A Bayesian framework for inferring dynamic intercellular interactions from time-series single-cell data

Cameron Y Park, Shouvik Mani, Nicolas Beltran-Velez, et al.

Genome Res. published online September 5, 2024
Access the most recent version at doi:10.1101/gr.279126.124

A Bayesian framework for inferring dynamic intercellular interactions from time-series single-cell data

Cameron Park[1,2,12]✉ Shouvik Mani[2,3,12] Nicolas Beltran-Velez[3]
Katie Maurer[4,5,6] Teddy Huang[4,9] Shuqiang Li[4,9] Satyen Gohil[4]
Kenneth J. Livak[4] David A. Knowles[3,7,8,10] Catherine J. Wu[4,5,6,]
Elham Azizi[1,2,3,11]✉

[1]Department of Biomedical Engineering, Columbia University, New York, NY 10027, USA
[2]Irving Institute for Cancer Dynamics, Columbia University, New York, NY 10027, USA
[3]Department of Computer Science, Columbia University, New York, NY 10027, USA
[4]Department of Medical Oncology, Dana-Farber Cancer Institute, Boston, MA, USA
[5]Harvard Medical School, Boston, MA, USA
[6]Broad Institute of MIT and Harvard, Cambridge, MA, USA
[7]Data Science Institute, Columbia University, New York, NY
[8]New York Genome Center, New York, NY
[9]Translational Immunogenomics Laboratory, Dana-Farber Cancer Institute, Boston MA USA
[10]Department of Systems Biology, Columbia University, New York, NY
[11]Data Science Institute, Columbia University, New York, NY
[12]These authors contributed equally
cyp2111@columbia.edu ; ea2690@columbia.edu

Running title: DIISCO resolves dynamic cell-cell interactions

Abstract

Characterizing cell-cell communication and tracking its variability over time are crucial for understanding the coordination of biological processes mediating normal development, disease progression, and responses to perturbations such as therapies. Existing tools fail to capture time-dependent intercellular interactions and primarily rely on existing databases compiled from limited contexts. We introduce DIISCO, a Bayesian framework designed to characterize the temporal dynamics of cellular interactions using single-cell RNA-sequencing data from multiple time points. Our method utilizes structured Gaussian process regression to unveil time-resolved interactions among diverse cell types according to their co-evolution and incorporates prior knowledge of receptor-ligand complexes. We show the interpretability of DIISCO in simulated data and new data collected from T cells co-cultured with lymphoma cells, demonstrating its potential to uncover dynamic cell-cell crosstalk.

Introduction

Single-cell RNA sequencing (scRNA-seq) which measures gene expression at the resolution of individual cells is a powerful technology for elucidating heterogeneous cell types and states (Tirosh et al., 2016; Wilk et al., 2020; Azizi et al., 2018). The recent expansion of single-cell datasets profiling biological systems and longitudinal clinical cohorts over multiple time points offers an exciting opportunity to characterize the temporal dynamics of cell types and their underlying communication. However, there is an unmet need for computational frameworks that can effectively integrate single-cell data across time points while accounting for temporal dependencies. This need is particularly acute in longitudinal clinical data, where the timing of biopsies cannot be planned or controlled and often varies significantly across patients (Bachireddy et al., 2021).

Complex systems such as the tumor microenvironment are contanstly evolving, particularly with disease progression or therapeatic interventions. Various cancerous and non-cancerous cell types engage in interactions that lead to diverse treatment outcomes. Uncovering crosstalk between tumor and immune cells (Kumar et al., 2018) for example, can unravel immune dysfunction mechanisms. Furthermore, characterizing the dynamic nature of these interactions and their effect is crucial for understanding mechanisms driving response or resistance to therapies and how these mechanisms

can be leveraged to develop more effective therapies (Yofe et al., 2020; Sievers et al., 2023; Maurer et al., 2024).

Current approaches for predicting cell-cell interactions using scRNA-seq data rely on existing databases of interacting protein complexes, and they predict interactions based on expression levels of known receptor-ligand (R-L) pairs (Efremova et al., 2020; Browaeys et al., 2020; Jin et al., 2021; Li et al., 2023). However, these methods have several limitations. The effects and expression of R-L subunits can vary across different cell types, a context-dependent nuance often overlooked by existing tools. Reliance on databases also limits the discovery of novel interactions and rare understudied cell types. Finally, existing tools only predict static interactions. They are not capable of inferring dynamic time-varying interactions from integration of time-series datasets and do not obtain the strength or sign of interactions, e.g. inhibitory or activating crosstalk.

Here, we present DIISCO (Dynamic Intercellular Interactions in Single Cell transcriptOmics), an open-source tool (https://github.com/azizilab/DIISCO_public) for joint inference of cell type dynamics and communication patterns. DIISCO is a Bayesian framework that infers dynamic interactions in scRNA-seq data from non-uniformly sampled time points, incorporates prior knowledge on R-L complexes, and quantifies uncertainty in both its predictions and interpretations. Unlike previous methods, DIISCO can be trained on temporal data to learn how interactions change over time between samples. We demonstrate the performance of DIISCO on simulated data as well as newly generated data collected from chimeric antigen receptor (CAR) T cells interacting with lymphoma cells.

In brief, the key contributions of our work are: (1) Construction of a probabilistic framework for modeling temporal dynamics of diverse and interacting cell types in complex biological systems, through the integration of scRNA-seq datasets collected across non-uniformly sampled time points; (2) Development of first-in-class computational tool for predicting time-resolved cell-cell interactions, along with the strength and sign (activating or inhibitory) of interactions, empowering discoveries of cell communication associated with progression of biological processes or treatments; (3) A novel Bayesian framework for incorporation of prior knowledge on signaling protein complexes with time-series scRNA-seq data to improve identifiability of dynamic cell interactions and quantify uncertainty; (4) Demonstration of performance in simulated and experimental lymphoma-immune interaction data.

Results

Overview of DIISCO

DIISCO captures the temporal dynamics and interactions of cell types using scRNA-seq data from multiple time points. We define cell types or states as clusters of cells with similar gene expression profiles in scRNA-seq data. The time-series single-cell data are first merged to define unique cell types or states through typical clustering approaches (Wolf et al., 2018; Levine et al., 2015). The pooling of cells helps with improving statistical power in the detection of rare cell types. Then, the number of cells assigned to each cell type at each time point is computed and used as training data for DIISCO (**Fig. 1A**). The cell type dynamics can be in the form of cell counts at each time point or standardized proportions of each cell type, i.e. normalized by total cells in the sample collected at each time point. Using proportions is often useful in complex patient data where the sampling and quality of biopsies are inconsistent over time points (Maurer et al., 2024).

Our first assumption is that cell-cell communication, especially if associated with response to external perturbation, can be reflected in compositional shifts in cell states. In other words, a positive or negative temporal correlation between two cell states may be indicative of an activating or inhibitory interaction between them, respectively. Our second assumption is that differential gene expression of receptor-ligand complexes by these correlated cell types reinforces the likelihood of direct cell-cell communication. To improve predictive power in time intervals with low cell numbers, we integrate time-series data and encode temporal dependencies between interaction networks at different time-points. By simultaneously fitting the model to all detected cell states and all time-points, the algorithm effectively identifies the best sender cell state communicating with each receiver cell state, and leverages data observed in other time points, particularly the closest ones to improve prediction at a given timepoint. The output of DIISCO is two-fold. It includes predicted cell type dynamics which is helpful for assessing model fit, e.g. when tuning hyperparameters. Additionally, DIISCO outputs a time-series interaction matrix that reveals which cell-cell interactions are evolving with time (**Fig. 1A**).

At its core, DIISCO utilizes a generative model. To capture the temporal dynamics of cell types, i.e. changes in their frequency over time, while addressing the challenge of varying time intervals, we deploy Gaussian Process regression models (Wilson et al., 2011) (**Methods**). Gaussian pro-

cesses (GPs) provide sufficient flexibility with nonparametric function learning, and therefore do not require any knowledge of form or rate of changes of cell types, thus expanding their applicability. They have been proven successful in integrating time-series single-cell datasets (Lönnberg et al., 2017), and in identifying key cell states with distinct temporal dynamics shaping response to immunotherapy in human leukemia (Bachireddy et al., 2021). By utilizing GPs, which encode temporal dependencies between all pairs of time points, we expand the utility of the model not only to controlled experimental settings but also to scenarios with variable timing, e.g. patient data.

To encode dynamic interactions between cell types (**Fig. 1B**), we draw inspiration from Gaussian Process Regression Networks (GPRNs) which have been previously applied to gene expression dynamics as well as datasets in finance, physics, geostatistics, and air pollution (Wilson et al., 2011; Li et al., 2020). GPRNs are Bayesian multi-output regression networks that leverage the flexibility and interpretability of GPs as well as structural properties of neural networks. The ability to capture highly nonlinear and time-dependent correlations between outputs makes GPRNs ideal for learning complex cell-cell interactions and their variability over time (**Fig. 1C; Methods**). While previous methods have applied GPs to learning cell-type dynamics in time-series single-cell data (Bachireddy et al., 2021), our method is the first to utilize a GPRN framework for learning interactions between cell types, rather than learning the dynamics of each cell type independently.

Finally, DIISCO utilizes an optional prior based on the expression of receptor ligands (R-Ls) and learns interactions from both R-L genepairs and temporal dynamics of cell types, while reliance on R-Ls can be adjusted with the prior strength (**Fig. 1D; Methods**). Utilizing the R-Ls is beneficial for constraining the solution space of interactions and improving model identifiability and robustness.

DIISCO recovers time-resolved interactions in synthetic data

We first evaluated the performance of DIISCO on synthetic data by constructing a dataset with known dynamics and interactions (**Methods**). We observe an accurate modeling of temporal dynamics for 5 simulated cell states, even with volatile dynamics and non-uniform sampling of timepoints (**Fig. 2A,C**).

Through these simulations, we noted DIISCO's ability to uncover a wide range of interactions, including constant interactions, monotonically increasing or decreasing interactions, as well

as transient interactions (**Fig. 2B,D**). Distinguishing various forms of interactions is important in biological data, especially in studying the impacts of perturbations and disease progression, and in resolving the timing of critical crosstalk. In the case of disease progression, interactions between cell types may steadily increase or decrease over time (such as in $C4 \leftarrow C0$ and $C2 \leftarrow C1$), e.g. as malignant cells become more prominent or immune cells exhibit clonal expansion. Recapitulating these complex interactions by DIISCO demonstrates its potential to derive insights into dynamic cell communication.

Since sparse sampling of time-points can typically occur in real-world scenarios, particularly with clinical samples, we evaluated the robustness of our method to the number of time-points. We repeated the described experiment with removing a random subset of timepoints, ranging from 10-100% and observed sustained performance even with 70% of time-points removed (**Fig. 2E**). We also tested the scalability of DIISCO by increasing the number of clusters and time points (**Fig. 2F**), as well as analyzing the algorithmic complexity of DIISCO (**Methods**). Our empirical assessment of scalability confirmed exponential scaling with number of cell types and time points; however we still observe run times less than 5 minutes for 100 simulated clusters and 20 time points using a Google Cloud machine instance with 117 GB of RAM and 30 cores (**Fig. 2D**). This scalability suggests DIISCO would be able to provide insight into large-scale temporal cell atlases, should they become available.

To further benchmark DIISCO's ability to infer temporal dynamics and interactions, we compared DIISCO's performance to both a linear model and a rolling linear model (**Methods**). We also tested the importance of the prior, by comparing to models with and without a prior incorporated. By simulating various scenarios with differing time-points and noise (**Methods**), we confirmed that DIISCO accurately reconstructs temporal changes, and demonstrates superior performance in predicting interactions compared to baseline models without the prior (**Fig. 2G**, **Supplementary Tables 1-2**). The performance in predicting interactions between cell types at measured time-points becomes comparable to the rolling linear model if the DIISCO prior is incorporated, underscoring the impact of promoting sparsity in the interaction network. However, the baseline models are yet incapable of interpolating interactions between measured time-points, limiting the evaluation of interaction dynamics over time. They are also unable to identify R-Ls or gene programs correlated with time-resolved cell type interactions. Overall, DIISCO performs both tasks of inferring smooth

time-dependent interactions without assumptions about functional form and incorporating prior knowledge in an integrated Bayesian framework, while quantifying uncertainty and enabling downstream analysis of underlying mechanisms. This is an essential feature for biological applications and hypothesis generation for future investigations.

DIISCO characterizes dynamic interactions between CAR-T and lymphoma cells

To demonstrate an application of DIISCO to biological data, we generated single-cell data from a controlled *in vitro* experimental setting. Specifically, we co-cultured GFP-transduced anti-CD19 chimeric antigen receptor (CAR) T cells together with MEC1 cells (Stacchini et al., 1999) - a B cell (chronic lymphocytic leukemia [CLL]) leukemia cell line expressing *CD19* - as the CAR-T cell target (**Fig. 3A**; **Supplementary Fig. S1A-C**; **Methods**). We confirmed dose response activation (cytotoxicity) of T cells by quantifying the percentage of remaining MEC1 cells. This showed a reduction in MEC1 cell survival starting from $50-60\%$ at an effector-to-target ratio of 0.125 to survival of $<5\%$ at an effector-to-target ratio of 4 (**Supplementary Fig. S1D**). We profiled 4 biological replicates with scRNA-seq at 10 time-points spanning $24hrs$, thereby obtaining high-quality data for $49,283$ total cells (**Supplementary Table 3**). The data shows compositional shifts with time, motivating the use of this dataset as a test case for investigating temporal interactions with DIISCO (**Fig. 3B**).

We preprocessed this data by defining major cell types from the combination of all replicate experiments to increase statistical power. Through clustering and examining expression of curated gene sets, we identified 4 major cell types (metaclusters): cancer cells, exhausted $CD8^+$ T cells, activated $CD8^+$ T cells, and other $CD8^+$ T cells showing neither activated nor exhausted markers. MEC1 cancer cells were annotated based on the expression of *CD19* and *CD79A*. T cells were annotated based on the expression of *CD3E*, *CD3D*, and *CD8A*. Activated T cells were defined by expression of *CD69*, *CD27*, and *CD28*, while exhausted cells were defined by expression of *TIGIT*, *PDCD1*, *TGFB1*, and *LAG3* (**Fig. 3C-E**). Clusters that were positive for both T cell and MEC1 gene markers were removed as doublets. We acknowledge that doublets could also contain interacting cells. However, since we only obtain one measurement of receptor or ligand gene, our current method (as well as existing methods) cannot resolve them. Future extensions deploying mixture models could help deconvolve doublets.

After cell type annotation, proportions were calculated at each time point, and DIISCO was trained on each experiment separately (**Supplementary Fig. S2A,C**). To construct the prior for interactions, we obtained a set of 8,234 literature-curated receptor-ligand (RL) pairs from OmniPath, a database of cell signaling prior knowledge (Türei et al., 2016) and summarized the number of R-Ls with differential expression between each pair of cell types (**Fig. 3F**; **Methods**).

Applied to the CAR-T and MEC1 data, DIISCO learned the proportions of cell types, highlighting the expected decrease in cancer cells over time and the increase in exhausted T cells (**Fig. 4A; Supplementary Fig. S2A,C**). We also inferred interactions between cell types and specifically observed strong negative interactions predicted between cancer cells and exhausted CAR-T cells, as well as between activated and exhausted T cells (**Fig. 4B,C; Supplementary Fig. S2B,D**). Investigating the dynamics of interactions $W(t)$ over time, as expected, we predicted an increase in the strength of the interaction ($|W|$) with co-culture time (**Fig. 4C**) between exhausted T cells and MEC1 cells, as well as in exhausted T cells interacting with activated T cells, which was directionally specific. The former recapitulated the targeting of cancer cells by T cells and the latter may be due to phenotypic shifts from activated T cells becoming more exhausted or terminally differentiated over time (Bachireddy et al., 2021). When comparing DIISCO results across experiments, we note that the predicted interactions (W) between cell types averaged over time, remain consistent for the strongest positive and negative interactions **(Supplementary Fig. S2E**. Furthermore, the main interaction links between Exhausted T cell - MEC1 and Exhausted - Activated T cells exhibit the same time-varying dynamics across experiments, demonstrating DIISCO's robustness and reproducibility across biological replicates **(Supplementary Fig. S2F-G**). One benefit of DIISCO is its ability to provide uncertainty estimates in interaction predictions, and we note that the uncertainty, as defined by the 85% confidence interval, increases after 12 hours, due to increased sparsity in collected timepoints (**Supplementary Fig. S3**).

The ability to predict how interactions between exhausted T cells and MEC1 cells evolve with time is a unique feature of DIISCO, whereas other existing methods are not able to achieve such time-resolved predictions. Indeed, we observed a positive correlation between predicted interaction strength and co-culture time (**Fig. 4C**), in line with the cytotoxicity of T cells, i.e. their increased ability to kill cancer cells.

Additionally, we assessed the sensitivity of DIISCO to the interaction prior by changing the

interaction score threshold that we used for quantifying the number of possible R-L complexes associated with a pair of cell states. A lower threshold would lead to a larger solution space, while higher thresholds emphasize clusters with the largest number of expressed R-L genes. We noted that the average interaction strength predicted by DIISCO remained stable, excluding the interactions that drop out with higher thresholds (**Fig. 4D**). Importantly, the ranking of most strongest predicted interactions remained stable (**Fig. 4E**).

To assess the sensitivty of DIISCO to the number of cells, we downsampled the cells in this dataset and observed robust ranking of interactions even when cells are downsampled to 30% of the original numbers, with only minor changes when downsampling $40-50\%$ of cells **(Fig. 4F-H; Supplementary Fig. 4A)**. To additionally test sensitivity to user input, we also downsampled the time-points in this dataset (**Supplementary Fig. 4B**). DIISCO recovers similar dynamics even with 50% of time-points dropped. Tracking the average interaction strengths over time, we see robust estimations even with 30% of time points dropped (**Supplementary Fig. 4C**). This is also evident in the time-varying dynamics of these interactions (**Supplementary Fig. 4D,E**). We additionally sought to test sensitivity to clustering techniques, as users may utilize different pre-processing pipelines. We first tested DIISCO's performance on the individual Phenograph clusters without merging them into metaclusters. The results show MEC1 cells (clusters 10, 17, 22, 23) have a decrease in proportion after 10 hours post-co-culture (**Supplementary Fig. 5A**). Activated T cells (clusters 3, 12) show higher proportions early in the experiment and then decrease over time. This complements the dynamics of exhausted T cells (clusters 9, 15, 26) which are increasing over time.

Examining the average W interactions over time, C10→C26 and C10→C9 are the strongest negative interactions (**Supplementary Fig. 5B-C**) These are both MEC-Exhausted Tcell links consistent with our previous results at the level of metaclusters. The strongest positive interactions belong to C23→C9, C9→C23, and C26→C9. These are MEC-Exhausted T cell and Exhausted T cell -Exhausted T cell interactions. These results demonstrate DIISCO's robustness to various clustering methods and varying data quality. One benefit to combining clusters into metaclusters that define individual cell types, is the ability to remove any self-interactions. With multiple clusters representing exhausted T cells, we see interactions between different exhausted T cell clusters arise. Although there are cases where this may be important to study, DIISCO has been built to study

intercellular interactions between different cell types, instead of within them. We also demonstrate DIISCO's robustness to clustering method by re-procecssing the data following SCANPY's pipeline, including Leiden clustering to define cell types (**Methods**). We see similar results in both dynamics and predicted interactions (**Supplementary Fig. 5D-F**) .

DIISCO outperforms existing methods in identifying cell communication mediated by signaling molecules

To investigate mechanisms underlying the interaction between MEC1 cells and exhausted T cells in our *in vitro* experiment, we calculated the correlations between different R-L gene pairs and the learned dynamic interaction (**Fig. 5A**). For calibration, we first confirmed that $CD19$, the target of the engineered CAR, is highly correlated with the interaction from exhausted T cells to MEC1 cells, as expected (**Fig. 5A**). Examining the MEC1 $\rightarrow$ Exhausted interaction, *CD86-CTLA4* and *CD80-CTLA4* were the top R-L pairs ranked by ligand correlation (**Fig. 5A**). *CTLA4* is an established marker of T cell exhaustion, and *CD80* and *CD86* are costimulatory molecules that are expressed in malignant lymphocytes (including B cells) in several hematologic diseases (Dorfman et al., 1997; Vyth-Dreese et al., 1998; Delabie et al., 1993). Ranked by the highest average correlation, *ICAM1-IL2RA* was the top predicted R-L pair. *ICAM1* is a known regulator of inflammatory response and is elevated in CLL patients with more severe disease progression (Molica et al., 1995; Musolino et al., 1996; Bui et al., 2020).

Finally, we benchmarked the predicted DIISCO R-L interactions against other methods including CellChat (Jin et al., 2021) and CellphoneDB (Efremova et al., 2020). We applied CellPhoneDB and CellChat to log-normalized expression data from all time points. We excluded any interactions within the same cluster (i.e. MEC1 $\rightarrow$ MEC1). R-L pairs were filtered based on p-value ($p < 0.05$), and we identified 260 significant pairs using CellPhoneDB and 180 significant interactions using CellChat.

To compare these results to DIISCO-predicted interactions, we calculated the correlation between the predicted receptor and ligand expression over all time points. Our rationale is that if R-L protein complexes mediate cell-cell interaction, their expression levels on the corresponding sender and receiver cell types will likely change concomitantly. Otherwise, the genes are less likely to be involved in the cellular interaction and could rather be intrinsic markers of the cell states. We

observed a significantly higher Pearson's correlation for R-Ls in DIISCO-predicted interactions than those predicted by other methods (**Fig. 5B**). Since changes in the expression of receptor genes on the receiver cell type may not be linearly associated with the corresponding ligand expression on the sender, we also calculated Spearman's correlation, again demonstrating a significant improvement with DIISCO (**Fig. 5B**). Although all 3 methods identified the *CD80/CD86-CTLA4* interactions, only DIISCO predicted the *ICAM1-IL2RA* interaction, which is the one with the most time-varying expression. Note that temporal changes in expression of R-L genes is not used in the DIISCO prior. Our results thus confirm the ability of DIISCO to identify a monotonic increase in the expression of R-L pairs underlying cell-cell interactions.

Further examining the overlap between methods, we noted that the predicted interactions shared between DIISCO and CellChat are the most correlated by both Spearman's correlation and mutual information as another metric, while the interactions predicted by both DIISCO and CellphoneDB are most correlated using the Pearson's correlation test (**Fig. 5C**). Combined, these results confirm DIISCO's ability to identify cell-cell communication that can be supported by co-expressed R-L gene pairs, from time-series single-cell data.

As DIISCO learns interactions from the integration of time-series data, we further compared DIISCO to cellphoneDB applied to different time-points. Due to low numbers of cells in metaclusters in individual time-points, we were unable to apply cellphoneDB to each time point separately. Instead, we combined the first 5 and last 4 time points into 'early' and 'late' time intervals to obtain pseudo-dynamic interactions. Identifying R-L pairs that were only predicted in the 'late' bucket suggests that these interactions are growing in strength over time. Predicted R-L interactions between MEC1–Exhausted cells found with both DIISCO and the 'late' bucket in cellphoneDB include *TNF-ICOS*, *FAM3C-PDCD1*, and *FAM3C-LAMP1*. *ICOS* is known to be expressed on activated T cells, and it has been investigated as a costimulatory domain for CAR-T cell therapy((He et al., 2022)). High expression of *TNF* in CLL cells has been associated with poor prognosis, and *TNF* has been suggested as a potential target for CLL immunotherapies ((Dürr et al., 2018)). Even with the bucketing approach cellphoneDB did not predict *CD80/CD86-CTLA4* interactions, which have strong literature support for their importance in CLL ((Dorfman et al., 1997; Vyth-Dreese et al., 1998; Delabie et al., 1993)). CellphoneDB also did not identify the *ICAM1-IL2RA* interaction. This suggests that incorporating temporal dynamics is crucial for identifying R-L interactions most

relevant to the biological system.

Discussion

As longitudinal single-cell datasets become more prevalent, we believe DIISCO can provide a rigorous and flexible framework for characterizing cell-cell communication and its temporal variability, aiding in the understanding of disease progression and impact of treatments or perturbations in complex biological systems.

We demonstrate the performance of DIISCO in simulated data as well as cancer-immune cell interaction data. In particular, we show how the integration of time-series data and prior knowledge of protein complexes in a probabilistic model leads to achieving dynamic intercellular interactions. The ability to infer the sign and change in strength of interactions with time is not attianable with previous methods and is crucial for studying the impact of perturbations such as treatment. In addition, DIISCO provides insight into signaling mechanisms mediating dynamic interactions.

The advantage of the probabilistic framework is that with very few time points or long time intervals, DIISCO will present high uncertainty for estimates. We also show high accuracy in synthetic data for predictions even with sparse sampling of time points. While the interaction prior impacts model identifiability, we showed that the strongest inferred interactions are robust to the prior strength.

The choice of hyperparameters is essential for achieving interpretable dynamics. We thus provide comprehensive guidance for the choice of hyperparameters according to data quality and sampling rate in Supplementary Information Section D. Additionally, spurious correlations may emerge when using cell type proportions due to the constraint of having all points in the simplex. Supplementary Information section E includes a short discussion on this issue.

As DIISCO assumes that cellular interactions lead to shifts in the cell state composition in the system, it may not be an appropriate method if no changes in cell state proportions or numbers are observed over time. This may be the case with communication mediated by cells exhibiting static proportion. However, compositional shifts frequently occur with perturbations such as therapy or external stimuli, and are particularly relevant in clonally expanding lymphocytes driving immune response, e.g. in hematological malignancies such as leukemia. We recently showed that DIISCO can

resolve time-varying dynamics associated with response to therapy in single-cell data collected from relapsed leukemia patient samples (Maurer et al., 2024). This dataset involves 74 longitudinal bone marrow samples from 25 heterogeneous relapsed leukemia patients. Furthermore, the number of cells ranged from 144 to 14,584 high-quality cells across samples and our results demonstrate DIISCO's robustness to unbalanced datasets. DIISCO revealed a cascade of immune interactions centered around a cytotoxic T-cell subtype that expands with therapy, and further identified potential R-L complexes associated with predicted interactions.

To address limitations in relying on changes in cell type composition in other systems, DIISCO can potentially be adapted to learn temporal correlations between modules of co-expressed genes instead of cell state proportions. Furthermore, there is an inherent limitation in using R-L databases to learn interactions. Although DIISCO addresses this limitation through binarization, the lack of context-dependent information in these databases requires the user to interpret the predicted interactions and downstream gene networks relevant to the context. Identified interactions should also be validated through literature searches or follow-up experiments.

Future application of DIISCO to other clinical datasets can elucidate cell-cell communication underlying response or resistance to treatments such as immunotherapy. Deciphering the mechanisms of response and lack of response to therapy allows for more patient-specific therapeutics and offers the opportunity to reverse-engineer new therapies for improved efficacy.

We envision multiple directions for building upon this work. Notably, for clinical samples, incorporating priors that account for sample size (i.e., total cells) could improve robustness by down-weighting the influence of samples with very few cells. Recent advances in spatial transcriptomics also offer an opportunity to expand DIISCO and infer interactions from both temporal and spatial dynamics by integrating spatial co-localization information.

Methods

Notations

We assume that we have N cell types with their frequencies measured at time points $t_1, \ldots, t_N$. We define $y(t_i)$ as a K-dimensional vector of observations at time t_i where the k-th dimension corresponds to the frequency of the k-th cell type. Additionally, we assume we have a set of M

unobserved time points $t_{N+1}, \ldots, t_{N+M}$, placed anywhere on the time axis, for which we would like to infer (interpolate) the cell type values $y(t_{N+1}), \ldots, y(t_{N+M})$. We denote the set of all time points as $\mathcal{T} = \{t_1, \ldots, t_{N+M}\}$ and call $\mathcal{T}_u$ the set of unobserved time points $\mathcal{T}_u = \{t_{N+1}, \ldots, t_{N+M}\}$ and $\mathcal{T}_o$ the set of observed time points $\mathcal{T}_o = \{t_1, \ldots, t_N\}$. We use the convention of $\cdot_u$ and $\cdot_o$ to denote unobserved and observed variables respectively. In the remainder, for ease of exposition, we will refer to $y(t_i)$ as proportions with the understanding that either proportions or cell counts could be used, although the results should be interpreted in a different manner depending on the value used.

In addition to these features, we have a binary matrix Λ of size $K \times K$ where $\Lambda_{k,k'} = 1$ indicates that the k-th cell type might interact with the k'-th cell type and $\Lambda_{k,k'} = 0$ indicates that they don't. We allow for non-symmetric matrices to allow for directionality in interactions. In practice, we obtain this matrix from measurements of expression of known receptor-ligand complexes.

Model Specification

DIISCO is a generative model that assumes cell type frequencies, denoted as $\widehat{y}(t_i)$, are derived from the following process (**Fig. 1B**):

1. For every time point, we sample a set of latent features, $f(t_i) \in \mathbb{R}^K$ where every feature, i.e every coordinate of $f(t_i)$, is a Gaussian process. We call this set $\mathcal{F} = \{f(t_i) \mid t_i \in \mathcal{T}\}$ and let $\mathcal{F}_o$ and $\mathcal{F}_u$ denote the set of latent features at observed and unobserved time points respectively.

2. Similarly, we sample at each time point an interaction matrix $W(t_i) \in \mathbb{R}^{K \times K}$ where we also assume that each of the $K \times K$ coordinates of $W(t_i)$ is a Gaussian process across time. We call this set $\mathcal{W} = \{W(t_i) \mid t_i \in \mathcal{T}\}$ and let $\mathcal{W}_o$ and $\mathcal{W}_u$ denote the set of interaction matrices at observed and unobserved time points respectively.

3. Finally, we sample the standardized cell type proportions $\widehat{y}(t_i) \in \mathbb{R}^K$. We do this by sampling from a multivariate Gaussian with mean $W(t_i)f(t_i)$ and covariance $\sigma_y^2 I$. In other words:

$$\widehat{y}(t_i) = W(t_i)f(t_i) + \epsilon(t_i) \tag{1}$$

where $\epsilon(t_i) \sim \mathcal{N}(0, \sigma_y^2 I)$ represents a zero-centered Gaussian noise process. We use $\mathcal{Y} =$

$\{\widehat{y}(t_i) \mid t_i \in \mathcal{T}\}$ to denote the set of all standardized cell type proportions across all time points with $\mathcal{Y}_o$ and $\mathcal{Y}_u$ having the same meaning as before. We also allow for the inclusion of a bias term $b(t)$ such that $Y(t) = W(t)f(t) + b(t) + \epsilon(t)$ where $b(t)$ is centered at 0 and has a kernel like that of $W(t)$ but without any interaction regularization. For brevity, we omit it as it is equivalent to extending the model dimension by one, ignoring the last coordinate of $y(t)$, and setting the last coordinate of $f(t)$ to be a GP centered at 1 with infinite length-scale.

Figure 1C depicts this process graphically and Algorithm 1 details the generative process. The arrows $\mathcal{W}_o \to \mathcal{W}_u$ and $\mathcal{F}_o \to \mathcal{F}_u$ represent tractable distributions that can be computed analytically and can be sampled due to the properties of Gaussian processes. Thus, the model only requires tractable operations that lead to the joint distribution described. DIISCO currently ignores the original space that the data lies in and makes the simplifying assumption that $y(t)$ can be treated as a vector in $\mathbb{R}^k$ rather than on the simplex or the space of positive integers. Although these assumptions aid greatly in the interpretability of the latent variables, further improvements using distributions that truly match the real support of the data would lead to better uncertainty estimates and predictions.

Algorithm 1 Generative Process used by DIISCO

Input: Set of time points $\mathcal{T}$, Number of latent features K, Noise covariance σ_y^2.

Sampling of Latent Features
for each feature $k \in \{1, \ldots, K\}$ **do**
    Sample $f_k(t) \sim \mathcal{GP}(m_k^f(t), K_k^f(t_i, t_j))$
end for

Sampling of Interaction Matrix
for each feature $k \in \{1, \ldots, K\}$ **do**
    for each feature $k' \in \{1, \ldots, K\}$ **do**
        Sample interaction $W_{k,k'}(t) \sim \mathcal{GP}(0, K_{k,k'}^W(t_i, t_j))$
    end for
end for

Sampling of Observed Values
Sample a K-dimensional Gaussian noise process $\epsilon(t) \sim \mathcal{N}(0, \sigma_y^2 I_K)$
Compute standardized cell type proportions: $\hat{y}(t) \leftarrow W(t)f(t) + \epsilon(t)$

Model Interpretation.

The latent variable $W_{i,j}(t)$ represents the direction and effect of intercellular interaction (signaling communication) from cell type j to cell type i at time t. In particular, $W_{i,j}(t) > 0$ represents an activating effect, $W_{i,j}(t) < 0$ denotes inhibitory impact, and $|W_{i,j}(t)|$ reflects the strength of the interaction. $f(t)_i$ is a latent variable that represents the normalized proportion of cell type i at time t (**Fig. 1B**). DIISCO focuses on identifying interactions between different cell states. Thus, in constructing the prior for W, we penalize self-interactions $W_{i,i}$, such that $W_{i,j}$ can be interpreted as the impact of other cell types $j \neq i$ on the dynamics of cell type i. We justify why this interpretation is reasonable, given our choice of priors below.

Prior Distribution over $\mathcal{F}$.

Intuitively, we aim for the latent features to align closely with standardized cell type proportions so that $W(t)$, with suitable restrictions on the diagonal, learns to capture the interactions between cell types by predicting the standardized proportion of one cell type given the standardized proportions of the rest. We achieve this by first defining for every feature an auxiliary zero-mean Gaussian process (GP) on which we perform inference. Formally, for every $k \in \{1, \ldots, K\}$ we define an auxiliary GP with covariance function K^f given by:

$$K^f(t_i, t_j) = v_f \exp\left(-\frac{(t_i - t_j)^2}{2\tau_f^2}\right) + \sigma_f^2 \delta_{ij}. \tag{2}$$

where δ_{ij} is the Kronecker delta function, and v_f, τ_f, and σ_f are shared hyperparameters denoting the variance of the latent features, the length scale of the latent features and the noise of the latent features, respectively. We then update this auxiliary GPs using the observations $\widehat{y}_k(t_i)$ at the observed time points $t_i \in \mathcal{T}_o$ which provides us with a posterior Gaussian Processes $\mathcal{GP}(\mu_k^f, K_k^f)$ that we use as the prior for the Gaussian Process corresponding to feature $f_k(t)$ in Algorithm 1.

Prior Distribution over $\mathcal{W}$.

To further limit the solution space, and improve model robustness and interpretability, we set two constraints on the sampling process of W. First, we set off-diagonal elements to zero if the cluster pairs do not express any complementary receptor-ligand pairs and second, we zero out the diagonals

when a row has at least one other non-zero entry. In the case when we are dealing with proportions and not raw counts this also ensures that we avoid a trivial solution due to the $\sum_k y_k(t) = 1$. Formally, we achieve this by sampling $W(t)$ so that for $k, k' \in \{1, \ldots, K\}$.

$$W_{k,k'}(t) \sim \mathcal{GP}(0, K^W_{k,k'}(t,t)) \tag{3}$$

where

$$K^W_{k,k'}(t,t') = \begin{cases} v_W \exp\left(-\frac{(t-t')^2}{2\tau_W^2}\right) + \sigma_W^2 & \text{if } k \neq k' \\ 0 & \text{if } k = k' \text{ or } \Lambda_{k,k'} = 0 \end{cases} \tag{4}$$

and v_W, τ_W, and σ_W are shared hyper-parameters, retaining the same interpretations as before. To construct Λ, we scan for complementary expression of known RL pairs, i.e. when the sending cluster expresses a ligand at a defined threshold and the receiving cluster expresses the complementary receptor at a similarly defined threshold (**Fig. 1D**). This threshold is user-specified, as different genes are captured at different rates in sequencing data and a single threshold may not be an accurate representation of relative expression across all genes.

We then quantify the number of these interactions expressed by cluster pairs. The number of RL pairs does not necessarily determine the strength of interaction, and a single complementary RL pair could be biologically important. To account for this, we set a binarization threshold. The binarization threshold can be set by users and may be datatype-specific. A very high threshold leads to a prior that is too sparse and may lead to losing potentially relevant interactions, whereas a very low threshold may lead to spurious interactions. The threshold is user-specified and determines the sparsity of the prior during the binarization step.

To demonstrate the robustness of predicted interactions to the threshold, we performed additional sensitivity analysis by dropping 30-70% of the R-L pairs in the Omnipath R-L database. When we drop different percentages of R-L pairs, we obtain a similar distribution of interaction scores (Supplementary Fig. 6A-C). By specifying a threshold guided by the knee point of these distributions, we achieve the same prior using 30% or 70% of the data. By changing the threshold to account for downsampled or incomplete interactions, we can retain flexibility in the model and utilize a consistent prior. Unfortunately, similar to current interaction methods like CellphoneDB

or CellChat, if there is a bias in the database where certain cell types are underrepresented, we are unable to capture that in the predicted results. Future databases expanding these interactions can improve results obtained with DIISCO.

By binarizing the RL interaction matrix, we allow for more model flexibility while also limiting the solution space to exclude clusters with no complementary RL expression. This final binary matrix is defined as Λ, and as shown in equation 4, the value of $\Lambda_{k,k'}$ determines whether the model allows for a possible interaction between cell types k and k'. In practice, we build the prior using OmnipathDB, a database used by more traditional interaction methods like cellphoneDB ((Türei et al., 2016)). However, it is important to note that R-L databases lack context-dependent information and are often incomplete with many interactions having not been characterized yet in literature.

Inference

Due to the complex nature of the model, it is not possible to perform inference analytically. Instead, we deploy a two-step approximate inference method combining both analytic and approximate techniques.

Specifically, we: (1) perform approximate inference to obtain samples from an estimate of the posterior distribution $p(\mathcal{W}_o, \mathcal{F}_o \mid \mathcal{Y}_o)$; and (2) perform ancestral sampling and standard GP conditioning to obtain an approximation to the posterior distribution $p(\mathcal{W}_u, \mathcal{F}_u, \mathcal{Y}_u \mid \mathcal{Y}_o, \mathcal{W}_o, \mathcal{F}_o)$. Supplementary Information Section B provides a justification for this approach. Below we describe in detail how we perform each of these steps.

Approximate Inference.

To approximate $p(\mathcal{W}_o, \mathcal{F}_o \mid \mathcal{Y}_o)$ we use variational inference implemented using Pyro (Bingham et al., 2019), a probabilistic programming language written in Python. Within this framework, we define a parametrized distribution $q_\phi(\mathcal{W}_o, \mathcal{F}_o)$ and then optimize the parameters ϕ to maximize an estimate of the evidence lower bound (ELBO) (Blei et al., 2017) via a gradient descent algorithm. In the situation where the hyper-parameters are fixed, this is equivalent to minimizing the KL divergence $\mathrm{KL}\left(q_\phi(\mathcal{W}_o, \mathcal{F}_o) || p(\mathcal{W}_o, \mathcal{F}_o | \mathcal{Y}_o)\right)$.

In this case the ELBO is given by

$$\mathcal{L}(\phi) = \mathbb{E}_{q_\phi(\mathcal{W}_o, \mathcal{F}_o)} \left[\log p(\mathcal{Y}_o, \mathcal{W}_o, \mathcal{F}_o) - \log q_\phi(\mathcal{W}_o, \mathcal{F}_o) \right] \tag{5}$$

and we define $q_\phi(\mathcal{W}_o, \mathcal{F}_o)$ as a multivariate Gaussian factorized by:

$$q_\phi(\mathcal{W}_o, \mathcal{F}_o) = \prod_{k,k' \in \{1,\dots,K\}} \mathcal{N}(W_{k,k'}(\mathcal{T}_o) \mid \hat{\mu}^W_{k,k'}, \hat{\Sigma}^W_{k,k'}) \prod_{k \in \{1,\dots,K\}} \mathcal{N}(f_k(\mathcal{T}_o) \mid \hat{\mu}^f_k, \hat{\Sigma}^f_k) \tag{6}$$

where $\hat{\mu}^W_{k,k'} \in \mathbb{R}^n$ and $\hat{\Sigma}^W_{k,k'} \in \mathbb{R}^{n \times n}$ are the mean and covariance of $W_{k,k'}(\mathcal{T}_o)$, $\hat{\mu}^f_k$ and $\hat{\Sigma}^f_k$ are the mean and covariance of $f_k(\mathcal{T})$, $W_{k,k'}(\mathcal{T}_o)$ denotes the vector of values of $W_{k,k'}(t)$ for $t \in \mathcal{T}_o$ and $f_k(\mathcal{T}_o)$ denotes the vector of values of $f_k(t)$ for $t \in \mathcal{T}_o$.

In other words, we assume that the latent features and the interaction matrix are independent of each other, but they are dependent across timepoints, and this dependency is captured by the mean and covariance of the Gaussian distribution. In practice, we parameterize the Cholesky decomposition of the covariance matrix rather than the covariance matrix itself and use Adam (Kingma and Ba, 2014) to optimize the parameters ϕ. Due to the amount of computation required for using this variational family, we also provide in the package the option to use a mean-field guide where each variable is fully independent from each other and is parameterized by a one-dimensional Normal distribution, where $q_\phi(\mathcal{W}_o, \mathcal{F}_o) = \prod_{z \in \mathcal{W}_o \cup \mathcal{F}_o} q(z|\mu_z, \sigma_z)$. We refer to the family with the time-wise dependency structure as the partially factorized variational family, and the one with a full diagonal covariance as the fully factorized variational family.

To perform optimization, we used an expectation with the form $\mathbb{E}_{q_\phi(Z)}[h_\phi(Z)]$ where $Z = (\mathcal{W}_o, \mathcal{F}_0)$ and h represents the function inside of the expectation in equation 5. This is problematic for taking the gradient with respect to ϕ because it appears both in the distribution with respect to which we are taking the expectation and in h. However, using the fact that if $L = \text{Cholesky}(\Sigma)$, $\epsilon \sim \mathcal{N}(0, I_d)$ and $z = \mu + \Sigma\epsilon$ then z must be distributed as $\mathcal{N}(\mu, \Sigma)$, it is easy to see that one can rewrite the expectation as $E_{\mathcal{D}(\epsilon)}[h(z(\epsilon, \phi))]$ where $\mathcal{D}$ is a multivariate normal distribution with identity covariance and $z(\epsilon, \phi)$ maps this random variable to the equivalent Z values using the above approach. We use this reparametrization trick (Kingma and Welling, 2022) to construct and estimate the gradient of the ELBO. Supplementary Information Section C outlines the inference

algorithm and provides details of ancestral sampling.

Algorithmic Complexity and Scalability

An important consideration for Gaussian Process-based models is their scalability and time complexity. For DIISCO, using either variational family, the computational complexity of approximate inference is of the order $O(|\mathcal{T}_o|^3 K^2)$ per gradient step and the exact conditioning step leads to a computational complexity of $O(K^2(|\mathcal{T}_o|^3 + |\mathcal{T}_o|^2|\mathcal{T}_u|))$ for either sampling or computing the means. Supplementary Information Section F details which exact operations lead to the above bounds.

Nevertheless, for most of the datasets that could appear in practice the algorithmic complexity is not an issue, specially when using the fully factorized normal guide. To demonstrate this, we ran DIISCO on various simulated datasets with different sizes of $\mathcal{T}_o$ and K to benchmark the performance. Specifically, we constructed the datasets by sampling the K Gaussian Processes with half of the processes being independently sampled and the other half being a linear combination of the independent ones. The values for the linear combination were also sampled from a GP with 10 times the length scale of the independent ones, which was set to 1. To simulate sparsity for every non-independent processes we selected at random a number between 0 and floor$(K/2)$ and then proceeded to zero out that many values in the linear combination. The observed time points were sampled uniformly between 0 and 10. We then ran DIISCO for up to 50,000 iterations with early stopping if there were no improvements for over a $1,000$ iterations. We conducted these experiments on a Google cloud machine instance with 117 GB of RAM and 30 cores. **Figure 2F** shows that the run-time remains under 5 minutes even when fitting to up to 100 clusters and 20 time points, which are reasonable dimensions for comprehensive time-series single-cell datasets.

Benchmarking with Simulated Data

To understand the behavior of DIISCO under different scenarios, we conducted a benchmarking experiment using simulated data. We specifically aimed to evaluate the ability to predict interactions with nonzero strength. Details of this process are provided below.

Data Generation To define ground truth interactions for W, we simulated 7 non-zero dynamic interactions among 5 cell types using the equations below. To include cases of isolated cell types

that do not interact with others, we also allow non-zero diagonals in the prior for a subset of cell types. This can alternatively be captured with a bias term. In particular, we assume:

$$
\begin{aligned}
W_{0,0}(t) &= 1 & W_{2,1}(t) &= \cos\left((\mathrm{t}-10)/30\right)\\
W_{1,1}(t) &= 1 & W_{3,0}(t) &= \cos\left((\mathrm{t}+20)/40\right)\\
W_{4,0}(t) &= \cos\left(\mathrm{t}/30\right) & W_{3,1}(t) &= \sin\left((\mathrm{t}+20)/40\right)\\
& & W_{4,2}(t) &= \cos\left((\mathrm{t}+10)/20\right)
\end{aligned}
$$

Any interactions not listed above were set to 0. Data points were generated using the following equations:

$$
\begin{aligned}
c_0(t) &= \frac{1}{2}\cos\left((\mathrm{t}+5)/3\right) + \cos\left(\mathrm{t}/5\right) + \epsilon_1(t) & c_2(t) &= W_{2,1}c_1(t) + \epsilon_3(t)\\
c_1(t) &= \sin\left((\mathrm{t}+1)/2\right) + \cos\left(\mathrm{t}/3\right) + \cos\left((\mathrm{t}+0.5)/3\right) + \epsilon_2(t) & c_3(t) &= W_{3,0}c_0(t) + W_{3,1}c_1(t) + \epsilon_4(t)\\
& & c_4(t) &= W_{4,2}c_2(t) + W_{4,0}c_0(t) + \epsilon_5(t)
\end{aligned}
$$

where ϵ is a noise term drawn for each time point for each cluster, and $\epsilon_i(t)$ represents noise processes. We run these simulations with two different epsilons, with a standard deviation equal to 0.1 and 0.5 to represent differences in data quality. We create the dataset by sampling 100 points uniformly between 0 and 100 and then randomly downsampling the data to a range between 10-90 time points. Based on this generative process, c_0 and c_1 are independent and act on c_2, c_3, and c_4 with different interaction strengths dependent on a time-varying W. We set the following elements of the prior matrix:

$$
\Lambda_{c_0,c_0} = \Lambda_{c_0,c_1} = \Lambda_{c_0,c_2} = \Lambda_{c_0,c_3} = \Lambda_{c_0,c_4} = \Lambda_{c_1,c_1} = \Lambda_{c_1,c_2} = \Lambda_{c_1,c_3} = \Lambda_{c_2,c_4} = 1 \tag{7}
$$

All other elements of Λ are set to 0. To address potential issues with co-linearity, any cluster pairs that were correlated but not causally related through the defined equations were set to 0. We follow the guide in Supplementary Information Section D. For fitting the model we let $\tau_F = 5$, $\sigma_f = 0.1$, $v_f = 3$, $\sigma_w = 0.01$, $v_w = 20$, $\sigma_y = 0.1$ and we place a gamma hyperprior on the lengthscale of W with 90 percent of the mass between 10 to 30. For this simple example, we use the DiagonalNormal guide and use Adam (Kingma and Ba, 2014) with a learning rate of 0.005

Baseline Models We compare DIISCO against two baselines.

1. **Linear Model**: We fit a linear model on the centered version of the data. For a fair comparison, we only regress $y_k(t)$ on all $y_{k'}(t)$ such that $\Lambda_{k,k'} = 1$. We also provide a version without using the prior.

2. **Rolling Linear Model**: As with the pure linear model, we fit a linear model on the centered version of the data with regressors chosen via Λ. However, unlike the previous model, at every observed time point t_{obs}, we fit a separate linear model containing only the data points in the set $\{y_k(t_i) \mid 1 \leq i \leq N, |t - t_{\text{obs}}| < \epsilon\}$ or the closest $n_{\text{min-closest}}$ data points, whichever set is larger. $n_{\text{min-closest}}$ and ϵ are hyperparameters chosen by the user. We also provide a version without using the prior we construct.

We chose these models as they capture particular aspects of DIISCO, such as linear dependency structure and, in the case of the rolling model, changing dependence through time. However, the baseline models are limited in that they don't capture uncertainty in the predictions and don't include a priori biases on smoothness.

Converting Latent W Matrix to 0-1 Interactions The result of DIISCO is a continuous latent matrix $W(t)$. Because this matrix can be large as it scales with the square of the number of cells, it is important to find a method of discretizing it automatically to determine whether a cell is active or not. A helpful guideline for our experiment is to say that cell i interacts with another cell j at time t if $|\mathbb{E}[W_{i,j}(t)]| \geq C\hat{\sigma}_W$, where C is a user-determined constant and $\hat{\sigma}_W$ is the standard deviation of all values of the form $|\mathbb{E}[W_{k,k'}(t)]|$ with $k, k' \in \{1, \ldots, K\}$ and $t \in \{t_1, \ldots, t_N\}$. In our experiments, we found that setting $C = 1$ is a good heuristic, but this might vary depending on the sparsity assumptions that a practitioner might have about their data. We emphasize that this heuristic can only be applied if the data has been scaled and centered so that the magnitude of the coefficients can be interpreted meaningfully.

CAR-T Data Collection

CD19 CAR-T cells were generated by transducing healthy donor T cells with 3rd generation lentiviral vectors encoding a bicistronic construct containing either FMC63 CD19 scFv-CD28-CD3z

and green fluorescent protein (GFP) or FMC63 CD19 scFv-41BB-CD3z (**Supplementary Fig. S1A**)(Nicholson et al., 1997). Peripheral blood mononuclear cells were re-suspended at 2×10^6/ml and seeded at 1ml per 24 well plate and activated with CD3/CD28 Beads. The next day, fresh media was added with IL2 to a final concentration of $100 IU/ml$. Six hours later cells were harvested, counted and re-suspended at $0.6 \times 10^6/ml$ and $0.5ml$ was seeded into a 24 well plate pre-coated with retronectin (**Supplementary Fig. S1C**). $1.5ml$ of lentiviral supernatant was added to each well with fresh IL2 to a final concentration of 100IU/ml and spun for 40 minutes at 1000G. Two days later cells were harvested, re-suspended at $0.5 \times 10^6/ml$ with $50 IU/ml$ of IL2 and left to expand and split every 48 hours. Transduction efficiency was assessed by determining the percentage of GFP+ T cells.

Co-cultures of CD19 CAR-T cells and MEC1 cell line (ATCC), a Chronic Lymphocytic Leukemia (CLL) cell line that constitutively expresses CD19 were established at various effector-to-target ratios and at different time points (**Supplementary Table S3**). Co-cultures were harvested together and stained with hashing antibodies (Biolegend), normalized and prepared for single-cell RNA sequencing on the Chromium 10X platform. Experimental details for each of the 4 experiments found in (**Supplementary Table S3**).

CAR-T Data Preprocessing and Analysis

Each co-culture experiment was hashed by timepoint according to **Supplementary Table S3**. Demultiplexing was done in R as detailed below: All cells with less than 200 detected UMIs were removed as empty droplets. For each cell, hashtag oligos (HTO) were ranked according to detected counts and the standard deviation and mean for HTOs ranked 2-4 were calculated. Cells were removed if the difference between HTO2 and HTO3 were within 1 standard deviation and if HTO2 was within 15% of the mean of HTO2-4. Finally, the top 2 HTOs for each cell were used as cell identifiers and matched based on **Supplementary Table S3**.

All count data across all 4 experiments was aggregated and log-normalized as $(log(0.1+counts))$. Data was visualized by UMAP, based on PCA with 8 components. Clustering was performed on 15 PCA components, using Phenograph (Levine et al., 2015) and a KNN with n=15. 30 clusters were found, and then separated into 4 different metaclusters. MEC1 cells (clusters 1, 7 8, 10, 17, 22, 23, 24) were defined based on differential expression of *IKGC*, *CD79A*, *MS4A1*, and *CD19*. T cells were

defined based on differential expression of *CD3D* and *CD3E*. Exhausted T cells (clusters 0, 6, 9, 11, 15, 21, 26) were defined based on *TIGIT*, *CTLA4*, *PDCD1*, *IL10*, *TGFB1*, *LAG3*, and activated cells defined based on expression of *CD69*, *CD27* and *CD28*. Any *CD8+* cells that did not express activation or exhaustion signatures were grouped into a separate fourth category (clusters 19, 28). The following clusters were removed as doublets: 2, 11, 14, 16, 20. All other clusters were removed due to low library size.

To construct Λ, we first obtained a set of $8,234$ literature-curated receptor-ligand pairs from OmniPath, a database of cell signaling prior knowledge (Türei et al., 2016). For each cell type pair k, k' in the experiment, we quantified interactions as the number of differentially-expressed ligand genes in cell type k with their corresponding differentially-expressed receptor genes in cell type k' at each time point. These interaction count values were averaged across all time points in the experiment and then thresholded to obtain the binary interaction prior matrix Λ, where $\Lambda_{k,k'}$ signifies whether cell types k and k' can interact apriori. The threshold was chosen with a data-driven approach according to the knee-point of sorted values.

CAR-T Data Preprocessing - Leiden comparison

Count data for all CAR-T replicates was aggregated and filtered according to the SCANPY scRNA pre-processing tutorial (`https://scanpy.readthedocs.io/en/stable/tutorials/basics/clustering.html`).Cells expressing <100 genes were removed, and any genes expressed in <3 cells were removed. Scrublet was then used, via scanpy.external.pp.scrublet() function to remove doublets. Data was log normalized. The kNN neighbor graph was built on 20 PCA components and 10 neighbors, and Leiden clustering was then used to define clusters with the resolution set to 0.5. Clusters were then sorted into cell types based on gene expression. T cells were assigned based on expression of *CD3E*, and *CD3D*. Clusters 5 and 6 were designated as Exhausted T cells based on the expression of *TIGIT*, *CTLA4*, *PDCD1*, *IL10*, *TGFB1*, and *LAG3*. Clusters 7, 9 ,17 were assigned as Activated T cells based on expression of *CD69*, *CD27*, and *CD28*. Clusters 0, 3, 4, 11, 14, 15 were designated as MEC1 cells based on expression of *IGKC*, *CD79A*, *MS4A1*, and *CD19*,. Clusters 1, 2, 13, 16 were designated as Other CD8+ T cells, and clusters 8, 10 were labeled as additional doublets and removed due to contradictory expression of both T cell and MEC1 markers. Proportions were calculated by normalizing to total numbers of cells in clusters.

DIISCO Design Parameters

CAR-T Experiments. Hyperparameters used for all co-culture experiments are as follows: $\sigma_F = 0.5$, $\sigma_W = 0.1$, $\sigma_Y = 0.5$, $\tau_F = 6.5$, $\tau_W = 6.5$, $v_F = 1$, $v_W = 1$. Threshold on number of interactions used for prior chosen based on midpoint for histogram of values.

For all CAR-T DIISCO runs, the model was trained with 10,000 iterations and a learning rate of 0.00005, and the `MultivariateNormalFactorized` guide. For predictions, 10,000 time points were sampled, and all confidence intervals were calculated as the region between the 86th and 16th percentile calculated on all predicted time points. Experiment B was excluded in comparison due to differences in Effector:Target ratio compared to all other experiments (**Supplementary Table S3**).

We used Experiment C to further evaluate the sensitivity of DIISCO to the following parameters:

1. Number of cells: Cells were subsampled by randomly removing 10%, 30%, 50%, 70%, and 90% of cells from the original set of Experiment C cells. For each subsample, metacluster proportions were calculated and a separate DIISCO model was trained.

2. Number of timepoints: Separate DIISCO models were trained on the original Experiment C metacluster proportions after randomly removing proportions at 10%, 20%, 30%, 40%, and 50% of the observed timepoints.

3. Clustering resolution: Cluster proportions were calculated for the higher-resolution PhenoGraph clusters instead of metaclusters. A new interaction prior was constructed based on cluster-level R-L expression. A new DIISCO model was trained using the cluster-level proportions and interaction prior.

4. Clustering method: Same as **3**, but cluster proportions were calculated for Leiden clusters instead of PhenoGraph clusters, as described in Section CAR-T Data Preprocessing - Leiden comparison.

R-L Comparison

CellPhoneDB version 2.0.0 was applied to generate interactions in experiment C. To filter predicted interactions, we only reported receptor-ligand pairs that had non-zero counts in the respective sender

and receiver clusters and $p<0.05$. For running CellChat in RStudio, we used the default human CellChat database (R Core Team, 2022; Jin et al., 2021). For both methods, all time points were aggregated to increase statistical power. To compare results between all 3 methods, we focused on the MEC→Exhausted interaction. Spearman's and Pearson's correlations were calculated with the `SciPy` package, and mutual information was calculated using `sklearn` package for the union of all R-L pairs predicted by any of the three methods.

We further ran CellphoneDB on 'early' and 'late' time points grouped together. 'Early' was defined as the first 5 time points and 'late' was defined as the last 4 time points. We then ran cellphoneDB and processed the predicted interactions as mentioned above. To compare to DIISCO, we limited predicted interactions to ones that were in either the 'early' or 'late' set, not both. We then compared the predicted R-L pairs from DIISCO and from these 2 bucketed groups, identifying R-L pairs shared between both methods and ones unique to either method.

Data Access

The DIISCO package and notebooks for generating the figures in this manuscript are publicly accessible at GitHub (`https://github.com/azizilab/DIISCO_public`) and as Supplemental Code. All raw and processed sequencing data generated in this study have been submitted to the NCBI Gene Expression Omnibus (GEO; https://www.ncbi.nlm.nih.gov/geo/) under accession number GSE255888.

Conflict of Interest

C.J.W. is an equity holder of BioNTech and receives research funding from Pharmacyclics.

Acknowledgements

We are thankful to David Blei and Dana Pe'er for helpful feedback and discussions. We are also thankful to Gitte Aasbjerg for help and guidance in demultiplexing the CAR-T data. This work was made possible by support from the National Institute of Health (NIH) NCI grants R00CA230195, P01CA229092, Leukemia & Lymphoma Society grant SCOR-22937-22, and the MacMillan Family and the MacMillan Center for the Study of the Non-Coding Cancer Genome at the New York

Genome Center. C.P. was supported by the Columbia University Kaganov Fellowship. K.M. was supported by the Richard K. Lubin Family Foundation Fellowship.

Contributions

C.Y.P, S.M, E.A conceived and supervised the study. C.Y.P, S.M, N.B, E.A designed and developed DIISCO. N.B designed pyro implementation of DIISCO. C.Y.P, S.M, N.B, K.M, E.A, C.J.W analyzed and interpreted results. S.G, S.L, T.H designed and performed all *in vitro* experiments and sequencing. D.A.K, K.J.L, C.J.W, and E.A offered project guidance and troubleshooting.

References

Aitchison J. 2003. *The statistical analysis of Compositional Data*. Blackburn Press.

Azizi E, Carr AJ, Plitas G, Cornish AE, Konopacki C, Prabhakaran S, Nainys J, Wu K, Kiseliovas V, Setty M, et al.. 2018. Single-cell map of diverse immune phenotypes in the breast tumor microenvironment. *Cell* **174**: 1293–1308.

Bachireddy P, Azizi E, Burdziak C, Nguyen VN, Ennis CS, Maurer K, Park CY, Choo ZN, Li S, Gohil SH, et al.. 2021. Mapping the evolution of t cell states during response and resistance to adoptive cellular therapy. *Cell reports* **37**.

Bingham E, Chen JP, Jankowiak M, Obermeyer F, Pradhan N, Karaletsos T, Singh R, Szerlip P, Horsfall P, and Goodman ND. 2019. Pyro: Deep universal probabilistic programming. *The Journal of Machine Learning Research* **20**: 973–978.

Blei DM, Kucukelbir A, and McAuliffe JD. 2017. Variational inference: A review for statisticians. *Journal of the American statistical Association* **112**: 859–877.

Browaeys R, Saelens W, and Saeys Y. 2020. Nichenet: modeling intercellular communication by linking ligands to target genes. *Nature methods* **17**: 159–162.

Bui TM, Wiesolek HL, and Sumagin R. 2020. Icam-1: A master regulator of cellular responses in inflammation, injury resolution, and tumorigenesis. *Journal of Leucocyte Biology* **108**: 787–799.

Delabie J, Ceuppens J, Vandenberghe P, de Boer M, Coorevits L, and De Wolf-Peeters C. 1993. The b7/bb1 antigen is expressed by reed-sternberg cells of hodgkin's disease and contributes to the stimulating capacity of hodgkin's disease-derived cell lines. *Blood* **82**: 2845–2852.

Dorfman D, Schultze J, Shahsafaei A, Michalak S, Gribben J, Freeman G, Pinkus G, and Nadler L. 1997. In vivo expression of b7-1 and b7-2 by follicular lymphoma cells can prevent induction of t-cell anergy but is insufficient to induce significant t-cell proliferation. *Blood* **90**: 4297–4306.

Dürr C, Hanna B, Schulz A, Lucas F, Zucknick M, Benner A, Clear A, Ohl S, Öztürk S, Zenz T, et al.. 2018. Tumor necrosis factor receptor signaling is a driver of chronic lymphocytic leukemia that can be therapeutically targeted by the flavonoid wogonin. *Haematologica* **103**: 688–697.

Efremova M, Vento-Tormo M, Teichmann SA, and Vento-Tormo R. 2020. Cellphonedb: inferring cell–cell communication from combined expression of multi-subunit ligand–receptor complexes. *Nature protocols* **15**: 1484–1506.

He Y, Vlaming M, van Meerten T, and Bremer E. 2022. The implementation of tnfrsf co-stimulatory domains in car-t cells for optimal functional activity. *Cancers* **14**: 299.

Jin S, Guerrero-Juarez CF, Zhang L, Chang I, Ramos R, Kuan CH, Myung P, Plikus MV, and Nie Q. 2021. Inference and analysis of cell-cell communication using cellchat. *Nature communications* **12**: 1088.

Kingma DP and Ba J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* .

Kingma DP and Welling M. 2022. Auto-encoding variational bayes.

Kumar MP, Du J, Lagoudas G, Jiao Y, Sawyer A, Drummond DC, Lauffenburger DA, and Raue A. 2018. Analysis of single-cell rna-seq identifies cell-cell communication associated with tumor characteristics. *Cell reports* **25**: 1458–1468.

Levine J, Simonds E, Bendall S, Davis K, El-ad D, Tadmor M, Litvin O, Fienberg H, Jager A, Zunder E, et al.. 2015. Data-driven phenotypic dissection of aml reveals progenitor-like cells that correlate with prognosis. *Cell* **162**: 184–197.

Li J, Hubisz MJ, Earlie EM, Duran MA, Hong C, Varela AA, Lettera E, Deyell M, Tavora B, Havel JJ, et al.. 2023. Non-cell-autonomous cancer progression from chromosomal instability. *Nature* pp. 1–9.

Li S, Xing W, Kirby M, and Zhe S. 2020. Scalable variational gaussian process regression networks. *arXiv preprint arXiv:2003.11489* .

Lönnberg T, Svensson V, James KR, Fernandez-Ruiz D, Sebina I, Montandon R, Soon MS, Fogg LG, Nair AS, Liligeto UN, et al.. 2017. Single-cell rna-seq and computational analysis using temporal mixture modeling resolves th1/tfh fate bifurcation in malaria. *Science immunology* **2**: eaal2192.

Maurer K, Park C, Mani S, Borji M, Penter L, Jin Y, Zhang JY, Shin C, Brenner JR, Southard J, et al.. 2024. Coordinated immune cell networks in the bone marrow microenvironment define the graft versus leukemia response with adoptive cellular therapy. *bioRxiv* .

Molica S, Dattilo A, Mannella A, Levato D, and Levato L. 1995. Expression on leukemic cells and serum circulating levels of intercellular adhesion molecule-1 (icam-1) in b-cell chronic lymphocytic leukemia: implications for prognosis. *Leukemia research* **19**: 573–580.

Musolino C, Alonci A, Allegra A, Bellomo G, Spatari G, Pernice F, Squadrito G, Quartarone C, Tringali O, and Quartarone M. 1996. Circulating levels of intercellular adhesion molecule-1 (icam-1) and soluble interleukin 2 receptors (il-2r) in patients with b cell chronic lymphocytic leukaemia. *Rivista Europea per le Scienze Mediche e Farmacologiche= European Review for Medical and Pharmacological Sciences= Revue Europeenne Pour les Sciences Medicales et Pharmacologiques* **18**: 113–118.

Nicholson I, Lenton K, Little D, Decorso T, Lee F, Scott A, Zola H, and Hohmann A. 1997. Construction and characterisation of a functional cd19 specific single chain fv fragment for immunotherapy of b lineage leukaemia and lymphoma. *Molecular Immunology* **34**: 1157–1165.

R Core Team. 2022. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Rasmussen CE and I WCK. 2008. *Gaussian processes for machine learning*. MIT Press.

Sievers C, Craveiro M, Friedman J, Robbins Y, Yang X, Bai K, Nguyen A, Redman JM, Chari R, Soon-Shiong P, et al.. 2023. Phenotypic plasticity and reduced tissue retention of exhausted tumor-infiltrating t cells following neoadjuvant immunotherapy in head and neck cancer. *Cancer Cell* **41**: 887–902.

Stacchini A, Aragno M, Vallario A, Alfarano A, Circosta P, Gottardi D, Faldella A, Rege-Cambrin G, Thunberg U, Nilsson K, et al.. 1999. Mec1 and mec2: two new cell lines derived from b-chronic lymphocytic leukaemia in prolymphocytoid transformation. *Leukemia research* **23**: 127–136.

Tirosh I, Izar B, Prakadan SM, Wadsworth MH, Treacy D, Trombetta JJ, Rotem A, Rodman C, Lian C, Murphy G, et al.. 2016. Dissecting the multicellular ecosystem of metastatic melanoma by single-cell rna-seq. *Science* **352**: 189–196.

Türei D, Korcsmáros T, and Saez-Rodriguez J. 2016. Omnipath: guidelines and gateway for literature-curated signaling pathway resources. *Nature Methods* **13**: 966–967.

Vyth-Dreese F, Boot H, Dellemijn T, Majoor D, Oomen L, Laman J, Van Meurs M, De Weger R, and De Jong D. 1998. Localization in situ of costimulatory molecules and cytokines in b-cell non-hodgkin's lymphoma. *Immunology* **94**: 580–586.

Wilk AJ, Rustagi A, Zhao NQ, Roque J, Martínez-Colón GJ, McKechnie JL, Ivison GT, Ranganath T, Vergara R, Hollis T, et al.. 2020. A single-cell atlas of the peripheral immune response in patients with severe covid-19. *Nature medicine* **26**: 1070–1076.

Wilson AG, Knowles DA, and Ghahramani Z. 2011. Gaussian process regression networks. *arXiv preprint arXiv:1110.4411* .

Wolf FA, Angerer P, and Theis FJ. 2018. Scanpy: large-scale single-cell gene expression data analysis. *Genome biology* **19**: 1–5.

Yofe I, Dahan R, and Amit I. 2020. Single-cell genomic approaches for developing the next generation of immunotherapies. *Nature medicine* **26**: 171–177.

Figure Captions

Figure 1: Overview of DIISCO framework. **(A)** General workflow of the DIISCO algorithm including inputs and outputs. Cell type proportions are computed from scRNA-seq data in each time-point. Expression of R-L complexes is incorporated to obtain time-resolved interactions between cell types. **(B)** Network diagram of model framework. **(C)** Graphical model of DIISCO, including all hyperparameters. Latent variables are depicted as white circles, and observed variables are depicted as yellow-shaded circles. **(D)** Process for incorporating domain knowledge on R-L interactions as a prior for cell-cell interactions into DIISCO.

Figure 2: Simulated data and results. **(A)** Simulated data for 5 cell types generated using the process described in section 3.1 with noise parameter set to 0.1 and 40 sampled timepoints. Non-uniform sampling of timepoints increases complexity. Dots represent datapoints for each of the 5 cell types. Orange line is the average learned DIISCO dynamics after fitting the data over 1000 samplings. Shaded region shows 95% confidence intervals. No normalization was used. **(B)** Interactions inferred by DIISCO between the cell types shown in (a). Ground truth W interactions are shown in dotted lines while DIISCO predictions are solid lines. Only dynamic interactions with $W(t) > 0.1$ in at least one time point t for either true W or learned W are shown. **(C)** Same as in **A**, except noise parameter set to 0.5. **(D)** Same as in **B** except learned interactions for the noisier dataset shown in (C). **(E)** Method robustness to number of timepoints: R^2 calculated between inferred and ground-truth $W(t)$. Mean and SD across 5 iterations are shown. **(F)** Runtime for varying number of clusters and timepoints. **(G)** Table demonstrating DIISCO performance compared to Linear Model (LM) and Rolling Linear Model (RLM) as described in **Methods** for 40 timepoints. Metrics for different numbers of time points from 10-90 can be found in Supplementary Tables 1 and 2. Each model was tested over 10 iterations and mean and standard deviation are shown for each metric across all iterations. Model acronyms denote the following: LM-PRIOR = Linear Model with prior. LM = Linear Model. RLM-PRIOR = Rolling Linear Model with prior. RLM = Rolling Linear Model. Model details can be found in **Methods**. Comparison metrics used are as follows: R^2_Y, R^2_W: R^2 value comparing predictions to ground truth for dynamics (Y) or interactions (W). Higher is better. $RMSE\_Y$, $RMSE\_W$: Root mean squared error for dynamics

(Y) or interactions (W). Lower is better. AUC: Area under ROC curve. Higher is better. AUPRC: Area under Precision-Recall curve. Higher is better. F1: Max F1 score. Higher is better. AUC, AUPRC, and F1 scores calculated comparing predicted interactions to ground truth interactions.

Figure 3: CAR-T experimental setup and data preprocessing. **(A)**. Overview of data collection process and analysis of co-culture experiments outlined in **Supplementary Table 1**). **(B-D)** 2D UMAP projection of 9082 single-cell transcriptomes from Experiment C colored by time point (B), cluster and cell type assignment (C), and average expression of different gene set markers (D) used to define metaclusters. **(E)** Heatmap of average cluster expression of individual genes used in gene set analysis, Z-scored across all clusters. **(F)** Prior constructed for DIISCO input, generated as explained in **Methods**.

Figure 4: DIISCO performance on CAR-T Experiment C. (A) Temporal dynamics of cell types. Data points represent the measured proportion of cell types over time, and DIISCO output is shown as inferred mean (solid line) and 85% confidence intervals (shaded area) of y for each cell type. Additional experiments are shown in **Supplementary Fig. 2**. (**B-C**) Inferred W interaction matrix at two example time points (B) and over the entire co-culture time window (C). Network node size reflects cell type proportion at that timepoint; arrows and their width designate the direction and strength of inferred interactions respectively. (**D**) Average interaction strength between clusters as a function of varying prior thresholds. The threshold indicates the minimum number of R-L interactions between two cluster pairs needed to denote a 1 in the prior matrix. (**E**) Rank of average interaction strengths as a function of varying prior thresholds. (**F**) Average interaction score with downsampling total cell number. (**G-H**) W interaction over time between Exhausted T cell and MEC1 cells (G) and Exhausted-Activated T cells (H) for different percentages of dropped cells.

Figure 5: Comparing predicted R-L pairs with existing methods (A) R-L pairs highly correlated with DIISCO predicted $W_{Exhausted \leftrightarrow MEC}$ interaction. Expression of ligands are shown in blue, receptors in orange (left axis), and interaction strength in red (right axis). **(B)** Comparison of DIISCO predicted R-L pairs with those predicted using Cellchat (CC) or CellphoneDB (CDB).

Correlations between paired R-Ls shown for the $W_{MEC \rightarrow ExhaustedTcell}$ link. p-values calculated with MWU test. (**C**) Evaluating temporal correlation of predicted RLs using each method: DIISCO, CellphoneDB (CDB) and Cellchat (CC). A + B refer to RL pairs predicted by method A and B. Table shows 3 different metrics calculated for all R-L pairs: Pearson's correlation, mutual information (MI), and Spearman's correlation. Mean values +/- SD are quantified, with the top 3 highest scores in **bold** and the highest score underlined.

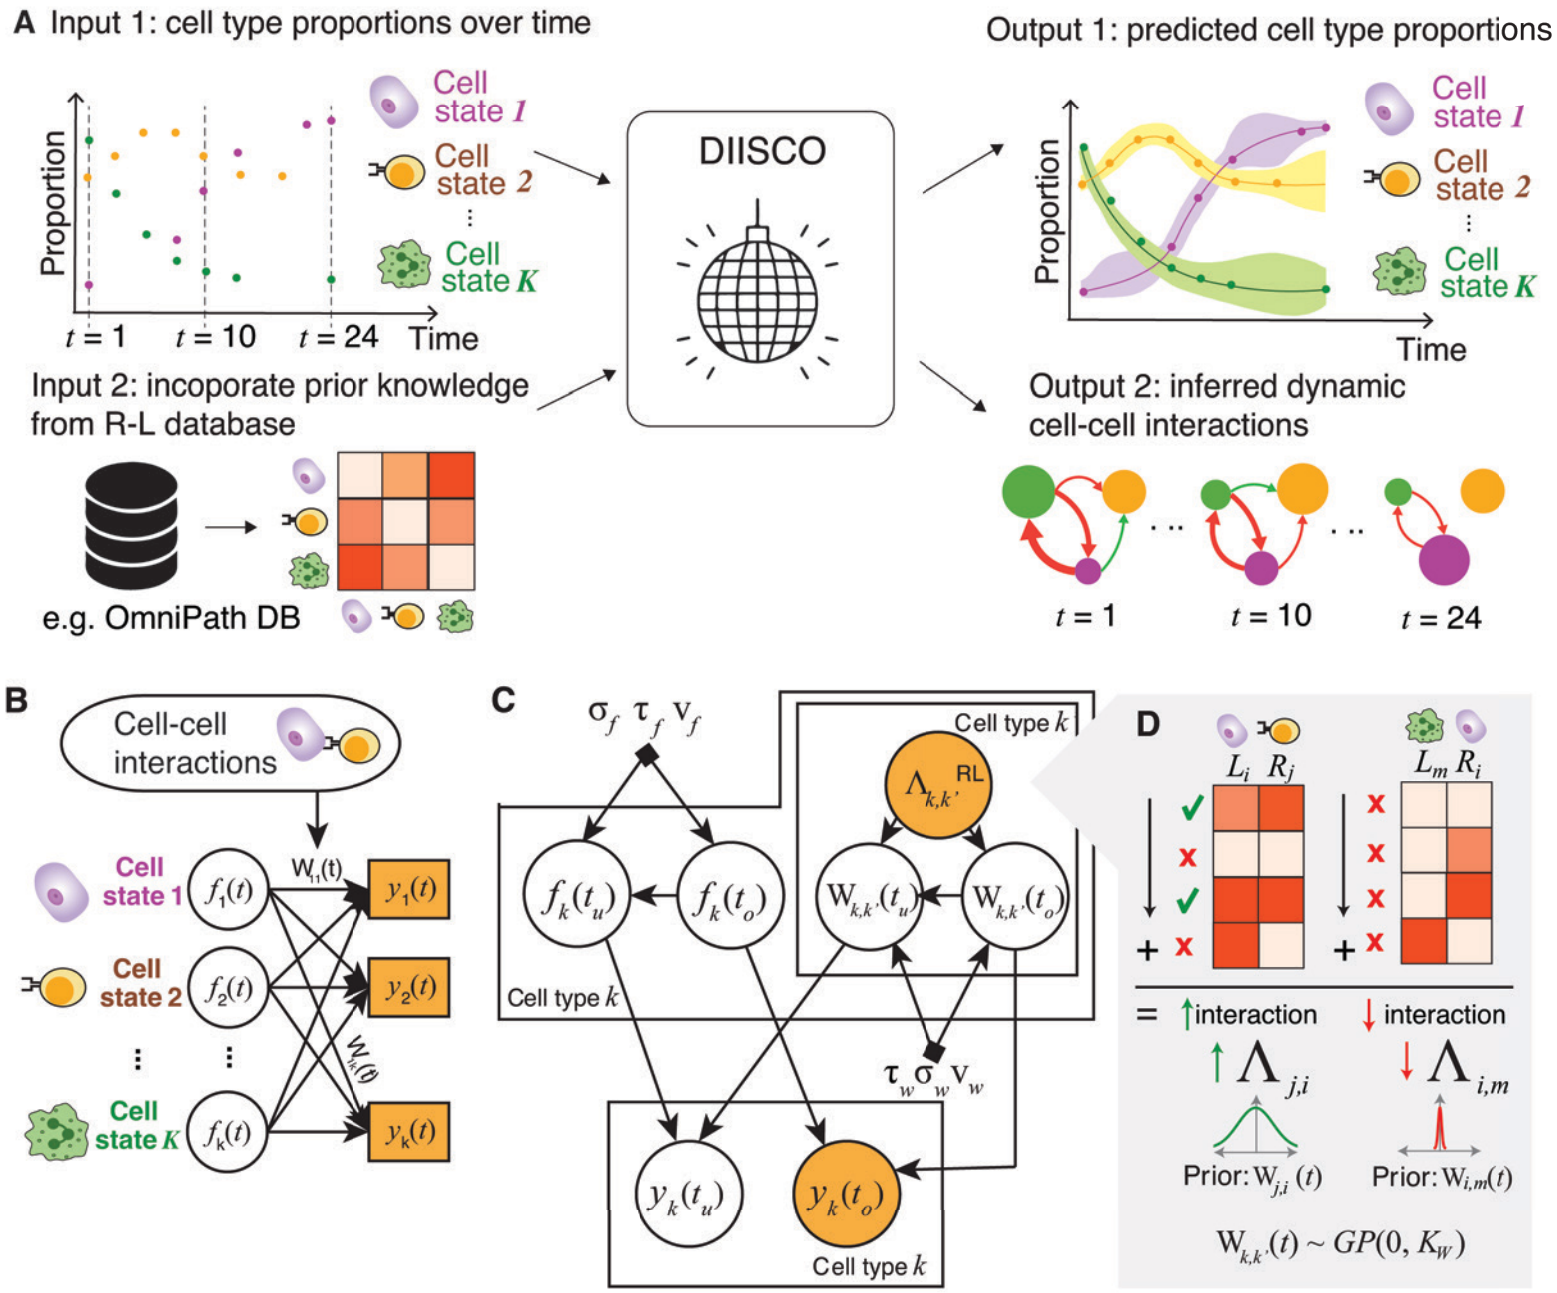
A Input 1: cell type proportions over time
Proportion
Time
t = 1
t = 10
t = 24
Cell state 1
Cell state 2
Cell state K
Input 2: incoporate prior knowledge from R-L database
e.g. OmniPath DB
DIISCO
Output 1: predicted cell type proportions
Output 2: inferred dynamic cell-cell interactions
B
Cell-cell interactions
C
Cell type k
D
interaction
Prior: W_{j,i}(t)
Prior: W_{i,m}(t)
W_{k,k'}(t) ~ GP(0, K_W)

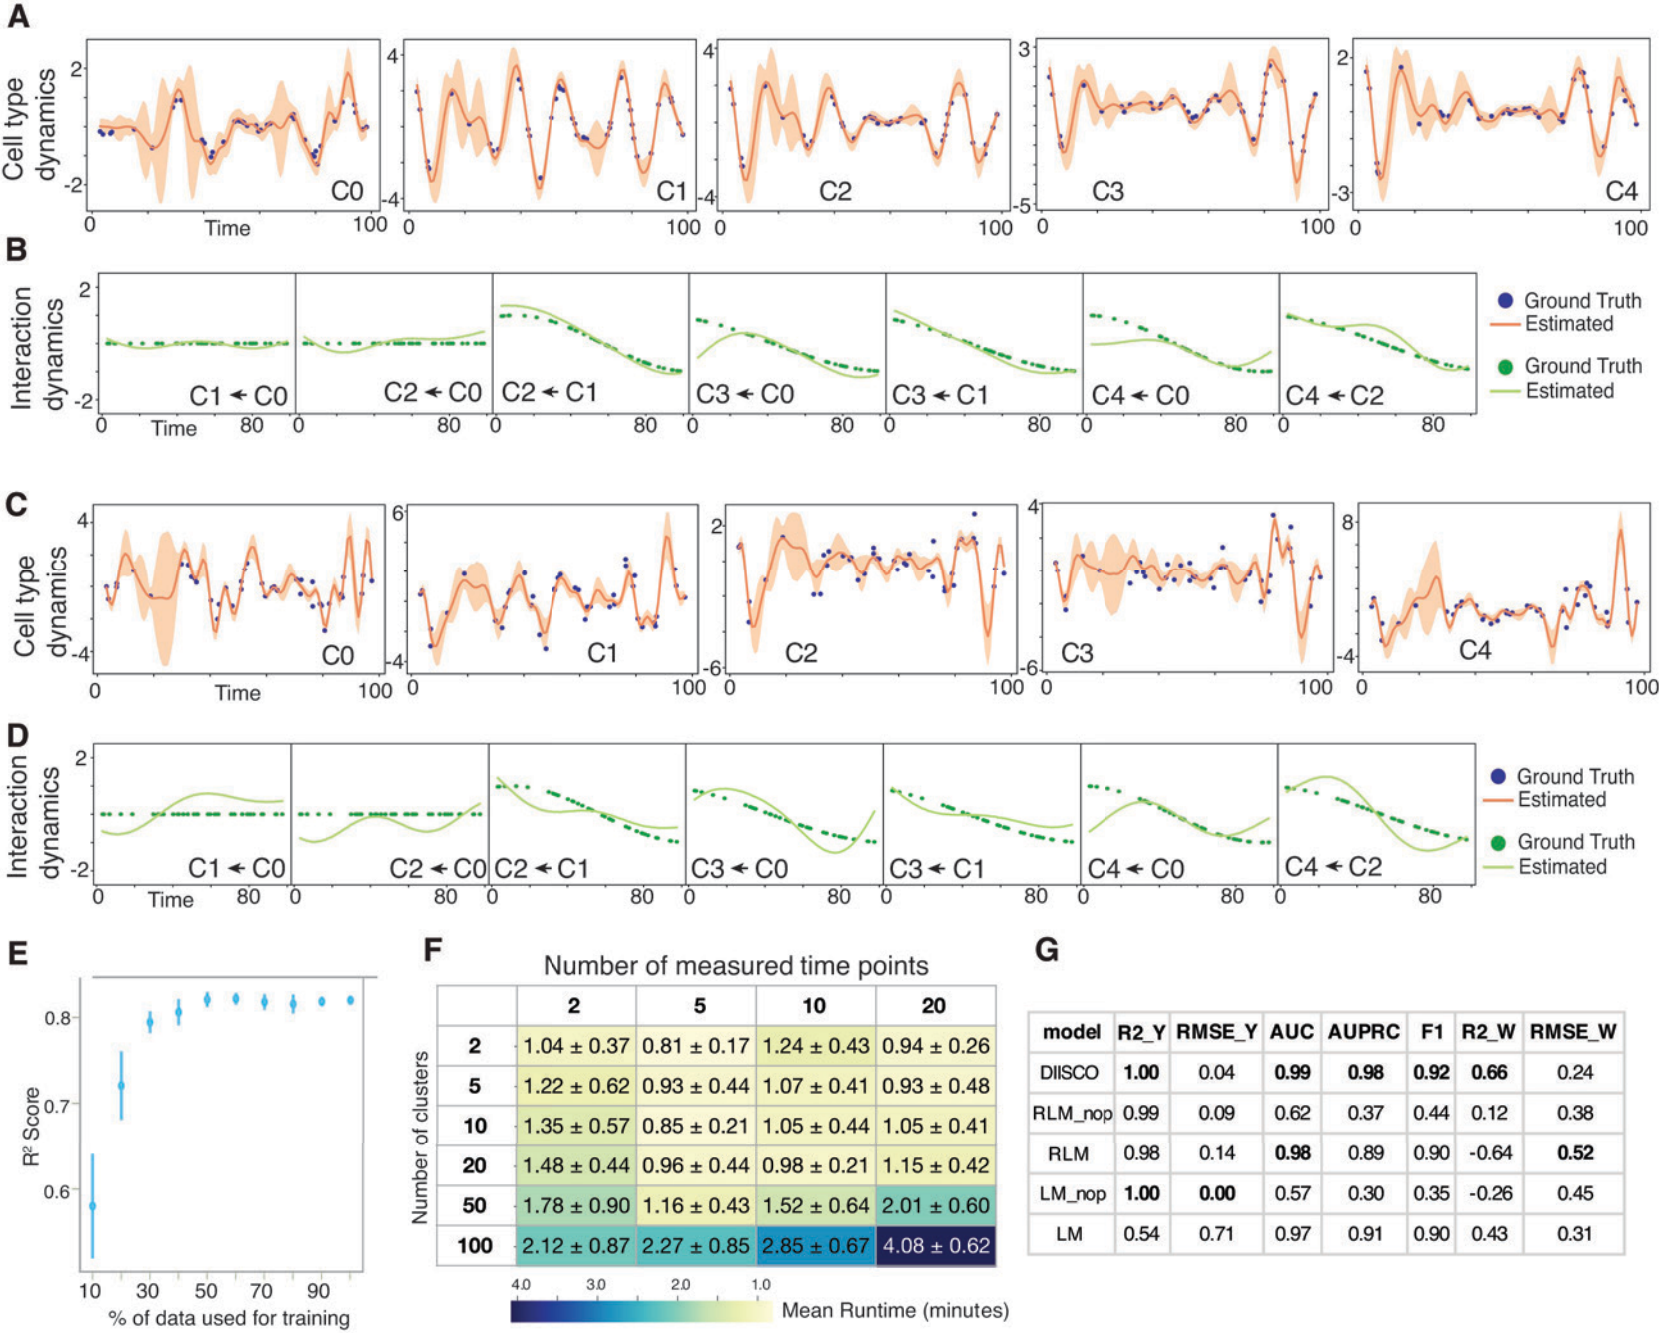

	2	5	10	20
2	1.04 ± 0.37	0.81 ± 0.17	1.24 ± 0.43	0.94 ± 0.26
5	1.22 ± 0.62	0.93 ± 0.44	1.07 ± 0.41	0.93 ± 0.48
10	1.35 ± 0.57	0.85 ± 0.21	1.05 ± 0.44	1.05 ± 0.41
20	1.48 ± 0.44	0.96 ± 0.44	0.98 ± 0.21	1.15 ± 0.42
50	1.78 ± 0.90	1.16 ± 0.43	1.52 ± 0.64	2.01 ± 0.60
100	2.12 ± 0.87	2.27 ± 0.85	2.85 ± 0.67	4.08 ± 0.62

G

model	R2_Y	RMSE_Y	AUC	AUPRC	F1	R2_W	RMSE_W
DIISCO	**1.00**	0.04	**0.99**	**0.98**	**0.92**	**0.66**	0.24
RLM_nop	0.99	0.09	0.62	0.37	0.44	0.12	0.38
RLM	0.98	0.14	**0.98**	0.89	0.90	-0.64	**0.52**
LM_nop	**1.00**	**0.00**	0.57	0.30	0.35	-0.26	0.45
LM	0.54	0.71	0.97	0.91	0.90	0.43	0.31

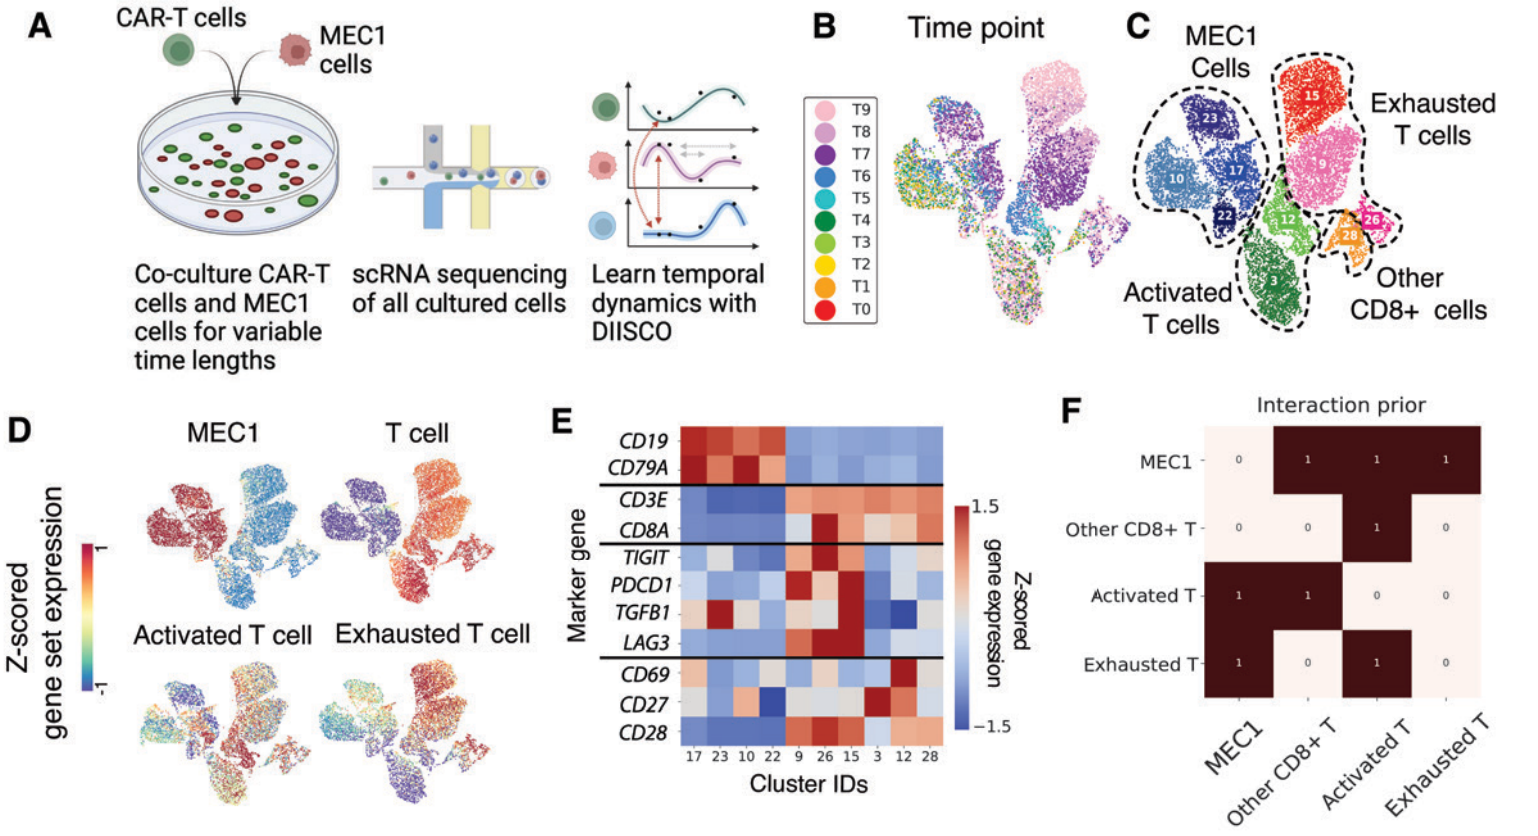
A
CAR-T cells
MEC1 cells
Co-culture CAR-T cells and MEC1 cells for variable time lengths
scRNA sequencing of all cultured cells
Learn temporal dynamics with DIISCO
B
Time point
T9
T8
T7
T6
T5
T4
T3
T2
T1
T0
C
MEC1 Cells
Exhausted T cells
Activated T cells
Other CD8+ cells
23
10
17
22
15
9
12
3
26
28
D
MEC1
T cell
Activated T cell
Exhausted T cell
Z-scored gene set expression
1
-1
E
Marker gene
CD19
CD79A
CD3E
CD8A
TIGIT
PDCD1
TGFB1
LAG3
CD69
CD27
CD28
17 23 10 22 9 26 15 3 12 28
Cluster IDs
1.5
−1.5
Z-scored gene expression
F
Interaction prior
MEC1
Other CD8+ T
Activated T
Exhausted T
0 1 1 1
0 0 1 0
1 1 0 0
1 0 1 0
MEC1
Other CD8+ T
Activated T
Exhausted T

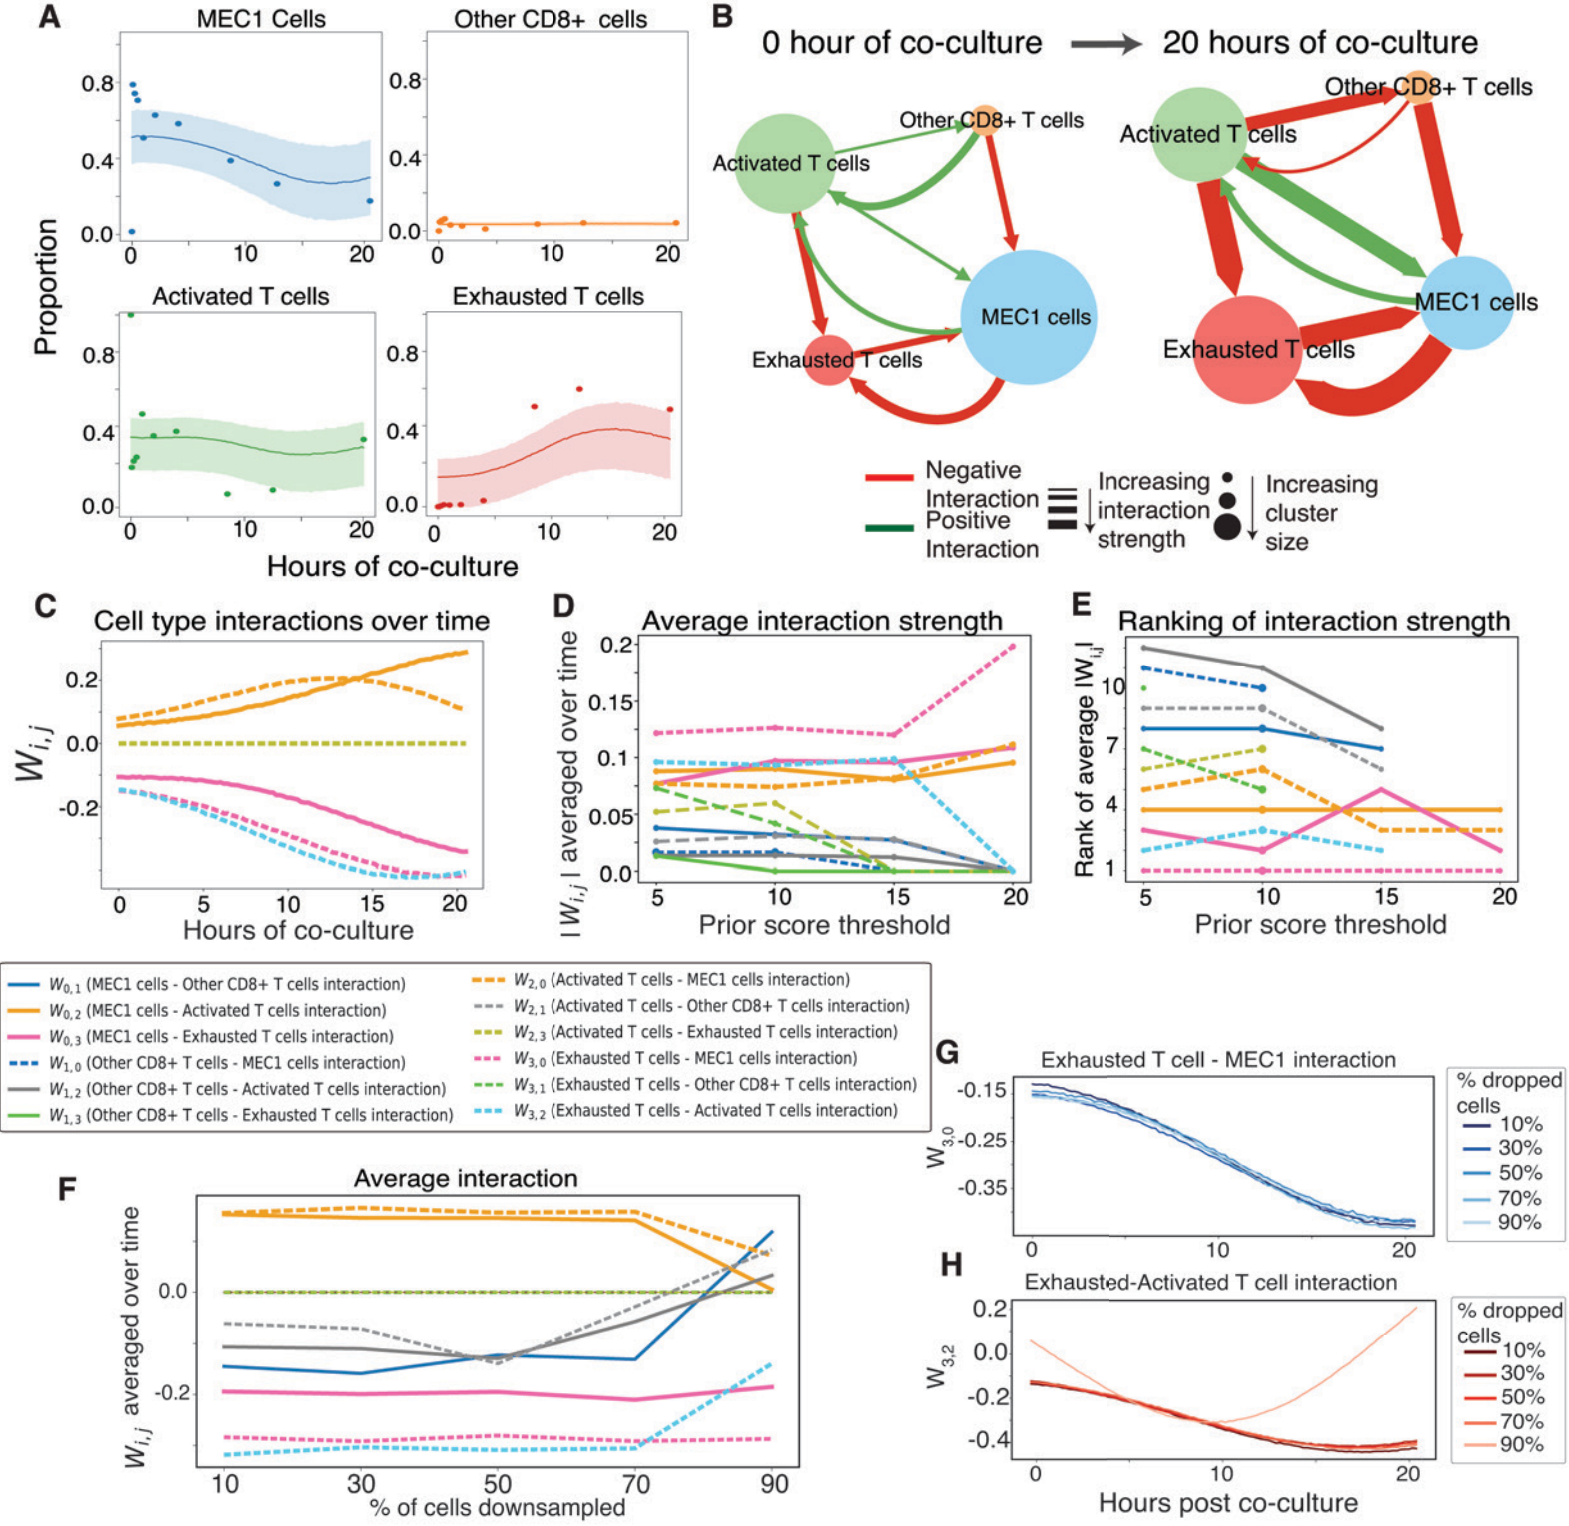
A
MEC1 Cells
Other CD8+ cells
Activated T cells
Exhausted T cells
Proportion
Hours of co-culture
B
0 hour of co-culture
20 hours of co-culture
Other CD8+ T cells
Activated T cells
MEC1 cells
Exhausted T cells
Negative Interaction
Positive Interaction
Increasing interaction strength
Increasing cluster size
C
Cell type interactions over time
$W_{i,j}$
Hours of co-culture
D
Average interaction strength
$|W_{i,j}|$ averaged over time
Prior score threshold
E
Ranking of interaction strength
Rank of average $|W_{i,j}|$
Prior score threshold
$W_{0,1}$ (MEC1 cells - Other CD8+ T cells interaction)
$W_{0,2}$ (MEC1 cells - Activated T cells interaction)
$W_{0,3}$ (MEC1 cells - Exhausted T cells interaction)
$W_{1,0}$ (Other CD8+ T cells - MEC1 cells interaction)
$W_{1,2}$ (Other CD8+ T cells - Activated T cells interaction)
$W_{1,3}$ (Other CD8+ T cells - Exhausted T cells interaction)
$W_{2,0}$ (Activated T cells - MEC1 cells interaction)
$W_{2,1}$ (Activated T cells - Other CD8+ T cells interaction)
$W_{2,3}$ (Activated T cells - Exhausted T cells interaction)
$W_{3,0}$ (Exhausted T cells - MEC1 cells interaction)
$W_{3,1}$ (Exhausted T cells - Other CD8+ T cells interaction)
$W_{3,2}$ (Exhausted T cells - Activated T cells interaction)
F
Average interaction
$W_{i,j}$ averaged over time
% of cells downsampled
G
Exhausted T cell - MEC1 interaction
$W_{3,0}$
% dropped cells
10%
30%
50%
70%
90%
H
Exhausted-Activated T cell interaction
$W_{3,2}$
% dropped cells
10%
30%
50%
70%
90%
Hours post co-culture

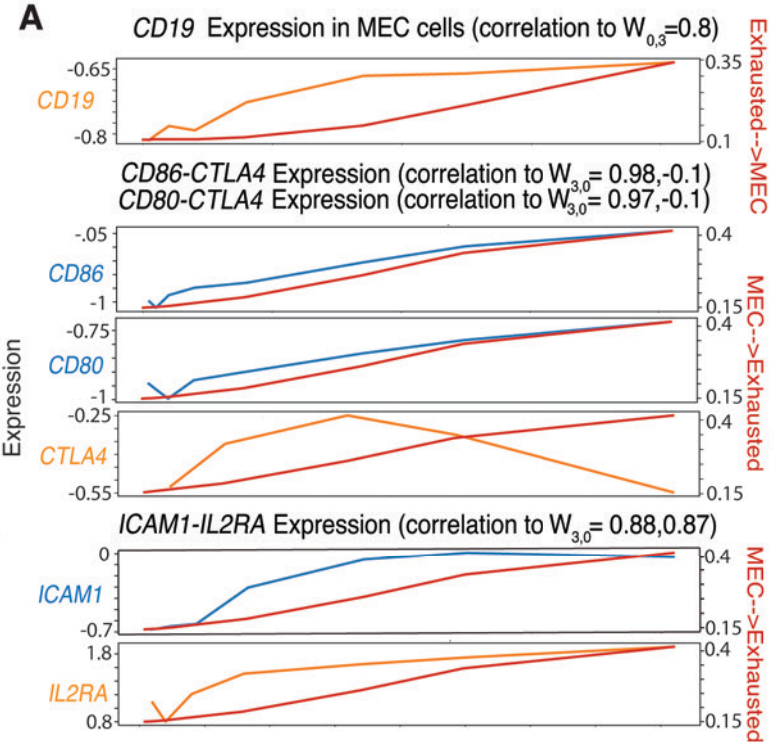

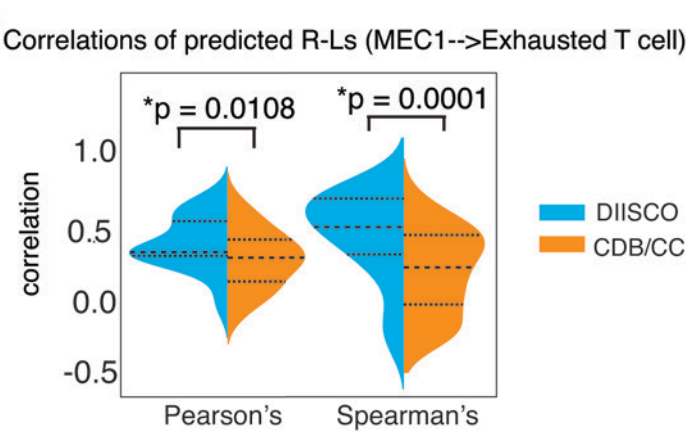

C

	Pearson's ± std	MI ± std	Spearman's ± std
CC	0.356 +- 0.133	**0.186 +- 0.099**	0.042 +- 0.333
CC + CDB	0.136 +- 0.173	0.067 +- 0.09	0.125 +- 0.114
CDB	0.345 +- 0.232	**0.112 +- 0.112**	0.339 +- 0.253
DIISCO	**0.462 +- 0.197**	0.083 +- 0.114	**0.39 +- 0.317**
DIISCO + CC	**0.381 +- 0.186**	**0.252 +- 0.082**	**0.641 +- 0.303**
DIISCO + CDB	**0.569 +- 0.201**	0.073 +- 0.084	**0.482 +- 0.265**