S.W., D.A.L., and E.D.S. conceptualized and supervised this study. S.W., M.A.A., and E.D.S. performed formal analysis and investigation. S.W. wrote all the original code used in this study. D.A.L. and E.D.S. acquired funding. All authors contributed to writing the original draft and revision of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to Douglas A. Lauffenburger or Eduardo D. Sontag.