Learning transcriptional and regulatory dynamics driving cancer cell plasticity using neural ODE-based optimal transport

Alexander Tong[1,2,3*], Manik Kuchroo[4,5*], Shabarni Gupta[6], Aarthi Venkat[7], Beatriz P. San Juan[6], Laura Rangel[6], Brandon Zhu[8], John G. Lock[9], Christine L. Chaffer[6,10§], Smita Krishnaswamy[1,11§],

[1]Department of Computer Science, Yale University, New Haven, CT, USA, [2]Department of Computer Science and Operations Research, University of Montreal, Montreal, QC, Canada, [3]Mila – Quebec AI Institute, Montreal, QC, Canada, [4]Department of Neuroscience, Yale University, New Haven, CT, USA, [4]Department of Medicine, Massachusetts General Hospital, Boston, MA, USA, [6]Kinghorn Cancer Centre, Garvan Institute of Medical Research, Darlinghurst 2010, Australia, [7]Computational Biology and Bioinformatics Program, Yale University, New Haven, CT, USA., [8]Yale College, New Haven, CT, USA, [9]School of Biomedical Sciences, University of New South Wales, Kensington, New South Wales, Australia, [10]St Vincent's Clinical School, University of New South Wales Medicine, University of New South Wales, Darlinghurst 2010, Australia, [11]Department of Genetics, Yale University, New Haven, CT, USA

[*] These authors contributed equally. [§] These authors jointly supervised the work.

Abstract

While single-cell technologies have allowed scientists to characterize cell states that emerge during cancer progression through temporal sampling, connecting these samples over time and inferring gene-gene relationships that promote cancer plasticity remains a challenge. To address these challenges, we developed TrajectoryNet, a neural ordinary differential equation network that learns continuous dynamics via interpolation of population flows between sampled timepoints. By running causality analysis on the output of TrajectoryNet, we compute rich and complex gene-gene networks that drive pathogenic trajectories forward. Applying this pipeline to scRNAseq data generated from *in vitro* models of breast cancer, we identify and validate a refined CD44^{hi}EPCAM^{+}CAV1^{+} marker profile that improves the identification and isolation of cancer stem cells (CSCs) from bulk cell populations. Studying the cell plasticity trajectories emerging from this population, we identify comprehensive temporal regulatory networks that drive cell fate decisions between an epithelial-to-mesenchymal (EMT) trajectory, and a mesenchymal-to-epithelial (MET) trajectory. Through these studies, we identify and validate estrogen related receptor alpha as a critical mediator of CSC plasticity. We further apply TrajectoryNet to an *in vivo* xenograft model and demonstrate it's ability to elucidate trajectories governing primary tumor metastasis to the lung, identifying a dominant EMT trajectory that includes elements of our newly-defined temporal EMT regulatory network. Demonstrated here in cancer, the TrajectoryNet pipeline is a transformative approach to uncovering temporal molecular programs operating in dynamic cell systems from static single-cell data.

Introduction

Cell state plasticity provides a mechanism for cancer cells to rapidly and dynamically evolve in a manner that facilitates primary tumor growth, metastasis and the development of therapy-resistant disease. While the advent of single-cell technologies have allowed detailed characterization of static cell states within tumors, further technological development is required to elucidate the mechanisms governing dynamic cell state transitions that may not only span days, months or years, but also create a variety of cell states shaping the tumor heterogeneity that drives disease progression. Furthermore, the transcriptional networks used by individual cells to undergo dynamic functional changes have been difficult to dissect due to the high dimensional nature of the data, as well as the computational challenge of resolving cellular trajectories over extended periods of time from static snapshot single-cell data. If these issues were addressed, it would be possible to use longitudinal patient samples to gain an unprecedented insight into the mechanisms governing metastasis and therapy-resistant disease, both of which have eluded scientists for decades.

Accordingly, we developed a cell dynamics pipeline centered around TrajectoryNet [1], an algorithm for interpolating continuous dynamics from static snapshot data. Here, we not only improve the ability of TrajectoryNet to learn cellular trajectories with enhanced capabilities of modeling cellular proliferation/death with an auxiliary proliferation network, but also add tools to identify the gene networks underlying these trajectories. TrajectoryNet is built on the new deep learning paradigm of neural ordinary differential equations (ODEs). In this paradigm, the neural network learns a time-varying derivative parameterized by the network weights and biases, and uses an ODE solver to integrate the derivative to compute the function at various points in time. Unlike the discrete dynamic models that are created by recurrent neural networks, neural ODEs can be trained to learn and interpolate dynamics continuously in time when given time course training data. Thus, these models combine the natural universality and trainability of neural networks with the capability of the ODE to model dynamics. When analyzing single cell data, however, we do not actually have cellular dynamics on any individual cell, since they are

destroyed by the scRNAseq experiment; thus, we train TrajectoryNet to match populations over time using continuous normalizing flows. As continuous normalizing flows are not properly constrained to generate biologically plausible dynamics, we add regularizations that bias the network towards more energy-efficient and plausible paths. While this regularization helps TrajectoryNet match population distributions across time, in practice it actually gives continuously normalizing flows of each individual trajectory, thus recovering individual cellular trajectories, spanning long ranges in time and gaps in the cellular manifolds between samples. Furthermore, since cells divide and die, we equip TrajectoryNet with an auxiliary network that learns proliferation rates of cells in order to reflect realistic cellular dynamics. As TrajectoryNet trains a model that learns cellular evolution over time, we can project single cells into future time points, effectively imputing continuous gene expression profiles. Most critically, we combine these continuous single cell gene expression dynamics with causality analysis and public gene regulatory databases to create trajectory-specific transcriptional networks that reveal the key transcription factors and genes that drive a trajectory forward. Through rigorous comparisons, we show that TrajectoryNet not only learns cellular trajectories better than other techniques, but also more accurately infers known gene regulatory relationships than established methods that do not incorporate cellular dynamics.

We showcase our method on time-lapsed single cell data generated in *in vitro* and *in vivo* models to explore cancer cell state plasticity. First, we apply TrajectoryNet to a 3D *in vitro* tumorsphere model system [2] to resolve temporal transcriptional dynamics that govern cancer stem cell (CSC) fate decisions. Using TrajectoryNet's computed proliferation rate, we identify a new subset of CD44^{hi}EPCAM^{+}CAV1^{+} CSCs enriched for tumor-initiating potential, demonstrating the power of TrajectoryNet to accurately identify specialized cells within a heterogeneous population. Tracking these cells over time, we define an epithelial-to-mesenchymal trajectory (EMT) that drives CSCs towards a more mesenchymal cell state, and a mesenchymal-to-epithelial trajectory (MET) that drives CSCs into an epithelial non-CSC state. Using our causality framework, we build the first comprehensive temporally regulated transcriptional networks underlying each of these trajectories.

In cancer, the MET regulatory network is poorly understood. Hence, we showcase the ability of the TrajectoryNet pipeline to reveal novel gene regulatory networks by identifying estrogen receptor related alpha (ESRRA), an orphan nuclear receptor, as a key checkpoint in CSC fate determination and driver of the MET in cancer. We show that modulating ESRRA alters the transcriptional circuitry defining the epithelial and mesenchymal cell states, and highlight its dynamic regulation throughout the *in vitro* tumorsphere-forming assay. We further demonstrate that the temporally regulated gene regulatory networks identified by the TrajectoryNet pipeline accurately predict gene expression dynamics in scRNAseq data capturing primary tumor metastasis to the lung over a 16 week period *in vivo*. Here, TrajectoryNet reveals a dominant EMT trajectory in the *in vivo* metastasis data, with dynamic regulation of the newly-defined core EMT regulatory network. Together, this manuscript demonstrates that TrajectoryNet is an effective tool to model dynamic cell state transitions from snapshot single cell data and can identify genes critical in cell state plasticity.

The remainder of this manuscript is organized as follows: 1) First we present our cell dynamics pipeline featuring the TrajectoryNet Neural ODE network for inferring cellular—and associated gene—dynamics from single-cell data, and Granger causality analysis for building networks from these dynamics. 2) Through rigorous comparisons, we demonstrate that TrajectoryNet is not only able to infer trajectories significantly better than other methods [3–5], but can also more accurately infer gene regulatory relationships than methods that do not account for cellular dynamics [6]. 3) Next, we introduce our unique dataset consisting of a 5-timepoint *in vitro* tumorsphere assay that begins with CD44^{hi} cells containing a CSC population, measured with scRNA-seq, and tracks their evolution over 30 days. 4) Next, we apply our cell dynamics pipeline to this dataset to derive temporal trajectories. From these trajectories we identify 5 clusters of co-expressed genes based on their temporal dynamics. 5) The TrajectoryNet-derived cellular trajectories allow us to derive a refined CSC marker set which includes

EPCAM and CAV1 in addition to previously known CD44. 6) Next, we compare two types of cellular trajectories that emerge from these CSCs, one set of trajectories that we call MET trajectories as they have a more epithelial fate, while the other set, we call EMT trajectories as they have a more mesenchymal fate. 7) We then apply our causality tools to gene dynamics from these trajectories to derive the core gene regulatory circuitry of both the EMT and MET. 8) From these trajectories we identify ESRRA as a major driver of the MET transition, which could be responsible for creating the cell state heterogeneity required for secondary tumor formation. We conduct experiments using siRNA and ESRRA antagonist treatment of 2D cultured CD44^{hi} cells, showing that they indeed lead to an increase in the epithalial marker CDH1 expression. 9) Finally, we test the relevance of our findings in an *in vivo* system using a xenograft model where we measure primary tumors and matched lung metastases with single cell and spatial RNA-seq technologies. This experiment replicates our refined stem cell signature as well as elements of the EMT regulatory circuit.

Results

Overview of the Cell Dynamics Pipeline

Timelapsed single-cell data (Figure 1A) poses significant challenges from a computational perspective. In cancer for example, cells undergo significant and rapid transcriptional changes in response to natural tumor evolution, the stressors of the metastatic cascade or therapeutic insults, where each individual timepoint represents a distribution of cells that largely do not overlap in cellular state with previous or subsequent timepoints. Under these circumstances, previously presented techniques for inferring dynamics such as RNA velocity [4] and pseudotime [5] fail as these approaches rely on well-sampled manifolds without large gaps in cell state space or time. Most critically, however, these approaches do not help users understand the gene regulatory dynamics necessary to drive these trajectories forward. Extracting longitudinal dynamics and understanding their underlying transcriptional drivers from disconnected static snapshot measurements can be challenging as there are few methods of interpolation between irregularly spaced or distant distributions and learn gene-gene regulatory dynamics. To address this knowledge-gap, we developed TrajectoryNet (Figure 1B), a tool for studying cellular differentiation dynamics from time-lapsed single cell transcriptomic data. Notably, TrajectoryNet produces trajectories for each individual cell instead of a single averaged trajectory, as is the case with pseudotime, further allowing user to build sophisticated transcriptional networks by leveraging causality metrics and known gene-gene relationships. Here, we present a pipeline that allows users to: i) learn single cell trajectories from time lapsed single cell data (Figure 1Ci), ii) interpolate continuous gene trends for each trajectory (Figure 1Cii) and iii) learn complex transcriptional networks from trajectories of interest (Figure 1Ciii).

TrajectoryNet learns the dynamics of cell state from cellular snapshot data (Figure 1A). Here we highlight its key features and adaptations for use in this context, and elaborate in the Methods section. The heart of the model is a single neural network f with parameters θ which takes as input a current cell state $z(t)$ and the current time t and outputs the instantaneous change in cell state with respect to time $\frac{dz(t)}{dt}$. This network f tells us how each cell evolves instantaneously and continuously. To track its progress over longer time scales, we can use standard ordinary differential equation (ODE) solvers to integrate its position over time. Thus, the neural ODE system is an ensemble system consisting of a neural network computing a derivative and an ODE solver that integrates the derivative to compute a per cell time-trajectory. While the original neural ODE paper made the key contribution of showing how to train such a system, by using a method called adjoint sensitivity, here we utilize this framework to learn

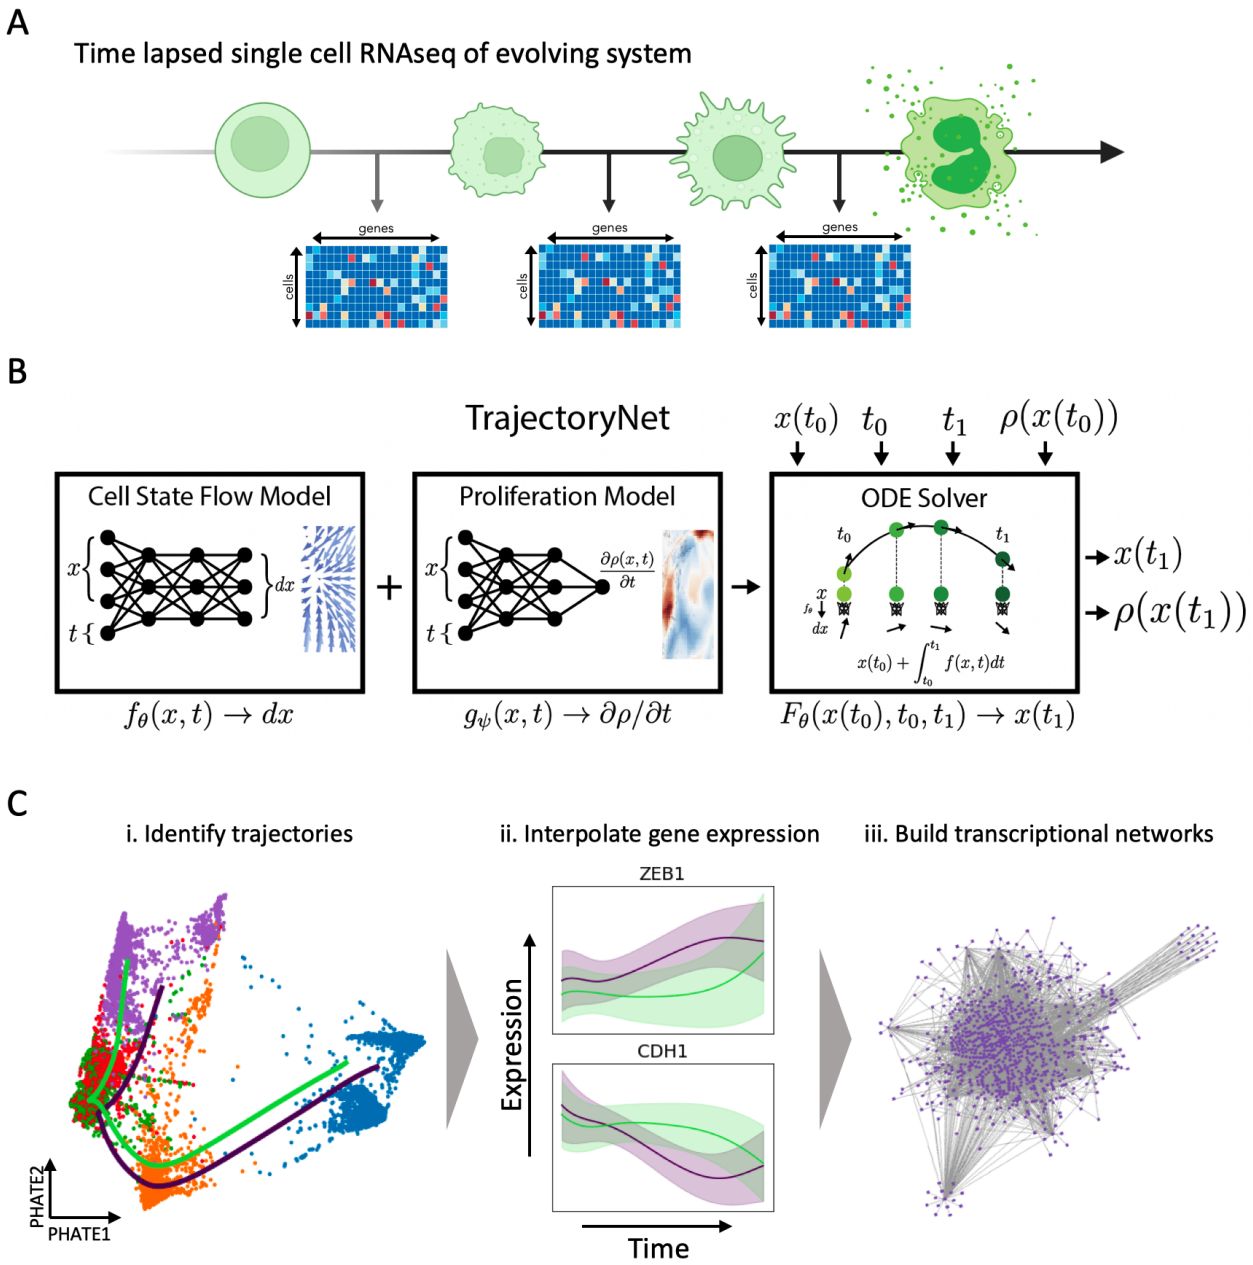

Figure 1: Overview of TrajectoryNet Algorithm. **(A)** Data for TrajectoryNet is generated by capturing timelapsed single-cell data (RNA-seq, ATAC-seq, CITE-seq, etc.) on an evolving system. **(B)** Illustration of the TrajectoryNet model. TrajectoryNet uses an adaptive ODE solver to integrate cell state over time in order to calculate trajectories and relative proliferation rates. Time series single cell data produces disconnected distributions over a developmental time. TrajectoryNet interpolates disconnected distributions to continuously infer transcriptional dynamics as well as cell growth and death via unbalanced dynamic optimal transport. **(C)** Overview of TrajectoryNet analysis framework: i) identify continuous trajectories from disconnected time-lapsed data; ii) interpolate continuous gene expression dynamics based on terminal cell population; iii) Build transcriptonal networks using causality analysis and public gene regulatory network databases.

cellular population dynamics.

$$\hat{z}(t_1) = F_\theta(z(t_0), t_0, t_1) = z(t_0) + \int_{t_0}^{t_1} f_\theta(z(t), t)dt \tag{1}$$

If we had longitudinal data—data that tracked a particular cell over time (Figure 1A)—then we could

directly apply a matching loss L between our estimated cell at time t_1, $\hat{z}(t_1)$ and the true cell position $z(t_1)$ and apply gradient descent on $\frac{\partial L(\hat{z}(t_1), z(t_1))}{\partial \theta}$. However, TrajectoryNet only assumes access to snapshot data, and so does not measure the cell $z(t_0)$ at t_1. Therefore we need a distribution level loss, for which we turn to the framework of *normalizing flows*.

In a normalizing flow, a random variable z_0 is transformed via a bijective function $z_1 = f(z)$ then the change in probability of z_1, $p(z_1)$ is computed through a change of variables formula

$$\log(p(z_1)) = \log(p(z_0)) + \log\left(\left|\det \frac{df}{dz_0}\right|\right) \tag{2}$$

In the original formulations of normalizing flows, discrete neural network layers would transform probability distributions from z_0 to z_1, to z_2 and so forth, gradually transforming one distribution of cells to another via reversible operations. Neural ODEs allow a continuous version of this kind of normalizing flow using an *instantaneous change of variables formula.* If $z(t)$ is a continuous random variable that is transformed via differential equation $d(z(t))/dt = f(z(t), t)$, then the change in probability is computed as $\frac{d\log(p(z(t))}{dt} = tr\left(\frac{df}{dt}\right)$. This calculated probability at each t_k, where we have data available is penalized by a KL divergence penalty $D_{\mathrm{KL}}(p(z(t_k)||q(z(t_k)))$ for known q.

While continuous normalizing flows provide a way to transform one probability distribution to another, there is no reason that such flows need to be biologically plausible, as is needed to simulate the cellular differentiation process. Indeed flows can be convoluted ways of altering one distribution via a series of circuitous paths. In order to constrain this to the realm of biological plausibility, we add a penalty that we believe to be true of biophysical systems: energy efficiency. In other words, we believe that when cellular populations transform they do so in an efficient way [3]. This penalty is enforced by penalizing the magnitude of the derivative computed by the neural network at each measured timestep $|\frac{df}{dt}|$ of the cellular differentiation process. Interestingly, this regularization allows the neural network to perform a *dynamic optimal transport*, i.e., an optimal transport between the populations where the paths of transport (rather than just the final displacements) are energy optimal. This neural ODE design of dynamic optimal transport is indeed the key contribution of TrajectoryNet.

We also add a penalty to constrain the flow to the observable cellular manifold. This penalty restricts cellular paths to flow along regions of cellular density, as determined by the distance to the K-nearest neighboring cell. Finally, because cells are not simply transported but also proliferate and die, we altered TrajectoryNet to perform an *unbalanced optimal transport* which is discussed in Methods.

The objective of TrajectoryNet can be decomposed into four parts:

$$L(x(T), \theta, \theta') = \underbrace{\overbrace{-\log(M(x(T), T))}^{\text{Distribution Matching}} + L_{energy} + \overbrace{L_{balance} + L_{proliferation}}^{\text{Proliferation penalties}}}_{\text{Unbalanced Dynamic OT}} + \overbrace{L_{density}}^{\text{Biological path priors}} \tag{3}$$

- The first term is a distribution matching term, which minimizes the divergence between the predicted distribution at time T and the true distribution at time T. This ensures that the model matches the data at each measured timepoint.

- The second term, L_{energy} integrates over the length of the path that the cell took from time 0 to time T to ensure that cellular paths are energy efficient. This is a key penalty in our framework that provably renders TrajectoryNet into a dynamic optimal transport framework. Of all the possible set of paths that link distributions together we chose the one that minimizes the total path energy. Here, we parameterize the energy as the sum of squared path lengths

$$\inf_f \int_{x(0)\sim\mu} \int_{\mathcal{T}} \|f(x(t), t)\|^2 dt dx. \tag{4}$$

Considering just these first two terms our objective is now equivalent to a relaxed optimal transport problem where we minimize Eq. 4 subject to the continuity equation $\partial_t \rho + \nabla \cdot (\rho f) = 0$ enforcing mass is neither created nor destroyed, and the boundary distributions match with $\rho(x, 0) = \mu$ and a relaxed $\rho(x, 1) \approx \nu$ where $\rho(x, t)$ is the density at time t.

- The third and fourth components are responsible for modeling cell proliferation and cell death help train a proliferation network g to transform the balanced OT problem in Eq. 4 into an unbalanced one. Intuitively, $L_{balance}$ makes sure that the total mass of the system is not changing so that $\int_{x \in \mathcal{C}} M(x, t) dx \approx 1$ for all t, and $L_{proliferation}$ penalizes the (squared) cost of mass creation and destruction $\int_{x(0) \sim \mu} \int_{\mathcal{T}} g(x(t), t)^2 dt dx$ enforcing how close to a balanced optimal transport the solution is. These first four terms can be seen as solving a dynamic optimal transport problem

$$\min_{\tilde{\mu}} W(\tilde{\mu}, \nu) + \lambda D_{\mathrm{KL}}(\mu, \tilde{\mu}) \tag{5}$$

which can be restated in the dynamic formulation for a vector flow field f and a scalar growth field g, as minimizing

$$\inf_{f,g} \int_{x(0) \sim \mu} \int_{\mathcal{T}} \left(\|f(x(t), t)\|^2 + g(x(t), t)^2 \right) dt dx \tag{6}$$

subject to the continuity equation $\partial_t \rho + \nabla \cdot (\rho f) - \rho g = 0$ which enforces that mass is neither created or destroyed after taking into account a scalar addition ρg, and subject to $\rho(x, 0) = \mu$ and $\rho(x, 1) \approx \nu$.

- The final term ensures that the flow remains on the cellular manifold. While we generally assume that cells on average take short paths, cells are also thought to transition over a low dimensional manifold, an assumption which is also made in graph-based methods where paths are restricted to transitions between observed cells. In contrast, our flows are continuous in gene space and time, nevertheless we still find it useful to encourage paths that are near data. We use a thresholded k-nearest-neighbors penalty which regularizes the sum of distances to the k nearest neighbors

$$\sum_k \max(0, \text{k-min}\left(\{\|x(t) - z\| : z \in \mathcal{X}\}\right) - h). \tag{7}$$

Using these four terms yields an optimization similar to dynamic unbalanced optimal transport through direct optimization of the drift and proliferation fields. This direct optimization allows us to impose additional biological priors on the cell paths. Additional descriptions of these terms, the TrajectoryNet algorithm and comparisons to other techniques can be found in methods. For rigorous benchmarking studies comparing TrajectoryNet to other techniques on synthetic (Figure S1) and real single cell data (Figure S2), please see methods.

Generating continuous gene expression data from trajectories

Once we have a trained model $f_{\hat{\theta}}$ as a dynamic optimal transport on single-cell RNA-sequencing timecourse data, we can use this model to predict the path of a cell. Given a cell with state x at time t_i $c = (x, t_i)$ we can integrate the state forward and backwards in time according to the ODE defined by $\frac{dx}{dt} = f_{\hat{\theta}}$, to get the corresponding cell state at arbitrary timepoints. Specifically, we can estimate the cell state $\hat{x}_{t_j}$ according to our model $f_{\hat{\theta}}$ at any time t_j:

$$\hat{x}(t_j) = x(t_i) + \int_{t_i}^{t_j} f_{\hat{\theta}}(x(t), t) dt \tag{8}$$

Furthermore, this equation can be numerically approximated using an ODE solver for each cell individually, or parallelized over multiple cells. In practical terms, this allows users to infer the continuous gene expression dynamics for a given cell or cluster of cells based on its state $x(t_i)$ at its measured time t_i.

Learning transcriptional networks with Granger causality

We combine continuous gene expression trends with novel time lapsed causality analysis as well as public databases to learn the transcriptional networks underlying cellular trajectories. While many methods that infer gene-gene networks, like correlation or mutual information methods [7], assume that the underlying cellular states are static, we believe that incorporating cellular dynamics information is critical to inferring the true underlying transcriptional network. As TrajectoryNet outputs complete transcriptional trends for each cell across time, we can infer better gene-gene relationships across time using Granger causality. Granger causality goes beyond correlations in time series by testing whether a variable (or set of variables) X forecasts a variable Y. More specifically, a variable X *Granger-causes* a variable Y if predictions of Y based on its past values and on the past values of X are better than the predictions of Y based on only its own past values. To determine the likely gene regulatory structure, we use Granger causality analysis on the average TrajectoryNet trajectories for groups of cells. We determine that there is a directed edge from a gene X to gene Y if gene X Granger-causes Y. We score this edge based on the Granger p-value signed by the direction of influence of X on Y. This is determined by the sign of the coefficient on X in the linear regression of Y based on the past values of X and Y. If this coefficient is positive, we take this to mean that X has a likely up-regulatory effect on Y. Conversely, if the linear regression coefficient of X is negative, we take this to mean X has a down-regulatory effect on Y (Figure S3A). By taking the log transformed Granger test statistic and adding a sign for direction of effect between a transcription factor and a target gene, we compute an interaction strength score that determines whether a transcription factor X may have an effect on a particular target gene Y. By aggregating signed Granger analysis results across large numbers of transcription factors and genes, we can identify transcriptional programs responsible for defining key trajectories.

We note that Granger causality is not the same as "true causality" as it is only evidence of X preceding Y, and there are many instances where this precedence may be incidental where changing X does not change Y. We, therefore, combine these results with known gene-gene relationships found in the Transcriptional Regulatory Relationships Unraveled by Sentence-based Text mining (TRRUST v2) database [8]. While these databases are also imperfect, the integration of TrajectoryNet, causality analysis and a public gene regulatory network better approximates the underlying transcriptional network. In this instance, we first identify the key driver transcription factors using a total Granger causal score (TGCS):

$$\mathrm{TGCS}(\mathrm{gene}_i) = \sum_j -\log\left[\mathrm{Granger\text{-}pvalue}(\mathrm{gene}_i, \mathrm{gene}_j) + \epsilon\right]$$

with the idea that key drivers will strongly Granger-cause many other genes. We show that TGCS outperforms other methods which assume static cellular states at inferring gene regulatory interactions from time series data (Figure S3B-D). Following TGCS scoring we then use the TRUSST v2 database to generate a seed network (using Cytoscape 3.9.0). We then use this seed network to extract the first-degree nodes of the top drivers of the trajectory (based on TGCS score) and their adjacent edges. We further prune and weigh the edges of this network with Granger scores to create higher fidelity networks (see methods). We do however, caution that Granger scores do not prove causality and that these should be regarded as candidates for further validation. For a more in-depth explanation of our gene network inference comparisons on synthetic data please see Methods.

Time-resolved single cell data for learning molecular drivers of cell fate decisions

We chose to apply the TrajectoryNet pipeline to study CSC plasticity, a major problem in cancer biology today. CSCs are highly specialized cells that drive tumor initiation and metastasis [9]. One of their key biological features is to use cell state plasticity to self-renew, or differentiate into various non-CSC progeny to create tumor heterogeneity. To date, the temporal regulation and biological networks underlying CSC plasticity have not been resolved.

To investigate the cellular and transcriptional dynamics behind CSC plasticity and potentially identify new strategies to control it, we suspended the CD44^{hi} CSCs in a 3-dimensional tumorsphere assay. This population of heterogeneous CD44^{hi} cells are enriched for CSCs that reside in a hybrid epithelial-mesenchymal state. In this assay, CD44^{hi} CSCs traverse at least two distinct trajectories: one where they self-renew and/or move towards a more mesenchymal CSC state via an EMT, and one where they differentiate into epithelial CD44^{lo} non-CSC progeny via an MET [2]. Note that both the trajectories we term as EMT and MET start from the CSC state and go towards a *more epithelial* or *more mesenchymal* state (Figure 2A). To measure the range of cellular states created, we collected tumorspheres at days 2, 12, 18, and 30 and performed single cell RNA sequencing (scRNA-seq) (Figure 2B). We also measured scRNA-seq on the unsorted day 0 population. This dataset, comprising of 16,983 genes measured in 17,983 cells across 5 timepoints, creates a unique opportunity to study CSC transitions and better understand when and how cell fate decisions are made.

TrajectoryNet learns transcriptional dynamics driving cancer cell plasticity

We applied TrajectoryNet to the time-lapsed single cell tumorsphere measurements. By projecting TrajectoryNet inferred trajectories on top of a PHATE visualization [10] of the time-lapsed single cell data, we observe a diverse set of trajectories starting at day 0 and ending at day 30 (Figure 2C). Visualizing TrajectoryNet's inferred growth or proliferation rate, we identify subsets of cells that are disproportionately more likely to contribute to the following time point (Figure 2D). This inferred proliferation rate identifies rapidly dividing cells that contribute to the emergence of distinct phenotypic cell states during tumorsphere development. Next, we compute an EMT signature score based on the average expression of genes known to play a role in EMT [11] and visualize this score on our PHATE embedding. Using this analysis, we can clearly identify a heterogeneous range of epithelial non-CSCs (red) and hybrid epithelial-mesenchymal CSCs (blue) across timepoints (Figure 2E). When we take this score into account and analyze the day 0 timepoint, which is the single cell data corresponding to the unsorted 2D cultured HCC38 cells, we observe an epithelial cell fraction representing CD44^{lo} cells (red dots) and the hybrid epithelial-mesenchymal CD44^{hi} fraction (blue dots). Critically, the trajectories defining cells that progress through to day 30 originate in cells within the CD44^{hi} fraction of the day 0 population, thus demonstrating the ability of TrajectoryNet to correctly identify tumorsphere cells-of-origin.

As TrajectoryNet learns the continuous gene expression dynamics, we can interpolate the complete expression dynamics of individual genes across the tumorsphere time-course. Visualizing these genes shows clear co-expression dynamics, allowing us to cluster the interpolated expression-vs-time trends of all trajectories into 5 co-expression clusters using kmeans (Figure 2F-G). These gene co-expression clusters are ordered temporally, with Cluster 1 describing genes that have highest expression early in tumorsphere formation during day 0, and, Cluster 4 describing genes that have highest expression late in tumorsphere formation during day 30. Finally, Cluster 5 describes a set of genes that have high expression both early and late in the transition (Figure 2F-G). A well-defined core EMT circuitry is critical to CSC plasticity [12], yet little is known about the temporal regulation of these transcription factors in dynamic systems. To demonstrate that TrajectoryNet can resolve gene-specific dynamics, we show that the core EMT transcription factors are not simply co-expressed and co-regulated at the same time,

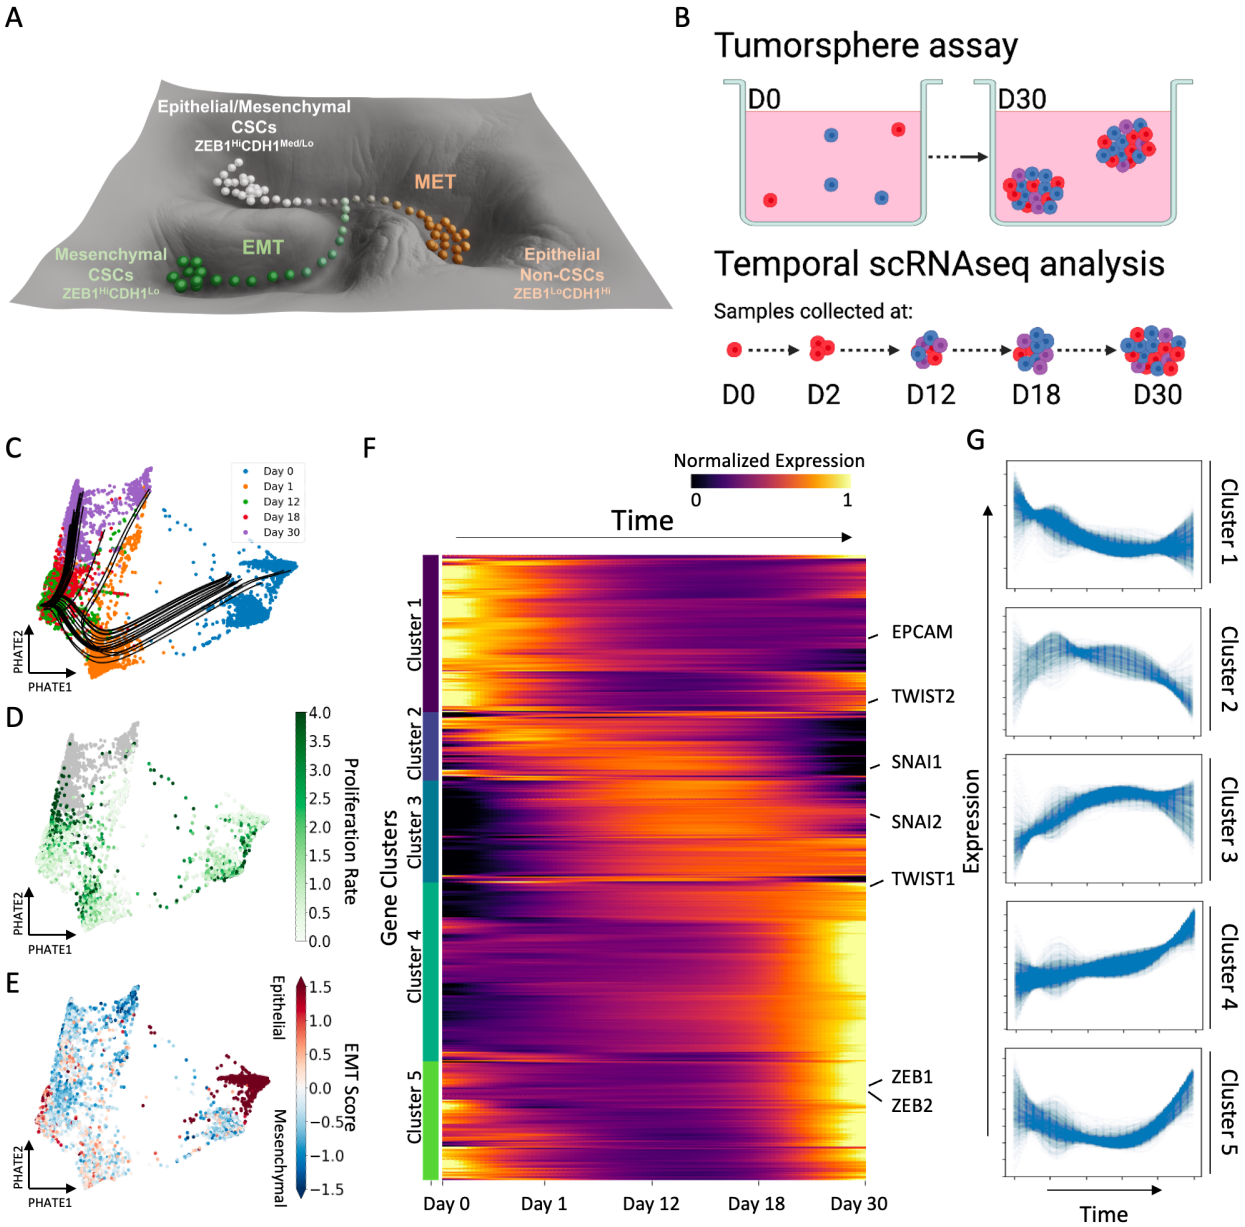

Figure 2: Overview of Tumorsphere dataset. **(A)** Schematic illustrating dynamic transitions via the EMT and MET between highly tumorigenic and metastatic CD44^{hi}ZEB1^{hi}CDH1^{lo} CSCs and poorly tumorigenic CD44^{lo}ZEB1^{lo}CDH1^{hi} epithelial cells. **(B)** Illustration of the tumorsphere protocol and single cell RNAseq experiment. CD44^{hi} CSCs are seeded in single cell suspension at day 1. By day 30, ten percent of single CD44^{hi} cells seeded produce three-dimensional heterogenenous tumorspheres. **(C)** PHATE [10] embedding of time-lapsed scRNAseq data generated by the tumorsphere assay described in Figure 2B. Samples are colored by timepoint of data acquisition. Trend lines are created by TrajectoryNet. **(D)** Visualization of TrajectoryNet inferred proliferation rate. **(E)** Visualization of EMT gene expression score [11] on PHATE embedding. **(F)** Visualization of inferred continuous gene expression dynamics over time. Genes have been clustered into 5 groups based on their gene expression dynamics (see Methods). Key EMT genes (*EPCAM*, *TWIST1/2*, *SNAI1/2* and *ZEB1/2*) have been indicated on heatmap. **(G)** Visualization of inferred continuous gene expression dynamics in each gene cluster shown in Figure 2C. Each line represents a single gene trend over time.

rather, they exhibit distinct temporal expression patterns where *SNAI1* peaks in Cluster 1, *SNAI2* and *TWIST1* peak in Cluster 3, while *ZEB1* and *ZEB2* in Cluster 5 have peak expression early and late across the tumorsphere time course. These data demonstrate that TrajectoryNet can delineate the

complex temporal regulatory interactions that initiate and maintain cancer plasticity dynamics.

A refined CSC marker profile to improve identification of the tumorsphere cell-of-origin

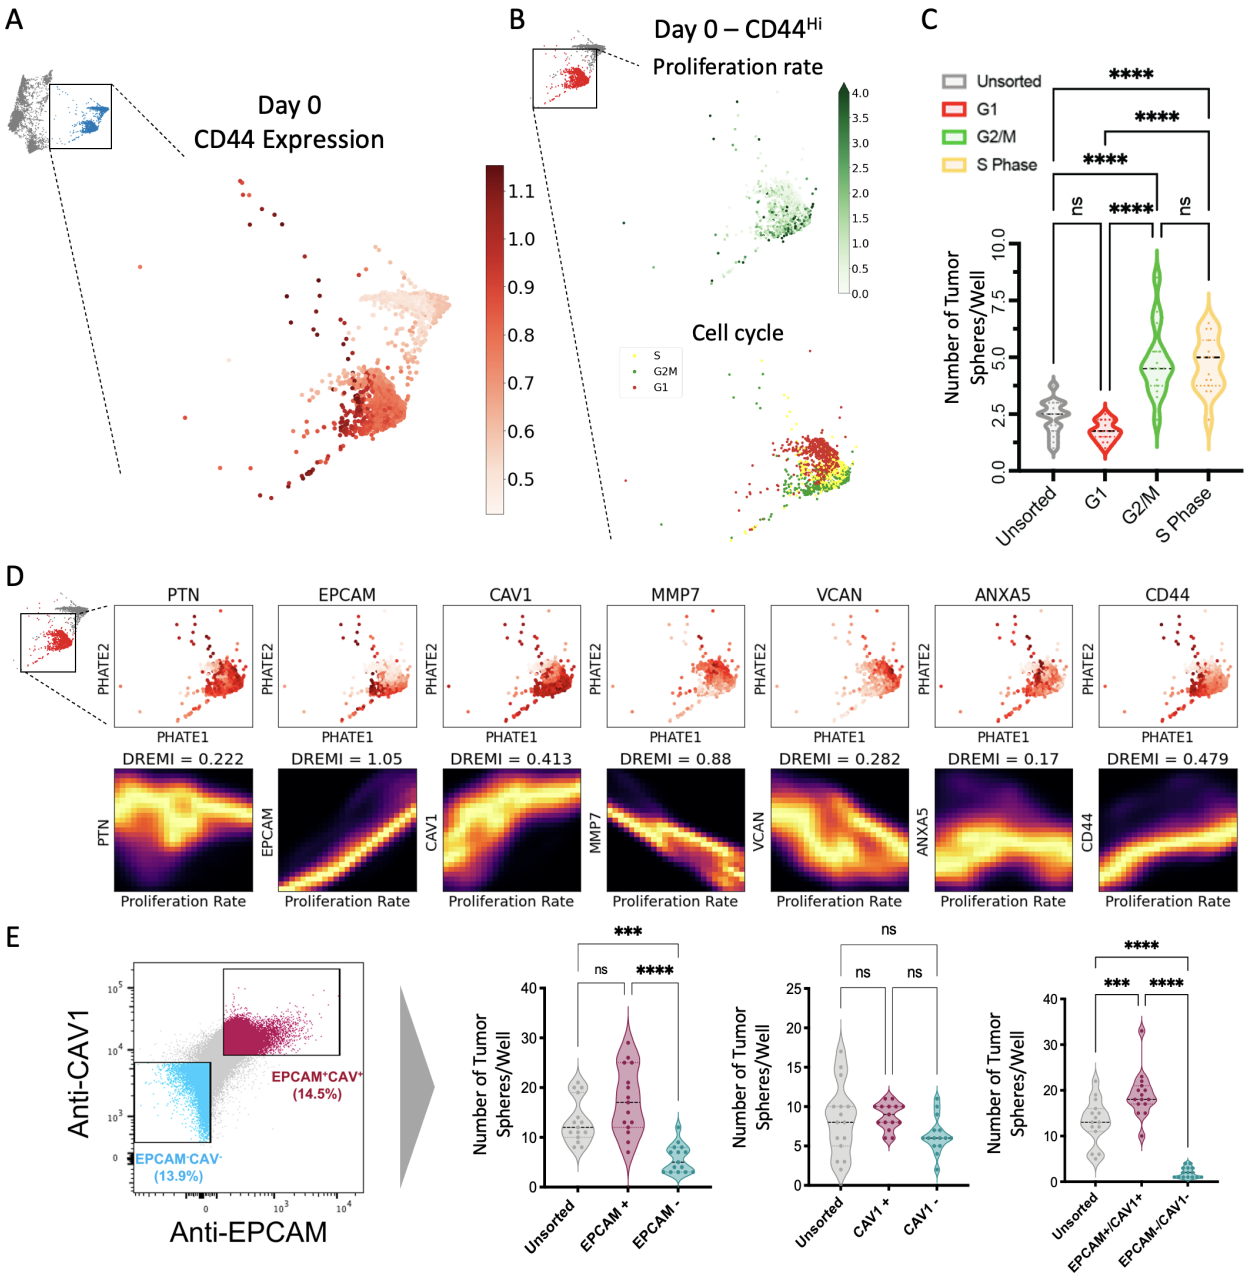

Figure 3: Refining identification of the tumorsphere cell-of-origin **(A)** PHATE visualization of day 0 scRNAseq timepoint with CD44 expression highlighted. **(B)** Zoom in on day 0 CD44^{hi} population with TrajectoryNet proliferation rate (top) and inferred cell cycle state (bottom) [13] visualized. **(C)** Flow cytometry based sorting of HCC38 CD44^{hi} cells infected with the FUCCI cell cycle sensor system to isolate cells at different stages of the cell cycle. Each cell cycle isolate is measured for *in vitro* tumorsphere-initiating potential. **(D)** Visualization in CD44^{hi} population of key differentially expressed cell surface markers and their DREMI [6] association scores with TrajectoryNet-inferred proliferation rate. **(E)** Flow cytometry isolated EPCAM$^{+/-}$ and CAV1$^{+/-}$ populations are measured for tumorsphere-initiating potential.

Cell surface markers are used to enrich for breast CSCs from heterogeneous cell populations,

e.g., combinations of CD44^{hi}, EPCAM^{lo}, CD104^{+} and CD24^{+} [14–16], yet only a subpopulation of those isolated cells actually contain tumor-initiating properties. The tumorsphere assay is an *in vitro* surrogate assay that measures tumor initiation capacity, where only CSCs are capable of seeding a tumorsphere when placed as single cells in suspension culture [2]. To improve the identification of the tumorsphere cell-of-origin we first use the output of TrajectoryNet's novel proliferation model to identify the subpopulation of CD44^{hi} cancer stems cells that disproportionately initiate tumorspheres. At the day 0 time point comprising bulk 2D cultured cells we observe CD44^{hi} and CD44^{lo} cell populations as expected (Figure 3A). Using TrajectoryNet, we identified a CD44^{hi} cell population with a high proliferation rate and computationally predicted their cell cycle phase [13] (Figure 3B). This analysis reveals a significant overlap between TrajectoryNet's inferred proliferation rate and cells in the G2/M and S cell cycle stage, suggesting that those cells differentially contribute to tumorsphere development (Figure 3B). To test this hypothesis, we isolated CD44^{hi} cells from each cell cycle phase by flow cytometry using the FUCCI cell cycle sensor system [17] and tested each population for tumorsphere-forming potential *in vitro*. As predicted, CD44^{hi} cells that reside in the S and G2/M phase are significantly enriched for tumorsphere-forming potential compared to unsorted or G1 phase CD44^{hi} cells (Figure 3C).

We then use TrajectoryNet to identify new cell surface markers that improve and refine CSC isolation. We perform differential expression analysis between CD44^{hi} and CD44^{lo} cells and identify surface markers highly expressed in the CD44^{hi} population whose expression overlaps with S/G2 phase cells including *PTN*, *EPCAM*, *CAV1*, *MMP7*, *VCAN* and *ANAX5* (Figure 3D). Using Density Re-sampled Estimation of Mutual Information (DREMI) [6] to measure non-linear associations between the expression of these markers and proliferation rate, we confirmed that each gene associates either positively or negatively with proliferation (Figure 3D). This analysis shows that EPCAM (epithelial cell adhesion molecule) is, as predicted from previous literature, enriched in this population, and newly, that CAV1 (caveolin 1) positively associates with the TrajectoryNet predicted proliferation rate within the CD44^{hi} day 0 subpopulation (Figure 3D).

To validate that EPCAM and CAV1 further enrich for tumorsphere-forming cells within the CD44^{hi} fraction, we isolated single EPCAM^{+} cells, single CAV1^{+} cells, and double EPCAM^{+}CAV1^{+} cells from the CD44^{hi} day 0 subpopulation using flow cytometry. After measuring tumorsphere-forming potential, we found that neither EPCAM nor CAV1 alone further enriched for tumorsphere-forming potential, however, double EPCAM^{+}CAV1^{+} cells formed significantly more tumorspheres (1.7 fold; $p<0.05$), while EPCAM^{+}CAV1^{-} cells showed significantly reduced (10 fold; $p<0.001$) than unsorted CD44^{hi} cells (Figure 3E). Thus, TrajectoryNet enabled the identification of biological properties (cell cycle) and a refined CD44^{hi}EPCAM^{+}CAV1^{+} CSC marker profile that improve the identification and isolation of CSCs from heterogeneous cell populations.

Tracing CSC plasticity along the EMT and MET trajectories

To define the transcriptional trajectories that drive CSCs through the EMT towards a mesenchymal cell state or the MET trajectory towards an epithelial cell state, we zoomed into the day 30 sample. Louvain clustering [18] identified three cell clusters - one apoptotic cell cluster (defined by high expression of mitochondrial genes), one epithelial, and one mesenchymal cell cluster (defined by the EMT signature score) (Figure 4A) [11]. We confirmed the existence of the epithelial and mesenchymal cell states within the tumorspheres by immunofluorescent co-staining for ZEB1 (mesenchymal) or CDH1 (epithelial) protein expression at day 7 and day 28 (Figure 4B). This data shows fairly uniform ZEB1^{hi}CDH1^{lo} expressing cells at day 7, consistent with expansion of the CSC pool at this early timepoint. In contrast, we see a striking dichotomy of cells emerging by day 28, where some cells retain the mesenchymal ZEB1^{hi}CDH1^{lo} state, while others have transitioned to the ZEB1^{lo}CDH1^{hi} epithelial cell state. These data indicate that cell fate decision between the EMT and MET trajectories emerges after day 7.

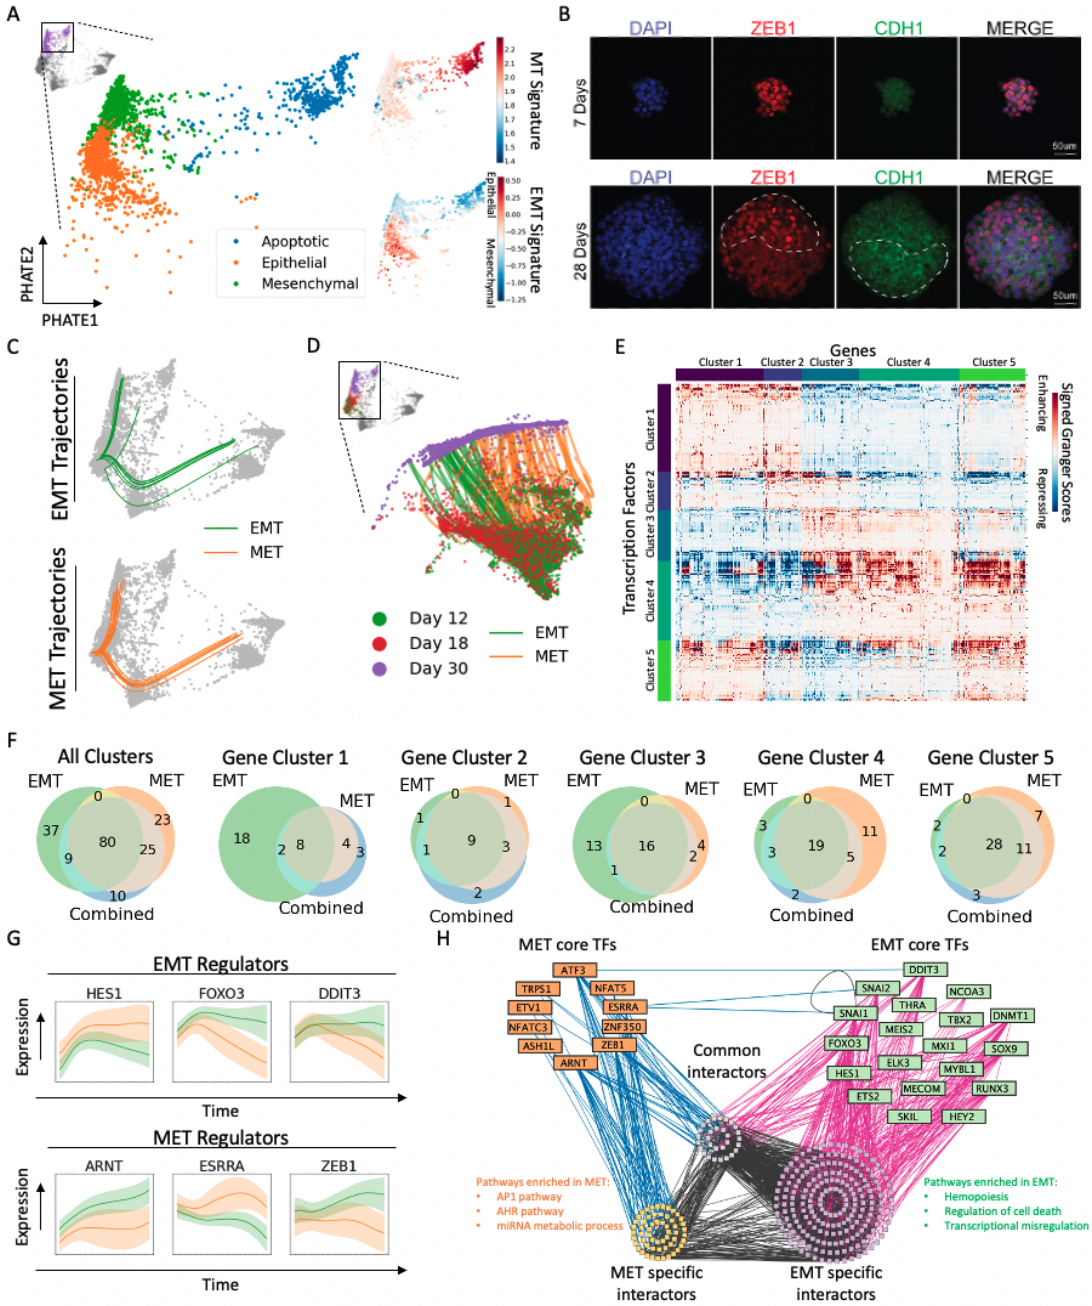

Figure 4: Comparing gene regulation in the EMT and MET trajectories. **(A)** PHATE visualization of day 30 scRNAseq timepoint with three populations highlighted corresponding to epithelial (orange), mesenchymal (green) and apoptotic (blue) populations (left). Populations were computed with Louvain clustering [18] and were identified using mitochondrial (MT) and EMT gene signatures (right). **(B)** Microscopy visualization of CDH1 and ZEB1 at day 7 and day 28 of the tumorsphere assay. Dashed white lines highlight distinct populations with high ZEB1 (mesenchymal) or CDH1 (epithelial) expression. **(C)** Tracing individual trajectories of cells that undergo EMT (green) and MET (orange). **(D)** 3D PHATE visualization of day 12, 18, and 30 showing the divergence of the EMT (green) and MET (orange) trajectories highlighting their mapping to discrete cell populations at day 30. **(E)** Heatmap of signed Granger values between 5,273 genes and 461 transcription factors organized by gene clusters from Figure 2C for combined trajectories. Red denotes a strong enhancing relationship between a transcription factor and a target gene, while blue denotes a strong repressive relationship. **(F)** Top regulatory transcription factors from EMT, MET and combined trajectories are separated by gene cluster (as computed in Figure 2C). The overlap between regulatory transcription factors from each of the EMT, MET and combined trajectories. **(G)** Visualizing inferred continuous gene expression dynamics of representative TFs regulating the EMT and MET. **(H)** Visualisation of gene regulatory networks of core EMT transcription factors (green) and MET transcription factors (orange) highlighting cross-talk through common gene interactors (grey nodes). The core MET transcription factors share regulatory relationships with a hub of MET-specific interactors (yellow nodes). Similarly, the EMT core transcription factors share regulatory relationships with a hub of EMT-specific interactors (purple).

To differentiate between the two cell fates, one more epithelial and one more mesenchymal, we trace the cellular trajectories starting from the two fates backward. We call one set of trajectories *EMT trajectories*, and the other *MET trajectories*, noting as before that they both start from the same hybrid CSC state (Figure 4C). Starting from CSCs, we show divergence of the EMT and MET trajectories between day 12-30 (Figure 4D). To elucidate the gene regulatory network underlying these trajectories, we identified the top 5,273 most highly and variably expressed genes (genes that were above the fiftieth percentile in their dispersion and expression scores) as well as the top 461 most highly and variably expressed transcription factors, which were chosen by the same metrics. We performed causality analysis using our signed Granger framework as described previously for every transcription factor-gene combination in the unpartitioned trajectory (Figure 4E), as well as the individual EMT and MET trajectories (Figure S4). Visualizing these signed Granger values across Combined, MET and EMT trajectories, we see clear enhancing regulatory relationships between transcription factors expressed in one gene cluster with genes expressed in the same and the following clusters. Furthermore, we also identify concomitant repressive regulatory relationships with genes expressed in previous clusters. For instance, transcription factors in Cluster 2 have a strong enhancing regulatory relationships with genes in Cluster 2 and 3 while having repressive regulatory relationships with genes in Clusters 1 (Figures 4E).

Using the TGCS framework for each of the combined, EMT and MET trajectories, we found the top 100 transcription factors that had the most regulatory effect within their respective trajectory. From these analyses we identify the top 80 highly regulatory transcription factors shared between the combined, EMT and MET trajectories (Figure 4F), suggesting that these factors are required for survival within the tumorsphere regardless of cell fate. We identify 37 highly regulatory core transcription factors unique to the EMT trajectory and 23 highly regulatory core transcriptional factors unique to the MET trajectory (Figure 4F and Supplementary Figure S5A,B). For example, in the EMT trajectory, *HES1* is down-regulated, while *FOXO3* and *DDIT3* expression remain high (Figure 4G). In the MET trajectory, *ARNT* and *ZEB1* remain low, while *ESRRA* levels follow an arc trajectory that increases early in the trajectory. then decreases at the end point when the epithelial state emerges (Figure 4G). Interestingly, 18 of 37 EMT-specific genes appear in Cluster 1, while the majority of MET-specific genes (11 of 23) appear later in Cluster 4 (Figure 4F). This data confirms the earlier observation that commitment to the MET trajectory occurs later in tumorsphere development (Figure 4B).

We use the EMT database (dbEMT 2.0 [19]) to determine if the core EMT and MET transcription factors identified by the TrajectoryNet pipeline have been implicated in the EMT. Interestingly, only 12 of the 37 EMT transcription factors (*SKIL, FOXO3, ELK3, SOX9, NCOA3, DNMT1, ETS2, SNAI1, RUNX3, HES1, SNAI2, TBX2*), and only 4 of the 23 MET transcription factors (*ESRRA, ETV1, TRPS1, ZEB1*), are in the database. These data show that the TrajectoryNet pipeline has uncovered a majority of new transcriptions yet to be associated with these cancer cell plasticity programs.

We therefore built gene regulatory networks around the EMT and MET core transcription factors using the TRRUSTv2 database (Figure 4H and Supplementary Figure S6) (see Methods). Of note, the EMT core hub (green) regulates a unique set of downstream EMT interactors (purple), while the MET core hub (orange) regulates a different set of unique downstream MET interactors (yellow). There is also a common set of downstream interactors (grey node) between the EMT and MET core hubs, suggesting that those genes lie at the interface of the two transcriptional circuits and may therefore play a causal role in cell-fate decisions between the two trajectories. Interestingly, there are minimal direct regulatory interactions between the core EMT and MET transcription factors themselves. Notable exceptions are the MET transcription factor *ATF3* regulation of *DDIT3*, and the MET transcription factor *ESRRA* regulation of the *SNAI1* and *SNAI2* transcription factors (Figure 4H).

Analysis of the top significantly enriched pathways ($p \leq 0.05$) within the MET subcluster (comprising 98 genes made up of the core MET transcription factors, unique downstream MET interactors, and, common interactors) were AP1 pathway, AHR pathway, and regulation of miRNA metabolic process.

Top pathways ($p \leq 0.05$) enriched in the EMT subcluster (comprising 251 genes made up of the core EMT transcription factors, downstream EMT interactors, and, common interactors) were hemopoiesis, regulation of cell death and transcriptional misregulation in cancer pathways (Figure 4H and Supplementary Figure S6). These analyses resolve, for the first time, the distinct gene regulatory networks and signalling pathways required to drive CSC plasticity along the EMT versus MET trajectory.

Validating the temporal MET gene regulatory network

The MET is critical for driving CSC differentiation into epithelial non-CSC progeny to create the heterogeneity required for robust tumor growth [20]. Yet to date, there is no comprehensive regulatory network that describes the MET gene regulatory network. We therefore used TrajectoryNet and the Granger Causality analysis to build and validate one based on our data (see Methods). Briefly, we used the TRRUST v2 database to extract and visualise the gene regulatory relationships of the 23 core transcriptional factors regulating the MET trajectory (Figure 5A) [21]. In this network, the core transcription factors are annotated by solid rectangles (Figure 5B). Using the EMT database (dbEMT 2.0 [19]), we show that the core MET transcription factors not only regulate known EMT genes (circles), but also novel genes yet to be associated with the EMT (diamonds) (Figure 5B) .

By color-coding and organizing the core transcription factors and their associated gene regulatory networks by Gene Cluster number, we visualise the temporal cross-talk between early time point Clusters 2 (green) and 3 (blue) and late time point Clusters 4 (orange) and 5 (red) (Figure 5B). Notably, ATF3 is involved at the earliest time point of the MET initiation (Gene Cluster 2) followed by *ESRRA*, *ETV1*, *NFATC2* (Gene Cluster 3), *NFAT5*, *ASH1L* (Gene Cluster 4), and then *ZEB1*, *ARNT*, TRSP1 and *ZNF350* (Gene Cluster 5). We also show that the Gene Cluster nodes commonly interact with four factors (*JUN*, *AHR*, *AR*, *ESR1*) (white circles), suggesting that these genes may serve as intermediary signaling nodes between the gene clusters (Figure 5B).

Uniquely, *ESRRA* and *PAX9* (Gene Cluster 3) expression peak midway through the MET trajectory, then decrease as the epithelial cell state emerges, whereas all other core MET transcription factors are down-regulated (Figure S5A). Due to poorly defined regulatory relationships of *PAX9* in public data we were unable to determine its direct regulatory impact on the MET network. We did, however, observe clear interactions for *ESRRA*. Thus, the fact that *ESRRA* expression peaks early in the MET trajectory (Cluster 3), and that ESRRA directly suppresses *SNAI1* and *SNAI2*, suggests that ESRRA may have two critical functions in the MET trajectory: (i) to initiate the MET transcriptional program through it's direct downstream interactors in the MET circuitry, and (ii) directly suppress the core EMT transcriptional circuitry via down-regulation of *SNAI1* and *SNAI2*. Accordingly, we delved further into the regulatory role of ESRRA as a putative instigator of the MET trajectory (Figure 5C).

While initially regarded as an orphan nuclear receptor, recent work has shown that ESRRA is closely related to the estrogen receptor and is commonly associated with poor outcome and poor prognosis in triple-negative breast cancer [22]. Reducing ESRRA expression via antagonists or gene knockdown decreases cell proliferation and tumorigenicity [22,23]. Importantly, ESRRA directly alters the expression of EMT transcription factors *SNAI1/2* to regulate the EMT in triple-negative breast cancers [24–27]. Accordingly, based on published interactions and gene-expression trends within the MET network, we identified a module (Figure 5C) wherein ESRRA directly (*SNAI1*, *SNAI2*) and indirectly (*ZEB1*) suppresses core EMT transcription factors to enable the upregulation of CDH1 and emergence of the epithelial cell state, whilst also modulating AHR- and AR-dependent pathways that may be critical for the emergence of the epithelial cell state. Indeed we have shown that AR is critical for induction of the EMT that creates new CSC populations in response to chemotherapeutic insults [28]. Furthermore, ESRRA interacts with several core intermediary nodes, JUN, AR, ESR1, and with core transcription factors in Cluster 5 (AHR, ARNT) (Figure 5C). Together, these data suggest the ESRRA sub-network

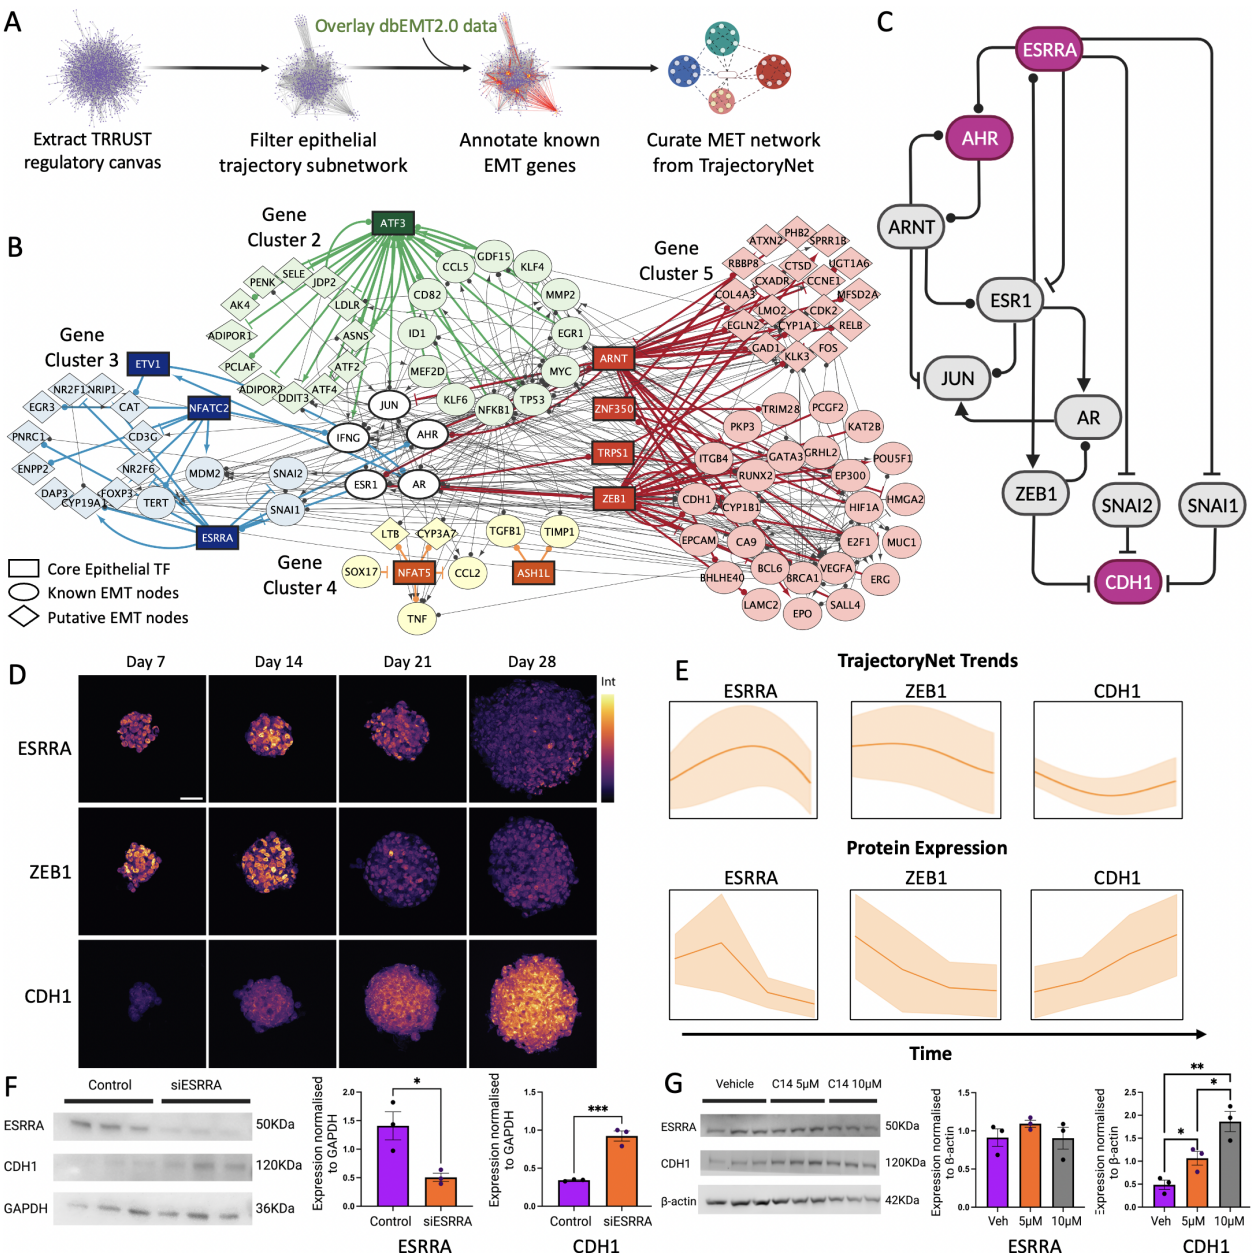

Figure 5: MET temporal gene network identification and validation. **(A)** Schematic of the workflow and filtering strategy used to curate the temporal MET gene regulatory network using TRRUST v2 database (Transcriptional Regulatory Relationships Unraveled by Sentence-based Text mining [8]), overlayed with genes associated with EMT (Epithelial Mesenchymal Transition Gene Database, dbEMT 2.0 [19]). **(B)** The resultant MET network comprises transcription factors identified by TrajectoryNet (rectangles) across Gene Clusters 1-5. Known E/M plasticity genes are highlighted by ellipses, and novel E/M plasticity genes are marked by diamonds. **(C)** Visualization of ESRRA gene regulatory module within panel B limited to known gene interactions with ESRRA and the known EMT genes (CDH1, ZEB1, SNAI1/2). Genes upregulated in the MET trajectory are highlighted in pink and genes downregulated are highlighted in grey. **(D)** Visualization of ESRRA, ZEB1 and CDH1 by immunofluorescence staining at 4 timepoints in 3D tumorspheres. **(E)** Visualization of TrajectoryNet interpolated gene trends as well as ground truth protein trends for ESRRA, ZEB1 and CDH1. **(F)** Effect of ESRRA knockdown (siRNA) on CDH1 expression in HCC38 CD44^{hi} cells by western blot. **(G)** Effect of ESRRA inhibition using C14 on CDH1 expression in HCC38 CD44^{hi} cells by western blot.

is a regulatory circuitry traversing temporal gene clusters to drive the MET trajectory and fortify the emergence of the epithelial cell state.

To test this hypothesis, we analysed ESRRA protein expression across the 3D tumorsphere time-course by immunofluorescence (IF). We also analyzed protein levels of ZEB1, to mark the mesenchymal cell state, and CDH1, to mark epithelial cells (Figure 5D). These studies confirm critical temporal regulation of ESRRA during the tumorpshere time-course. First we note that ESRRA is co-expressed with ZEB1 at day 7 demonstrating that it is present in the CSC population. However, we observe a marked increase in ESRRA expression from day 7 to day 14. Interestingly, ZEB1 down-regulation follows peak ESRRA expression. ESRRA then decreases for the remainder of the time-course, such that at day 28 ESRRA and ZEB1 are down-regulated, while CDH1 is up-regulated marking the presence of epithelial cell state. These protein expression findings are confirmed when comparing TrajectoryNet inferred expression dynamics with the protein (IF) trends (Figure 5E). Together, these data confirm unique temporal gene expression patterns not previously identified, whereby ESRRA expression is required for the CSC state, yet an increase in ESRRA expression triggers a down-regulation of EMT transcription factors SNAI1, SNAI2, and ZEB1, thereafter, ESRRA itself is downregulated to enable CDH1 upregulation and the emergence of the epithelial cell state. To confirm a causal relationship of ESRRA expression with the emergence of the epithelial cell state as observed at day 28, we demonstrate that siRNA and or ESRRA antagonist treatment of 2D cultured CD44^{hi} cells lead to a significant increase in CDH1 expression ($p \leq 0.05$) (Figure 5F-G and Figure S7).

Thus, TrajectoryNet enabled identification of a core regulatory unit defining the first comprehensive temporal MET regulatory network in breast cancer with ESRRA confirmed as a CSC marker and early key regulator of cell fate decisions towards the epithelial cell state.

Validating TrajectoryNet gene expression dynamics from primary tumors to lung metastases

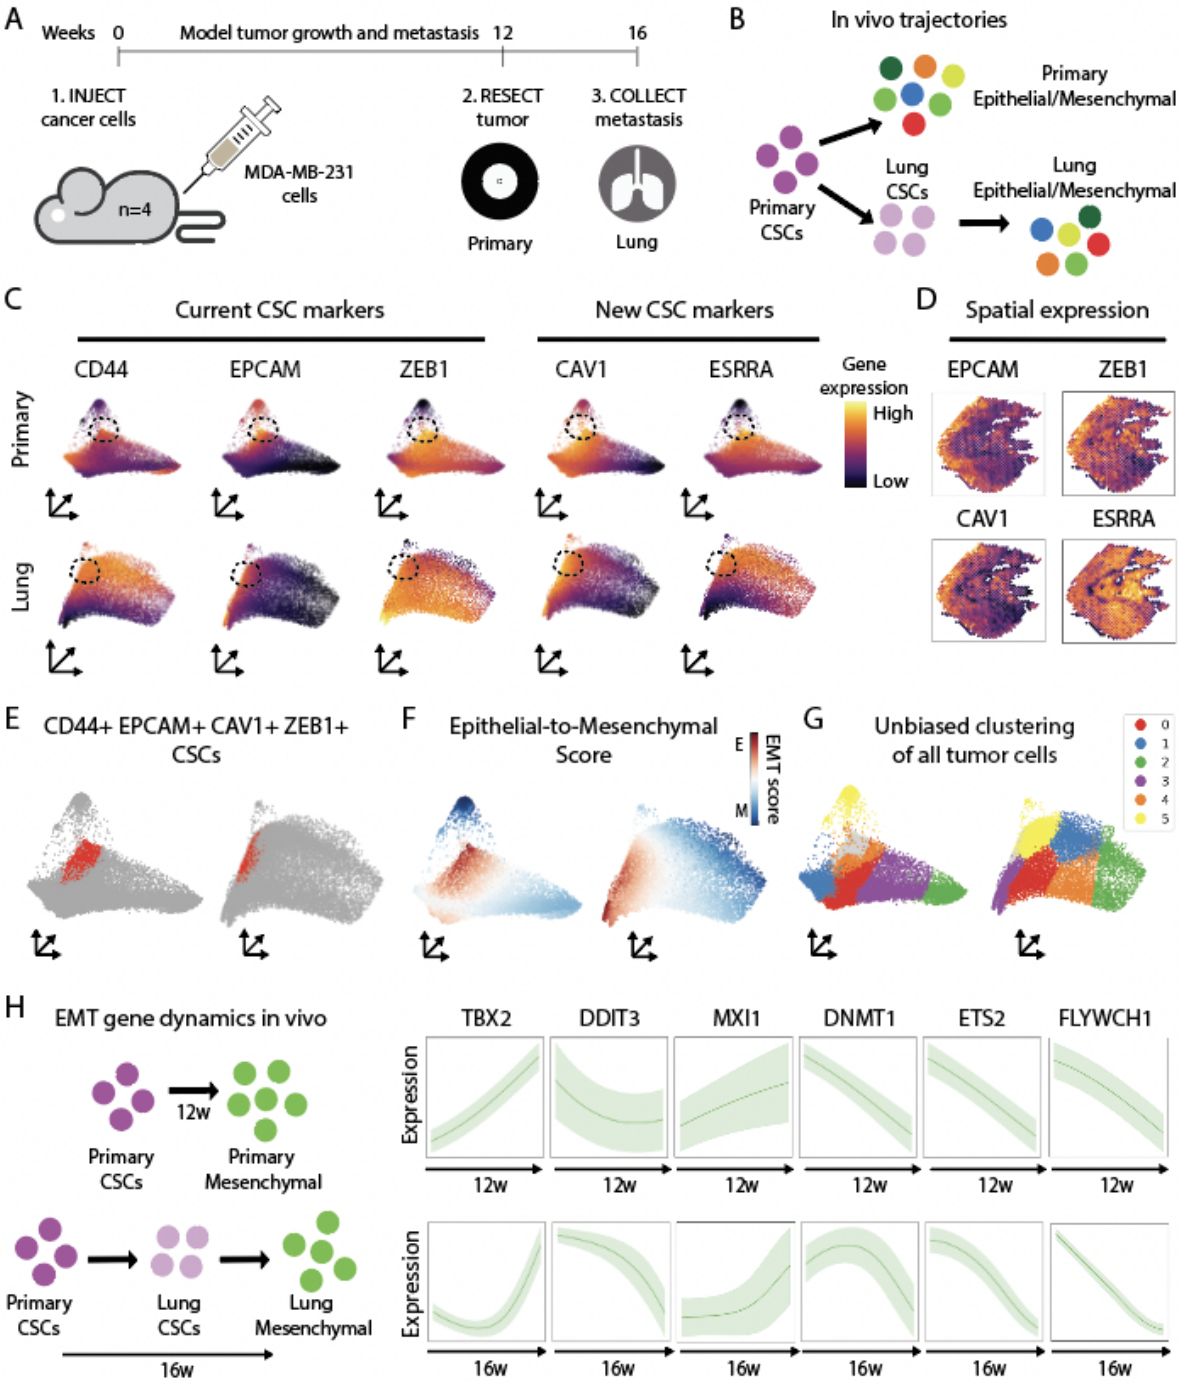

Figure 6: Validating cancer cell plasticity trajectories *in vivo*. **(A)** Schematic of *in vivo* scRNAseq experimental workflow on primary tumors implanted in the mammary fat pad with matched spontaneous lung metastases from the xenograft MDA-MB-231 triple-negative breast cancer model. **(B)** Schematic of learned trajectories from scRNAseq data developing 1) within a primary tumor and 2) from a primary tumor to lung metastasis. **(C)** Expression of current and new CSC markers shown on scRNAseq data from primary tumor and matched lung metastasis, and **(D)** on matched spatial transcriptomic profiling of the primary tumor. **(E)** Identification of primary tumor and lung metastasis CSCs (red) based on high expression of *CD44*, *EPCAM*, *CAV1*, and *ZEB1*. **(F)** EMT score of primary tumors and lung metastases [11] **(G)** Clustering of primary tumors and lung metastases into six subpopulations. Gray cluster represents CSCs. Green cluster represents emergent mesenchymal subpopulation. **(H)** Transcriptional dynamics of core EMT transcription factors in primary tumor trajectory (top row) and primary tumor to lung metastasis trajectory (bottom row)

In this manuscript, we follow the processes of EMT and MET using *in vitro* 5-timepoint data sampling. While this allows us to follow the dynamics of cancer cell plasticity carefully in time, we also wanted to establish the relevance of these dynamics to *in vivo* systems. To do this, we used a xenograft model of metastatic triple-negative breast cancer. We performed scRNAseq and spatial transcriptomics on primary tumors and matched lung metastases from four mice bearing the MDA-MB-231 xenograft. Following orthotopic implantation in the mammary fat pad, primary tumors were resected at 12 weeks, lung metastases at 16 weeks, and both samples were analyzed thereafter by scRNAseq (Chromium 10X) and spatial transcriptomics (Visium 10X, primary tumors only) (Figure 6A). We embedded the scRNA-seq data into 3 dimensions using PCA and applied TrajectoryNet to map two trajectories. Firstly, we mapped the differentiation of CSCs within a primary tumor, and secondly, we mapped the trajectories of CSCs originating in the primary tumor that migrate to the lung to form a metastasis (Figure 6B).

To validate the refined CD44^{hi}EPCAM^{+}CAV1^{+} CSC marker profile, we first visualized regions of the primary tumors and lung metastases expressing high levels of established CSC markers *CD44*, *EPCAM* and *ZEB1* (dotted circles) (Figure 6C). Supporting the validation of CAV1 as a new cell surface marker to refine identification of the CSC state, we confirm CAV1 is positively correlated in both samples with CSC markers CD44 (primary ρ=0.19, lung ρ=0.40), EPCAM (primary ρ=0.86, lung ρ=0.97), and ZEB1 (primary ρ=0.92, lung ρ=0.48) (Figure 6C). To further validate these findings, we integrated the primary tumor scRNAseq data with matched spatial transcriptomic data using scMMGAN [29], and show their spatial overlap in primary tumors (Figure 6D). The co-localization of these CSC markers supports the identification of regions of primary and lung metastases enriched for cells residing in the refined CSC state (Figure 6E) that was identified by the prior cell-of-origin analysis (Figure 3D).

Next, we applied the EMT score [11] to identify regions of the tumors enriched for the epithelial or mesenchymal cell state (Figure 6F). This analysis shows that CSCs reside in areas of the tumor enriched for epithelial markers, reflecting their hybrid epithelial/mesenchymal phenotype. Radiating out from the CSC-enriched area we observe a predominant emergence of mesenchymal cells, suggesting that CSCs in this *in vivo* model predominantly follow an EMT trajectory. This is not surprising given that MDA-MB-231 cells retain their CD44^{hi} state *in vivo*, which is associated with a more mesenchymal phenotype [14]. We then performed unbiased clustering to identify the region of the primary tumors and lung metastases associated with the most mesenchymal cell state (green) to highlight the emergent EMT trajectory (Figure 6G).

Finally, we examined the dynamics within the *in vivo* EMT trajectories of genes associated with components of the newly-defined core EMT hub from Figure 4H, including *TBX2*, *DDIT3*, *MXI1*, *DNMT1*, *ETS2* and *FLYWCH1*. We observe complementary temporal gene expression patterns associated with the EMT trajectory within the primary tumor (and highlight their unique spatial distributions within the primary tumor; Figure S8), and within the primary to lung metastasis trajectory (Figure 6H). For example, TBX2 is linearly up-regulated in the primary tumor trajectory, whereas, in the primary to lung metastasis trajectory, its expression is initially flat at the beginning, reflecting the dynamics of the primary CSC to lung CSC portion of the trajectory, and thereafter linearly increases reflecting the lung CSC to lung mesenchymal portion of the trajectory. These dynamics are also reflected in the continuous expression of *DDIT3*, *MXI1*, *DNMT1* and *ETS2* (Figure 6H). These analyses demonstrate that the refined EMT core circuitry, that was identified in an *in vitro* model system, is conserved across distinct triple-negative breast cancer cell lines (HCC38 versus MDA-MB-231), from *in vitro* to *in vivo systems*, and across tumors growing in vastly different tissues (mammary fat pad versus lung tissue).

Together these analyses show that the TrajectoryNet pipeline identifies critical markers enabling improved classification and identification of the CSC state, and, defines biological trajectories that define distinct CSC trajectories and ultimately, cell fates, across disparate timepoints and organs in *in vitro* and *in vivo* model systems. Accordingly, the ability to characterize and quantify a tumor's CSC component as a measure of tumor aggressiveness, coupled with the discovery of temporal gene regulatory strategies to

drive CSCs towards different cell fates, is a critical advance for the discovery of temporal gene regulatory networks underlying these critical biological processes.

Discussion

Dynamic cell state evolution is a hallmark of many diseases. In cancer, dynamic transitions manifest as changes from relatively benign cell states to highly aggressive, proliferative and therapy-resistant cell states. Until now, it has been difficult to study these dynamics comprehensively because single-cell and other high throughput data modalities have largely captured static snapshot measurements. While learning dynamics has often been a holy grail for this type of data, previous efforts have significant shortcomings. Pseudotime [5] and RNA velocity [4] methods are useful in certain contexts, however they are not able to compute trajectories at the single cell level, interpolate between distant or disconnected distributions of cells, or learn transcriptional networks underlying cellular state transitions. To address the knowledge-gap in learning data dynamics, we developed the TrajectoryNet pipeline, which implements a breakthrough neural network paradigm called neural ODE to learn a *dynamic optimal transport* between time-lapsed single cell populations. We then use the gene dynamics information from TrajectoryNet to perform time-lagged Granger causality analysis to learn dynamic transcriptional networks that drive trajectories forward. We demonstrate that TrajectoryNet is not only able to infer trajectories significantly better than other OT [3], splicing [4] and graph based [5] methods, but also can more accurately infer gene regulatory relationships than methods that do not account for cellular dynamics, like correlation and mutual information [6] based techniques.

Here, we applied TrajectoryNet to identify the cellular initiators and molecular drivers of CSC plasticity using a triple-negative breast cancer model. First, we show that TrajectoryNet enables improved characterization and isolation of cells residing in the CSC state. Using the tumorsphere assay, an *in vitro* surrogate of tumor-initiating potential, we trace cells at the end of the trajectory back to their cell-of-origin within the starting population of cells, thereby identifying the proliferative subpopulation of cells most likely to represent the CSC population that initiates tumorspheres. Applying this strategy, TrajectoryNet reveals a refined CD44^{hi}EPCAM^{+}CAV1^{+} CSC marker profile that we validate to identify cells significantly enriched for tumorsphere-forming potential above the CD44^{hi}EPCAM^{+} markers previously described [9]. This strategy can now be applied in other settings, for example, to define cell populations that drive site-specific metastasis, or to identify cells that emerge in therapy-resistant cell states following cytotoxic treatments. These types of analyses will have profound impact on the future development of therapeutic strategies to specifically identify and target aggressive subpopulations of cancer cells.

Importantly, we demonstrate the ability of the TrajectoryNet pipeline to define detailed temporal analysis of transcription factors and their associated targets regulating dynamic CSC plasticity programs in 3D *in vitro* tumorsphere models and in *in vivo* models of breast cancer metastasis. Accordingly, we identify the first comprehensive gene regulatory networks defining the EMT trajectory that drive CSCs towards a mesenchymal cell fate, and the MET trajectory that drives CSC differentiation into epithelial non-CSC progeny. Given that the details of the MET program were largely unknown, we identified an MET subnetwork for further validation, with ESRRA as a potential upstream initiator. We show that ESRRA is present in CSCs, and that its expression peaks to initiate the MET then regresses to enable the emergence of the non-CSC state. Indeed, we show that the ESRRA regulatory network comprises layered and complex gene regulation acting to both initiate an epithelial program, potentially through regulation of AHR:ARNT signalling pathway [30], whilst simultaneously suppressing the mesenchymal program through direct and indirect regulation of the core EMT circuitry (*SNAI1*, *SNAI2*, *ZEB1* and *CDH1*). Consistent with our findings, previous studies have shown that ESRRA silencing results in the up-regulation of CDH1 in MDA-MB-231 cell lines (TNBC model) and HEC-1A cells (human endometrial

adenocarcinoma model) [31].

Thus, the identification of the ESRRA subnetwork as a strategy to regulate CSC plasticity could lead to the validation of novel targets that prevent metastatic outgrowth by driving CSC differentiation into a non-CSC state that is sensitive to chemotherapy. Future studies validating the role of the ESRRA network in metastasis and chemotherapy-resistance will shed light on the importance of these findings in the clinical setting [32–34].

Finally, we highlight the ability of the TrajectoryNet pipeline to reveal gene regulatory networks in complex cellular transitional models, demonstrated here spanning a 16 week *in vivo* time course. We show temporal regulation of a subset of newly-defined regulatory components of the EMT, highlighting complementary gene expression dynamics that hold true across the *in vitro* and *in vivo* models. The TrajectoryNet pipeline is a breakthrough method that can unravel complex temporal genomic networks governing cell state dynamics. Here we have demonstrated the ability of TrajectoryNet to reveal novel dynamic cancer gene networks, however, our pipeline can be used to study any type of dynamic transition captured via single cell technology over multiple timepoints including stem cell differentiation, response to therapeutic interventions, or infections. As longitudinal data becomes increasingly available and critical in the study of dynamic systems, TrajectoryNet will be applied to a broader set of datasets and types.

Author Contributions

Alexander Tong and Manik Kuchroo jointly led this work under the supervision of Christine Chaffer and Smita Krishnaswamy who were collectively responsible for writing the manuscript. Smita Krishnaswamy, Alexander Tong, Manik Kuchroo, Aarthi Venkat and Brandon Zhu were responsible for algorithm development, running computational comparisons and applying TrajectoryNet to biological datasets under the supervision of Smita Krishnaswamy. Beatriz Perez San Juan was responsible for single cell experiments, FACS analysis and cell sorting, and tumorsphere asssays, Laura Rangel was responsible for tumorsphere assays and fluorescent staining, Shabarni Gupta was responsible for network development, analysis and validation all under the supervision of Christine Chaffer. John Lock performed the analysis of immunofluorescence experiments.

Declaration of Interests

Dr. Krishnaswamy is a visiting professor at META Fundamental AI Research.
Dr. Christine Chaffer is the Founder and Managing Director of Kembi Therapeutics.

Methods

TrajectoryNet Analysis Details

Overview of TrajectoryNet

TrajectoryNet [1] learns a continuous normalizing flow between cross-sectional snapshot measurements in such a way that the flow represents biologically plausible paths of development and differentiation. The key to obtaining paths that adhere to our previous knowledge of cell development was to penalize a continuous normalizing flow [35], implemented as a neural ODE, such that it performs dynamic optimal transport [36]. This enables us to efficiently solve the problem of cells evolving over time as a transport problem which is penalized over path lengths of the cells (See Eq. 3).

Neural Ordinary Differential Equation Model

TrajectoryNet is built on a neural ordinary differential equation (neural ODE) framework [37]. As previously defined, a neural ODE defines a neural network f with parameters θ which defines a time derivative for the flow at every point and time. Thus for any time t the flow is defined by $\frac{\partial x(t)}{\partial t} = f_\theta(x(t), t)$. At some initial time t_0, for any $x(t)$, we have that $x(t) = F_{\theta,\text{solver}}(x_{t_0}, t) = x(t_0) + \int_{t_0}^{t} f_\theta(x(t), t)dt$. This integral could be computed using Euler integration where at a fixed set of times $t_0, t_1, \ldots t_T$ we calculate $x_{t_{i+1}} = x_{t_i} + (t_{i+1} - t_i) f_\theta(x_{t_i}, t_i)$ incrementally. However, this can accumulate significant error depending on the dynamics, therefore we use ODE solvers with lower error and or adaptive step sizes. An important question is how to compute the loss with respect to the parameters θ. For this we use the adjoint method. Given a loss function $L(F_{\theta,\text{solver}}(x_{t_0}, t), x_t)$ then we can take the gradient with respect to the parameters θ using the adjoint dynamics $a(t) = \partial L / \partial x(t)$. By running the adjoint dynamics, we can compute $\partial L / \partial \theta$ without having to store the intermediate steps of the forward-time ODE solver, which greatly reduces memory usage.

An important note here is that F is approximately invertible due to the deterministic dynamics. However, the time derivative determined by the network f need not be invertible, and is representable by any standard neural network architecture. This is in contrast with the architectures used in invertible neural networks that directly represent an invertible function, and must be severely restricted for fast Jacobian computation [38].

Continuous Normalizing Flows

While TrajectoryNet models the flow of cells over time, it is trained from cross-sectional population data. At a set of timepoints $\mathcal{T} = \{t_1, t_2, \ldots, t_T\}$ we have a corresponding set of discrete samples $\mathcal{X} = \{X_{t_1}, X_{t_2}, \ldots, X_{t_T}\}$ from the population of viable cells at that time. Importantly, as with other destructive measurement techniques, a cell present at time t_i is destroyed and is cannot be measured at time t_{i+1}. Instead, we must infer the most likely dynamics of single cells from the population that is present at each timepoint. To accomplish this we formulate the optimization as a *continuous normalizing flow* (CNF), which transforms an initial known at time t_0 to match an empirical distribution at time t_1. Since $x_{t_1} = F(x_{t_0}, t_1)$ is approximately invertible by definition, and we can efficiently calculate the divergence of the probability distribution then given a probability distribution $P(t_0)$ we can find $P(t_1) = F_\#(\rho(t_0))$ where $F_\#$ denotes the pushforward operator of F, i.e. $F_\#(\rho(t_0))(B) = \rho(t_0)(F^{-1}(B))$ for any subset $B \subseteq \mathbb{R}^d$. A continuous normalizing flow also defines a likelihood at every point x_{t_1}. We can calculate $\log P(x_{t_1})$ for any x_{t_1} using the change of variables, and in particular the instantaneous change of

variables as derived in [37].

$$\log \rho(x_{t_1}, t_1) = \log \rho(x_{t_0}, t_0) + \int_{t_0}^{t_1} -\mathrm{tr}\left(\frac{df}{dx(t)}\right) \tag{9}$$

where $p(x, t_0) := \mathcal{N}(0, 1)$ is the density of a standard Gaussian distribution. This log-likelihood definition allows the CNF to be trained using maximum likelihood, i.e. trained using the loss

$$\mathbb{E}_{x \sim \rho(x_{t_1}, t_1)} \log \rho(x, t_1) \tag{10}$$

Continuous Normalizing Flows with Multiple Target Timepoints

Previously continuous normalizing flows were applied to a single timepoint with a known distribution at time t_0 and a single empirical distribution at time t_1. This can be used for generative modeling of complex distributions such as distributions of images [37] or medical charts [39]. Here, we would like to model a series of distributions over time with one smooth time varying flow f_θ. In this work we applied a few tricks to train this longer flow over time. One way to accomplish this would be to integrate points sampled from $x_{t_T} \sim X_{t_T}$ to $X_{t_{T-1}}$ and so on backwards to X_{t_1} and at each time step apply a distribution matching loss between empirical distributions $F^{-1}(x_{t_T}, t_i)$ and $x_{t_i} \sim X_{t_i}$ at $i \in [1, T)$. Such a loss is used in [40, 41]. However, distribution type losses are difficult to calculate and to optimize, and while there is significant ongoing work improving the stability, efficiency, and flexibility of such losses, there are still significant drawbacks.

Instead, we tackle this by integrating samples from all timepoints in parallel backwards to a timepoint t_0 where we can define a continuous probability density function $\rho(t_0)$. With a single integration we are able to calculate, and therefore maximize, $\log(\rho(t_i)(x_{t_i}))$ for all $i \in [1, T]$.

Optimal Transport

Informally, optimal transport considers the problem of moving one pile of mass to another at minimum cost. Optimal transport considers a *ground distance* defined between points and lifts this to distances between measures. Define the metric measure space $\mathcal{M} = (\mathcal{X}, d)$, then the squared 2-Wasserstein distance between two measures defined on $\mathcal{M}$, μ, ν is

$$W(\mu, \nu)_2^2 = \inf_{(pi \in \Pi(\mu, \nu)} \int d(x, y) d\pi(x, y) \tag{11}$$

where $\Pi(\mu, \nu)$ is the set of joint distributions with marginals μ and ν respectively and $\int_{\mathcal{X}} \mu = \int_{\mathcal{X}} \nu = 1$. The optimal transport distance can be thought of the minimal total cost of moving mass piled at μ to mass piled at ν. This optimization is in general difficult to compute for measures over a continuous space such as $\mathcal{X} \subseteq \mathbb{R}^n$. We will show how to approximately solve a dynamic version of the optimization with TrajectoryNet in the following sections.

Unbalanced Optimal Transport

A common variant of optimal transport is unbalanced optimal transport where the mass equality constraint is relaxed. Intuitively, there may be cases where for some additional cost, in addition to transporting mass, we can "teleport" mass for some additional cost. The unbalanced transport problem balances between this transport and teleportation cost for each mass. The unbalanced optimal transport problem is defined with respect to some ϕ-divergence D_ϕ often taken to be the Kullback-Leibler divergence

(D_{KL}), but in principle can be used with more general ϕ divergences. Given D_ϕ and a regularization parameters λ_μ, λ_ν we define the unbalanced optimal transport problem as

$$\mathrm{UW}(\mu, \nu, \lambda) = \inf_{\pi \in \Pi(\mu', \nu')} c(x, y) d\pi(x, y) + \lambda_\mu D_\phi(\mu, \mu') + \lambda_\nu D_\phi(\nu, \nu') \tag{12}$$

where $\lambda = (\lambda_\mu, \lambda_\nu)$ controls the relative cost of transportation vs. teleportation. To model cell division and cell death, we will use unbalanced transport.

Dynamic Optimal Transport

Dynamic optimal transport adds a time component to the static version of optimal transport. Adding a time component is useful as it gives a way to link the theory of optimal transport to that of dynamical systems and in particular fluid dynamics [36]. Instead of optimizing over the matching between distributions, the dynamic formulation considers an optimization over time-parameterized paths. For a given interval $[t_0, t_1]$ with a source distribution μ at time t_0 and a target distribution ν at time t_1, we define a time-dependent probability distribution $p_t(x)$ and a time dependent vector field $f(x, t)$, such that if the probability distribution evolves according to the continuity equation

$$\partial_t \rho + \nabla \cdot (\rho f) = 0 \tag{13}$$

for $t_0 < t < t_1$ and $x \in \mathbb{R}^d$, and the conditions

$$\rho(\cdot, t_0) = \mu, \quad \rho(\cdot, t_1) = \nu \tag{14}$$

we can relate the L^2 2-Wasserstein distance to (ρ, f) in the following way

$$W(\mu, \nu)_2^2 = \inf_{(\rho, f)} (t_1 - t_0) \int_{\mathbb{R}^d} \int_{t_0}^{t_1} \rho(x, t) \| f(x, t) \|^2 dt dx \tag{15}$$

In other words, a velocity field $f(x, t)$ with minimum L^2 norm that transports mass at μ to mass at ν when integrated over the time interval is the optimal plan for an L^2 Wasserstein distance. This can be seen by noticing that the optimal paths for each point pair (x_0, x_1) are geodesics and that $\inf_f \int_t \| f(x, t) \|^2 = d(x_0, x_1)^2$. This problem is in general challenging to solve, particularly in the high dimensional case where a discretization of space-time (such as used in [42] is impractical. We use CNFs to parameterize and learn the vector field f which approximates dynamic optimal transport.

Dynamic Optimal Transport with a Continuous Normalizing Flow

Continuous normalizing flows (CNFs) traditionally model an unknown data density by mapping it to a normal distribution using a continuous-time invertible function. While these are traditionally used only to "normalize" a single distribution, there is nothing in principle stopping us from mapping from $t_1 \to t_0 \to t_{norm}$ where we have a known normal density at t_{norm}. By adding these additional steps and an energy regularization, we can show that TrajectoryNet can approximate dynamic OT. Instead of the hard boundary constraint $\rho(\cdot, t_0) = \mu$ and $\rho(\cdot, t_1) = \nu$ as in Dynamic OT, we have a hard constraint at t_0, but relax the constraint at t_1 minimizing $D_{\mathrm{KL}}(\rho(\cdot, t_1), \nu$. Using the standard Lagrangian formulation, for large enough penalization these forms are equivalent. This is formalized in the following Theorem.

Theorem 1 (Theorem 4.1 [1]). *With time varying field $f(x,t) : \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}^d$ and density $\rho(x,t) : \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}^+$ such that $\int \rho(x,t)dx = 1$ for all $t_0 \leq t \leq t_1$ and subject to the continuity Eq. 13. There exists a sufficiently large λ such that*

$$W(\mu,\nu)_2^2 = (t_1 - t_0) \inf_{(P,f)} \mathop{\mathbb{E}}_{x_0 \sim \mu} \left[\int_{t_0}^{t_1} \|f(x(t),t)\|_2^2 \, dt \right] + \lambda KL(\rho(\cdot,t_1) \parallel \nu); \quad \textit{s.t. } \rho(\cdot,t_0) = \mu$$

By parameterizing our field f with a neural network, this theorem allows us to optimize for the dynamic optimal transport flows with deep learning techniques. For a full proof see [1], we provide a proof sketch here. This result can be seen by viewing λ as a Lagrange multiplier over the relaxing the constraint $\rho(\cdot,t_1) = \nu$. Here we can relax this constraint with any proper divergence, but chose the KL-Divergence because of its link to maximum likelihood. As $\lambda \to \infty$, $D_{\mathrm{KL}} \to 0$ approximately satisfying the equality constraint. See [1] for a more detailed analysis.

Modelling continuous trajectories with TrajectoryNet

We model cells as points in a continuous state-space $\mathcal{C} \subset \mathbb{R}^d$ over a time interval $\mathcal{T} = [0, t_k] \subset \mathbb{R}$. We assume that cells evolve via a drift through state space described by an ordinary differential equation (ODE):

$$dx(t) = f(x(t),t)dt \tag{16}$$

where $x(t) \in \mathcal{C}$ denotes the state of a cell at time t and f is a time varying vector field. Additionally, cells may grow or die according to a time varying proliferation rate modeled as $g(x,t) : (\mathcal{C},\mathcal{T}) \to \mathcal{R}$. $g(x,t) > 1$ represents a state where on average cells are proliferating and conversely $g(x,t) < 1$ denote states and times where cells are dying. At a population level, this can be modeled according to the continuity equation:

$$\partial_t M + \nabla \cdot (Mf) - Mg = 0 \tag{17}$$

which holds at every point $(x,t) \in (\mathcal{C},\mathcal{T})$.

We would like to learn the underlying dynamics from a cross-sectional observation model. We have a set of discrete samples $\bar{X} = \{X(t_i) : i \in [1,k]\}$ of the measurements from the continuous population of cells $\rho(t_i)$ present at k discrete timepoints in $[t_1, t_2, \ldots, t_k] \subset \mathcal{T}$, where $X(t_i)$ consists of n observations $\{x_j(t_i)\}_{j=1}^n$ in some space $\mathcal{X}$. Here, because of the destructive nature of the measurements, we assume if a cell is observed at t_i it is not observed at any $t_j \neq t_i$.

Nevertheless, the goal of TrajectoryNet is to learn the underlying dynamics of this space described by f and g in equation Eq. 17. To accomplish this, we parameterize f and g by neural networks f_θ and $g_{\theta'}$, and then train using a KL-divergence objective implemented using maximum likelihood in the style of *continuous normalizing flows* [37].

The full objective of TrajectoryNet can be decomposed into four parts as appears in Eq. 3 repeated here for clarity:

$$L(x(T),\theta,\theta') = \underbrace{\overbrace{D_{\mathrm{KL}}(M(x(T)||T))}^{\text{Distribution Matching}} + L_{energy} + \overbrace{L_{balance} + L_{proliferation}}^{\text{Proliferation penalties}}}_{\text{Unbalanced Dynamic OT}} + \overbrace{L_{density}}^{\text{Biological path priors}} \tag{18}$$

where

$$D_{\mathrm{KL}}(M(x(T),T)||\mu) = \log M(x(0),0) + \int_0^T \left[\log g(x(t),t) - Tr\left(\frac{\partial f(x(t),t)}{\partial x(t)}\right)\right] dt \tag{19}$$

$$L_{energy} = \lambda_{energy} \int_0^T \|f(x(t),t)\|_2^2 dt + \lambda_j \int_t \|J_f(\tilde{x})\|_F^2 \tag{20}$$

$$L_{balance} = \lambda_{balance} \int_0^T \left(1 - \int_{x\in\mathcal{C}} g(x(t),t)M(x(t),t)dx\right)^2 \tag{21}$$

$$L_{proliferation} = \lambda_{proliferation} \int_0^T g(x(t),t)^2 dt \tag{22}$$

$$L_{density} = \lambda_{density} \max(0, \text{min-k}\big(\{\|x(t)-z\| : z \in \bar{\mathcal{X}}\}\big) - \tau). \tag{23}$$

Performing Unbalanced Dynamic Optimal Transport Using A Proliferation Rate Neural Network

To model cell proliferation and death we need to model the system as an unbalanced optimal transport problem. The standard optimal transport problem is balanced in that all mass in the source distribution is moved to match the target distribution, however this does not model cell proliferation and cell death. The *unbalanced* optimal transport problem allows for the creation and destruction of mass (at a cost) thereby modelling cell proliferation and cell death. For a source measure μ and a target measure ν over a domain $\mathcal{Y} \subset \mathbb{R}^d$ let $\Pi(\mu,\nu)$ denote the set of all joint measures on $\mathcal{Y} \times \mathcal{Y}$ whose marginals are μ and ν. Then the squared Wasserstein-Fischer-Rao (WFR) metric is

$$WFR^2(\mu,\nu) := \frac{1}{2} \inf_{\rho,f,g} \int_{\mathcal{T}} \int_{\mathcal{C}} \left(\|f(x,t)\|^2 + g(x,t)^2\right) d\rho(x,t)dt \tag{24}$$

where $\rho(x,t)$ is a time-dependent density that interpolates between μ and ν, f and g are time-dependent vector and scalar fields respectively modeling the drift, proliferation and destruction of mass, and ρ, f, and g satisfy the continuity equation

$$\partial_t \rho + \nabla \cdot (\rho f) - \rho g = 0 \tag{25}$$

which matches Eq. 17 substituting the general measure M for a density ρ. This enforces balance of mass at a particular point in space time (x,t) i.e. that the change in mass $\partial_t \rho$ balances with the amount of mass leaving the point $\nabla \cdot (\rho f)$ and the amount of mass growth at that location ρg. The WFR is related to static formulations of unbalanced optimal transport in [43,44]. While the dynamic formulation is interesting theoretically, it presents numerical challenges, and thus most work has focused on discrete and static unbalanced transport. Previous work has explored the benefits of the unbalanced formulation in terms of robustness, generalization, and more accurate modeling of cell trajectories [3], however this method could only be applied to Euclidean geodesics, where we are able to learn more complex geodesics based on other knowledge about the system ($L_{density}$). These works focus on the discrete case, where the optimization is over a fixed number of cells and cannot be easily generalized to new cells. While the ancestors, descendants, and proliferation rate for measured cells are modeled, these parameters cannot be easily extended to any unmeasured cell states. TrajectoryNet learns continuously parameterized neural networks which allow us to optimize for the proliferation rate directly with our unbalanced transport portion of the loss in Eq. 18.

We now describe how the distribution matching, energy, balance, and proliferation losses approximate the WFR unbalanced dynamic optimal transport problem. The TrajectoryNet unbalanced OT loss can be thought of as a relaxation of the WFR optimization where the two constraints $\rho(x,T) = \nu(x)$, and

$\int_{\mathcal{C}} \rho(x,t) = 1$ are relaxed as regularizations. As the KL divergence term $D_{\mathrm{KL}}(\cdot||\cdot) \to 0$ the population distribution at time T approaches the observed density ν. Similarly, as the $L_{balance}$ term goes to zero, the constraint $\int_{\mathcal{C}} \rho(x,t) = 1$ is satisfied. The WFR distance can be formulated in dynamic terms following [43, 44] as

$$WFR^2(\mu,\nu) := \frac{1}{2} \inf_{f,g} \int_{x(0)\sim\mu} \int_{\mathcal{T}} \left(\|f(x(t),t)\|^2 + g(x(t),t)^2 \right) dt dx \tag{26}$$

such that $\int_{\mathcal{C}} \rho(x,t) = 1$ and $D_{\mathrm{KL}}(x(T),\nu) = 0$. This can be thought of as minimizing the cost following the average single point over time rather than minimizing the cost with respect to a fixed position in space time. From Eq. 26 it is easier to see the connection to our TrajectoryNet loss. L_{energy} and $L_{proliferation}$ make up the main terms within the integral and $D_{\mathrm{KL}}(\cdot||\cdot) \to 0$ and $L_{balance}$ enforce the constraints. Since our loss is evaluated by drawing samples from μ and integrating backwards in time if all of these losses are minimized with the correct relative λ regularizations weights such that $D_{\mathrm{KL}}(\cdot||\cdot) \to 0$ and $L_{balance}$ go to zero while L_{energy} and $L_{proliferation}$ are minimized, then we have computed the WFR distance.

Practical Considerations for Efficient Computation

We can approximate the first part of this continuous time equation using a Riemann sum as

$$\mathop{\mathbb{E}}_{x_0\sim\mu} \left[\int_{t_0}^{t_1} \|f_\theta(x,t)\|^2 \, dt \right] = \sum_{x\sim P_{t_0}} \sum_{i=t_0}^{t_1} \Delta t_i \, \|f_\theta(x,t)\|^2 \, . \tag{27}$$

This requires a forward integration using a standard ODE solver to compute as shown in [37]. If we consider the case where the divergence is small, then this can be combined with the standard backwards pass for even less added computation. Instead of penalizing $\|f_\theta(x,t)\|^2$ on a forward pass, we penalize the same quantity on a backwards pass.

We can also compute the change in the data log-likelihood using an augmented neural ODE. Here we add an additional dimension that represents the relative change in log probability over time.

$$\frac{d}{dt} \begin{bmatrix} x(t) \\ \log p(x(t),t)) \end{bmatrix} = \begin{bmatrix} f(x(t),t)) \\ \log \left[g(x(t),t) - \exp \left[\sum_{i=1}^{D} \frac{\partial f_i(x(t),t))}{dx_i} \right] \right] \end{bmatrix} \tag{28}$$

This means that the TrajectoryNet loss can be computed in a single integration of an ODE solver in $d+1$ dimensions.

For the energy regularization L_{energy}, in practice we found that both a penalty on the Jacobian or additional training noise helped to get straight paths with a lower energy regularization λ_e similarly to [45]. We found that a value of λ_e large enough to encourage straight paths, but unsurprisingly also shortens the paths undershooting the target distribution. To counteract this, we add a penalty on the norm of the Jacobian of f as used in [46, 47]. Since f represents the derivative of the path, this discourages paths with high local curvature, and can be thought of as penalizing the second derivative (acceleration) of the flow. Our energy loss is then

$$L_{energy}(x) = \lambda_e \int_t \|f(\tilde{x},t)\|^2 + \lambda_j \int_t \|J_f(\tilde{x})\|_F^2, \tag{29}$$

where $\|J_f(x)\|_F^2$ is the Frobenius norm of the Jacobian of f. Without energy regularization TrajectoryNet paths are unconstrained. However, with energy regularization we approach the paths of the optimal map. Our energy loss gives control over how much to penalize indirect, high energy paths.

Optimal transport is traditionally performed between a source and target distribution. Extensions to a series of distributions is normally done by performing optimal transport between successive pairs of distributions as in [3]. This creates flows that have discontinuities at the sampled times, which may be undesirable when the underlying system is smooth in time as in biological systems. The dynamic model approximates dynamic OT for two timepoints, but by using a single smooth function to model the whole series the flow becomes the minimal cost *smooth* flow over time.

Enforcing Transport on a Manifold

Methods that display or perform computations on the cellular manifold often include an implicit or explicit way of normalizing for density and model data geometry. PHATE [48] uses scatter plots and adaptive kernels to display the geometry. We would like to constrain our flows to the manifold geometry but not to its density. We penalize the flow such that it is always close to at least a few measured points across all timepoints through the following function:

$$L_{density}(x, t_d) = \sum_k \max(0, \text{min-k}\big(\{\|x(t_d) - z\| : z \in \mathcal{X}\}\big) - h) \tag{30}$$

This can be thought of as a loss that penalizes points until they are within h Euclidean distance of their k nearest neighbors. We use $h = 0.1$ and $k = 5$ in all of our experiments. We evaluate $L_{density}$ on an interpolated time $t_d \in (t_0, t_k)$ every batch.

Training

For simplicity, the neural network architecture of TrajectoryNet consists of three fully connected layers of 64 nodes with leaky ReLU activations. It takes as input a cell state and time and outputs the derivative of state with respect to time at that point. To train a continuous normalizing flow we need access to the density function of the source distribution. Since this is not accessible for an empirical distribution we use an additional Gaussian at t_0, defining $\rho_{t_0}(\cdot) = \mathcal{N}(0, 1)$, the standard Gaussian distribution, where $P_t(x)$ is the density function at time t.

For a training step we draw samples $x_{t_i} \sim \mathcal{X}_{t_i}$ for $i \in \{1, \ldots, k\}$ and calculate the loss with a single backwards integration of the ODE. While there are a number of ways to computationally approximate these quantities, we use a parallel method to iteratively calculate the $\log P_{t_i}$ based on $\log P_{t_{i-1}}$. To make a backward pass through all timepoints we start at the final timepoint, integrate the batch to the second to last timepoint, concatenate these points to the samples from the second to last timepoint, and continue till t_0, where the density is known for each sample. We note that this can compound the error especially for later timepoints if k is large or if the learned system is stiff, but gives significant speedup during training. To sample from P_{t_i} we first sample $\hat{x}_{t_0} \sim P_{t_0}$ then use the adjoint method to perform the integration $\hat{x}_{t_i} = \hat{x}_{t_0} + \int_{t_0}^{t_i} f_\theta(x(t), t)dt$. This is similar to other continuous normalizing flows where the model is trained backwards in time, but we use the invertible dynamics of the ODE to sample.

Pre-training of the Growth Network To pretrain the growth network, we use a simple and computationally efficient method that adapts discrete static unbalanced optimal transport to our framework in the continuous setting. We train a network $g(x, t) : \mathbb{R}^d \times [0, 1] \to \mathbb{R}^+$, which takes as input a cell state and time pair and produces a proliferation rate of a cell at that time. This is trained to match the result from discrete optimal transport. We note that adding proliferation rate regularization in this way does not guarantee conservation of mass. We could normalize $M(x)$ to be a probability distribution during training, e.g., as $\rho(x) = M(x)/\sum_{x \in \mathbb{R}^d} M(x)$. However, this now requires an integration over $\mathbb{R}^d$, which

is too computationally costly. Instead, we use the equivalence of the maximum likelihood formulation over a fixed proliferation function g and normalize it after the network is trained.

Comparison of TrajectoryNet Algorithm

Next we benchmark our TrajectoryNet algorithm on both simulated and real single-cell data. We compare TrajectoryNet to RNA Velocity [4, 49] and Diffusion Pseudotime [5].

Existing methods in cellular trajectory learning attempt to infer a trajectory within one timepoint [4, 5, 50], or interpolate linearly between two timepoints [3, 40], but TrajectoryNet can interpolate non-linearly using information from more than two timepoints. TrajectoryNet has advantages over existing methods in that it:

1. can interpolate by following the manifold of observed entities between measured timepoints, thereby solving the static-snapshot problem,
2. can create continuous-time trajectories of individual entities, giving researchers the ability to follow an entity in time.

Comparing TrajectoryNet to other trajectory inference tools on synthetic data

To test the ability of TrajectoryNet to follow a manifold we generated two toy datasets, an arch dataset representing a single trajectory and a branching tree dataset representing a diverging structure Figure S1(a). In both cases the datasets are designed to represent low-dimensional (1D) curved manifolds in a higher dimensional ambient space (2D). Both datasets are split into three timepoints X_0, $X_{1/2}$, and X_1 at times t_0, $t_{1/2}$ and t_1 respectively. t_0 and t_1 are supplied as training data and data at $t_{1/2}$ is held out as test data. In Figure S1(b) we compared the EMD between the predicted data at $t_{1/2}$, $\hat{X}_{1/2}$ with the ground truth data $X_{1/2}$. Here we compare the TrajectoryNet approach to six baseline methods. The prev baseline, which predicts the previous timepoint as $t_{1/2}$ i.e. $\hat{X}_{1/2} = X_0$. The next baseline, which predicts the next timepoint as $t_{1/2}$ i.e. $\hat{X}_{1/2} = X_1$. The optimal transport interpolant (OT), which predicts the McCann interpolant of the exact optimal transport plan between X_0 and X_1 as used in [3]. The random transport interpolant, which predicts the McCann interpolant of a random transport plan between X_0 and X_1. An RNA-velocity method $\mathrm{RNA_v}$, which takes the ground truth instantaneous velocity dX_0 of data at X_0, finds the optimal t^* which minimizes $\min_t \mathrm{EMD}(X_1, X_0 + t dX_0)$ then predicts $\hat{X}_{1/2} = X_0 + \frac{t}{2} dX_0$. Finally, we compare TrajectoryNet to a standard continuous normalizing flow model [37] CNF with energy regularization but without additional manifold regularization. We find TrajectoryNet performs the best in terms $\mathrm{EMD}(X_{1/2}, \hat{X}_{1/2})$ followed by the random and optimal transport McCann interpolations. This is because TrajectoryNet is able to follow the manifold while other methods are not.

Next, we evaluate the learned trajectories and vectorfields qualitatively to check the ability of various methods to follow the cell manifold. As seen in Figure S1(c) the paths for TrajectoryNet are able to follow the curve of the manifold as compared to an energy regularized CNF, which follows the straight line paths i.e. the dynamic optimal transport paths for a Euclidean geometry. While these methods define an interpretable vector field everywhere, our other baselines do not. In Figure S1(c) we visualize the imputed distribution of standard OT in red, a random interpolation in purple, and RNA-velocity in brown.

Comparing TrajectoryNet to RNA-velocity and Pseudotime on real scRNAseq data

TrajectoryNet along with other optimal transport methods contrast in a few notable ways from RNA-velocity or Pseudotime-based approaches to inferring cell time. In Figure S2 we show the differences in

these approaches on our tumorsphere dataset. TrajectoryNet learns a distribution flow $\rho(x, t)$ where the marginal $p(x, t = t_i)$ approximates the observed distribution at time t_i. This contrasts to a Pseudotime approach which assigns a Pseudotime (or developmental time to each cell. While there are a large variety of Pseudotime inference algorithms [50], many of them do not take the observed cell time into account. Here we compare against the popular Diffusion Pseudotime algorithm [5]. In Figure S2A we qualitatively evaluate the inferred pseudotime of different algorithms by coloring the PHATE embedding of our tumorsphere data by the inferred cell time (normalized to between lie between zero and one) for TrajectoryNet, scVelo, and Diffusion Pseudotime. Interestingly, scVelo reverses the ordering of the cells, assigning cells at day 0 a later Pseudotime than all other timepoints on average. We suspect this is because the days here are relatively far apart, and RNA-velocity analysis can only predict cell trajectories over shorter time scales. We can also see that Diffusion Pseudotime misorders the timepoints. While cells for day 0, 2, 12, and 18 are correctly ordered, cells from day 30 are assigned pseudotime values that are on average earlier than those in days 12 and 18. TrajectoryNet in contrast is able to infer trajectories over longer time scale by taking observed cell time into account and (optionally) RNA-velocity estimates at each cell.

We next evaluate this quantitatively in Figure S2B. Here we plot the distribution of inferred psuedotime for each day for different methods. TrajectoryNet orders the days correctly while RNAVelocity and Diffusion Pseudotime methods do not order the days monotonically.

In this dataset we did not incorporate RNA velocity data into the TrajectoryNet loss because we found it unreliable. In Figure S2C we show the scVelo inferred velocity flow plots separately on each day on three different 2D projections: A principal component analysis (PCA) projection, a UMAP [51] projection, and a PHATE [48] projection. We color cells by cell phase for reference. We find that depending on the projection, the inferred order of cells is different. This suggests RNA-velocity type analysis may not be reliable in this dataset.

Comparing TrajectoryNet Pipeline and Total Granger Causal Score inference against static gene network inference methods on synthetic data

Our gene network inference platform first infers the trajectories of individual cells via TrajectoryNet then performs total granger causality score (TGCS) analysis on an average trajectory for each endpoint. Our belief is that accurately inferring the dynamics information with TrajectoryNet will more accurately predict gene-gene interactions than approaches that do not include this information. To establish that this approach can recover the ground truth structure on simulated single-cell data with a known regulatory structure, we compare TGCS analysis to three other static methods for inferring regulatory structure from single-cell data, DREMI [7], Pearson correlation, and Spearman correlation across genes. We use the BoolODE package [52] to simulate multiple dataset structures including linear, cyclic, and bifurcating structures (Figure S3B). We test how well the predicted regulatory structures match the true regulatory structure of the simulation across experiments using area under the receiver operator characteristic curve (AUCROC). For these experiments we exclude evaluation of the diagonal, assuming that there are self loops in a time varying system. Furthermore, we evaluate whether there is an edge, and do not differentiate between positive and negative regulatory relationships. Since we are curious about the performance of our TGCS here, we use the ground truth simulated trajectories and ignore the post-processing steps that add additional noise to the data. Across unique runs of each algorithm on 100 randomly generated data of each type, we find that our TGCS analysis consistently infers gene-gene regulatory relationships more accurately than other approaches (Figure S3C-D).

Other Computational Methods Details

Single-cell RNA sequencing and pre-processing

scRNA-seq data from 5 samples, were processed with 10X and CellRanger pipeline according to the following steps. Sample demultiplexing and read alignment to the NCBI reference GRCh38 was completed to map reads using CellRanger. We prefiltered using parameters in scprep (v1.0.3, https://github.com/KrishnaswamyLab/scprep). Cells that contained at least 1,000 unique transcripts were kept for further analysis to generate a cell by gene matrix containing 17,983 cells and 16,983 genes. Normalization was performed using default parameters with L1 normalization, adjusting total library side of each cell to 10,000. Any cell expressing mitochondrial genes greater than 10% of their overall transcriptome were removed. Raw data files for scRNA-seq data will be available for download through GEO under an accession number to be assigned with no restrictions on publication.

in vivo scRNA-seq sequencing and pre-processing

scRNA-seq data from four replicates for the primary tumor and lung metastasis were aligned to GRCh38 and read mapping was completed with CellRanger. We retained cells that mapped strongly to the human genome and expressed at least 1000 genes. We then retained genes expressed in at least 20 cells and filtered the data to remove cells with fewer than 2000 counts and more than 5000 total counts. Finally, we library size normalized and square-root transformed the data. For each tissue, we used an MNN kernel to build a cellular graph with batch correction between the replicates. We then used MAGIC [53] to transform the MNN graph into the data space and ran TruncatedSVD to reduce the dimensionality to 100.

For calculating the epithelial-to-mesenchymal score, we adapted the approach from [11] to calculate gene correlation with the mean expression of epithelial markers EPCAM, CDH3, CTNNB1, KRT8, KRT18, as CDH1 was filtered out in above quality control steps.

in vivo spatial sequencing and pre-processing

Spatial transcriptomics was conducted using 10X Visium Spatial Gene Expression Slide and Reagent Kit, 16 rxns (PN-1000184), according to the protocol detailed in document CG000239RevD for the TNBCs and CG000239RevE for the Xenografts, available in 10X Genomics demonstrated protocols.

Next, we leveraged scMMGAN [29] to generate single-cell expression values for each spatial voxel in the same data space as the single-cell data. With scMMGAN, we use a generator consisting of three internal layers of 128, 256, and 512 neurons with batch norm and leaky rectified linear unit activations after each layer, and a discriminator consisting of three internal layers with 1,024, 512, and 256 neurons with the same batch norm and activations except with minibatching after the first layer. We use the geometry-preserving correspondence loss with a coefficient of 10, cycle-loss coefficient of 1, learning rate of 0.0001, and batch size of 256.

Transcriptional network generation

Transcriptional networks were built to visualise direct and indirect regulatory interactions of the EMT and MET trajectory using core transcription factors (transcription factors) identified by TrajectoryNet. Briefly, TRUSST v2 (Transcriptional Regulatory Relationships Unraveled by Sentence-based Text mining (https://www.grnpedia.org/trrust/)) database was used as the canvas for all known regulatory relationships across the human genome. Transcriptional interactions were filtered for the 23-core MET and 37-core EMT transcription factors derived for Gene Clusters 1-5. The resultant network (Fig 4H)

was organised to enable the visualisation of (i) direct interactions between the core MET and EMT transcription factors, (ii) common nodes that share regulatory relationships with the core MET and EMT transcription factors as well as (iii) unique nodes that shared direct regulatory relationships with either the core MET or core EMT transcription factors.

To develop a detailed transcriptional network describing the MET, the TRUSST v2 regulatory canvas was used to filter direct regulatory relationships of the 23 core MET transcription factors. Nine of the 23 transcription factors had known regulatory relationships on TRUSST v2, which allowed us to map the genes regulated by these core nodes. Next, the Epithelial Mesenchymal Gene Database, dbEMT 2.0 (http://www.dbemt.bioinfo-minzhao.org/) was used to distinguish known EMT genes (diamond) from novel EMT genes (ellipse) in the network. The resultant network was then organised based on the temporal gene clusters (2-5) and their regulatory relationships to observe direct and indirect crosstalk across the clusters forming the MET network (Fig 5C). All networks were built using Cytoscape 3.9.0 (https://cytoscape.org/).

Enrichment analysis

Gene lists for the MET and EMT sub-networks (Fig S8A, E) were extracted using Cytoscape 3.9.0. These gene lists were subjected to analysis on Metascape (https://metascape.org) to identify enriched biochemical pathways [54].

Software Availability

The TrajectoryNet package, as implemented in python, is available for download with a guided tutorial on the Krishnaswamy Lab Github page: https://github.com/KrishnaswamyLab/Cell-Dynamics-Pipeline.

Biological Methods Details

Cell Culture

HCC38 cells were purchased from ATCC and cultured in RPMI containing 10% (v/v) FBS (PS; 5.000 units penicillin and 5 mg streptomycin/ml in H_2O, Sigma Aldrich, cat no. P4333). Cells were routinely tested to confirm the absence of mycoplasma contamination. All cell line-specific media were supplemented with 1% (v/v) penicillin-streptomycin.

Tumorsphere assay

Single-cell suspensions were plated in ultra-low attachment 96-well plates (Corning # CLS3474, New York, USA) at low densities optimized to ensure tumorspheres arose from single anchor-independent cells. HCC38 $CD44^{hi}$ cells were seeded in 100 μl at 100 cells/well. Cell-line specific serum-free media was supplemented with 1% (v/v) penicillin/streptomycin, 20 ng/ml EGF, 20 ng/ml FGFb, 4 μg/mL heparin, 1x B27, and 1% (v/v) methylcellulose (Sigma-Aldrich). Fresh media was topped up every 5 days by adding 50μl per well of the appropriate tumorsphere media. Tumorspheres were counted at day 30 under 4x magnification and averaged 10 tumorspheres/well

For 3D immunostaining, tumorspheres were fixed with 10% formalin (Australian Biostain Pty Ltd) for 1 hour at RT in a gently rocking rotator and washed in TBS (3 x 15 min). Tumorspheres were then permeabilized with 100% methanol for 10 minutes at 4℃ and washed in TBS (3 x 15 min). Tumorspheres were then blocked in TBS 5% BSA, 10% Horse Serum and 0.1% Triton O/N at 4 degrees with rotation. Following incubation with primary antibodies ESRRA (Cell Signalling Technology, E1G1J, 1:200, Cat no.

13826); ZEB1 (Santa Cruz, #H-102, 1:200); CDH1 (BD Biosciences/(36/E-Cadherin, Cat. no. 610181, 1:200) diluted in blocking buffer O/N at 4℃ in rotation, spheres were washed (4 x 30 min) in TBS and stained with the appropriate fluorophore-conjugated secondary antibodies (1:500) (Ms Cy3 (#M30010), Rb 647 (#A32722), Ms 488 (#A11001), Rb 488 (#711546152), ThermoFisher Scientific) and DAPI O/N at 4℃. Tumorspheres were washed with TBS (4 x 30 min) then resuspended in 20μl of mounting media (ProLong Diamond Antifade Mounting Media (ThermoFisher Scientific)) and mounted between a glass slide and a coverslip spaced by tape.

Labelled tumorspheres were imaged using confocal microscopy (Leica DMI 6000 SP8 with 40x (NA 1.3) or 63x (NA 1.4) oil objectives or a Nikon A1R confocal with 20x Plan Apochromat air objective (NA 0.75) at 2x zoom using an HD25 resonance scanner) using identical acquisition settings (optimised per protein marker) for all time points. Quantitative image analysis was performed using CellProfiler (v4.2.1, [55]) to segment individual cells via maximum cross-entropy-based threshold detection of nuclei (DAPI) and cell bodies (sum of all channels). Mean intensity per cell was measured for each marker, with per cell image and quantitative data integrated via a Knime software (v4.6.4, [56]) for visualisation [57,58].

FACS analysis

Bulk HCC38 cell lines were cultured in 2D tissue culture dishes. For isolating CD44^{hi} cells from this bulk population, cells were trypsinized and stained with CD44 antibody (BD Biosciences anti-human CD44-PE-cy7 (1:800)) for 25 min at 4C. CD44^{hi} cells were sorted on BD Aria III. Sorted cells were cultured in media supplemented with 0.1% (v/v) gentamicin and 1% (v/v) antibiotic-antimycotic for at least two passages to avoid contamination. Multiple rounds of FACS enrichment were performed on these expanded cultures until pure CD44^{hi} populations were isolated. To identify EPCAM$^{+/-}$, CAV1$^{+/-}$, EPCAM/CAV1$^{+/+}$ populations, HCC38 CD444^{hi} were subjected to FACS sorting using Anti-CD326 (EPCAM) (Invitrogen #53-8326-42) and Anti-Caveolin 1 (BD Biosciences). Data acquisition was performed using BD Aria III and FACSDiva software (BD Biosciences and data analysis was performed using Flowjo X10.7.1.

ESSRA knockdown and validation

Preliminary validation of the MET network included an in-vitro knockdown of ESRRA in HCC38 CD44^{hi} cells (n=3 biological replicates). Predesigned siRNA specific to ESRRA and scrambled siRNA were purchased from Integrated DNA Technologies, USA (TriFECTa® RNAi Kit, Design ID hs.Ri.ESRRA.13). HCC38 CD44^{hi} cells were seeded in 24 well plates at a density of 9000 cells/well and ESRRA knockdown was performed using 10nM of pooled siRNA (hs.Ri.ESRRA.13.1, hs.Ri.ESRRA.13.2 and hs.Ri.ESRRA.13.3) using Lipofectamine RNAiMax (ThermoFisher Scientific, USA) as per the manufacturer's protocol. Cells were harvested for protein extraction 48hrs post siRNA transfection to study the downstream effect on CDH1 expression upon ESRRA knockdown using western blot.

In parallel, HCC38 CD44^{hi} cells were also treated with an ESRRA inhibitor ((2-Aminophenyl)(1-(3-isopropylphenyl)-1H-1,2,3-triazol-4-yl)methanone) (BLD Pharm, China, Cat no. BD01201330) [referred to as Compound 14 (C14) here on] at 5 and 10μM concentrations. Vehicle controls were treated with DMSO. Cells were harvested 48 hours post-treatment and processed for protein extraction which were studied for changes in ESRRA and CDH1 expression using western blot (n=3 biological replicates).

Western Blotting

Proteins were extracted from control and treated cells using ice cold modified RIPA buffer (50 mM Tris-HCl pH 7.5, 150 mM NaCl, 1 mM EDTA, 1 mM EGTA, 1% Triton X-100, 0.1% SDS with supplemented with 1X protease and phosphatase inhibitor cocktails). The lysates were sonicated using a QSONICA

Q55 probe sonicator at 50 kHz for 20 seconds in an ice bath. This whole cell lysate was centrifuged at 14,000 x g for 10min at 4℃ and stored in -80℃ until further use. Proteins were quantified using the Pierce™ BCA Protein Assay Kit (Cat. 23227, Thermo Fisher Scientific, USA) as per the manufacturer's instructions. 20μg of whole cell lysates were separated on 1D SDS-PAGE using NuPAGE™ 4 to 12%, Bis-Tris (Cat. NP0322BOX, ThermoFisher Scientific, USA). Proteins separated on the gel were transferred onto nitrocellulose membrane using Trans-Turbo Transfer system (BioRad Laboratories, USA) at 1.3V for 7min. The membrane was blocked using 5% milk in TBS 1hr. The membrane was next incubated overnight at 4℃ with primary antibodies against the protein of interest [1: 1000 ESRRA (Cell Signalling Technologies, USA, Cat. no. 13826), 1:1000 CDH1 (Cell Signalling Technologies, USA, Cat. no. 3195S), 1:5000 GAPDH (Cell Signalling Technologies, USA, Cat. no. 97166S), 1:5000 β-Actin (Cell Signalling Technologies, USA, Cat. no. 3700S)]. After incubation with the primary antibody, membranes were washed with 1X TBST for 3 x 10min on a rocker at room temperature. Next, membranes were incubated with appropriate HRP-conjugated secondary antibodies. Washing steps were repeated, and the membrane was developed using ECL substrate (Western Lightning™ Ultra, Perkin Elmer or Clarity Western ECL Substrate) and scanned using Fusion FX Vilber Lourmat scanner. Each Western blot experiment was performed using three biological replicates, to calculate the statistical significance (p-value) of relative fold-change in expression. GAPDH or β-Actin was used to normalize the relative fold-change expression value of ESRRA and CDH1. The signal intensity of the bands in western blots was quantified using Image Studio Lite version 5.2 (LI-COR Biotechnology, USA).

References

[1] 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* (2020).

[2] Chaffer, C. L. *et al.* Normal and neoplastic nonstem cells can spontaneously convert to a stem-like state. *Proceedings of the National Academy of Sciences* **108**, 7950–7955 (2011). URL https://doi.org/10.1073/pnas.1102454108.

[3] Schiebinger, G. *et al.* Optimal-Transport Analysis of Single-Cell Gene Expression Identifies Developmental Trajectories in Reprogramming. *Cell* **176**, 928–943.e22 (2019).

[4] La Manno, G. *et al.* RNA velocity of single cells. *Nature* **560**, 494–498 (2018).

[5] Haghverdi, L., Büttner, M., Wolf, F. A., Buettner, F. & Theis, F. J. Diffusion pseudotime robustly reconstructs lineage branching. *Nature Methods* **13**, 845–848 (2016).

[6] Krishnaswamy, S. *et al.* Conditional density-based analysis of t cell signaling in single-cell data. *Science* **346**, 1250689–1250689 (2014).

[7] Krishnaswamy, S. *et al.* Conditional density-based analysis of t cell signaling in single-cell data. *Science* **346**, 1250689 (2014).

[8] Han, H. *et al.* TRRUST v2: an expanded reference database of human and mouse transcriptional regulatory interactions. *Nucleic Acids Research* **46**, D380–D386 (2017). URL https://doi.org/10.1093/nar/gkx1013.

[9] Al-Hajj, M., Wicha, M. S., Benito-Hernandez, A., Morrison, S. J. & Clarke, M. F. Prospective identification of tumorigenic breast cancer cells. *Proceedings of the National Academy of Sciences* **100**, 3983–3988 (2003). URL https://doi.org/10.1073/pnas.0530291100.

[10] Moon, K. R. *et al.* Visualizing structure and transitions in high-dimensional biological data. *Nature Biotechnology* **37**, 1482–1492 (2019).

[11] Chakraborty, P., George, J. T., Tripathi, S., Levine, H. & Jolly, M. K. Comparative study of transcriptomics-based scoring metrics for the epithelial-hybrid-mesenchymal spectrum. *Frontiers in Bioengineering and Biotechnology* **8** (2020). URL https://doi.org/10.3389/fbioe.2020.00220.

[12] Yang, J. *et al.* Guidelines and definitions for research on epithelial–mesenchymal transition. *Nature Reviews Molecular Cell Biology* **21**, 341–352 (2020).

[13] Tirosh, I. *et al.* Dissecting the multicellular ecosystem of metastatic melanoma by single-cell RNA-seq. *Science* **352**, 189–196 (2016). URL https://doi.org/10.1126/science.aad0501.

[14] Bierie, B. *et al.* Integrin-β4 identifies cancer stem cell-enriched populations of partially mesenchymal carcinoma cells. *Proceedings of the National Academy of Sciences* **114**, E2337–E2346 (2017). URL https://doi.org/10.1073/pnas.1618298114.

[15] Al-Hajj, M., Wicha, M. S., Benito-Hernandez, A., Morrison, S. J. & Clarke, M. F. Prospective identification of tumorigenic breast cancer cells. *Proc. Natl. Acad. Sci. U. S. A.* **100**, 3983–3988 (2003).

[16] Chaffer, C. L. *et al.* Poised chromatin at the zeb1 promoter enables breast cancer cell plasticity and enhances tumorigenicity. *Cell* **154**, 61–74 (2013).

[17] Sakaue-Sawano, A. *et al.* Visualizing spatiotemporal dynamics of multicellular cell-cycle progression. *Cell* **132**, 487–498 (2008).

[18] Blondel, V. D., Guillaume, J.-L., Lambiotte, R. & Lefebvre, E. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* **2008**, P10008 (2008).

[19] Zhao, M., Liu, Y., Zheng, C. & Qu, H. dbEMT 2.0: An updated database for epithelial-mesenchymal transition genes with experimentally verified information and precalculated regulation information for cancer metastasis. *Journal of Genetics and Genomics* **46**, 595–597 (2019). URL https://doi.org/10.1016/j.jgg.2019.11.010.

[20] Castaño, Z. *et al.* Il-1β inflammatory response driven by primary breast cancer prevents metastasis-initiating cell colonization. *Nature cell biology* **20**, 1084 (2018).

[21] Shannon, P. *et al.* Cytoscape: A software environment for integrated models of biomolecular interaction networks. *Genome Research* **13**, 2498–2504 (2003). URL https://doi.org/10.1101/gr.1239303.

[22] Berman, A. Y. *et al.* ERRα regulates the growth of triple-negative breast cancer cells via s6k1-dependent mechanism. *Signal Transduction and Targeted Therapy* **2** (2017). URL https://doi.org/10.1038/sigtrans.2017.35.

[23] Ma, J.-H. *et al.* STAT3 targets ERR-α to promote epithelial-mesenchymal transition, migration, and invasion in triple-negative breast cancer cells. *Mol. Cancer Res.* **17**, 2184–2195 (2019).

[24] De Luca, A. *et al.* Mitochondrial biogenesis is required for the anchorage-independent survival and propagation of stem-like cancer cells. *Oncotarget* **6**, 14777–14795 (2015).

[25] Wu, Y.-M. *et al.* Inhibition of ERRα suppresses epithelial mesenchymal transition of triple negative breast cancer cells by directly targeting fibronectin. *Oncotarget* **6**, 25588–25601 (2015).

[26] Huang, J.-W. *et al.* Effects of estrogen-related receptor alpha (ERRα) on proliferation and metastasis of human lung cancer A549 cells. *J. Huazhong Univ. Sci. Technolog. Med. Sci.* **34**, 875–881 (2014).

[27] Chen, S., Ye, J., Kijima, I., Kinoshita, Y. & Zhou, D. Positive and negative transcriptional regulation of aromatase expression in human breast cancer tissue. *J. Steroid Biochem. Mol. Biol.* **95**, 17–23 (2005).

[28] San Juan, B. P. *et al.* Targeting phenotypic plasticity prevents metastasis and the development of chemotherapy-resistant disease. *medRxiv* (2022). URL https://www.medrxiv.org/content/early/2022/03/21/2022.03.21.22269988. https://www.medrxiv.org/content/early/2022/03/21/2022.03.21.22269988.full.pdf.

[29] Amodio, M. *et al.* Single-cell multi-modal GAN reveals spatial patterns in single-cell data from triple-negative breast cancer. *Patterns (N Y)* **3**, 100577 (2022).

[30] Mulero-Navarro, S. & Fernandez-Salguero, P. M. New trends in aryl hydrocarbon receptor biology. *Frontiers in Cell and Developmental Biology* **4** (2016).

[31] Yoriki, K. *et al.* Estrogen-related receptor alpha induces epithelial-mesenchymal transition through cancer-stromal interactions in endometrial cancer. *Scientific Reports* **9** (2019). URL https://doi.org/10.1038/s41598-019-43261-z.

[32] Berman, A. Y. *et al.* ERRα regulates the growth of triple-negative breast cancer cells via S6K1-dependent mechanism. *Signal Transduct. Target. Ther.* **2**, 17035 (2017).

[33] Carey, L. A. *et al.* The triple negative paradox: primary tumor chemosensitivity of breast cancer subtypes. *Clin. Cancer Res.* **13**, 2329–2334 (2007).

[34] Bray, F. *et al.* Global cancer statistics 2018: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA Cancer J. Clin.* **68**, 394–424 (2018).

[35] Grathwohl, W., Chen, R. T. Q., Bettencourt, J., Sutskever, I. & Duvenaud, D. FFJORD: Free-form Continuous Dynamics for Scalable Reversible Generative Models. In *ICLR* (2019). 1810.01367.

[36] Benamou, J.-D. & Brenier, Y. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik* **84**, 375–393 (2000).

[37] Chen, R. T. Q., Rubanova, Y., Bettencourt, J. & Duvenaud, D. Neural Ordinary Differential Equations. In *Advances in Neural Information Processing Systems 31* (2018). 1806.07366.

[38] Rezende, D. J. & Mohamed, S. Variational Inference with Normalizing Flows. In *Proceedings of the 32nd International Conference on Machine Learning*, vol. 37, 1530–1538 (2015). 1505.05770.

[39] Rubanova, Y., Chen, R. T. Q. & Duvenaud, D. Latent ODEs for Irregularly-Sampled Time Series. *arXiv:1907.03907 [cs, stat]* (2019). 1907.03907.

[40] Yang, K. D. & Uhler, C. Scalable Unbalanced Optimal Transport Using Generative Adversarial Networks. In *7th International Conference on Learning Representations*, 20 (2019).

[41] Hashimoto, T. B., Gifford, D. K. & Jaakkola, T. S. Learning Population-Level Diffusions with Generative Recurrent Networks. In *Proceedings of the 33rd International Conference on Machine Learning*, 2417–2426 (2016).

[42] Papadakis, N., Peyré, G. & Oudet, E. Optimal Transport with Proximal Splitting. *SIAM Journal on Imaging Sciences* **7**, 212–238 (2014). 1304.5784.

[43] Liero, M., Mielke, A. & Savaré, G. Optimal Entropy-Transport problems and a new Hellinger–Kantorovich distance between positive measures. *Inventiones mathematicae* **211**, 969–1117 (2018).

[44] Chizat, L., Peyré, G., Schmitzer, B. & Vialard, F.-X. Unbalanced optimal transport: Dynamic and Kantorovich formulations. *Journal of Functional Analysis* **274**, 3090–3123 (2018).

[45] Finlay, C., Jacobsen, J.-H., Nurbekyan, L. & Oberman, A. M. How to train your neural ODE: The world of Jacobian and kinetic regularization. *ICML* (2020). 2002.02798.

[46] Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y. & Manzagol, P.-A. Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion. *Journal of Machine Learning Research* 3371–3408 (2010).

[47] Rifai, S., Vincent, P., Muller, X., Glorot, X. & Bengio, Y. Contractive Auto-Encoders: Explicit Invariance During Feature Extraction. In *Proceedings of the 29th International Conference on Machine Learning*, 833–840 (2011).

[48] Moon, K. R. *et al.* Visualizing structure and transitions in high-dimensional biological data. *Nature Biotechnology* **37**, 1482–1492 (2019).

[49] Bergen, V., Lange, M., Peidli, S., Wolf, F. A. & Theis, F. J. Generalizing RNA velocity to transient cell states through dynamical modeling. *BioRxiv 820936* (2019).

[50] Saelens, W., Cannoodt, R., Todorov, H. & Saeys, Y. A comparison of single-cell trajectory inference methods. *Nature Biotechnology* **37**, 547–554 (2019).

[51] McInnes, L., Healy, J. & Melville, J. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction (2020). 1802.03426.

[52] Pratapa, A., Jalihal, A. P., Law, J. N., Bharadwaj, A. & Murali, T. M. Benchmarking algorithms for gene regulatory network inference from single-cell transcriptomic data. *Nature Methods* **17**, 147–154 (2020).

[53] van Dijk, D. *et al.* Recovering gene interactions from single-cell data using data diffusion. *Cell* **174**, 716 – 729.e27 (2018).

[54] Zhou, Y. *et al.* Metascape provides a biologist-oriented resource for the analysis of systems-level datasets. *Nature Communications* **10** (2019). URL https://doi.org/10.1038/s41467-019-09234-6.

[55] Carpenter, A. E. *et al.* Cellprofiler: image analysis software for identifying and quantifying cell phenotypes. *Genome Biology* **7**, R100 (2006). URL https://doi.org/10.1186/gb-2006-7-10-r100.

[56] Berthold, M. R. *et al.* KNIME: The konstanz information miner. In *Data Analysis, Machine Learning and Applications*, 319–326 (Springer Berlin Heidelberg, 2008). URL https://doi.org/10.1007/978-3-540-78246-9_38.

[57] Bryce, N. S. *et al.* High-content imaging of unbiased chemical perturbations reveals that the phenotypic plasticity of the actin cytoskeleton is constrained. *Cell Systems* **9**, 496–507.e5 (2019). URL https://doi.org/10.1016/j.cels.2019.09.002.

[58] Lock, J. G. *et al.* Visual analytics of single cell microscopy data using a collaborative immersive environment. In *Proceedings of the 16th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and its Applications in Industry* (ACM, 2018). URL https://doi.org/10.1145/3284398.3284412.

[59] Granger, C. W. J. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* **37**, 424 (1969). URL https://doi.org/10.2307/1912791.

Supplemental Information

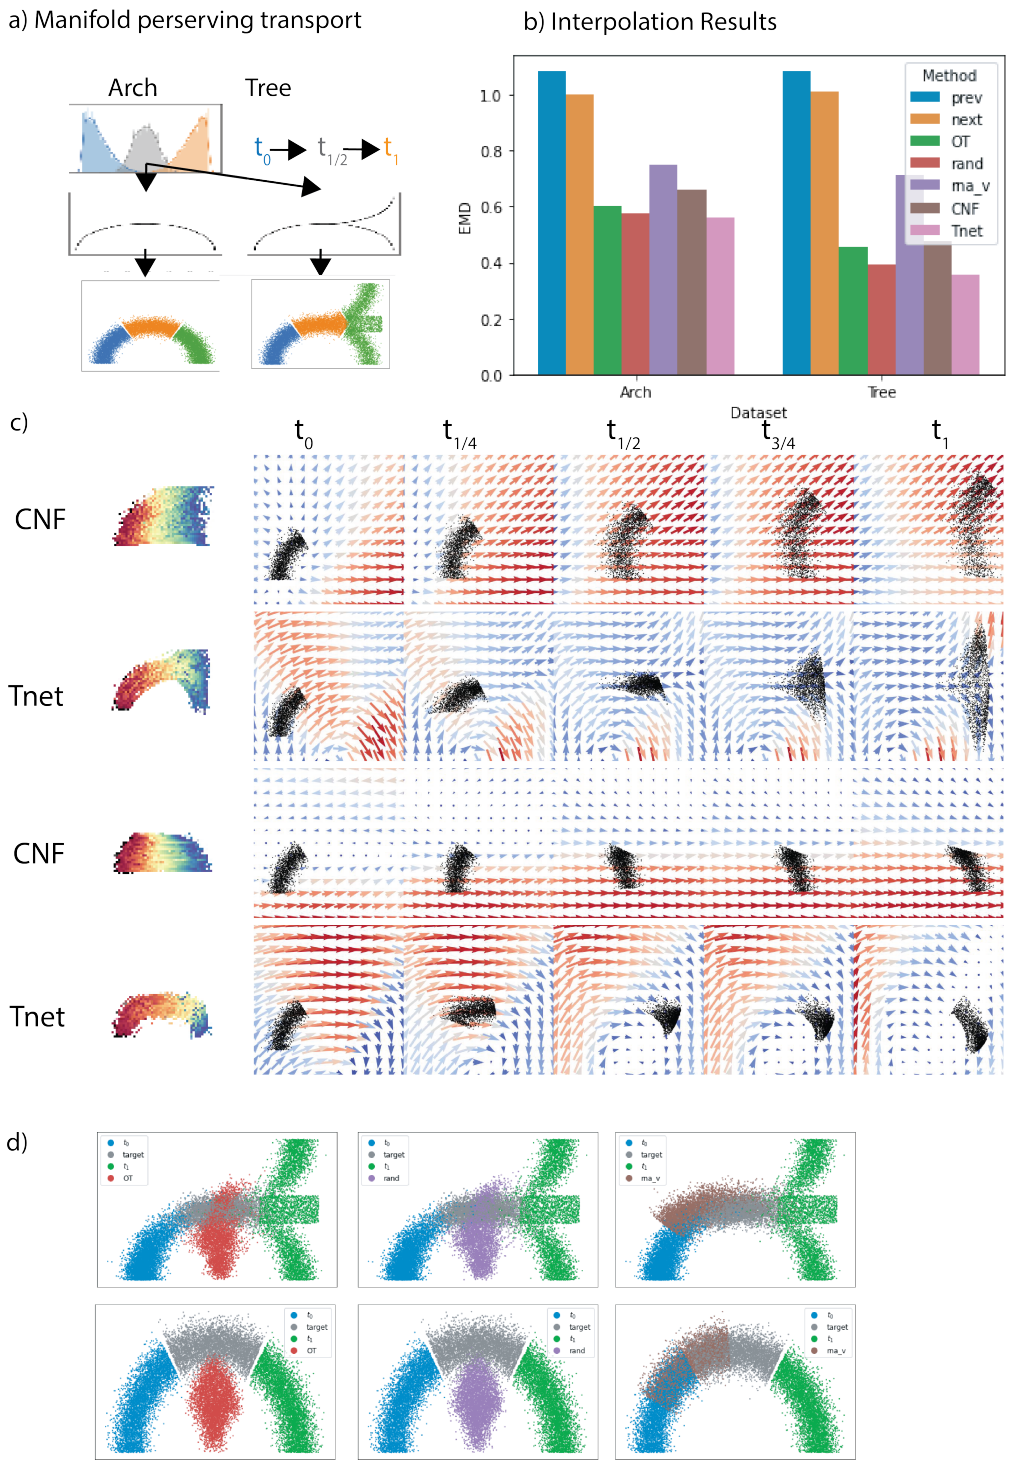

Supplementary Figure S1: TrajectoryNet comparisons on synthetic data of different structures. (A) Depiction of 1D manifolds with non-branching and branching. Models are trained on data from timepoints t_0 (blue) and t_1 (orange) with the goal of accurately interpolating ground truth data at $t_{1/2}$. **(B)** Comparing TrajectoryNet with other methods at predicting $t_{1/2}$. Earth Mover's Distance between the predicted data at $t_{1/2}$ and the ground truth distribution at $t_{1/2}$ shown across models, with lower distances indicating a more accurate prediction of $t_{1/2}$. **(C)** Visualizing individual paths projected forward from t_0 across methods. **(D)** Comparing interpolated and ground truth $t_{1/2}$ for the three other methods (OT, random, RNA velocity) that do not create paths to visualize.

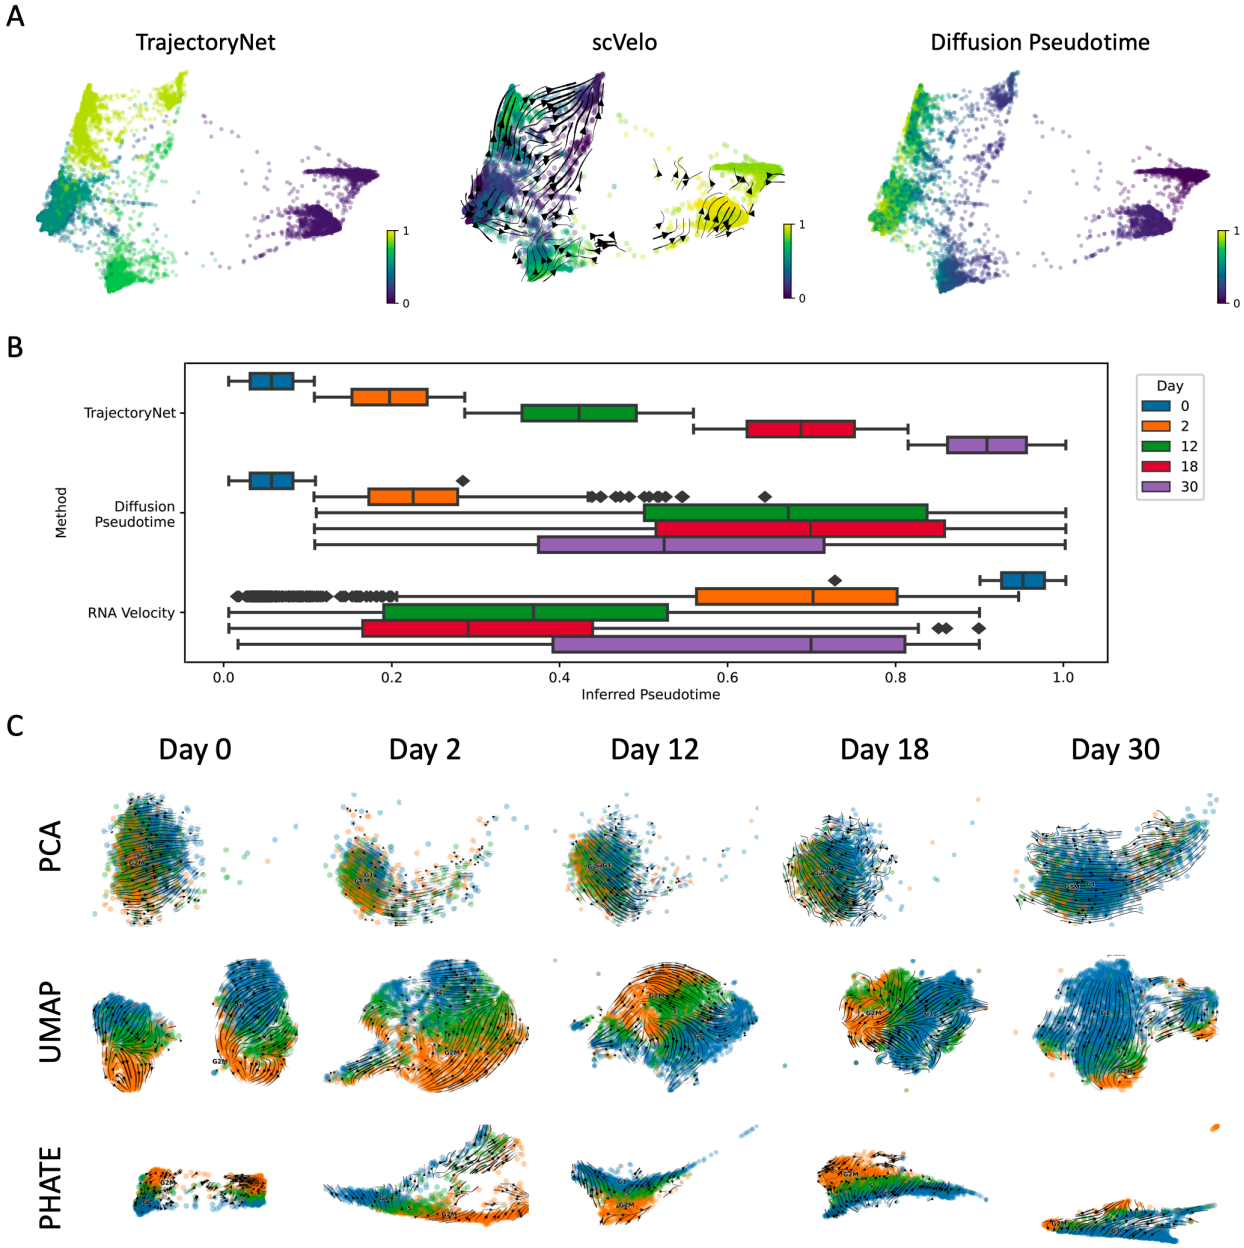

Supplementary Figure S2: Ordering of cells by TrajectoryNet, scVelo, and Diffusion Pseudotime. (A) Cells colored by (from left to right) TrajectoryNet, scVelo, and Diffusion inferred time between zero (purple) and one (yellow). scVelo represents RNA-Velocity based methods and Diffusion Pseudotime represents graph-based pseudotime inference methods. **(B)** Plots inferred cell time vs. ground truth observation time. We see that TrajectoryNet is the only method where the inferred time correlates with the ground truth observation time. **(C)** Visualization of scVelo inferred velocity streams across timepoints and embeddings. Each embedding is colored by the inferred cell state (S: Green), (G1: Blue) and (G2M: Orange). We see that the scVelo inferred velocity streams are inconsistent between embeddings.

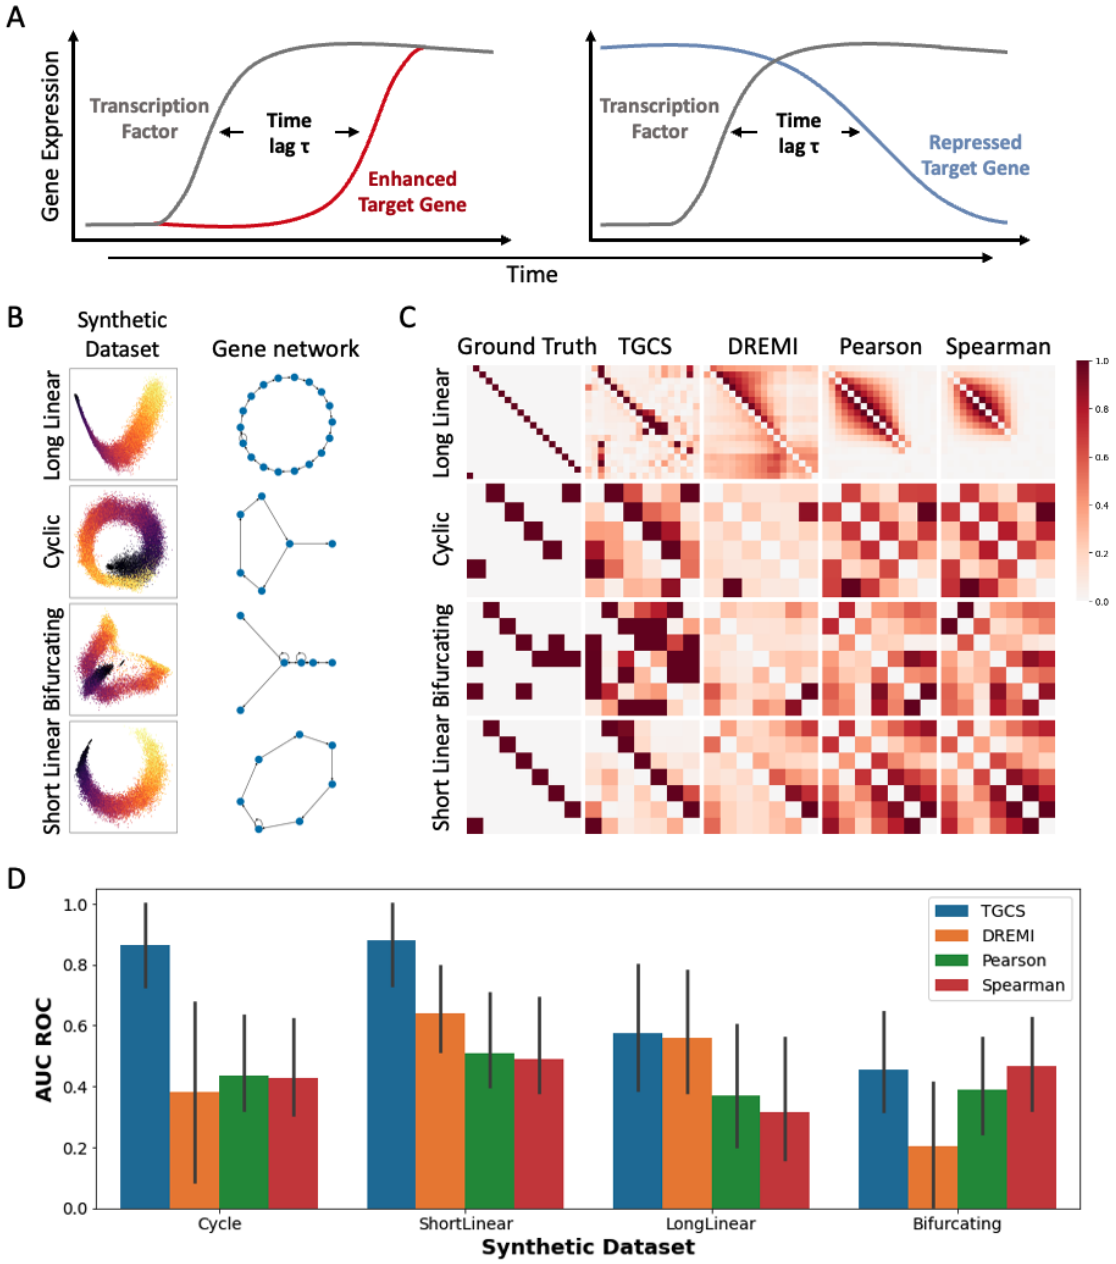

Supplementary Figure S3: Dynamic versus static inference of gene regulatory interactions. (A) Schematic overview of Granger analysis. Signed Granger values are computed using Granger causality inference [59] between a potentially regulatory transcription factor and a potentially regulated target gene based on τ time lag in regulatory effect. The Granger values are signed based on the direction of effect with an enhancing relationship having a $+$ sign and a repressive relationship having a $-$ sign. **(B)** Visualizations of synthetic datasets, in PC dimensions, and gene networks used to evaluate our method, total Granger causal score (TGCS) in part **(C)**. **(C)** Visual comparison of the performance of TGCS, against three methods of gene network inference, DREMI [7], Pearson correlation and Spearman correlation, on four synthetic datasets with different known ground truth regulatory interaction structures as specified in BoolODE [52]. **(D)** Numerical comparison of TGCS against DREMI, Pearson correlation and Spearman correlation at identifying ground truth gene network structure using area under the receiver operator characteristic curve (AUC ROC). Here higher scores indicate that a method is more accurately able to identify ground truth gene network structure. Results were averaged across 100 runs with one standard deviation error bars visualized.

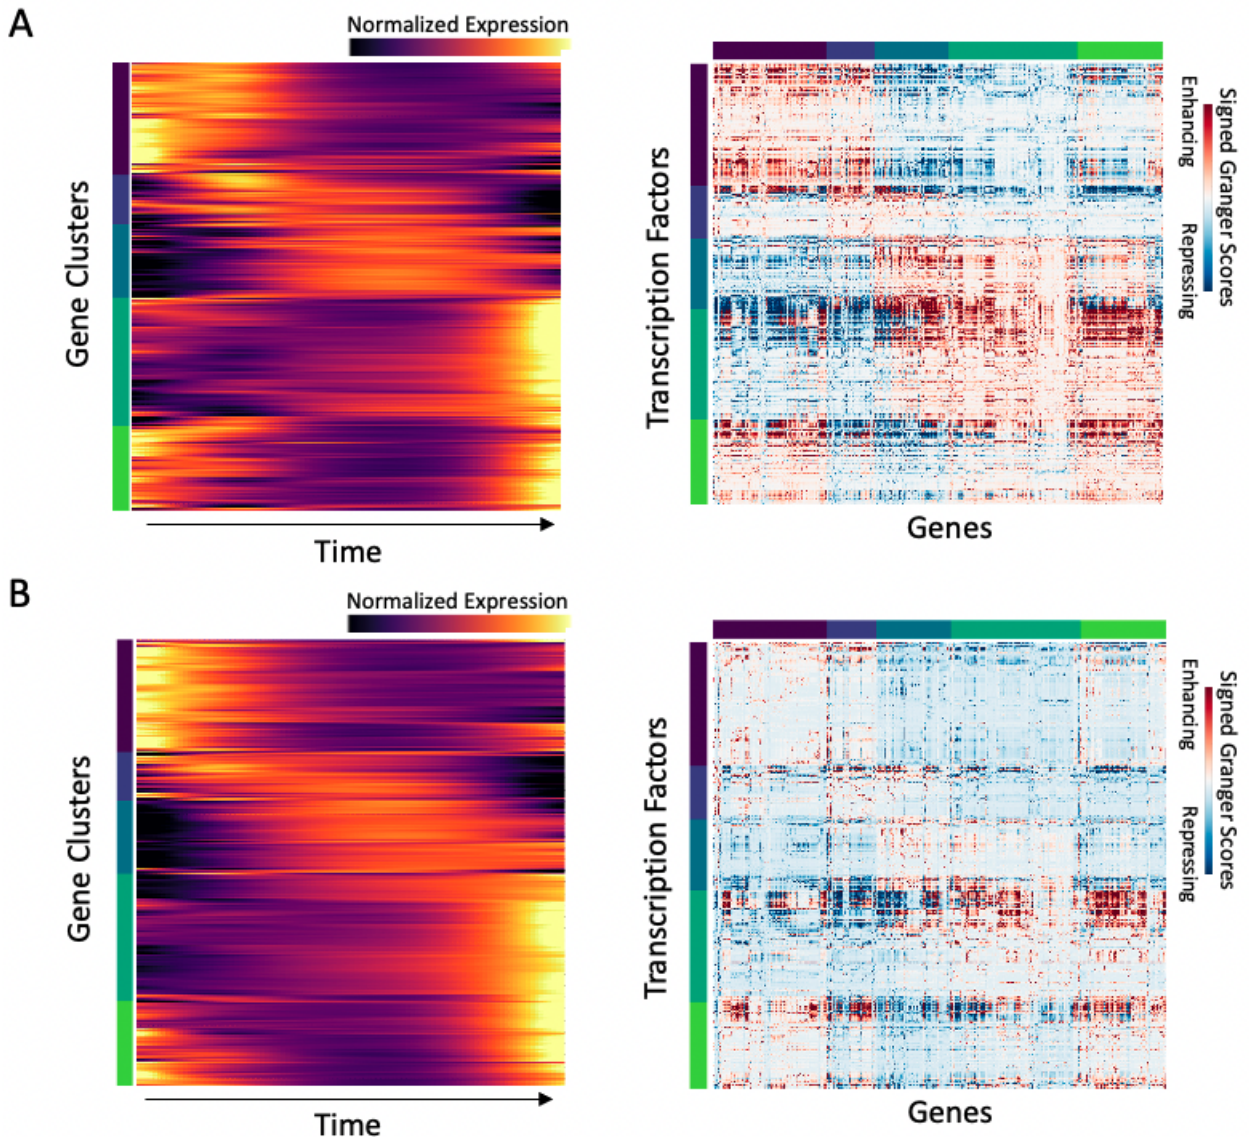

Supplementary Figure S4: Gene dynamics and gene network calculations for EMT and MET dynamics. **(A)** Recomputed gene expression dynamics based on trajectories that terminate in mesenchymal cellular cluster identified in Figure 4A (left). Signed Granger analysis between 5,273 genes and 461 transcription factors gene trends of cells that terminate in mesenchymal cluster (right). Red denotes a strong enhancing relationship between a transcription factor and a target gene, while blue denotes a strong repressive relationship. **(B)** Recomputed gene expression dynamics based on trajectories that terminate in epithelial cellular cluster identified in Figure 4A (left). Signed Granger analysis between 5,273 genes and 461 transcription factors gene trends of cells that terminate in epithelial cluster (right). Red denotes a strong enhancing relationship between a transcription factor and a target gene, while blue denotes a strong repressive relationship.

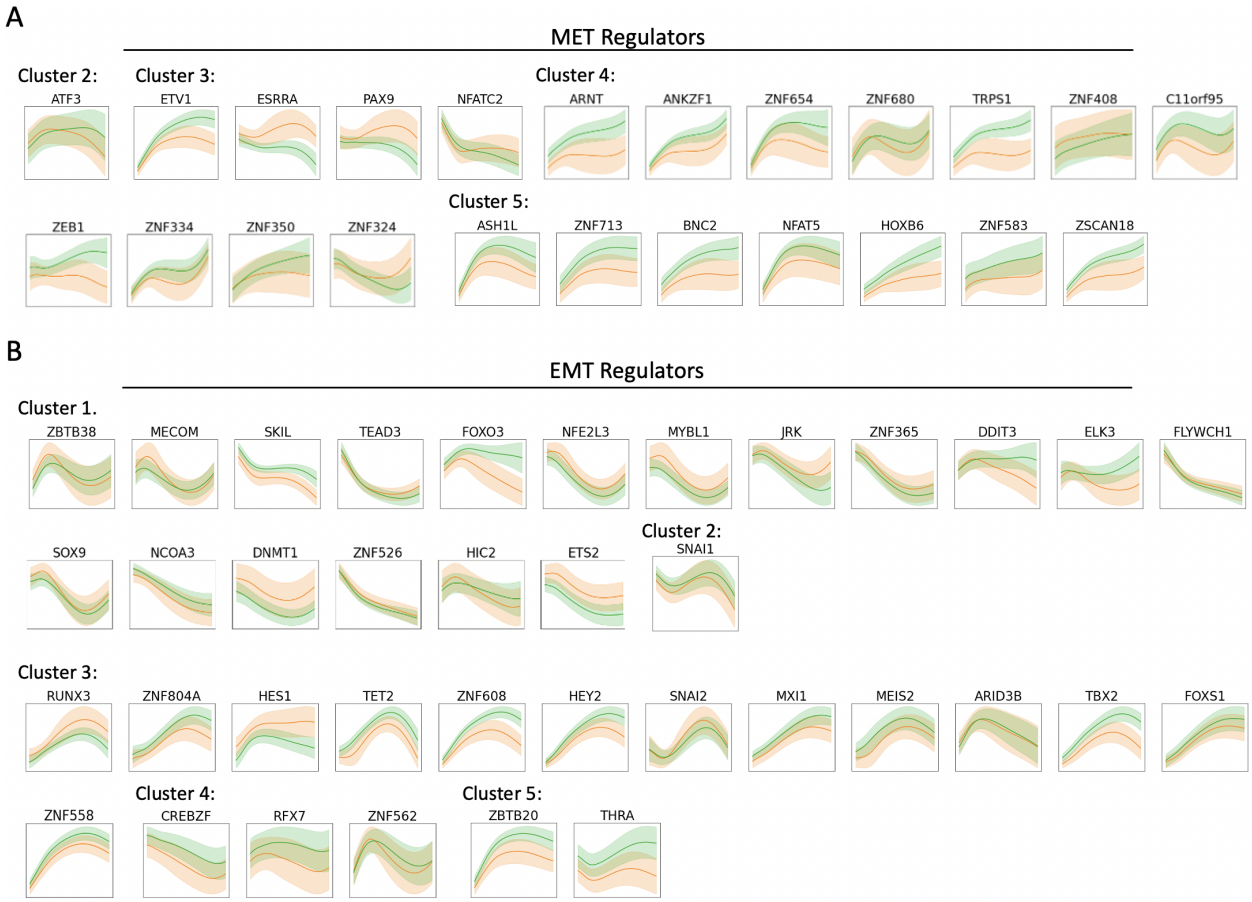

Supplementary Figure S5: Visualizing expression dynamics of key mesenchymal and epithelial regulatory genes. (A) Top MET-specific regulatory transcription factors separated by gene clusters: i) cluster 2, ii) cluster 3, iii) cluster 4, iv) cluster 5. **(B)** Top EMT-specific regulatory transcription factors separated by gene clusters: i) cluster 1, ii) cluster 2, iii) cluster 3, iv) cluster 4, v) cluster 5.

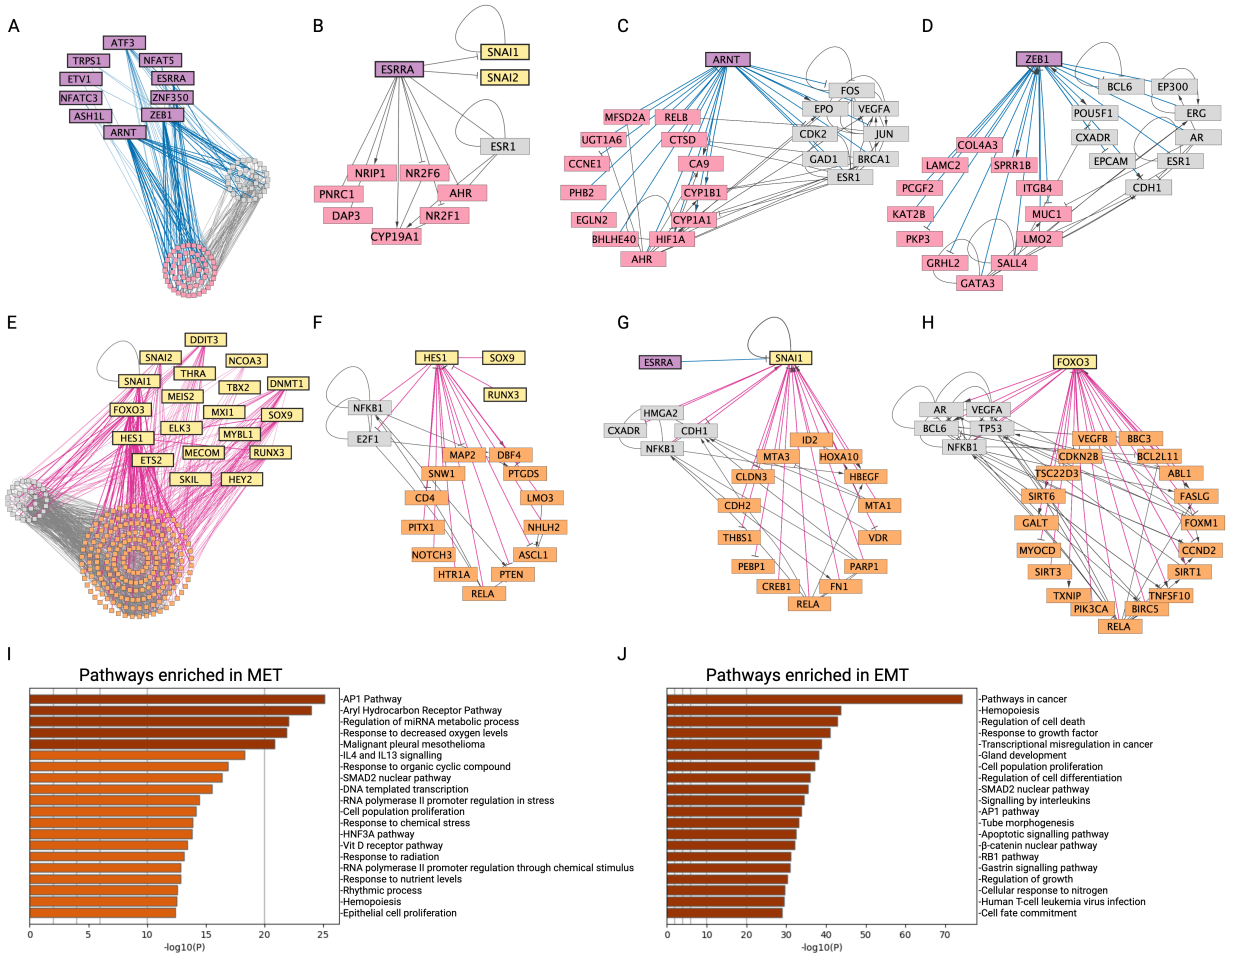

Supplementary Figure S6: Visualising the extended EMT and MET subnetwork along with the key pathways they regulate. (A) MET subnetwork comprising the core MET transcription factors (orange), MET specific interactors (yellow and interactors common with EMT subnetwork (grey). Example gene regulatory relationships of core MET transcription factors ESRRA **(B)**, ARNT **(C)**, ZEB1 **(D)**. **(E)** EMT subnetwork comprising the core EMT transcription factors (green), EMT specific interactors (purple) and interactors common with EMT core transcription factors (grey). Example gene regulatory relationships of core EMT transcription factors HES1 **(F)**, SNAI1 **(G)**, FOXO3 **(H)**. **(I)** Pathways enriched in the MET subnetwork. **(J)** Pathways enriched in the EMT subnetwork

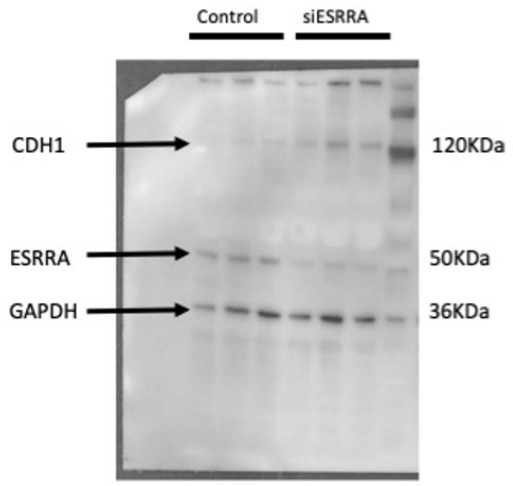

(a) Western-Blot Scan of CDH1 after siControl and siESRRA treatment

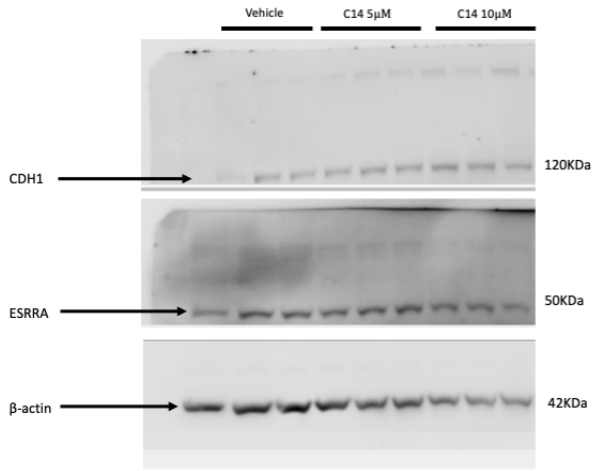

(b) Western-Blot Scan of CDH1 after control or C14 treatment

Supplementary Figure S7: Validating predicted CDH1 expression dependence on ESRRA using siRNA knock-down and C14 antagonist. (A) Western-blot scans of ESRRA with GAPDH as reference using siRNA knockdown of ESRRA. Quantification can be seen in (Figure 5F). **(B)** Western-blot scans of ESRRA with β-actin as reference using C14 antagonist of ESRRA. Quantification can be seen in (Figure 5G).

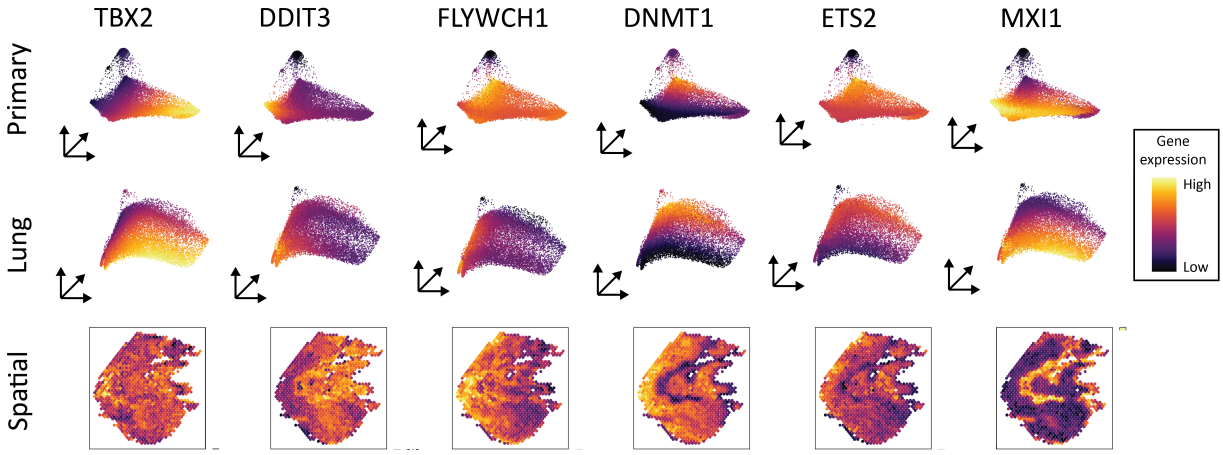

Supplementary Figure S8: Spatial expression and dynamics of EMT network genes within *in vivo* datasets. Visualization of EMT-related genes in scRNAseq data of primary tumor and lung metastasis, and their spatial localization in spatial transcriptomic (Visium 10x) data of a primary tumor.