

FlowMap: Geometry–Dynamics Consistent Embedding of RNA Velocity for Interpretable Cellular Trajectories

Jingyuan Hu¹, Luca Pinello², Rong Ma¹

1 Interpretation of RNA Velocity

RNA velocity provides a quantitative description of how a cell’s molecular profile changes over time, inferred from the coupling between unspliced (U) and spliced (S) mRNA abundances. The kinetic model for a single gene is:

$$\begin{aligned}\frac{dU(t)}{dt} &= \alpha - \beta U(t), \\ \frac{dS(t)}{dt} &= \beta U(t) - \gamma S(t),\end{aligned}\tag{1.1}$$

where α , β , and γ are gene-specific kinetic parameters.

For a given cell, the RNA velocity vector

$$\frac{d\mathbf{S}}{dt} = \left(\frac{dS_1}{dt}, \frac{dS_2}{dt}, \dots, \frac{dS_m}{dt} \right)^\top \in \mathbb{R}^m$$

represents the instantaneous change in expression across all p genes.

Manifold view and role of the Jacobian. Intuitively, a cell’s state can be described by a small set of intrinsic coordinates $\mathbf{y} \in \mathbb{R}^d$ that capture its biological identity and position along developmental or functional trajectories. Even though gene expression lives in a high-dimensional space \mathbb{R}^m , not every combination of expression levels is biologically valid — cells occupy only a smooth, low-dimensional surface (manifold) within this space. The mapping $\psi : \mathbb{R}^d \rightarrow \mathbb{R}^p$ takes a point in this intrinsic *state space* to its corresponding gene expression profile.

The Jacobian $\mathbf{J}_\psi(\mathbf{y})$ tells us how small changes in state coordinates translate into changes in gene expression. It defines the local *tangent space* — the set of all gene expression changes that are consistent with a valid movement along the manifold. When we take an observed RNA velocity vector $\frac{d\mathbf{S}}{dt}$ and project it onto this tangent space, we filter out components that would move the cell off the manifold (i.e., biologically implausible directions) and keep only those that reflect genuine state changes. This projection step both denoises the velocity vector and constrains it to represent an evolution of the cell’s underlying state.

Chain rule connection. By the multivariate chain rule, RNA velocity (change in expression per unit time) factors into:

$$\frac{d\mathbf{S}}{dt} = \frac{\partial \mathbf{S}}{\partial \mathbf{y}} \cdot \frac{d\mathbf{y}}{dt} = \mathbf{J}_\psi(\mathbf{y}) \frac{d\mathbf{y}}{dt},$$

where $\mathbf{J}_\psi(\mathbf{y}) \in \mathbb{R}^{p \times d}$ is the Jacobian of ψ , with entries $\partial S_i / \partial y_j$.

Expanded in matrix form:

$$\underbrace{\begin{bmatrix} \frac{dS_1}{dt} \\ \frac{dS_2}{dt} \\ \vdots \\ \frac{dS_m}{dt} \end{bmatrix}}_{\substack{\text{change in expression} \\ \text{change in time}}} = \underbrace{\begin{bmatrix} \frac{\partial S_1}{\partial y_1} & \dots & \frac{\partial S_1}{\partial y_d} \\ \frac{\partial S_2}{\partial y_1} & \dots & \frac{\partial S_2}{\partial y_d} \\ \vdots & \ddots & \vdots \\ \frac{\partial S_m}{\partial y_1} & \dots & \frac{\partial S_m}{\partial y_d} \end{bmatrix}}_{\substack{\text{change in expression} \\ \text{change in cell state}}} \underbrace{\begin{bmatrix} \frac{dy_1}{dt} \\ \frac{dy_2}{dt} \\ \vdots \\ \frac{dy_d}{dt} \end{bmatrix}}_{\substack{\text{change in cell state} \\ \text{change in time}}}.$$

Recovering state-space velocity. Conceptually, we want to answer the question: *what change in cell state would best explain the observed change in gene expression?* The Jacobian $\mathbf{J}_\psi(\mathbf{y})$ acts as a bridge between the two spaces: it tells us how a small movement in state space translates into changes in gene expression. By using this bridge in reverse, we can map the observed gene expression change back into a change in state coordinates, giving us the intrinsic velocity $\frac{d\mathbf{y}}{dt}$ that best explains the data.

Formally, we solve:

$$\widehat{\frac{d\mathbf{y}}{dt}} = \arg \min_{\mathbf{v} \in \mathbb{R}^d} \left\| \mathbf{J}_\psi(\mathbf{y}) \mathbf{v} - \frac{d\mathbf{S}}{dt} \right\|_2^2,$$

whose closed-form solution uses the Moore–Penrose pseudoinverse:

$$\widehat{\frac{d\mathbf{y}}{dt}} = \mathbf{J}_\psi(\mathbf{y})^+ \frac{d\mathbf{S}}{dt}.$$

In a 2-D embedding, $\frac{d\mathbf{y}}{dt}$ gives the best local linear approximation of how the cell’s state is evolving, yielding a denoised and interpretable trajectory in the intrinsic state space.

2 FlowMap embedding method

2.1 Overview of the FlowMap framework

The core rationale behind FlowMap is to model the gene expression space as a *differentiable manifold*, where RNA velocity defines a continuous vector field over this manifold. By learning a smooth, parametric mapping between the low-dimensional embedding and the high-dimensional expression space, FlowMap preserves the manifold’s geometric structure and enables direct, quantitative analysis of cellular dynamics.

Common embedding methods such as t-SNE, UMAP, or PHATE are powerful for preserving local or global geometry, yet they do not retain an explicit mapping to the original space. Using these embeddings as a starting point, FlowMap reconstructs them as a differentiable manifold, providing a major advantage: the ability to map between embedding and gene space, recover tangent spaces, and quantify feature contributions in a consistent framework. This differentiable structure expands the interpretability of the embedding beyond visualization while preserving the strengths of the original method, including:

- **Parametric mapping:** reconstructing the manifold explicitly via a spline function, yielding a smooth, analytical representation of the data and enabling direct mapping between embedding coordinates and gene expression space.

- **Interpretability:** providing a transparent alternative to graph-based or heuristic embeddings, with well-defined derivatives that support velocity projection and gene-level attribution.
- **Geometric understanding:** enabling the analysis of local and global structure, including curvature, fixed points, and branching topology.

Guided by these goals, FlowMap introduces several key components:

- Incorporating directional information through a velocity-aware distance metric that captures both spatial proximity and local flow.
- Reconstructing the manifold with a spline mapping, enabling smooth back-projection and tangent space computation.
- Refining embeddings via joint optimization of positions and vector alignment to minimize reconstruction error and velocity mismatch.
- Analyzing manifold geometry to classify fixed-point dynamics (e.g., sources, sinks, saddles) and identify genes aligned with inferred trajectories.

FlowMap takes as input a gene expression matrix $X \in \mathbb{R}^{N \times p}$ and a corresponding RNA velocity matrix $V \in \mathbb{R}^{N \times p}$ for N cells and p genes. The goal is to construct a low-dimensional embedding that is geometrically consistent with the RNA velocity-defined dynamics.

The procedure consists of the following steps:

1. **Velocity-aware distance construction.** For each pair of cells, we compute a phase distance that quantifies their relative position along the RNA velocity flow (Section 2.2). This phase distance is combined with Euclidean distance in expression space using a weighting parameter α , yielding a velocity-aware distance matrix.
2. **Low-dimensional embedding.** A neighborhood graph is constructed from the combined distance matrix, and a low-dimensional embedding is obtained using a standard manifold learning algorithm such as UMAP or t-SNE. This produces an initial representation that reflects both expression similarity and directional progression.
3. **Smooth manifold reconstruction.** To obtain an explicit and differentiable mapping between the embedding and gene expression space, we fit a spline that maps embedding coordinates to the original expression space (or a reduced PCA representation; Section 2.3). This reconstruction defines a smooth parametric manifold underlying the embedding.
4. **Tangent-space velocity projection.** Using the spline mapping, we compute the tangent space of the reconstructed manifold at each embedded point (Section 2.4). RNA velocity vectors are projected onto these tangent spaces, ensuring that inferred transcriptional changes are consistent with the geometry of the embedded state space.
5. **Optional joint refinement.** FlowMap optionally supports a joint optimization procedure that adjusts embedding coordinates and projected velocities by minimizing expression reconstruction error and velocity mismatch (Section 2.5). This step is not required for geometric consistency and is disabled by default.

Together, these steps produce a low-dimensional representation in which cellular states and RNA velocity are embedded in a geometrically coherent manner, with velocity vectors constrained to lie along the reconstructed manifold.

2.2 Velocity-aware distance metric

In the context of a continuous vector field, the *phase* of a point can be understood as its position along a flow trajectory, analogous to a timestamp measured in the intrinsic coordinates of the field. Points with the same phase lie on level sets orthogonal to the direction of flow, and moving along the flow direction changes phase while preserving other coordinates on the manifold. For RNA velocity data, phase provides a natural way to quantify temporal progression between cells, separating differences aligned with the primary direction of flow from those arising in directions orthogonal to it.

To compute the phase distance between two nearby cells, FlowMap estimates their relative position along the local velocity flow by allowing each cell to move along its RNA velocity vector and measuring how much progression is required for the two states to align. Intuitively, if two cells lie at different phases along the same trajectory, they can be brought closer in expression space by sliding them forward or backward along the flow; the amount of sliding reflects their separation along the dynamical direction. Cells that are synchronized along the flow have similar optimal offsets and thus small phase distance, whereas cells that are ahead of or behind one another along a trajectory require different displacements to align and therefore have large phase distance. A formal definition of the resulting phase distance, along with its interpretation as a symmetric first-order time difference, is provided in Supplementary Notes 1.1–1.2.

Formally, let $x_i, x_j \in \mathbb{R}^p$ denote the gene expression profiles of cells i and j , with corresponding RNA velocity vectors $v_i, v_j \in \mathbb{R}^p$. We define their optimal alignment by solving

$$(\tau_1^*, \tau_2^*) = \arg \min_{(\tau_1, \tau_2)} \|x_i + \tau_1 v_i - x_j - \tau_2 v_j\|^2, \quad (2.1)$$

which finds the pair of displacements along the velocity directions that minimizes the distance between the two states after perturbation. The scalars τ_1 and τ_2 can be interpreted as local progression times along the flow required to best align the two cells.

This optimization admits a closed-form solution obtained by solving a 2×2 linear system for each pair of cells. Details of the derivation, numerical implementation, and stability considerations are provided in Supplementary Note 1.3.

To capture both spatial proximity and this notion of dynamical phase, FlowMap computes the *pairwise phase distance* between cells using both their gene expression profiles and local velocity vectors (Supplementary Note 1). This phase distance is then combined with the standard Euclidean distance to define a *velocity-aware distance*:

$$D_{\text{combined}}^2(i, j) = D_{\text{Euclidean}}^2(i, j) + \alpha D_{\text{phase}}^2(i, j), \quad (2.2)$$

where $D_{\text{Euclidean}}(i, j)$ is the standard Euclidean distance between cells i and j in expression space, $D_{\text{phase}}(i, j)$ is their estimated phase distance, and α controls the relative weighting between spatial and dynamical separation.

The resulting metric is small for cells that are both close in expression space and synchronized in their progression along the flow, and large for cells that are separated along the trajectory. By preserving this combined distance in the embedding, FlowMap aligns the low-dimensional geometry with the underlying cellular dynamics.

Starting from the velocity-aware distance matrix, FlowMap constructs a neighborhood graph by connecting each cell to its k nearest neighbors (KNN) based on the combined metric. This graph serves as the input to manifold learning algorithms such as UMAP or t-SNE, enabling embeddings that jointly preserve spatial similarity and temporal progression derived from the vector field.

2.3 Manifold reconstruction

A key limitation of common embedding methods such as t-SNE or UMAP is the absence of an explicit, differentiable mapping between the low-dimensional embedding and the original high-dimensional gene expression space. These methods typically rely on cell–cell transition probabilities or neighborhood graphs, which are effective for visualization but limit interpretability: without a mapping function, it is not possible to directly recover local tangent spaces, quantify feature contributions, or project inferred trajectories back into gene space.

FlowMap addresses this limitation by reconstructing the embedding as a smooth *differentiable manifold* using a spline-based function. This approach preserves the geometric structure of the original embedding while acting as an effective denoising step: by enforcing smoothness, the spline filters out high-frequency measurement noise and retains only the dominant, coherent patterns of variation that drive the cellular dynamics. In this way, the learned manifold provides a low-dimensional approximation of the high-dimensional data that captures the main dynamical trends, enabling both interpretable geometry and reduced noise sensitivity. By endowing the embedding with an explicit, parametric mapping, the reconstructed surface supports rigorous geometric analysis, enabling downstream computations such as curvature, geodesics, and local basis vectors.

Formally, let $\mathbf{X} \in \mathbb{R}^{N \times p}$ denote the matrix of gene expression profiles for N cells and p genes, and $\mathbf{Y} \in \mathbb{R}^{N \times d}$ their corresponding low-dimensional coordinates ($d \ll p$, typically $d = 2$). We seek a smooth mapping $\psi : \mathbb{R}^d \rightarrow \mathbb{R}^p$ such that

$$\psi(\mathbf{y}_i) \approx \mathbf{x}_i, \quad i = 1, \dots, N,$$

where \mathbf{x}_i and \mathbf{y}_i are the expression profile and embedding coordinates of cell i , respectively.

We model ψ as a vector-valued spline in a reproducing kernel Hilbert space (RKHS),

$$\psi(\mathbf{y}) = \sum_{j=1}^N \mathbf{c}_j K(\mathbf{y}, \mathbf{y}_j) + A\mathbf{y} + b,$$

where $K(\cdot, \cdot)$ is a positive-definite kernel, $\{\mathbf{c}_j\}$ are coefficient vectors, and $A \in \mathbb{R}^{p \times d}$, $b \in \mathbb{R}^p$ define the affine component.

The coefficients are obtained by solving the regularized regression problem

$$\hat{\psi} = \arg \min_{\psi \in \mathcal{H}} \sum_{i=1}^N \|\psi(\mathbf{y}_i) - \mathbf{x}_i\|^2 + \lambda \|\psi\|_{\mathcal{H}}^2,$$

where \mathcal{H} denotes the RKHS associated with K , and $\lambda > 0$ controls the smoothness–accuracy trade-off. Details of the linear system solution and numerical implementation are provided in Supplementary Note 2.1-2.2.

A common choice for K is the class of *polyharmonic spline* kernels [1, 2], defined as radial basis functions of the form

$$K(\mathbf{y}, \mathbf{y}') = \phi(\|\mathbf{y} - \mathbf{y}'\|),$$

where

$$\mathbf{y}, \mathbf{y}' \in \mathbb{R}^d$$

are embedding coordinates and d denotes the embedding dimension. The radial function

$$\phi : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$$

is given by

$$\phi(r) = \begin{cases} r^{2m-d} \log r, & \text{if } 2m - d \text{ is even,} \\ r^{2m-d}, & \text{if } 2m - d \text{ is odd,} \end{cases}$$

where m denotes the spline order and satisfies

$$m > \frac{d}{2}.$$

Higher values of m produce smoother interpolating functions by penalizing higher-order derivatives of the reconstructed manifold.

Thin-plate spline as a special case. When $d = 2$ and the kernel is chosen as the biharmonic radial basis function

$$K(\mathbf{y}, \mathbf{y}') = \phi(\|\mathbf{y} - \mathbf{y}'\|), \quad \text{with } \phi(r) = r^2 \log r,$$

the above formulation reduces to the classical thin-plate spline (TPS), where the RKHS norm corresponds to a bending energy penalty.

Smoothing parameter selection. The parameter λ controls the effective complexity of the reconstructed surface by trading off fidelity to the data and smoothness of the mapping. Smaller values of λ allow the spline to closely interpolate the observed data, while larger values enforce a smoother, lower-variance surface. In practice, λ is selected using generalized cross-validation (GCV), which provides a data-driven estimate of the optimal bias–variance trade-off (Supplementary Note 2.3).

2.4 Tangent space and velocity projection

Given the smooth mapping

$$\psi : \mathbb{R}^d \rightarrow \mathbb{R}^p$$

from embedding to gene expression space, the *tangent space* at an embedded point \mathbf{y} is the d -dimensional subspace of \mathbb{R}^p spanned by the columns of the Jacobian

$$\mathbf{J}_\psi(\mathbf{y}) \in \mathbb{R}^{p \times d} :$$

$$T_{\mathbf{y}}\mathcal{M} = \text{span} \left\{ \frac{\partial \psi}{\partial y_1}(\mathbf{y}), \dots, \frac{\partial \psi}{\partial y_d}(\mathbf{y}) \right\}.$$

The Jacobian $\mathbf{J}_\psi(\mathbf{y})$ provides a local linear approximation of the manifold, mapping embedding coordinates to gene expression space. A closed-form expression for $\mathbf{J}_\psi(\mathbf{y})$ under the spline model is given in Supplementary Note 2.4.

Given a high-dimensional velocity vector

$$\mathbf{v}_i \in \mathbb{R}^p,$$

we estimate its low-dimensional representation by projecting it onto the local tangent space. This is formulated as the least-squares problem

$$\mathbf{u}_i = \arg \min_{\mathbf{u} \in \mathbb{R}^d} \|\mathbf{J}_\psi(\mathbf{y}_i)\mathbf{u} - \mathbf{v}_i\|^2,$$

which seeks the embedding-space vector whose pushforward best matches the observed velocity.

The solution is given by

$$\mathbf{u}_i = \mathbf{J}_\psi^+(\mathbf{y}_i)\mathbf{v}_i,$$

where $\mathbf{J}_\psi^+(\mathbf{y}_i)$ denotes the Moore–Penrose pseudoinverse. This yields a low-dimensional velocity field

$$\{\mathbf{u}_i\} \subset \mathbb{R}^d$$

defined on the embedding.

To obtain a smooth and interpretable vector field, we fit a **second spline function**

$$\nu : \mathbb{R}^d \rightarrow \mathbb{R}^d$$

to the projected velocities $\{\mathbf{u}_i\}$. Unlike the manifold reconstruction ψ , which maps embedding coordinates to gene expression space, ν operates entirely within the embedding and models the velocity field directly in low-dimensional coordinates. The resulting function defines a continuous vector field in the embedding space that is consistent with the geometry induced by the reconstructed manifold.

2.5 Embedding optimization algorithm

The tangent space formulation naturally suggests that the embedding can be refined to better align with the underlying dynamics. While standard embedding methods (e.g., UMAP, t-SNE) are constructed solely from positional relationships, they do not explicitly enforce consistency with RNA velocity. As a result, the embedding may distort dynamical structure even when local neighborhoods are preserved.

FlowMap addresses this by refining an initial embedding using both positional and velocity information. Let ψ denote the spline mapping from embedding coordinates to gene expression space, and ν denote the low-dimensional velocity field. We optimize the embedding coordinates $\{y_i\}$ by minimizing the joint objective

$$\min_{\psi, \nu, \{y_i\}} (L_\psi + \alpha L_\nu + \lambda(\|\psi\|_{\mathcal{H}}^2 + \|\nu\|_{\mathcal{H}}^2)),$$

where

$$L_\psi = \sum_i \|x_i - \psi(y_i)\|^2,$$

and

$$L_\nu = \sum_i \|v_i - J_\psi(y_i)\nu(y_i)\|^2.$$

Here, L_ψ enforces reconstruction of the observed gene expression profiles, while L_ν enforces consistency between the observed high-dimensional velocities v_i and their projection through the Jacobian $J_\psi(y_i)$. The smoothness terms

$$\|\psi\|_{\mathcal{H}}^2 \quad \text{and} \quad \|\nu\|_{\mathcal{H}}^2$$

are RKHS norms associated with the spline model, controlling the regularity of both the manifold and the velocity field.

By jointly minimizing these terms, the embedding is iteratively adjusted to reduce mismatches between observed expression and inferred dynamics, yielding a representation that is both geometrically faithful and dynamically consistent.

Algorithm 1: Iterative Embedding Refinement

Algorithm 1 Iterative refinement of embedding and velocity field

- 1: **Input:** Expression data $\{x_i\}$, velocities $\{v_i\}$
- 2: **Initialize:** $\{y_i\}$ using a standard embedding (e.g., UMAP, t-SNE)
- 3: **repeat**
- 4: Fit spline mappings ψ and ν
- 5: Compute gradients:

$$\begin{aligned}\frac{\partial L_\psi}{\partial y_i} &= -2(x_i - \psi(y_i))J_\psi(y_i), \\ \frac{\partial L_\nu}{\partial y_i} &= -2(v_i - J_\psi(y_i)\nu(y_i))\left(H_\psi(y_i)\nu(y_i) + J_\psi(y_i)J_\nu(y_i)\right),\end{aligned}$$

- 6: Update $y_i \leftarrow y_i - \eta \nabla_{y_i} L$
 - 7: **until** convergence
 - 8: **Output:** Refined embedding $\{y_i^*\}$, mappings ψ and ν
-

This procedure refines the initial embedding by repositioning cells in regions where velocity information is inconsistent with the current geometry, improving alignment between cellular states and their inferred dynamics.

3 Interpreting differentiable velocity fields with FlowMap

3.1 Gene-level consistency of the reconstructed manifold and velocity field

To assess how well the reconstructed manifold captures the underlying cellular dynamics, we evaluate consistency between observed and reconstructed quantities for both expression and RNA velocity. Although this analysis is typically performed at the gene level, the same framework can also be applied to principal components or other feature representations.

Let i index cells and j index features (e.g., genes or principal components). Given embedding coordinates y_i , the reconstructed expression is

$$\hat{x}_i = \psi(y_i),$$

and the reconstructed velocity is

$$\hat{v}_i = J_\psi(y_i)\nu(y_i),$$

where ψ denotes the fitted manifold spline and ν the fitted velocity spline in embedding coordinates.

For each feature j , we quantify consistency using both the coefficient of determination (R^2) and Pearson correlation (r) between observed and reconstructed values across cells.

For expression:

$$\begin{aligned}R_{\text{expr},j}^2 &= 1 - \frac{\sum_i (x_{ij} - \hat{x}_{ij})^2}{\sum_i (x_{ij} - \bar{x}_j)^2}, \\ r_{\text{expr},j} &= \text{CORR}(x_{ij}, \hat{x}_{ij}),\end{aligned}$$

where

$$\bar{x}_j = \frac{1}{N} \sum_i x_{ij}.$$

Because the manifold reconstruction is fit independently for each feature through a regression model with an affine component,

$$R_{\text{expr},j}^2$$

is typically non-negative and approaches 1 when the reconstructed manifold accurately captures the feature variation.

For velocity:

$$R_{\text{vel},j}^2 = 1 - \frac{\sum_i (v_{ij} - \hat{v}_{ij})^2}{\sum_i (v_{ij} - \bar{v}_j)^2},$$

$$r_{\text{vel},j} = \text{corr}(v_{ij}, \hat{v}_{ij}),$$

where

$$\bar{v}_j = \frac{1}{N} \sum_i v_{ij}.$$

As described previously, the velocity reconstruction is obtained by first projecting observed velocities onto the local tangent space to obtain embedding-space velocities u_i , fitting the smooth velocity spline ν , and then mapping the smoothed velocities back to expression space through the manifold Jacobian:

$$\hat{v}_i = J_\psi(y_i)\nu(y_i).$$

Because reconstructed velocities are constrained to lie in the tangent space of the manifold,

$$R_{\text{vel},j}^2$$

may become negative when the projected velocity explains less variance than the mean-centered baseline.

Genes or features with high consistency in both expression and velocity are interpreted as dynamically aligned with the reconstructed manifold, whereas features with low consistency are less compatible with the inferred vector field structure. Because velocity reconstruction depends on both the manifold geometry and the projected vector field, velocity consistency provides a stricter measure of dynamical coherence than expression reconstruction alone.

These quantities provide a feature-level diagnostic for identifying components that contribute most strongly to the learned dynamical geometry.

3.2 Gene-expression gradients and velocity alignment

Because FlowMap reconstructs a smooth mapping between the embedding space and the ambient gene-expression space, local gene-expression gradients can be computed directly from the manifold Jacobian.

Let

$$x = \psi(y),$$

where

$$x = (x_1, \dots, x_p)$$

represents gene-expression coordinates and

$$y = (y_1, \dots, y_d)$$

denotes the embedding coordinates.

The Jacobian of the manifold mapping is

$$J_{\psi(y)} = \frac{\partial \psi}{\partial y}(y) = \begin{bmatrix} \frac{\partial \psi_1}{\partial y_1} & \dots & \frac{\partial \psi_1}{\partial y_d} \\ \vdots & \ddots & \vdots \\ \frac{\partial \psi_p}{\partial y_1} & \dots & \frac{\partial \psi_p}{\partial y_d} \end{bmatrix}.$$

Each row of the Jacobian defines the local gradient of a gene-expression coordinate with respect to the embedding:

$$\nabla_y \psi_k(y) = \left(\frac{\partial \psi_k}{\partial y_1}, \dots, \frac{\partial \psi_k}{\partial y_d} \right).$$

Because the embedding inherits a Riemannian metric from the reconstructed manifold, gene gradients are compared to the velocity field using the pullback metric

$$g(y) = J_{\psi(y)}^\top J_{\psi(y)}.$$

The corresponding Riemannian gradient is

$$\nabla_g \psi_k(y) = g(y)^{-1} \nabla_y \psi_k(y).$$

Given the fitted velocity spline $\nu(y)$, we quantify alignment between the local gene gradient and the inferred flow using the metric-aware cosine similarity

$$A_k(y) = \frac{\nabla_g \psi_k(y)^\top g(y) \nu(y)}{\sqrt{\nabla_g \psi_k(y)^\top g(y) \nabla_g \psi_k(y)} \sqrt{\nu(y)^\top g(y) \nu(y)}}.$$

Positive alignment indicates that the gene increases along the direction of cellular progression, whereas negative alignment indicates decreasing expression along the flow. Values near zero correspond to genes whose local gradients are approximately orthogonal to the velocity field and therefore weakly associated with the local dynamical trajectory.

3.3 Vector field geometry

RNA velocity as a flow on a Riemannian manifold. We model cellular dynamics as a continuous flow defined on a smooth, low-dimensional state manifold \mathcal{M} . Let \mathcal{M} be a d -dimensional manifold equipped with a Riemannian metric g , which defines inner products on the tangent space $T_x \mathcal{M}$ at each point $x \in \mathcal{M}$.

RNA velocity induces a vector field

$$v(x) \in T_x \mathcal{M},$$

and the corresponding dynamics on the manifold satisfy

$$\frac{dx(t)}{dt} = v(x(t)).$$

FlowMap represents the dynamics in embedding coordinates $y \in \mathbb{R}^d$ using the fitted velocity spline

$$\nu(y) \in \mathbb{R}^d,$$

which defines the embedding-space dynamics

$$\frac{dy(t)}{dt} = \nu(y(t)).$$

The manifold geometry is determined locally by the Jacobian of the reconstruction mapping ψ , which induces the pullback metric

$$g(y) = J_{\psi(y)}^\top J_{\psi(y)}.$$

This metric describes how distances, angles, and local flow directions in the embedding correspond to the geometry of the reconstructed expression manifold. Consequently, even locally linear velocity fields in embedding coordinates may correspond to curved trajectories in expression space. Exact analysis of the dynamics therefore depends on both the velocity field and the local manifold geometry, making global solutions difficult to obtain in closed form.

To analyze the dynamics near fixed points, we therefore locally flatten the manifold geometry and approximate the flow using a locally Euclidean linear system.

Local flattening and linearization of the dynamics. To analyze the flow near a fixed point, we linearize the dynamics in locally Euclidean coordinates. The Riemannian metric $g(y^*)$ induces a local inner product on $T_{x^*}\mathcal{M}$. Let

$$g(y^*) = LL^\top$$

denote the Cholesky decomposition of the metric tensor. The transformation

$$z = L^\top(y - y^*)$$

defines locally flattened coordinates in which the metric becomes the identity matrix. In these coordinates, the flow equation can be approximated by its first-order Taylor expansion,

$$\frac{dz}{dt} \approx Az, \quad A = L^\top J_\nu(y^*)L^{-\top},$$

where $J_\nu(y^*)$ is the Jacobian of the embedding velocity field evaluated at the fixed point. This procedure locally normalizes the metric structure, reducing the problem to a standard Euclidean linear system and yielding a stable representation of the local dynamics.

Classification of fixed points via linearized dynamics. The qualitative behavior of the flow near x^* is determined by the eigenvalues of the linear operator A . In particular, the signs of the real parts characterize local stability, while complex eigenvalues indicate rotational or oscillatory behavior. The main classes of fixed points are summarized in Table 1.

Table 1: Classification of fixed points from the eigenvalues of the linearized operator A .

Eigenvalue condition	Classification	Interpretation
All eigenvalues real with $\text{Re}(\lambda_i) < 0$	Stable node (sink)	Stable terminal or attractor state
All eigenvalues real with $\text{Re}(\lambda_i) > 0$	Unstable node (source)	Divergent progenitor or repulsive state
Real eigenvalues with mixed signs in $\text{Re}(\lambda_i)$	Saddle	Branching or commitment state
Complex eigenvalues with $\text{Re}(\lambda_i) < 0$	Spiral in	Damped oscillatory convergence
Complex eigenvalues with $\text{Re}(\lambda_i) > 0$	Spiral out	Expanding oscillatory dynamics
Complex eigenvalues with $\text{Re}(\lambda_i) \approx 0$	Center	Approximately conservative or cyclic dynamics
At least one eigenvalue with $ \text{Re}(\lambda_i) \approx 0$	Degenerate	Marginal or weakly identifiable dynamics

Because this classification is invariant under smooth local coordinate transformations, it provides an intrinsic and robust characterization of local cellular dynamics on the manifold.

Practical implementation via local Jacobian regression. In principle, both the Riemannian metric $g(x)$ and the Jacobian of the coordinate vector field $J_u(y)$ can be computed in closed form from the fitted spline representations of the expression manifold and projected velocity field. These analytic derivatives provide exact local geometric quantities and are well defined everywhere on the reconstructed manifold. In practice, however, direct use of closed-form Jacobians leads to unstable estimates near fixed points. Because the vector field is inferred from noisy, sparsely sampled data, higher-order derivatives amplify local interpolation artifacts, resulting in irregular Jacobians that overfit small-scale fluctuations rather than capturing biologically meaningful dynamics.

To obtain robust estimates of local flow geometry, we therefore adopt a data-driven local linearization strategy. For each candidate fixed point y^* , we identify a neighborhood $\mathcal{N}(y^*)$ consisting of its k nearest neighbors in the embedding space. Within this neighborhood, the vector field is approximated by a first-order linear model,

$$\nu(y_i) \approx J_\nu(y^*)(y_i - y^*), \quad y_i \in \mathcal{N}(y^*),$$

with the constraint that $u(y^*) \approx 0$. We estimate the Jacobian $J_u(y^*)$ by solving a weighted least-squares regression problem,

$$\min_J \sum_{y_i \in \mathcal{N}(y^*)} w_i \|\nu(y_i) - J(y_i - y^*)\|^2,$$

where the weights w_i decrease smoothly with distance from y^* . This local regression effectively smooths the Jacobian estimate, averages over measurement noise, and avoids over-interpreting high-frequency variation introduced by spline interpolation.

The resulting locally averaged Jacobian is then combined with the metric tensor at y^* to perform the metric-aware linearization described above. This approach yields stable, interpretable estimates of fixed-point geometry while preserving the intrinsic structure of the underlying manifold and vector field.

3.4 Least-action path formulation for trajectory inference

To infer differentiation trajectories connecting fixed points to terminal cell states, we adopt a least-action path (LAP) framework motivated by Dynamo [3]. Let

$$\frac{d\mathbf{x}(t)}{dt} = \mathbf{v}(\mathbf{x}(t))$$

denote the ordinary differential equation (ODE) associated with the reconstructed RNA velocity vector field, where $\mathbf{x}(t) \in \mathbb{R}^d$ represents the cellular state at time t , and $\mathbf{v}(\mathbf{x})$ specifies the local direction and magnitude of state change. Solutions of this ODE generate flow curves (trajectories or paths) that describe continuous progression through the cellular state space.

In practice, observed cellular trajectories may deviate from ideal flow curves because of noise, stochasticity, sparse sampling, or imperfections in the estimated vector field. To identify trajectories that remain maximally consistent with the inferred dynamics, we define an action functional over candidate paths:

$$\mathcal{A}[\mathbf{x}] = \int_0^T \|\dot{\mathbf{x}}(t) - \mathbf{v}(\mathbf{x}(t))\|^2 dt, \tag{3.1}$$

where $\dot{\mathbf{x}}(t)$ denotes the instantaneous tangent vector of the trajectory. This functional measures the discrepancy between the trajectory direction and the local velocity field along the path. Minimizing the action therefore favors trajectories that closely follow the vector field while remaining globally smooth and dynamically coherent.

Given a source state \mathbf{x}_s and a target state \mathbf{x}_t , corresponding to fixed points or terminal cellular states, the least-action path is defined as the solution of the variational problem

$$\mathbf{x}^*(t) = \arg \min_{\mathbf{x}(t)} \mathcal{A}[\mathbf{x}], \quad \text{subject to } \mathbf{x}(0) = \mathbf{x}_s, \mathbf{x}(T) = \mathbf{x}_t. \tag{3.2}$$

Numerically, the trajectory is discretized into a finite set of control points, and the action is minimized using iterative gradient-based optimization. The resulting path approximates the most dynamically consistent transition trajectory between the two cellular states under the reconstructed vector field.

In FlowMap, least-action trajectories are computed directly within the low-dimensional geometry-preserving embedding. The velocity field entering the action functional is the smooth tangent-space-constrained vector field reconstructed from the FlowMap manifold. Consequently, inferred trajectories remain consistent with both the manifold geometry and the denoised continuous dynamics represented by the embedding.

3.5 Curvature analysis

Decomposition of the acceleration vector for RNA-velocity Let

$$x(t) \in \mathcal{M}$$

denote a trajectory on the cellular state manifold evolving according to the RNA velocity field

$$\frac{dx(t)}{dt} = v(x(t)),$$

where

$$v(x(t)) \in T_{x(t)}\mathcal{M}.$$

Differentiating along the trajectory gives the acceleration vector

$$a(t) = \frac{d}{dt}v(x(t)).$$

Let

$$e_1(t) = \frac{v(x(t))}{\|v(x(t))\|_g}$$

denote the unit flow direction, and let

$$\{e_1(t), e_2(t), \dots, e_d(t)\}$$

be an orthonormal basis for the tangent space

$$T_{x(t)}\mathcal{M}.$$

Because the manifold is embedded in the ambient gene-expression space \mathbb{R}^p , the tangent space admits an orthogonal complement given by the normal space. Let

$$\{n_1(t), \dots, n_{p-d}(t)\}$$

be an orthonormal basis for the normal space of the embedded manifold at $x(t)$.

The acceleration vector admits the orthogonal decomposition

$$a(t) = a_{\text{flow}}(t) + a_{\text{steer}}(t) + a_{\text{surface}}(t),$$

where

$$a_{\text{flow}}(t) = g(a(t), e_1(t)) e_1(t)$$

is the component parallel to the local flow direction,

$$a_{\text{steer}}(t) = \sum_{i=2}^d g(a(t), e_i(t)) e_i(t)$$

is the component tangent to the manifold but orthogonal to the flow direction, and

$$a_{\text{surface}}(t) = \sum_{\alpha=1}^{p-d} \langle a(t), n_\alpha(t) \rangle n_\alpha(t)$$

is the component orthogonal to the embedded manifold in the ambient expression space.

The flow component a_{flow} measures local acceleration or deceleration along the trajectory. The steering component a_{steer} captures directional turning of the flow within the manifold and reflects intrinsic curvature of the cellular dynamics. The surface component a_{surface} measures deviation of the dynamics away from the locally embedded manifold surface in the ambient expression space.

Because these components are mutually orthogonal,

$$\|a(t)\|_g^2 = \|a_{\text{flow}}(t)\|_g^2 + \|a_{\text{steer}}(t)\|_g^2 + \|a_{\text{surface}}(t)\|_g^2.$$

Interpretation of the acceleration components. Each term in the decomposition of the acceleration vector captures a distinct aspect of the local dynamics of an embedded flow curve.

The flow component $a_{\text{flow}}(t)$ lies along the instantaneous velocity direction and quantifies changes in the *speed* of progression along the trajectory. A nonzero $a_{\text{flow}}(t)$ indicates local acceleration or deceleration along the same path without altering the direction of motion. Consequently, this component reflects temporal rescaling of the dynamics rather than geometric bending of the trajectory in state space.

The steering component $a_{\text{steer}}(t)$ lies within the tangent space of the manifold but is orthogonal to the flow direction. It captures changes in direction that occur *along the manifold itself*, corresponding to steering or turning within the intrinsic geometry of the learned surface. When $a_{\text{steer}}(t) = 0$, the trajectory follows a geodesic of the manifold, representing the straightest possible path constrained to the surface. Nonzero values of $a_{\text{steer}}(t)$ therefore reflect intrinsic curvature of the flow induced by the underlying biological dynamics.

The surface component $a_{\text{surface}}(t)$ is orthogonal to the manifold and measures how the trajectory bends relative to the ambient gene-expression space in which the manifold is embedded. This term captures extrinsic curvature arising from the geometry of the reconstructed manifold itself rather than from directional changes within the surface. Intuitively, $a_{\text{surface}}(t)$ reflects how the trajectory responds to the “terrain” of the manifold, analogous to motion over a curved landscape embedded in a higher-dimensional space.

In practice, low-dimensional RNA-velocity visualizations primarily reveal the tangent components of the dynamics, namely the flow and steering terms. The flow component describes local acceleration or deceleration along trajectories, while the steering component appears visually as turning or redirection of the velocity field within the embedding. By contrast, the surface component cannot generally be identified directly from the embedding visualization alone. Because the embedding represents a curved manifold projected into a low-dimensional coordinate system, deviations orthogonal to the manifold are not visually separable from the curvature of the surface itself. Recovering the surface component therefore requires an explicit differentiable reconstruction of the embedding manifold together with its tangent and normal spaces, which in FlowMap is provided by the spline reconstruction.

Together, these components disentangle speed modulation, intrinsic steering, and extrinsic bending of cellular trajectories, providing a geometrically interpretable decomposition of local dynamics in the FlowMap embedding.

Computing acceleration components Let

$$x(t) \in \mathcal{M}$$

denote a trajectory on the cellular state manifold, and let

$$y(t) \in \mathbb{R}^d$$

denote its embedding coordinates. FlowMap represents the manifold using a smooth mapping

$$\psi : \mathbb{R}^d \rightarrow \mathbb{R}^p, \quad x(t) = \psi(y(t)),$$

where p is the dimension of the ambient gene-expression space.

The dynamics on the manifold satisfy

$$\frac{dx(t)}{dt} = v(x(t)),$$

while the projected dynamics in embedding coordinates satisfy

$$\frac{dy(t)}{dt} = u(y(t)).$$

The velocity in expression space is therefore

$$v(x(t)) = J_{\psi(y(t))} u(y(t)),$$

where

$$J_{\psi(y(t))} = \frac{\partial \psi}{\partial y}(y(t))$$

is the Jacobian of the manifold mapping evaluated at $y(t)$.

The ambient acceleration vector is

$$a(t) = \frac{d}{dt}v(x(t)) = \frac{d}{dt} [J_{\psi(y(t))}u(y(t))].$$

Equivalently, for each expression coordinate $k = 1, \dots, p$,

$$a_k(t) = \frac{d}{dt} \left[\sum_{i=1}^d \frac{\partial \psi_k}{\partial y_i}(y(t)) u_i(y(t)) \right].$$

Applying the product rule gives

$$a_k(t) = \sum_{i,j=1}^d \frac{\partial^2 \psi_k}{\partial y_i \partial y_j}(y(t)) u_i(y(t)) u_j(y(t)) + \sum_{i,j=1}^d \frac{\partial \psi_k}{\partial y_i}(y(t)) \frac{\partial u_i}{\partial y_j}(y(t)) u_j(y(t)).$$

The first term captures acceleration induced by local curvature of the reconstructed expression manifold, while the second term captures local variation of the embedded velocity field.

To decompose the acceleration vector, we first project it onto the tangent space of the embedded manifold:

$$a_{\Gamma}(t) = J_{\psi(y(t))} \left[J_{\psi(y(t))}^{\top} J_{\psi(y(t))} \right]^{-1} J_{\psi(y(t))}^{\top} a(t).$$

The surface component is defined as the residual orthogonal to the tangent space,

$$a_{\text{surface}}(t) = a(t) - a_{\Gamma}(t).$$

Let

$$T(t) = \frac{v(x(t))}{\|v(x(t))\|}$$

denote the unit flow direction in expression space. The flow component is

$$a_{\text{flow}}(t) = \langle a_{\Gamma}(t), T(t) \rangle T(t),$$

and the steering component is

$$a_{\text{steer}}(t) = a_{\Gamma}(t) - a_{\text{flow}}(t).$$

Thus, the acceleration decomposes as

$$a(t) = a_{\text{flow}}(t) + a_{\text{steer}}(t) + a_{\text{surface}}(t).$$

Curvature and acceleration. Curvature quantifies how rapidly a trajectory changes direction along the flow, after removing changes in speed. The acceleration of an embedded flow curve decomposes into (motivated in Supplementary Note 3.1)

$$a(t) = a_{\text{flow}}(t) + a_{\text{steer}}(t) + a_{\text{surface}}(t),$$

where a_{flow} is parallel to the velocity, a_{steer} is tangent to the manifold but orthogonal to the velocity, and a_{surface} is normal to the reconstructed manifold.

Let

$$\|v(x(t))\|$$

denote the magnitude of the velocity vector in expression space. Since curvature measures directional change rather than acceleration along the same direction, the flow component is not included in the curvature. The steering curvature is defined as

$$\kappa_{\text{steer}}(t) = \frac{\|a_{\text{steer}}(t)\|}{\|v(x(t))\|^2},$$

which measures turning of the trajectory within the manifold. The surface curvature is defined as

$$\kappa_{\text{surface}}(t) = \frac{\|a_{\text{surface}}(t)\|}{\|v(x(t))\|^2},$$

which measures bending of the trajectory induced by the embedding of the reconstructed manifold in ambient expression space.

Because a_{steer} and a_{surface} are orthogonal, the total curvature of the ambient trajectory is

$$\kappa_{\text{total}}(t) = \sqrt{\kappa_{\text{steer}}(t)^2 + \kappa_{\text{surface}}(t)^2} = \frac{\sqrt{\|a_{\text{steer}}(t)\|^2 + \|a_{\text{surface}}(t)\|^2}}{\|v(x(t))\|^2}.$$

Thus, curvature separates intrinsic directional steering of the cellular flow from extrinsic bending of the reconstructed expression manifold. The flow component a_{flow} instead measures speed modulation along the same trajectory and is therefore reported separately from curvature. Additional geometric motivation and the relationship to intrinsic and extrinsic surface curvature are provided in Supplementary Notes 3.1–3.2.

Normalization by the squared velocity magnitude removes dependence on traversal speed: faster motion along the same geometric path produces larger acceleration but not larger curvature. Consequently, curvature reflects how sharply the trajectory bends in expression space rather than the absolute magnitude of transcriptional change.

References

- [1] Jean Duchon. Splines minimizing rotation-invariant semi-norms in sobolev spaces. In *Constructive theory of functions of several variables: proceedings of a conference held at oberwolfach April 25–May 1, 1976*, pages 85–100. Springer, 2006.
- [2] Grace Wahba. *Spline models for observational data*. SIAM, 1990.
- [3] Xiaojie Qiu, Yan Zhang, Jorge D Martin-Rufino, Chen Weng, Shayan Hosseinzadeh, Dian Yang, Angela N Pogson, Marco Y Hein, Kyung Hoi Joseph Min, Li Wang, et al. Mapping transcriptomic vector fields of single cells. *Cell*, 185(4):690–711, 2022.