



Published in final edited form as:

Nat Biotechnol. 2023 March ; 41(3): 387–398. doi:10.1038/s41587-022-01476-y.

Multi-omic single-cell velocity models epigenome–transcriptome interactions and improves cell fate prediction

Chen Li¹, Maria C. Virgilio^{1,2}, Kathleen L. Collins^{2,3,4}, Joshua D. Welch^{1,5,✉}

¹Department of Computational Medicine and Bioinformatics, University of Michigan, Ann Arbor, MI, USA.

²Cellular and Molecular Biology Program, University of Michigan, Ann Arbor, MI, USA.

³Department of Microbiology and Immunology, University of Michigan, Ann Arbor, MI, USA.

⁴Department of Internal Medicine, University of Michigan, Ann Arbor, MI, USA.

⁵Department of Computer Science and Engineering, University of Michigan, Ann Arbor, MI, USA.

Abstract

Multi-omic single-cell datasets, in which multiple molecular modalities are profiled within the same cell, offer an opportunity to understand the temporal relationship between epigenome and transcriptome. To realize this potential, we developed MultiVelo, a differential equation model of gene expression that extends the RNA velocity framework to incorporate epigenomic data. MultiVelo uses a probabilistic latent variable model to estimate the switch time and rate parameters of chromatin accessibility and gene expression and improves the accuracy of cell fate prediction compared to velocity estimates from RNA only. Application to multi-omic single-cell datasets from brain, skin and blood cells reveals two distinct classes of genes distinguished by whether chromatin closes before or after transcription ceases. We also find four types of cell states: two states in which epigenome and transcriptome are coupled and two distinct decoupled states. Finally, we identify time lags between transcription factor expression and binding site

✉ **Correspondence and requests for materials** should be addressed to Joshua D. Welch. welchjd@umich.edu.

Author contributions

M.C.V. and K.L.C. generated the 10x Multiome HSPC data. J.D.W. conceived the idea of multi-omic extension of RNA velocity. C.L. and J.D.W. developed and implemented the method, performed data analyses and wrote the manuscript. All authors read and approved the final manuscript.

Competing interests

The authors declare no competing interests.

Code availability

MultiVelo is implemented in Python. The package is available on GitHub (<https://github.com/welch-lab/MultiVelo>), PyPI and Bioconda.

Reporting summary

Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Additional information

Extended data are available for this paper at <https://doi.org/10.1038/s41587-022-01476-y>.

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41587-022-01476-y>.

Peer review information *Nature Biotechnology* thanks Yuanhua Huang and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

accessibility and between disease-associated SNP accessibility and expression of the linked genes. MultiVelo is available on PyPI, Bioconda and GitHub (<https://github.com/welch-lab/MultiVelo>).

The regulation of gene expression from DNA to RNA to protein is a key process governing cell fates. Coordinated, sequential gene expression changes underlie the developmental processes by which cells specialize. Increasingly, high-throughput single-cell sequencing techniques are being applied to reveal these sequential gene expression changes. However, because experimental measurement destroys the cell, only temporal snapshot measurements are available, and it is not possible to observe the same individual cell changing over time.

Computational approaches can use single-cell snapshots to infer sequential gene expression changes during developmental processes. For example, cell trajectory inference algorithms^{1–5} use pairwise cell similarities to map cells onto a ‘pseudotime’ axis corresponding to predicted developmental progress. However, trajectory inference based on similarity cannot predict the directions or relative rates of cellular transitions. Methods for inferring RNA velocity^{6,7} address these limitations by fitting a system of differential equations that describes the directions and rates of transcriptional changes using spliced and unspliced transcript counts. A recent paper further extended the RNA velocity framework to include gene expression and protein measurements from the same cells but used the steady-state assumption to estimate parameters and thus did not estimate latent time values for each cell⁸. Single-cell epigenome values have also been used individually to infer future directions of cell differentiation, but these approaches did not incorporate gene expression^{9,10}.

Single-cell multi-omic measurements provide an opportunity to incorporate epigenomic data into mechanistic models of transcription. For example, new technologies such as SNARE-seq¹¹, SHARE-seq⁹ and 10x Genomics Multiome can quantify both RNA and chromatin accessibility in the same cell. The epigenome and transcriptome both change during cellular differentiation, and thus, the temporal snapshots in single-cell multi-omic datasets potentially reveal the interplay among these molecular layers.

Existing RNA velocity models assume that the transcription rate of a gene is uniform throughout the induction phase of gene expression. However, epigenomic changes play a key role in regulating gene expression, such as tightening or loosening the chromatin compaction of promoter and enhancer regions^{12–16}. For example, a transition from euchromatin to heterochromatin reduces the rate of transcription at that locus because transcriptional machinery cannot access the DNA. Therefore, a more realistic model would reflect the influence of enhancer and promoter chromatin accessibility on transcription rate.

We present MultiVelo, a computational approach for inferring epigenomic regulation of gene expression from single-cell multi-omic datasets. We extend the dynamical RNA velocity model to incorporate multi-omic measurements to more accurately predict the past and future state of each cell, jointly infer the instantaneous rate of induction or repression for each modality and determine the extent of coupling or time lag between modalities.

Results

MultiVelo: a differential equation model of gene expression incorporating chromatin accessibility

MultiVelo describes the process of gene expression as a system of three ordinary differential equations (ODEs) characterized by a set of switch time and rate parameters (Fig. 1a). This model represents a deliberately simplified view of gene expression in which the complex effects of chromatin modifiers, pioneer factors and transcription factors (TFs) are abstracted into rate constants. The time-varying levels of chromatin accessibility (c), unspliced pre-mRNA (u) and spliced mature mRNA (s) are related by ODEs describing the rates of chromatin opening (α_o) and closing (α_c), RNA transcription (α), RNA splicing (β) and RNA degradation or nuclear export (γ). We assume that chromatin opening rapidly leads to full accessibility and similarly that chromatin closing rapidly leads to full inaccessibility. Each gene has distinct rate parameters describing its unique kinetics. We assume that the transcription rate is proportional to the chromatin accessibility $c(t)$ and thus is time-varying, and we model the distinct phases or states k that a cell traverses as its time t advances. There are two states each for chromatin accessibility (c) and RNA (u, s): chromatin opening, chromatin closing, transcriptional induction and transcriptional repression. Each state begins at an associated switch time (t_c, t_i and t_r ; chromatin opening begins at $t_o = 0$) and converges to an associated steady-state value as $t \rightarrow \infty$. The rate parameters and switch times are estimated for each gene using the three-dimensional phase portrait of (c, u, s) triplets observed across a set of single cells. The state k and time t for each cell are determined by projecting the cell to the nearest point on the curve described by the ODEs.

Starting from these assumptions gives a mathematical model with two interesting qualitative properties. First, there are multiple mathematically feasible combinations of chromatin accessibility and RNA transcription states. That is, chromatin can be either opening or closing while transcription is being either induced or repressed. This means that multiple orders of events are possible: chromatin closing can occur either before or after transcriptional repression begins (Fig. 1b). We refer to the first ordering (chromatin closing begins before transcriptional repression) as model 1 and the second ordering as model 2.

The second interesting qualitative property is that two distinct types of discordance between chromatin accessibility and transcription can occur. At the beginning of the gene expression process, chromatin opens before transcription initiates. This creates a time interval during which $c(t)$ is positive but $u(t)$ and $s(t)$ are both zero (Fig. 1c). We refer to this phenomenon as priming. In addition, at the end of the gene expression process, chromatin closing and transcriptional repression can occur at different times. This creates a time interval in which chromatin accessibility and gene expression move in opposite directions (Fig. 1d), a phenomenon we refer to as decoupling.

MultiVelo infers and quantifies these phenomena of multiple orders and types of discordance through the ODE parameters estimated from single-cell data. First, the switch times (t_c, t_i , and t_r) indicate when chromatin closing, transcriptional induction and transcriptional repression begin. Thus, the lengths of priming and decoupling phases are estimated by the

model: $\Delta t_{\text{priming}} = t_i - t_o = t_i$ and $\Delta t_{\text{decoupling}} = t_r - t_c$. Furthermore, because each cell is assigned latent time (t) and latent state (k) values, MultiVelo determines whether each cell is in a primed, decoupled or coupled phase for each gene (Fig. 1e). Thus, we refer to the four possible states as primed (red), coupled on (orange), decoupled (green) and coupled off (blue). Second, the parameters fit by MultiVelo can be used to determine, for each gene, whether its observed (c, u, s) values are best fit by model 1 or model 2 (Fig. 1f,g). Intuitively, it is possible to distinguish these models because model 1 genes achieve their highest accessibility values during the transcriptional induction phase, whereas model 2 genes reach maximum accessibility during the transcriptional repression phase.

Because both the cell times and the ODE parameters are unknown, we use an iterative expectation-maximization algorithm to jointly estimate them. Briefly, we initialize both the parameters and cell times using heuristics derived from the steady-state assumption (Methods). Then, we iterate the following steps: (1) calculate the most likely time of each cell based on the current ODE parameter estimates, and (2) update the ODE parameters to maximize the data fit given the current time estimates.

We further derived a stochastic version of the MultiVelo model (Methods), which uses moment equations and the steady-state assumption to perform parameter estimation. Runtime and memory usage statistics are in Supplementary Table 1 and Methods.

MultiVelo distinguishes two models of gene expression regulation in embryonic mouse brain

We first applied MultiVelo to 10x Multiome data from the embryonic mouse brain embryonic day 18 (E18). MultiVelo accurately fits the observed chromatin accessibility, unspliced pre-mRNA and spliced mRNA counts across the population of brain cells, identifying 426 genes whose patterns fit the model with high likelihood. The resulting velocity vectors and latent time values inferred by MultiVelo accurately recover the known trajectory of mammalian cortex development. Specifically, radial glia (RG) cells in the outer subventricular zone give rise to neurons, astrocytes and oligodendrocytes^{17–19}. Cortical layers are formed in an inside-out fashion during neuron migration, with newborn cells moving to upper layers and older cells staying in deeper layers²⁰. RG cells can divide into intermediate progenitor cells (IPC) that serve as neural stem cells and further generate various mature excitatory neurons in different layers^{21,22}.

Incorporating both chromatin accessibility and gene expression improves the accuracy of velocity estimation compared to RNA-only models such as scVelo (Fig. 2a). In particular, the RNA-only model predicts biologically implausible backflows inside upper layer neurons (Fig. 2b). Cell cycle scores^{7,23} indicate that the developmental process begins with a cycling population (Fig. 2c) near RG, confirming the latent time inferred by MultiVelo.

We expect the addition of chromatin accessibility to be most helpful for distinguishing cell states where chromatin remodeling and gene expression are out of sync, such as when a gene's promoters and enhancers have begun to open but little transcription has occurred. Two clear examples are *Eomes* and *Tle4*, canonical markers of IPCs and deep layer neurons^{24–27}. RNA transcripts from these genes are highly expressed in only one or

two specific cell types. The remaining cells are densely clustered near the origin of the (u, s) phase portrait, making it difficult for RNA velocity methods to distinguish their relative order (Fig. 2d). However, the chromatin accessibility of these genes begins to rise before the gene expression, revealing gradual changes that are not visible from gene expression alone. To put it another way, incorporating chromatin allows us to infer 3D velocity vectors indicating each cell's predicted differentiation for each gene, better resolving cellular differences than the two-dimensional phase portraits from RNA alone.

MultiVelo identifies clear examples of genes that are best described by either model 1 or model 2 in this dataset. Comparing the phase portraits of the genes assigned to model 1 and model 2 shows clear differences in the timing of maximum chromatin accessibility, consistent with the model predictions (Fig. 2e). Model 1 genes such as *Satb2* reach maximum chromatin accessibility during the transcriptional induction phase (above the diagonal steady-state line on the phase portrait⁶), whereas the accessibility of model 2 genes like *Gria2* is highest during the transcriptional repression phase (below the diagonal steady-state line). The distinction between model 1 and model 2 is also evident when inspecting pairwise phase portraits of c, u and c, s (Fig. 2f). However, the models cannot be distinguished by inspecting the RNA information alone in a phase portrait of u, s ; the distinction requires the additional information from chromatin. We further investigated the model 1 and model 2 genes to see if they have any characteristic properties. The genes in the two classes do not differ significantly in their total expression or accessibility levels (Wilcoxon $P = 0.38$ and $= 0.32$). Gene ontology analysis showed that M2 genes are significantly enriched for terms related to the cell cycle, such as 'positive regulation of cell cycle', 'mitotic cell cycle' and 'regulation of cell cycle phase transition'. Furthermore, model 2 genes tend to achieve their highest spliced expression earlier in latent time than model 1 genes ($P = 9 \times 10^{-7}$, Wilcoxon rank-sum one-sided test; Fig. 2g). We hypothesize that cells may use model 2 for rapid, transient activation of genes that do not need to maintain expression, whereas model 1 may be useful for genes that need to be stably expressed.

We next looked at how often each type of gene expression kinetics (induction only, repression only, model 1 or model 2) occurred. Most of the highly variable genes show both induction and repression phases (a complete trajectory), and for genes that only have partial trajectories, induction-only phase portraits appear more often than repression-only (29.5% versus 2.4% of variable genes; Fig. 2h). Note that, because model 1 and model 2 make the same predictions during the induction phase, we cannot distinguish model 1 versus model 2 for induction-only genes. Among the genes with both an induction and repression phase, the majority are best explained by model 1 (41.4% of variable genes), while the remainder are best fit by model 2 (26.7% of variable genes). The fact that model 1 is more common is consistent with the expectation that chromatin state changes generally precede mRNA expression changes.

Whether genes have complete or partial kinetics, MultiVelo fits ODE parameters that describe the three-dimensional trajectory of their chromatin accessibility and gene expression dynamics (Fig. 2i). We also found that MultiVelo can recover very similar results

on this dataset even if computationally inferred multi-omic profiles from separate datasets, rather than 10x Multiome or SHARE-seq⁹, are used (Supplementary Fig. 1).

MultiVelo identifies epigenomic priming and decoupling in embryonic mouse brain

An exciting property of MultiVelo is its ability to quantify the discordance and concordance between chromatin accessibility and gene expression within differentiating cells. Specifically, MultiVelo infers switch time parameters that identify the intervals during which each gene is in one of the four possible states (primed, coupled on, decoupled and coupled off; Fig. 1e). We next investigated whether these inferred states and time intervals can accurately capture the interplay between epigenomic and transcriptomic changes in embryonic mouse brain cells.

MultiVelo identifies clear examples of each of the four states in the 10x Multiome data (Fig. 3a). For example, *Grin2b* is an induction-only gene with expression increasing toward the neuronal fate, so only induction states (primed and coupled on) were predicted for this gene (Fig. 3a, left). The phase portrait of *Nfix*, a model 1 gene, possesses a complete trajectory shape and was labeled with all four states (Fig. 3a, middle). Conversely, *Epha5* is a model 2 gene, and its accessibility continues to rise throughout the whole time range without an observed closing phase, so it only occupies the coupled on and decoupled states (Fig. 3a, right).

The state assignments can be confirmed qualitatively by plotting accessibility (c) and expression u and s on Uniform Manifold Approximation and Projection (UMAP) coordinates and examining them side-by-side (Fig. 3b). Visually, we observe that the colors of the c and u UMAP plots match when the state assignments are coupled on or coupled off, and the differences in color occur when the assigned states are primed or decoupled. For example, the largest discrepancy between *Robo2* RNA expression and chromatin accessibility occurs in the circled region, which is predicted to be in the decoupled state (Fig. 3b, top). *Robo2* is a model 1 gene; after chromatin closing begins, expression stays at a relatively high level, even though its accessibility has already experienced a drop toward the maturing neurons. Similarly, the accessibility of *Gria2* differs from RNA in the decoupled state (Fig. 3b, middle). The chromatin accessibility of *Gria2*, a model 2 gene, continues to increase beyond the transcriptional induction phase. Furthermore, the gene *Grin2b* shows a clear example of the chromatin priming phase, during which chromatin opens before RNA production (Fig. 3b, bottom).

Plotting c , u and s along the inferred time t for each gene allows us to inspect the state transitions in detail (Fig. 3c). First, the $u(t)$ and $s(t)$ values for *Robo2* show two inflection points during the transcriptional repression phase, corresponding to the transitions from coupled on to decoupled states and from decoupled to coupled off states (Fig. 3c, top). This pattern suggests that the distinct effects of chromatin closing and transcriptional repression are visible in $u(t)$ and $s(t)$. In other words, MultiVelo predicts that for *Robo2*, chromatin closing decreases the overall transcription rate as RNA level begins to drop immediately following the chromatin switch. The subsequent switch of transcription rate from positive to zero causes a second inflection, leading to even more rapid down-regulation of RNA

expression. The plots of $c(t)$, $u(t)$ and $s(t)$ for *Gria2* show the opposite trend: c continues to rise even after the switch to transcriptional repression, causing c and u to move in opposite directions during the decoupled state (Fig. 3c, middle). In *Grin2b*'s long priming phase $c(t)$ begins to rise while $u(t)$ and $s(t)$ stay at zero (Fig. 3c, bottom).

Because MultiVelo fits rate and switch time parameters for each gene, our analysis provides an opportunity to observe general trends in gene regulation. First, to determine whether the states of different genes are temporally coordinated, we counted the number of high-likelihood genes in each state per cell. There is indeed a cascade of state transitions through the neuronal clusters; multiple genes per cell are often simultaneously in the priming or decoupling states (Fig. 3d). Second, we looked for trends in the switch time and rate parameters. We placed each gene's induction/repression cycle on a time scale between 0 and 1 and found that the coupled-on and coupled-off states account for a larger proportion of the gene expression process than the primed and decoupled states (Fig. 3e). This finding makes sense, because even if genes experience some level of decoupling and time lag between the two modalities, chromatin accessibility and gene expression should still be generally correlated^{28–31}. The median primed interval length is 21% of the overall time, and the median decoupled interval length is 19% of the overall time. Furthermore, we can rank genes by how long their priming and decoupling intervals are to find examples of discordance between accessibility and expression (Extended Data Fig. 1d). Additionally, we found that chromatin generally opens and closes at similar rates; the median ratio between inferred chromatin closing rate (α_{cc}) and chromatin opening rate (α_{co}) is almost exactly 1 (Fig. 3f).

MultiVelo quantifies epigenomic priming in SHARE-seq data from mouse hair follicle

A recent study⁹ used SHARE-seq to investigate the rapid proliferation of transit-amplifying cells (TACs) in hair follicle tissue, which give rise to several mature effector cells, including inner root sheath (IRS) and layers of hair shaft: cuticle, cortical layer and medulla³². When applied to this dataset, MultiVelo correctly identified direction of differentiation from TACs to IRS and hair shaft cells (Fig. 4a), consistent with the diffusion map³³ analysis reported in the initial paper⁹. Latent time predicted the TACs to be the root cells, agreeing with biological expectation, whereas velocity analysis using RNA alone failed to capture the hair-shaft differentiation direction (Fig. 4b). When using a similar number of high-likelihood genes, the latent time inferred by MultiVelo has a Spearman correlation coefficient of 0.51 with the pseudotime inferred by Palantir⁴ in the SHARE-seq paper⁹, higher than scVelo's 0.44. We also observed more induction-only and fewer model 2 genes in this dataset compared to mouse brain (Fig. 4c).

One of the key results of the original SHARE-seq paper was the identification of genes where promoter and enhancer chromatin accessibility presaged gene expression, a phenomenon the authors termed 'chromatin potential'. The clearest example of this phenomenon was *Wnt3*, which encodes a paracrine signaling molecule and is important in controlling hair growth³⁴. Indeed, UMAP plots colored by accessibility, and unspliced and spliced mRNA expression show a clear time delay across modalities (Fig. 4d). We next examined the other genes identified in the SHARE-seq paper. Our fit models show that

MultiVelo faithfully captured the dynamics of each gene and provide clear illustrations of priming and decoupling regions (Fig. 4e). For instance, *Wnt3* and *Dsc1* show induction-only patterns and a priming state at the beginning, whereas *Cux1*, *Dlx3* and *Cobll1* have both induction and repression states with a short decoupling period in the middle.

To further quantify the temporal relationships among accessibility, unspliced expression and spliced expression, we used dynamic time warping (DTW)³⁵ to align the time series values for each molecular layer. DTW nonlinearly warps two time series to maximize their similarity and identify possible lagged correlation. DTW results on *Wnt3* show that the optimal warping function maps each point on the *c* time series forward in time, consistent with chromatin accessibility preceding gene expression (Fig. 4f, top). Unspliced and spliced expression show a similar pattern but with a shorter time delay (Fig. 4f, middle). Because DTW maps each time point on the earlier curve to a time point on the later curve, the time lag at each point in time can be computed by subtracting the times of the matched points (Fig. 4f, bottom). This analysis shows that both the delay between *c* and *s* and the delay between *u* and *s* remain positive throughout the observed time. In addition, the delay between *c* and *s* is longer than the delay between *u* and *s* throughout the observed range, with the maximum *c* and *s* delay reaching 0.6 (out of a total time range of 1).

MultiVelo reveals early epigenomic and transcriptomic changes in human HSPCs

Hematopoietic progenitors consist of stem-like cell populations that rapidly and continuously differentiate into various intermediate and mature blood cell types with progressively reduced self-renewal potential as they enter more lineage-restricted states^{30,36}.

We obtained 11,605 high-quality cells post-filtering with both single-nucleus RNA-sequencing and assay for transposase-accessible chromatin with sequencing (ATAC-seq) data. Using previously described marker genes^{37–40}, we identified clusters resembling many of the populations of early blood development (Extended Data Fig. 2a), including hematopoietic stem cells (HSC), multipotent progenitors (MPP), lymphoid-primed multipotent progenitors (LMPP), granulocyte-macrophage progenitors (GMP) and megakaryocyte-erythrocyte progenitors (MEP). We also identified clusters resembling early granulocytes, erythrocytes, dendritic cells (DC) and platelets.

Blood cell differentiation is a challenging system to model with RNA velocity⁴¹. Nevertheless, we find that incorporating chromatin information improves the local consistency and biological accuracy of predicted cell directions in our hematopoiesis dataset (Fig. 5a). In comparison, velocity vectors inferred from RNA alone do not accurately reflect the known differentiation hierarchy of hematopoietic stem and progenitor cells (HSPCs). As with the mouse brain, MultiVelo predicts model 1 to be more common than model 2 in this dataset; induction only is the second most common gene class (Fig. 5b). The median lengths of observed primed and decoupled intervals are shorter than those of the coupled phases (Fig. 5c). These patterns are consistent with what we observed in the mouse brain dataset, suggesting a possible common underlying biological mechanism.

As with the mouse brain dataset, model 2 genes in the HSPC dataset are significantly enriched for gene ontology terms related to the cell cycle. The terms ‘regulation of mitotic

cell cycle', 'regulation of mitotic metaphase/anaphase transition' and 'regulation of mitotic sister chromatid separation' are all enriched in model 2 genes at false discovery rate < 0.002 . If we examine the separate trajectories toward myeloid, erythroid and platelet lineages, many G2/M phase marker genes²³ show clear model 2 patterns, with highest chromatin accessibility after expression begins to drop (examples shown in Fig. 5d). The gene models fit by MultiVelo reveal many examples of priming (Fig. 5e). Several terminal cell-type-specific markers show induction-only dynamics with an increase in chromatin accessibility followed by increasing gene expression (*AZU1* in GMP, *HBD* in erythrocytes, *HDC* in granulocytes, *LYZ* in dendritic cell progenitors and *PF4* in the megakaryocyte progenitors direction)^{40,42}. In HSPCs, we again see some clear examples of long priming periods, such as in *LYZ* and *PF4*.

Plotting velocities allows us to examine local chromatin and RNA trends in more detail (Fig. 5f,g). Although the chromatin shows most potential (highest velocity) at the beginning for these genes, for RNA, stem cell populations such as HSCs, multipotent progenitors, megakaryocyte-erythrocyte progenitors and GMPs show increased potential during their differentiation process toward one lineage. More differentiated cell types lose the ability to maintain such potential and gradually approach equilibrium (zero velocity), even though expression is still increasing somewhat. Note that even though the overall expression elevates, and velocities stay positive, local acceleration can still switch signs.

We further assessed MultiVelo's performance using an additional HSPC sample from a second time point (Extended Data Fig. 2d).

MultiVelo relates TFs, polymorphic sites and gene expression in developing human brain

We next applied MultiVelo to a recently published 10x Multiome dataset from developing human cortex⁴³. As with the embryonic mouse brain dataset, MultiVelo inferred velocity vectors consistent with known patterns of brain cell development (Fig. 6a). MultiVelo correctly inferred a cycling population of cells near RG as the cell type earliest in latent time. In contrast, velocity vectors inferred without chromatin information predicted incongruous backflows in IPCs and upper layer excitatory neurons (Fig. 6b).

As with the mouse brain dataset, we identified clear examples of both model 1 and model 2 genes (Fig. 6c and Extended Data Fig. 3a), though fewer genes are predicted to follow model 2 in the human dataset (Fig. 6d).

A key benefit of MultiVelo is its ability to place cells onto a latent time scale inferred from both chromatin and expression data. We reasoned that latent time can identify time lags between expression and accessibility of loci other than just those immediately near a gene. For example, latent time can be used to calculate the length of time between the expression of a TF and the accessibility of its binding sites (Fig. 6e and Extended Data Fig. 3b,c). To do this, we used chromVar⁴⁴ to calculate, for each cell, the total accessibility of the peaks with binding sites for each TF, subsetting to only the TFs variably expressed in the dataset. We then used DTW³⁵ to align the time series expression of each TF with the accessibility of its binding sites. This process revealed a consistent pattern, in which the time of the highest RNA expression of the TF preceded the time of corresponding high accessibility of

downstream targets. UMAP plots colored by TF expression and binding site accessibility visually confirmed this pattern. The median time lag across all expressed TFs was positive, indicating TF expression precedes binding site accessibility in most cases (Fig. 6f). We cannot conclusively determine the mechanisms underlying these time lags without additional data. However, post-transcriptional and post-translational regulation, factors that affect the activity of chromatin remodeling complexes, and intercellular signaling could all contribute to this phenomenon.

Latent time inferred by MultiVelo is also useful for relating the chromatin accessibility of disease-related variant loci to the expression of nearby genes. We collected a list of 6,968 single-nucleotide polymorphisms (SNPs) and their linked genes implicated by genome-wide association studies of psychiatric diseases, including bipolar disorder and schizophrenia. We further subset these SNPs to those overlapping chromatin accessibility peaks linked to the genes fit by our model (a total of 757 SNPs). Many of these variants occur near neuronal TFs and other developmentally important genes. We then calculated the chromatin accessibility, per cell, of a 400-bp window centered on each SNP. Using MultiVelo's latent time, we determined the time of maximum accessibility for each SNP and the time lag between SNP accessibility and the maximum expression of its linked gene (Fig. 6g). This analysis revealed three major groups of SNPs, distinguished by whether their maximum accessibility occurred early or late in latent time and before or after the expression of the linked gene. UMAP plots of the SNP accessibility and linked gene expression confirm that these groups of SNPs have qualitatively distinct profiles. These groupings are relevant for understanding the functions of the SNPs; for example, an SNP that is accessible only early in latent time likely plays a bigger role in developing cells than in fully differentiated cells. Similarly, a SNP whose accessibility precedes a gene's expression is more likely to participate in regulating its expression than an SNP whose accessibility lags behind.

Discussion

In summary, MultiVelo models temporal chromatin accessibility and gene expression levels and quantifies the length of priming and decoupling intervals in which chromatin accessibility and gene expression are temporarily out of sync. Our model accurately fits single-cell multi-omic datasets from embryonic mouse brain, mouse dorsal skin, fetal human brain and human HSPCs. We find that incorporating chromatin accessibility data improves the overall accuracy of velocity estimates, with the largest differences in early stem cells undergoing rapid epigenomic changes. In our view, the most exciting new direction opened by MultiVelo is the ability to relate epigenomic and transcriptional changes during differentiation. For example, our model identifies two classes of genes that differ in the relative order of chromatin closing and transcriptional repression, and we find clear examples of both mechanisms across all of the tissues we investigated.

Our velocity estimates can be combined with methods such as trajectoryNet⁴⁵, WaddingtonOT⁴⁶ or VeloAE⁴⁷ to predict global dynamics. An interesting direction for future work is extending the approach to incorporate additional steps of the gene expression process, such as TF binding and chromatin looping. We anticipate that MultiVelo will

provide insights into epigenomic regulation of gene expression across a range of biological settings, including normal cell differentiation, reprogramming and disease.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41587-022-01476-y>.

Methods

Previous approaches: RNA velocity

In the original RNA velocity model, the proposed system of differential equations for RNA splicing is as follows:

$$\frac{du}{dt} = \alpha(t) - \beta(t)u(t) \quad (1)$$

$$\frac{ds}{dt} = \beta(t)u(t) - \gamma(t)s(t) \quad (2)$$

where u is unspliced RNA, s is spliced RNA and α, β, γ are transcription, splicing and degradation rate, respectively. Assuming constant transcription and degradation rates, the rate equation parameters can be normalized by β and are reduced to

$$\frac{du}{dt} = \alpha - u(t) \quad (3)$$

$$\frac{ds}{dt} = u(t) - \gamma' s(t) \quad (4)$$

In steady-state cell populations, the amount of spliced mRNA does not change: $\frac{ds}{dt} = 0$.

Therefore, $\gamma' = \frac{u}{s}$ and $\alpha = u$. The ratio γ' can be calculated using a simple linear regression that fits cells with expression values in upper and lower quantiles. RNA velocity is then defined as $v = \frac{ds}{dt}$.

Bergen et al.⁷ developed a dynamical RNA velocity model (scVelo) by extending the original equations to include time and cell state latent variables, capturing transient states between steady states:

$$\frac{du(t)}{dt} = \alpha^{(k)} - \beta u(t) \quad (5)$$

$$\frac{ds(t)}{dt} = \beta u(t) - \gamma s(t) \quad (6)$$

where k indicates one of the four transcription states: induction ($k = 1$), repression ($k = 0$) and two associated steady states ($k = ss1$ and $k = ss0$).

This system of differential equations can be solved analytically as follows:

$$u(t) = u_0 e^{-\beta\tau} + \frac{\alpha^{(k)}}{\beta} (1 - e^{-\beta\tau}) \quad (7)$$

$$s(t) = s_0 e^{-\gamma\tau} + \frac{\alpha^{(k)}}{\gamma} (1 - e^{-\gamma\tau}) + \frac{\alpha^{(k)} - \beta u_0}{\gamma - \beta} (e^{-\gamma\tau} - e^{-\beta\tau}) \quad (8)$$

where u_0 and s_0 are initial values, and $\tau = t - t_0^{(k)}$ is the time interval from the start of the induction or repression state.

The analytical solution converges to the steady-state values as $\tau \rightarrow \infty$:

$$(u_\infty^{(k)}, s_\infty^{(k)}) = \left(\frac{\alpha^{(k)}}{\beta}, \frac{\alpha^{(k)}}{\gamma} \right). \quad (9)$$

Because the equations involve the latent time variable τ , scVelo uses an expectation-maximization algorithm to iteratively estimate latent time and the parameters of the ODE $\theta = (\alpha^{(k)}, \beta, \gamma)$, as well as state starting time $t_0^{(k)}$. Cells are assigned to latent times by approximately inverting the ODE solution.

Differential equation model of gene expression incorporating chromatin accessibility

To incorporate chromatin accessibility measurements into a differential equation model of gene expression, we assume that the rate of transcription for a gene is influenced by the accessibility of its promoter and enhancers. For simplicity, we model a single value c , which is the sum of accessibility at the promoter and linked peaks for a gene. Unlike gene expression, which can theoretically grow without bound, it is possible in principle for chromatin to be fully open or fully closed at a particular locus. Thus, we normalize chromatin accessibility to $[0, 1]$ and assume that c approaches 1 with rate of change proportional to the current accessibility level c (proportionality constant $\alpha_{co} > 0$) during the opening phase and approaches 0 with rate of change proportional to c (proportionality constant $\alpha_{cc} > 0$) during the closing phase. Our biological motivation for this mathematical formulation can be summarized as follows: impulses of remodeling signals cause chromatin to begin opening or closing rapidly at first. However, biochemical constraints such as the structures of histone complexes and their intermolecular interactions gradually slow the rate of opening or closing so that c asymptotically approaches full accessibility or inaccessibility (Extended Data Fig. 4a). Empirically, we find that the observed $c(t)$ values in single-cell multi-omic dataset show this qualitative behavior (Extended Data Fig. 4b). We define a new system of differential equations to reflect these modeling assumptions:

$$\frac{dc(t)}{dt} = -\alpha_c c(t) \text{ or } \frac{dc(t)}{dt} = \alpha_{co} - \alpha_{co} c(t). \quad (10)$$

If we assume that the chromatin opening and closing kinetics are mirror images of each other, only a single chromatin rate parameter $\alpha_c > 0$ is required, and the system of equations simplifies to:

$$\frac{dc(t)}{dt} = k_c \alpha_c - \alpha_c c(t) \quad (11)$$

$$\frac{du(t)}{dt} = \alpha^{(k)} c(t) - \beta u(t) \quad (12)$$

$$\frac{ds(t)}{dt} = \beta u(t) - \gamma s(t), \quad (13)$$

where

$$k_c = \begin{cases} 1, & \text{if chromatin is opening} \\ 0, & \text{if chromatin is closing} \end{cases}.$$

Note that we write the ODEs in terms of a single chromatin rate α_c purely for notational simplicity. In the MultiVelo package and in all of our experiments in the paper, we fit separate chromatin opening and closing rates.

As with the RNA velocity model, we define chromatin velocity as $\frac{dc}{dt}$. The parameter k_c allows for different dynamics during chromatin opening ($k = 1$) and chromatin closing ($k = 0$), analogous to how the transcription rate α_k in the dynamical RNA velocity model varies between transcriptional induction and repression phases ($k = 1$ and $k = 0$). The system of differential equations can be solved analytically to obtain:

$$c(t) = k_c - (k_c - c_0)e^{-\alpha_c t} = c_0 e^{-\alpha_c t} + k_c(1 - e^{-\alpha_c t}) \quad (14)$$

$$u(t) = u_0 e^{-\beta t} + \frac{\alpha^{(k)} k_c}{\beta} (1 - e^{-\beta t}) + \frac{(k_c - c_0) \alpha^{(k)}}{\beta - \alpha_c} (e^{-\beta t} - e^{-\alpha_c t}) \quad (15)$$

$$\begin{aligned} s(t) &= s_0 e^{-\gamma t} + \frac{\alpha^{(k)} k_c}{\gamma} (1 - e^{-\gamma t}) \\ &+ \frac{\beta}{\gamma - \beta} \left(\frac{\alpha^{(k)} k_c}{\beta} - u_0 - \frac{(k_c - c_0) \alpha^{(k)}}{\beta - \alpha_c} \right) (e^{-\gamma t} - e^{-\beta t}), \\ &+ \frac{\beta}{\gamma - \alpha_c} \frac{(k_c - c_0) \alpha^{(k)}}{\beta - \alpha_c} (e^{-\gamma t} - e^{-\alpha_c t}) \end{aligned} \quad (16)$$

where c_0, u_0 and s_0 are the initial values of one of the four states, and $\tau = t - t_0$ is the time interval from the start of that state. Note that the analytical solution is the same even if we assume different opening and closing rates, if we simply use

$$\alpha_c = \begin{cases} \alpha_{co}, & \text{if } k_c = 1 \\ \alpha_{cc}, & \text{if } k_c = 0 \end{cases}$$

Note that we write the ODEs in terms of a single chromatin rate α_c purely for notational simplicity. In the MultiVelo package and in all of our experiments in the paper, we fit separate chromatin opening and closing rates.

Similar to RNA velocity, the origin of the trajectory is $(0, 0, 0)$ (whether observed or not), and initial values of the next state can be obtained by solving the expected values at the switch interval using equations for the previous state. The range of chromatin values is restricted to $[0,1]$ to span from fully closed to fully open chromatin accessibility. As such, the hypothetical steady states for chromatin accessibility $c_\infty^{(k_c)}$, as time approaches infinity on each interval, is simply 0 for closing state and 1 for opening state. The steady-state values for each state become

$$(c_\infty^{(k_c)}, u_\infty^{(k)}, s_\infty^{(k)}) = (k_c, \frac{\alpha^{(k)}k_c}{\beta}, \frac{\alpha^{(k)}k_c}{\gamma}). \quad (17)$$

Because the model includes separate latent variables for chromatin state k_c and RNA state k , there are multiple potential orders of chromatin remodeling states and transcription states. We label these possible orders as model 0 (M0), model 1 (M1) and model 2 (M2):

$$\begin{aligned} \text{M0: } & (k_c = 1, k = 0) \rightarrow (k_c = 0, k = 0) \rightarrow (k_c = 0, k = 1) \rightarrow (k_c = 0, k = 0) \\ \text{M1: } & (k_c = 1, k = 0) \rightarrow (k_c = 1, k = 1) \rightarrow (k_c = 0, k = 1) \rightarrow (k_c = 0, k = 0) \\ \text{M2: } & (k_c = 1, k = 0) \rightarrow (k_c = 1, k = 1) \rightarrow (k_c = 1, k = 0) \rightarrow (k_c = 0, k = 0) \end{aligned}$$

We reason that it is biologically implausible for chromatin to be closed when transcription initiates, because it is difficult or impossible for a gene with inaccessible chromatin to be transcribed¹⁶. Thus, we implement the capability to fit model 0 if desired, but fit only model 1 and model 2 by default. Model 1 and model 2 are both biologically plausible, and these different orders have biologically meaningful interpretations. We refer to model 1 as delayed transcriptional repression and model 2 as delayed chromatin repression. Within each model, a trajectory is defined by a set of eight core parameters θ , including three phase switching time points (transcriptional initiation time t_o , chromatin closing time t_c and transcriptional repression time t_r) and five rate parameters (chromatin opening rate α_{co} , chromatin closing rate α_{cc} , transcription rate α , splicing rate β and RNA degradation rate γ). There is also a fourth possible switch time t_o at which chromatin opening begins, but by excluding model 0, we can assume that $t_o = 0$ for all genes.

Model likelihood

We can formulate a probabilistic model to calculate the likelihood of the observed data for a gene under particular ODE parameters θ . To do this, we simply assume that the observations are independent and identically distributed and that the residuals are also normally distributed with mean given by the deterministic ODE solution and diagonal covariance. Because we scale the c , u and s values, we can further assume that the variance is the same in all directions. That is, if we define the ODE prediction as $\mathbf{f}(t_i, \theta) = \hat{\mathbf{x}}_i = (\hat{c}_i, \hat{u}_i, \hat{s}_i)$, then the distribution of the observed data $\mathbf{x}_i = (c_i, u_i, s_i)$ for each gene is

$$\mathbf{x}_i \sim \mathcal{N}(\mathbf{f}(t_i, \theta), \sigma^2 \mathbf{I}). \quad (18)$$

The negative log likelihood of all n observations is then

$$-\log \mathcal{L}(\theta) = \frac{3}{2} \log(2\pi\sigma^2) + \frac{1}{2n\sigma^2} \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{f}(t_i, \theta)\|^2. \quad (19)$$

We can infer the ODE parameters θ by maximum likelihood estimation, which is equivalent to minimizing the mean-squared error (MSE). The maximum likelihood estimate of σ^2 is the sample variance of the residuals along each coordinate. We can then rank genes by their likelihood to identify the genes best fit by the ODE model. We can also determine which model best explains the c , u , s values observed for a particular gene by comparing the MSE under model 1 and model 2.

Parameter estimation and latent time inference by expectation maximization

Both the cell times t and the ODE parameters are unknown, so we perform expectation-maximization to simultaneously infer them. The E-step involves determining the expected value of latent time for each cell given the current best estimate of the ODE parameters. Because inverting the three-dimensional ODEs analytically is not straightforward, we perform this time estimation by finding the time whose ODE prediction is nearest each data point, selecting the time from a vector of uniformly spaced time points (Implementation details section). In the M-step, we find the ODE parameters that maximize the data likelihood (equivalent to minimizing MSE) given the current time estimates for each cell. We use the Nelder-Mead⁴⁸ simplex algorithm to minimize MSE in each iteration.

Model predetermination and distinguishing genes with partial and complete dynamics

A gene does not have to complete a full trajectory within the measured cell population. In fact, for differentiating cells, we found that it is not uncommon for a gene to possess only an induction or repression phase, especially for differentially expressed cell-type marker genes. The three types of gene expression patterns (induction only, repression only and complete trajectory) can be directly inferred before fitting a model, thus avoiding ambiguous assignments near RNA phase transition points.

We used a combination of two methods for this purpose. The first method directly results from the assumptions of RNA velocity⁶: given a steady-state fit, cells in the induction phase

reside above the fit steady-state line while cells in the repression phase reside below the steady-state line. Thus, the ratio of sum of squared distances of cells on either side of the steady-state line is an indicator that can be used to determine the direction of the trajectory.

The second method incorporates low-dimensional coordinates (for example, from principal-component analysis or UMAP) as global information. We use UMAP coordinates by default, because these are often precomputed for visualization. Assuming that a gene possesses a complete trajectory, then at lower quantiles of its unspliced-spliced phase portrait, these cells are expected to have a bimodal pairwise distance pattern in the low-dimensional representation. Such a bimodal pattern indicates dissimilar populations, as some of these cells are in the early phase of induction, whereas the others have reached the late phase of repression. In contrast, for partial trajectories, cells at lower quantiles of the RNA phase portrait will have similar low-dimensional coordinates. Similarly, the unimodal or bimodal pattern can also be derived from the assumption that noise is normally distributed along the trajectory given by the ODE solution. We thus used a Gaussian mixture model to test if the distribution of pairwise distances among cells in a gene's lower quantile region is unimodal or bimodal, designating the trajectory being partial or complete, respectively. To be classified as a complete trajectory, the distance of the means between two Gaussians under bimodal distribution must exceed the globally measured variation (one standard deviation by default) of all pairwise distances on the low-dimensional coordinates for cells that express that gene, and the weight of the second, usually smaller Gaussian must pass a certain threshold (0.2 by default). The final assignment of partial or complete trajectory utilizes a combination of both methods (steady-state line ratio and bimodality), with the first method given priority (Extended Data Fig. 4g).

Additionally, whether a gene is better explained by model 1 or model 2 can be determined without actually fitting parameters under both models. To see how, note that the chromatin closing phase precedes transcriptional repression in model 1 but succeeds transcriptional repression in model 2. This implies that the highest chromatin accessibility values occur during the transcriptional induction phase for model 1 genes but during the repression phase for model 2 genes. Thus, the ratio of top chromatin values across the steady-state line can be used to determine whether each gene is best described by model 1 or model 2 before actually fitting the parameters. We implement this model predetermination as a default to speed up computation, but users can alternatively opt to fit both models and compare their losses instead (Extended Data Fig. 4h).

Parameter initialization

Parameters specifically related to RNA (α , β , γ and the RNA switch time interval) are initialized based on steady-state model as in scVelo. The rescaling factor for chromatin accessibility is initialized to 1, as the maximum observed accessibility is likely some value between 0 and 1. Other parameters can be found in the Implementation details section below.

We also initialize a scale factor for u . Here, we show that its value is closely related to the roundness of the unspliced-spliced portrait under steady-state assumptions. First, u and s

are both normalized to the range $[0, 1]$. Next, points of steady-state rate are found on the induction phase:

$$\begin{aligned}\frac{\alpha - \beta u_1}{\beta u_1 - \gamma s_1} &= \gamma \\ \frac{\alpha - u_1}{u_1 - \gamma s_1} &= \gamma \\ \alpha - u_1 &= \gamma u_1 - \gamma^2 s_1,\end{aligned}\tag{20}$$

$$\begin{aligned}u_1 &= \frac{\alpha + \gamma^2 s_1}{\gamma + 1} \\ u_1 &= \frac{a + a^2 s_1}{a + 1}\end{aligned}$$

where a is an unknown scalar and equals to the expected maximum of rescaled u . Similarly, on the repression phase,

$$\begin{aligned}\frac{-\beta u_2}{\beta u_2 - \gamma s_2} &= \gamma \\ \frac{-u_2}{u_2 - \gamma s_2} &= \gamma \\ -u_2 &= \gamma u_2 - \gamma^2 s_2.\end{aligned}\tag{21}$$

$$\begin{aligned}u_2 &= \frac{\gamma^2 s_2}{\gamma + 1} \\ u_2 &= \frac{a^2 s_2}{a + 1}\end{aligned}$$

Then, if we assume $u_1 = u_2 = \frac{1}{2}$ of maximum unspliced count, meaning the line connecting u_1 and u_2 is parallel to the s axis and at the same time crosses the middle point of u (due to symmetry), then

$$\begin{aligned}a + a^2 s_1 &= a^2 s_2 \\ S_2 - S_1 &= \frac{1}{a}.\end{aligned}\tag{22}$$

The rescale factor for u is therefore $s_2 - s_1$ around middle of u when s is normalized to range of $[0, 1]$. $u/(1/a) = a \times u$ and s are then used to initialize other parameters. Note that value of a is then further optimized during fitting (Extended Data Fig. 4i).

Implementation details

Estimating latent time.—A key implementation detail is how to estimate each cell's latent time given the ODE solution from the current parameters. Inverting the ODE solution is analytically challenging due to the complexity arising from a system of three ODEs. Thus, rather than pursuing an exact or approximate analytical solution to calculate time, we simply

maintain a set of anchor points uniformly spaced in time. For each cell, we then identify the nearest anchor point and assign the cell's time to the time of the anchor point. In more detail, we calculate the (c, u, s) values of the ODE solution at a specified number of uniformly distributed time points. Then we calculate pairwise distances from the observed cells to these anchor points. The shortest distance represents the residuals to the inferred trajectory, and the time of the anchor point is assigned to the cell. We found that 500–1,000 points are sufficient to capture the full trajectory dynamics. We restrict the time range to span from 0 to 20 h, consistent with scVelo's default setting.

Parameter initialization.—After determining trajectory direction and model to fit, expression values are shifted so that the minimum value starts from zero, and then they are scaled but not centered. RNA rate parameters are initialized based on the steady-state model: α is initialized as the mean of top-percentile u values to represent a gene's overall transcription potential⁷. The splicing rate β is initialized to 1, consistent with the steady-state model heuristic, and the degradation rate γ is obtained through linear regression of the top-percentile (u, s) values⁶. Chromatin rate α_c is initialized as $-\log(1 - c_{high})/t_{cc}$, where c_{high} is the mean accessibility of those cells with accessibility above average of all cells for that gene and t_{cc} is the chromatin closing switch time in the current grid search iteration. We initialize the RNA switch-off time using the explicit time-inversion procedure described in scVelo's method. To initialize the RNA switch-on time and chromatin switch-off time, we search over a grid of times 2 hours apart. The best initial switch time combinations are chosen based on MSE loss.

Estimating ODE parameters.—To fit and optimize parameters, we minimize the negative log likelihood (equivalent to MSE loss) using the Nelder-Mead downhill simplex method⁴⁸, implemented in the `scipy` minimize function. The Nelder-Mead algorithm performs a series of transformations on the model parameters, including reflection, shrinking and expansion to improve the fitting results. When fitting induction-only trajectories, only the first two phases (chromatin priming phase and coupled induction phase) are aligned to observations. When fitting repression-only trajectories, only the latter two phases are fit. To improve convergence speed, we minimize with respect to subsets of parameters at any time, holding the others fixed. This is similar to a block coordinate descent strategy. Within each iteration, we first update parameters exclusive to c , then parameters related to u and finally parameters affecting s . We found that five to ten iterations are sufficient for convergence in most cases. To ensure that the switch times occur in the proper order (for example, transcriptional induction precedes transcriptional repression), we opted to use switch intervals rather than switch time points as actual parameters. Thus, a model is guaranteed to be valid if all parameters are positive, with no other constraints needed.

Assigning cells to anchor points.—The trajectory constructed using a set of rate parameters is represented by a set of uniformly distributed anchor time points. By using the uniform distribution, we assume cells have equal prior probability to be measured at any given time point. The local sparsity of cells is determined by model parameters. We used KD-tree⁴⁹ from `scipy` to search for the closest anchor to each observation and its corresponding distance. Using anchor points also allows the model to mimic

the expected local sparsity of cells along the fit trajectories by encouraging anchors to concentrate near where cells concentrate in order to reduce small distance offsets caused by discrete representation of the trajectory. The main fitting steps are implemented in `multivelo.recover_dynamics_chrom` function.

Latent time normalization.—After fitting the models, because genes with partial fit trajectories result in a shorter total observed time range, violating the assumption that all genes share one time scale, the rate parameter set and the switch times are scaled down and up, respectively, so that time ranges from 0 to 20 hours. (Note that multiplying the time and dividing the rates by the same constant will result in identical trajectories.) This ensures that the time parameters from all genes are comparable. Switch times are shifted backward in time if the observable start of the trajectory happens later than 0 hour.

Other details.—The optimized rate parameters and time assignments are plugged back into the system of ODEs to obtain velocities for chromatin accessibility, unspliced RNA and spliced RNA for each cell. Our multi-omic velocity method is implemented in python. Many internal functions in our method have been accelerated with Numba. Distances, time assignments and velocity vectors are smoothed among nearest neighbors to mitigate the effect of measurement stochasticity.

Because multi-omic velocity is an upstream extension of the original RNA velocity model, it can be easily reduced to the RNA-only model by setting chromatin to be fully open (constant of 1) throughout the entire trajectory. Fitting this RNA-only model is then very similar to running the multi-omic model, but there will be no notion of the model 1 and model 2 distinction.

Likelihood ratio test for identifying genes with significant decoupling

We derived a likelihood ratio test (LRT) to determine whether a given gene has a statistically significant decoupling interval. Whether adding the decoupling phase significantly improves the likelihood of fit can be examined with a LRT. In this case, the reduced model has one fewer parameter (the length of the decoupling interval) compared to the full model. We use the following test statistic:

$$\lambda_{LR} = -2n(\ell(\theta_0) - \ell(\hat{\theta})), \quad (23)$$

where n is the number of cells, ℓ denotes the sample-size normalized log likelihood (Model likelihood section), θ_0 is the null/reduced parameter set without the decoupling interval and $\hat{\theta}$ contains the alternative/full model parameter set. By Wilks' theorem⁵⁰, the distribution of this test statistic can be asymptotically approximated by a χ^2 distribution with degrees of freedom equal to the difference in the number of free parameters between models (one in this case):

$$\lambda_{LR} \xrightarrow{D} \chi_1^2 \quad (24)$$

Because the decoupling interval primarily affects the fit of the chromatin data, which contributes only a small fraction of the overall likelihood, we used only the likelihood of chromatin fit to compute the test statistics. A significantly low P value indicates that the decoupling phase helps the model fit the chromatin significantly more accurately and that the null hypothesis of no decoupling interval can be rejected. This test is implemented in the `multivelo.LRT_decoupling` function.

Multi-omic stochastic velocity model

Building on the idea of the stochastic model in the `scVelo` paper, we developed a stochastic multi-omic velocity model and a parameter estimation strategy based on the steady-state assumption. A key idea of our stochastic model is that chromatin ‘breathes’ by rapidly switching between binary open and closed states⁵¹; at any instant of time, chromatin is either open or closed, but over a finite time interval, the chromatin accessibility can be interpreted as a probability between 0 and 1. Thus, we can model the instantaneous chromatin state using a Markov process with transition parameters for moving from a closed to an open state and vice versa. We further assume that the transcription rate at any moment in time t is the product of a rate constant α and the current chromatin accessibility c_t . Thus, transcription happens only during ‘bursts’ of chromatin accessibility, consistent with some experimental evidence that enhancers play a key role in transcriptional bursting⁵².

More formally, we model the changes in chromatin accessibility over time as a transition graph between two states $S = \{closed, open\}$ with transition probabilities $\mathbf{p} = (\alpha_{cc}, \alpha_{co})$ (Supplementary Fig. 2a). The stationary distribution and the expected value of the Markov chain are then given by

$$P(c_t) = \left[\frac{\alpha_{cc}}{\alpha_{cc} + \alpha_{co}} \quad \frac{\alpha_{co}}{\alpha_{cc} + \alpha_{co}} \right] \quad (25)$$

$$E[c_t] = \frac{\alpha_{co}}{\alpha_{cc} + \alpha_{co}}. \quad (26)$$

We can further write a system of stochastic ODEs describing the probability of c , u and s increasing or decreasing during an infinitesimal time step dt :

$$\begin{aligned} P(c_{t+dt} = c_t + 1) &= \alpha_{co}[c_t = 0]dt \\ P(c_{t+dt} = c_t - 1) &= \alpha_{cc}[c_t = 1]dt \\ P(c_{t+dt} = c_t) &= ((1 - \alpha_{co})[c_t = 0] + (1 - \alpha_{cc})[c_t = 1])dt. \\ P(u_{t+dt} = u_t + 1, s_{t+dt} = s_t) &= \alpha c_t dt \\ P(u_{t+dt} = u_t - 1, s_{t+dt} = s_t + 1) &= \beta u_t dt \\ P(u_{t+dt} = u_t, s_{t+dt} = s_t - 1) &= \gamma s_t dt \end{aligned} \quad (27)$$

Although this system cannot be readily solved analytically, we can derive moments in closed form. Following a similar argument to the one given in the `scVelo` paper, we obtain the following equations:

$$\frac{d\langle u_i^2 \rangle}{dt} = \alpha\langle c_i \rangle + 2\alpha\langle c_i u_i \rangle + \beta\langle u_i \rangle - 2\beta\langle u_i^2 \rangle \quad (28)$$

$$\frac{d\langle s_i^2 \rangle}{dt} = \beta\langle u_i \rangle + 2\beta\langle u_i s_i \rangle + \gamma\langle s_i \rangle - 2\gamma\langle s_i^2 \rangle, \quad (29)$$

where the $\langle \rangle$ operator denotes expectation. Bergen et al. showed that if we further assume that the system is at steady state, we can use these moment equations to estimate the ratio between transcription and degradation rates⁷. The resulting parameter estimates depend on a covariance between u and s (second moment), which confers additional robustness compared to the deterministic steady-state parameter estimates. Comparing our stochastic model with scVelo's, the moments now additionally depend on the chromatin state c . Thus, to perform parameter estimation, we assign cells into open or closed chromatin states for each gene and use the open cells to estimate parameters. In practice, we assign cells with chromatin accessibility one standard deviation above the mean to the open state. These cells are then used to estimate the transcription/degradation ratio by linear regression as in scVelo. This strategy will result in a shift in the steady-state location, particularly for genes with long decoupling intervals.

When we compare the hypothetical steady-state locations proposed in the MultiVelo dynamical model, $(c_\infty^{(k)}, u_\infty^{(k)}, s_\infty^{(k)}) = (k_c, \alpha^{(k)}k_c/\beta, \alpha^{(k)}k_c/\gamma)$, and RNA velocity $(u_\infty^{(k)}, s_\infty^{(k)}) = (\alpha^{(k)}/\beta, \alpha^{(k)}/\gamma)$, one can also see that they are directly affected by k_c , an indicator for open chromatin state. In addition, depending on whether a gene is model 1 or model 2, the steady-state location shifts earlier or later in time compared to the RNA-only model. Supplementary Fig. 2b shows examples of how the steady state changes in practice when we use the multi-omic stochastic model. The steady-state methods are implemented in the function `multivelo.velocity_chrom`.

Post-fitting analyses

Bergen et al.⁷ have developed great downstream analyses methods for RNA velocity in the scVelo toolkit. Because our method is a direct extension of the dynamical model to multi-omic data, many of scVelo's methods can be applied with only a change of arguments. Our main method replaces the scVelo functions `tl.recover_dynamics` and `tl.velocity`. In this paper, scVelo's `tl.velocity_graph` with total-normalized spliced velocity vectors computed from our multi-omic method was used to obtain a transition matrix between cells based on cosine similarity between a cell's velocity vector and expression differences. We used `pl.velocity_embedding_stream` to embed and plot velocity streams onto UMAP coordinates. Computation of global latent time among cells and genes is implemented in `tl.latent_time`.

We performed DTW using the `dtw` R package^{53,54}. First, the accessibilities or expressions of cells were aggregated to 20 equal-sized bins based on either their gene time (for *Wnt3* in the skin dataset) or latent time (for human brain motifs) and then maximum-normalized to the same range of [0, 1]. For motifs, a rolling mean of three bins was applied to the RNA and motif counts to smooth the curves. Next, we added a zero to each end of the time series

to ensure that the starting and ending values of each time series matched. Then we used DTW to find the best alignment (local for *Wnt3* or global for motifs) between the two time series with Euclidean distance penalty. We then calculated time lags by simply subtracting the times of the aligned points. When many-to-one mappings occurred in global alignments, we averaged the time lags across all points mapped to the same time. For SNP time analysis, both the SNP accessibilities and log RNA expressions were aggregated to 100 equal-sized bins. We then calculated the time lag as the time difference between the time bins with highest values in the two modalities.

Generation of simulated data

We first simulated genes independently to assess MultiVelo's ability to recover the underlying parameters and model 1 versus model 2 distinctions (Extended Data Fig. 5). In this analysis, 1,000 genes were simulated with various rate parameters, switch times, time sequences and models (1 and 2). $\alpha_c = \alpha_o, \alpha, \beta$ and γ values were generated from multivariate log-normal distributions with mean $[-2, 2, 0, 0]$ and variance $[0.5, 1, 0.3, 0.3]$, with a small covariance of 0.01 between α_c, α and β . Four switch intervals were random chosen from $[1,4], [1,9], [1,9]$ and $[1,9]$ and scaled to give a time range from 0 to 20 hours. The model (model 1 versus model 2) was sampled uniformly at random. Cell times were sampled from a Poisson distribution. Noise was added to each cell with diagonal covariances of $[\max(c)^2/90, \max(u)^2/90, \max(s)^2/90]$. The accuracy of loss-based and predetermined model decisions were separately computed.

We next performed a simulation to test MultiVelo's ability to recover latent time under varying noise levels and differing numbers of complete, induction and repression genes (Supplementary Fig. 3). In this analysis, 2,000 cells were simulated with times equally spaced between 0 and 20. The four rate parameters (α_c, α, β and γ) were sampled from log-normal distributions with mean $[-2.5, 2, -0.5, 0]$ and variance of $[0.3, 1, 0.3, 0.3]$. We chose these mean and variance values to mimic the data distribution observed in the 10x mouse brain dataset. Each of the 500 genes was then randomly assigned to complete, induction only or repression only based on the specified sampling ratio among the three types. The noise was added to each cell so that the variance-mean ratio of each modality was matched between simulation and the mouse brain. Here, the variance was obtained using the signed distances to trajectory anchors, and mean expressions of cells above $0.05 \times \max(\text{modality})$ were computed. The variance of chromatin-wise noise, which serves as an indicator of ATAC-seq stochasticity and sparsity, was scaled by a factor of 0.2 or 5. After constructing the cell-gene matrix, UMAP was computed for each setting without further preprocessing steps, as the simulated dataset is designed to resemble the real data after normalization and smoothing.

Working with unpaired RNA and ATAC datasets

To test whether the model can work for scRNA and snATAC datasets derived from separate cells, we used the anchors algorithm from the Seurat package⁴². We treated the RNA and ATAC measurements in the 10x mouse brain dataset as separately sequenced modalities. Although the RNA, unspliced and spliced matrices stayed unaltered, the gene-aggregated ATAC matrix was imputed and paired to the RNA cells by the anchor transfer method

implemented in the FindTransferAnchors and TransferData functions. This new ATAC-seq matrix was then input into MultiVelo. We found that this procedure gave the most similar and clean three-dimensional phase portraits to the ones from the true cell pairs (Supplementary Fig. 1). Transferring RNA into ATAC space instead gave worse results (not shown).

Development environment and runtime

The Python package was developed on Arch Linux with Intel Core i7-9750H (12 threads) and 32 GB memory. We summarized the runtime and memory usage statistics in Supplementary Table 1. The main dynamics recovery function finished running in parallel in 40 min, 69 min, 124 min and 40 min for the four biological datasets tested (mouse brain, mouse skin, human HSPC and human brain). The maximum python memory usage statistics when running the multivelo.recover_dynamics_chrom function were 857 MiB, 1,602 MiB, 2,921 MiB and 1,100 MiB. The memory increments when running the main function were 481.5 MiB, 1,136.5 MiB, 2,293.2 MiB and 660.6 MiB. The memory usage usually goes up with the number of threads requested. Downsampling the cells or lowering the number of anchors can further reduce runtime and memory.

Preprocessing of data, weighted nearest neighbors (WNNs) and smoothing

10x embryonic E18 mouse brain.—Filtered expression matrix for ATAC-seq, feature linkage file, as well as position-sorted RNA alignment (BAM) file of E18 mouse embryonic brain data of around 5,000 cells were downloaded from 10x Genomics website (CellRanger ARC 1.0.0). Total, unspliced and spliced RNA reads were separately quantified using the Velocity run10x command. The resulting loom file was read into python as an AnnData object and preprocessed with scanpy and scVelo to perform filtering, normalization and nearest neighbor assignment. Next, clusters were computed using the Leiden⁵⁵ algorithm. Cell types were manually annotated based on expression of known marker genes^{56–59}. We then excluded interneurons, Cajal-Retzius and microglia cell populations for our downstream analyses, because these cell types are not actively differentiating. We then reprocessed the raw counts of subset clusters, which consists of more than 3,000 remaining cells, with scVelo. The unspliced and spliced reads were neighborhood smoothed (averaged) by scVelo's pp.moments method with 30 principal components among 50 neighbors. The downloaded feature linkage file contains correlation information for gene-peak pairs of genomic features across cells. We first collected all distal putative enhancer peaks (not in promoter or gene body regions) with 0.5 correlation with either promoter accessibility or gene expression that were annotated to the same gene or within 10 kilobase pairs of that gene. We then aggregated these enhancer peaks with 10x annotated promoter peaks for the corresponding genes, as a single chromatin accessibility modality to boost chromatin signal. These aggregated accessibility values were then normalized using the term frequency-inverse document frequency (TF-IDF) method²⁸. (Note that during fitting, chromatin values are normalized to [0, 1], so using other total-count based normalization will produce identical results.)

Due to the increased sparsity of ATAC-seq data, the neighborhood graph and clustering results based solely on peaks is often noisy and unreliable. Seurat group recently

developed a method to compute neighborhood assignments for simultaneously measured multi-modality data in the Seurat V4 toolkit, which they called weighted nearest neighbor, or WNN⁶⁰. The WNN method learns weights of each cell in either modality based on its predictive power by neighboring cells in each of the modalities, so that both RNA and ATAC information can be incorporated when assigning neighbors. We used 50 WNNs obtained from Seurat for each cell to smooth the aggregated and normalized chromatin peak values. Our WNN analysis followed the recommended steps in Seurat V4 vignette for 10x RNA + ATAC. We thus obtained three matrices containing chromatin accessibility, unspliced and spliced counts. Shared cell barcodes and genes were filtered among matrices and resulted in 3,365 cells and 936 highly variable genes, and these matrices were then used for dynamical modeling.

SHARE-seq mouse skin (hair follicle) data.—The quantified ATAC-seq expression matrix, raw ATAC-seq fragments file and cell annotations of SHARE-seq mouse skin dataset⁹ were downloaded from the Gene Expression Omnibus (GEO). The RNA alignment BAM file as well as UMAP coordinates for TAC, IRS, medulla and hair shaft cuticle/cortex cell populations used in the SHARE-seq manuscript were obtained directly from the authors. We run Velocity to quantify unspliced and spliced counts, and the RNA AnnData object was further preprocessed with scanpy/scVelo for the four cell types of interest. In R, the chromatin fragment file was used to construct a gene activity matrix by aggregating peaks onto gene coordinates using the GeneActivity function in Signac⁶¹. Domain of regulatory chromatin (DORC) is defined as chromatin regions that contain clusters of peaks that are highly correlated with gene expressions in SHARE-seq's analysis. A list of computed DORCs coordinates was downloaded from its supplementary material section. These coordinates were output to the bed format, and we extracted fragments together with their corresponding cell barcodes that overlap with these DORCs regions. A peak expression matrix for DORCs was constructed with LIGER's⁶² makeFeatureMatrix method. The gene activity and DORCs counts were then merged in python to form a single chromatin modality. Similar to brain data, this matrix underwent TF-IDF normalization and WNN smoothing. A total of 6,436 cells and 962 genes participated in the downstream analyses. When computing the Spearman correlation with pseudotime based methods, only genes with likelihood higher than 0.07 were kept, resulting in 140 velocity genes. This filtering step ensures a fair comparison with scVelo by using a small set of high-quality genes.

Human HSPCs.—Purified human CD34⁺ cells were purchased from the Fred Hutch Hematology Core B. Freshly thawed cells from a single donor were either immediately prepared for single-cell processing or maintained at 37°C with 5% CO₂ in Stemspan II medium supplemented with 100 ng ml⁻¹ stem cell factor, 100 ng ml⁻¹ thrombopoietin, 100 ng ml⁻¹ Flt3 ligand (all from Stemcell Technologies) and 100 ng ml⁻¹ insulin-like growth factor binding protein 2 (R&D Systems) for seven days. HSPCs were prepared according to the manufacturer's '10x Genomics Nuclei Isolation Single Cell multi-ome ATAC + Gene Expression Sequencing' demonstrated protocol. Briefly, cells were washed in PBS supplemented with 0.04% BSA and filtered with a Flowmi Cell Strainer (Bel-Art) (day 0) or sorted using the Sony SH800 cell sorter (Sony Biotechnologies) (day 7). Nuclei

were isolated following the ‘Low Cell Input Nuclei Isolation’ sub-protocol and immediately processed using the Chromium Next GEM Single Cell Multiome + Gene Expression kit.

10x filtered expression matrices, Velocyto computed unspliced and spliced counts and feature linkage and peak annotation files from CellRanger ARC 2.0.0 were read into python to construct RNA and ATAC AnnData objects. Filtering, normalization and variable-gene selection were performed following scVelo’s online tutorial. Because HSPCs are rapidly proliferating, we noticed systematic differences in cell cycle stage across the set of cells. The cell cycle scores for both G2M and S phases, computed using scVelo’s `tl.score_genes_cell_cycle` function were then regressed out of the RNA expression matrices with scanpy’s `pp.regress_out` function (Extended Data Fig. 2b). Note that the regression did not change unspliced and spliced counts. Then gene expression scaling was performed. ATAC peaks were aggregated and normalized using the same procedure as described for the 10x mouse brain. Joint filtering between RNA and ATAC resulted in 11,605 cells and 1,000 genes. RNA expression was smoothed by scVelo’s `pp.moments` with 30 principal components and 50 neighbors. Leiden found 11 clusters. Cell types were assigned based on canonical HSPC markers^{63–67}. The chromatin accessibility matrix was WNN smoothed with 50 neighbors computed using Seurat. Then the RNA and ATAC objects were input to our dynamical function with default parameters. We relaxed the likelihood threshold for velocity genes (used for computing the velocity graph) to 0.02 compared to the default of 0.05 due to noisiness of this dataset.

To find complete genes in each of the lineages from HSC toward GMPs (myeloid), erythrocytes and platelets, we subset cells of each specific lineage and select known complete genes as those genes that have higher unspliced and spliced expressions in the progenitor populations leading to each of the terminal cell types. We then ran the model predetermination algorithm based on peak chromatin accessibility as described in the previous section. The genes predicted as model 1 and model 2 for each lineage are then merged with duplicates removed, and we performed gene ontology enrichment analysis (GORilla⁶⁸) using all sequenced genes as the background set.

Preprocessed bulk chromatin immunoprecipitation with sequencing (ChIP-seq) peaks of H3K4me3, H3K4me1 and H3K27ac for CD34⁺ HSPCs were downloaded from GEO:GSE70677⁶⁹. Peaks were mapped to genes with Homer⁷⁰. Known complete genes in the myeloid and erythroid lineages were grouped together, and predicted model 1 and model 2 genes were extracted. Scores of peaks associated with the same genes were aggregated. Wilcoxon rank-sum test was used to compute significance.

The day 0 Multiome HSPC sample was integrated with the day 7 sample (Fig. 5) using the Seurat⁴² anchors workflow (FindIntegrationAnchors and IntegrateData) to remove technical differences due to batch effects. Each of the full RNA expression matrices, unspliced and spliced matrices, as well as gene-aggregated ATAC-seq matrix, were integrated and imputed independently between the two samples. A total of 2,000 RNA genes and 5,000 ATAC genes were used for integration with the day 7 sample as the reference dataset. The two raw ATAC-seq peak matrices were integrated using the IntegrateEmbeddings method in Signac⁶¹, and together with the integrated full RNA matrix, the WNNs were found. These outputs were

then read into python and processed using the standard MultiVelo procedure. The UMAP resulting from the Seurat anchors RNA integration was used for plotting (Extended Data Fig. 2d). Seurat cluster results were used for cell types.

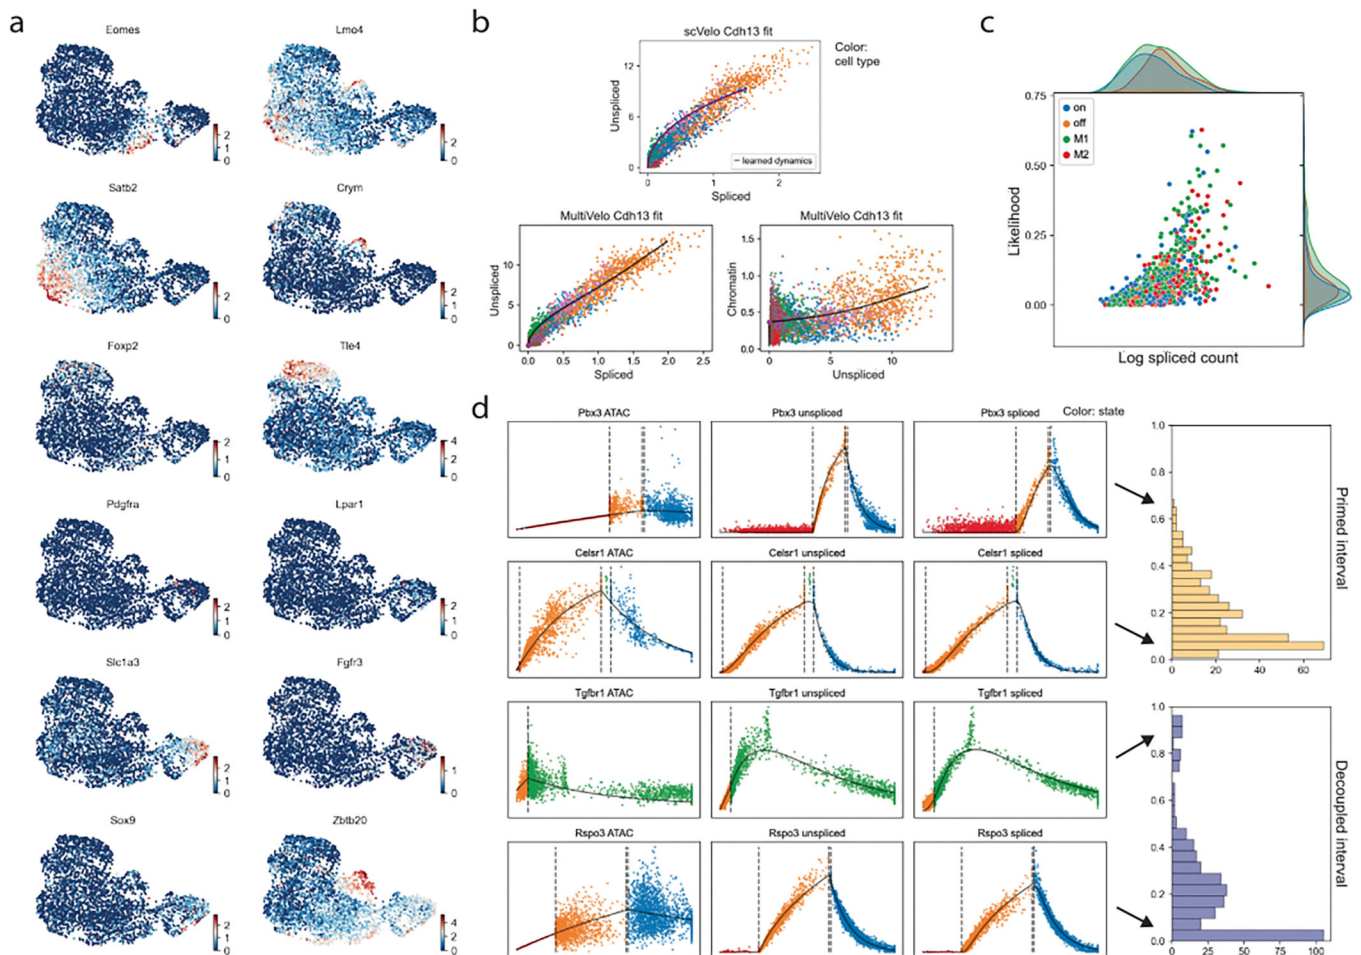
Human cerebral cortex.—We obtained the multi-omic RNA, unspliced, spliced and ATAC-seq peak files from the authors. The ATAC peak matrix contains consensus peaks of nonoverlapping uniform 500 bp length. After initial clustering, we observed a severe batch effect in one of the three samples. We thus decided to remove this third sample and perform all downstream analyses with the two remaining samples (dc2r2_r1 and dc2r2_r2). We renamed the clusters from the original paper as follows based on marker gene expression: RG → RG/Astro, nIPC/GluN1 → nIPC/ExN, GluN3 → ExM, GluN2 → ExUp, GluN4 and GluN5 → ExDp⁵⁷. Peaks were annotated to genes with Homer⁷⁰. We considered peaks within 10,000 bp of transcription start sites as promoter peaks. A list of peak-gene links and correlations were downloaded from the supplementary material and aggregated to promoter peaks if the correlation exceeded 0.4. After filtering the RNA and ATAC matrices, 4,693 cells and 919 genes were left and input to model fitting.

TF motif profiles were computed with chromVAR⁴⁴ on the JASPAR2020 database⁷¹ using all consensus peaks. The background-corrected deviation z-scores were used as normalized motif accessibilities, and the values were smoothed with WNN. Then TF genes appearing in the variable-gene list (after internal filtering by the dynamical function) were extracted for time-lag analysis, which resulted in 30 known motifs.

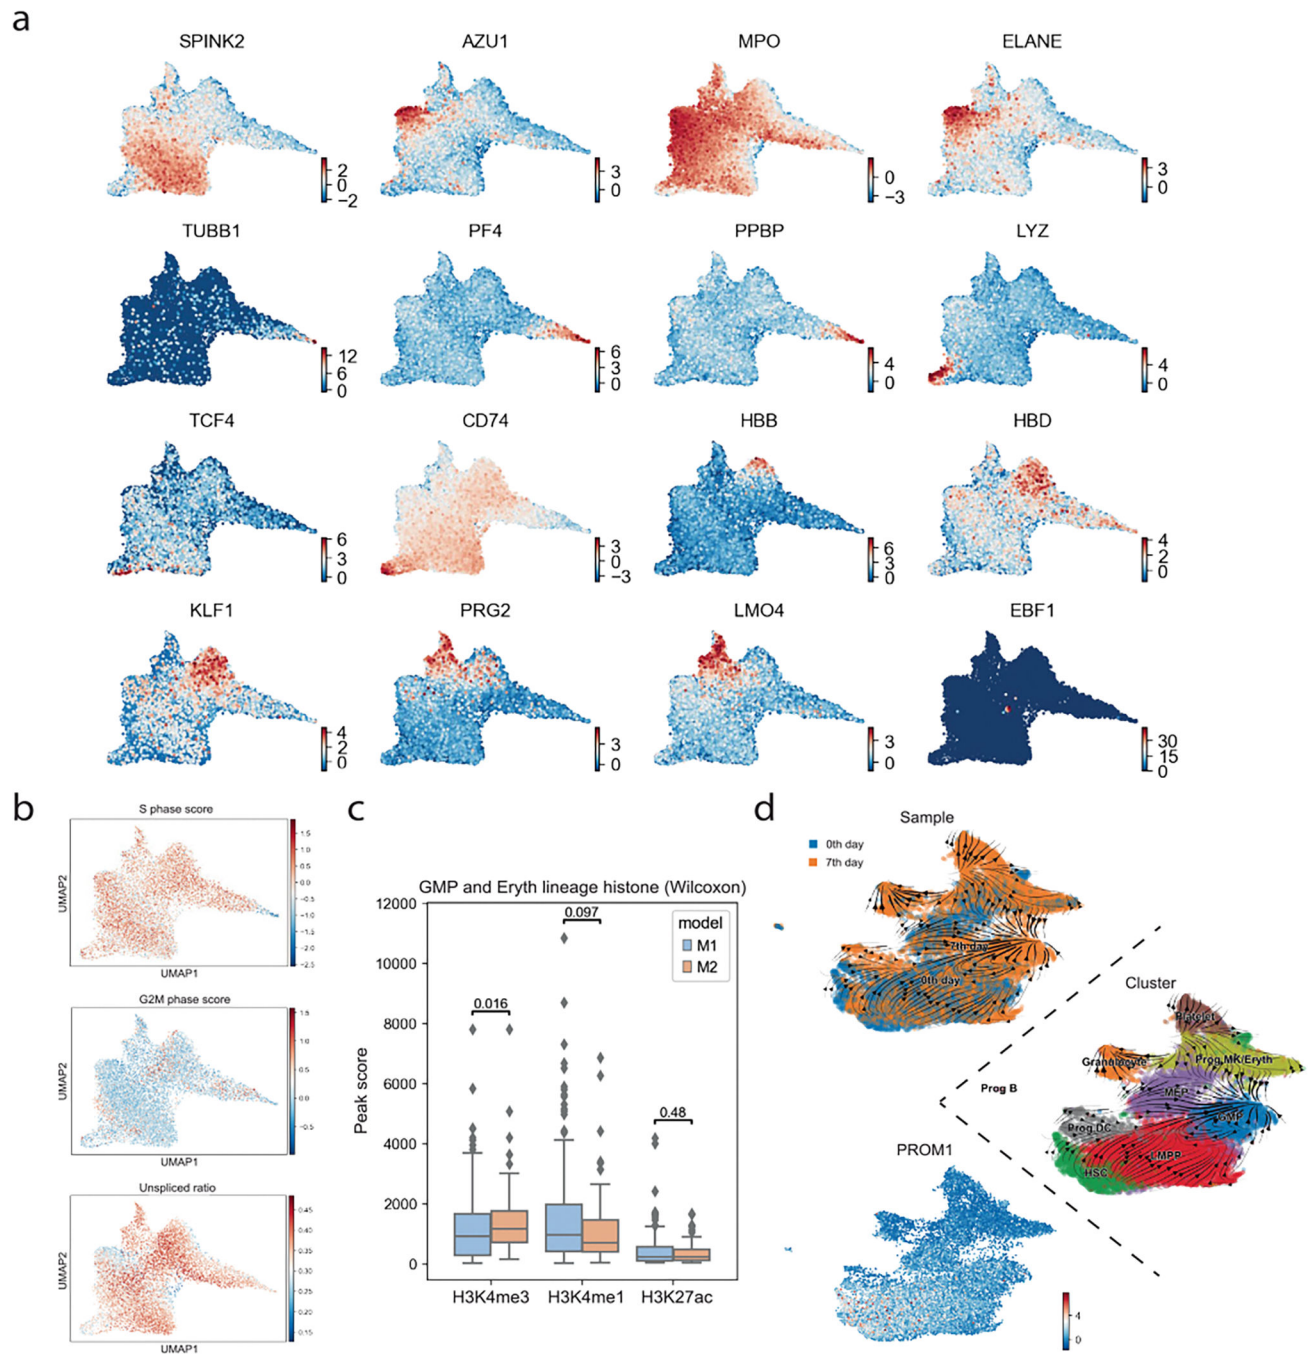
All mental or behavioral disorder-associated SNPs (EFO_0000677) were downloaded from the Ensembl GWAS Catalog. The list contains 6,968 SNPs, and filtering for overlap with consensus peaks linked to the top genes resulted in 757 SNPs. Each SNP's accessibility was quantified as the count of all ATAC fragments that overlap a 400-bp bin centered on the SNP location. The accessibility matrix was normalized by library size and smoothed by WNN neighbors.

Downstream target genes for *EGR1*, *EOMES*, *FOXP2* and *PBX3* were found through a literature search. A number of these targets were confirmed by CHIP-seq experiments^{72–75}. The time delay between expression of a TF and expression of its downstream targets was quantified with DTW.

Extended Data

**Extended Data Fig. 1 | Additional figures for mouse brain dataset.**

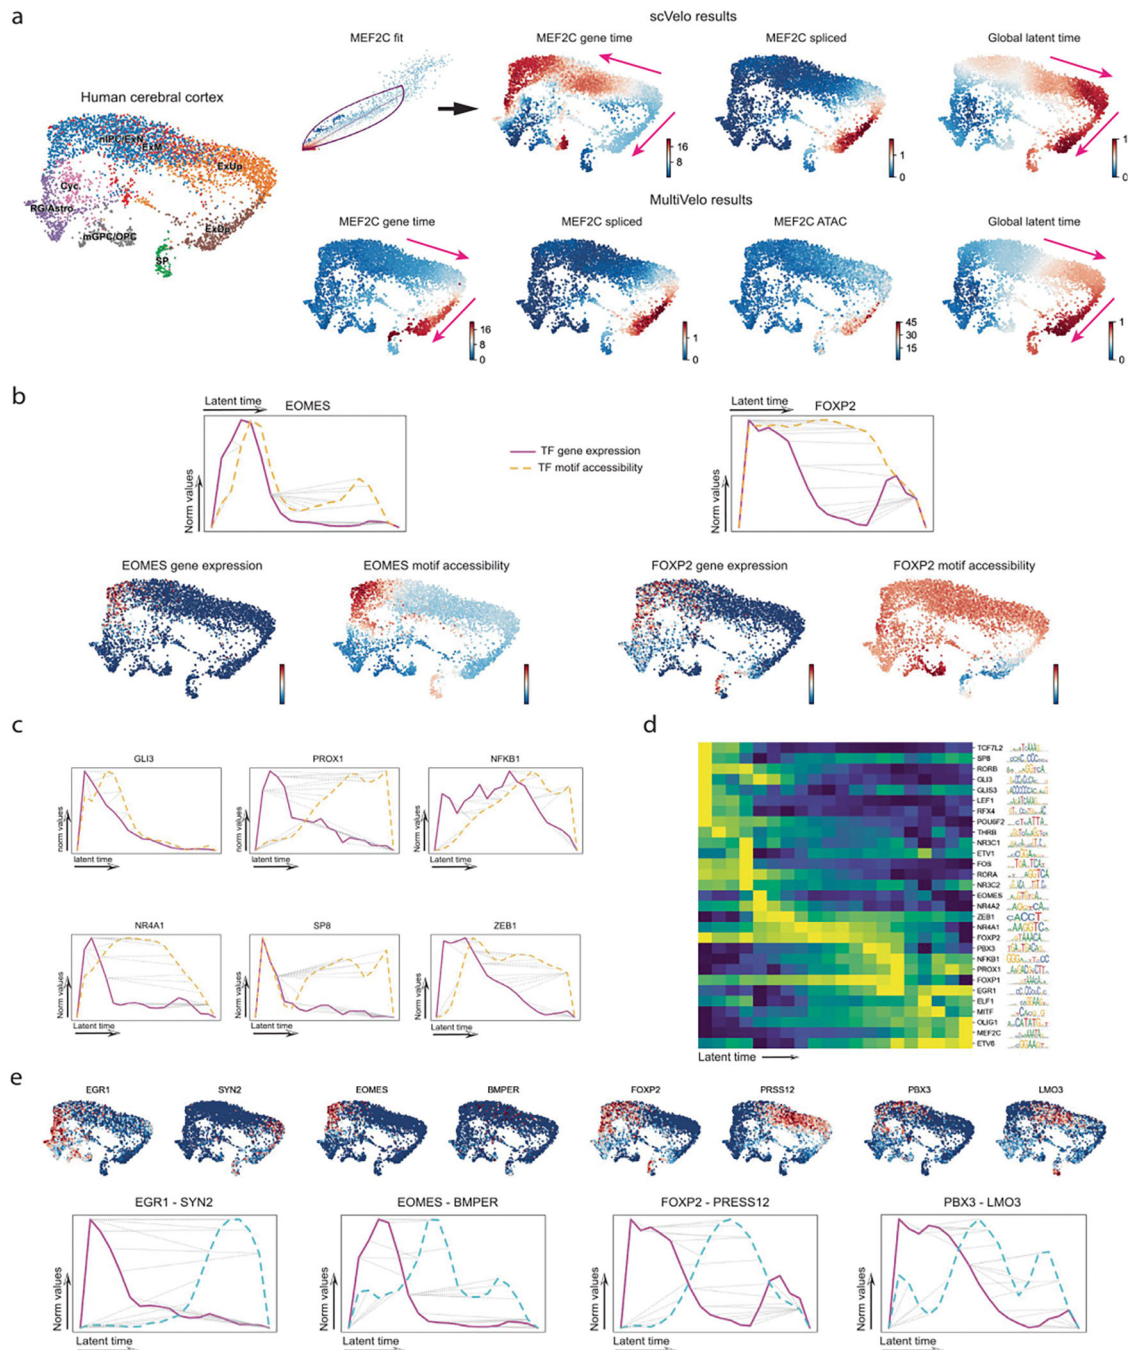
a, Canonical marker gene expression for embryonic mouse brain cell types. **b**, Comparison of *Cdh13* fits from scVelo and MultiVelo. An elevating transcription rate due to opening of chromatin produces a more linear fit and better captures the observed phase portrait. **c**, Scatterplot of gene likelihood against log total spliced count. Gene likelihood is not significantly affected by model assignment or trajectory type. Likelihood does increase with spliced count, as this usually indicates higher quality or highly variable genes. **d**, Switch times can be used to rank genes by the length of priming and decoupling intervals. Each range is scaled to 1 with outliers ($n = 1$) removed. **Top two rows**: Histogram of priming intervals. *Pbx3* and *Celsr1* possess short and long priming phases, respectively. **Bottom two rows**: Histogram of decoupled intervals. While *Rspo3* has a short decoupling phase with few cells within, *Tgfb1*'s decoupling phase extends from RNA induction to RNA repression, and up to the end of the trajectory.



Extended Data Fig. 2 |. Additional figures for HSPC dataset.

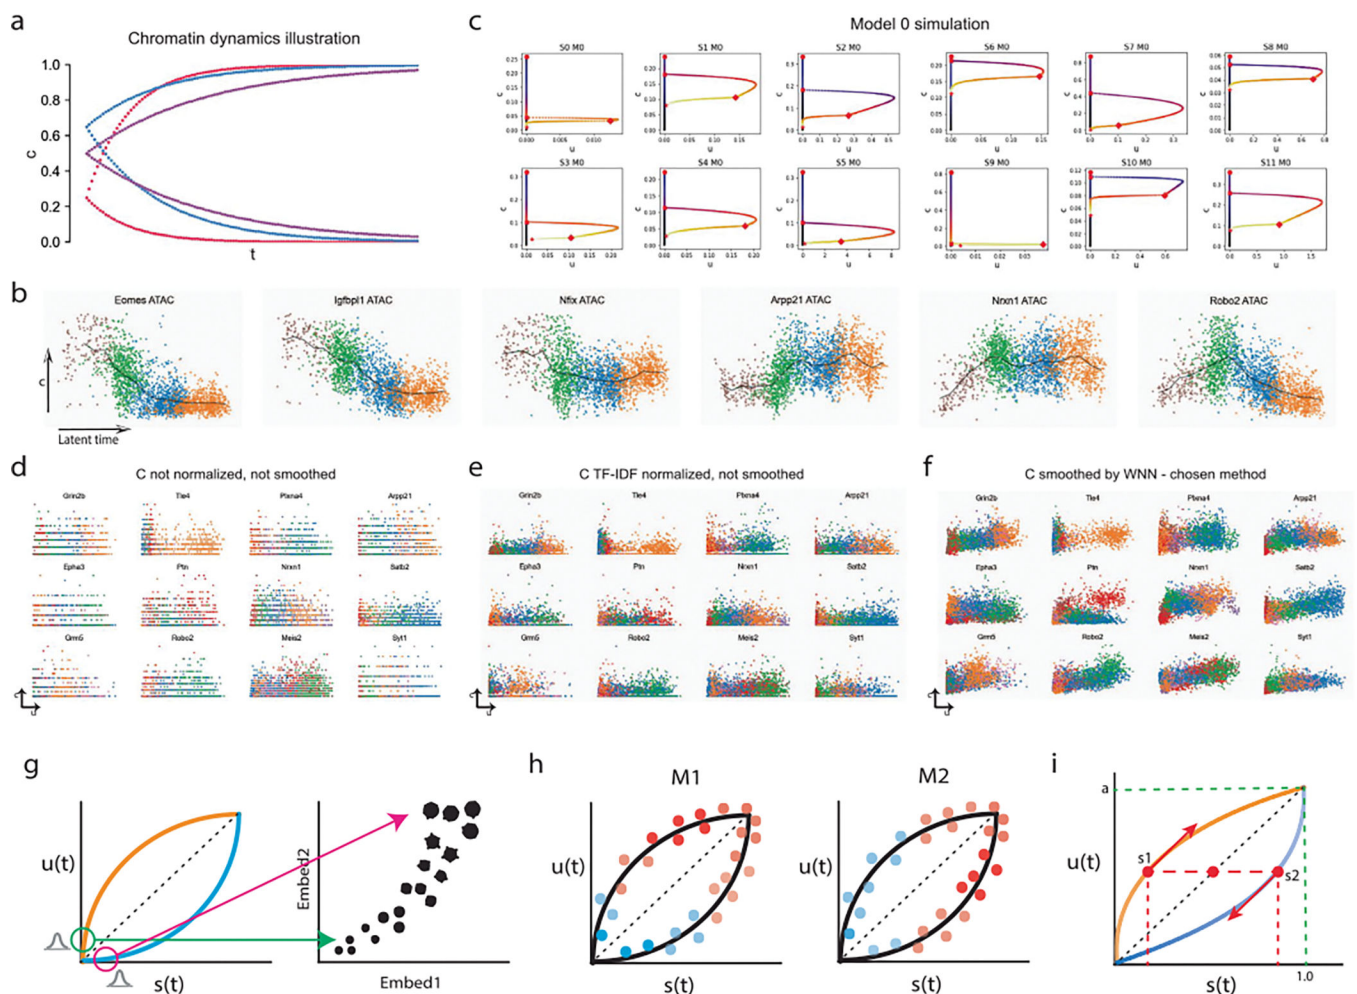
a. Canonical marker gene expression for HSPCs. **b.** Cell cycle (S phase and G2M phase) scores and total unspliced ratio ($U/(U+S)$) plotted on UMAP coordinates. These factors were regressed out of the total RNA expression (but not the unspliced and spliced counts) during the preprocessing step as they do not appear to be cell-type or lineage specific. **c.** Box plots of histone modification levels from bulk ChIP-seq of FACS-purified HSCs (center line, median; box, Q1 and Q3; whiskers, 1.5x IQR; points, outliers). Each point in the box plot represents the sum of histone modification signal at chromatin accessibility peaks linked

to a Model 1 or Model 2 gene. P-values are from a one-sided Wilcoxon rank-sum test. **d**, Velocity stream plot from MultiVelo analysis of Day 0 and Day 7 HSPC samples (**Top**). The majority of arrows go from Day 0 stem cells toward more differentiated Day 7 cells. UMAP coordinates colored by cell-type labels (**Middle**). UMAP coordinates colored by expression of CD133 (*PROM1*), an HSPC marker (**Bottom**).



Extended Data Fig. 3 | Additional figures for human brain dataset.

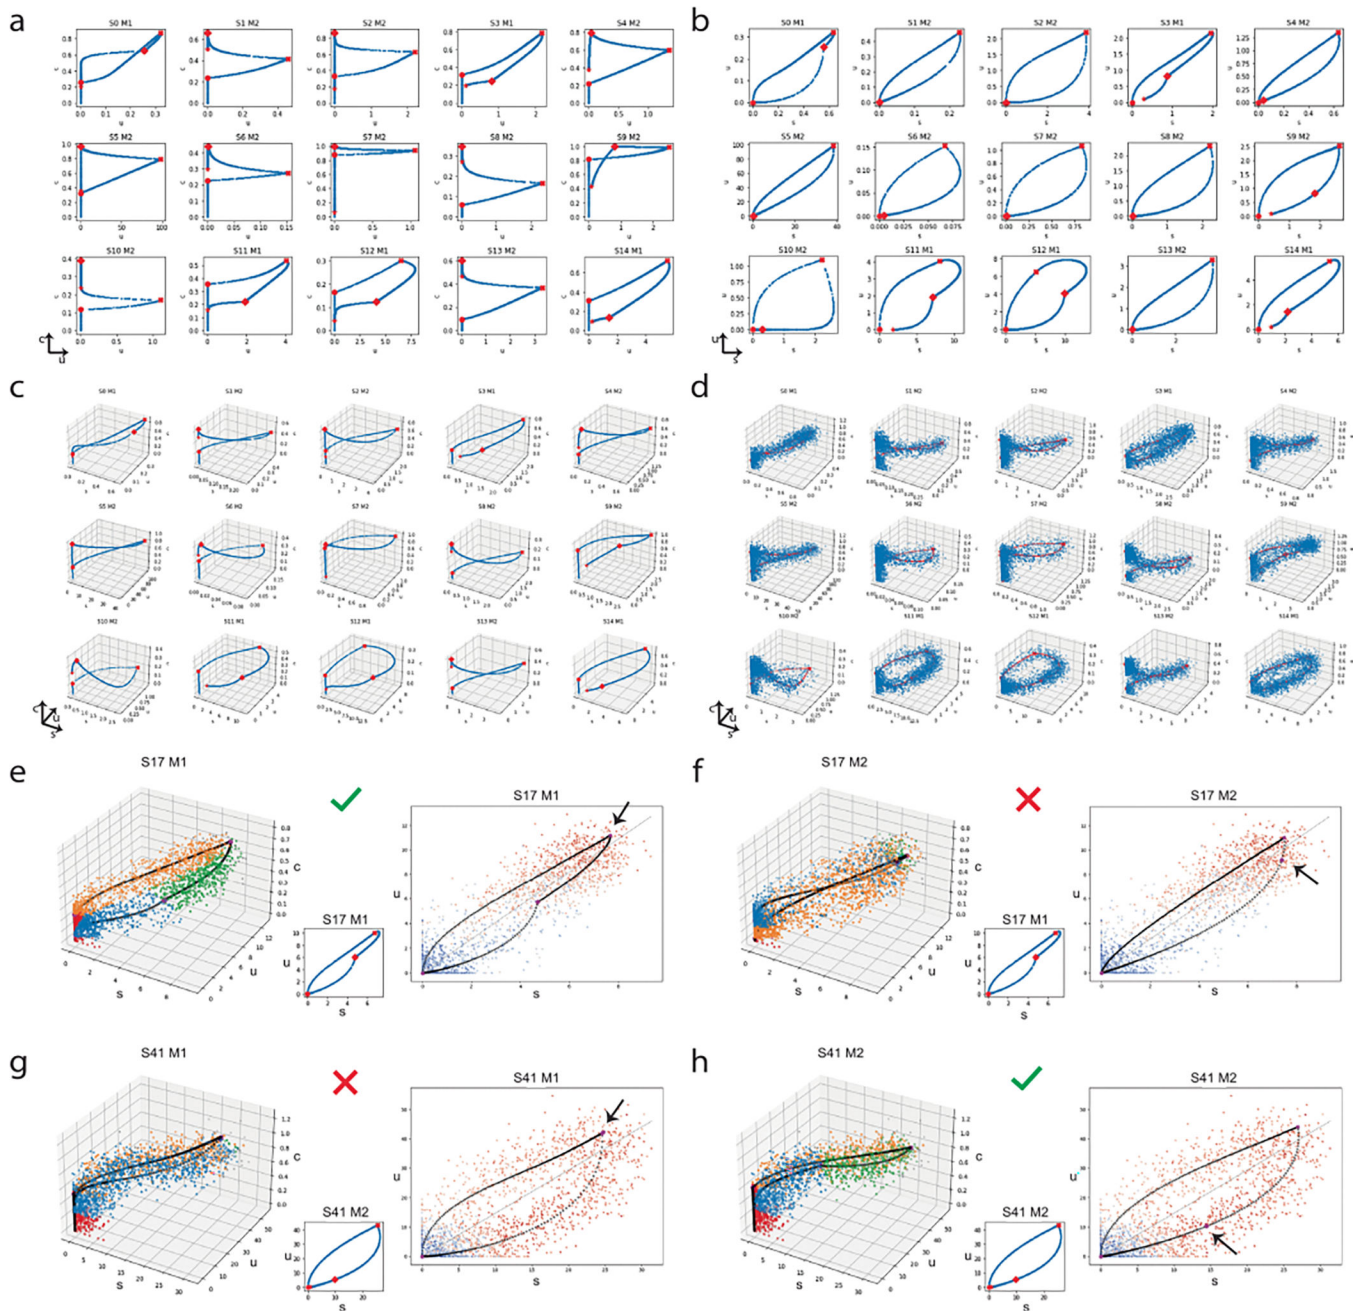
a, Validation of the direction of *MEF2C*. **Left**: UMAP with cell types. **Top**: scVelo's *MEF2C* fit produces inconsistency between gene time and global latent time. **Bottom**: MultiVelo's results show consistent progression from nIPC to deeper layer (ExDp). **b**, DTW and UMAP results for *EOMES* and *FOXP2* transcription factors. **c**, Additional motif DTW alignment results showing time lags between TF gene expression and corresponding motif accessibility. **d**, The accessibility of TF motifs binned across latent time. The latent time scale was split into 20 equalized bins, and the average motif accessibility of cells in each bin was computed and plotted. The motif sequence logos (downloaded from jaspar2020.genereg.net) are shown next to the TF names. **e**, Time-lag analysis of transcription factors and the expression of their validated downstream target genes. **Top**: UMAP plots colored by TF and target gene expression. **Bottom**: Line plots of TF and target gene expression, with correspondences from DTW alignment shown as dotted lines. Magenta: TFs. Cyan: target genes.



Extended Data Fig. 4 | Chromatin dynamics, Model 0, the necessity of chromatin preprocessing, and pre-fitting illustrations.

a, Chromatin dynamics illustration: chromatin opening and closing are modeled as asymptotically approaching fully opened (1) or fully closed (0) starting from any initial

value. **b**, Chromatin accessibility change as a function latent time inferred by scVelo using only the RNA portion of the 10X multiome mouse brain dataset (colored by mouse brain cell types). Black lines connect the mean accessibilities within 20 equal-sized windows. The shapes of the ATAC trends are qualitatively very similar to the ODE model we propose. **c**, Simulation of Model 0 samples. The long delay between chromatin closing and transcription initiation is unlikely to happen in real biological systems. In the rare cases when high chromatin accessibility but low expression or high expression but low accessibility pattern is observed, it is likely due to technical issues such as dropout or background noise. **d**, The need for normalization as a preprocessing step for ATAC-seq. **e**, The need for smoothing as a preprocessing step for ATAC-seq. **f**, Chromatin accessibility results after peak-to-gene aggregation, TF-IDF normalization, and WNN smoothing. It is the same as Fig. S2E. **g**, Illustration of bi-modal expression pattern for complete genes. Cells at the lower quantile can be far apart in low-dimensional embedded expression space. **h**, Simplified illustration of model predetermination reasoning. The highest chromatin accessibility region appears in different RNA phases in M1 and M2 genes. **i**, Illustration of the internal unspliced modality rescaling factor initialization.



Extended Data Fig. 5 |. Simulation study to assess parameter estimation and model determination.

A total of 1000 genes were simulated with various parameters for both model 1 and model 2. **a**, C-U view of noiseless simulations of 2000 time-points in the 0–20 hr range. **b**, U-S view of noiseless simulations from A. **c**, 3D view of noiseless simulations from A. **d**, Noise added to simulated points to mimic real data. **e**, **f**, Model 1 and Model 2 fits for the same simulated gene (S17). The likelihood is higher under Model 1, consistent with the ground truth. **e**, **Left**: 3D view of the fit Model 1 trajectory colored by states, along with predicted switch time points. **Middle**: simulation with ground-truth switch times. **Right**: U-S view

of fitted trajectory colored by $\log(c)$. **f**, Similar to **e**, but the fit shown is for Model 2 (the incorrect model). **g**, **h**, Model fits for simulated gene S41, similar to **e** and **f**, but this time, Model 2 is the ground truth model. MultiVelo correctly identifies the sample to be Model 2 with accurate switch time estimations. The model assignments of 985/1000 samples were correctly predicted based on likelihood.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

This work was supported by National Institutes of Health grants R01AI149669 to K.L.C. and J.D.W., R01HG010883 to J.D.W., F31AI155047 to M.C.V., training grants T32GM070449 to C.L. and T32GM007315 to M.C.V. and additional funding provided by the Rackham Regents Fellowship to M.C.V. The HSPC sample was provided by Cooperative Center of Excellence in Hematology grant DK106829. We thank J. Li, S.C.J. Parker, Y. Gu and members of the Collins lab for helpful discussions.

Data availability

10x embryonic mouse brain dataset can be accessed at the 10x website at <https://www.10xgenomics.com/resources/datasets/fresh-embryonic-e-18-mouse-brain-5-k-1-standard-1-0-0>.

SHARE-seq⁹ mouse skin dataset can be found at the GEO (GSE140203).

Human brain multi-ome dataset⁴³ can be found at the GEO (GSE162170) and the authors' GitHub page.

ChIP-seq peaks for bulk CD34+ HSPC⁶⁹ were downloaded from the GEO (GSE70677).

The processed files of the newly sequenced 10x Multiome HSPC samples are available at the GEO (GSE209878). Raw sequences were uploaded to dbGaP phs002915.v1.p1 under restricted access due to patient privacy concerns. Source data are provided with this paper.

References

1. Cao J et al. The single-cell transcriptional landscape of mammalian organogenesis. *Nature* 566, 496–502 (2019). [PubMed: 30787437]
2. Street K et al. Slingshot: cell lineage and pseudotime inference for single-cell transcriptomics. *BMC Genom.* 19, 477 (2018).
3. Ji Z & Ji H Tscan: pseudo-time reconstruction and evaluation in single-cell RNA-seq analysis. *Nucleic Acids Res.* 44, 117 (2016). [PubMed: 26350215]
4. Setty M et al. Characterization of cell fate probabilities in single-cell data with Palantir. *Nat. Biotechnol.* 37, 451–460 (2019). [PubMed: 30899105]
5. Welch JD, Hartemink AJ & Prins JF SLICER: inferring branched, nonlinear cellular trajectories from single cell RNA-seq data. *Genome Biol.* 17, 106 (2016). [PubMed: 27215581]
6. La Manno G et al. RNA velocity of single cells. *Nature* 560, 494–498 (2018). [PubMed: 30089906]
7. Bergen V, Lange M, Peidli S, Wolf FA & Theis FJ Generalizing RNA velocity to transient cell states through dynamical modeling. *Nat. Biotechnol.* 38, 1408–1414 (2020). [PubMed: 32747759]

8. Gorin G, Svensson V & Pachter L Protein velocity and acceleration from single-cell multiomics experiments. *Genome Biol.* 21, 39 (2020). [PubMed: 32070398]
9. Ma S et al. Chromatin potential identified by shared single-cell profiling of RNA and chromatin. *Cell* 183, 1103–1116 (2020). [PubMed: 33098772]
10. Tedesco M et al. Chromatin velocity reveals epigenetic dynamics by single-cell profiling of heterochromatin and euchromatin. *Nat. Biotechnol.* 40, 235–244 (2021). [PubMed: 34635836]
11. Chen S, Lake BB & Zhang K High-throughput sequencing of the transcriptome and chromatin accessibility in the same cell. *Nat. Biotechnol.* 37, 1452–1457 (2019). [PubMed: 31611697]
12. Weake VM & Workman JL Inducible gene expression: diverse regulatory mechanisms. *Nat. Rev. Genet.* 11, 426–437 (2010). [PubMed: 20421872]
13. Rada-Iglesias A et al. A unique chromatin signature uncovers early developmental enhancers in humans. *Nature* 470, 279–283 (2011). [PubMed: 21160473]
14. Ostuni R et al. Latent enhancers activated by stimulation in differentiated cells. *Cell* 152, 157–171 (2013). [PubMed: 23332752]
15. Lara-Astiaso D et al. Chromatin state dynamics during blood formation. *Science* 345, 943–949 (2014). [PubMed: 25103404]
16. Klemm SL, Shipony Z & Greenleaf WJ Chromatin accessibility and the regulatory epigenome. *Nat. Rev. Genet.* 20, 207–220 (2019). [PubMed: 30675018]
17. Merkle FT, Tramontin AD, García-Verdugo JM & Alvarez-Buylla A Radial glia give rise to adult neural stem cells in the subventricular zone. *Proc. Natl. Acad. Sci. U.S.A.* 101, 17528–17532 (2004). [PubMed: 15574494]
18. Hansen DV, Lui JH, Parker PRL & Kriegstein AR Neurogenic radial glia in the outer subventricular zone of human neocortex. *Nature* 464, 554–561 (2010). [PubMed: 20154730]
19. Pollen AA et al. Molecular identity of human outer radial glia during cortical development. *Cell* 163, 55–67 (2015). [PubMed: 26406371]
20. Nadarajah B & Parnavelas JG Modes of neuronal migration in the developing cerebral cortex. *Nat. Rev. Neurosci.* 3, 423–432 (2002). [PubMed: 12042877]
21. Englund C et al. Pax6, Tbr2, and Tbr1 are expressed sequentially by radial glia, intermediate progenitor cells, and postmitotic neurons in developing neocortex. *J. Neurosci.* 25, 247–251 (2005). [PubMed: 15634788]
22. Mayer S et al. Multimodal single-cell analysis reveals physiological maturation in the developing human neocortex. *Neuron* 102, 143–158 (2019). [PubMed: 30770253]
23. Tirosh I et al. Dissecting the multicellular ecosystem of metastatic melanoma by single-cell RNA-seq. *Science* 352, 189–196 (2016). [PubMed: 27124452]
24. Arnold SJ et al. The T-box transcription factor Eomes/Tbr2 regulates neurogenesis in the cortical subventricular zone. *Genes Dev.* 22, 2479–2484 (2008). [PubMed: 18794345]
25. Vasistha NA et al. Cortical and clonal contribution of Tbr2 expressing progenitors in the developing mouse brain. *Cereb. Cortex* 25, 3290–3302 (2014). [PubMed: 24927931]
26. McEvelly RJ, de Diaz MO, Schonemann MD, Hooshmand F & Rosenfeld MG Transcriptional regulation of cortical neuron migration by POU domain factors. *Science* 295, 1528–1532 (2002). [PubMed: 11859196]
27. Zahr SK et al. A translational repression complex in developing mammalian neural stem cells that regulates neuronal specification. *Neuron* 97, 520–537 (2018). [PubMed: 29395907]
28. Cusanovich DA et al. A single-cell atlas of in Vivo mammalian chromatin accessibility. *Cell* 174, 1309–1324 (2018). [PubMed: 30078704]
29. Lake BB et al. Integrative single-cell analysis of transcriptional and epigenetic states in the human adult brain. *Nat. Biotechnol.* 36, 70–80 (2018). [PubMed: 29227469]
30. Ranzoni AM et al. Integrative single-cell RNA-seq and ATAC-seq analysis of human developmental hematopoiesis. *Cell Stem Cell* 28, 472–487.e7 (2021). [PubMed: 33352111]
31. Jia G et al. Single cell RNA-seq and ATAC-seq analysis of cardiac progenitor cell transition states and lineage settlement. *Nat. Commun.* 9, 4877 (2018). [PubMed: 30451828]
32. Zhang B & Hsu Y-C Emerging roles of transit-amplifying cells in tissue regeneration and cancer. *Wiley. Interdiscip. Rev. Dev. Biol.* 6, e282 (2017).

33. Coifman RR et al. Geometric diffusions as a tool for harmonic analysis and structure definition of data: diffusion maps. *Proc. Natl. Acad. Sci. U.S.A.* 102, 7426–7431 (2005). [PubMed: 15899970]
34. Millar SE et al. WNT signaling in the control of hair growth and structure. *Dev. Biol.* 207, 133–149 (1999). [PubMed: 10049570]
35. Berndt DJ & Clifford J Using dynamic time warping to find patterns in time series. In *Proceedings of the 3rd International Conference on Knowledge Discovery and Data Mining, AAAIWS'94*, 359–370 (AAAI Press, 1994).
36. Buenrostro JD et al. Integrated single-cell analysis maps the continuous regulatory landscape of human hematopoietic differentiation. *Cell* 173, 1535–1548 (2018). [PubMed: 29706549]
37. Orkin SH & Zon LI Hematopoiesis: an evolving paradigm for stem cell biology. *Cell* 132, 631–644 (2008). [PubMed: 18295580]
38. Görgens A et al. Revision of the human hematopoietic tree: granulocyte subtypes derive from distinct hematopoietic lineages. *Cell Rep.* 3, 1539–1552 (2013). [PubMed: 23707063]
39. Laurenti E & Göttgens B From haematopoietic stem cells to complex differentiation landscapes. *Nature* 553, 418–426 (2018). [PubMed: 29364285]
40. Pellin D et al. A comprehensive single cell transcriptional landscape of human hematopoietic progenitors. *Nat. Commun.* 10, 2395 (2019). [PubMed: 31160568]
41. Bergen V, Soldatov RA, Peter V K & Theis FJ RNA velocity—current challenges and future perspectives. *Mol. Syst. Biol.* 17, e10282 (2021). [PubMed: 34435732]
42. Stuart T et al. Comprehensive integration of single-cell data. *Cell* 177, 1888–1902.e21 (2019). [PubMed: 31178118]
43. Trevino AE et al. Chromatin and gene-regulatory dynamics of the developing human cerebral cortex at single-cell resolution. *Cell* 184, 5053–5069.e23 (2021). [PubMed: 34390642]
44. Schep AN, Wu B, Buenrostro JD & Greenleaf WJ chromVAR: inferring transcription-factor-associated accessibility from single-cell epigenomic data. *Nat. Methods* 14, 975–978 (2017). [PubMed: 28825706]
45. Tong A, Huang J, Wolf G, van Dijk D & Krishnaswamy S Trajectorynet: a dynamic optimal transport network for modeling cellular dynamics. In *Proceedings of the 37th International Conference on Machine Learning, ICML'20 (JMLR.org, 2020)*.
46. Schiebinger G et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell* 176, 928–943 (2019). [PubMed: 30712874]
47. Qiao C & Huang Y Representation learning of RNA velocity reveals robust cell transitions. *Proc. Natl. Acad. Sci. U.S.A.* 118, e2105859118 (2021). [PubMed: 34873054]
48. Nelder JA & Mead R A simplex method for function minimization. *Comput. J.* 7, 308–313 (1965).
49. Bentley JL Multidimensional binary search trees used for associative searching. *Commun. ACM* 18, 509–517 (1975).
50. Wilks SS The large-sample distribution of the likelihood ratio for testing composite hypotheses. *Ann. Math. Stat.* 9, 60–62 (1938).
51. Kim J, Sheu KM, Cheng QJ, Hoffmann A & Enciso G Stochastic models of nucleosome dynamics reveal regulatory rules of stimulus-induced epigenome remodeling. *Cell Rep.* 40, 111076 (2022). [PubMed: 35830792]
52. Larsson AJM et al. Genomic encoding of transcriptional burst kinetics. *Nature* 565, 251–254 (2019). [PubMed: 30602787]
53. Giorgino T Computing and visualizing dynamic time warping alignments in R: the dtw package. *J. Stat. Softw.* 31, 1–24 (2009).
54. Tormene P, Giorgino T, Quaglioni S & Stefanelli M Matching incomplete time series with dynamic time warping: an algorithm and an application to post-stroke rehabilitation. *Artif. Intell. Med.* 45, 11–34 (2009). [PubMed: 19111449]
55. Traag VA, Waltman L & van Eck NJ From Louvain to Leiden: guaranteeing well-connected communities. *Sci. Rep.* 9, 5233 (2019). [PubMed: 30914743]
56. Batiuk MY et al. Identification of region-specific astrocyte subtypes at single cell resolution. *Nat. Commun.* 11, 1220 (2020). [PubMed: 32139688]

57. Polioudakis D et al. A single-cell transcriptomic atlas of human neocortical development during mid-gestation. *Neuron* 103, 785–801.e8 (2019). [PubMed: 31303374]
58. Cahoy JD et al. A transcriptome database for astrocytes, neurons, and oligodendrocytes: a new resource for understanding brain development and function. *J. Neurosci. Res.* 28, 264–278 (2008).
59. Hochgerner H, Zeisel A, Lönnerberg P & Linnarsson S Conserved properties of dentate gyrus neurogenesis across postnatal development revealed by single-cell RNA sequencing. *Nat. Neurosci.* 21, 290–299 (2018). [PubMed: 29335606]
60. Hao Y et al. Integrated analysis of multimodal single-cell data. *Cell* 184, 3573–3587.e29 (2021). [PubMed: 34062119]
61. Stuart T, Srivastava A, Madad S, Lareau CA & Satija R Single-cell chromatin state analysis with Signac. *Nat. Methods* 18, 1333–1341 (2021). [PubMed: 34725479]
62. Welch JD et al. Single-cell multi-omic integration compares and contrasts features of brain cell identity. *Cell* 177, 1873–1887 (2019). [PubMed: 31178122]
63. Karamitros D et al. Single-cell analysis reveals the continuum of human lympho-myeloid progenitor cells. *Nat. Immunol.* 19, 85–97 (2018). [PubMed: 29167569]
64. Yáñez A et al. Granulocyte-monocyte progenitors and monocyte-dendritic cell progenitors independently produce functionally distinct monocytes. *Immunity* 47, 890–902.e4 (2017). [PubMed: 29166589]
65. Villani A-C et al. Single-cell RNA-seq reveals new types of human blood dendritic cells, monocytes, and progenitors. *Science* 356, eaah4573 (2017). [PubMed: 28428369]
66. Grajkowska LT et al. Isoform-specific expression and feedback regulation of E protein TCF4 control dendritic cell lineage specification. *Immunity* 46, 65–77 (2017). [PubMed: 27986456]
67. Paul F et al. Transcriptional heterogeneity and lineage commitment in myeloid progenitors. *Cell* 163, 1663–1677 (2015). [PubMed: 26627738]
68. Eden E, Navon R, Steinfeld I, Lipson D & Yakhini Z GOrilla: a tool for discovery and visualization of enriched GO terms in ranked gene lists. *BMC Bioinform.* 10, 48 (2009).
69. Romano O et al. Transcriptional, epigenetic and retroviral signatures identify regulatory regions involved in hematopoietic lineage commitment. *Sci. Rep.* 6, 24724 (2016). [PubMed: 27095295]
70. Heinz S et al. Simple combinations of lineage-determining transcription factors prime cis-regulatory elements required for macrophage and B cell identities. *Mol. Cell* 38, 576–589 (2010). [PubMed: 20513432]
71. Fornes O et al. JASPAR 2020: update of the open-access database of transcription factor binding profiles. *Nucleic Acids Res.* 48, D87–D92 (2019).
72. Koldamova R et al. Genome-wide approaches reveal EGR1-controlled regulatory networks associated with neurodegeneration. *Neurobiol. Dis.* 63, 107–114 (2014). [PubMed: 24269917]
73. Tosic J et al. Eomes and Brachyury control pluripotency exit and germ-layer segregation by changing the chromatin state. *Nat. Cell Biol.* 21, 1518–1531 (2019). [PubMed: 31792383]
74. Spiteri E et al. Identification of the transcriptional targets of FOXP2, a gene linked to speech and language, in developing human brain. *Am. J. Hum. Genet.* 81, 1144–1157 (2007). [PubMed: 17999357]
75. Golonzhka O et al. Pbx regulates patterning of the cerebral cortex in progenitors and postmitotic neurons. *Neuron* 88, 1192–1207 (2015). [PubMed: 26671461]

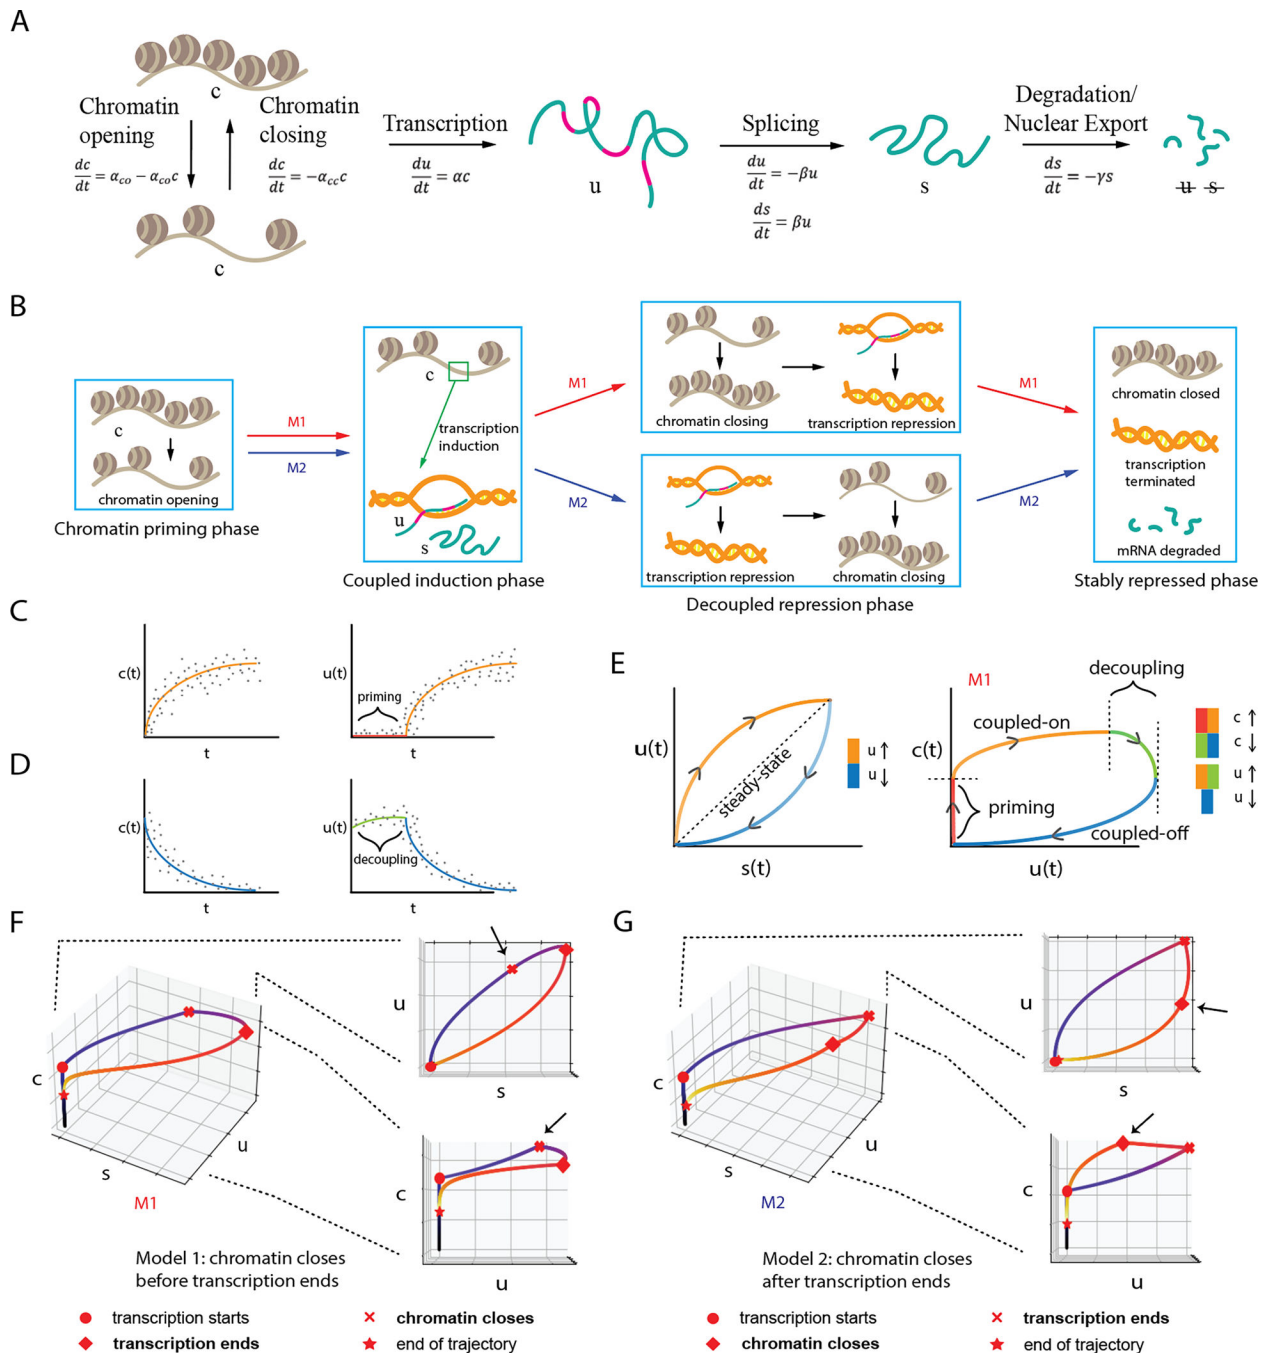


Fig. 1 | Schematic of MultiVelo approach.

a. System of three ODEs summarizes the temporal relationship among c , u and s values during the gene expression process. **b.** Two different models (abbreviated as M1 and M2) describe two potential orderings of chromatin and RNA state changes. Chromatin accessibility starts to drop before transcriptional repression begins in M1, and the reverse happens in M2. **c.** Priming occurs when chromatin opens before transcription initiates. **d.** Decoupling occurs when chromatin closing and transcription repression begin at different times (example shown for model 1). **e.** Phase portraits predicted by the ODE model, showing

the four possible states each gene can occupy. Gene expression and chromatin accessibility are coupled in the orange and blue states and decoupled in the red and green states. **f,g**, Simulated (c, u, s) values for a model 1 (f) and a model 2 (g) gene.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

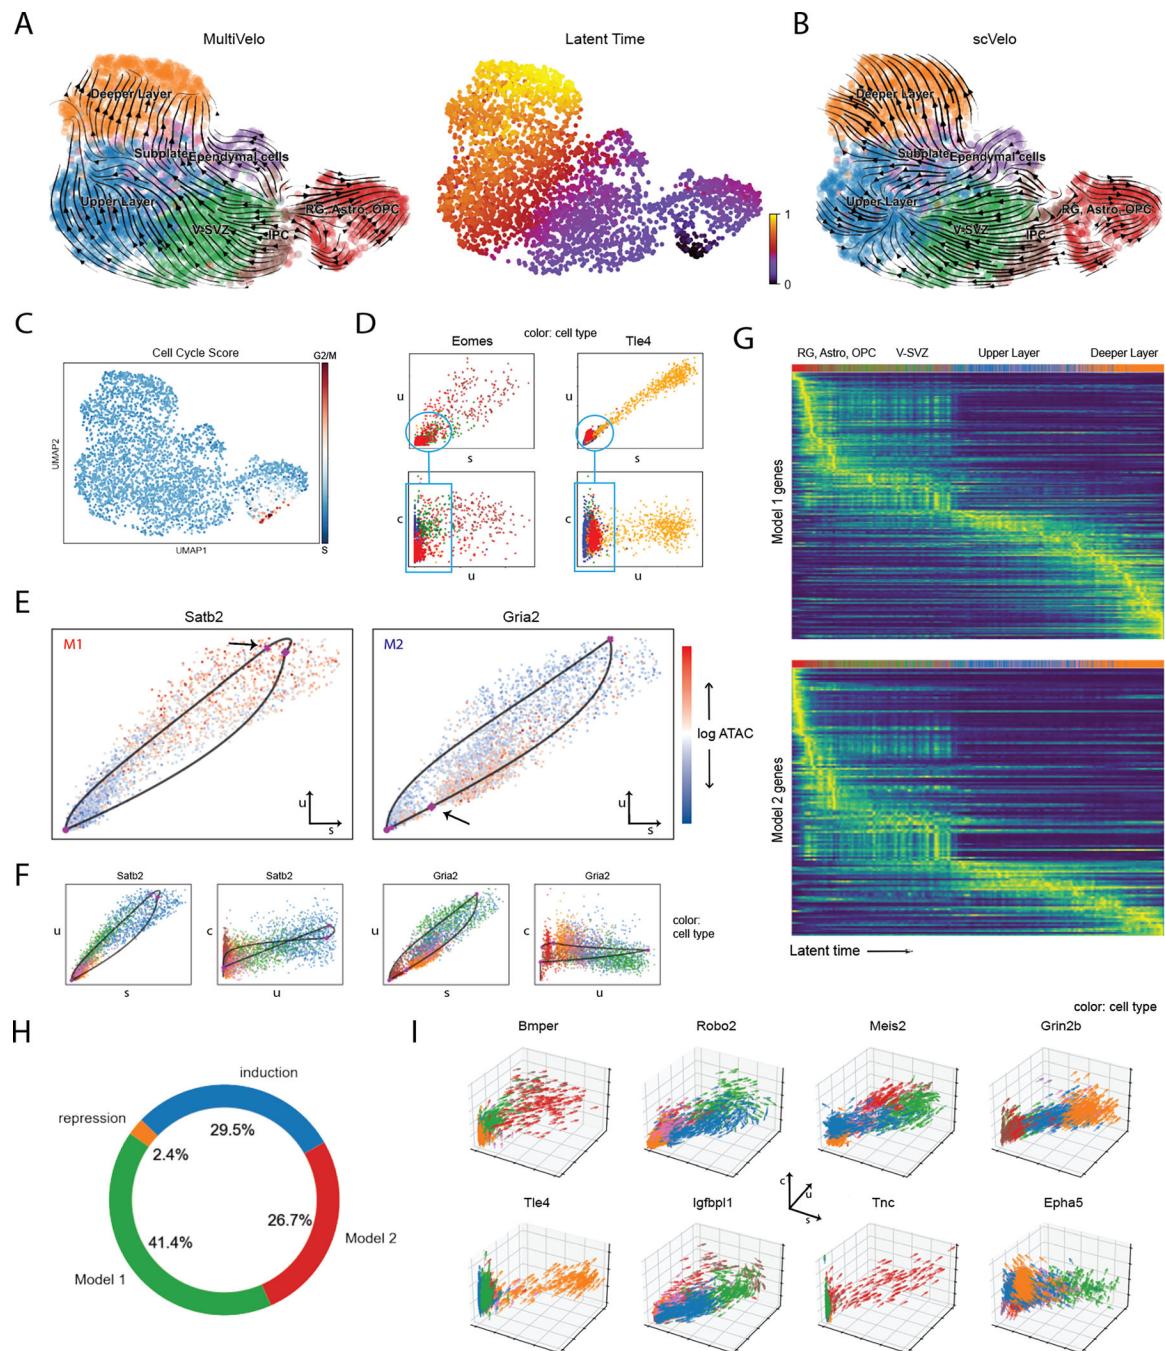


Fig. 2 | MultiVelo reveals two distinct mechanisms of gene regulation.

a, UMAP coordinates with stream plot of velocity vectors (left) and latent time (right) from MultiVelo. OPC, oligodendrocyte progenitor cells; Astro, astrocytes; V-SVZ, ventricular–subventricular zone. **b**, Stream plot of velocity vectors estimated from RNA only by scVelo. **c**, Cell cycle score indicating active dividing and cycling population (arrow). **d**, Chromatin values better separate differentiating cells when chromatin opening precedes transcription. **e**, RNA phase portraits (u versus s) colored by c values show clear differences between model 1 (left) and model 2 (right) genes. **f**, Additional phase portraits for the genes shown

in panel e. **g**, Heatmaps of model 1 and model 2 gene expressions as a function of latent time. Color represents smoothed spliced counts. Model 2 genes tend to achieve highest expression earlier in latent time than model 1 genes. **h**, Relative proportion of each type of kinetics across all fit genes ($n = 865$). Note that genes with partial kinetics (induction only or repression only) cannot be identified as model 1 or model 2. **i**, MultiVelo predicts 3D velocity vectors, which can be visualized as three-dimensional arrow plots.

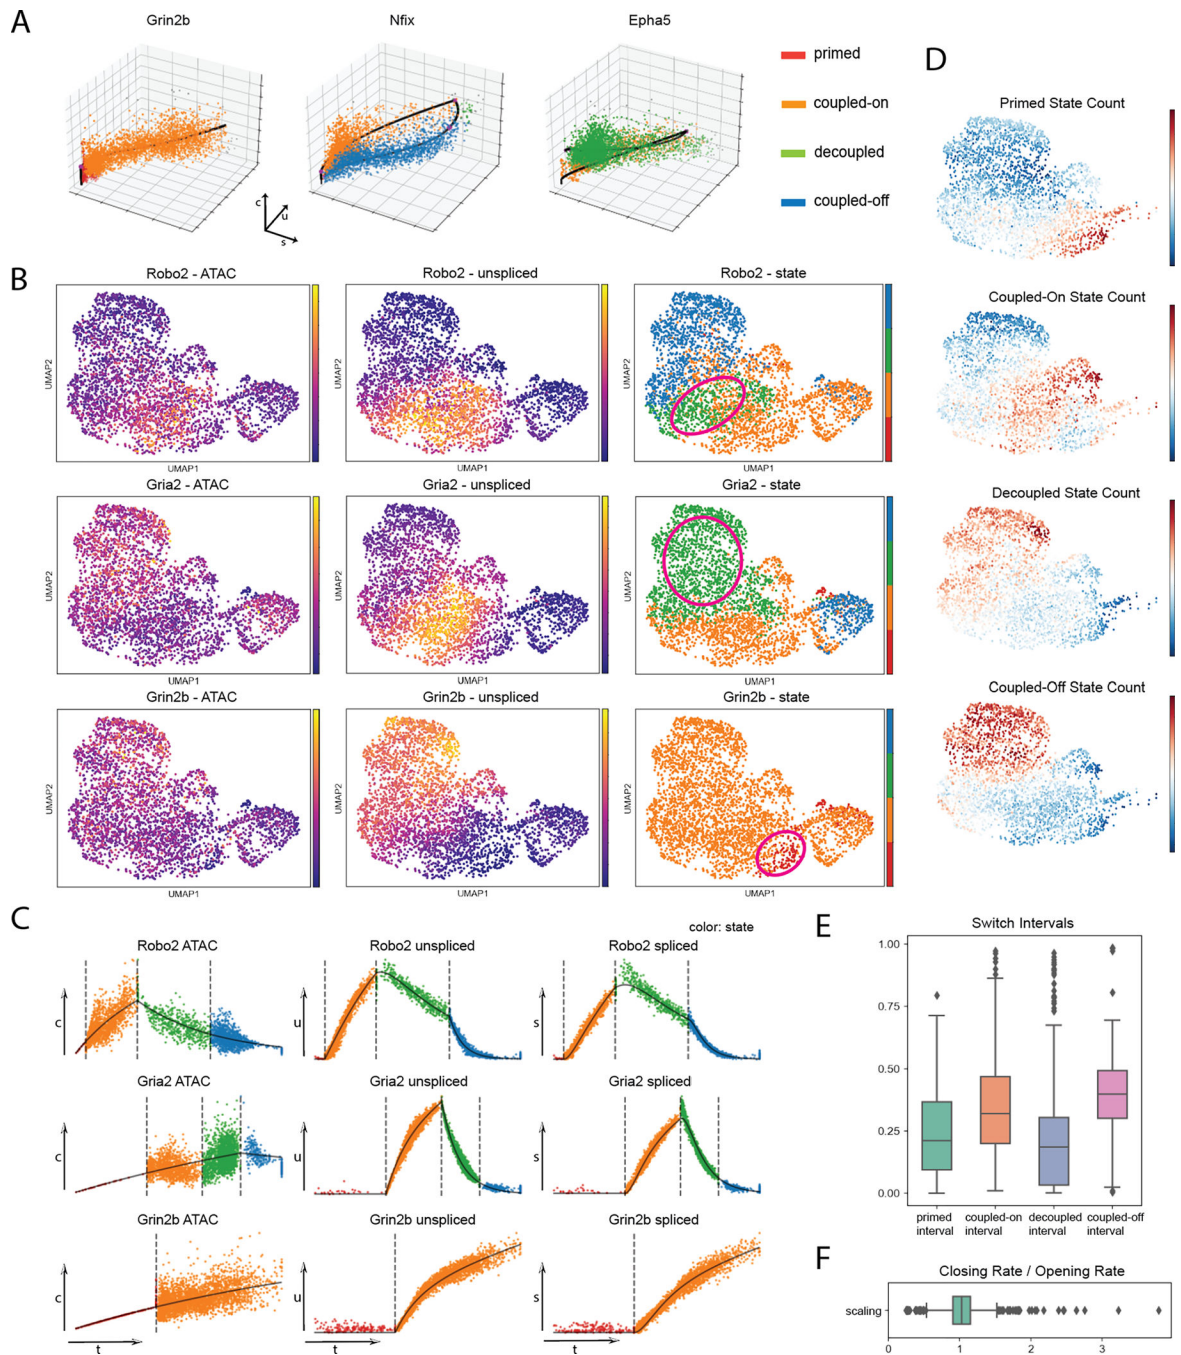


Fig. 3 | MultiVelo captures epigenomic priming and decoupling in embryonic mouse brain.
a, Three-dimensional phase portraits overlaid with MultiVelo fits (solid lines) and inferred states (colors). Each point represents the (c, u, s) values observed for one gene in one cell. **b**, UMAP plots colored by c (left), u (middle) and state assignments (right) for genes predicted by MultiVelo to have notable priming or decoupling intervals. Regions with priming or decoupling are circled. **c**, Observed values for c (left), u (middle) and s (right) plotted as a function of latent time and colored by state assignment. Vertical lines indicate inferred switch times. **d**, UMAP plots colored by the number of genes in each cell assigned to each

of the four states. **e**, Box plots summarizing the lengths of each of the four states across all fit genes (center line, median; box, Q1 and Q3; whiskers, $1.5 \times \text{IQR}$; points, outliers). **f**, Box plot summarizing the ratio between chromatin closing rate α_{cc} and opening rate α_{co} across all fit genes.

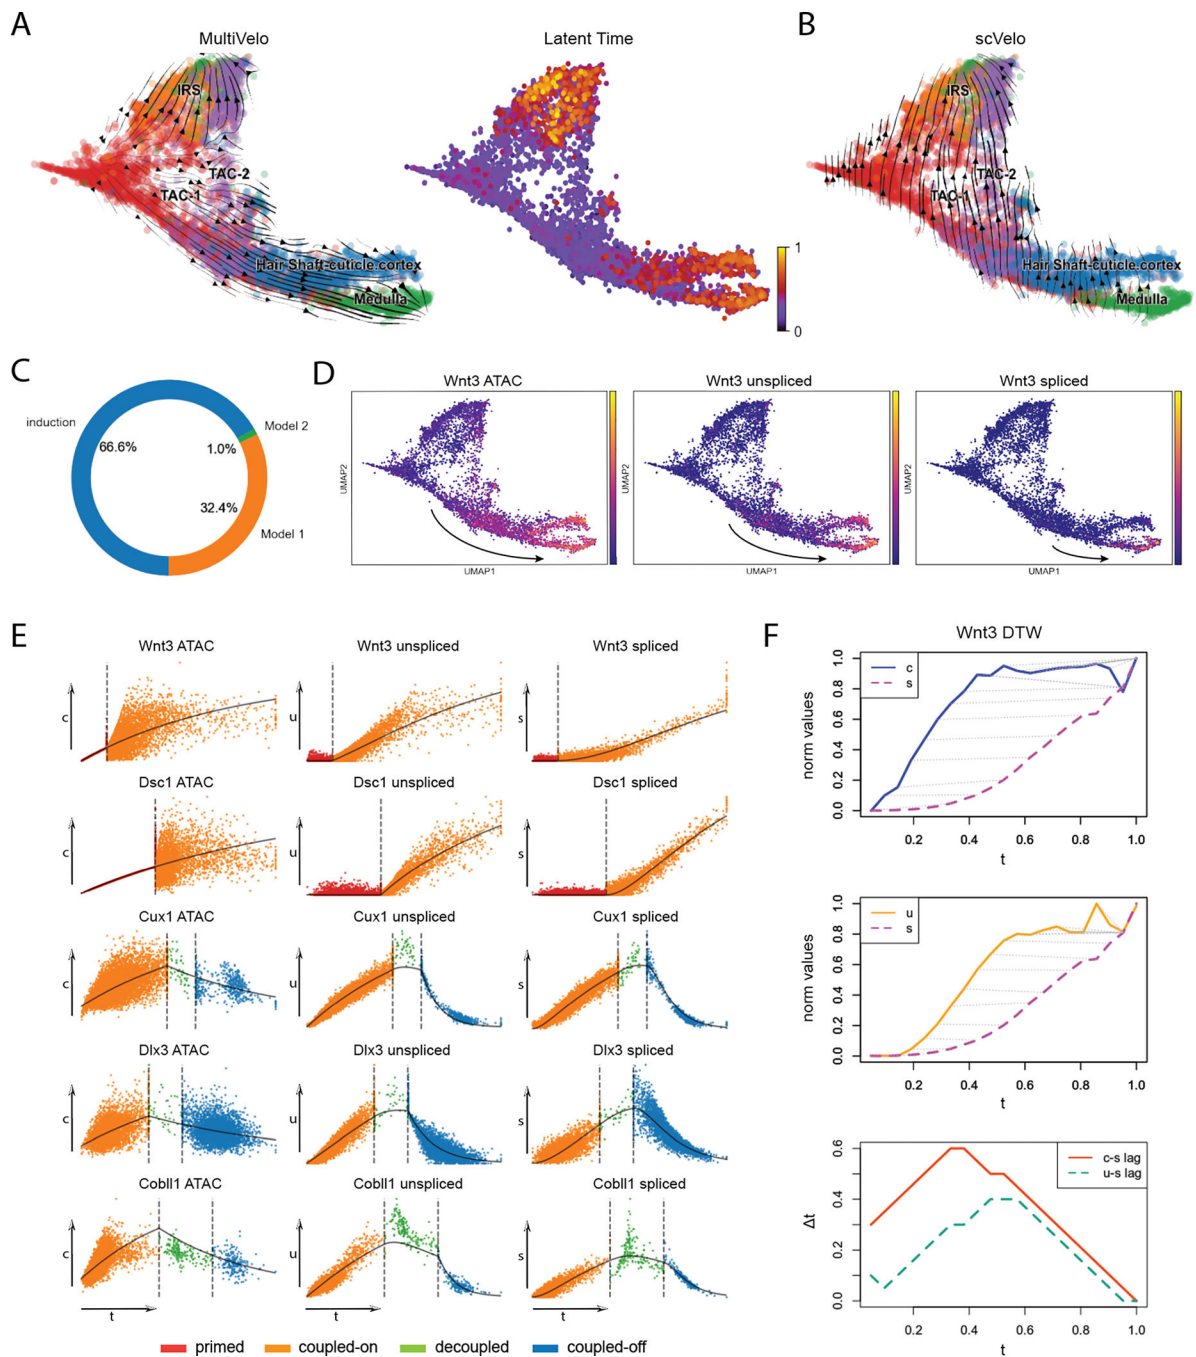


Fig. 4 | MultiVelo quantifies epigenomic priming in mouse skin.

a, UMAP coordinates with stream plot of velocity vectors (left) and latent time (right) from MultiVelo. **b**, Velocity stream plot from RNA-only model (scVelo). **c**, Relative proportion of each type of kinetics across all fit genes ($n = 960$). **d**, UMAP coordinates colored by c (left), u (middle) and s (right) values for *Wnt3*. **e**, Examples of genes showing priming or decoupling. Observed c (left), u (middle) and s (right) values plotted as a function of latent time and colored by state assignment. Vertical lines indicate inferred switch times. **f**, Dynamic time warping alignment of c and s values (top) and u and s values (middle) for

Wnt3. Dotted gray lines indicate corresponding time points after alignment. Bottom panel shows instantaneous time lags computed by subtracting times of aligned time points from the previous two panels. Norm values, normalized values.

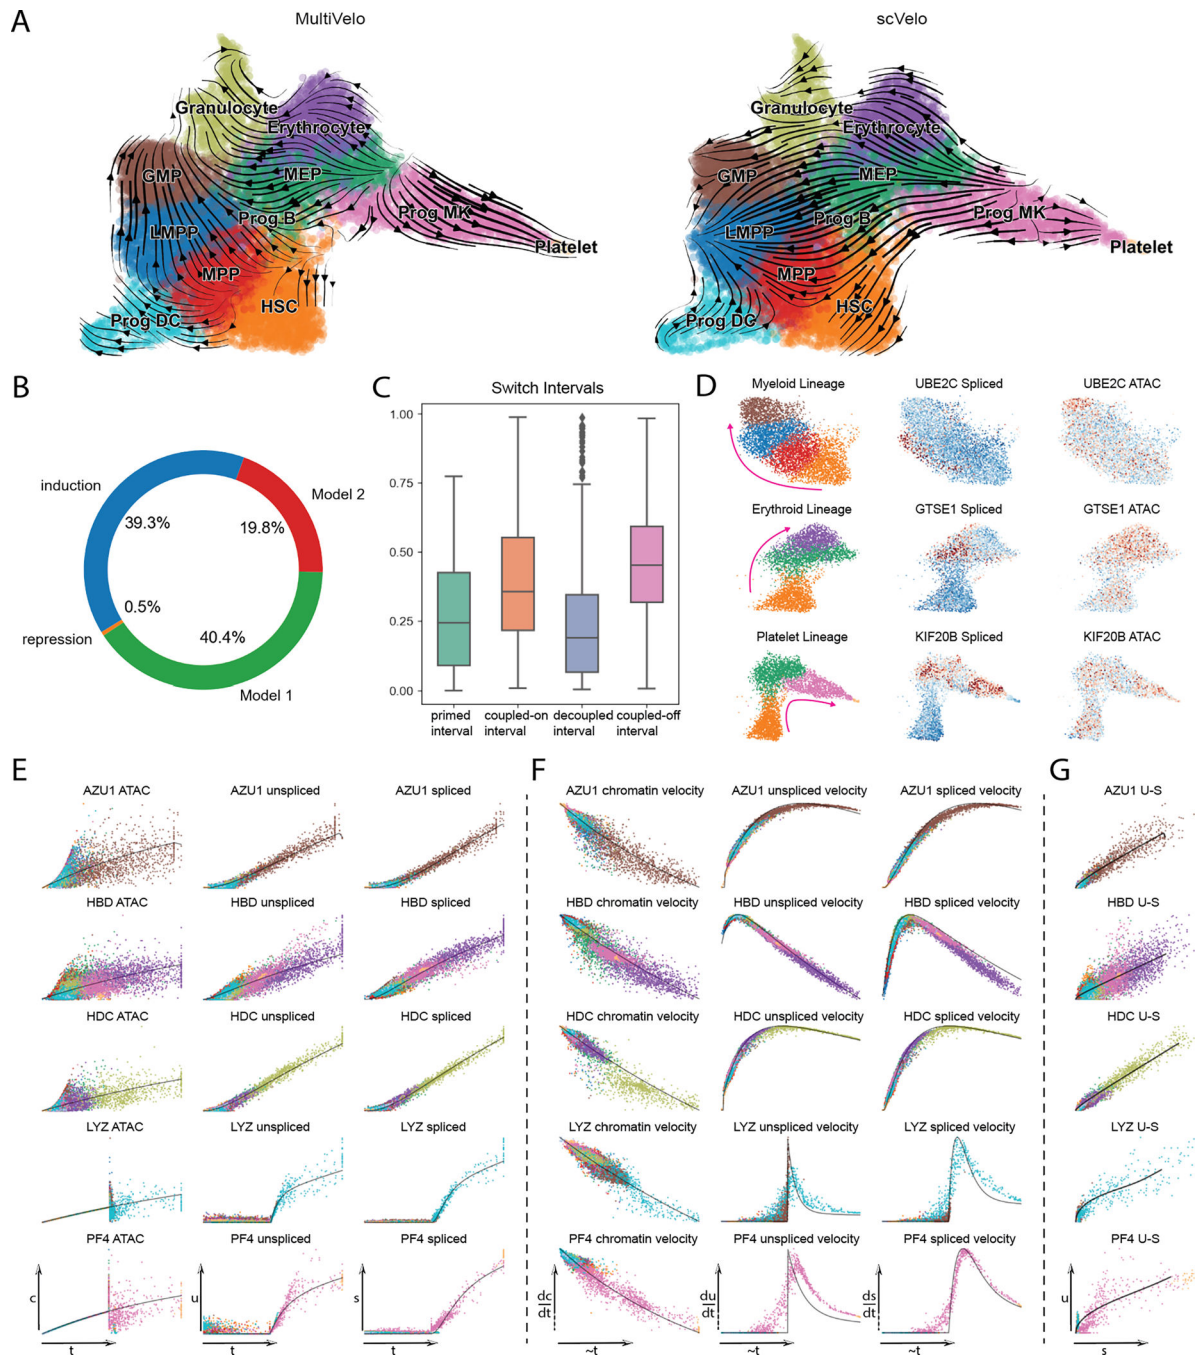


Fig. 5 | MultiVelo identifies priming in HSPCs.

a, UMAP coordinates with stream plot of velocity vectors inferred by MultiVelo (left) and an RNA-only model (scVelo; right). Cell types were annotated based on marker gene expression (Extended Data Fig. 2a). Prog., progenitors; MK, megakaryocyte. **b**, Relative proportion of each type of kinetics across all fit genes ($n = 936$). **c**, Box plots summarizing the lengths of each of the four states across all fit genes (see Fig. 3e). **d**, Several G2/M cell cycle phase markers show model 2 expression pattern toward different lineages. **e**, Examples of genes showing priming or decoupling. Observed c , u and s values plotted as a function

of latent time and colored by cell type. **f**, Corresponding velocity vectors of the same genes as in panel e. Cell velocities and times have been smoothed by RNA neighbors. Note that all velocity values are non-negative, and the lowest velocities are not necessarily at 0. dc/dt , chromatin velocity; du/dt , unspliced velocity; ds/dt , spliced velocity. **g**, RNA phase portraits of the same genes as in panels e and f.

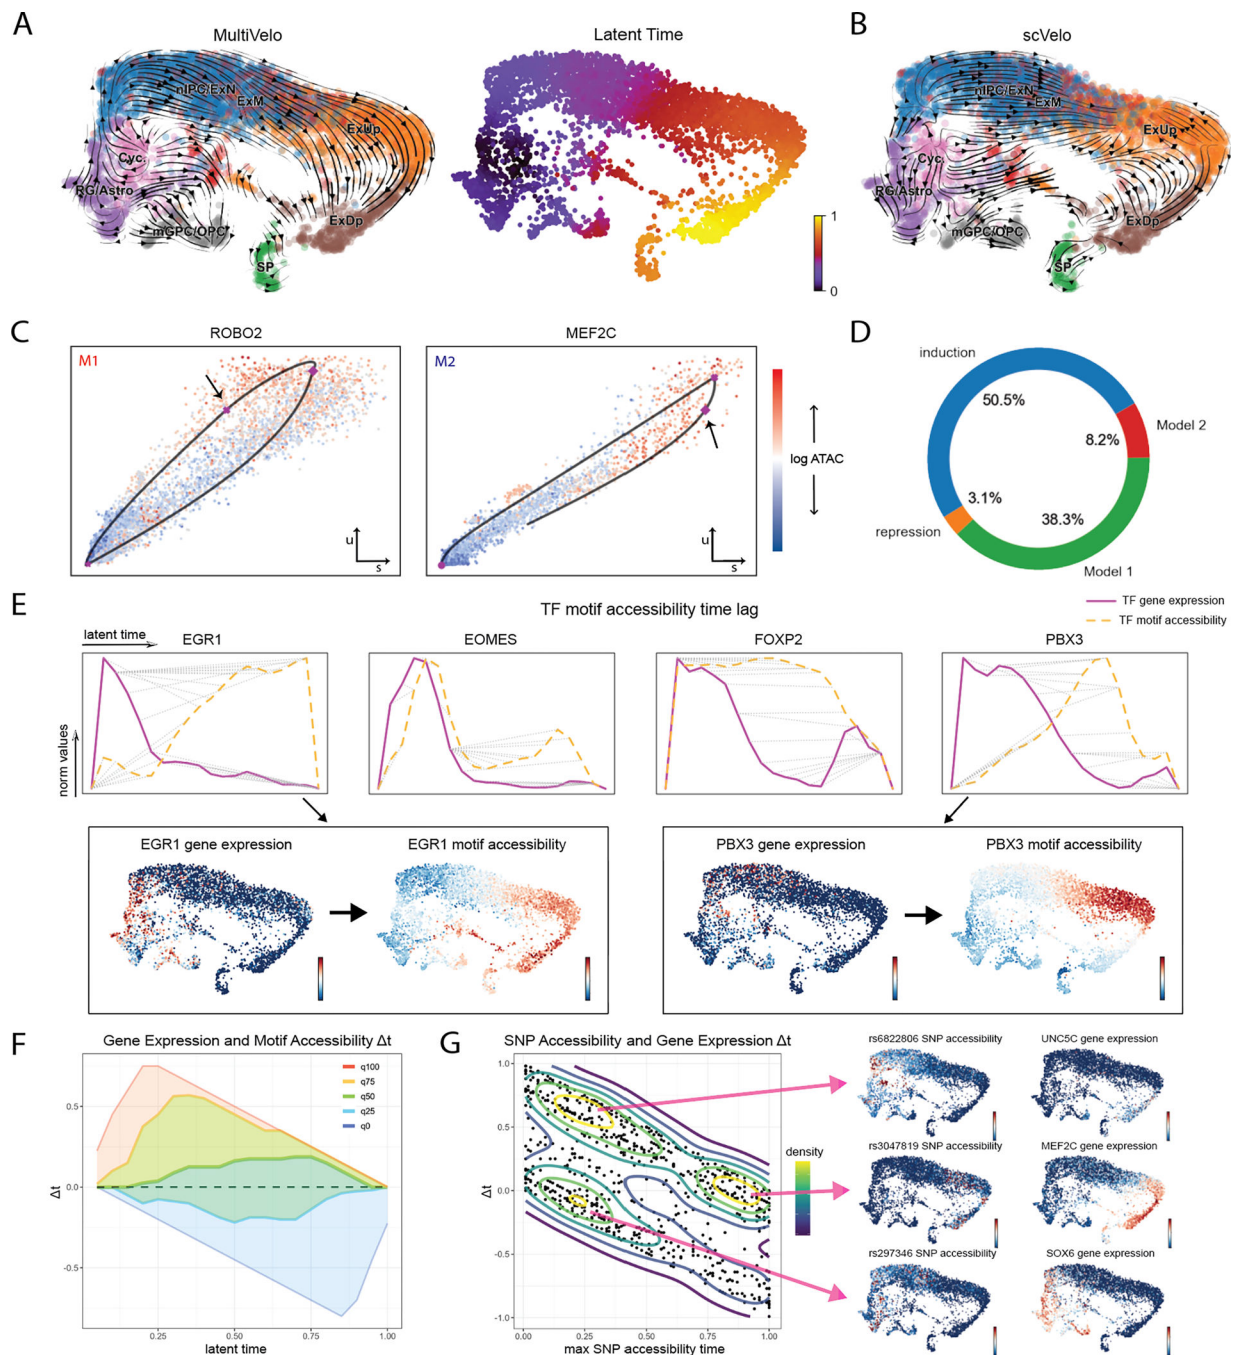


Fig. 6 | MultiVelo infers epigenome and transcriptome dynamics in fetal human brain.
a, UMAP coordinates with stream plot of velocity vectors (left) and latent time (right) from MultiVelo. nIPC/ExN, intermediate progenitor cells/newborn excitatory neurons; ExUp, upper-layer neurons; SP, subplate; ExM, maturing neurons; RG/Astro, radial glia/astrocytes; ExDp, deep-layer neurons; Cyc., cycling progenitors; mGPC/OPC, multipotent glial progenitor cells/oligodendrocyte progenitor cells. **b**, Velocity streamplot from RNA-only model (scVelo). **c**, RNA phase portraits (u versus s) colored by c values show clear differences between model 1 (*ROBO2*) and model 2 (*MEF2C*) genes. Arrows indicate where

chromatin closing begins. **d**, Relative proportion of each type of kinetics across all fit genes ($n = 747$). **e**, Dynamic time warping alignment of TF gene expression and the accessibility of predicted binding sites for four TFs. Dotted gray lines indicate corresponding time points after alignment. Inset UMAPs colored by TF expression and motif accessibility are shown for two of the TFs, *EGR1* and *PBX3*. **f**, Quantiles (q) of TF motif time lags inferred by DTW across all expressed TFs. The median time lag across TFs is positive at most times, indicating that TF expression generally precedes motif accessibility. **g**, Classification of SNPs according to the relationship between maximum accessibility time and time of maximum linked gene expression. The contour lines indicate density, and three main groups of SNPs are visible. Inset UMAP plots are shown for one example SNP from each group.