LOW-RANK INTERPRETABLE CELL-CELL HIDDEN INTERACTIONS FROM EMBEDDINGS

Anonymous authors
Paper under double-blind review

ABSTRACT

Multicellular organisms rely on continuously changing cell–cell interactions that govern critical biological processes as cells modify their internal states and trajectories in space over time. Studying these interactions is critical to understand development, homeostasis, and disease progression. Live-cell imaging provides a unique opportunity to directly observe these dynamical events; however, current computational approaches often fail to model complex, time-varying events involving diverse populations and spatial contexts. Here, we present `LICCHIE`, a model designed to infer time-changing, feature-based cell-cell interactions, applicable across systems and conditions. Our approach represents each cell with a dynamic multi-feature vector, and interactions are modeled as spatially constrained, directed influences between cell pairs, evolving over time. We optimize the model using an iterative scheme balancing data fidelity, interactions smoothness, and low-rank sparse structure. We validated `LICCHIE`'s ability to capture cellular interactions across populations in a controlled synthetic setting, and applied it to real-world 3D live-cell imaging of patient-derived tumor organoids to (1) identify components with interpretable structures that capture interaction type and directionality, and (2) suggest modulation strategies that may accelerate NK polarization and tumor cell death.

1 INTRODUCTION

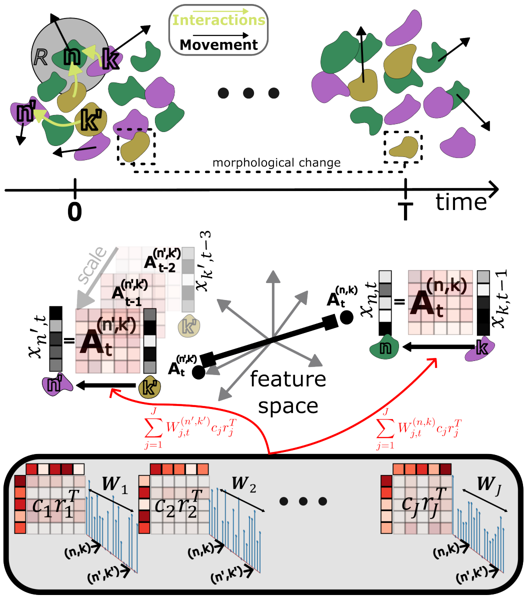

Figure 1: *`LICCHIE` learns cell-cell interactions from temporal trajectories.* **Top**, Cells from diverse populations can adapt their morphology, internal state, and location in response to cross-cell interactions and environmental effects. **Middle**, Interaction matrices, $\{\boldsymbol{A}_t^{(n,k)}\}$, representing the effect of cell k on cell n at time t, capture pairwise time-changing linear effects via cell features, constrained on feature-space distance for smooth, interpretable changes. **Bottom**, While the underlying matrices $\boldsymbol{A}_t^{(n,k)}$ may vary in space, time, and cell pairs, they emerge from a set of fixed rank-1 processes representing source $\{\boldsymbol{c}_j\}$ and target $\{\boldsymbol{r}_j\}$ effects. These components are modulated via flexible weights to shape the overall transition matrix $\boldsymbol{A}_t^{(n,k)} = \sum_{j=1}^{J} W_{j,t}^{(n,k)} \boldsymbol{c}_j \boldsymbol{r}_j^T$, capturing complexity while maintaining parsimony.

Cell–cell interactions are central to the function and organization of multicellular systems (Su et al., 2024). Disruption or mis-regulation of communication is implicated in cancer, autoimmune disease, and developmental disorders (Armingol et al., 2021; Liu et al., 2024). Capturing how one cell's state and behavior influence another in real time is challenging since interactions span spatial and temporal scales, and depend on cell type and state.

Studies based on static snapshots reveal a partial story. Single-cell sequencing technologies have enabled systematic inference of putative ligand–receptor activity and co-variation patterns, offering a powerful view of communication at scale (Wilk et al., 2024; Heumos et al., 2023). However, such data are typically dissociated and time-agnostic, and methods built on them suffer from an inherent limitation: they recover statistical associations from static measurements rather than the dynamics of how cells influence each other in real time (Armingol et al., 2024; 2021; Wagner et al., 2016). As a result, they fail to reveal how cells change as a result of interactions and struggle to expose directionality, temporal delays, or state-dependent effects that are essential to mechanistic understanding (Weinreb et al., 2018).

Live-cell imaging: a closer observational proxy for cellular dynamics. Live-cell imaging provides a means to overcome many of the challenges in studying cell–cell interactions. By tracking individual cells over time, such data allow observation of cell states before, during and after interactions, and thus offer a window into the dynamic consequences of cellular communication (Cuny et al., 2022). Such longitudinal data provides ground truth for understanding how cells respond to given interactions and interventions, allowing for advanced therapeutic applications (Alieva et al., 2023). Hence, there is a need for novel computational models that can: (1) detect interactions, (2) quantify their strength, and (3) reveal how they promote changes in cell state.

Our contributions. We introduce **LICCHIE**—**L**ow-rank, **I**nterpretable **C**ell-**C**ell **H**idden **I**nteractions from **E**mbeddings—a data-driven method to infer time-varying, feature-based interaction rules from live-cell imaging. We represent each cell with a dynamic feature vector encompassing spatial and internal states. Direct interactions between source–target cell pairs are modeled as temporally-evolving linear transitions over features, active only within a spatial radius between cells and regularized to vary smoothly within the feature space. To balance expressivity with interpretability, `LICCHIE` decomposes pairwise interaction matrices into a few rank-1 components, with time- and cell-pair-varying weights capturing context, while keeping the components biologically readable. We jointly estimate the interaction matrices, shared components, and sparse weights in an iterative approach (Fig. 1). More precisely, we present:

- A modeling framework that detects direct, time- and cell-varying interactions at the feature level (`LICCHIE`).
- A low-rank decomposition into shared, interpretable source/target motifs with sparse, context-varying weights.
- Validation on synthetic data and 3D tumor–NK imaging, recovering time- and cell-varying interactions with interpretable structure.
- Identification of interactions that can influence NK polarization and tumor death using engineered observations, capturing unseen data points.

2 BACKGROUND

Representation learning and static graphs. Deep learning architectures, such as autoencoders have been used to discover interaction-relevant structures in single-cell omics but often yield embeddings that are hard to interpret mechanistically for interactions (Alessandri et al., 2021; Ternes et al., 2022). Graph neural networks encode cells as nodes and potential interactions as edges, yet typically presume a predefined (often symmetric) graph and still obscure direct, dynamic effects; attention models can highlight dependencies but do not directly encode interaction signals (Lazaros et al., 2024; Tang et al., 2023). Together, these tools are not tailored to direct, time-varying, feature-level effects required for dynamic inferences.

Latent dynamical systems. In neuroscience, switching linear dynamical systems capture abrupt regime changes (Linderman et al., 2017; Nassar et al., 2018), and decomposed LDS variants capture smoothly time-evolving interactions (Mudrik et al., 2024; Chen et al., 2024). However, these models do not directly address live-cell specifics: distance-modulated effects, transient visibility, strong state changes (e.g., death/polarization), division/proliferation, and motility that reshapes neighborhoods—motivating an imaging-first, interpretable model at per-cell feature resolution.

Requirements for an imaging-based interpretable model. We seek a framework that (i) models direct, feature-level influences between individual cells; (ii) captures non-stationarities over time and allows for other sources of interaction variability (e.g., distances, temporal resolutions), (iii) is spatially constrained and smooth in feature space for robustness; (iv) uncovers the core interpretable structures underlying the dynamics to support biological understanding of the interactions; and (v) does not rely on persistent cell identities, by operating directly in feature space. `LICCHIE` is designed around exactly these requirements, as summarized in Figure 1 and detailed in Section 4.

3 PROBLEM FORMULATION

Observations and notation. We observe a system of N interacting cells over T frames, where each cell $n = 1 \dots N$ is defined via a position vector (in 2D/3D space) $\{\boldsymbol{\psi}_t^{(n)}\}_{t=1}^{T} \in \mathbb{R}^s$, and a multi-feature, time-varying vector $x_t^{(n)} \in \mathbb{R}^{m^{p(n)}}$ capturing cell-specific features, defined by the biological system observed. A system can consist of one or more populations, such that $p(n)$ indexes the population type of cell n; therefore, $m^{p^{(n)}}$ represents the number of features for population $p^{(n)}$ of cell n (complete notation provided in Table S1).

Feature space learning. Cellular features are transiently measurable; for example, certain cells may only be visible in a subset of frames due to occlusion, movement, or other limitations. Hence, it is desirable to avoid tracking individual cells but instead operate in cell feature space, learning how features influence interaction dynamics regardless of specific cell identity.

Locality of interactions. Interactions are assumed to be local in space; hence we restrict attention to a radius $d_{n,n'}(t) \leq R$ (Fig. 1 top).

Learning objective. We seek to recover biologically meaningful, time-varying interaction rules from transient multi-population observations, avoiding fixed cell identities, towards revealing the mechanisms through which cellular networks govern biological functions and disease processes.

4 THE `LICCHIE` APPROACH

Feature-evolution model. Looking to identify how multi-cellular interactions drive changes in cell-states, we capturing interactions via feature evolution–modeling the state vector $\boldsymbol{x}_t^{(n)}$ of cell n at time t as the weighted sum of learnable interaction matrices capturing influences from previous time-point, namely,

$$\boldsymbol{x}_t^{(n)} = \sum_{k=1}^{N} \mathbf{1}_{d_{n,k}(t) \leq R} d_{n,k}(t) \boldsymbol{A}_t^{(n,k)} \boldsymbol{x}_{t-1}^{(k)}. \tag{1}$$

For clarity of presentation, we hereafter omit the distance-driven reweighting. Notably, in the above formulation feature-level effects are captured by linear relations. This approach was deliberately chosen, prioritizing direct interpretability of the results–deriving understandable interaction rules over model complexity, and justified via empirical validations (App. A.1).

The interaction matrices. A matrix $\boldsymbol{A}_t^{(n,k)} \in \mathbb{R}^{m \times m}$ represents the effect of source cell k on target cell n at time t, conditioned on their previous states. An entry $\left[\boldsymbol{A}_t^{(n,k)}\right]_{i,j}$ captures the influence of feature j of cell k on feature i of cell n, from $t-1$ to t (e.g., how much the size, i.e., feature i of cell n at time t, was impacted by the position, i.e., feature j, of cell k, at time $t-1$,

Fig. 1 middle). Here, we assume that all cells share the same feature set, yet the model naturally extends to population-specific features (App. A.2).

Cross-interaction similarity. Naturally, pairs with similar stacked state vectors induce similar interaction maps. With that, if the source vector $\boldsymbol{x}_t^{(k)}$ lies in the null space of a matrix $\widetilde{\boldsymbol{A}}$, then $\boldsymbol{A}_t^{(n,k)} + \widetilde{\boldsymbol{A}}$ is an equivalent solution to $\boldsymbol{A}_t^{(n,k)}$ in terms of fit—thus applying regularization can improve inference by enforcing consistency and smoothness across parameter changes (App. A.3). We employ this by constraining distances between $\boldsymbol{A}_t^{(n,k)}$ with the respective distance in feature space,

$$\|\boldsymbol{A}_t^{(n,k)} - \boldsymbol{A}_t^{(n',k')}\|_2^2 \le \epsilon\big(\delta_{\boldsymbol{A}_t^{(n,k)},\boldsymbol{A}_t^{(n',k')}}\big), \tag{2}$$

where $\varepsilon(\cdot)$ decreases with the distance $\delta_{\boldsymbol{A}_t^{(n,k)},\boldsymbol{A}_t^{(n',k')}}(t)$ between the stacked state vectors of the two pairs at time t. In practice implemented as a *soft, distance-weighted penalty* on interaction matrices distances.

Low-rank interaction-features regularization. To reflect a limited number of shared biological processes, we impose a low-rank structure by expressing each $\boldsymbol{A}_t^{(n,k)}$ via a restricted set of global rank-1 components with pair-/time-specific weights,

$$\boldsymbol{A}_t^{(n,k)} = \sum_{j=1}^{J} W_{j,t}^{(n,k)} \boldsymbol{c}_j \boldsymbol{r}_j^\top = \sum_{j=1}^{J} W_{j,t}^{(n,k)} \underbrace{\boldsymbol{M}_j}_{\boldsymbol{c}_j \boldsymbol{r}_j^\top} \tag{3}$$

where $\{\boldsymbol{c}_j\}_{j=1}^J$ and $\{\boldsymbol{r}_j\}_{j=1}^J$ are the set of underlying components such that each $\boldsymbol{c}_j \in \mathbb{R}^m$ and $\boldsymbol{r}_j \in \mathbb{R}^m$, spans each $\{\boldsymbol{A}_t^{(n,k)}\}_{t,n,k}$, and each term $\boldsymbol{M}_j := \boldsymbol{c}_j \boldsymbol{r}_j^\top$ is a rank-1 matrix constructed via the outer product of two m-dimensional vectors. Each $\boldsymbol{W}^{(n,k)} \in \mathbb{R}^{J \times T}$ is the corresponding sparse per time-point weight matrix, such that $W_{j,t}^{(n,k)}$ is the contribution of components $\boldsymbol{c}_j, \boldsymbol{r}_j$ to the interaction from k to n at time t (Fig. 1 bottom).

The multi-feature vector. While the set of observed features is defined by the data in hand, the induced multi-feature vector, $x_t^{(n)} \in \mathbb{R}^{m_{p(n)}}$, can be defined based on the task in hand and vary from engineered, pre-defined, interpretable features based on observed attributes (e.g. spatial, kinematic and available phenotypic data, Stirling et al. (2021)) to a latent space embedding of the observation (e.g. a deep learning representation, Moshkov et al. (2024)). In what follows, we focus on the former, as it allows for direct interpretation of observed feature relations. Importantly, when considering the latter, the setting can uncover dependencies and importance of different latent dimensions, desirable for explainability of AI models (Saranya & Subhashini, 2023).

Learning low-rank interactions via optimization. Balancing fidelity with the assumptions above, we solve:

$$\begin{aligned}
\mathcal{L} = \sum_{n=1}^{N}\sum_{k=1}^{N} \Bigg[& \|\boldsymbol{x}_t^{(n)} - (\mathbf{1}_{d_{n,k}(t) \le R}) \boldsymbol{A}_t^{(n,k)} \boldsymbol{x}_{t-1}^{(k)}\|_2^2 + \lambda_1 \|\boldsymbol{A}_t^{(n,k)} - \sum_{j=1}^{J} W_{j,t}^{(n,k)} \boldsymbol{c}_j \boldsymbol{r}_j^T\|_F^2 \\
& + \lambda_2 \sum_{\substack{\boldsymbol{A}' \in \{\boldsymbol{A}_{t'}^{(n',k')}\}_{n',k',t'} \\ \boldsymbol{A}' \neq \boldsymbol{A}_t^{(n,k)}}} \Big[\|(\mathbf{1}_{d_{n,k}(t) \le R}) e^{-\sigma \delta(\boldsymbol{A}_t^{(n,k)}, \boldsymbol{A}')} (\boldsymbol{A}_t^{(n,k)} - \boldsymbol{A}')\|_F^2 \Big] \\
& + \lambda_3 \|\mathrm{vec}(\boldsymbol{A}_t^{(n,k)})\|_1 + \lambda_4 \|\boldsymbol{W}_{:,t}^{(n,k)}\|_1 \Bigg],
\end{aligned} \tag{4}$$

where R is the interaction radius, σ is a parameter controlling the decay of similarity with distance in feature space, $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ are scalar regularization weights that balance the different terms. $\boldsymbol{A}_t^{(n,k)}, \boldsymbol{A}'$ denote interaction matrices corresponding to feature pairs such that $\delta(\boldsymbol{A}_t^{(n.k)}, \boldsymbol{A}')$ is the distance between these pairs in the feature space. The first term ensures fidelity to the observed dynamics by minimizing prediction error, the second term enforces a low-rank structure by approximating each $\boldsymbol{A}$ as a weighted sum of shared components, the third term encourages smoothness by penalizing differences between nearby interactions in feature space, and the two final terms impose sparsity on the interactions and corresponding weights (see discussion on hyper-parameters in App. A.4).

We fit the model using an iterative approach, iterating between estimating the interaction matrices $\{\boldsymbol{A}_t^{(n,k)}\}$ via LASSO Tibshirani (1996), the shared components $\{\boldsymbol{c}_j, \boldsymbol{r}_j\}$, and their weights $\{\boldsymbol{W}_{:t}^{(n,k)}\}$ using PARAFAC (Harshman et al., 1970), alternating between the following steps.

1. **Update interaction matrices $\boldsymbol{A}$:** Given the current components and weights, infer each interaction matrix by minimizing the reconstruction and smoothness losses (for notation clarity, we denote in the equation $\widehat{\boldsymbol{A}}_t^{(n,k)} := \boldsymbol{A}$)

$$\widehat{\boldsymbol{A}}_t^{(n,k)} = \arg\min_{\boldsymbol{A}} \|\boldsymbol{x}_t^{(n)} - \sum_{k=1}^{N} (\mathbf{1}_{d_{n,k}(t) \leq R}) \boldsymbol{A} \boldsymbol{x}_{t-1}^{(k)}\|_2^2 + \lambda_1 \left\| \boldsymbol{A} - \sum_j W_{j,t} \boldsymbol{c}_j \boldsymbol{r}_j^\top \right\|_2^2 + \lambda_2 \sum_{\boldsymbol{A}' \neq \boldsymbol{A}} e^{-\sigma \delta(\boldsymbol{A}, \boldsymbol{A}')} \|\boldsymbol{A} - \boldsymbol{A}'\|_2^2 + \lambda_3 \|\mathrm{vec}(\boldsymbol{A})\|_1 \tag{5}$$

   where λ_1, λ_2, and λ_3 are hyper-parameters balancing low-rank fidelity, interaction smoothness, and sparsity level.

2. **Update shared components $\boldsymbol{c}_j, \boldsymbol{r}_j$:** With fixed weights and $\{\boldsymbol{A}_t^{(n,k)}\}$, optimize the low-rank factors to best approximate the interaction matrices across all pairs and times:

$$\{\widehat{\boldsymbol{c}}_j, \widehat{\boldsymbol{r}}_j\}_{j=1}^{J} = \arg \min_{\{\boldsymbol{c}_j, \boldsymbol{r}_j\}} \sum_{n,k,t} \left\| \boldsymbol{A}_t^{(n,k)} - \sum_{j=1}^{J} W_{j,t}^{(n,k)} \boldsymbol{c}_j \boldsymbol{r}_j^\top \right\|_2^2 \tag{6}$$

3. **Update weights $\forall t : \{\widehat{\boldsymbol{W}}_{:t}^{(n,k)}\}_{n,k}$:** Given the current components and $\{\boldsymbol{A}_t^{(n,k)}\}$, update the weights for each time step by minimizing the reconstruction error with sparsity constraints:

$$\widehat{\boldsymbol{W}}_{:,t}^{(n,k)} = \arg \min_{\boldsymbol{W}_{:t}^{(n,k)}} \left\| \boldsymbol{A}_t^{(n,k)} - \sum_{j=1}^{J} W_{j,t}^{(n,k)} \boldsymbol{c}_j \boldsymbol{r}_j^\top \right\|_2^2 + \lambda_4 \|\boldsymbol{W}_{:,t}^{(n,k)}\|_1 \tag{7}$$

   where λ_4 controls the sparsity of the number of active components in the $k \to n$ interaction at time t.

`LICCHIE` can thus identify interactions between individual cells based on their features (i.e., with good resolution and not mean-field) while considering all cells within the interaction radius, thus providing a coherent and unified representation in the feature space. Notably, environmental effects that are not caused by other cells, as well as the natural evolution of a cell, can be captured via the self–self interaction matrix of a cell with itself and integrated into the summation of all interactions, which enables disentangling these effects. The steps are summarized in Algorithm 1 and Fig. 1, optimization details and complexity analysis are provided in App. A.5.

5 EXPERIMENTS

5.1 `LICCHIE` RECOVERS TRUE COMPONENTS IN SYNTHETIC DATA

Setup. We simulated two populations (25 and 10 cells, 10 features each) evolving over T=50 frames (Fig. 2a). Eight ($J = 8$) rank-1 ground-truth motifs $\{M\star_j\}_{j=1}^{8}$ were generated to be orthogonal under the Frobenius inner product, with per-pair, per-time weights that are 20% sparse. Initial features came from population-specific Gaussians; cell motion followed trigonometric trajectories with population-specific amplitudes and frequencies; Gaussian noise was added each step. Edges were constructed by kNN ($k = 10$) on t_0 features and pruned per frame by a spatial radius R. At each time, targets were updated by a weighted sum of sources within the radius. To prevent divergence, we scaled updates to keep the spectral radius below 1.

Algorithm 1 LICCHIE: Low-rank interpretable cell-cell hidden interactions via embedding model

```
Require: Observed cell features $\{\boldsymbol{x}_t^{(n)}\}$, interaction radius $R$, number of components $J$, regularization weights $\lambda_1, \lambda_2, \lambda_3$, feature-space decay $\sigma$
Ensure: Interaction matrices $\{\boldsymbol{A}_t^{(n,k)}\}$, shared components $\{\boldsymbol{c}_j, \boldsymbol{r}_j\}$, weights $\boldsymbol{W}$
Initialize $\boldsymbol{A}_t^{(n,k)}$, $\boldsymbol{c}_j, \boldsymbol{r}_j$, $\{\boldsymbol{W}_{:t}^{(n,k)}\}_{n,k,t}$ randomly from normal i.i.d distribution
repeat
    Update interaction matrices $\{\boldsymbol{A}_t^{(n,k)}\}$:
    for each cell pair $(n, k)$ and time $t$ do
        Infer $\boldsymbol{A}_t^{(n,k)}$        ▷ equation 5, balances low-rank, smoothness, and sparsity
    end for
    Update shared components $\boldsymbol{c}_j, \boldsymbol{r}_j$        ▷ equation 6
    Update weights $\{\boldsymbol{W}_{:t}^{(n,k)}\}_{n,k,t}$        ▷ equation 7
until convergence or maximum iterations reached
Return $\{\boldsymbol{A}_t^{(n,k)}\}_{n,k,t}, \{\boldsymbol{c}_j, \boldsymbol{r}_j\}_j, \{\boldsymbol{W}_{:t}^{(n,k)}\}_{n,k,t}$
```

Baselines. We consider two baseline approaches (a) a *global time-fixed linear* model, inferring a constant matrix $\widetilde{\boldsymbol{A}}$ via least squares on features from all cells and (b) *per-cell linear dynamics*, fitting a per-cell linear dynamics model using all observed interactions, yielding $\widetilde{\boldsymbol{A}}_n^*$ for each target cell $n = 1 \dots N$. Of note, alternative deep learning methods do not produce directly interpretable, comparable components of the underlying interactions and thus are not suitable for this comparison.

Results. We first validate `LICCHIE`'s inferred interactions via their predictive power; we observe strong performance in prediction of cell-type identity of either source (0.86) or target (0.87) cells, or both jointly (0.83, Fig. 2c). Next, in comparison to baselines, `LICCHIE` produced interaction structures closer to the ground truth construction $\{\boldsymbol{M}\star_j\}_{j=1}^{8}$ (Fig. 2d) and lower reconstruction error, both per-pair (distribution in Fig. 2e) and on average across pairs (Fig. 2f).

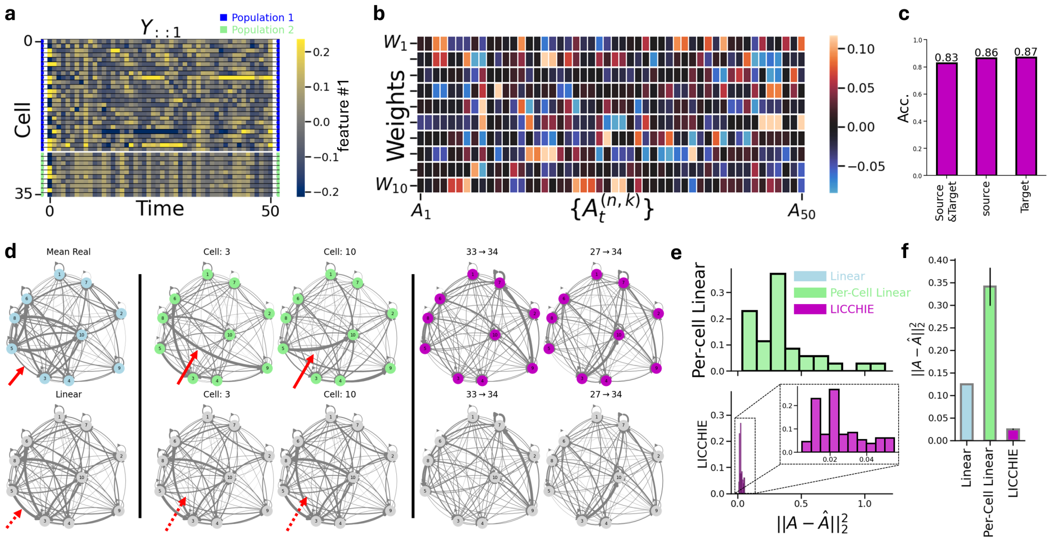

Figure 2: *Synthetic data overview and results.* **a**, Examples of generated features illustrating population differences (blue/green side bars). **b**, Ground-truth weights show sparse structure. **c**, Accuracy results for `LICCHIE`'s cell type prediction using the inferred interactions (logistic regression with 5-fold cross-validation); predictions of source and target (0.83), or source (0.86) and target (0.87) independently. **d**, Rank-1 motif inference (top row; left to right: linear, per-cell linear, `LICCHIE`) vs. ground truth (bottom row). Red arrows highlight discrepancies between ground truth (dotted) and baseline-identified networks (solid). **e**, Per-pair MSE distributions vs. ground truth for the per-cell linear baseline (top) and `LICCHIE` (bottom). **f**, Average MSE for the global time-fixed linear baseline, per-cell dynamics, and `LICCHIE` (additional metrics Fig. S2)

5.2 `LICCHIE` UNCOVERS INTERPRETABLE TUMOR–NK INTERACTIONS

Data and features. We applied `LICCHIE` to 3D live-cell confocal co-cultures of patient-derived gastric tumor organoids with primary human NK cells (Liu et al., 2024), presenting variability in cell morphology and intensity distribution (Fig. 3a). Imaging spanned 8–12 h at 4-min cadence with 27–30 z-slices per time point, and includes multiple conditions . Using a representative experiment we extracted spatial, morphology, and intensity features. We normalize the features such that each feature, at population level (NK or tumor cells), is normally distributed, $f_i^{(p)} \sim \mathcal{N}(\mu = 0, \sigma = 1)$, implying that all features, although representing geometric properties, attain negative values (App. A.6, Table S2). We fitted `LICCHIE` with J=10 components.

Components and structure. `LICCHIE` identified components that reveal diverse interaction patterns, including *localized targets* (e.g., $\boldsymbol{M}_2, \boldsymbol{M}_4$), *source-emerging* (e.g., $\boldsymbol{M}_1, \boldsymbol{M}_{10}$), self *feedback-like* interactions (e.g., $\boldsymbol{M}_3$) and multi-step interactions, including direct and indirect influences (e.g., $\boldsymbol{M}_6$, Fig. S3).

Interaction motifs can capture indicators of cell state, such as NK cell polarization. Upon polarization, NK cells undergo a characteristic morphological shift from round to elongated (Fig. 3a), a process necessary for their cytotoxic activity. Analysis of motif 2 ($\boldsymbol{M}_2$) and the respective weights ($\boldsymbol{W}_2$) suggests that it captures predominant interactions targeting NK cells, with opposing effects depending on the source cell (Fig. 3b,c), specifically identifying that the length of the major axis (light green) of a target NK cell–a measure of its longest dimension that reflects its polarization status–is predicted by diverse source features. This suggests that in a complex co-culture system, the NK cell length (indication of its polarization) is not a singular, homogeneous event but is instead coupled to diverse source cell states.

Motif 5 ($\boldsymbol{M}_5$) has a sparse interaction matrix (Fig. 3b) and predominantly captures NK-to-tumor interactions ($\boldsymbol{W}_5$, Fig. 3c). This motif indicates that source cell mean intensity and sphericity, are the main predictive features of two primary target features: cell aspect ratio and a tumor cell identity indicator. Therefore, $\boldsymbol{M}_5$ suggests that interaction with a subset of NK cells distinguished by their brightness and sphericity is indicative for tumor-cell elongation. Interestingly, motif 6 ($\boldsymbol{M}_6$) also centers on cell aspect ratio through different means, including direct and indirect influences of several features. For example, cell sphericity (light green) directly and indirectly through cell height (gray) influences cell aspect ratio. This underscores `LICCHIE`'s ability to disentangle multiple processes with overlapping actors while distinguishing the underlying mechanisms.

Lastly, motif 7 ($\boldsymbol{M}_7$) highlights the centrality of source cell size in shaping multiple target features, and the respective weights ($\boldsymbol{W}_7$) indicate that $\boldsymbol{M}_7$ mostly negatively regulates tumor-to-NK interactions. A particularly notable target feature is cell surface area. In NK cells, an increased surface area could indicate the formation of a lytic immunological synapse (Orange, 2008). This may indicate an impact of interactions between tumor cells of specific size and the process of NK-cell spreading prior to synapse formation.

Interaction-type specificity. Component activations differ across interaction classes. As shown in Fig. 3c, components display class-specific activation patterns across tumor→tumor, NK→NK, NK→tumor, and tumor→NK pairs. Interesting relations observed include:

- *Opposing weights between cell classes*. Component 6 (weighted by $\boldsymbol{W}_6$) presents positive weights for same-cell interactions and negative weights for opposing cell classes. Mechanistically, similar shapes may promote symmetric contact geometry or aligned polarity, resulting in higher positive w_6 values for same-cell pairs, whereas mismatched shapes reduce contact or produce asymmetry, leading to negative $\boldsymbol{W}_6$ values for different-cell pairs.
- *Single interaction dominance*. Components 1 and 6 ($\boldsymbol{W}_1$, $\boldsymbol{W}_6$), show clear dominance of a single interaction type (e.g., NK → NK), which may reflect specialized functional role.
- *Source disentangling*. Components 1 and 2 ($\boldsymbol{W}_1$, $\boldsymbol{W}_2$) capture predominant interactions targeting tumor and NK cells respectively, with opposing weight signs given the source.
- *Cell-specific regulation*. Component 3 ($\boldsymbol{W}_3$) is mainly associated with interactions between tumor and NK cells (rather than homogeneous cell-type interactions.

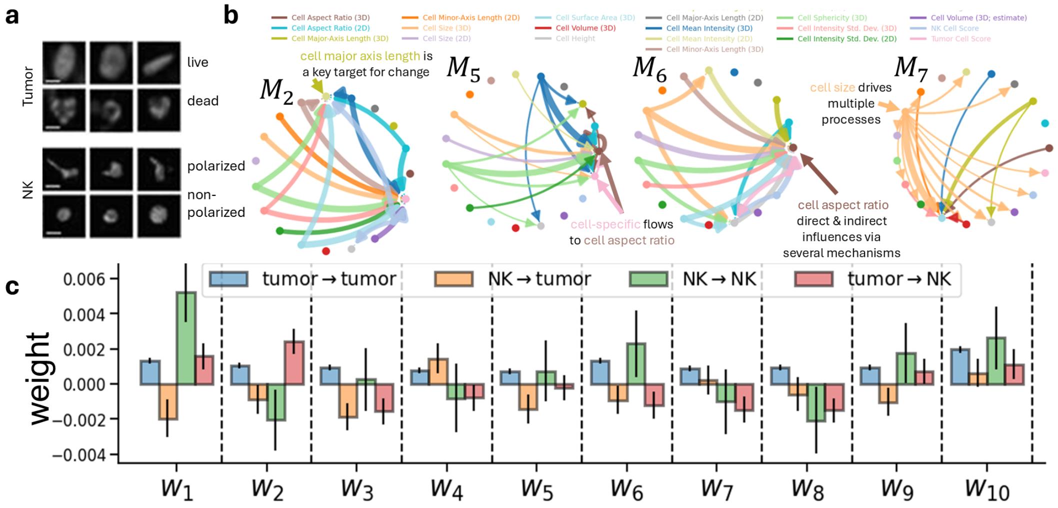

Figure 3: `LICCHIE` *identifies meaningful cellular interactions in real-world data.* **a**, Representative images illustrate morphology/intensity differences between phenotypes (Liu et al., 2024). **b**, Examples of `LICCHIE`'s rank-1 motifs (J=10; full set in Fig. S3). **c**, Component weights by interaction class (tumor→tumor, NK→NK, NK→tumor, tumor→NK) show class-specific activation patterns.

5.3 FROM MOTIFS TO PHENOTYPES: USING ENGINEERED CELLS FOR PREDICTION

Weight sweeps on unseen data. A unique property of `LICCHIE` is the globality of the rank-1 components, $\{M_j\}_{j=1}^{J}$ that enables generalization to unseen data that can drive critical biological downstream processes. We assessed whether individual components are associated with increased NK polarization or tumor-cell death by modulating each component's weight while keeping source features within the observed feature space.

Procedure. For each M_j, $j = 1 \ldots J$ component, we gradually increased its weight from 0 to 2 in 20 steps and applied it to various combinations of source features. We then assigned each inferred target a polarization score (for NK cells) or a death score (for tumor cells) using a distance-weighted k-nearest neighbor classifier (k=150, total # samples from all time points = $20,437$) based on Euclidean distance in feature space and normalized by phenotype frequency (Fig. 4a).

Components and cell-state relations. Different components yielded distinct death/polarization score trajectories with non-linear changes as weights increased (Fig. 4a). For tumor-cell death, some components (e.g., M_1, M_5, M_6, Fig. S3) showed increasing scores with weight, whereas others (e.g., M_2, M_3, M_8) showed decreases, often most prominently around weight ≈ 1. Overall, tumor death scores exhibited phase-transition-like effects (non-smooth changes), consistent with the acute nature of cell death. NK-cell polarization trends were more variable; for example, M_7 increased polarization at higher weights, while M_5 decreased it. In some cases (e.g., M_5), tumor-death and NK-polarization scores moved in opposite directions as weight increased.

Prediction from inferred weights. Using weights inferred on unseen frames, we trained a cross-validation logistic regression to predict target phenotypes. Although phenotypes are not used during training, weight-only features predicted phenotypes on held-out data (Fig. 4b), indicating that component activations capture phenotype-relevant interaction signals.

Source-feature extremes. We also examined the impact of source features on NK polarization under maximal activation (W=2). For each source feature, we compared scores at its empirical minimum vs. maximum (others fixed at their means). We found that source properties give rise to both diversity and correlations in the potential effect of each component. For example, when component M_1 is applied to a source feature vector with a minimized cell aspect ratio, it increases the polarization score of the target. In contrast, when applied to the maximum cell aspect ratio, the target cell's polarization score remains near zero. Other components, such as M_4 or M_9, show similar target polarization scores under both minimized and maximized cell aspect ratio values, sug-

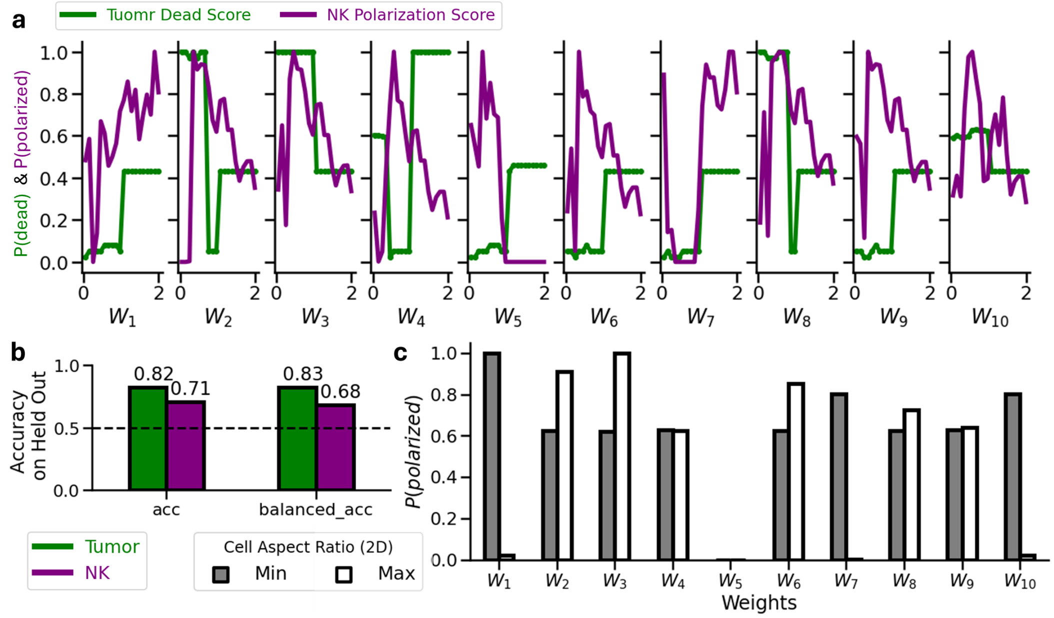

Figure 4: *Probing phenotype effects via weight modulation*. **a**, Gradually modulating the weights on unseen data points reveals diverse encoding mechanisms for NK polarization (magenta) and tumor cell death (green). **b**, Accuracy and balanced-accuracy for phenotype prediction using inferred component weights on held-out data. **c**, Effect of extreme source-feature values on target NK polarization.

gesting that this feature may not be critical in these circuits for driving changes in target polarization (Fig. 4c). Altogether, these results highlight the potential of `LICCHIE` to guide the design of future interventions when applied to additional datasets.

6 DISCUSSION AND FUTURE WORK

We presented `LICCHIE`, an interpretable method to study cell-cell interactions capturing modular, feature-specific cell–cell interactions. The cell state feature representations are user defined; allowing for problem specific optimization. Over these, `LICCHIE` reveals interpretable motifs showing how features act as targets or sources within modular cell–cell interactions through coordinated, overlapping, and co-modulatory effects. Applied to live-cell imaging data of tumor–NK co-cultures, `LICCHIE` revealed key patterns of cellular communication–identified components with distinct feature relations–revealing the modular biological regulation that integrates diverse features into phenotypic outcomes. An important advantage of `LICCHIE` is that it can address the challenge of modeling interactions from separate, asynchronous measurements while preserving data-specific details. In particular, its global vectors $\boldsymbol{c}_j$ and $\boldsymbol{r}_j$ allow new data to be analyzed quickly by reusing the components and adjusting only the weights, avoiding full-model retraining. Notably, `LICCHIE` naturally handles data imbalance by its ability to re-weight interactions based on the prevalence of each population and interaction type, ensuring adequate representation for all interactions.

Limitations and future directions. The model assumes linear or additive effects, which may not fully capture certain non-linear dynamics; this design choice promotes interpretability, and found valid in practice through reconstruction accuracy (see discussion in App. A.1). Future work can incorporate non-linear activation functions $f(\cdot)$ to each interaction pair to improve flexibility. Here, interactions beyond the defined radius are not captured; this limitation can be easily addressed by applying a kernel and sampling from an appropriate probability distribution that varies with distance. Further biological interpretation of components requires additional supervision via experimental data that can be obtained by applying `LICCHIE` to additional systems with dedicated measurements. Lastly, using temporal mapping tools (e.g., optimal transport, Klein et al. (2025)), `LICCHIE` can be applied to spatio-temporal single-cell datasets, providing insight into interactions in gene feature space.

REFERENCES

Luca Alessandri, Francesca Cordero, Marco Beccuti, Nicola Licheri, Maddalena Arigoni, Martina Olivero, Maria Flavia Di Renzo, Anna Sapino, and Raffaele Calogero. Sparsely-connected autoencoder (sca) for single cell rnaseq data mining. *NPJ systems biology and applications*, 7(1):1, 2021.

Maria Alieva, Amber KL Wezenaar, Ellen J Wehrens, and Anne C Rios. Bridging live-cell imaging and next-generation cancer treatment. *Nature Reviews Cancer*, 23(11):731–745, 2023.

Tomas Andersson, Lars Alfredsson, Henrik Källberg, Slobodan Zdravkovic, and Anders Ahlbom. Calculating measures of biological interaction. *European journal of epidemiology*, 20(7):575–579, 2005.

Erick Armingol, Adam Officer, Olivier Harismendy, and Nathan E Lewis. Deciphering cell–cell interactions and communication from gene expression. *Nature Reviews Genetics*, 22(2):71–88, 2021.

Erick Armingol, Hratch M Baghdassarian, and Nathan E Lewis. The diversification of methods for studying cell–cell interactions and communication. *Nature Reviews Genetics*, 25(6):381–400, 2024.

Robert H Boling Jr. Toward state-space models for biological populations. *Journal of theoretical biology*, 40(3):485–506, 1973.

Fred Brauer and Christopher Kribs. *Dynamical systems for biological modeling*. CRC press, 2016.

Arijit Chakrabarti and Jayanta K Ghosh. Aic, bic and recent advances in model selection. *Philosophy of statistics*, pp. 583–605, 2011.

Yenho Chen, Noga Mudrik, Kyle A Johnsen, Sankaraleengam Alagapan, Adam S Charles, and Christopher Rozell. Probabilistic decomposed linear dynamical systems for robust discovery of latent neural dynamics. *Advances in Neural Information Processing Systems*, 37:104443–104470, 2024.

Andreas P Cuny, Fabian P Schlottmann, Jennifer C Ewald, Serge Pelet, and Kurt M Schmoller. Live cell microscopy: From image to insight. *Biophysics Reviews*, 3(2), 2022.

Richard A Harshman et al. Foundations of the parafac procedure: Models and conditions for an "explanatory" multi-modal factor analysis. *UCLA working papers in phonetics*, 16(1):84, 1970.

Lukas Heumos, Anna C Schaar, Christopher Lance, Anastasia Litinetskaya, Felix Drost, Luke Zappia, Malte D Lücken, Daniel C Strobl, Juan Henao, Fabiola Curion, et al. Best practices for single-cell analysis across modalities. *Nature Reviews Genetics*, 24(8):550–572, 2023.

Dominik Klein, Giovanni Palla, Marius Lange, Michal Klein, Zoe Piran, Manuel Gander, Laetitia Meng-Papaxanthos, Michael Sterr, Lama Saber, Changying Jing, et al. Mapping cells through time and space with moscot. *Nature*, 638(8052):1065–1075, 2025.

Konstantinos Lazaros, Dimitris E Koumadorakis, Panagiotis Vlamos, and Aristidis G Vrahatis. Graph neural network approaches for single-cell data: a recent overview. *Neural Computing and Applications*, 36(17):9963–9987, 2024.

Scott Linderman, Matthew Johnson, Andrew Miller, Ryan Adams, David Blei, and Liam Paninski. Bayesian learning and inference in recurrent switching linear dynamical systems. In *Artificial intelligence and statistics*, pp. 914–922. PMLR, 2017.

Bodong Liu, Yanting Zhu, Zhenye Yang, Helen HN Yan, Suet Yi Leung, and Jue Shi. Deep learning–based 3d single-cell imaging analysis pipeline enables quantification of cell–cell interaction dynamics in the tumor microenvironment. *Cancer Research*, 84(4):517–526, 2024.

Nikita Moshkov, Michael Bornholdt, Santiago Benoit, Matthew Smith, Claire McQuin, Allen Goodman, Rebecca A Senft, Yu Han, Mehrtash Babadi, Peter Horvath, et al. Learning representations for image-based profiling of perturbations. *Nature communications*, 15(1):1594, 2024.

Noga Mudrik, Yenho Chen, Eva Yezerets, Christopher J Rozell, and Adam S Charles. Decomposed linear dynamical systems (dlds) for learning the latent components of neural dynamics. *Journal of Machine Learning Research*, 25(59):1–44, 2024.

Josue Nassar, Scott W Linderman, Monica Bugallo, and Il Memming Park. Tree-structured recurrent switching linear dynamical systems for multi-scale modeling. *arXiv preprint arXiv:1811.12386*, 2018.

Jim Nilsson and Tomas Akenine-Möller. Understanding ssim. *arXiv preprint arXiv:2006.13846*, 2020.

Jordan S Orange. Formation and function of the lytic nk-cell immunological synapse. *Nature Reviews Immunology*, 8(9):713–725, 2008.

Toby A Patterson, Len Thomas, Chris Wilcox, Otso Ovaskainen, and Jason Matthiopoulos. State–space models of individual animal movement. *Trends in ecology & evolution*, 23(2):87–94, 2008.

A Saranya and R Subhashini. A systematic review of explainable artificial intelligence models and applications: Recent developments and future trends. *Decision analytics journal*, 7:100230, 2023.

David R Stirling, Anne E Carpenter, and Beth A Cimini. Cellprofiler analyst 3.0: accessible data exploration and machine learning for image analysis. *Bioinformatics*, 37(21):3992–3994, 2021.

Jimeng Su, Ying Song, Zhipeng Zhu, Xinyue Huang, Jibiao Fan, Jie Qiao, and Fengbiao Mao. Cell–cell communication: new insights and clinical implications. *Signal transduction and targeted therapy*, 9(1):196, 2024.

Wenzhuo Tang, Hongzhi Wen, Renming Liu, Jiayuan Ding, Wei Jin, Yuying Xie, Hui Liu, and Jiliang Tang. Single-cell multimodal prediction via transformers. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pp. 2422–2431, 2023.

Luke Ternes, Mark Dane, Sean Gross, Marilyne Labrie, Gordon Mills, Joe Gray, Laura Heiser, and Young Hwan Chang. A multi-encoder variational autoencoder controls multiple transformational features in single-cell image analysis. *Communications biology*, 5(1):255, 2022.

Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1):267–288, 1996.

Allon Wagner, Aviv Regev, and Nir Yosef. Revealing the vectors of cellular identity with single-cell genomics. *Nature biotechnology*, 34(11):1145–1160, 2016.

Caleb Weinreb, Samuel Wolock, Betsabeh K Tusi, Merav Socolovsky, and Allon M Klein. Fundamental limits on dynamic inference from single-cell snapshots. *Proceedings of the National Academy of Sciences*, 115(10):E2467–E2476, 2018.

Aaron J Wilk, Alex K Shalek, Susan Holmes, and Catherine A Blish. Comparative analysis of cell–cell communication at single-cell resolution. *Nature Biotechnology*, 42(3):470–483, 2024.

Eva Yezerets, Noga Mudrik, and Adam S Charles. Decomposed linear dynamical systems (dlds) models reveal instantaneous, context-dependent dynamic connectivity in c. elegans. *Communications Biology*, 8(1):1218, 2025.

A APPENDIX

A.1 THE VALIDITY OF THE LINEARITY ASSUMPTION

Linear approximations are widely used in dynamics because they simplify analysis and control. Around an equilibrium point or operating region, many nonlinear systems can be locally approximated with linear models (i.e., via Taylor expansion). Accordingly, locally linear models, including linear state-space models, auto-regressive models, GLMs, and linearized ODEs, are standard for modeling, e.g., movement Patterson et al. (2008), population dynamics Boling Jr (1973), and diverse biological interactions Yezerets et al. (2025); Brauer & Kribs (2016); Andersson et al. (2005).

A unique advantage of maintaining a linear transition matrix $\boldsymbol{A}$ in biological systems is interpretability: in our framework each entry $\boldsymbol{A}_{ij}$ directly reflects the effect of feature i on feature j, and can be linked to the observation space. `LICCHIE` extends typical linear systems model to be flexible distance-varying time-changing cell-specific locally changing dynamics, and thus is expressive enough to capture the complexity of temporally evolving cell–cell interactions, while preserving this interpretability: each entry in a rank-1 component corresponds to a source–target effect between cells.

To account for non-linear dynamics one can consider non-linear activations, e.g. $\boldsymbol{x}_t = \sigma(\boldsymbol{A}_t)\boldsymbol{x}_{t-1}$. However, including these prevents direct interpretation of the results, necessary for biological discovery.

A.2 POPULATION-SPECIFIC FEATURE SPACE

The `LICCHIE` framework can naturally extend to account for multiple cell populations exhibiting distinct feature subsets. To do so, the fixed component sets $\{\boldsymbol{c}_j\}$ and $\{\boldsymbol{r}_j\}$ are replaced by population-specific sets $\{\boldsymbol{c}_j^{(p)}\}_{j=1}^J$ and $\{\boldsymbol{r}_j^{(p)}\}_{j=1}^J$, where p indexes the population. Within each set, the vector dimensions are consistent, but they may differ between populations.

For within-population interactions (e.g., population p), the $\boldsymbol{A}$ matrices are modeled using the components of the corresponding p set. For cross-population interactions, the sets are chosen according to the participating populations: for an interaction from a cell in population p to a cell in population p', we use $\{\boldsymbol{r}_j\}$ from p and $\{\boldsymbol{c}_j\}$ from p', and vice versa. The resulting interaction matrix may not be square if the number of features differs between populations.

Notably, the above extension preserves the interpretability afforded by the component structure while allowing flexibility across nuanced population dynamics.

A.3 AVOIDING EQUIVALENT SOLUTIONS THROUGH DISTANCE CONSTRAINTS

Learning interactions for all feature and spatial combinations is intractable and highly non-interpretable for scientific purposes. Hence, we leverage our assumption of cross-interaction similarity in $\boldsymbol{A}$ matrices that represent similar source-target feature distributions.

When fitting dynamics, similarities in source and target vector pairs should ideally result in similar transition matrices (i.e., if (n,k) is close to (n',k'), then their interactions $\boldsymbol{A}_t^{(n,k)}$ and $\boldsymbol{A}_t^{(n',k')}$ should be similar). Yet, when fitting dynamics in an unconstrained way, if the source vector $\mathbf{x}_t^k$ lies in the null space of any matrix $\widetilde{\boldsymbol{A}}$, then $\boldsymbol{A}_t^{(n,k)} + \widetilde{\boldsymbol{A}}$ is an equivalent solution to $\boldsymbol{A}_t^{(n,k)}$ in terms of fit. This means that we must identify a solution that is consistent across pair, enforcing smoothness across parameter changes.

Hence, we introduced the distance constraint; forcing distances between $\{\boldsymbol{A}\}$ matrices to follow distances in the corresponding feature space, i.e.,

$$\|\boldsymbol{A}_t^{(n,k)} - \boldsymbol{A}_t^{(n',k')}\|_2^2 \leq \epsilon\big(\delta_{(n,k),(n',k')}(t)\big)$$

where ϵ is some function that depends on $\delta_{(n,k),(n',k')}(t)$ refers to the distance between stacked feature vectors of interaction pairs (n,k) and (n',k') at time t.

A.4 OBJECTIVE HYPER-PARAMETERS

The `LICCHIE` objective includes a minimal set of hyper-parameters introduced to balance representation sparsity (λ_3, λ_4), smoothness (λ_2), data fidelity, and low-rank approximation (λ_1). In addition, the framework provides the user with the flexibility to set the maximum rank of each interaction, J, and the effective radius of interaction R; both are quantities that should be set with respect to the characteristics of the biological system studied, i.e., accounting for the total number of features and the representative length scale.

In practice, hyper-parameters can be tuned by running a parameter search (e.g., grid search) while optimizing an information criterion that balances degrees of freedom and model complexity, such as AIC (Akaike Information Criterion) or BIC (Bayesian Information Criterion, Chakrabarti & Ghosh (2011)).We provide guidance and further intuition for tuning these,

- **Low-rank constraint (J, λ_1)**: These control how strongly the $\boldsymbol{A}$'s follow the low-rank approximation. When J is larger, or when $\boldsymbol{A}$ is naturally close to low-rank, this term should converge to 0, and the sensitivity to λ_1 will be minimal (since its multiplier is close to zero). For highly complex, high-dimensional, and fast-varying dynamics, the low-rank Frobenius norm may not be small. In such cases, we recommend monitoring the rank-1 reconstruction over iterations (calculated automatically within the model). If the low-rank reconstruction remains low across iterations, users should increase both J and λ_1.
- **Interactions smoothness (λ_2)**: The smoothness of interactions in feature space, mainly preventing null-space solutions from causing abrupt switches (see App. A.3), is controlled by λ_2. In practice, a small value $\lambda_2 < 1$ is usually sufficient. We recommend examining the smoothness of interactions under gradually changing parameters to confirm that the identified interactions align with biological or system-specific assumptions.
- **Sparsity constraint (λ_3, λ_4)**: Sparsity is induced by these by modulating the number of zero or near-zero entries. λ_3 promotes sparsity in the interaction matrices, facilitating a clearer understanding of the functionality of each interaction. λ_4 encourages interactions to be composed a limited number of $\{\boldsymbol{c}_j\}$ and $\{\boldsymbol{r}_j\}$ components, valuable for,
    1. Understanding the modular role of each $\{\boldsymbol{c}_j, \boldsymbol{r}_j\}$ pair in driving the overall interaction,
    2. Identifying commonalities and differences between interaction types (if all components are used in all interactions, distinguishing differences is harder), and
    3. Reducing noise captured by other components, yielding a more robust solution.
- **Effective interaction range (R)**: The radius, R, defines a disk around a cell i in which interaction are considered, i.e., cells within this disk are considered interacting with it. This shall be set with respect to the system's length scales and desired scope of interactions one wishes to study. e.g., one can choose to limit the analysis to first-order interactions, accounting for cells in immediate neighborhood or larger value for long-range effects.

A.5 OPTIMIZATION AND COMPLEXITY

To fit the model's components we use LASSO (Tibshirani, 1996), estimating the interaction matrices, $\{\boldsymbol{A}_t^{(n,k)}\}$, and Parallel Factor Analysis (PARAFAC, Harshman et al. (1970)) for the shared components $\{\boldsymbol{c}_j, \boldsymbol{r}_j\}$, and their weights $\{\boldsymbol{W}_{:t}^{(n,k)}\}$. PARAFAC is a method to decompose high-dimensional tensor data (multi-way) into underlying, independent components, extending the Principal Component Analysis (PCA) model to more than two dimensions. This provides a unique solution that allows for the recovery of pure components.

The complexity of the optimization is as follows; let N denote the number of cells, T the number of time points, d the number of features, k the average number of neighbors (within radius), $S = k - 1$ the number of sources, R the number of outer iterations, C the number of PARAFAC components, I the number of PARAFAC ALS iterations, M the number of past A's maintained for similarity search, and $E \approx T \cdot N \cdot k$ the total number of interactions.

Pre-processing (neighbors within radius).

- Naive all-pairs distances (per time): $O(T \cdot N^2)$ time, $O(1)$ extra memory.

- With a spatial index (grid/kd-tree): build + query costs $O(T \cdot N \log N + T \cdot E)$ time. The neighbor graph requires $O(E)$ memory.

Per target per time (one cell as target).

- Build wide $\boldsymbol{A} \in \mathbb{R}^{d \times (S \cdot d)}$: $O(S \cdot d^2)$.
- Least-squares solve $x_{\text{target}} = \boldsymbol{A}\, x_{\text{sources}}$: $O(S^2 \cdot d^3)$ time, $O(S \cdot d^2)$ memory.
- Maintain M most similar past $\boldsymbol{A}$'s: similarity search over history $H \approx R$, costing $O(H \cdot S \cdot d^2)$ (naive).
- Decompose each A (reshaped as a tensor $d \times S \times d$) with PARAFAC/CP: ALS per $\boldsymbol{A}$ costs $O(I \cdot C \cdot S \cdot d^2)$ time, $O(C \cdot (d + S + d))$ memory.

Totals per outer iteration (across all targets and times; $\sim T \cdot N$ problems).

- LS solves: $O(T \cdot N \cdot S^2 \cdot d^3)$.
- Similarity searches: $O(T \cdot N \cdot R \cdot S \cdot d^2)$ (if $H \approx R$).
- PARAFAC: $O(T \cdot N \cdot I \cdot C \cdot S \cdot d^2)$.

Overall complexity.

$$O\big(T \cdot N \cdot \big(S^2 d^3 + (R + I \cdot C) S d^2 + \log N + k\big)\big) .$$

A.6 THE TUMOR-NK DATASET

We used publicly available data from (Liu et al., 2024), accessed via Zenodo. The complete dataset includes 3D live-cell imaging datasets of gastric tumor organoids co-cultured with primary human Natural Killer (NK) cells. The dataset includes multiple sessions recorded asynchronously under varying experimental conditions (see (Liu et al., 2024) for more details). We selected a representative random session, 'GX048-TO + NK cell (IL-15)', in which NK cells were stimulated with IL-15.

A.7 TABLES

Table S1: Summary of defined notations.

Notation	Meaning
N	Number of interacting cells
T	Total duration of observation (number of time points)
s	Dimensionality of the space $\mathcal{R} \subset \mathbb{R}^s$
$\boldsymbol{\psi}_t^{(n)}$	Position of cell n at time t in $\mathbb{R}^s$ (spatial coordinates)
$\boldsymbol{x}_t^{(n)}$	State (feature) vector of cell n at time t in $\mathbb{R}^{m^{p^{(n)}}}$
$p^{(n)}$	Population label/type of cell n
$\boldsymbol{P} = \{p^{(n)}\}_{n=1}^N$	Set of population identifiers for all cells
$m^{p^{(n)}}$	Number of features for population $p^{(n)}$; if constant across all cells, denote as m
$m_{\text{spatial}}^{p^{(n)}}$	Number of spatial features for population $p^{(n)}$
$m_{\text{internal}}^{p^{(n)}}$	Number of internal features for population $p^{(n)}$
$d_{n,n'}(t)$	Spatial distance between cells n and n' at time t, e.g., $\\|\boldsymbol{r}_t^{(n)} - \boldsymbol{r}_t^{(n')}\\|$
$\delta_{n,n'}(t)$	Feature distance between cells n and n' at time t, e.g., $\\|\boldsymbol{x}_t^{(n)} - \boldsymbol{x}_t^{(n')}\\|$
$\delta_{(n,k),(n',k')}(t)$	Feature distance between stacked state vectors of cell pairs (n, k) and (n', k') at time t.
$\delta_{A,A'}(t)$	Feature distance between As
R	Interaction radius (hyperparameter)
$\mathbf{1}_{d_{n,n'}(t) \leq R}$	Indicator: 1 if distance between n and n' at time t is $\leq R$, else 0

Table S2: Summary of cell features.

Feature	**Description**
Cell Aspect Ratio (2D)	Ratio of major to minor axis lengths in 2D projection.
Cell Aspect Ratio (3D)	Ratio of principal axis lengths in 3D reconstruction.
Cell Major-Axis Length (2D)	Length of the longest axis in 2D projection.
Cell Major-Axis Length (3D)	Length of the longest axis in 3D reconstruction.
Cell Minor-Axis Length (2D)	Length of the shortest axis in 2D projection.
Cell Minor-Axis Length (3D)	Length of the shortest axis in 3D reconstruction.
Cell Height	Extent of the cell along the z-axis.
Cell Size (2D)	Area of the cell in 2D projection.
Cell Size (3D)	Approximate size based on 3D reconstruction.
Cell Surface Area (3D)	Total exposed surface area of the 3D cell.
Cell Volume (3D)	Computed volume of the cell in 3D.
Cell Volume (3D; estimate)	Estimated cell volume when full reconstruction is not available.
Cell Sphericity (3D)	Measure of how closely the shape approaches a sphere.
Cell Mean Intensity (2D)	Average intensity of the cell in 2D projection.
Cell Mean Intensity (3D)	Average intensity of the cell in 3D reconstruction.
Cell Intensity Std. Dev. (2D)	Standard deviation of pixel intensities in 2D.
Cell Intensity Std. Dev. (3D)	Standard deviation of voxel intensities in 3D.
NK Cell Score	Quantitative score indicating NK cell features.
Tumor Cell Score	Quantitative score indicating tumor cell features.

Table S3: Synthetic parameters table.

synthetic parameters	
num_components	8
interaction_radius	10.0
num_features_each_cell	10
distribution_low_comps	normal
sparsity_low_comps	False
T	50
n_distances_to_calculate	100
sparsity_percent	0.2
noise_low_rank_reco	0
noise_linear_reco	0
normalize_rows_columns	True
distribution_features	normal
update_of_features_if_no_effect	keep
fixed_A	False
scaling_features_method	standard
n_cell_type_rows	2
sigma_diff_graph	0.3
n_cells	35
building_w_thres	69
closest_As_to_consider	10
decor_components	True
include_plotting	True
n_unique_cells	2
with_timescales	False
with_gradient	False
interactions_style	tumor_nk
correlated_features_control	False
include_self_effect	True
distance_thres	0.168793

A.8 SUPPLEMENTARY FIGURES

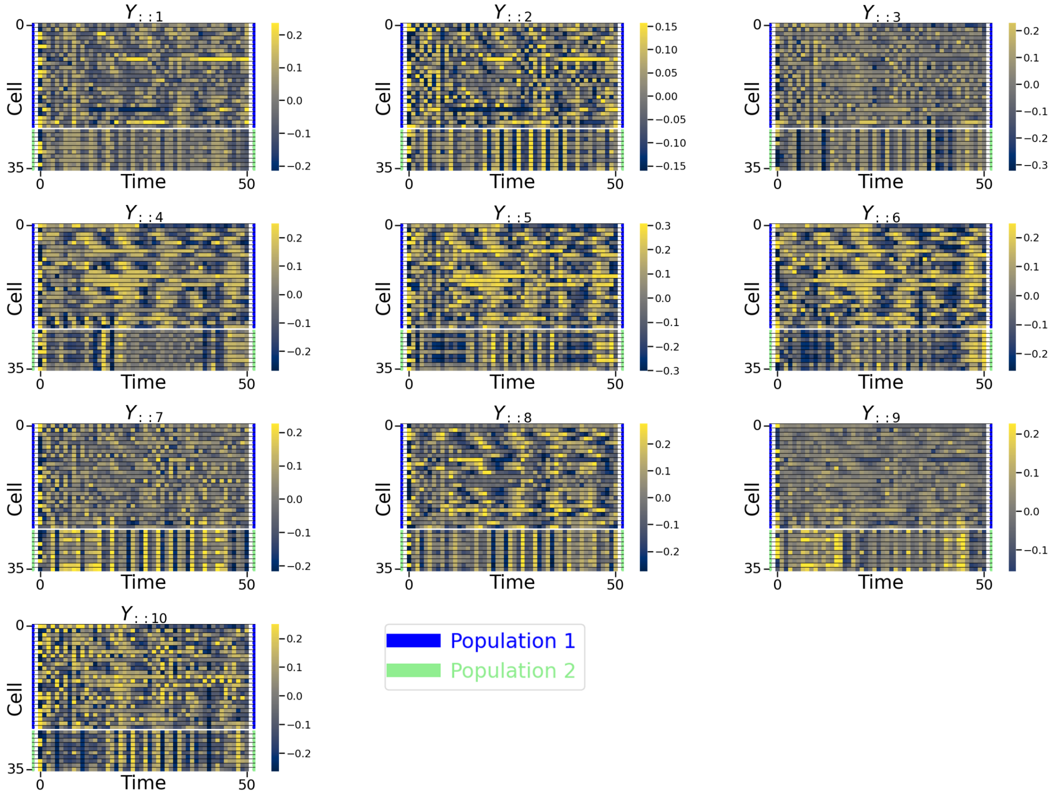

Figure S1: *Overview of the generated synthetic data.* Each subplot presents values of a given feature over all cells across time.

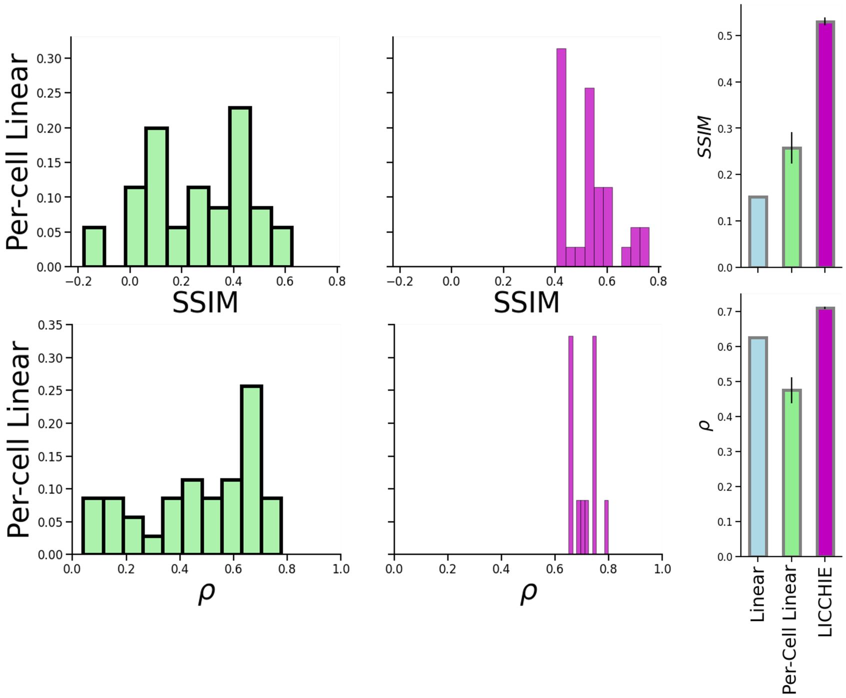

Figure S2: *Additional evaluations over synthetic data using linear dynamics.* (left) Evaluation of linear dynamics dynamics models using (1) correlation of identified interactions with the ground truth as well as (2) SSIM (Nilsson & Akenine-Möller, 2020) score. (right) comparison to `LICCHIE`.

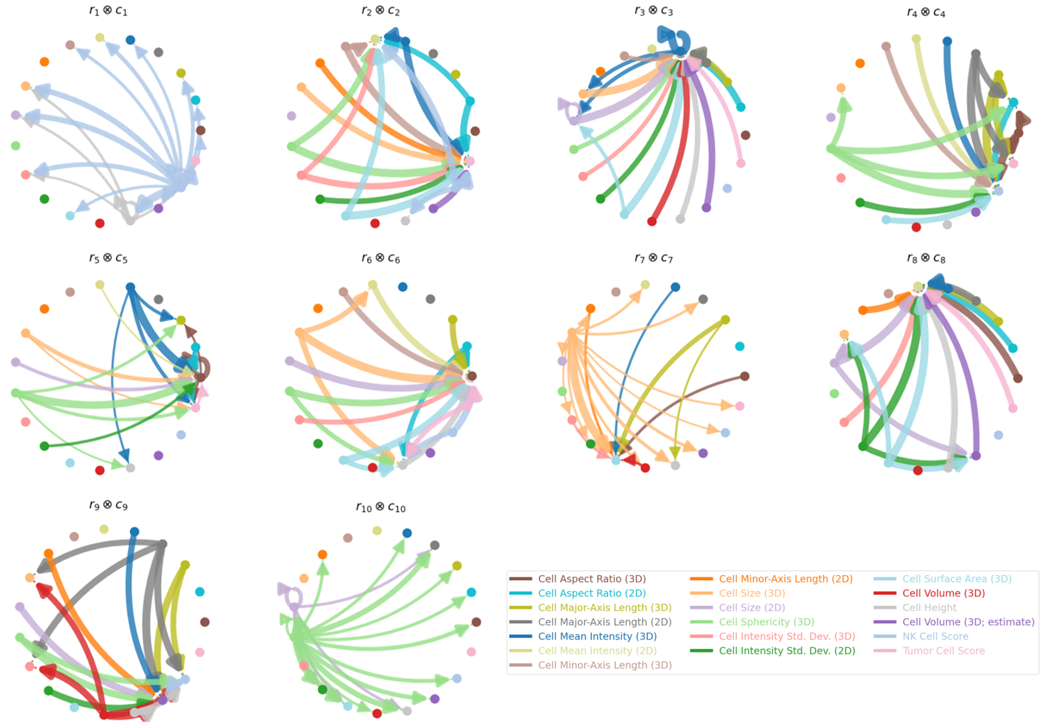

Figure S3: *Identified components by* `LICCHIE` *reveal variability in interaction structures*. Structures include *localized targets* (e.g., M_2, M_4), *source-emerging* (e.g., M_1, M_{10}), self *feedback-like* interactions (e.g., M_3) and multi-step interactions, including direct and indirect influences (e.g., M_6).

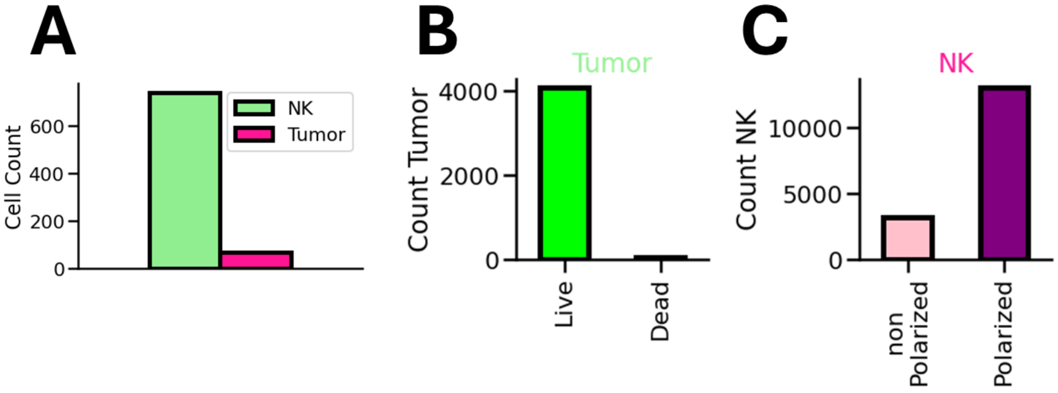

Figure S4: *Tumor-NK data overview.* **a**, Total number of observations for tumor cells (across all time points). **b**, Total number of observations for NK cells (across all time points). **c**, Number of unique NK vs tumor cells.