

Dissecting and steering cell dynamics using spatially-informed RNA velocity with veloAgent

Vishvak Raghavan^{1,2,3,5}, Brent Yoon^{1,2,5}, Gregory J Fonseca⁴, Yue Li^{1,3} & Jun Ding^{1,2,3} 

Abstract

RNA velocity enables inference of cell state transitions from single-cell transcriptomics by modeling transcriptional dynamics from spliced and unspliced mRNA. However, existing methods overlook spatial context and struggle to scale to large datasets, limiting insights into tissue organization and dynamic processes. We introduce veloAgent, a deep generative and agent-based framework that estimates gene- and cell-specific transcriptional kinetics while integrating spatial information through agent-based simulations of local microenvironments. By leveraging both molecular and spatial cues, veloAgent improves velocity accuracy and achieves sublinear memory scaling, enabling efficient analysis of large and multi-batch spatial datasets. A distinctive feature of veloAgent is its in silico perturbation module, which allows targeted manipulation of spatial velocity vectors to simulate regulatory interventions and predict their impact on cell fate dynamics. These capabilities position veloAgent as a scalable and versatile framework for dissecting spatially resolved cellular dynamics and guiding cell fate manipulation across diverse biological processes.

Subject Category Computational Biology

<https://doi.org/10.1038/s44320-026-00213-w>

Received 8 January 2026; Revised 9 April 2026;

Accepted 16 April 2026

Published online: 6 May 2026

Introduction

Dynamic changes in gene expression, including both transcriptional and post-transcriptional regulation are fundamental to determining a cell's identity and function (Kouno et al, 2013; Pascual-Ahuir et al, 2020). Single-cell RNA sequencing (scRNA-seq) has revolutionized the field of transcriptomics by enabling high-resolution profiling of gene expression at the single-cell level, thereby facilitating unprecedented insights into cellular heterogeneity, gene regulatory dynamics, and developmental trajectories (Kashima et al, 2020; Kolodziejczyk et al, 2015). In parallel, advances in spatial transcriptomics have made it possible to preserve spatial context while quantifying gene expression,

enabling researchers to investigate tissue organization, niche-specific expression programs and cell–cell interactions (Vahid et al, 2023; Williams et al, 2022; Zhao et al, 2024). Despite its quantitative precision and comprehensive coverage, these technologies only provide a temporally static snapshot of cellular states rather than a trajectory of cellular behavior. This limitation constrains our ability to reconstruct dynamic processes such as differentiation, and lineage commitment (Tritschler et al, 2019). Addressing these limitations, RNA velocity has emerged as a transformative framework for inferring future cellular states by leveraging the relative abundances of unspliced and spliced mRNA transcripts to model ongoing gene expression kinetics (Bergen et al, 2020; La Manno et al, 2018). This approach not only enables the extrapolation of transcriptional dynamics and trajectory predictions but also offers critical insights into cellular decision-making processes with broad applications in developmental biology, and disease progression.

Several computational frameworks have been developed to infer RNA velocity, including cellDancer (Li et al, 2024), scVelo (Bergen et al, 2020), SIRV (Abdelaal et al, 2024), DeepVelo (Cui et al, 2024), VeloVAE (Gu et al, 2022), VeloVI (Gayoso et al, 2024), and UniTVelo (Gao et al, 2022). These methods introduce various strategies to improve accuracy, such as incorporating dynamical systems modeling, deep generative modeling, and kinetic parameter estimation. However, existing RNA velocity methods face several limitations that constrain our ability to fully understand development and disease progression. Many existing methods still assume uniform transcriptional kinetics across genes or cells, an oversimplification that erases the cell-to-cell heterogeneity that drives development, disease onset and progression. While more advanced models such as cellDancer estimate gene- and cell-specific kinetics and yield biologically plausible results, their heavy computational demands make them impractical for the large, multi-batch datasets increasingly central to development or disease atlas-scale studies (Li et al, 2024). Perhaps, most critically, nearly all current approaches ignore spatial context, relying instead on embeddings of transcriptional similarity. This omission ignores vital information from tissue microenvironments and cell–cell interactions, leading to inferred trajectories that cross anatomical boundaries or contradict known tissue structure (Armingol et al, 2021; Asp et al, 2020; Zheng et al, 2023). This loss of spatial fidelity prevents us from uncovering how tumors invade, how immune cells coordinate in tissue, or how stem cells rebuild organs (Hunter et al, 2021; Oliveira

¹School of Computer Science, McGill University, Montreal, QC, Canada. ²Meakins-Christie Laboratories, Translational Research in Respiratory Diseases Program, Research Institute of the McGill University Health Centre, Montreal, QC, Canada. ³Mila - Quebec Artificial Intelligence Institute, Montreal, QC, Canada. ⁴Khalifa University, Abu Dhabi, United Arab Emirates. ⁵These authors contributed equally: Vishvak Raghavan, Brent Yoon. ✉E-mail: jun.ding@mcgill.ca

et al, 2025; Tang et al, 2025). Furthermore, no current framework allows researchers to computationally perturb key genes or pathways and predict how such perturbations reshape cellular trajectories. Such a tool would be especially powerful in regenerative medicine and disease intervention, where rapid in silico tests could help prioritize experiments, highlight therapeutic targets, and significantly reduce the time and cost of discovery pipelines. Together, these gaps limit not just technical accuracy but our ability to translate single-cell maps into insights about developmental and disease biology. This limitation has also constrained our understanding of how cellular dynamics unfold within the structural and spatial context of native tissues.

To address these challenges, we introduce *veloAgent*, a hybrid framework consisting of three distinct components: a variational autoencoder (VAE), a gene interaction-informed neural network, and an agent-based model (ABM) to infer spatial RNA velocities. The VAE learns robust latent representations from noisy spliced and unspliced mRNA counts, providing a unified and denoised embedding of transcriptional states. The neural network guided by gene–gene interaction networks derived from the STRING database encodes cell- and gene-specific transcriptional kinetics, while operating in a fixed-dimensional, gene-centric space that scales efficiently to a large number of cells (Szkarczyk et al, 2023). The ABM then leverages spatial transcriptomics data to simulate each cell as an autonomous agent positioned in its native tissue context, whose future state is influenced by neighboring cells within its local environment. ABMs have previously been used in spatial transcriptomics to model interactions between cells, supporting our choice of this modeling paradigm (Raghavan et al, 2025). This integrative design enables *veloAgent* to model spatial RNA velocity, capturing how gene expression dynamics unfold across both time and space. By embedding tissue architecture directly into the inference process, we bridge the divide between molecular kinetics and spatial organization, transforming RNA velocity from an abstract representation into a spatially coherent model of cell-state transitions.

We demonstrate the utility of *veloAgent* across a range of spatial transcriptomics platforms with varying spatial resolutions, including Visium (Stahl et al, 2016), Stereo-seq (Chen et al, 2022), HybISS (Gyllborg et al, 2020), and seqFISH (Lohoff et al, 2022). Unlike existing RNA velocity methods, *veloAgent* does not simply infer trajectories from just the transcriptome, but directly integrates spatial coordinates, enabling analysis of velocity vectors onto tissue architecture. This uniquely allows researchers to resolve dynamic cellular programs within native niches, revealing transition pathways that respect anatomical boundaries and microenvironmental organization rather than producing over-smoothed, anatomically implausible fields. *veloAgent* is also designed for scale with its gene-centric framework that makes velocity analysis practical on atlas-scale and multi-batch datasets. Finally, *veloAgent* introduces an innovation not previously available in velocity analysis: in silico perturbation of transcriptional kinetics. This enabling virtual screening and prioritization of interventions, generating testable hypotheses that reduce experimental cost and accelerate translational discovery. By integrating transcriptomic kinetics, gene-network priors, and spatial context, *veloAgent* enables both mechanistic interpretation and in silico perturbation, defining a new paradigm for data-driven hypothesis generation and therapeutic target discovery.

Results

Overview of *veloAgent*

veloAgent introduces a hybrid deep learning and agent-based modeling (ABM) framework for estimating RNA velocity at single-cell resolution by integrating gene expression data with prior gene–gene interaction networks and spatial information. Unlike existing RNA velocity frameworks, *veloAgent* jointly models transcriptional kinetics and spatial context within a unified deep generative-agent-based model hybrid. As illustrated in Fig. 1, the methodology proceeds through four main stages. First (Fig. 1A), a variational autoencoder (VAE) encodes each cell's spliced and unspliced mRNA counts into a unified latent representation. This step reduces dimensionality while preserving key transcriptional features that reflect dynamic cellular states. Second (Fig. 1B), the resulting latent vectors are processed through a biologically constrained gene–gene deep neural network (DNN) to infer cell- and gene-specific kinetic parameters—transcription rate (α), splicing rate (β), and degradation rate (γ). The network architecture is guided by protein–protein interactions from the STRING database, enforcing biologically grounded sparsity: each gene's output depends only on its annotated interactors. This design enables the model to capture meaningful regulatory dependencies during kinetic inference. The inferred parameters are then incorporated into differential equations to compute initial RNA velocity estimates.

The third stage (Fig. 1C) refines the estimated velocity field by enforcing local coherence. Each cell's velocity vector is aligned with those of its transcriptionally similar neighbors, minimizing directional noise and enhancing trajectory continuity. Finally (Fig. 1D), for datasets with spatial transcriptomics measurements, *veloAgent* integrates an ABM to further refine velocities in their spatial context. The ABM incorporates intercellular distance, local cell density, and expression similarity to simulate spatial constraints, yielding biologically realistic and spatially informed velocity fields in structured tissue environments. Beyond velocity inference, *veloAgent* provides multiple outputs for downstream analysis—including temporal and spatial velocity projections—and features a perturbation module that enables in silico steering of RNA velocity directions to model potential gene perturbations (Fig. 1E). Together, these components establish *veloAgent* as a unified framework for reconstructing and perturbing cellular dynamics in both temporal and spatial contexts.

veloAgent enables spatially informed RNA velocity inference

veloAgent extends RNA velocity analysis by embedding spatial information directly into the inference process, allowing trajectories to better reflect tissue architecture and neighborhood organization. Through an agent-based spatial refinement module, *veloAgent* integrates intercellular proximity, local density, and expression similarity to couple neighboring cells during velocity estimation. This spatial coupling promotes locally aligned velocity directions and suppresses noise arising from isolated or contradictory cell states. As a result, *veloAgent* produces coherent and biologically interpretable velocity fields that conform to underlying tissue structure and resolve directional ambiguities commonly observed in conventional methods.




Figure 1. Overview of the veloAgent framework.

veloAgent is a hybrid deep learning and agent-based modeling framework for estimating RNA velocity and decoding cellular dynamics by integrating spliced and unspliced transcriptomics with spatial information. The model operates in four main phases, followed by versatile downstream applications. (A) Phase 1: Latent representation learning. Spliced and unspliced gene expression matrices are separately encoded using variational autoencoders (VAEs). The resulting embeddings are concatenated into a shared latent space Z , which captures underlying cellular states while minimizing reconstruction error. (B) Phase 2: Cell- and gene-specific kinetic rate inference. Kinetic parameters—transcription (α), splicing (β), and degradation (γ) rates—are inferred for each gene in each cell by passing the latent vector Z through a gene-informed neural network. The network is trained to maximize cosine similarity between predicted and observed RNA velocity vectors across cells. (C) Phase 3: Velocity vector refinement. Cell-level velocity vectors are refined to enforce local coherence in the latent space, aligning similar cells toward consistent velocity directions and improving trajectory continuity. (D) Phase 4: Spatial RNA velocity refinement. When spatial transcriptomics data are available, an agent-based model further refines velocity vectors by incorporating spatial priors, weighting the influence of neighboring cells based on intercellular distance and expression similarity. (E) Downstream applications. veloAgent supports multiple downstream analyses, including temporal and spatially informed velocity estimation, projection of velocity fields onto spatial tissue maps, and in silico perturbation experiments to identify candidate interventions that may reverse disease-associated cell state transitions.

To evaluate its performance, we applied veloAgent to four representative spatial transcriptomics datasets—mouse brain (Stereo-seq), breast cancer, chicken heart, and mouse brain (HybISS). Across all datasets, veloAgent achieved superior performance in both qualitative (Fig. 2A–H) and quantitative (Fig. 2I–K) benchmarks, demonstrating that spatial context provides a more faithful representation of tissue-level dynamics. In most regions, veloAgent and SIRV yielded broadly consistent trends; however, veloAgent clearly outperformed in complex or transitional zones where spatial organization was critical. In the Stereo-seq mouse brain, for example, SIRV produced unclear or contradictory velocities around fiber-tract cells, whereas veloAgent captured coherent transitions from hippocampus, thalamus sensory-motor (DORsm), and polymodal association (DORpm) regions, consistent with known projections of fiber-tract neurons (Fig. 2A; red circles) (Ferguson and Gao, 2015; Guido, 2018; Qiu et al, 2024; Senova et al, 2020). In breast cancer, SIRV incorrectly directed flows from tumor cells toward normal cells, whereas veloAgent corrected these mispredictions, yielding biologically consistent trajectories (Fig. 2B; red circles) (Sahai et al, 2020). In the chicken heart, veloAgent corrected SIRV's reversed trajectories at the ventricle-atrium boundary and valve cluster (Fig. 2C; red circles) (Tallquist, 2020). Similarly, in the HybISS brain, SIRV misdirected neuron velocities toward neuroblasts, whereas veloAgent correctly modeled neurons as terminal states (Fig. 2D; red circles) (Xu et al, 2020). Collectively, these results highlighted veloAgent's ability to recover biologically consistent transitions where other approaches failed.

We next validated the spatial fidelity of inferred velocities using well-established regional marker genes that define terminal or tissue-specific populations (Fig. 2E–H). In each dataset, veloAgent's velocity fields aligned closely with the known spatial domains of these markers, whereas SIRV frequently produced diffuse or misplaced activity. In the Stereo-seq mouse brain, *Mbp* velocity was confined to fiber tracts, consistent with its role in oligodendrocyte maturation and myelination (Fig. 2E) (Xu et al, 2020). In breast cancer, *ERBB2* velocity localized correctly within tumor regions but was erroneously assigned to macrophages by SIRV, contradicting its established role as an oncogene and therapeutic target (Fig. 2F) (Yu and Hung, 2000). In the chicken heart, *NPPA* velocity was restricted to atria under veloAgent yet aberrantly extended into ventricles with SIRV, inconsistent with its function as an atrial differentiation marker distinguishing atria from ventricles (Fig. 2G) (Li, 2022). In the HybISS mouse brain, *Nrg1* velocity was specifically confined to neurons, consistent with its role in neuronal

development, while SIRV failed to capture any localized signal (Fig. 2H) (Rajebhosale et al, 2024). These spatial-marker-based evaluations confirm that veloAgent faithfully captures known anatomical domains and improves the spatial interpretability of RNA velocity.

Quantitative benchmarking against seven leading RNA-velocity models (cellDancer, SIRV, DeepVelo, veloVAE, veloVI, UniTVelo, scVelo) confirmed these gains (Fig. 2I–K). Evaluated across three complementary metrics—fate probability (Bergen et al, 2020), cross-boundary direction correctness (Qiao and Huang, 2021), and velocity confidence (Qiao and Huang, 2021)—veloAgent achieved the highest or near-highest scores in all datasets. For cross-boundary direction correctness, veloAgent was on average 92.03% (Stereo-seq brain), 92.02% (breast cancer), 55.52% (chicken heart), and 12.15% (HybISS brain) better than all other methods. Specifically, veloAgent matched veloVI in Stereo-seq brain (0.8447 vs. 0.8590, n.s) and outperformed all methods in the other datasets: breast cancer (0.6148 vs. 0.2849, cellDancer; FDR = 3.588×10^{-60}), chicken heart (0.6742 vs. 0.6193, cellDancer; FDR = 2.923×10^{-3}), and HybISS brain (0.1237 vs. -0.0209, SIRV; FDR = 3.880×10^{-3}). For fate probability, veloAgent consistently showed stronger or comparable predictive power in comparison to other methods and on average, veloAgent achieved 64.97% (Stereo-seq brain), 65.87% (breast cancer), 58.12% (chicken heart), and 33.32% (HybISS brain) higher predictive power than all competing methods. Specifically, veloAgent consistently demonstrated stronger or comparable performance: Stereo-seq brain (0.9995 vs. 0.9919, veloVI; FDR = 6.990×10^{-132}), breast cancer (0.9841 vs. 0.6309, DeepVelo; FDR < 1.0×10^{-308}), chicken heart (0.9779 vs. 0.8520, cellDancer; FDR < 1.0×10^{-308}), and HybISS brain (0.8324 vs. 0.6231, cellDancer; FDR = 1.071×10^{-51}). For velocity confidence, veloAgent achieved on average 35.61% (Stereo-seq brain), 34.92% (breast cancer), 31.23% (chicken heart), and 28.25% (HybISS brain) higher internal consistency than all competing methods. Specifically, veloAgent maintained near-perfect internal consistency (≥ 0.998) across datasets: Stereo-seq brain (0.9995 vs. 0.8704, DeepVelo; FDR < 1.0×10^{-308}), breast cancer (0.9987 vs. 0.8304, SIRV; FDR < 1.0×10^{-308}), chicken heart (0.9999 vs. 0.6349, veloVAE; FDR < 1.0×10^{-308}), and HybISS brain (0.9987 vs. 0.9524, cellDancer; FDR < 1.0×10^{-308}). These improvements reflect more accurate modeling of transition boundaries, terminal fates, and noise-robust velocity estimation.

Although principally designed for spatial data, veloAgent also applies to conventional scRNA-seq datasets, maintaining advantages over existing approaches (Appendix Figs. S1–2, Appendix




Figure 2. Performance benchmarking of veloAgent for spatially informed RNA velocity inference.

This figure evaluates veloAgent's ability to infer RNA velocity in spatial transcriptomics datasets, comparing its performance to existing methods using both qualitative and quantitative benchmarks. (A–D) RNA velocity vectors computed by veloAgent (left) and SIRV (right), projected onto UMAP embeddings of four developmental datasets: (A) mouse brain (Stereo-seq); (B) breast cancer; (C) chicken heart; and (D) mouse brain (HybISS). Red circles highlight key regions of disagreement between methods. (E–H) Velocity fields of marker genes expressed in terminal cell populations, comparing veloAgent (top) to SIRV (bottom): (E) *Mbp* in mouse brain (Stereo-seq); (F) *ERBB2* in breast cancer; (G) *NPPA* in chicken heart; (H) *Nrg1* in mouse brain (HybISS). Lighter colors (green and yellow) represent higher magnitudes of velocity, whereas darker colors (purple) represent lower magnitudes of velocity. veloAgent demonstrates stronger alignment with expected terminal regions. (I) Cross boundary scores across benchmarked methods (veloAgent, cellDancer, SIRV, DeepVelo, veloVAE, veloVI, UniTVelo, scVelo), measuring the correctness of predicted transitions across annotated cluster boundaries, based on ground-truth developmental labels. Higher scores reflect better agreement with known trajectories. (J) Fate probability scores across methods, quantifies the likelihood that a cell will progress toward a defined terminal state based on inferred RNA velocity. (K) Velocity confidence scores across benchmarked methods, quantifying how well predicted velocities align with local gene expression structure. Higher scores indicate more reliable predictions. Boxplots in (I–K) represent the median, first quartile and third quartile with error bars extending to the furthest value within 1.5 of the interquartile range (IQR). Each panel displays four boxplots for Mouse Brain (Stereo-seq, $n = 7765$), Breast Cancer ($n = 4628$), Chicken Heart ($n = 1965$), and Mouse Brain (HybISS, $n = 1656$), where n denotes the number of cells in each dataset. Statistical significance calculated using Mann-Whitney U-test with FDR correction. Significance levels are indicated as follows: n.s, not significant ($P \geq 0.05$); * ($P < 0.05$); ** ($P < 0.01$); *** ($P < 0.001$); and **** ($P < 0.0001$). Exact P -values can be found in Appendix Table S1.

Table S2). Extended benchmarking across additional datasets and metrics (Appendix Figs. S3–5) further demonstrate its robustness, positioning veloAgent as a powerful framework for reconstructing cellular dynamics with spatial fidelity and quantitative precision.

Spatially informed RNA velocity reveals tissue-coherent cell-state transitions

Most existing RNA velocity methods infer transcriptional dynamics solely in latent transcriptomic space, treating velocity as an abstract trajectory independent of tissue structure. In contrast, veloAgent integrates spatial priors into RNA velocity inference through an ABM, embedding velocity vectors directly onto tissue coordinates to capture how cellular transitions are shaped by native architecture. Importantly, the projected velocity vectors represent inferred transcriptional state transitions and do not imply physical cell migration within the tissue (Fig. 3A–D). This integration enables spatially coherent and biologically grounded reconstruction of cell-state dynamics. Across four spatial transcriptomics datasets, veloAgent produced velocity fields that were directionally coherent, spatially continuous, and extended along broader developmental transition tracks than competing methods (Fig. 3A–D; Appendix Fig. S6). By enforcing local spatial coherence through ABM simulations, veloAgent revealed how tissue organization constrains and guides cellular trajectories—an effect absent from models that neglect spatial context.

At the gene level, veloAgent more accurately captured the spatial dynamics of key driver genes than other methods. For terminal-state marker genes, velocity magnitudes were tightly localized to expected regions, reflecting transcriptional activation or repression near terminal states (Fig. 3E,F; Appendix Figs. S7 and 8). Although expression patterns were comparable across methods, their velocity profiles differed markedly: veloAgent consistently yielded sharper and more biologically meaningful velocity patterns. For example, in the Stereo-seq mouse brain, *Mbp* velocity was concentrated within fiber tracts, consistent with its role in oligodendrocyte myelination (Xu et al, 2020) (Fig. 3E,F; left), whereas cellDancer incorrectly predicted high velocity outside these regions. In the breast cancer dataset, *ERBB2* velocity was tightly restricted to cancer cell populations by veloAgent (Fig. 3E,F; middle left), whereas SIRV lacked this specificity despite *ERBB2* being a known oncogene driving tumor progression (Yu and Hung, 2000). Similarly, in the HybISS mouse brain, *Nrg1* velocity localized to neurons in line with its role in neuronal development (Rajebhosale et al, 2024) (Fig. 3E,F;

right), while SIRV and DeepVelo produced diffuse or nonspecific signals. Even when velocity and expression were imperfectly aligned, such as for *NPPA* in the chicken heart (Fig. 3E,F; middle right), veloAgent still best reflected biology by capturing its role as an atrial differentiation marker (Li, 2022).

Quantitatively, veloAgent consistently achieved the highest scores across cross-boundary direction, spatial velocity confidence, and fate probability metrics among all datasets (Fig. 3G; Appendix Fig. S9), confirming superior alignment between predicted velocity and spatially localized transitions. For cross boundary direction, veloAgent achieved on average 49.57% (Stereo-seq brain), 54.30% (breast cancer), 72.74% (chicken heart), and 8.75% (HybISS brain) higher accuracy than all other methods, consistently matching or surpassing competing models across datasets.

Specifically, veloAgent achieved superior performance relative to the top 3 next-best methods: Stereo-seq brain (0.5648 vs. 0.5594, veloVI; n.s) (0.5648 vs. 0.5492, DeepVelo; n.s) (0.5648 vs. −0.1164, UniTVelo; FDR = 6.845×10^{-35}), breast cancer (0.3816 vs. 0.2806, DeepVelo; FDR = 3.192×10^{-2}) (0.3816 vs. 0.00654, scVelo; FDR = 2.646×10^{-5}) (0.3816 vs. 0.0162, UniTVelo; FDR = 1.490×10^{-9}), chicken heart (0.5744 vs. 0.4714, DeepVelo; FDR = 3.631×10^{-2}) (0.5744 vs. 0.4511, SIRV; FDR = 4.325×10^{-2}) (0.5744 vs. −0.0759, UniTVelo; FDR = 5.793×10^{-9}), and HybISS brain (0.1937 vs. 0.1273, cellDancer; FDR = 2.041×10^{-2}); (0.1937 vs. 0.1228, DeepVelo; FDR = 2.041×10^{-2}) (0.1937 vs. 0.1199, SIRV; FDR = 2.041×10^{-2}) showing that it can consistently match or outperform other velocity inference methods.

To further assess spatial fidelity, we introduced two analog metrics—spatial velocity confidence and spatial fate probability—representing spatially mapped extensions of their conventional counterparts. These two metrics are spatially-aware and help specifically assess spatial accuracy in addition to transcriptomic performance. On average, veloAgent achieved 46.30% (Stereo-seq brain), 13.77% (breast cancer), 29.69% (chicken heart), and 30.36% (HybISS brain) higher predictive power in spatial fate probability than competing methods (Appendix Fig. S9B). Specifically, against the second-best method, veloAgent achieved: Stereo-seq brain (0.7250 vs. 0.5954, veloVI; FDR < 1.0×10^{-308}), breast cancer (0.5648 vs. 0.5370, SIRV; FDR = 1.351×10^{-2}), chicken heart (0.8331 vs. 0.7926, cellDancer; FDR = 8.071×10^{-101}), and HybISS brain (0.7769 vs. 0.7280, cellDancer; FDR = 5.249×10^{-40}). veloAgent demonstrated on average 36.63% (Stereo-seq brain), 36.22% (breast cancer), 33.03% (chicken heart), and 35.04% (HybISS brain)




Figure 3. Spatial RNA velocity inference by veloAgent.

veloAgent's ability to infer and embed RNA velocity onto spatial coordinates is compared against competing methods across multiple developmental spatial transcriptomics datasets. (A–D), Spatial RNA velocity fields computed by veloAgent (top) and SIRV (bottom), overlaid on tissue coordinate maps for: (A) mouse brain (Stereo-seq); (B) breast cancer; (C) chicken heart; and (D) mouse brain (HybISS). Arrows indicate predicted transcriptional state transitions inferred from RNA velocity and do not represent physical movement of cells across spatial locations. (E) Spatial distribution of spliced expression for key marker genes associated with terminal cell states: *Mbp* in mouse brain (Stereo-seq), *ERBB2* in breast cancer; *NPPA* in chicken heart, *Nrg1* in mouse brain (HybISS), visualized across spatial coordinates. (F) Spatial RNA velocity fields for the same marker genes, comparing veloAgent (top left), SIRV (top right), cellDancer (bottom left) and DeepVelo (bottom right). veloAgent captures clearer and more localized velocity patterns toward terminal regions, consistent with expected developmental trajectories. Lighter colors (green and yellow) represent higher magnitudes of velocity, whereas darker colors (purple) represent lower magnitudes of velocity. (G) Cross-boundary direction scores comparing veloAgent and competing methods for spatial velocity inference, measuring the correctness of predicted transitions across spatially localized cell state boundaries. Boxplots show consistently higher accuracy for veloAgent across datasets, Mouse Brain (Stereo-seq, $n = 7765$), Breast Cancer ($n = 4628$), Chicken Heart ($n = 1965$), and Mouse Brain (HybISS, $n = 1656$), where n denotes the number of cells in each dataset. Boxplots represent the median, first quartile and third quartile with error bars extending to the furthest value within 1.5 of the interquartile range. Statistical significance calculated using Mann-Whitney U-test with FDR correction. Significance levels are indicated as follows: n.s., not significant ($P \geq 0.05$); * ($P < 0.05$); ** ($P < 0.01$); *** ($P < 0.001$); and **** ($P < 0.0001$). Exact P -values can be found in Appendix Table S3.

higher spatial velocity confidence than all other methods (Appendix Fig. S9A). In direct comparison to the other models: Stereo-seq brain (0.9999 vs. 0.8576, DeepVelo; FDR = 4.406×10^{-248}), breast cancer (0.9998 vs. 0.9776, cellDancer; FDR = 2.576×10^{-74}), chicken heart (0.9999 vs. 0.8859, SIRV; FDR = 5.397×10^{-51}), and HybISS brain (0.9999 vs. 0.9188, cellDancer; FDR = 3.157×10^{-105}). veloAgent maintained superior performance across both measures, reinforcing its ability to generate accurate and spatially consistent velocity fields (Appendix Fig. S9). Collectively, these results demonstrate that by integrating spatial priors through ABM, veloAgent transforms RNA velocity estimation from a purely temporal abstraction into a spatially coherent model of tissue dynamics, enabling more accurate and interpretable studies of cellular transitions and their underlying gene programs.

veloAgent reveals lineage-specific driver genes and transcription factors

To determine whether the improved RNA velocity estimates from veloAgent translate into biologically meaningful gene-level insights, we leveraged CellRank as a downstream framework that operates on velocity fields (Lange et al, 2022). Unlike prior analyses where CellRank was applied to conventional RNA velocity, here it is driven by veloAgent's spatially and temporally coherent velocity vectors, allowing us to assess how enhanced spatial modeling improves downstream fate and driver-gene inference. For each dataset (HybISS mouse brain, chicken heart, and Stereo-seq mouse brain), we computed fate probabilities and ranked genes along pseudotime within each lineage. These ranked genes were visualized as pseudotime heatmaps, where expression was scaled and genes were ordered by the pseudotime at which smoothed expression peaked (Fig. 4A,D,G).

Across all datasets, veloAgent revealed coherent cascades of early regulators, mid-stage transition factors, and late differentiation markers that followed expected developmental orderings. In the HybISS mouse brain, veloAgent-derived pseudotime captured the sequential activation of regulators driving neuronal differentiation (Fig. 4A). Early *En1* expression aligned with progenitor-like states, consistent with its role in embryonic patterning, where loss of *En1* causes mid/hindbrain defects by E9 (Simon et al, 2005). *Dlk1* peaked at mid-pseudotime, mirroring pro-neural factors (*Ngn2*, *Mash1*, *NeuroD*) known to precede neuronal maturation (Surmacz et al, 2012). Late-stage genes such as *Lhx5* and *Klhl14* were upregulated in terminal neuronal populations; *Lhx5* regulates hippocampal neuron

differentiation and migration, while *Klhl14* directs corticospinal axon targeting (Miquelajàregui et al, 2015; Miquelajàregui et al, 2010; Pillai et al, 2007; Sahni et al, 2021; Zhao et al, 1999). *Pax8* is also a known critical transcription factor involved in mid-hindbrain development (Fig. 4C) (Giacco et al, 2025). These results demonstrate that veloAgent reconstructs the temporal hierarchy of transcriptional regulators underlying neurodevelopment.

In the chicken heart, veloAgent distinguished transient valve progenitors enriched for *DICER1* from terminal valve cells expressing ECM genes (*COL1A1/2*, *SPARC*) (Fig. 4D). This ordering aligns with reports that *DICER1* activity is essential for early valve development, whereas ECM expression supports structural maturation and stability of the valves later in development (Acharya et al, 2012; Deckx et al, 2019; Yan et al, 2022). *IFI30* was also identified, with previous studies linking it with cardiomyocytes (Fig. 4F) (Li et al, 2022). In the Stereo-seq mouse brain, veloAgent aligned early regulators such as *Cdk8*, mid-phase factors including *Quaking* (*Qk*), and late markers such as *Mbp* (Fig. 4G). *Cdk8*, a Mediator-complex component, is essential for early mouse development (Galbraith et al, 2010; Westerling et al, 2007), while *Qk* isoforms stabilize *p27^Kip1^* mRNA to promote oligodendrocyte maturation (Larocque et al, 2005), and *Mbp* marks terminal differentiation during myelination (Fig. 4I) (Landry et al, 1996). Together, these cascades highlight veloAgent's ability to delineate lineage-specific transcriptional programs across distinct spatial and developmental contexts.

To assess the biological relevance of the inferred driver genes, we performed Gene Ontology (GO) enrichment analysis on the top fate-associated genes for each lineage (Fig. 4B,E,H). Enriched pathways were strongly consistent with known developmental programs. In the chicken heart (Fig. 4E), enriched terms included circulatory-system development (FDR = 1.4×10^{-12} , 40 genes), anatomical-structure morphogenesis (FDR = 8.7×10^{-11} , 59 genes), and mesenchymal development (FDR = 4.9×10^{-11} , 21 genes), implicating *POSTN* and *COL1A1* as central regulators of cardiac morphogenesis and collagen signaling (Conway and Molkentin, 2008; Deng et al, 2025). In the HybISS mouse brain (Fig. 4B), enriched terms such as neuron differentiation (FDR = 1.9×10^{-37} , 57 genes), nervous-system development (FDR = 2.6×10^{-36} , 66 genes), and generation of neurons (FDR = 1.9×10^{-37} , 58 genes) confirmed neural-lineage specificity, featuring *Lhx5* and *Phox2b* as canonical regulators (Pattyn et al, 1999; Pillai et al, 2007). In the Stereo-seq brain (Fig. 4H), enriched pathways included nervous-system development and synaptic signaling, with *Qk* and *Mbp*




Figure 4. Improved driver gene and transcription factor analysis enabled by accurate RNA velocity inference from veloAgent.

Accurate velocity estimates from veloAgent enhance downstream analyses such as identification of cell-type-specific driver genes and regulatory transcription factors across multiple spatial and temporal datasets. (A) Heatmap showing normalized spliced gene expression ordered by pseudotime for driver genes identified by CellRank in neuron cell states of the HybISS mouse brain dataset. (B) Gene Ontology (GO) enrichment for neuron driver genes from the HybISS mouse brain dataset. (C) Spatial distribution of spliced gene expression for the top correlated driver gene, *Pax8* (top), and pseudotime projected onto the UMAP embedding (bottom). (D) Heatmap of normalized spliced gene expression ordered by pseudotime for driver genes identified by CellRank in valve cell states of the chicken heart development dataset. (E) GO enrichment analysis for valve driver genes in the chicken heart dataset. (F) Spatial distribution of spliced expression for the top correlated driver gene, *IFI30* (top), and pseudotime projection on the UMAP embedding (bottom). (G) Heatmap of normalized spliced gene expression ordered by pseudotime for driver genes identified by CellRank in fiber tract cell states of the Stereo-seq mouse brain dataset. (H) GO enrichment analysis for fiber tract driver genes. (I) Spatial distribution of spliced gene expression for the top correlated driver gene, *Mbp* (top), and pseudotime projection on the UMAP embedding (bottom). (J) Heatmaps of transcription factor activity scores derived from SCENIC analysis across cell states in the mouse brain (HybISS, left), mouse brain (Stereo-seq, middle), and mouse organogenesis (right) datasets.

emerging as key regulators of oligodendrocyte maturation and myelin formation (Landry et al, 1996; Larocque et al, 2005). These enrichment patterns confirm that the driver genes inferred from veloAgent-based velocity recapitulate lineage-regulatory programs across tissues.

To further examine upstream transcriptional regulation, we applied SCENIC (Aibar et al, 2017) to infer gene-regulatory networks and identify active transcription factors (TFs) across the HybISS, Stereo-seq, and mouse-organogenesis datasets (Fig. 4J). Because SCENIC's predictions depend directly on the quality of the input expression and velocity information, improvements in RNA velocity modeling should translate into clearer and more biologically coherent TF-activity patterns. The predicted TFs showed strong concordance with known lineage regulators. In the HybISS dataset, *Foxj1*, *Lmx1a*, *Msx1*, and *Msx2* were recovered, consistent with their established roles in choroid-plexus patterning (Liddel, 2015; Venugopalan et al, 2008) (Fig. 4J; left); *Lmx1a*, for instance, directs cells toward roof-plate and choroid-plexus lineages rather than neuronal fates originating from the rhombic lip (Chizhikov et al, 2010). In the Stereo-seq brain, *Foxn3*, *Nfia*, and *Tcf12* were predicted (Fig. 4J; middle); *Nfia* and *Tcf12* are established regulators of fiber-tract formation, while *Foxn3*, though not previously linked, may play an indirect regulatory role and represents a potential novel TF for further investigation (das Neves et al, 1999; Samaan et al, 2010; Singh et al, 2022). In the mouse-organogenesis dataset, *Foxf1* and *Gata4* were identified in association with the definitive endoderm and cardiomyocyte lineages, respectively (Fig. 4J; right); *Foxf1* mediates endoderm-mesoderm signaling, and *Gata4* is a key regulator of cardiac gene expression and differentiation (Madison et al, 2009; Malek Mohammadi et al, 2017).

These results demonstrate that veloAgent not only infers spatially and temporally consistent cell-state dynamics but also enables biologically validated and accurate identification of lineage-specific driver genes and transcription factors across diverse tissues.

veloAgent facilitates in silico perturbation of RNA velocity to reveal causal regulators and therapeutic targets


A capability absent in all existing RNA velocity inference methods—including scVelo, cellDancer, and DeepVelo—is the ability to perform *in silico perturbations* directly within the learned dynamical system. veloAgent fills this gap by explicitly parameterizing gene-specific transcription (α), splicing (β), and degradation (γ) kinetics, enabling mechanistic simulation of how gene-level

perturbations propagate through the cellular velocity field. By computationally silencing a gene's transcription rate, veloAgent emulates a virtual knockout and recomputes the resulting flow of cell states, predicting how the loss of that gene would redirect developmental or disease trajectories in real tissue space. This approach transforms RNA velocity from a descriptive framework into a causal, predictive system capable of identifying regulatory drivers and therapeutic targets.

Perturbations were performed by systematically setting each gene's transcription rate (α) to zero in the trained model, without retraining, and propagating the change through the dynamical equations to recalculate RNA velocity for all cells (Fig. 5A). Because the model learns gene-specific transcription, splicing, and degradation kinetics as coupled parameters, altering a post-training preserves the learned regulatory dependencies and propagates mechanistically through the velocity field rather than numerically rescaling expression. The resulting perturbed velocity fields were projected onto the UMAP embedding to visualize directional shifts. To quantify the magnitude and orientation of these shifts, we calculated the cross-boundary direction (CBD) score, which measures the change in flow vectors across boundaries separating progenitor and terminal states. A positive score indicates that silencing a gene pushes trajectories toward terminal fates, whereas a negative score reflects diversion away from them. Genes with high absolute CBD values thus act as key regulators of fate progression.

We first applied this framework to a developmental mouse brain dataset profiled by Stereo-seq. The unperturbed velocity field showed expected progression toward fiber-tract cells. Silencing top-ranked genes (Fig. 5B; top), identified by high |CBD| values, reversed this directionality, implying their causal roles in guiding differentiation toward the fiber-tract fate (Fig. 5C). Gene Ontology enrichment of the top perturbed genes revealed significant overrepresentation of neurogenesis, axonogenesis, and synaptic-organization pathways (Fig. 5D), confirming that the perturbation-derived regulators align with known developmental programs.

We next tested whether this causal-perturbation framework could uncover regulatory dependencies in disease. Using a Visium human breast cancer spatial-transcriptomics dataset, the unperturbed velocity field indicated trajectories converging toward malignant clusters. Perturbing top-scoring genes markedly redirected the velocity vectors away from these clusters (Fig. 5E; Appendix Fig. S10), suggesting suppression of malignant progression. The most strongly perturbed genes (Fig. 5B, bottom) included *CARD14* and *ERBB2* (Dumbrava et al, 2019; Lim et al, 2020)—well-established oncogenes whose inhibition reversed trajectory



directionality. Spatial maps confirmed their enrichment within malignant cell populations (Fig. 5F). Kaplan–Meier survival analyses showed that patients with high expression of these genes experienced significantly poorer overall survival (Fig. 5G,H).

Notably, *ERBB2* is a validated therapeutic target of multiple FDA-approved drugs, including trastuzumab and trastuzumab deruxtecan (Modi et al, 2020; Swain et al, 2023), underscoring the translational fidelity of our in silico predictions.

Figure 5. In-silico perturbation analysis enabled by veloAgent identifies potential therapeutic targets.

veloAgent enables in silico perturbation by systematically silencing transcription to assess the impact of gene-specific regulatory loss on cellular dynamics. (A) Schematic of the perturbation pipeline, where the transcription rate (α) of each gene is individually set to zero to simulate gene knockout or drug-induced repression, followed by RNA velocity recalculation. (B) Top 5% of perturbed genes ranked by the change in cross-boundary score of the velocity following perturbation in mouse brain developmental data (top) and human breast cancer data (bottom). (C) Comparison of perturbed versus original velocity projections in the developmental dataset, with red circles highlighting regions of substantial deviation. (D) GO enrichment analysis of the top 5% most perturbed genes, revealing biological processes associated with developmental regulation. (E) Velocity projection comparison in the breast cancer dataset showing notable shifts following transcriptional perturbation. (F) UMAP projections showing expression of candidate therapeutic targets, *CARD14* and *ERBB2*, each enriched in cancer-associated cell populations. (G) Survival analysis comparing patient outcomes between high- and low-expression groups for the combined set of significantly perturbed genes in the breast cancer dataset. (H) Individual Kaplan-Meier survival analyses for *CARD14* and *ERBB2*, demonstrating prognostic relevance based on expression level.

To assess whether veloAgent's perturbation framework generalizes beyond transcriptional regulation, we further simulated perturbations of the splicing rate (β) for each gene. These β -perturbations altered velocity directionality in a manner consistent with post-transcriptional control, typically inducing smoother and temporally delayed trajectory shifts relative to α -perturbations (Appendix Fig. S11). This result demonstrates that veloAgent can model both transcriptional and splicing-driven dynamics within the same mechanistic framework, further validating its capacity for causal simulation.

By enabling gene-level kinetic perturbations within a learned velocity model, veloAgent provides a scalable, mechanistic, and hypothesis-driven alternative to physical perturbation assays. Together, these analyses establish veloAgent as a bridge between mechanistic modeling and functional discovery: by linking gene-specific kinetic parameters to emergent changes in cellular dynamics, veloAgent enables causal inference of regulatory genes, nominates candidate therapeutic targets such as *CARD14* and *ERBB2*, and offers a generalizable computational platform for virtual perturbation experiments across developmental and disease contexts.

veloAgent is scalable and computationally robust

veloAgent's architecture, which scales with the number of genes rather than the number of cells, provides substantial gains in computational and memory efficiency compared to existing RNA velocity models. This design choice is critical because the number of genes in each organism is essentially fixed, whereas single-cell datasets continue to expand exponentially with improvements in sequencing throughput. In contrast, most frameworks—including cellDancer—scale with cell count, leading to rapidly increasing memory and runtime requirements that limit their applicability for atlas-level analyses involving hundreds of thousands or millions of cells.

To systematically benchmark scalability, we progressively duplicated the HybISS mouse brain dataset to generate datasets ranging from 50,000 to one million cells. While cellDancer's memory usage scaled near cubically, reaching ~120 GB at one million cells, veloAgent exhibited near-linear scaling, requiring only ~37 GB under identical conditions (Fig. 6A; Appendix Fig. S12). Similarly, veloAgent completed velocity inference substantially faster, maintaining nearly constant per-cell runtime as dataset size increased (Fig. 6B; Appendix Fig. S13). cellDancer was selected as the primary comparator because it is the most competitive deep-learning-based RNA velocity framework and offers a direct architectural contrast to veloAgent's gene-scaled design. Other frameworks such as scVelo were excluded from this

analysis because their smaller network architectures preclude meaningful efficiency comparisons, though they were included in biological benchmarking elsewhere in this study. Full experimental and hardware details are provided in the Methods section.

Beyond computational scalability, veloAgent enhances biological robustness through its agent-based modeling (ABM) module, which integrates spatial context to refine velocity fields. The ABM takes temporally inferred velocity estimates and couples them with spatial coordinates of neighboring cells to produce spatially coherent velocity directions. Notably, this module is fully modular and can be applied to improve the outputs of other RNA velocity methods without retraining. To directly assess the contribution of spatial information, we applied the ABM module to velocity embeddings generated by existing RNA velocity methods, incorporating spatial priors post hoc. Across all datasets, ABM integration markedly increased median velocity confidence for all methods (Fig. 6C). For example, relative to scVelo adding the ABM module led to the following performance improvements (Fig. 6C): mouse brain (Stereo-seq, 0.9943 vs. 0.0677; FDR $< 1.0 \times 10^{-308}$), breast cancer (0.6920 vs. 0.0929; FDR $< 1.0 \times 10^{-308}$), chicken heart (0.9959 vs. 0.3550; FDR $< 1.0 \times 10^{-308}$), and mouse brain (HybISS, 0.9794 vs. 0.3166; FDR $< 1.0 \times 10^{-308}$). These consistent gains demonstrate that the ABM functions as a general-purpose enhancement, improving spatial coherence while preserving veloAgent's computational efficiency.

To evaluate model robustness, we conducted ablation experiments on spatial transcriptomics datasets using four complementary metrics: cross-boundary direction, fate probability, velocity confidence, and in-cluster coherence. Removing any of veloAgent's key components—the VAE encoder (Phase 1), gene-gene interaction network (Phase 2), neighborhood adjustment (Phase 3), or ABM refinement (Phase 4)—significantly reduced performance, whereas the complete model consistently achieved the highest accuracy (Fig. 6D–F; Appendix Figs. S14 and S15). Removal of the ABM resulted in the largest performance reduction across spatial transcriptomics datasets, underscoring that the observed performance of veloAgent is primarily driven by incorporation of spatial priors rather than solely by the VAE or underlying gene-gene neural network (Fig. 6D–F). Ablation analyses on scRNA-seq-only datasets (without ABM) yielded consistent results (Appendix Fig. S16).

Discussion

In this study, we present veloAgent, a deep generative and agent-based model (ABM) for spatial RNA velocity estimation that incorporates




Figure 6. Ablation study and computational efficiency of veloAgent.

(A) Comparison between veloAgent and cellDancer regarding peak CPU memory usage (top) for increasing number of cells subsampled from HybISS mouse brain dataset. (B) Comparison of computational runtime on CPU for increasing number of cells, subsampled from HybISS mouse brain dataset between veloAgent and cellDancer. (C) Comparative evaluation of velocity confidence scores before and after applying veloAgent's agent-based model (ABM) for spatial refinement on top of each method's temporal velocity estimates. Original estimates from each method are contrasted with their spatially adjusted counterparts, showing consistent improvement when spatial data is incorporated. (D–F) To assess the contribution of key model components, an ablation study was performed by systematically removing each module from veloAgent, including the ABM module. (D) Cross-boundary accuracy, (E) velocity confidence and (F) fate probability metrics across four configurations: full veloAgent, without latent-space smoothing, without gene-informed network architecture, without neighbor adjustment and without incorporating spatial transcriptomics data using agent-based modeling. Performance drops observed in ablated models validate the importance of each component. Boxplots in (C–F) represent the median, first quartile and third quartile with error bars extending to the furthest value within 1.5 of the interquartile range (IQR). Each panel displays four boxplots for Mouse Brain (Stereo-seq, $n = 7765$), Breast Cancer ($n = 4628$), Chicken Heart ($n = 1965$), and Mouse Brain (HybISS, $n = 1656$), where n denotes the number of cells in each dataset. Statistical significance calculated using Mann-Whitney U-test with FDR correction. Significance levels are indicated as follows: n.s., not significant ($P \geq 0.05$); * ($P < 0.05$); ** ($P < 0.01$); *** ($P < 0.001$); and **** ($P < 0.0001$). Exact P -values can be found in Appendix Table S4.

spatial transcriptomics data. Most existing RNA velocity frameworks treat cells as independent points in expression space, neglecting the tissue architecture and interactions between cells that influence cell state transitions. veloAgent bridges this gap by combining a VAE for robust latent representation, a gene-regulatory-aware neural network that incorporates known molecular interactions, and an ABM that simulates each cell's spatial microenvironment. This integration ensures that velocity estimates are not only informed by transcriptional dynamics but also constrained by spatial neighborhood structure, improving both accuracy and interpretability. Applied to developmental and disease datasets, veloAgent outperforms existing methods across quantitative benchmarks and captures literature-supported marker gene dynamics in their correct spatial contexts. Beyond trajectory inference, veloAgent's *in silico* perturbation module predicts how changes in gene regulation or signaling alter cell fate, enabling targeted hypothesis generation in developmental biology oncology and regenerative medicine.

veloAgent offers several key advances over existing RNA velocity frameworks. First, it is designed for scalability through a gene-centric architecture, where model complexity scale with the number of genes—typically fixed for an organism—rather than the number of cells. This contrasts with cell-centric methods such as cellDancer, whose complexity grows with cell count, making large-scale analysis computationally expensive. By scaling with genes instead of cells, veloAgent achieves sublinear memory usage and can handle atlas-scale datasets with millions of cells. This is particularly important given the recent rise in single-cell experimental throughput, which is expected to continue growing. Unlike models that rely on global assumptions or cell-centric designs, our approach generalizes across tissues and experimental platforms without compromising interpretability or performance (Bergen et al, 2021). Second, veloAgent uniquely integrates spatial transcriptomics data into the velocity inference process through its ABM module. In the ABM, individual cells are treated as agents and their velocities are refined using predefined local interaction rules that incorporate neighboring cells, spatial proximity, and transcriptional similarity. Although this module is not a full multi-step agent-based simulation framework, it preserves the central ABM principle of updating individual cellular states through local interactions. Spatial context is often overlooked in RNA velocity frameworks, yet it plays a fundamental role in shaping lineage boundaries, localized gene expression patterns, and environmental signals that guide cell behavior (Feng et al, 2025; Ru et al, 2023). By simulating the spatial microenvironments of each cell, veloAgent embeds tissue architecture into its RNA velocity estimations,

ensuring that transcriptionally similar and spatially proximal cells are inferred to follow coherent trajectories. In contrast, existing approaches either omit spatial information entirely or use it only for post hoc projection (e.g., SIRV), which can visualize velocities on a tissue map but cannot influence the underlying estimation. veloAgent unifies both, incorporating spatial constraints during velocity inference and projecting the resulting vectors back into physical space to reveal region-specific linear trajectories and niche-dependent differentiation. Third, veloAgent's modular ABM design allows it to function as a universal booster, seamlessly integrating with existing RNA velocity frameworks. By coupling their outputs with spatial context, the ABM refines and harmonizes velocity fields across tissue regions, substantially enhancing spatial coherence and biological interpretability without retraining of the original models. Finally, beyond trajectory inference, veloAgent introduces an *in silico* perturbation framework that allows targeted manipulation of RNA velocity vectors to simulate the effects of gene regulatory or signaling interventions. *In silico* perturbation has been employed in other areas in single-cell genomics outside of RNA velocity (Jolasun et al, 2025; Raghavan et al, 2025; Zhao et al, 2025; Zheng et al, 2025). Our *in silico* perturbation enables rational cell state steering and identification of candidate therapeutic targets. As a case study, we show that veloAgent can recover key breast cancer linked genes whose perturbation could reverse disease trajectories, including *ERBB2*, which has existing FDA-approved inhibitors. This opens new avenues for therapeutic discovery, such as identifying candidate drug targets or designing intervention strategies for regenerative medicine and cell reprogramming.

While veloAgent offers significant advancements, it also faces limitations that warrant future investigation. Since Visium spots may capture transcripts from multiple cells, inferred dynamics represent averaged transcriptional states. However, consistent results across higher-resolution platforms such as Stereo-seq and HybISS suggest that veloAgent captures biologically meaningful transition signals that are robust to spatial granularity. Secondly, the current perturbation framework models single gene interventions for simplicity and does not capture the complexity of combinatorial perturbations or the interplay of multiple regulatory pathways. Additionally, integrating other omics modalities such as ATAC-seq and proteomics could further enhance both velocity estimation as well as perturbation modeling by providing a more complete view of cell state and regulatory control. Another limitation arises from incomplete prior knowledge in interaction databases such as STRING. For genes without annotated interactions, we connect them to all other gene nodes in the gene–gene

neural network, allowing the model to learn potential regulatory relationships from the data. While this ensures that such genes are not excluded from the model, it may introduce biologically implausible connections and additional noise. Nevertheless, this design represents a practical compromise given the incompleteness of current interaction databases and is consistent with approaches used in prior work that combine biological priors with data-driven learning. A further limitation is that veloAgent, like other RNA velocity methods, operates at the gene level and does not capture exon-level transcriptional dynamics. Incorporating exon-level quantification, as explored in DOLPHIN (Song et al, 2025), could enable more precise modeling of transcriptional regulation and RNA velocity estimates. Finally, while veloAgent is designed for scalability, it has not yet been applied to a true atlas dataset with millions of cells, representing an important step to fully demonstrate its large-scale applicability.

veloAgent defines a scalable, spatially aware, and biologically interpretable framework that extends RNA-velocity analysis into previously inaccessible domains. Its scalability supports emerging cell-atlas projects, and its perturbation module bridges computational modeling with experimental design. The framework provides a foundation for predictive, intervention-ready models of cellular behavior, with the potential to accelerate discovery, advance developmental understanding, and guide therapeutic innovation.

Methods

Reagents and tools table

Reagent/ Resource	Reference or Source	Identifier or Catalog Number
Experimental models		
N/A	N/A	N/A
Recombinant DNA		
N/A	N/A	N/A
Antibodies		
N/A	N/A	N/A
Oligonucleotides and other sequence-based reagents		
N/A	N/A	N/A
Chemicals, Enzymes and other reagents		
N/A	N/A	N/A
Software		
Python v3.8.19		
cellDancer	https://github.com/ GuangyuWangLab2021/cellDancer	
scVelo v0.3.2	https://github.com/theislab/scvelo	
DeepVelo v0.2.5	https://github.com/bowang-lab/ DeepVelo	
VeloVI v0.3.1	https://github.com/YosefLab/velovi	
VeloVAE	https://github.com/welch-lab/velovae	
UniTVelo v0.2.4	https://github.com/StatBiomed/ UniTVelo	

Reagent/ Resource	Reference or Source	Identifier or Catalog Number
VeloAE v0.2.0	https://github.com/qiaochen/VeloAE	
CellRank v2.0.5	https://github.com/scverse/cellrank	
SCENIC v0.12.1	https://github.com/aertslab/SCENIC	
Scanpy v1.9.8	https://github.com/scverse/scanpy	
Other		

Single-cell sequencing data preprocessing

Spliced and unspliced count matrices were obtained using the standard velocity pipeline (La Manno et al, 2018). For datasets in which these matrices were not already available, we generated them using velocity with standard/default settings; otherwise, we used the publicly available processed matrices directly.

We processed the single-cell RNA-sequencing data using the scVelo pipeline. We began by filtering out cells with less than 20 genes. We then selected the top 2000 highly variable genes and normalized against the library size and transformed into log space. We then computed a nearest-neighbor graph (with 30 neighbors) in principal component analysis space (30 principal components). Finally, we calculated the first-order moment of spliced and unspliced counts for subsequent velocity estimation steps.

We used spatial transcriptomics datasets from multiple technologies, including Visium, Stereo-seq, SeqFISH+, and HybISS. We note that Visium’s poly-A capture chemistry may lead to reduced intronic capture efficiency compared to single-cell RNA-sequencing, potentially resulting in increased sparsity in unspliced counts. In veloAgent, this limitation is partially alleviated through a combination of latent embedding and the agent-based modeling (ABM) refinement module. The latent representation provides a denoised and compact encoding of transcriptional states, while the ABM module aggregates information from neighboring cells based on spatial proximity and transcriptional similarity, thereby stabilizing velocity estimation under sparse conditions.

veloAgent framework

The veloAgent model consists of three sub-components, a variational autoencoder (VAE), a deep neural network, and an agent-based model (ABM).

Phase 1: Variational autoencoder

veloAgent uses a VAE designed to represent high-dimensional single-cell gene expression data within a lower-dimensional embedding space. Specifically, our VAE incorporates two separate encoders, each sharing an identical structure composed of three fully connected hidden layers. These hidden layers progressively decrease in dimension, evenly spaced in thirds, ultimately reducing the input to the latent dimension Z . Each encoder, q_θ , independently process the spliced (q_θ^S) and unspliced (q_θ^U) gene expression,

transforming them into distinct low-dimensional embeddings. The outputs from both encoders are concatenated to form a unified representation, $h = \text{concat}(q_\theta^S(\mathbf{X}_S), q_\theta^U(\mathbf{X}_U))$. Two MLPs are then employed to determine the mean vectors, $\mu_z = f_{\mu_\theta}(\mu_z|h)$, and standard deviation vectors $\sigma_z = f_{\sigma_\theta}(\sigma_z|h)$ of the latent representation, assuming a multivariate normal distribution as a prior, $\mathbf{Z} \sim \mathcal{N}(\mu_z, \sigma_z^2)$. The approximated posterior distribution is represented as $q_\theta(\mathbf{Z}|\mathbf{X}_S, \mathbf{X}_U)$. The latent representation is particularly valuable, as they are designed to preserve critical biological information while reducing noise and less relevant features. This enables more coherent and accurate estimations of single-cell velocities. Two corresponding decoders (p_ψ^S, p_ψ^U) mirror the encoders architectures reconstructing the original cell-by-gene matrices for splices and unspliced counts from $\mathbf{Z}$. The VAE loss function is given by the following equation:

$$\mathcal{L}_{VAE} = D_{KL}(q_\theta(\mathbf{Z}|\mathbf{X}_S, \mathbf{X}_U)||p(\mathbf{Z})) - \mathbb{E}_{q_\theta(\mathbf{Z}|\mathbf{X}_S, \mathbf{X}_U)} [\log p_\phi^S(\mathbf{X}_S|\mathbf{Z}) + \log p_\psi^U(\mathbf{X}_U|\mathbf{Z})] \quad (1)$$

Phase 2: Velocity rate estimation using deep neural network

The transcriptional kinetics of an individual gene, g , are modeled by a system of two ordinary differential equations:

$$\frac{d\mathbf{u}_{i,g}(t)}{dt} = \alpha_{i,g}(t) - \beta_{i,g}(t) \cdot \mathbf{u}_{i,g}(t) \quad (2)$$

$$\frac{d\mathbf{s}_{i,g}(t)}{dt} = \beta_{i,g}(t) \cdot \mathbf{u}_{i,g}(t) - \gamma_{i,g}(t) \cdot \mathbf{s}_{i,g}(t) \quad (3)$$

where $\mathbf{u}_{i,g}(t)$ and $\mathbf{s}_{i,g}(t)$ represent the concentrations of the premature and mature mRNAs for cell i , gene g at time t and $\alpha_{i,g}$, $\beta_{i,g}$, $\gamma_{i,g}$ represent the transcription, splicing and degradation rates at time t , respectively. Previous methods assume that parameters α , β , and γ are either constant or universally shared across all genes and cells. However, this assumption inadequately models the continuously evolving and heterogeneous cell subpopulations.

To accurately estimate α , β , and γ , we utilize a neural network that takes as input the latent representation $\mathbf{Z}$ generated by the VAE in phase 1. The first hidden layer has dimensionality equal to the number of genes (G) and connects the latent representation to gene-level feature through dense fully connected weights. The second hidden layer is also of size G , but instead of dense connections as in a standard multilayer perceptron, it incorporates sparse connections informed by gene-gene interaction data from the STRING database (Szklarczyk et al, 2023). Specifically, when a corresponding interaction exists between two genes in the STRING, a trainable connection is introduced between their corresponding nodes in the model. We incorporate all available interactions from STRING to maximize the use of prior biological knowledge and do not filter edges based on confidence score, as this also led to better performance than restricting the model to only high-confidence interactions (Appendix Fig. S17B,C). Genes without known connections are connected to all other gene nodes, resulting in a fully connected node for that gene. This design ensures that genes lacking documented interactions are not excluded from the model and allows the network to learn potential regulatory relationships directly from the data (Raghavan et al, 2025).

We found that the majority of genes in the datasets were also present in STRING (Appendix Fig. S17A). Following this biologically informed layer, the network proceeds through two additional fully connected layers, expanding first to $2G$ and then to $3G$, where G is the number of genes. The final layer is structured such that the first G units correspond to transcription rate (α), the second G units correspond to splicing rate (β), and the third G units correspond to degradation rate (γ). The output provides independent rate parameters for each gene in every cell, resulting in cell and gene-specific estimates of kinetic rates.

The neural network is trained using a loss function based on cosine similarity to compare the predicted RNA velocity with observed velocity vectors, similar to in cellDancer (Li et al, 2024). Given the outputs α , β , and γ , the predicted spliced ($\hat{\mathbf{s}}$) and unspliced ($\hat{\mathbf{u}}$) counts are obtained using the ordinary differential equations. The velocity vector for each cell is defined as the change from its observed counts to its predicted counts:

$$\mathbf{v}_i = (\hat{\mathbf{u}}_i(t) - \mathbf{u}_i(t), \hat{\mathbf{s}}_i(t) - \mathbf{s}_i(t)), \quad (4)$$

$$\mathbf{v}'_i = (\mathbf{u}_i(t) - \mathbf{u}_i(t), \mathbf{s}_i(t) - \mathbf{s}_i(t)) \quad (5)$$

Where $\{i'\}$ is the collection of neighboring cells of cell i ; and i' is the neighboring cell that minimizes the cosine similarity loss function. To assess the quality of the predicted velocity, we compute the cosine similarity between the velocity of cell i , $\mathbf{v}_i$, and the velocity vectors of neighboring cells:

$$\mathcal{L}_{NN} = \sum_{i=1}^N [1 - \max_{j \in \mathcal{N}(i)} \frac{\mathbf{v}_i \cdot \mathbf{v}'_i}{\|\mathbf{v}_i\| \|\mathbf{v}'_i\|}] \quad (6)$$

where N is the number of cells and $\mathcal{N}(i)$ is the neighborhood of cell i , and j is a neighboring cell of i . We take the maximum cosine similarity between the cell's predicted velocity and the velocity of all its neighbors.

Phase 3: Cell velocity adjustment

In this phase, we refine each cell's predicted velocity to produce a more consistent and confident estimate by leveraging information from neighboring cells. Specifically, we adjust each cell's velocity to align with those of its neighbors by minimizing the cosine similarity between their velocity vectors. This alignment step enhances the stability and reliability of the velocity estimates by encouraging local consistency in the velocity field. The underlying assumption is that cells within a local neighborhood in gene expression space tend to follow similar dynamic patterns. The loss function used for this alignment is formally defined below, where N is the number of cells, $\mathcal{N}(i)$ is the neighborhood of cell i , and $j \in \mathcal{N}(i)$ is a neighboring cell of i :

$$\mathcal{L}_{NBR} = - \sum_{i=1}^N \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \frac{\mathbf{v}_i \cdot \mathbf{v}_j}{\|\mathbf{v}_i\| \|\mathbf{v}_j\|} \quad (7)$$

The overall loss function of veloAgent combines three components: the VAE reconstruction loss $\mathcal{L}_{VAE}$, the rate prediction loss from the neural network $\mathcal{L}_{NN}$, and the neighborhood coherence loss $\mathcal{L}_{NBR}$. Together, these form the total training

objective of the model, expressed as:

$$\mathcal{L} = \mathcal{L}_{VAE} + \mathcal{L}_{NN} + \tau \mathcal{L}_{NBR} \quad (8)$$

Agent-based model for spatial RNA velocity modeling

To incorporate spatial context into RNA velocity estimation, we developed an agent-based module in which each cell is treated as an agent and its velocity is refined using predefined local interaction rules based on neighboring cells within a defined spatial radius. Although this module is not a full multi-step agent-based simulation framework, it follows the core principle of agent-based modeling by updating individual cellular states through local interactions. ABMs have previously been used to model cellular dynamics in single-cell and spatial transcriptomics data, capturing complex interactions among cells in a spatially aware manner (Raghavan et al, 2025). The model refines velocity estimates based on three key factors: local cell density, inverse distance, and RNA expression similarity. To construct the ABM, we first define the states and behaviors of the agent, where each agent represents a cell in the biological system. Agent states capture the properties and dynamic features of each cell, while agent rules define how each agent behaves.

Cell agent state description. Since our primary goal is to study RNA velocity in the context of spatially organized tissues, each cell agent is characterized by features relevant of both its transcriptional state and its spatial interactions. These include gene expression profiles, RNA velocity vectors, and spatial coordinates. We also include metadata such as cell type or tissue annotations, batch ID etc.

RNA velocity adjustment behavior rules. The agent behaviors are defined by rules governing how each cell's velocity is modulated by its neighbors. For a given cell i , the neighborhood $N(i)$ comprises all cells located within a specified radius r . The updated velocity of cell i is computed as a weighted combination of its own velocity and those of its neighbors, reflecting three biologically motivated influence terms. First, neighboring cells with similar expression profiles are expected to exhibit similar transcriptional dynamics. This relationship is captured by a similarity weight (**SW**) defined as

$$\mathbf{SW}_{ij} = \sigma_{ABM} \cdot \frac{\text{cosine}(\mathbf{x}_i, \mathbf{x}_j)}{\sum_{k \in N(i)} \text{cosine}(\mathbf{x}_i, \mathbf{x}_k)} \quad (9)$$

Here, σ_{ABM} is a scaling constant and $\mathbf{x}_i$ is the expression profile of cell i .

Second, denser local environments exert greater collective influence on a cell's velocity, quantified by a density weight (**DW**)

$$\mathbf{DW}_i = \frac{\rho_i - \min(\rho)}{\max(\rho) - \min(\rho)}, \quad (10)$$

where ρ_i denotes the local cell density for cell i , computed as the number of neighboring cells within the predefined radius. We observed that not including the density weight reduced model performance (Appendix Fig. S18).

Third, spatial proximity modulates intercellular influence, as nearby cells are more likely to affect one another through diffusion-limited molecular signaling that alters gene expression. This effect is modeled using spatial inverse-distance weighting (**SpW**), assigning

stronger influence on closer neighbors.

$$\mathbf{SpW}_{ij} = (1 - \sigma_{ABM}) \cdot \frac{1/d_{ij}^{\tau_{\text{spatial}}}}{\sum_{k \in N(i)} 1/d_{ik}^{\tau_{\text{spatial}}}}, \quad (11)$$

where d_{ij} is the Euclidean distance between cells i and j , and $\tau_{\text{spatial}} = 2$, for 2D data.

The refined velocity ($\mathbf{RV}'$) for each cell i is then obtained by combining its intrinsic and neighborhood influences according to

$$\mathbf{RV}'_i = \mathbf{DW}_i \cdot \mathbf{RV}_i + \sum_{j \in N(i)} (\mathbf{SpW}_{ij} \cdot \mathbf{SW}_{ij}) \mathbf{RV}_j \quad (12)$$

$\mathbf{RV}_i$ is the velocity of cell i after phase 3. This formulation ensures that transcriptionally similar and spatially proximal cells exhibit coherent directional dynamics, while preserving biological heterogeneity across distinct tissue regions. Through this spatial refinement, the ABM transforms RNA velocity into a tissue-consistent and biologically interpretable model of cellular state transitions.

In silico perturbation

To evaluate the contribution of individual genes to RNA velocity and trajectory inference, we implemented an in silico perturbation framework within veloAgent. Perturbation is applied to either the transcription rate parameter (α) or the splicing rate parameter (β) for a given gene, but not both simultaneously. In the main analyses (Fig. 5) presented in this study, perturbations were performed by modifying α , whereas additional experiments modifying β are reported in Appendix Fig. S11. For each gene, a virtual loss-of-function was simulated by setting its transcription rate parameter (α) or splicing rate parameter (β) to zero for all cells, effectively removing its transcriptional or splicing influence from the learned kinetic system. When α is modified, the remaining kinetic parameters (β and γ) remain fixed, and similarly, when β is perturbed α and γ remain unchanged. The model was then used to recompute RNA velocity estimates ($v^{(-g)}$), allowing direct assessment of how the perturbation altered predicted cellular dynamics. Since perturbation is applied after the gene–gene interaction network has already inferred the kinetic parameters α , β , and γ , the neural network is not retrained during perturbation, and all learned model parameters remain unchanged.

The impact of each perturbation was quantified using the cross-boundary direction (CBD) metric. This metric quantifies how well velocity vectors maintain the correct directionality between transitioning cell types (e.g., cell type A to cell type B), reflecting the gene's influence (I_g) on the inferred dynamic processes $\text{CBD}(\mathbf{v}, d(A, B))$. The change in CBD following perturbation, $\text{CBD}(v^{(-g)}, d(B, A))$, reflects the influence of gene g on the inferred dynamic process. A gene-level importance score was computed as

$$I_g = \frac{1}{N} \sum_{i=1}^N \text{CBD}(\mathbf{v}_i^{(-g)}, d(B, A)) \quad (13)$$

where N denotes the total number of cells. This score represents the extent to which perturbing a specific gene alters trajectory

directionality, thus identifying regulators that drive or constrain lineage progression. Statistical significance was determined using permutation testing to identify genes whose perturbation caused a significant disruption in velocity alignment.

To assess the biological relevance of these computational perturbations, we performed Reactome pathway enrichment analysis for the mouse brain developmental dataset, confirming that top-ranked perturbed genes were enriched for neurogenesis- and differentiation-related pathways. For the breast cancer dataset, we validated candidate therapeutic genes by examining their association with patient survival using Kaplan–Meier (KM) analysis, linking predicted regulatory effects to clinical outcomes.

Benchmarking

We benchmarked veloAgent against seven state-of-the-art RNA velocity frameworks encompassing both conventional and spatially aware approaches. Among these, SIRV is the only method that also uses spatial transcriptomics data for spatial velocity projection. Comparisons were performed on scRNA-seq datasets for all methods except SIRV and on spatial transcriptomics datasets for all methods.

Performance was evaluated using three established metrics: velocity confidence, cross-boundary direction, and fate probability accuracy. Velocity confidence measures the local cosine similarity between a cell's velocity vector and those of its neighbors, reflecting the global coherence and reliability of the inferred velocity field. Cross-boundary direction assesses whether velocity vectors point in the biologically expected direction across annotated lineage transitions, quantifying directional accuracy at cell-type boundaries. Fate probability accuracy evaluates how well each method captures the commitment of cells toward terminal states based on their inferred RNA velocities. Fate probabilities were computed by constructing velocity-based transition matrices and quantifying the agreement between inferred and known terminal states.

To specifically assess spatial performance, we extended velocity confidence and fate probability into spatially aware variants, termed spatial velocity confidence (Qiao and Huang, 2021) and spatial fate probability (Lange et al, 2022). These metrics compute coherence and fate commitment using spatial neighborhoods rather than transcriptomic nearest neighbors, thereby evaluating whether inferred velocities are consistent with tissue architecture. This extension enables direct quantification of spatial accuracy beyond gene expression similarity alone.

Ablation

To evaluate the contribution of each major component in the veloAgent framework, we conducted systematic ablation experiments in which individual modules were selectively removed or simplified. Four variants were tested: the VAE, the gene–gene interaction network, the neighbor refinement mechanism, and the agent-based model (ABM). In the veloAgent-VAE variant, the learned VAE was replaced with a fixed principal component analysis (PCA) projection of the spliced and unspliced counts, eliminating the ability to learn a nonlinear latent space that captures complex cellular structure. In the veloAgent-GeneNet variant, the biologically informed gene–gene interaction layer was substituted with a fully connected layer, removing the inductive bias derived from known molecular relationships and thereby

reducing the biological interpretability of the learned dependencies. The veloAgent-NBR variant ablated the neighborhood refinement step, which aligns velocity vectors across transcriptionally similar cells to ensure local directional consistency. The veloAgent-ABM variant excluded spatial refinement, preventing the model from leveraging spatial transcriptomics data to enhance tissue-level coherence and directional accuracy. All ablation variants were evaluated using the same three metrics—cross-boundary correctness, velocity confidence, and fate probability—to quantify the contribution of each component.

Driver gene analysis

To identify key driver genes underlying cellular differentiation trajectories, we performed a pseudotime analysis using RNA velocity-based ordering. After computing the velocity graph, we derived pseudotime estimates for each cell using `scv.tl.velocity_pseudotime`. We then visualized the expression dynamics of genes of interest (GOIs) across pseudotime using a heatmap. The selected GOIs were identified from CellRank (Lange et al, 2022) lineage analysis as putative drivers of specific fate transitions. The heatmap was sorted by velocity pseudotime to capture temporal expression trends, highlighting the sequential activation and repression of lineage-associated genes along differentiation paths.

To further validate the relevance of lineage-driving genes identified through CellRank, we extracted all genes significantly associated with a specific terminal fate based on their correlation with fate probabilities. We then performed Gene Ontology (GO) enrichment analysis on this set of genes to evaluate the biological processes and pathways most enriched within each fate, thereby confirming the functional relevance of the identified gene programs.

Transcription factor analysis

To investigate the transcriptional regulators associated with terminal fates, we performed transcription factor (TF) inference using SCENIC (Aibar et al, 2017) pipeline. Transcription factor–target gene relationships were first inferred from gene co-expression patterns and then refined by assessing the enrichment of transcription factor binding motifs in the regulatory regions of target genes, using a curated database of mouse DNA motifs derived from experimental evidence. Regulon activity was quantified using AUCell and summarized by cell type to reveal distinct regulatory programs associated with specific cellular identities.

Hyperparameters

All hyperparameters used were fixed across datasets. The VAE was trained using Adam optimizer with a learning rate of 1×10^{-2} for 1000 iterations (or when early stopping was reached) and a latent embedding dimension of 5. The neural network for kinetic parameter estimation was trained using Adam optimization with a learning rate of 1×10^{-4} for 500 iterations (or when early stopping is reached). The number of nearest neighbors for the neural network and the neighborhood refinement was set to 30. The neighborhood velocity refinement module was run for 50 iterations. τ which controls the strength of the regularization of the neighborhood refinement loss is by default set to 0.2, as this value gave the best performance (Appendix Fig. S19). For the ABM.

Spatial velocity refinement was performed over 100 update steps. The spatial interaction radius r was set to a value for each dataset that led to each cell having roughly between 15 and 20 neighbours on average. The transcriptional similarity weight (SW) is controlled by σ_{ABM} which was by default set to 0.7 as this value gave the best performance (Appendix Fig. S20). The degree of spatial freedom τ_{spatial} was set to 2, since the spatial transcriptomics data was derived from 2-dimensional slices.

Computer configurations

Computational experiments were conducted on a high-performance computing server running Ubuntu 20.04.6 LTS with Linux kernel 5.4.0-216-generic (x86_64). The system was equipped with two Intel® Xeon® Platinum 8160 CPUs (96 logical cores, 2.10 GHz), 500 GB of RAM, and 8 GPUs (four NVIDIA Tesla M10 GPUs, 8 GB each; and four NVIDIA A16 GPUs, 15 GB each).

Data availability

All scRNA-seq data in this study were downloaded publicly. For the pancreatic endocrinogenesis and erythroid lineage of the mouse gastrulation datasets, we retrieve them from the Bergen et al in the scVelo study (Bergen et al, 2020). The mouse hippocampal dentate gyrus neurogenesis data was obtained from the study by La Manno et al (2018). The HybISS developing mouse brain, developing chicken heart, and SeqFISH mouse organogenesis spatial datasets were taken from the SIRV study (Abdelaal et al, 2021). The Stereo-seq mouse brain data was obtained from STOmicsDB publicly available dataset collection from the study by Chen et al (2021). The human breast cancer data was obtained from 10X Genomics published datasets <https://www.10xgenomics.com/datasets/fresh-frozen-visium-on-cytassist-human-breast-cancer-probe-based-whole-transcriptome-profiling-2-standard>. veloAgent is implemented in Python and available at <https://github.com/mcgilldinglab/veloAgent>.

The source data of this paper are collected in the following database record: [biostudies:S-SCDT-10_1038-S44320-026-00213-w](https://biostudies.org/studies/S-SCDT-10_1038-S44320-026-00213-w).

Expanded view data, supplementary information, appendices are available for this paper at <https://doi.org/10.1038/s44320-026-00213-w>.

Peer review information

A peer review file is available at <https://doi.org/10.1038/s44320-026-00213-w>

References

- Abdelaal T, Grossouw LM, Pasterkamp RJ, Lelieveldt BP, Reinders MJ, Mahfouz A (2024) SIRV: spatial inference of RNA velocity at the single-cell resolution. *NAR Genomics Bioinforma* 6:lqae100
- Acharya A, Baek ST, Huang G, Eskiocak B, Goetsch S, Sung CY, Banfi S, Sauer MF, Olsen GS, Duffield JS (2012) The bHLH transcription factor Tcf21 is required for lineage-specific EMT of cardiac fibroblast progenitors. *Development* 139:2139–2149
- Aibar S, González-Blas CB, Moerman T, Huynh-Thu VA, Imrichova H, Hulselmans G, Rambow F, Marine J-C, Geurts P, Aerts J (2017) SCENIC: single-cell regulatory network inference and clustering. *Nat Methods* 14:1083–1086
- Armingol E, Officer A, Harismendy O, Lewis NE (2021) Deciphering cell-cell interactions and communication from gene expression. *Nat Rev Genet* 22:71–88
- Asp M, Bergenstråhle J, Lundeberg J (2020) Spatially resolved transcriptomes—next generation tools for tissue exploration. *BioEssays* 42:1900221
- Bergen V, Lange M, Peidli S, Wolf FA, Theis FJ (2020) Generalizing RNA velocity to transient cell states through dynamical modeling. *Nat Biotechnol* 38:1408–1414
- Bergen V, Soldatov RA, Kharchenko PV, Theis FJ (2021) RNA velocity—current challenges and future perspectives. *Mol Syst Biol* 17:MSB20210282
- Chen A, Liao S, Cheng M, Ma K, Wu L, Lai Y, Qiu X, Yang J, Xu J, Hao S (2022) Spatiotemporal transcriptomic atlas of mouse organogenesis using DNA nanoball-patterned arrays. *Cell* 185:1777–1792.e1721
- Chizhikov VV, Lindgren AG, Mishima Y, Roberts RW, Aldinger KA, Miesegaes GR, Currle DS, Monuki ES, Millen KJ (2010) Lmx1a regulates fates and location of cells originating from the cerebellar rhombic lip and telencephalic cortical hem. *Proc Natl Acad Sci USA* 107:10725–10730
- Conway SJ, Molkentin JD (2008) Periostin as a heterofunctional regulator of cardiac development and disease. *Curr Genomics* 9:548–555
- Cui H, Maan H, Vladoiu MC, Zhang J, Taylor MD, Wang B (2024) DeepVelo: deep learning extends RNA velocity to multi-lineage systems with cell-specific kinetics. *Genome Biol* 25:27
- das Neves L, Duchala CS, Godinho F, Haxhiu MA, Colmenares C, Macklin WB, Campbell CE, Butz KG, Gronostajski RM (1999) Disruption of the murine nuclear factor IA gene (Nfia) results in perinatal lethality, hydrocephalus, and agenesis of the corpus callosum. *Proc Natl Acad Sci USA* 96:11946–11951
- Deckx S, Johnson DM, Rienks M, Carai P, Van Deel E, Van der Velden J, Sipido KR, Heymans S, Papageorgiou A-P (2019) Extracellular SPARC increases cardiomyocyte contraction during health and disease. *PLoS ONE* 14:e0209534
- Deng Y, He Y, Xu J, He H, Zhang M, Li G (2025) Cardiac Fibroblasts regulate myocardium and coronary vasculature development via the collagen signaling pathway. *eLife* 13:RP102305
- Dumbrava EEI, Balaji K, Raghav K, Hess K, Javle M, Blum-Murphy M, Ajani J, Kopetz S, Broadus R, Routbort M (2019) Targeting ERBB2 (HER2) amplification identified by next-generation sequencing in patients with advanced or metastatic solid tumors beyond conventional indications. *JCO Precis Oncol* 3:1–12
- Feng J, Janečková E, Guo T, Ziaei H, Zhang M, Geng JJ, Cha S, Araujo-Villalba A, Liu M, Ho T-V (2025) High-resolution spatial transcriptomics and cell lineage analysis reveal spatiotemporal cell fate determination during craniofacial development. *Nat Commun* 16:4396
- Ferguson BR, Gao W-J (2015) Development of thalamocortical connections between the mediodorsal thalamus and the prefrontal cortex and its implication in cognition. *Front Hum Neurosci* 8:1027
- Galbraith MD, Donner AJ, Espinosa JM (2010) CDK8: a positive regulator of transcription. *Transcription* 1:4–12
- Gao M, Qiao C, Huang Y (2022) UniTVelo: temporally unified RNA velocity reinforces single-cell trajectory inference. *Nat Commun* 13:6586
- Gayoso A, Weiler P, Lotfollahi M, Klein D, Hong J, Streets A, Theis FJ, Yosef N (2024) Deep generative modeling of transcriptional dynamics for RNA velocity analysis in single cells. *Nat Methods* 21:50–59
- Giacco A, Iervolino S, Cioffi F, Peluso T, Mercurio G, Roberto L, de Rosa V, Cammarota M, Varricchio S, Staibano S (2025) Brain abnormalities in young single-and double-heterozygote mice for both Nkx2-1-and Pax8-null mutations. *Mol Neurobiol* 62:4023–4041

- Gu Y, Blaauw D, Welch JD (2026) Bayesian inference of RNA velocity incorporating timepoints, lineage bifurcations, and count data. *PLoS Comput Biol* 22:e1014060
- Guido W (2018) Development, form, and function of the mouse visual thalamus. *J Neurophysiol* 120:211–225
- Gyllborg D, Langseth CM, Qian X, Choi E, Salas SM, Hilscher MM, Lein ES, Nilsson M (2020) Hybridization-based in situ sequencing (HybISS) for spatially resolved transcriptomics in human and mouse brain tissue. *Nucleic Acids Res* 48:e112
- Hunter MV, Moncada R, Weiss JM, Yanai I, White RM (2021) Spatially resolved transcriptomics reveals the architecture of the tumor-microenvironment interface. *Nat Commun* 12:6278
- Jolasun Y, Song K, Zheng Y, Wang J, Fonseca GJ, Eidelman DH, Ding J (2025) SIDISH integrates single-cell and bulk transcriptomics to identify high-risk cells and guide precision therapeutics through in silico perturbation. *Nat Commun* 6:1127
- Kashima Y, Sakamoto Y, Kaneko K, Seki M, Suzuki Y, Suzuki A (2020) Single-cell sequencing techniques from individual to multiomics analyses. *Exp Mol Med* 52:1419–1427
- Kolodziejczyk AA, Kim JK, Svensson V, Marioni JC, Teichmann SA (2015) The technology and biology of single-cell RNA sequencing. *Mol Cell* 58:610–620
- Kouno T, de Hoon M, Mar JC, Tomaru Y, Kawano M, Carninci P, Suzuki H, Hayashizaki Y, Shin JW (2013) Temporal dynamics and transcriptional control using single-cell gene expression analysis. *Genome Biol* 14:R118
- La Manno G, Soldatov R, Zeisel A, Braun E, Hochgerner H, Petukhov V, Lidschreiber K, Kastri ME, Lönnerberg P, Furlan A (2018) RNA velocity of single cells. *Nature* 560:494–498
- Landry CF, Ellison JA, Pribyl TM, Campagnoni C, Kampf K, Campagnoni AT (1996) Myelin basic protein gene expression in neurons: developmental and regional changes in protein targeting within neuronal nuclei, cell bodies, and processes. *J Neurosci* 16:2452–2462
- Lange M, Bergen V, Klein M, Setty M, Reuter B, Bakhti M, Lickert H, Ansari M, Schniering J, Schiller HB (2022) CellRank for directed single-cell fate mapping. *Nat methods* 19:159–170
- Larocque D, Galarneau A, Liu H-N, Scott M, Almazan G, Richard S (2005) Protection of p27Kip1 mRNA by quaking RNA binding proteins promotes oligodendrocyte differentiation. *Nat Neurosci* 8:27–33
- Li M, Wang J, Qian Q, Chen B, Zhang Z, Barroso E, Zhang Y, Hall DD, Abel ED, Song L-S (2022) Gamma-interferon-inducible lysosomal thiol reductase maintains cardiac immuno-metabolic homeostasis in heart failure. Preprint at <https://doi.org/10.1101/2022.07.03.498477>
- Li S, Zhang P, Chen W, Ye L, Brannan KW, Le N-T, Abe J-i, Cooke JP, Wang G (2024) A relay velocity model infers cell-dependent RNA velocity. *Nat Biotechnol* 42:99–108
- Li Y (2022) Cellular Mechanism of Mouse Atrial Development. *Open J Regen Med* 11:1–24
- Lidell SA (2015) Development of the choroid plexus and blood-CSF barrier. *Front Neurosci* 9:32
- Lim J-Y, Kim SW, Kim B, Park S-J (2020) Knockdown of CARD14 inhibits cell proliferation and migration in breast cancer cells. *Anticancer Res* 40:1953–1962
- Lohoff T, Ghazanfar S, Missarova A, Koulina N, Pierson N, Griffiths JA, Bardot ES, Eng C-H, Tyser RC, Argelaguet R (2022) Integration of spatial and single-cell transcriptomic data elucidates mouse organogenesis. *Nat Biotechnol* 40:74–85
- Madison BB, McKenna LB, Dolson D, Epstein DJ, Kaestner KH (2009) FoxF1 and FoxL1 link hedgehog signaling and the control of epithelial proliferation in the developing stomach and intestine. *J Biol Chem* 284:5936–5944
- Malek Mohammadi M, Kattih B, Grund A, Froese N, Korf-Klingebiel M, Gigina A, Schrameck U, Rudat C, Liang Q, Kispert A (2017) The transcription factor GATA4 promotes myocardial regeneration in neonatal mice. *EMBO Mol Med* 9:265–279
- Miquelajàuregui A, Sandoval-Schaefer T, Martínez-Armenta M, Pérez-Martínez L, Cárabez A, Zhao Y, Heide M, Alvarez-Bolado G, Varela-Echavarría A (2015) LIM homeobox protein 5 (Lhx5) is essential for mamillary body development. *Front Neuroanat* 9:136
- Miquelajàuregui A, Varela-Echavarría A, Ceci ML, García-Moreno F, Ricano I, Hoang K, Frade-Pérez D, Portera-Cailliau C, Tamariz E, De Carlos JA (2010) LIM-homeobox gene Lhx5 is required for normal development of Cajal-Retzius cells. *J Neurosci* 30:10551–10562
- Modi S, Saura C, Yamashita T, Park YH, Kim S-B, Tamura K, Andre F, Iwata H, Ito Y, Tsurutani J (2020) Trastuzumab deruxtecan in previously treated HER2-positive breast cancer. *N Engl J Med* 382:610–621
- Oliveira MFD, Romero JP, Chung M, Williams SR, Gottscho AD, Gupta A, Pilipauskas SE, Mohabbat S, Raman N, Sukovich DJ (2025) High-definition spatial transcriptomic profiling of immune cell populations in colorectal cancer. *Nat Genet* 57:1512–1523
- Pascual-Ahuir A, Fita-Torró J, Proft M (2020) Capturing and understanding the dynamics and heterogeneity of gene expression in the living cell. *Int J Mol Sci* 21:8278
- Pattyn A, Morin X, Cremer H, Goridis C, Brunet J-F (1999) The homeobox gene Phox2b is essential for the development of autonomic neural crest derivatives. *Nature* 399:366–370
- Pillai A, Mansouri A, Behringer R, Westphal H, Goulding M (2007) Lhx1 and Lhx5 maintain the inhibitory-neurotransmitter status of interneurons in the dorsal spinal cord. *Development* 134:357–366
- Qiao C, Huang Y (2021) Representation learning of RNA velocity reveals robust cell transitions. *Proc Natl Acad Sci USA* 118:e2105859118
- Qiu S, Hu Y, Huang Y, Gao T, Wang X, Wang D, Ren B, Shi X, Chen Y, Wang X (2024) Whole-brain spatial organization of hippocampal single-neuron projectomes. *Science* 383:eadi9198
- Raghavan V, Zheng Y, Li Y, Ding J (2025) Harnessing agent-based frameworks in CellAgentChat to unravel cell-cell interactions from single-cell and spatial transcriptomics. *Genome Res* 35:1646–1663
- Rajebhosale P, Jone A, Johnson KR, Hofland R, Palarpalar C, Khan S, Role LW, Talmage DA (2024) Neuregulin1 nuclear signaling influences adult neurogenesis and regulates a schizophrenia susceptibility gene network within the mouse dentate gyrus. *J Neurosci* 44:e0063-24.2024
- Ru B, Huang J, Zhang Y, Aldape K, Jiang P (2023) Estimation of cell lineages in tumors from spatial transcriptomics data. *Nat Commun* 14:568
- Sahai E, Astsaturov I, Cukierman E, DeNardo DG, Egeblad M, Evans RM, Fearon D, Gretchen FR, Hingorani SR, Hunter T (2020) A framework for advancing our understanding of cancer-associated fibroblasts. *Nat Rev Cancer* 20:174–186
- Sahni V, Itoh Y, Shnyder SJ, Macklis JD (2021) Crim1 and Kelch-like 14 exert complementary dual-directional developmental control over segmentally specific corticospinal axon projection targeting. *Cell Rep* 37
- Samaan G, Yugo D, Rajagopalan S, Wall J, Donnell R, Goldowitz D, Gopalakrishnan R, Venkatachalam S (2010) Foxn3 is essential for craniofacial development in mice and a putative candidate involved in human congenital craniofacial defects. *Biochem Biophys Res Commun* 400:60–65
- Senova S, Fomenko A, Gondard E, Lozano AM (2020) Anatomy and function of the fornix in the context of its potential as a therapeutic target. *J Neurol, Neurosurg Psychiatry* 91:547–559
- Simon HH, Scholz C, O'Leary DD (2005) Engrailed genes control developmental fate of serotonergic and noradrenergic neurons in mid-and hindbrain in a gene dose-dependent manner. *Mol Cell Neurosci* 28:96–105
- Singh A, Mahesh A, Noack F, Cardoso de Toledo B, Calejari F, Tiwari VK (2022) Tcf12 and NeuroD1 cooperatively drive neuronal migration during cortical development. *Development* 149:dev200250

- Song K, Zheng Y, Zhao B, Eidelman DH, Tang J, Ding J (2025) DOLPHIN advances single-cell transcriptomics beyond gene level by leveraging exon and junction reads. *Nat Commun* 16:6202
- Ståhl PL, Salmén F, Vickovic S, Lundmark A, Navarro JF, Magnusson J, Giacomello S, Asp M, Westholm JO, Huss M (2016) Visualization and analysis of gene expression in tissue sections by spatial transcriptomics. *Science* 353:78–82
- Surmacz B, Noisa P, Risner-Janiczek JR, Hui K, Ungless M, Cui W, Li M (2012) DLK1 promotes neurogenesis of human and mouse pluripotent stem cell-derived neural progenitors via modulating Notch and BMP signalling. *Stem Cell Rev Rep* 8:459–471
- Swain SM, Shastry M, Hamilton E (2023) Targeting HER2-positive breast cancer: advances and future directions. *Nat Rev Drug Discov* 22:101–126
- Szklarczyk D, Kirsch R, Koutrouli M, Nastou K, Mehryary F, Hachilif R, Gable AL, Fang T, Doncheva NT, Pyysalo S (2023) The STRING database in 2023: protein-protein association networks and functional enrichment analyses for any sequenced genome of interest. *Nucleic Acids Res* 51:D638–D646
- Tallquist MD (2020) Developmental pathways of cardiac fibroblasts. *Cold Spring Harb Perspect Biol* 12:a037184
- Tang XT, Chen LV, Zhou BO (2025) Resolving the spatial organization of fetal liver hematopoiesis by SeekSpace. *Cell Regen* 14:15
- Tritschler S, Büttner M, Fischer DS, Lange M, Bergen V, Lickert H, Theis FJ (2019) Concepts and limitations for learning developmental trajectories from single cell genomics. *Development* 146:dev170506
- Vahid MR, Brown EL, Steen CB, Zhang W, Jeon HS, Kang M, Gentles AJ, Newman AM (2023) High-resolution alignment of single-cell and spatial transcriptomes with CytoSPACE. *Nat Biotechnol* 41:1543–1548
- Venugopalan SR, Amen MA, Wang J, Wong L, Cavender AC, D'Souza RN, Akerlund M, Brody SL, Hjalt TA, Amendt BA (2008) Novel expression and transcriptional regulation of FoxJ1 during oro-facial morphogenesis. *Hum Mol Genet* 17:3643–3654
- Westerling T, Kuuluvainen E, Maäkelä TP (2007) Cdk8 is essential for preimplantation mouse development. *Mol Cell Biol* 27:6177–6182
- Williams CG, Lee HJ, Asatsuma T, Vento-Tormo R, Haque A (2022) An introduction to spatial transcriptomics for biomedical research. *Genome Med* 14:68
- Xu H, Dzhashiashvili Y, Shah A, Kunjamma RB, Weng Y-I, Elbaz B, Fei Q, Jones JS, Li YI, Zhuang X (2020) m6A mRNA methylation is essential for oligodendrocyte maturation and CNS myelination. *Neuron* 105:293–309.e295
- Yan S, Peng Y, Lu J, Shakil S, Shi Y, Crossman DK, Johnson JrWH, Liu S, Rokosh DG, Lincoln J (2022) Differential requirement for DICER1 activity during the development of mitral and tricuspid valves. *J Cell Sci* 135:jcs259783
- Yu D, Hung M-C (2000) Overexpression of ErbB2 in cancer and ErbB2-targeting strategies. *Oncogene* 19:6115–6121
- Zhao B, Su T-Y, Wang J, Islam QS, Song K, Huang SK, Allez M, Fonseca GJ, Baglione CJ, Ding J (2025) DENetwork unveils non-differentially expressed genes with functional relevance across conditions through information flow perturbation. *Nucleic Acids Res* 53:gkaf1350
- Zhao S-H, Ji X-Y, Yuan G-Z, Cheng T, Liang H-Y, Liu S-Q, Yang F-Y, Tang Y, Shi S (2024) A bibliometric analysis of the spatial transcriptomics literature from 2006 to 2023. *Cell Mol Neurobiol* 44:50
- Zhao Y, Sheng HZ, Amini R, Grinberg A, Lee E, Huang S, Taira M, Westphal H (1999) Control of hippocampal morphogenesis and neuronal differentiation by the LIM homeobox gene *Lhx5*. *Science* 284:1155–1158
- Zheng SC, Stein-O'Brien G, Boukas L, Goff LA, Hansen KD (2023) Pumping the brakes on RNA velocity by understanding and interpreting RNA velocity estimates. *Genome Biol* 24:246
- Zheng Y, Schupp JC, Adams T, Clair G, Justet A, Ahangari F, Yan X, Hansen P, Carlon M, Cortesi E (2025) A deep generative model for deciphering cellular dynamics and in silico drug discovery in complex diseases. *Nat Biomed Eng* 9:1–26

Acknowledgements

This work is supported by grants from the Canadian Institutes of Health Research (CIHR) [PJT-180505 to JD]; the Fonds de recherche du Québec—Santé (FRQS) [295298 to JD, 295299 to JD, 366764 to JD]; the Natural Sciences and Engineering Research Council of Canada (NSERC) [RGPIN2022-04399 to JD]; and the Meakins-Christie Chair in Respiratory Research [to JD]; VR received support from the Canada First Research Excellence Fund and the Fonds de recherche du Québec awarded to the D2R Initiative at McGill University. This research was enabled in part by support provided by Calcul Québec (calculquebec.ca) and the Digital Research Alliance of Canada (alliancecan.ca).

Author contributions

Vishvak Raghavan: Conceptualization; Data curation; Software; Formal analysis; Visualization; Methodology; Writing—original draft; Writing—review and editing. **Brent Yoon:** Conceptualization; Software; Formal analysis; Methodology; Writing—original draft. **Gregory J Fonseca:** Resources; Writing—review and editing. **Yue Li:** Supervision; Writing—review and editing. **Jun Ding:** Conceptualization; Resources; Supervision; Funding acquisition; Methodology; Writing—original draft; Project administration; Writing—review and editing.

Source data underlying figure panels in this paper may have individual authorship assigned. Where available, figure panel/source data authorship is listed in the following database record: [biostudies:S-SCDT-10_1038-S44320-026-00213-w](https://biostudies.org/studies/S-SCDT-10_1038-S44320-026-00213-w).

Disclosure and competing interests statement

The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. Creative Commons Public Domain Dedication waiver <http://creativecommons.org/public-domain/zero/1.0/> applies to the data associated with this article, unless otherwise stated in a credit line to the data, but does not extend to the graphical or creative elements of illustrations, charts, or figures. This waiver removes legal barriers to the re-use and mining of research data. According to standard scholarly practice, it is recommended to provide appropriate citation and attribution whenever technically possible.

© The Author(s) 2026