A Deep Learning Framework for Estimating Cell-specific Kinetic Rates of RNA Velocity

Haotian Cui [*1 2 3] **Hassaan Maan** [*2 3 4] **Michael D. Taylor** [4 5 6 7] **Bo Wang** [1 2 3 5]

Abstract

Existing RNA velocity estimation methods rely on strong assumptions of predefined dynamics and cell-agnostic constant transcriptional rates, which are often violated in complex and heterogeneous single-cell RNA sequencing (scRNA-seq) data. To overcome these limitations, we present *DeepVelo*, a novel method that estimates *cell-specific* dynamics of transcriptional kinetics using a Graph Convolution Neural Network (GCN) model and generalizes RNA velocity to cell populations containing time-dependent kinetics and multiple lineages. We applied DeepVelo to complex developmental datasets, including dentate gyrus and hindbrain neurogenesis, and demonstrated its ability to disentangle multifaceted kinetics. DeepVelo infers time-varying cellular rates of transcription, splicing and degradation, recovers each cell's stage in the underlying differentiation process, and detects putative driver genes regulating these processes. By relaxing the constraints of previous techniques, DeepVelo facilitates the study of more complex differentiation and lineage decision events in heterogeneous scRNA-seq data. The DeepVelo package is available at https://github.com/bowang-lab/DeepVelo.

1. Introduction

The concept of RNA velocity refers to the rate of change of mRNA abundance in a cell, which reflects varying levels of RNA processing and degradation. Current velocity estimation methods leverage the abundance and ratio between unspliced pre-messenger RNAs and spliced mature messenger RNAs to infer changes in gene expression dynamics. Since unspliced mRNAs can be distinguished in common single-cell RNA sequencing (scRNA-seq) protocols (La Manno et al., 2018), estimating dynamic RNA velocity using only static sequencing libraries is feasible.

The original RNA velocity approach (La Manno et al., 2018) utilized the assumption that the observed transcriptional phases in scRNA-seq last long enough to reach both an apex of induction and a quiescent steady-state equilibrium. A more recent approach, scVelo (Bergen et al., 2020), attempted to generalize the *steady-state* assumption by replacing it with *four transcriptional states* and a dynamical model. Both of these existing techniques assume each gene follows a pre-defined trajectory depicted by constant, cell-agnostic kinetic rates of RNA dynamics. This workflow implies that each gene goes through a shared velocity trajectory across all celltypes, and limits the application in complex multi-lineage systems (Bergen et al., 2021; Gorin et al., 2022).

To resolve these limitations, we propose *DeepVelo*, a deep neural network based method that models RNA velocity without pre-defined kinetic patterns. Empowered by deep Graph Convolutional Networks (GCN), DeepVelo infers **gene-specific and cell-specific** RNA transcription, splicing, and degradation rates. Compared with existing techniques that use cell-agnostic parameters (La Manno et al., 2018; Bergen et al., 2020), DeepVelo is able to model RNA velocity for dynamics of high complexity - particularly for cell populations with highly heterogeneous celltypes and multiple lineages. We demonstrate the efficacy of DeepVelo on multiple developmental scRNA-seq datasets and find that DeepVelo yields more consistent velocity estimates and accurately identifies transcriptional states compared to existing models. DeepVelo exceeds the capacity of existing models in realistic single-cell datasets with multiple trajectories/lineages.

[*]Equal contribution [1]Department of Computer Science, University of Toronto, Toronto, ON Canada [2]Vector Institute, Toronto, ON, Canada [3]Peter Munk Cardiac Centre, University Health Network, Toronto, ON, Canada [4]Department of Medical Biophysics, University of Toronto, Toronto, ON, Canada [5]Department of Laboratory Medicine and Pathobiology, University of Toronto, Toronto, ON, Canada [6]Departments of Surgery, University of Toronto, Toronto, ON Canada [7]Division of Neurosurgery, The Hospital for Sick Children, Toronto, Ontario, Canada. Correspondence to: Bo Wang <Bo.Wang@uhnresearch.ca>.

The 2022 ICML Workshop on Computational Biology.

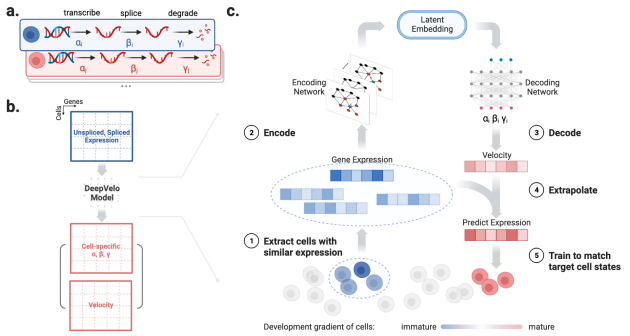

Figure 1. **Overview of the DeepVelo pipeline and velocity prediction method.** (a) DeepVelo estimates cell-specific transcription (α_i), RNA splicing (β_i) and RNA degradation rates (γ_i). (b) Overview of the velocity analysis pipeline using DeepVelo. (c) Overview of the DeepVelo GCN model.

2. Methods

2.1. The differential processes of RNA velocity

For each cell, the dynamics of RNA transcription, splicing, and degradation (Fig.1a) can be approximated as the following differential processes

$$\begin{aligned} \frac{du(t)}{dt} &= \alpha_{i,g} - \beta_{i,g} u(t), \\ \frac{ds(t)}{dt} &= \beta_{i,g} u(t) - \gamma_{i,g} s(t). \end{aligned} \tag{1}$$

where $\alpha_{i,g}, \beta_{i,g}, \gamma_{i,g}$ are the kinetic rates for cell i and gene g. t denotes a time coordinate in the developmental dynamics. Unspliced immature mRNA is first generated by transcription of DNA and then post-transcriptionally modified and spliced into mature mRNA. The dynamics of unspliced RNA abundance, $\frac{du(t)}{dt}$, are modeled by the first differential equation of Eq.1, where $\alpha_{i,g}$ and $\beta_{i,g}$ denote the rates of transcription and splicing, respectively. Similarly, the second equation models the dynamics of spliced RNA abundance, and $\gamma_{i,g}$ denotes the rate for RNA degradation. These kinetic rates are intrinsically cell-specific since there is a high degree of variability in transcriptional dynamics between cells (Hsu & Moses; Larsson et al., 2019). **However, previous velocity estimation techniques assume global constant kinetic rates across cells.**.

2.2. The DeepVelo model and cell-specific prediction

Given the unspliced gene counts $u(t)$ and spliced gene counts $s(t)$ for individual cells obtained in a scRNA-seq experiment, DeepVelo estimates the **RNA velocity**, $\frac{ds(t)}{dt}$, by modeling cell and gene-specific coefficients $\alpha_{i,g}, \beta_{i,g}, \gamma_{i,g}$ in Eq.1 using a Graph Convolutional Neural Network (GCN). As opposed to previous techniques (La Manno et al., 2018; Bergen et al., 2020), DeepVelo models the coefficients per each cell and each gene (Fig.1c), providing more faithful RNA velocity estimates for individual cells.

Graph convolutional network (GCN). The GCN learns node embeddings based on message passing along the graph edges (Kipf & Welling, 2016). Given a graph of nodes V and adjacency matrix A, a multi-layer neural network is constructed on the graph with the following layer-wise propagation rule:

$$H^{(l+1)} = \sigma(\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} H^{(l)} W^{(l)}), \tag{2}$$

where $H^{(l)}$ denotes the node feature vectors at the l-th layer, $\tilde{A} = A + I_N$ is the adjacency matrix, $\tilde{D}$ is the diagonal degree matrix such that $\tilde{D}_{ii} = \sum_j \tilde{A}_{ij}$, $W^{(l)}$ is the layer-specific trainable parameter matrix, and σ is the activation function (e.g. ReLU).

As input, we build a nearest neighbor graph based on the expression values of sequenced cells $G = (\mathcal{V}, \mathcal{E})$. The vertex $\mathfrak{v}_i \in \mathcal{V}$ in the graph denotes the expression reads of cell i, which includes the spliced and unspliced gene expression values $\mathfrak{v}_i = [s_i, u_i]$. A cell i is connected to cell j (i.e. $\mathcal{E}_{ij} = 1$) if cell j is one of its nearest neighbors. The DeepVelo GCN model consists of stacked graph convolution layers (Eq.2).

Cell-specific prediction. The output of the final layer, H^L, is processed by a fully connected neural network, which then yields the estimated velocity parameters $\alpha \in \mathbb{R}^{N \times D}$, $\beta \in \mathbb{R}^{N \times D}$ and $\gamma \in \mathbb{R}^{N \times D}$ for all N cells and D genes. Finally, the estimated **RNA velocity** $v_i \in \mathbb{R}^D$ for each cell is computed as

$$\tilde{v}_i = \beta_i u_i - \gamma_i s_i, \tag{3}$$

where β_i and γ_i are the i-th row in β and γ, u_i and s_i are the unspliced and spliced reads of cell i. In summary, DeepVelo predicts a cell's velocity vector and extrapolates the cell state to match future states extracted from the sequencing data (Fig.1c).

Training DeepVelo. To optimize the parameters in the DeepVelo model, we introduce a probabilistic objective function to calculate the difference between the estimated future cellular state $s_i + \tilde{v}_i$ and the possible future cell states. For each cell ,i, in the sequenced population Ω, we assumed there exists a "$t+1$" neighborhood, $\tilde{\mathcal{N}}_{i,t+1}$, such that cells in the neighborhood contain gene expression values similar to possible future cell i at developmental time $t+1$. This *continuity* assumption can be expressed as the following equation,

$$\frac{1}{\Omega} \sum_{i \in \Omega} \left[s_i + v_i - \sum_{j \in \mathcal{N}_{i,t+1}} s_j P(i \to j) \right] \approx 0, \tag{4}$$

where $i \to j$ denotes that cell i develops at time $t+1$ into a cell that has similar gene expression values as cell j,

and $P(i \to j)$ is the probability of this event. We use this relation (Eq.4) to derive the loss function to optimize the estimated $\tilde{v}_i$, as follows,

$$\mathcal{L}_+ = \frac{1}{\Omega} \sum_{i \in \Omega} \left[s_i + \tilde{v}_i - \sum_{j \in \tilde{\mathcal{N}}_i} s_j P_{c+} (i \to j) \right]^2, \quad (5)$$

where $\tilde{v}_i$, $\tilde{\mathcal{N}}_i$ and P_{c+} are estimated by DeepVelo. To be specific, $\tilde{\mathcal{N}}_i$ contains a large pool of candidate target cells for each cell i selected by nearest-neighbor search in the principal component analysis (PCA) reduction of gene expression values. We used the top 30 neighbors by default for all experiments. P_{c+} uses the cosine similarity between $\tilde{v}_i$ and $s_j - s_i$ at each iteration. The probabilities for $i \to j$ with positive cosine similarities are set to a uniform distribution, and the probabilities are set to zero for the others with non-positive similarities.

Apart from $\mathcal{L}_+$, we use $\mathcal{L}_-$ to model target cells at time $t-1$ similarly,

$$\mathcal{L}_- = \frac{1}{\Omega} \sum_{i \in \Omega} \left[s_i - \tilde{v}_i - \sum_{j \in \tilde{\mathcal{N}}_i} s_j P_{c-} (i \leftarrow j) \right]^2, \quad (6)$$

and we add a direction term of Pearson correlation coefficients to encourage the correct sign of $\tilde{v}_i$ values,

$$\mathcal{L}_{Pearson} = -\left(\lambda_u corr(\tilde{v}_i, u_i) + \lambda_s corr(\tilde{v}_i, -s_i)\right). \quad (7)$$

Overall, DeepVelo is optimized iteratively by minimizing $\mathcal{L}_c = \mathcal{L}_+ + \mathcal{L}_- + \mathcal{L}_{Pearson}$. This promotes the predicted RNA velocities to match the observed target cells in the sequenced datasets.

3. Results

3.1. DeepVelo recovers complex transcriptional dynamics for individual cells

To test the ability of DeepVelo to identify complex kinetics, we utilized a neurogenesis scRNA-seq dataset of the developing mouse dentate gyrus (Hochgerner et al., 2018) consisting of tissue samples from two time points, P12 and P35 (postnatal day 12 and 35), collected by droplet-based single-cell RNA sequencing. After pre-processing, we calculated the RNA velocities using the proposed DeepVelo model and the dynamical model from scVelo (Bergen et al., 2020). The velocity plots show the projection of calculated velocity vectors onto the UMAP (McInnes et al., 2018)-based embedding of the scRNA-seq data. In the velocity estimates (Fig.2a), DeepVelo accurately recovers ground-truth developmental directions aligned with existing literature. In particular, the granule cell lineage dominates the main structure in the plot, where neuroblast cells develop into immature and mature granule cells. The directions of DeepVelo velocity estimates between celltypes follow ground-truth patterns validated by Hochgerner et al. (2018).

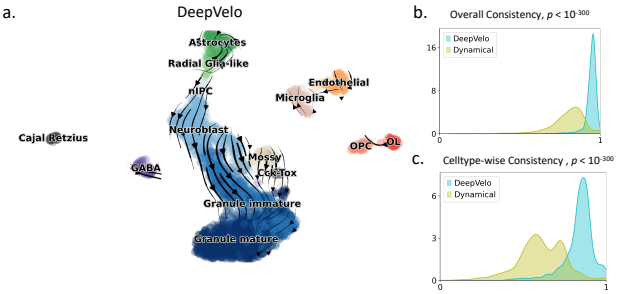

Figure 2. (a) Velocity plot for dentate gyrus neurogenesis (Hochgerner et al., 2018) using DeepVelo. (b) The histogram of the overall consistency scores. (c) The histogram of the celltype-wise consistency scores. DeepVelo achieves higher scores in (b,c) compared to the dynamic (scVelo) model (Bergen et al., 2020).

Further, DeepVelo quantitatively shows better performance through higher consistency of RNA velocity values. To compute consistency of velocity estimates, we (1) compute the average cosine similarity of the velocity vector of each cell to its neighbors, which defines the overall consistency, and (2) since the overall consistency could be biased toward over-smoothened estimations, we also propose a celltype-wise consistency to complement the overall score. The celltype-wise consistency computes the average cosine similarity of each cell's velocity to all velocity vectors of the same celltype. For both metrics, DeepVelo shows significant improvements over the scVelo dynamical method, with higher average consistency scores (Mann-Whitney U Test $p < 1 \times 10^{-300}$, Fig.2b,c).

3.2. Cell-specific kinetic rate estimates enable accurate quantification of time-dependant and branching gene dynamics

Due to the cell-specific estimation of ($\alpha_{i,g}, \beta_{i,g}, \gamma_{i,g}$ in Eq.1), DeepVelo for the first time provides a profile of individual kinetic rates for each cell. In Fig.3a, we show the UMAP projection of all cell-specific kinetic rates of 2930 cells of the dentate gyrus dataset from the previous section. Although DeepVelo is **unaware of celltype labels during training**, the learned kinetic rates are naturally clustered into groups corresponding to ground-truth celltypes. Further, clusters of cells from the same lineage (e.g. the outlined granule lineage) are positioned closely compared to other cells. Overall, the similarity of learned cell-specific kinetic rates reflects the biological similarity of cells at both the celltype and lineage levels.

Velocity-associated kinetic rates across cells may vary for

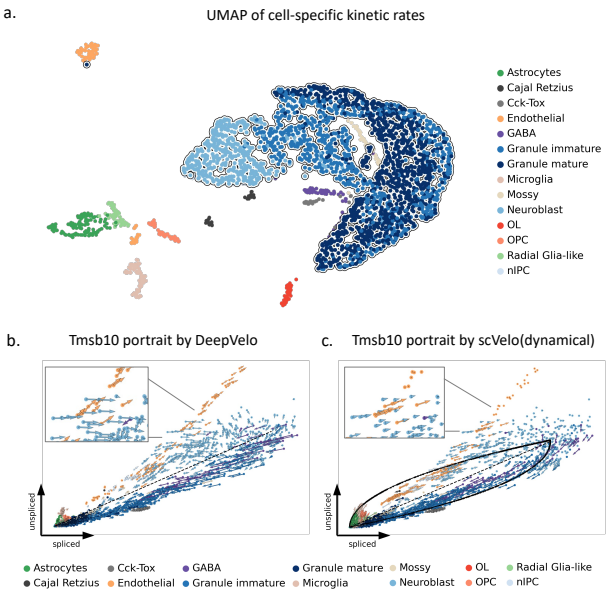

Figure 3. (a) UMAP of cell-specific kinetic rates show clusters consistent with celltypes. (b,c) Velocity vectors in the unspliced/spliced phase portrait of *Tmsb10*, computed by DeepVelo and the scVelo dynamical model respectively.

genes undergoing dynamic regulation involving multiple processes. These varying kinetics are often misinterpreted in existing velocity methods (Bergen et al., 2021). This stems from the fact that previous methods used constant cell-agnostic coefficients for modeling the kinetic rates in (Eq.1). In contrast, DeepVelo provides estimates for different celltypes and cell-states by introducing cell-specific kinetic rates, leading to more biologically accurate velocity estimation in multi-faceted systems.

In the mouse dentate gyrus dataset (Hochgerner et al., 2018), *Tmsb10* shows multiple kinetic regimes and trajectories in the phase portrait of spliced and unspliced reads. The cells in the granule lineage (including neuroblast, granule immature and granule mature celltypes) form a cyclic trajectory. The endothelial cells, which are not a part of this lineage, also undergo dynamics for *Tsmb10* (Fig.3b,c). These two regimes are very likely regulated by different kinetic rates.

For the multi-faceted dynamics of *Tsmb10*, DeepVelo correctly predicted the patterns for both regimes of *Tsmb10* (Fig.3b). For the granule lineage, DeepVelo captures the direction of velocity from neuroblast cells to granule immature cells and then to granule mature cells. For the endothelial cells, the predicted velocity direction correctly points to the position of the same celltype with amplified spliced reads (Fig.3b - Zoom-in panel). In contrast to DeepVelo, scVelo forces the velocities to follow the assumed cyclic trajectory of the model (Fig.3c). As a result, although scVelo successfully captures the trajectory for the granule lineage, it incorrectly predicts the velocities of endothelial cells to lead towards differentiation of neuroblasts, (Fig.3c - Zoom-in panel). Overall, DeepVelo is capable of predicting celltype-specific velocity for cells within the same region, due to the advantage of cell-specific kinetic rate estimation (Fig.3b).

3.3. Predicting differentiation and driver genes in multi-lineage hindbrain developmental data

To test velocity methods in a complex setting with multiple lineages, we applied DeepVelo and scVelo on a mouse hindbrain development dataset comprising of the GABAergic and gliogenic lineages (Vladoiu et al., 2019) (Fig.4(a)). DeepVelo's ability to learn cell-specific kinetic rates accounts for the multi-faceted differentiation present in this multi-lineage system. The result of DeepVelo (Fig.4(b)) shows RNA velocity vectors over the developmental process from Neural stem cells to the differentiating GABA interneurons and gliogenic progenitors. We performed trajectory inference using *PAGA* (Wolf et al., 2019) over the velocity graph of DeepVelo, and found that DeepVelo was able to recapitulate the ground-truth differentiation pattern - specifically the branching between VZ progenitors and differentiating GABA interneurons, and VZ progenitors and gliogenic progenitors (Fig.4 (c)).

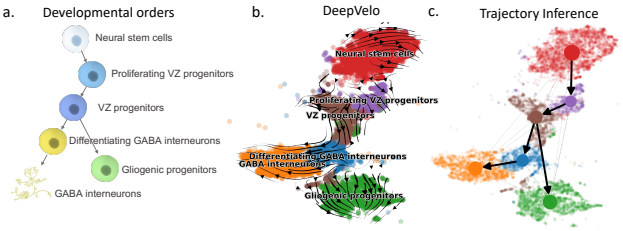

Figure 4. The (a) ground-truth trajectory, (b) DeepVelo velocity estimates, and (c) DeepVelo inferred trajectory for the mouse hindbrain dataset.

Further, we found that DeepVelo was able to predict important driver genes in the differentiation of the GABAergic and gliogenic lineages, and was able to do so at a significantly higher rate than the scVelo dynamical model (Table.1). Using the velocity estimates of DeepVelo, we utilized the *CellRank* (Lange et al., 2020) technique to estimate driver gene relevance. Analyzing the top 100 driver gene results, DeepVelo indicated higher enrichment for known marker genes from the Vladoiu et al. (2019) dataset, and annotated mouse transcription factors, which are both more likely to have a functional role in driving differentiation. Pathway enrichment analysis of the top predicted driver genes was performed using Gene Ontology (GO) and REACTOME pathway databases. The result for DeepVelo indicated significant recovery of relevant functional terms for neurogenesis

and development (as a percentage of enriched pathways) compared to scVelo for the GABAergic lineage, while the percentages were fairly similar for the gliogenic lineage.

Table 1. Driver-gene analysis results of the GABAergic and Gliogenic hindbrain development lineages.

	scVelo		DeepVelo	
Top 100 driver genes	**GABAergic**	**Gliogenic**	**GABAergic**	**Gliogenic**
Marker gene overlap	44	66	**86**	**81**
Transcription-factor overlap	10	5	**14**	**5**
Pathway relevance percentage	19.4	51.3	**62.9**	**58.3**

4. Conclusion

DeepVelo is a novel GCN framework for estimating RNA velocity that is not limited by assumptions of constant, cell-agnostic RNA transcription, splicing, and degradation rates. By estimating these rates at a cell-specific level, it outperforms current state-of-the-art techniques on challenging scRNA-seq datasets with more consistent velocity estimates, has the ability to generate multi-faceted velocities for each gene, and can model complex multi-lineage dynamics.

RNA velocity estimation remains a major challenge due to complex dynamics, limitations in modelling, and sparse/noisy RNA readouts in common sequencing protocols (Bergen et al., 2021; Gorin et al., 2022). DeepVelo undoes a major assumption of previous methods by taking into account cell-specific "bursting" RNA expression kinetics and the tendency of similar cells to share these kinetic patterns. Through this, DeepVelo offers a step towards a generalized and biologically accurate framework of RNA velocity estimation.

References

Bergen, V., Lange, M., Peidli, S., Wolf, F. A., and Theis, F. J. Generalizing rna velocity to transient cell states through dynamical modeling. *Nature biotechnology*, pp. 1–7, 2020.

Bergen, V., Soldatov, R. A., Kharchenko, P. V., and Theis, F. J. Rna velocity—current challenges and future perspectives. *Molecular systems biology*, 17(8):e10282, 2021.

Gorin, G., Fang, M., Chari, T., and Pachter, L. Rna velocity unraveled. *bioRxiv*, 2022.

Hochgerner, H., Zeisel, A., Lönnerberg, P., and Linnarsson, S. Conserved properties of dentate gyrus neurogenesis across postnatal development revealed by single-cell rna sequencing. *Nature neuroscience*, 21(2):290–299, 2018.

Hsu, I. S. and Moses, A. M. Stochastic models for single-cell data: Current challenges and the way forward. *The FEBS Journal*, 289(3):647–658.

Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

La Manno, G., Soldatov, R., Zeisel, A., Braun, E., Hochgerner, H., Petukhov, V., Lidschreiber, K., Kastriti, M. E., Lönnerberg, P., Furlan, A., et al. Rna velocity of single cells. *Nature*, 560(7719):494–498, 2018.

Lange, M., Bergen, V., Klein, M., Setty, M., Reuter, B., Bakhti, M., Lickert, H., Ansari, M., Schniering, J., Schiller, H. B., et al. Cellrank for directed single-cell fate mapping. *BioRxiv*, 2020.

Larsson, A. J., Johnsson, P., Hagemann-Jensen, M., Hartmanis, L., Faridani, O. R., Reinius, B., Segerstolpe, Å., Rivera, C. M., Ren, B., and Sandberg, R. Genomic encoding of transcriptional burst kinetics. *Nature*, 565(7738): 251–254, 2019.

McInnes, L., Healy, J., and Melville, J. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

Vladoiu, M. C., El-Hamamy, I., Donovan, L. K., Farooq, H., Holgado, B. L., Sundaravadanam, Y., Ramaswamy, V., Hendrikse, L. D., Kumar, S., Mack, S. C., et al. Childhood cerebellar tumours mirror conserved fetal transcriptional programs. *Nature*, 572(7767):67–73, 2019.

Wolf, F. A., Hamey, F. K., Plass, M., Solana, J., Dahlin, J. S., Göttgens, B., Rajewsky, N., Simon, L., and Theis, F. J. Paga: graph abstraction reconciles clustering with trajectory inference through a topology preserving map of single cells. *Genome biology*, 20(1):1–9, 2019.