

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 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 **Natural Killer** (NK) cells 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-left**, Interaction matrices, $\{A_t^{(n,k)}\}$, represent 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. **Middle-right**, time-varying vector $x_t^{(n)} \in \mathbb{R}^{m_p(n)}$, captures cell-specific attributes. The interaction matrices map features of nearby cells at $t - 1$ to source feature at time t , providing temporal consistency. **Bottom**, Underlying matrices $A_t^{(n,k)}$ emerge from a set of fixed rank-1 processes representing source $\{c_j\}$ and target $\{r_j\}$ effects, modulated via flexible weights to shape the overall transition matrix $A_t^{(n,k)} = \sum_{j=1}^J W_{j,t}^{(n,k)} c_j r_j^T$, capturing complexity while maintaining parsimony.

054 Cell–cell interactions are central to the function and organization of multicellular systems (Su et al.,
 055 2024). Disruption or mis-regulation of communication is implicated in cancer, autoimmune disease,
 056 and developmental disorders (Armingol et al., 2021; Liu et al., 2024). Capturing how one cell’s
 057 state and behavior influence another in real time is challenging since interactions span spatial and
 058 temporal scales, and depend on cell type and state.
 059

060 **Studies based on static snapshots reveal a partial story.** Single-cell sequencing technologies
 061 have enabled systematic inference of putative ligand–receptor activity and co-variation patterns,
 062 offering a powerful view of communication at scale (Wilk et al., 2024; Heumos et al., 2023). How-
 063 ever, such data are typically dissociated and time-agnostic, and methods built on them suffer from
 064 an inherent limitation: they recover statistical associations from static measurements rather than the
 065 dynamics of how cells influence each other in real time (Armingol et al., 2024; 2021; Wagner et al.,
 066 2016). As a result, they fail to reveal how cells change as a result of interactions and struggle to
 067 expose directionality, temporal delays, or state-dependent effects that are essential to mechanistic
 068 understanding (Weinreb et al., 2018).
 069

070 **Live-cell imaging: a closer observational proxy for cellular dynamics.** Live-cell imaging pro-
 071 vides a means to overcome many of the challenges in studying cell–cell interactions. By tracking
 072 individual cells over time, such data allow observation of cell states before, during and after inter-
 073 actions, and thus offer a window into the dynamic consequences of cellular communication (Cuny
 074 et al., 2022). Such longitudinal data provides ground truth for understanding how cells respond to
 075 given interactions and interventions, allowing for advanced therapeutic applications (Alieva et al.,
 076 2023). Hence, there is a need for novel computational models that can: (1) detect interactions, (2)
 077 quantify their strength, and (3) reveal how they promote changes in cell state.
 078

079 **Our contributions.** We introduce **LICCHIE**—Low-rank, Interpretable Cell-Cell Hidden
 080 Interactions from Embeddings—a data-driven method to infer time-varying, feature-based inter-
 081 action rules from live-cell imaging. We represent each cell with a dynamic feature vector encom-
 082 passing spatial and internal states. Direct interactions between source–target cell pairs are modeled
 083 as temporally-evolving linear transitions over features, active only within a spatial radius between
 084 cells and regularized to vary smoothly within the feature space. **This formulation enforces tem-**
 085 **poral consistency across consecutive time-points while operating in a generalizable feature space.**
 086 To balance expressivity with interpretability, LICCHIE decomposes pairwise interaction matrices
 087 into a few rank-1 components, with time- and cell-pair-varying weights capturing context, while
 088 keeping the components biologically readable. We jointly estimate the interaction matrices, shared
 089 components, and sparse weights in an iterative approach (Fig. 1). More precisely, we present:
 090

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

100 2 BACKGROUND

101 **Representation learning and static graphs.** Deep learning architectures, such as autoencoders
 102 have been used to discover interaction-relevant structures in single-cell omics but often yield embed-
 103 dings that are hard to interpret mechanistically for interactions (Alessandri et al., 2021; Ternes et al.,
 104 2022). Graph neural networks (GNNs) encode cells as nodes and potential interactions as edges, yet
 105 typically presume a predefined (often symmetric) graph and still obscure direct, dynamic effects;
 106 attention models can highlight dependencies but do not directly encode interaction signals (Lazaros
 107 et al., 2024; Tang et al., 2023). Together, these tools are not tailored to direct, time-varying, feature-
 108 level effects required for dynamic inferences.

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

114 **Requirements for an imaging-based interpretable model.** We seek a framework that (i) mod-
 115 els direct, feature-level influences between individual cells; (ii) captures non-stationarity over time
 116 and allows for other sources of interaction variability (e.g., distances, temporal resolutions), (iii) is
 117 spatially constrained and smooth in feature space for robustness; (iv) uncovers the core interpretable
 118 structures underlying the dynamics to support biological understanding of the interactions; and (v)
 119 does not rely on persistent cell identities, by operating directly in feature space. LICCHIE is de-
 120 signed around exactly these requirements (Sec. 4, Fig. 1).

122 3 PROBLEM FORMULATION

124 **Observations and notation.** We observe a system of N interacting cells over T frames, where
 125 each cell $n = 1 \dots N$ is defined via a position vector (in 2D/3D space) $\{\psi_t^{(n)}\}_{t=1}^T \in \mathbb{R}^s$, and a
 126 multi-feature, time-varying vector $x_t^{(n)} \in \mathbb{R}^{m^{p(n)}}$ capturing cell-specific features, defined by the
 127 biological system observed. **In what follows, we relate to the multi-dimensional feature vector of**
 128 **each cell at time t as the “cell-state”.** A system can consist of one or more populations, such that
 129 $p(n)$ indexes the population type of cell n ; therefore, $m^{p(n)}$ represents the number of features for
 130 population $p(n)$ of cell n (complete notation provided in Table A1).

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

135 **Feature space learning.** Cellular features are transiently measurable; for example, certain cells
 136 may only be visible in a subset of frames due to occlusion, movement, or other limitations Fig. S6.
 137 These properties make accurate long-term tracking an ongoing challenge (Maška et al., 2023). Thus,
 138 it is desirable to avoid reliance on error-prone global trajectories and instead focus on short-horizon
 139 correspondence—predicting targets at time t are from their spatial neighbors at time $t - 1$. Cellular
 140 identity and long-term relations are then preserved by operating in feature space—learned rules
 141 generalize across cells with similar states (Fig. 1 middle).

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

148 **Balancing predictive accuracy with interpretability.** We seek a model that not only maximizes
 149 fit but also maintains parsimony through: (i) limiting the number of learned parameters, and (ii) pro-
 150 viding components that directly map to observable space, enabling mechanistic interpretation of
 151 interactions.

153 4 THE LICCHIE APPROACH

155 **Feature-evolution model.** Looking to identify how multi-cellular interactions drive changes in
 156 cell-states, we describe interactions via feature evolution—modeling the state vector $x_t^{(n)}$ of cell n
 157 at time t as the weighted sum of learnable interaction matrices capturing influences from previous
 158 time-point, namely,

$$160 \quad x_t^{(n)} = \sum_{k=1}^N \mathbf{1}_{d_{n,k}(t) \leq R} d_{n,k}(t) \mathbf{A}_t^{(n,k)} x_{t-1}^{(k)}. \quad (1)$$

162 For clarity of presentation, we hereafter omit the distance-driven reweighting. Notably, in the above
 163 formulation feature-level effects are captured by linear relations. This approach was deliberately
 164 chosen, prioritizing direct interpretability of the results—deriving understandable interaction rules
 165 over model complexity, and justified via empirical validations (App. A). **Importantly, the above**
 166 **formulation enforces temporal consistency**, requiring reconstruction of cell identity at time t based
 167 on its state and local neighborhood at $t - 1$. Together with the model design choices, described
 168 shortly, this construction ensures that the interaction matrices depict genuine cell-cell dynamics,
 169 decoupling them from transient noisy events which may be captured in the data.

170
 171 **The interaction matrices.** A matrix $\mathbf{A}_t^{(n,k)} \in \mathbb{R}^{m \times m}$ represents the effect of source cell k on
 172 target cell n at time t , conditioned on their previous states. An entry $[\mathbf{A}_t^{(n,k)}]_{i,j}$ captures the in-
 173 fluence 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.,
 174 feature i of cell n at time t , was impacted by the position, i.e., feature j , of cell k , at time $t - 1$,
 175 Fig. 1 middle). Here, we assume that all cells share the same feature set, yet the model naturally
 176 extends to population-specific features (App. B).

177
 178 **Cross-interaction similarity.** Naturally, pairs with similar stacked state vectors induce similar
 179 interaction maps. With that, if the source vector $\mathbf{x}_t^{(k)}$ lies in the null space of a matrix $\tilde{\mathbf{A}}$, then
 180 $\mathbf{A}_t^{(n,k)} + \tilde{\mathbf{A}}$ is an equivalent solution to $\mathbf{A}_t^{(n,k)}$ in terms of fit—thus applying regularization can
 181 improve inference by enforcing consistency and smoothness across parameter changes (App. C).
 182 We employ this by constraining distances between $\mathbf{A}_t^{(n,k)}$ with the respective distance in feature
 183 space,
 184

$$185 \|\mathbf{A}_t^{(n,k)} - \mathbf{A}_t^{(n',k')}\|_2^2 \leq \varepsilon(\delta_{\mathbf{A}_t^{(n,k)}, \mathbf{A}_t^{(n',k')}}), \quad (2)$$

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

189
 190 **Low-rank interaction-features regularization.** To reflect a limited number of shared biological
 191 processes, we impose a low-rank structure by expressing each $\mathbf{A}_t^{(n,k)}$ via a restricted set of global
 192 rank-1 components with pair-/time-specific weights,

$$193 \mathbf{A}_t^{(n,k)} = \sum_{j=1}^J W_{j,t}^{(n,k)} \mathbf{c}_j \mathbf{r}_j^\top = \sum_{j=1}^J W_{j,t}^{(n,k)} \underbrace{\mathbf{M}_j}_{\mathbf{c}_j \mathbf{r}_j^\top} \quad (3)$$

194 where $\{\mathbf{c}_j\}_{j=1}^J$ and $\{\mathbf{r}_j\}_{j=1}^J$ are the set of underlying components such that each $\mathbf{c}_j \in \mathbb{R}^m$ and
 195 $\mathbf{r}_j \in \mathbb{R}^m$, spans each $\{\mathbf{A}_t^{(n,k)}\}_{t,n,k}$, and each term $\mathbf{M}_j := \mathbf{c}_j \mathbf{r}_j^\top$ is a rank-1 matrix constructed
 196 via the outer product of two m -dimensional vectors. Each $\mathbf{W}^{(n,k)} \in \mathbb{R}^{J \times T}$ is the corresponding
 197 sparse per time-point weight matrix, such that $W_{j,t}^{(n,k)}$ is the contribution of components $\mathbf{c}_j, \mathbf{r}_j$ to
 198 the interaction from k to n at time t (Fig. 1 bottom).

199
 200 **The multi-feature vector.** While the set of observed features is defined by the data in hand, the
 201 induced multi-feature vector, $\mathbf{x}_t^{(n)} \in \mathbb{R}^{m_{p(n)}}$, can be defined based on the task in hand and vary from
 202 engineered, pre-defined, interpretable features based on observed attributes (e.g. spatial, kinematic
 203 and available phenotypic data, Stirling et al. (2021)) to a latent space embedding of the observation
 204 (e.g. a deep learning representation, Moshkov et al. (2024)). In what follows, we focus on the former,
 205 as it allows for direct interpretation of observed feature relations. Importantly, when considering the
 206 latter, the setting can uncover dependencies and importance of different latent dimensions, desirable
 207 for explainability of AI models (Saranya & Subhashini, 2023).

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

$$\begin{aligned}
 219 \quad \mathcal{L} = & \sum_{n=1}^N \sum_{k=1}^N \left[\|\mathbf{x}_t^{(n)} - (\mathbf{1}_{d_{n,k}(t) \leq R}) \mathbf{A}_t^{(n,k)} \mathbf{x}_{t-1}^{(k)}\|_2^2 + \lambda_1 \|\mathbf{A}_t^{(n,k)} - \sum_{j=1}^J W_{j,t}^{(n,k)} \mathbf{c}_j \mathbf{r}_j^\top\|_F^2 \right. \quad (4) \\
 220 \quad & + \lambda_2 \sum_{\substack{\mathbf{A}' \in \{\mathbf{A}_t^{(n',k')}\}_{n',k',t'} \\ \mathbf{A}' \neq \mathbf{A}_t^{(n,k)}}} \left[\|(\mathbf{1}_{d_{n,k}(t) \leq R}) e^{-\sigma \delta(\mathbf{A}_t^{(n,k)}, \mathbf{A}')} (\mathbf{A}_t^{(n,k)} - \mathbf{A}')\|_F^2 \right] \\
 221 \quad & \left. + \lambda_3 \|\text{vec}(\mathbf{A}_t^{(n,k)})\|_1 + \lambda_4 \|\mathbf{W}_{:,t}^{(n,k)}\|_1 \right], \\
 222 \quad & \\
 223 \quad & \\
 224 \quad & \\
 225 \quad & \\
 226 \quad & \\
 227 \quad & \\
 228 \quad &
 \end{aligned}$$

229 where R is the interaction radius, σ is a parameter controlling the decay of similarity with distance
 230 in feature space, $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ are scalar regularization weights that balance the different terms.
 231 $\mathbf{A}_t^{(n,k)}, \mathbf{A}'$ denote interaction matrices corresponding to feature pairs such that $\delta(\mathbf{A}_t^{(n,k)}, \mathbf{A}')$ is the
 232 distance between these pairs in the feature space. The first term ensures fidelity to the observed
 233 dynamics by minimizing prediction error, the second term enforces a low-rank structure by approx-
 234 imating each \mathbf{A} as a weighted sum of shared components, the third term encourages smoothness
 235 by penalizing differences between nearby interactions in feature space, and the two final terms im-
 236 pose sparsity on the interactions and corresponding weights (see discussion on hyper-parameters in
 237 App. D).

238 We fit the model using an iterative approach, iterating between estimating the interaction mat-
 239 rices $\{\mathbf{A}_t^{(n,k)}\}$ via LASSO Tibshirani (1996), the shared components $\{\mathbf{c}_j, \mathbf{r}_j\}$, and their weights
 240 $\{\mathbf{W}_{:,t}^{(n,k)}\}$ using PARAFAC (Harshman et al., 1970), alternating between the following steps.
 241

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

$$\begin{aligned}
 246 \quad \widehat{\mathbf{A}}_t^{(n,k)} = & \arg \min_{\mathbf{A}} \|\mathbf{x}_t^{(n)} - \sum_{k=1}^N (\mathbf{1}_{d_{n,k}(t) \leq R}) \mathbf{A} \mathbf{x}_{t-1}^{(k)}\|_2^2 + \lambda_1 \left\| \mathbf{A} - \sum_j W_{j,t} \mathbf{c}_j \mathbf{r}_j^\top \right\|_2^2 \\
 247 \quad & + \lambda_2 \sum_{\mathbf{A}' \neq \mathbf{A}} e^{-\sigma \delta(\mathbf{A}, \mathbf{A}')} \|\mathbf{A} - \mathbf{A}'\|_2^2 + \lambda_3 \|\text{vec}(\mathbf{A})\|_1 \quad (5) \\
 248 \quad & \\
 249 \quad & \\
 250 \quad & \\
 251 \quad &
 \end{aligned}$$

252 where λ_1, λ_2 , and λ_3 are hyper-parameters balancing low-rank fidelity, interaction smooth-
 253 ness, and sparsity level.

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

$$\begin{aligned}
 256 \quad \{\widehat{\mathbf{c}}_j, \widehat{\mathbf{r}}_j\}_{j=1}^J = & \arg \min_{\{\mathbf{c}_j, \mathbf{r}_j\}} \sum_{n,k,t} \left\| \mathbf{A}_t^{(n,k)} - \sum_{j=1}^J W_{j,t}^{(n,k)} \mathbf{c}_j \mathbf{r}_j^\top \right\|_2^2 \quad (6) \\
 257 \quad & \\
 258 \quad & \\
 259 \quad & \\
 260 \quad &
 \end{aligned}$$

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

$$\begin{aligned}
 264 \quad \widehat{\mathbf{W}}_{:,t}^{(n,k)} = & \arg \min_{\mathbf{W}_{:,t}^{(n,k)}} \left\| \mathbf{A}_t^{(n,k)} - \sum_{j=1}^J W_{j,t}^{(n,k)} \mathbf{c}_j \mathbf{r}_j^\top \right\|_2^2 + \lambda_4 \|\mathbf{W}_{:,t}^{(n,k)}\|_1 \quad (7) \\
 265 \quad & \\
 266 \quad & \\
 267 \quad & \\
 268 \quad &
 \end{aligned}$$

269 where λ_4 controls the sparsity of the number of active components in the $k \rightarrow n$ interaction
 270 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. E.

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

Require: Observed cell features $\{x_t^{(n)}\}$, interaction radius R , number of components J , regularization weights $\lambda_1, \lambda_2, \lambda_3$, feature-space decay σ

Ensure: Interaction matrices $\{A_t^{(n,k)}\}$, shared components $\{c_j, r_j\}$, weights W

1: Initialize $A_t^{(n,k)}, c_j, r_j, \{W_{jt}^{(n,k)}\}_{n,k,t}$ randomly from normal *i.i.d* distribution

2: **repeat**

3: Up

4: **for** each cell pair (n, k) and time t **do**

5: Infer $A_t^{(n,k)}$ ▷ equ

6: **end for**

7: Update

8: **Update weights** $\{W_{:t}^{(n,k)}\}_{n,k,t}$

9: **until** convergence or maximum iterations reached

10: **Return** $\{A_t^{(n,k)}\}_{n,k,t}$, $\{c_j, r_j\}_j$, $\{W_{:t}^{(n,k)}\}_{n,k,t}$

Q1 Q2 Q3 Q4 Q5 Q6 Q7 Q8 Q9 Q10 Q11 Q12 Q13 Q14 Q15 Q16 Q17 Q18 Q19 Q20

5 EXPERIMENTS


5.1 LICCHIE RECOVERS TRUE COMPONENTS IN SYNTHETIC DATA

Setup. We simulate two populations (25 and 10 cells, 10 features each) evolving over $T=50$ frames (Fig. 2a). Eight ($J = 8$) rank-1 ground-truth motifs $\{M_j\}_{j=1}^8$ 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 k NN ($k = 10$) 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.

Baselines. We benchmark **LICCHIE** against four alternative methods spanning linear, switching-LDS, and GNN approaches: (i) **Global linear dynamics**: a single time-fixed matrix A fit by ordinary least squares (OLS) on all cells to predict x_{t+1} from x_t ; (ii) **Per-target linear dynamics**: one OLS model per target cell n , yielding A_n ; (iii) **Per-cell SLDS**: a single switching linear dynamical system fit per trajectory; and (iv) **GNN**: a distance-weighted spatial graph trained to predict x_{t+1} from x_t . (App. H). Methods (i)–(ii) produce explicit interaction matrices that enable comparison with **LICCHIE**’s output, yet are limited in their capacity to capture complex interactions that vary in space and time. Methods (iii)–(iv) lack comparable interpretable components (App. H). This limitation also prevents comparison to alternative deep learning methods.

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 linear baselines (methods i and ii), LICCHIE produced interaction structures closer to the ground truth construction $\{M_j\}_{j=1}^8$ (Fig. 2d). Visualization of the learned components from the per-cell SLDS (method iii, Fig. S3) and the GNN (method iv, Fig. S4) depicts that these approaches are ill-suited for the targeted task, with the former constrained by abrupt switching transitions (Fig. S3a), while the latter requires additional processing to approximate transitions via uninterpretable dense loadings (Fig. S4b,c). At last, LICCHIE

324 obtained the most accurate predictions for the interaction matrices, compared to all baselines, con-
 325 sidering the reconstruction error (Fig. 2e,f) and correlation with ground-truth (Fig. S2).
 326



346 **Figure 2: Synthetic data overview and results.** **a**, Examples of generated features illustrating popula-
 347 tion differences (blue/green side bars). **b**, Ground-truth weights show sparse structure. **c**, Accuracy
 348 results for LICCHIE’s cell type prediction using the inferred interactions (logistic regression with
 349 5-fold cross-validation); predictions of source and target (0.83), or source (0.86) and target (0.87)
 350 independently. **d**, Rank-1 motif inference (top row; left to right: linear, per-cell linear, LICCHIE)
 351 vs. ground truth (bottom row). Red arrows highlight discrepancies between ground truth (dotted)
 352 and baseline-identified networks (solid). **e**, Per-pair MSE distributions vs. ground truth for the per-
 353 cell linear baseline (top) and LICCHIE (bottom). **f**, Average MSE for the global time-fixed linear
 354 baseline, per-cell dynamics, per-cell SLDS, GNN, and LICCHIE (additional metrics in Fig. S2)

355 5.2 LICCHIE UNCOVERS INTERPRETABLE TUMOR–NK INTERACTIONS

356 **Data and biological relevance.** As our real-world application, we analyze patient-derived gas-
 357 tric tumor organoids co-cultured with primary human NK cells in 3D time-lapse confocal imaging
 358 (8–12 h, 4-min cadence, 27–30 z-slices per frame; Liu et al. (2024)). Organoid co-cultures pro-
 359 vide an established platform to study interactions between motile immune effectors and tumor cells
 360 in a controlled 3D microenvironment (Polak et al., 2024). These platforms allow investigation of
 361 contact-resolved immune synapses where brief, often serial engagement events underlie cytotoxic
 362 action (Vanherberghen et al., 2013; Dekkers et al., 2023) and imaging readouts from such co-culture
 363 conditions have proven translational value (Alieva et al., 2023). We hence apply LICCHIE to study
 364 time-varying, local influences between tumor-immune co-cultures.

365 **The extracted features.** As noted by Liu et al. (2024), cells in the data presenting variability in
 366 cell morphology and intensity distribution (Fig. 3a). Hence, we extracted spatial, morphology, and
 367 intensity features, relying on cell-level segmentation masks from a representative experiment. We
 368 normalize the features such that each feature, at population level (NK or tumor cells), is normally
 369 distributed, $f_i^{(p)} \sim \mathcal{N}(\mu = 0, \sigma = 1)$, implying that all features, although representing geometric
 370 properties, attain negative values (App. F, Table A2). We fitted LICCHIE with $J=10$ components.

372 **Components and structure.** LICCHIE identified components that reveal diverse interaction pat-
 373 terns, including *localized targets* (e.g., M_2, M_4), *source-emerging* (e.g., M_1, M_{10}), *self feedback-like*
 374 interactions (e.g., M_3) and multi-step interactions, including direct and indirect influences (e.g.,
 375 M_6 , Fig. S7).

377 Interaction motifs can capture indicators of cell state, such as NK cell polarization. Upon polar-
 378 ization, NK cells undergo a characteristic morphological shift from round to elongated (Fig. 3a), a

process necessary for their cytotoxic activity. Analysis of motif 2 (M_2) and the respective weights (\mathbf{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 (M_5) has a sparse interaction matrix (Fig. 3b) and predominantly captures NK-to-tumor interactions (\mathbf{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, 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 (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 (M_7) highlights the centrality of source cell size in shaping multiple target features, and the respective weights (\mathbf{W}_7) indicate that 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.




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. S7). **c**, Component weights by interaction class (tumor \rightarrow tumor, NK \rightarrow NK, NK \rightarrow tumor, tumor \rightarrow NK) show class-specific activation patterns.

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

- *Opposing weights between cell classes.* Component 6 (weighted by \mathbf{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 \mathbf{W}_6 values for different-cell pairs.
- *Single interaction dominance.* Components 1 and 6 (\mathbf{W}_1 , \mathbf{W}_6), show clear dominance of a single interaction type (e.g., NK \rightarrow NK), which may reflect specialized functional role.


432

- 433 • *Source disentangling*. Components 1 and 2 (W_1, W_2) capture predominant interactions
- 434 targeting tumor and NK cells respectively, with opposing weight signs given the source.
- 435 • *Cell-specific regulation*. Component 3 (W_3) is mainly associated with interactions between
- 436 tumor and NK cells (rather than homogeneous cell-type interactions).

437 5.3 FROM MOTIFS TO PHENOTYPES: USING ENGINEERED CELLS FOR PREDICTION

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

444 **Procedure.** For each M_j , $j = 1 \dots J$ component, we gradually increased its weight from 0 to
 445 2 in 20 steps and applied it to various combinations of source features. We then assigned each
 446 inferred target a polarization score (for NK cells) or a death score (for tumor cells) using a distance-
 447 weighted k -nearest neighbor classifier ($k=150$, total # samples from all time points = 20, 437) based
 448 on Euclidean distance in feature space and normalized by phenotype frequency (Fig. 4a).



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

476 **Predictive evaluation against baselines.** The biological relevance of models’ outputs is tested
 477 by assessing their predictive power to known biological properties. Comparing `LICCHIE` to two
 478 baseline methods (SLDS and GNN; Sec. 5.1), we evaluate accuracy of cell type prediction, inter-
 479 action types, and phenotypes, directly from the interaction matrices. These tests demonstrate that
 480 `LICCHIE` is the only method whose outputs could directly be used to obtain meaningful predictions
 481 (App. I, Fig. S8, Fig. S9, Fig. S10).

483 **Predictions from inferred weights.** Using weights inferred by `LICCHIE` on unseen frames, we
 484 trained a cross-validation logistic regression to predict target phenotypes. Although phenotypes
 485 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.

486 **Cell status relations in interactions.** Focusing on **LICCHIE**’s interpretable components we notice
 487 they yield distinct death/polarization score trajectories with non-linear changes as weights increased
 488 (Fig. 4a). For tumor-cell death, some components (e.g., M_1, M_5, M_6 , Fig. S7) showed increasing
 489 scores with weight, whereas others (e.g., M_2, M_3, M_8) showed decreases, often most prominently
 490 around weight ≈ 1 . Overall, tumor death scores exhibited phase-transition-like effects
 491 (non-smooth changes), consistent with the acute nature of cell death. NK-cell polarization trends
 492 were more variable; for example, M_7 increased polarization at higher weights, while M_5 decreased
 493 it. In some cases (e.g., M_5), tumor-death and NK-polarization scores moved in opposite directions
 494 as weight increased.

495 **Source-feature extremes.** We also examined the impact of source features on NK polarization
 496 under maximal activation ($W=2$). For each source feature, we compared scores at its empirical
 497 minimum vs. maximum (others fixed at their means). We found that source properties give rise to
 498 both diversity and correlations in the potential effect of each component. For example, when com-
 499 ponent M_1 is applied to a source feature vector with a minimized cell aspect ratio, it increases the
 500 polarization score of the target. In contrast, when applied to the maximum cell aspect ratio, the
 501 target cell’s polarization score remains near zero. Other components, such as M_4 or M_9 , show
 502 similar target polarization scores under both minimized and maximized cell aspect ratio values, sug-
 503 gesting that this feature may not be critical in these circuits for driving changes in target polarization
 504 (Fig. 4c). Altogether, these results highlight the potential of **LICCHIE** to guide the design of future
 505 interventions when applied to additional datasets.

507 6 DISCUSSION AND FUTURE WORK

508 We presented **LICCHIE**, an interpretable method to study cell-cell interactions capturing modular,
 509 temporally consistent, feature-specific cell-cell interactions. The PARAFAC rank-1 decomposition
 510 yields interpretable source-target components, making it explicit which sources influence which
 511 targets rather than obscuring them in complex, opaque patterns. The cell state feature representa-
 512 tions are user defined; allowing for problem specific optimization. Over these, **LICCHIE** reveals
 513 interpretable motifs capturing features role as targets or sources within modular cell-cell interac-
 514 tions through coordinated, overlapping, and co-modulatory effects. **LICCHIE**’s generality enables
 515 use across experimental systems—from isolated cells in a dish to complex 3D models—capturing
 516 the distinct cell-cell communication aspects they exhibit. The observed system shall guide the in-
 517 terpretations of the inferred interaction motifs. Applied to live-cell imaging data of tumor–NK
 518 co-cultures, **LICCHIE** revealed key patterns of cellular communication—identified components with
 519 distinct feature relations—revealing the modular biological regulation that integrates diverse features
 520 into phenotypic outcomes. An important advantage of **LICCHIE** is that it can address the challenge
 521 of modeling interactions from separate, asynchronous measurements while preserving data-specific
 522 details. Its global vectors c_j and r_j allow new data to be analyzed quickly by reusing the compo-
 523 nents and adjusting only the weights, avoiding full-model retraining. Notably, **LICCHIE** naturally
 524 handles data imbalance by its ability to re-weight interactions based on the prevalence of each pop-
 525 ulation and interaction type, ensuring adequate representation for all interactions.

526 **Limitations and future directions.** The model assumes linear or additive effects, which may not
 527 fully capture certain non-linear dynamics; this design choice promotes interpretability, and found
 528 valid in practice through reconstruction accuracy (see discussion in App. A). Future work can in-
 529 incorporate non-linear activation functions $f(\cdot)$ to each interaction pair to improve flexibility. Here,
 530 interactions beyond the defined radius are not captured; this limitation can be easily addressed by ap-
 531 plying a kernel and sampling from an appropriate probability distribution that varies with distance.
 532 Further biological interpretation of components requires additional supervision via experimental
 533 data that can be obtained by applying **LICCHIE** to additional systems with dedicated measurements.
 534 In addition, future experimental work coupled with **LICCHIE** could establish the mechanisms and
 535 causal relationships underlying the identified interaction motifs, for example through targeted per-
 536 turbations along effector-target axes. Moreover, by isolating the morphological features underly-
 537 ing such interaction motifs **LICCHIE** provides a framework to link dynamic imaging phenotypes
 538 with molecular information by layering in spatial and single-cell multi-omics in subsequent studies.
 539 Lastly, using temporal mapping tools (e.g., optimal transport, Klein et al. (2025)), **LICCHIE** can
 be applied to spatio-temporal single-cell datasets, providing insights into interactions over a richer
 cell-state representation space, which may be more directly applicable for therapeutic applications.

540 REFERENCES
541

542 Luca Alessandri, Francesca Cordero, Marco Beccuti, Nicola Licheri, Maddalena Arigoni, Martina
543 Olivero, Maria Flavia Di Renzo, Anna Sapino, and Raffaele Calogero. Sparsely-connected au-
544 toencoder (sca) for single cell rnaseq data mining. *NPJ systems biology and applications*, 7(1):1,
545 2021.

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

548 Tomas Andersson, Lars Alfredsson, Henrik Källberg, Slobodan Zdravkovic, and Anders Ahlbom.
549 Calculating measures of biological interaction. *European journal of epidemiology*, 20(7):575–
550 579, 2005.

552 Erick Armingol, Adam Officer, Olivier Harismendy, and Nathan E Lewis. Deciphering cell–cell
553 interactions and communication from gene expression. *Nature Reviews Genetics*, 22(2):71–88,
554 2021.

556 Erick Armingol, Hratch M Baghdassarian, and Nathan E Lewis. The diversification of methods
557 for studying cell–cell interactions and communication. *Nature Reviews Genetics*, 25(6):381–400,
558 2024.

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

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

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

566 Yenho Chen, Noga Mudrik, Kyle A Johnsen, Sankaraleengam Alagapan, Adam S Charles, and
567 Christopher Rozell. Probabilistic decomposed linear dynamical systems for robust discovery of
568 latent neural dynamics. *Advances in Neural Information Processing Systems*, 37:104443–104470,
569 2024.

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

573 Johanna F Dekkers, Maria Alieva, Astrid Cleven, Farid Keramati, Amber KL Wezenaar, Esmée J
574 van Vliet, Jens Puschhof, Peter Brazda, Inez Johanna, Angelo D Meringa, et al. Uncovering the
575 mode of action of engineered t cells in patient cancer organoids. *Nature biotechnology*, 41(1):
576 60–69, 2023.

578 Matthias Fey, Jinu Sunil, Akihiro Nitta, Rishi Puri, Manan Shah, Blaz Stojanovic, Ramona Ben-
579 dias, Alexandria Barghi, Vid Kocijan, Zecheng Zhang, Xinwei He, Jan Eric Lenssen, and Jure
580 Leskovec. PyG 2.0: Scalable Learning on Real World Graphs. *Temporal Graph Learning Work-
581 shop @ KDD*, 2025.

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

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

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

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

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

597

598 Scott Linderman, Benjamin Antin, David Zoltowski, and Joshua Glaser. SSM: Bayesian Learn-
 599 ing and Inference for State Space Models, October 2020. URL <https://github.com/lindermanlab/ssm>.

600

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

604

605 Martin Maška, Vladimír Ulman, Pablo Delgado-Rodriguez, Estibaliz Gómez-de Mariscal, Tereza
 606 Nečasová, Fidel A Guerrero Peña, Tsang Ing Ren, Elliot M Meyerowitz, Tim Scherr, Katharina
 607 Löffler, et al. The cell tracking challenge: 10 years of objective benchmarking. *Nature Methods*,
 608 20(7):1010–1020, 2023.

609

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

613

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

617

618 Noga Mudrik, Ryan Ly, Oliver Ruebel, and Adam S Charles. Creimbo: Cross-regional ensemble
 619 interactions in multi-view brain observations. The International Conference on Learning Repre-
 620 sentations, 2025.

621

622 Josue Nassar, Scott W Linderman, Monica Bugallo, and Il Memming Park. Tree-structured recurrent
 623 switching linear dynamical systems for multi-scale modeling. *arXiv preprint arXiv:1811.12386*,
 624 2018.

625

626 Jim Nilsson and Tomas Akenine-Möller. Understanding ssim. *arXiv preprint arXiv:2006.13846*,
 627 2020.

628

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

631

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

634

635 Roger Penrose. On best approximate solutions of linear matrix equations. In *Mathematical Pro-
 636 ceedings of the Cambridge Philosophical Society*, volume 52, pp. 17–19. Cambridge University
 637 Press, 1956.

638

639 Roel Polak, Elisa T Zhang, and Calvin J Kuo. Cancer organoids 2.0: modelling the complexity of
 640 the tumour immune microenvironment. *Nature Reviews Cancer*, 24(8):523–539, 2024.

641

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

644

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

647

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

651

652 Wenzhuo Tang, Hongzhi Wen, Renming Liu, Jiayuan Ding, Wei Jin, Yuying Xie, Hui Liu, and
 653 Jiliang Tang. Single-cell multimodal prediction via transformers. In *Proceedings of the 32nd*
 654 *ACM International Conference on Information and Knowledge Management*, pp. 2422–2431,
 655 2023.

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

651

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

654

655 Bruno Vanherberghe, Per E Olofsson, Elin Forslund, Michal Sternberg-Simon, Mohammad Ali
656 Khorshidi, Simon Pacouret, Karolin Guldevall, Monika Enqvist, Karl-Johan Malmberg, Ramit
657 Mehr, et al. Classification of human natural killer cells based on migration behavior and cytotoxic
658 response. *Blood, The Journal of the American Society of Hematology*, 121(8):1326–1334, 2013.

659

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

662

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

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702
703
704
705
706
Appendix707
708
709
710
711
712
A ON THE LINEARITY ASSUMPTION713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
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 (Brauer & Kribs, 2016; Mudrik et al., 2025; Andersson et al., 2005).726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
A unique advantage of maintaining a linear transition matrix \mathbf{A} in biological systems is interpretability: in our framework each entry \mathbf{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.726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
To account for non-linear dynamics one can consider non-linear activations, e.g. $\mathbf{x}_t = \sigma(\mathbf{A}_t)\mathbf{x}_{t-1}$. However, including these prevents direct interpretation of the results, necessary for biological discovery.726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
B POPULATION-SPECIFIC FEATURE SPACE726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
The LICCHIE framework can naturally extend to account for multiple cell populations exhibiting distinct feature subsets. To do so, the fixed component sets $\{c_j\}$ and $\{r_j\}$ are replaced by population-specific sets $\{c_j^{(p)}\}_{j=1}^J$ and $\{r_j^{(p)}\}_{j=1}^J$, where p indexes the population. Within each set, the vector dimensions are consistent, but they may differ between populations.726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
For within-population interactions (e.g., population p), the \mathbf{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 $\{r_j\}$ from p and $\{c_j\}$ from p' , and vice versa. The resulting interaction matrix may not be square if the number of features differs between populations.726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
Notably, the above extension preserves the interpretability afforded by the component structure while allowing flexibility across nuanced population dynamics.726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
C AVOIDING EQUIVALENT SOLUTIONS THROUGH DISTANCE CONSTRAINTS726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
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 \mathbf{A} matrices that represent similar source-target feature distributions.726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
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 $\mathbf{A}_t^{(n,k)}$ and $\mathbf{A}_t^{(n',k')}$ 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 $\tilde{\mathbf{A}}$, then $\mathbf{A}_t^{(n,k)} + \tilde{\mathbf{A}}$ is an equivalent solution to $\mathbf{A}_t^{(n,k)}$ in terms of fit. This means that we must identify a solution that is consistent across pair, enforcing smoothness across parameter changes.726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
Hence, we introduced the distance constraint; forcing distances between $\{\mathbf{A}\}$ matrices to follow distances in the corresponding feature space, i.e.,

726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
$$\|\mathbf{A}_t^{(n,k)} - \mathbf{A}_t^{(n',k')}\|_2^2 \leq \epsilon(\delta_{(n,k),(n',k')}(t))$$

726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
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 .

756 **D OBJECTIVE HYPER-PARAMETERS**
757758 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
759 addition, the framework provides the user with the flexibility to set the maximum rank of each interaction,
760 J , and the effective radius of interaction R ; both are quantities that should be set with
761 respect to the characteristics of the biological system studied, i.e., accounting for the total number
762 of features and the representative length scale.763 In practice, hyper-parameters can be tuned by running a parameter search (e.g., grid search) while
764 optimizing an information criterion that balances degrees of freedom and model complexity, such as
765 AIC (Akaike Information Criterion) or BIC (Bayesian Information Criterion, Chakrabarti & Ghosh
766 (2011)). We provide guidance and further intuition for tuning these,767
768

- **Low-rank constraint** (J, λ_1): These control how strongly the \mathbf{A} 's follow the low-rank
769 approximation. When J is larger, or when \mathbf{A} is naturally close to low-rank, this term
770 should converge to 0, and the sensitivity to λ_1 will be minimal (since its multiplier is close
771 to zero). For highly complex, high-dimensional, and fast-varying dynamics, the low-rank
772 Frobenius norm may not be small. In such cases, we recommend monitoring the rank-1
773 reconstruction over iterations (calculated automatically within the model). If the low-rank
774 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
775 preventing null-space solutions from causing abrupt switches (see App. C), is controlled
776 by λ_2 . In practice, a small value $\lambda_2 < 1$ is usually sufficient. We recommend examining
777 the smoothness of interactions under gradually changing parameters to confirm that the
778 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
779 zero or near-zero entries. λ_3 promotes sparsity in the interaction matrices, facilitating a
780 clearer understanding of the functionality of each interaction. λ_4 encourages interactions
781 to be composed a limited number of $\{c_j\}$ and $\{r_j\}$ components, valuable for,
782
 1. Understanding the modular role of each $\{c_j, r_j\}$ pair in driving the overall interaction,
 2. Identifying commonalities and differences between interaction types (if all compo-
783 nents 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
784 interaction are considered, i.e., cells within this disk are considered interacting with it.
785 This shall be set with respect to the system's length scales and desired scope of interactions
786 one wishes to study. e.g., one can choose to limit the analysis to first-order interactions,
787 accounting for cells in immediate neighborhood or larger value for long-range effects.

788789 **E OPTIMIZATION AND COMPLEXITY**
790791 To fit the model's components we use LASSO (Tibshirani, 1996), estimating the interaction
792 matrices, $\{\mathbf{A}_t^{(n,k)}\}$, and Parallel Factor Analysis (PARAFAC, Harshman et al. (1970)) for the shared
793 components $\{c_j, r_j\}$, and their weights $\{\mathbf{W}_{:t}^{(n,k)}\}$. PARAFAC is a method to decompose high-
794 dimensional tensor data (multi-way) into underlying, independent components, extending the
795 Principal Component Analysis (PCA) model to more than two dimensions. This provides a unique
796 solution that allows for the recovery of pure components.797 The complexity of the optimization is as follows; let N denote the number of cells, T the number of
798 time points, d the number of features, k the average number of neighbors (within radius), $S = k - 1$
799 the number of sources, R the number of outer iterations, C the number of PARAFAC components,
800 I the number of PARAFAC ALS iterations, M the number of past \mathbf{A} 's maintained for similarity
801 search, and $E \approx T \cdot N \cdot k$ the total number of interactions.802 **Pre-processing (neighbors within radius).**
803804

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

810 • With a spatial index (grid/kd-tree): build + query costs $O(T \cdot N \log N + T \cdot E)$ time. The
 811 neighbor graph requires $O(E)$ memory.
 812

813 **Per target per time (one cell as target).**

814 • Build wide $\mathbf{A} \in \mathbb{R}^{d \times (S \cdot d)}$: $O(S \cdot d^2)$.
 815 • Least-squares solve $x_{\text{target}} = \mathbf{A} x_{\text{sources}}$: $O(S^2 \cdot d^3)$ time, $O(S \cdot d^2)$ memory.
 816 • Maintain M most similar past \mathbf{A} 's: similarity search over history $H \approx R$, costing $O(H \cdot$
 817 $S \cdot d^2)$ (naive).
 818 • Decompose each \mathbf{A} (reshaped as a tensor $d \times S \times d$) with PARAFAC/CP: ALS per \mathbf{A} costs
 819 $O(I \cdot C \cdot S \cdot d^2)$ time, $O(C \cdot (d + S + d))$ memory.
 820

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

824 • LS solves: $O(T \cdot N \cdot S^2 \cdot d^3)$.
 825 • Similarity searches: $O(T \cdot N \cdot R \cdot S \cdot d^2)$ (if $H \approx R$).
 826 • PARAFAC: $O(T \cdot N \cdot I \cdot C \cdot S \cdot d^2)$.
 827

828 **Overall complexity.**

829
$$O(T \cdot N \cdot (S^2 d^3 + (R + I \cdot C) S d^2 + \log N + k)) .$$

832 **F THE TUMOR-NK DATASET**

834 We used publicly available data from (Liu et al., 2024), accessed via Zenodo. The complete dataset
 835 includes 3D live-cell imaging datasets of gastric tumor organoids co-cultured with primary human
 836 Natural Killer (NK) cells. The dataset includes multiple sessions recorded asynchronously under
 837 varying experimental conditions (see (Liu et al., 2024) for more details). We selected a representative
 838 random session, ‘GX048-TO + NK cell (IL-15)’, in which NK cells were stimulated with IL-15.
 839

840 **G TABLES**

843 Table A1: Summary of defined notations.

844 Notation	845 Meaning
846 N	846 Number of interacting cells
847 T	847 Total duration of observation (number of time points)
848 s	848 Dimensionality of the space $\mathcal{R} \subset \mathbb{R}^s$
849 $\psi_t^{(n)}$	849 Position of cell n at time t in \mathbb{R}^s (spatial coordinates)
850 $\mathbf{x}_t^{(n)}$	850 State (feature) vector of cell n at time t in $\mathbb{R}^{m^{p^{(n)}}}$
851 $p^{(n)}$	851 Population label/type of cell n
852 $\mathcal{P} = \{p^{(n)}\}_{n=1}^N$	852 Set of population identifiers for all cells
853 $m^{p^{(n)}}$	853 Number of features for population $p^{(n)}$; if constant across all cells, denote as m
854 $m_{\text{spatial}}^{p^{(n)}}$	854 Number of spatial features for population $p^{(n)}$
855 $m_{\text{internal}}^{p^{(n)}}$	855 Number of internal features for population $p^{(n)}$
856 $d_{n,n'}(t)$	856 Spatial distance between cells n and n' at time t , e.g., $\ \mathbf{r}_t^{(n)} - \mathbf{r}_t^{(n')}\ $
857 $\delta_{n,n'}(t)$	857 Feature distance between cells n and n' at time t , e.g., $\ \mathbf{x}_t^{(n)} - \mathbf{x}_t^{(n')}\ $
858 $\delta_{(n,k),(n',k')}(t)$	858 Feature distance between stacked state vectors of cell pairs (n, k) and (n', k') at time t .
859 $\delta_{A,A'}(t)$	859 Feature distance between A s
860 R	860 Interaction radius (hyperparameter)
861 $\mathbf{1}_{d_{n,n'}(t) \leq R}$	861 Indicator: 1 if distance between n and n' at time t is $\leq R$, else 0

864

865

Table A2: Summary of cell features.

866

867

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.

868

869

870

871

872

873

874

875

876

877

878

879

880

881

882

883

884

885

886

Table A3: Synthetic parameters table.

887

888

889

890

891

892

893

894

895

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

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

918

H BASELINE IMPLEMENTATION DETAILS

919

920

H.1 LINEAR DYNAMICAL SYSTEMS (LDS)

921

922 We fit a single, time-fixed feature-to-feature map $\mathbf{A} \in \mathbb{R}^{F \times F}$ by least squares over *all* ordered
923 source-target cell pairs and time points. The model predicts target features at the next frame from
924 source features at the current frame:

925
$$\hat{\mathbf{A}} = \arg \min_{\mathbf{A}} \sum_{(s, n, t): s \neq n} \left\| \mathbf{x}_{t+1}^{(n)} - \mathbf{A} \mathbf{x}_t^{(s)} \right\|_2^2,$$
926
927

928 with no intercept term and no spatial/radius filtering. Let $\mathcal{I} = \{(s_i, n_i, t_i)\}_{i=1}^M$ index sampled
929 ordered pairs and times, and define

930
$$\mathbf{X} = [\mathbf{x}_{t_1}^{(s_1)} \ \mathbf{x}_{t_2}^{(s_2)} \ \cdots \ \mathbf{x}_{t_M}^{(s_M)}] \in \mathbb{R}^{F \times M}, \quad \mathbf{Y} = [\mathbf{x}_{t_1+1}^{(n_1)} \ \mathbf{x}_{t_2+1}^{(n_2)} \ \cdots \ \mathbf{x}_{t_M+1}^{(n_M)}] \in \mathbb{R}^{F \times M}.$$
931
932

933 The least-squares estimate solves $\min_{\mathbf{A}} \|\mathbf{Y} - \mathbf{A} \mathbf{X}\|_F^2$ with closed-form

934
$$\mathbf{A} = \mathbf{Y} \mathbf{X}^\dagger,$$
935

936 where \mathbf{X}^\dagger is the Moore–Penrose pseudoinverse (Penrose, 1956), estimated by ‘sklearn.linalg’ pack-
937 age. $\mathbf{X}^\dagger = \mathbf{X}^\top (\mathbf{X} \mathbf{X}^\top)^{-1}$, hence $\mathbf{A} = \mathbf{Y} \mathbf{X}^\top (\mathbf{X} \mathbf{X}^\top)^{-1}$. This baseline yields a single, global,
938 time-fixed map shared across all pairs and frames. We included two versions of LDSs:

939

940 - Applied across all cells together (such that all cells can interact with all cells).
941 - Applied per cell, to enable higher expressivity, while limiting cross-cell interactions.

942

943

H.2 SWITCHING LINEAR DYNAMICAL SYSTEM (SLDS)

944

945 In SLDS, the system can be in one of K discrete states (modes), and in each mode, the dynamics
946 follow a linear state space model. Particularly, at each time step t :

947

948 1. **Discrete state (switch variable):**

949

950 - $z_t \in \{1, 2, \dots, K\}$ - which dynamical mode we are in
951 - Evolves as a Markov chain: $P(z_t | z_{t-1})$

952 *In LICCHIE however, multiple processes can be co-active simultaneously through the*
953 *decomposition $\mathbf{A}_t^{(n,k)} = \sum_{j=1}^J W_{j,t}^{(n,k)} \mathbf{c}_j \mathbf{r}_j^\top$.*

954

955 2. **Continuous latent state:**

956

957 - $\mathbf{x}_t \in \mathbb{R}^d$ - hidden continuous state
958 - In SLDS, dynamics depend on current mode z_t :

959
$$\mathbf{x}_t = \mathbf{A}_{z_t} \mathbf{x}_{t-1} + \mathbf{B}_{z_t} \mathbf{u}_t + \mathbf{w}_t \tag{8}$$

960 where $\mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_{z_t})$.

961 *In contrast LICCHIE captures explicit pairwise interactions: $\mathbf{x}_t^{(n)} =$*
962 *$\sum_{k=1}^N \mathbb{1}_{d_{n,k}(t) \leq R} \mathbf{A}_t^{(n,k)} \mathbf{x}_{t-1}^{(k)}$, with spatial constraints and cell-specific matrices.*

963

964 3. **Observations:**

965
$$\mathbf{y}_t = \mathbf{C}_{z_t} \mathbf{x}_t + \mathbf{D}_{z_t} \mathbf{u}_t + \mathbf{v}_t \tag{9}$$

966 where $\mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_{z_t})$

967

968 - \mathbf{A}_{z_t} : Transition matrix for mode z_t in SLDS
969 - \mathbf{B}_{z_t} : Input matrix
970 - \mathbf{C}_{z_t} : Emission/observation matrix
971 - $\mathbf{Q}_{z_t}, \mathbf{R}_{z_t}$: Process and observation noise covariances
- \mathbf{u}_t : Optional control input

972 *In contrast, `LICCHIE`’s rank-1 underlying vectors $\mathbf{c}_j, \mathbf{r}_j$ represent interpretable source and target*
 973 *effects respectively.*

974 In SLDS, the discrete state follows a Markov chain with transition matrix:

$$P(z_t = j | z_{t-1} = i) = \pi_{ij} \quad (10)$$

975 This allows the system to switch between different linear dynamics.

976 *In contrast, `LICCHIE`’s interactions evolve smoothly through distance-weighted regularization*
 977 $\|\mathbf{A}_t^{(n,k)} - \mathbf{A}_{t'}^{(n',k')}\|_2^2 \leq \epsilon(\delta)$, *avoiding discrete jumps.*

978 At last, when SLDS is applied to feature space, it learns feature-to-feature transitions but cannot dis-
 979 tinguish within-cell evolution from cross-cell interactions without scaling to dimension $\mathbb{R}^{N^m \times N^m}$.

980 We implemented Switching Linear Dynamical System (SLDS) baselines using the `ssm` pack-
 981 age (Linderman et al., 2020). The SLDS model was configured with $K = 10$ discrete latent states to
 982 match `LICCHIE`’s number of components ($J = 10$), and the latent dimensionality D was set equal
 983 to the number of observed features. We used Gaussian identity emissions (emissions=’gaussian_id’)
 984 to directly model the observed feature space without additional transformations. The model was fit
 985 using black-box variational inference (BBVI) with a mean-field variational posterior, running for
 986 500 iterations.

987 We fitted an independent SLDS for each cell trajectory (interactions are implicit within that cell’s
 988 feature space; Fig. S3)

989 Notably, since SLDS does not produce components that can co-occur at the same time point, it
 990 cannot separate the evolution of multiple driving forces acting together, nor does it account for unit-
 991 to-unit (e.g., cell-to-cell) distance. In addition, since we applied SLDS in feature space—where
 992 each latent transition represents feature-to-feature effects (i.e., nodes are features, not cells)—it
 993 does not naturally distinguish interactions within the same cell from interactions across different
 994 cells. Extending SLDS to capture these aspects would require increasing the number of nodes to
 995 $n \times k$, (n the number of cells and k the number of features), which would be extremely large and
 996 thus hinders interpretability.

1001 H.3 GRAPH NEURAL NETWORKS (GNNs)

1002 To contextualize `LICCHIE` within the broader landscape of methods for spatiotemporal cellular
 1003 data, we compared against a Graph Neural Network (GNN) baseline implemented in PyTorch Geo-
 1004 metric (Fey et al., 2025). Each frame t is represented as a graph whose nodes are cells and whose
 1005 edges connect spatial neighbors with distance-weighted strengths ($w_{n,k}(t) = \exp(-d_{n,k}(t)/R)$,
 1006 with R equivalent to the distance used for `LICCHIE`). A message-passing network is trained to
 1007 predict the feature set x_t from x_{t-1} by minimizing mean-squared error using Adam.

1008 We construct a GNN with 3 graph convolutional layers, 64 hidden channels, ReLU activations, batch
 1009 normalization, and dropout ($p = 0.1$). Training was performed with Adam optimizer at learning rate
 1010 0.001, batch size 16, and 100 epochs with early stopping (patience 20). Training was performed on
 1011 GPU using PyTorch 2.0+ and PyTorch Geometric 2.3+.

1012 Table A4 summarizes the key differences between the GNN approach and `LICCHIE`. Broadly,
 1013 GNNs learn implicit interactions through message passing in a black-box neural network— they
 1014 achieve high expressivity through non-linear transformations but require post-hoc explanation meth-
 1015 ods to interpret learned representations (Fig. S4). In contrast, `LICCHIE` explicitly recovers inter-
 1016 pretable interaction matrices $\mathbf{A}_t^{(n,k)}$ where each entry $[\mathbf{A}_t^{(n,k)}]_{i,j}$ quantifies the effect of source
 1017 feature j and cell k on target feature i of cell n . Further, `LICCHIE`’s low-rank decomposi-
 1018 tion $\mathbf{A}_t^{(n,k)} = \sum_{j=1}^J W_{j,t}^{(n,k)} \mathbf{c}_j \mathbf{r}_j^\top$ reveals modular interaction components that can be individually in-
 1019 terpreted and manipulated, enabling mechanistic hypothesis testing through targeted weight pertur-
 1020 bations—capabilities absent in standard GNN architectures.

1021 H.4 EVALUATIONS

1022

- 1023 **Quantitative:** accuracy of the inferred “interaction matrices” compared to ground truth construc-
 1024 tion. The “interaction matrices” per method are defined as:

1026

1027 Table A4: Comparison of GNN and **LICCHIE** approaches for learning cell-cell interactions.

1028

Aspect	GNN	LICCHIE
Interactions	Implicit via message passing	Explicit pairwise matrices $A_t^{(n,k)}$
Interpretability	Post-hoc explanation needed	Direct feature-to-feature effects
Structure	Black-box neural network	Low-rank with interpretable components
Discovery	Distributed representations	Explicit modular components $\{c_j, r_j\}$
Intervention	Hard to target mechanisms	Modulate component weights $W_{j,t}^{(n,k)}$

1035

1. Linear: a single feature-to-feature matrix; A for $\mathbf{x}_t \approx A\mathbf{x}_{t-1}$.
2. Per-cell Linear: per-cell feature-to-feature matrices; $A^{(n)}$ for $\mathbf{x}_t^{(n)} \approx A^{(n)}\mathbf{x}_{t-1}^{(n)}$.
3. Per-cell SLDS: per-cell transition matrices; $A_t^{(n)}$ for $\mathbf{x}_t^{(n)} \approx A_{z|t-1}^{(n)}\mathbf{x}_{t-1}^{(n)}$.
4. GNN: approximated transition matrices, propagating the GNN loading matrices from the feature space of one time point to the next via the feature vector.
5. **LICCHIE**: the learned interaction matrices.

- **Qualitative**: visualization of the methods’ outputs and comparison to expected ground-truth dynamical operators.

1046

I METHODS PERFORMANCE COMPARISON ON REAL-WORLD APPLICATION

1049

In order to compare **LICCHIE** to other methods in a real-world scenario, we conducted a series of evaluations to test each model’s ability to predict biologically meaningful properties of the tumor-NK interaction system. Namely, we rely on known properties of cells in the data—the cell type (tumor, NK) and changes in their viability status (tumor) or cell polarization (NK), and test the ability of the models’ outputs to accurately predict them.

1054

First, we define a consistent set of seven properties describing the interactions identified by each method. We used these properties (hereafter ‘output features’) to predict the involved cell types and their polarization or viability status over time. These output features include: (i) total influence (sum of all elements), (ii) maximum column sum (strongest influence on any target feature), (iii) maximum row sum (strongest influence from any source feature), (iv) diagonal sum (self-feature interactions), (v) Frobenius norm (overall matrix magnitude), and (vi, vii) maximum/minimum elements (extreme interaction strengths). These features were used as input to a logistic regression classifier with balanced class weights to test each method’s ability to capture distinct cell and interaction properties.

1063

To address class imbalance in the dataset, we employed two strategies: (1) balanced class weights in logistic regression, and (2) sub-sampling approaches where we resampled both classes to equal sizes using bootstrap resampling with replacement.

1066

Source cell type Prediction For each interaction at time t , we extracted the above output features to predict the identity of the source cell (NK vs. tumor) via logistic regression with cross validation. Notably, while the baselines accuracy was close to random, GNN (0.54) and SLDS (0.57), **LICCHIE** presented higher predictive power (0.71) and was found to be less sensitive to data imbalance (Fig. S8a). Furthermore, looking at the confusion matrices, it is apparent that baseline methods show severe prediction collapse with degenerate concentration in single classes. **LICCHIE** on the other hand, produces balanced predictions for both cell types, maintaining strong performance even with balanced sampling (**LICCHIE**: 0.7, GNN: 0.5, SLDS: 0.5; Fig. S8b).

1074

1075

Target cell type prediction Similarly to the previous setting we now predict the identity of the target cell. **LICCHIE** again demonstrates better performance (0.68) compared to baseline methods (GNN: 0.54, SLDS: 0.57; Fig. S8c), maintaining meaningful predictions across both classes while baselines exhibit prediction bias with near-zero off-diagonal elements. With balanced sampling, **LICCHIE** achieves accuracy of 0.71 compared to baselines depicting random-like results (GNN: 0.51, SLDS: 0.5).

1080
 1081 **Interaction type classification** We extended the binary classification to predict the specific
 1082 source→target cell type combination, yielding four classes (NK→NK, NK→Tumor, Tumor→NK,
 1083 and Tumor→Tumor). This test aims to test the models’ ability to capture different interaction pat-
 1084 terns dictated by identities of interacting cells. While baselines remain at close to random per-
 1085 formance, LICCHIE maintains good performance (LICCHIE: 0.60, GNN: 0.26, SLDS: 0.32;
 1086 Fig. S8a).

1087 **Cellular state transition prediction** At last, we turn into a temporally resolved test—checking the
 1088 ability to predict cellular states at both current time t and future time $t + 1$. For each interaction
 1089 matrix $A_t^{(n,k)}$ extracted by each method, we predicted the status of target cell n both at time t and
 1090 at time $t + 1$. We conducted separate analyses for NK cells (predicting polarization status: non-
 1091 polarized vs. polarized) and tumor cells (predicting viability: alive vs. dead). Only samples with
 1092 non-NaN states were included in the analysis. LICCHIE shows better performance on all prediction
 1093 tasks (Fig. S8a,d), achieving above-chance accuracy for the challenging NK polarization prediction
 1094 (t : 0.62, $t + 1$: 0.65) while baselines remain at random-like performance (GNN t : 0.23, $t + 1$: 0.22;
 1095 SLDS t : 0.23, $t + 1$: 0.23).

1096 To summarize, across all evaluation tasks LICCHIE maintains consistently high accuracy (> 0.6)
 1097 while baselines tend to follow random-like accuracy (Fig. S8a). The lower performance of the base-
 1098 line methods suggests that their learned interactions cannot fully capture these existing biological
 1099 components in the real-world data (or at least that a more elaborate post-processing is required to
 1100 extract the information from their outputs). This is in line with the synthetic evaluations performed,
 1101 and can be explained by considering the nature of these baselines; SLDS, while interpretable, pro-
 1102 duces transitions with abrupt switches between them (Fig. S9) that cannot capture the smoothly
 1103 time- and distance-varying nature of cell-to-cell interactions; GNN produces loading matrices that
 1104 require extra processing to recover approximate transitions, and these approximations change over
 1105 time in ways that are hard to interpret because they lack an underlying interaction basis (Fig. S10).

1106
 1107
 1108
 1109
 1110
 1111
 1112
 1113
 1114
 1115
 1116
 1117
 1118
 1119
 1120
 1121
 1122
 1123
 1124
 1125
 1126
 1127
 1128
 1129
 1130
 1131
 1132
 1133

J 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. A comparison of LICCHIE to all baselines (linear dynamics, per-cell linear approximation, per-cell SLDS, and GNN) considering (left) correlation of identified interactions with the ground truth, and (right) SSIM score (Nilsson & Akenine-Möller, 2020)




Figure S3: *Per-cell SLDS baseline components on synthetic data.* **a**, Extracted weights $\mathbf{W}_t^{(k)}$ for the respective transition matrix, $\mathbf{A}^{(n,k)}$ with $k \in [1, 10]$ and for $n \in [1, 10]$ (a subset of the total 35 cells), presented over time, $t \in [0, 50]$. **b**, Example of the set of per-cell SLDS transition matrices, $\mathbf{A}^{(n,k)}$ with $k \in [1, 10]$, and for $n \in [1, 10]$ (a subset of the total 35 cells) identified by SLDS.




Figure S4: *GNN-based modeling of the synthetic data.* **a**, Feature space reconstruction evaluation. **a.1** Per-cell reconstruction comparison across different cells (left: ground truth, right: reconstructed). **a.2** Per-feature reconstruction for two example features, comparing ground truth (h) with predictions (\hat{h}). **b**, The loading (weight) matrices of the three layers in the trained GNN; **b.1**, **b.2**, **b.3** respectively. **c**, Approximated transition matrices for first 5 time points, reconstructed using the loading matrices in subplot **(b)** and the feature values at each time point.





Figure S5: *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.






Figure S7: *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 S8: *Performance comparison of LICCHIE against baseline methods (SLDS, GNN) on real-world data using machine learning evaluation tasks.* **a**, Classification accuracy across different prediction tasks. For each task, we trained logistic regression classifiers with balanced class weights on interaction matrix features, comparing performance with and without sub-sampling (addressing class imbalance). **b-d**, Normalized confusion matrices showing prediction performance for: **(b)** source cell type classification, **(c)** target cell type classification, and **(d)** NK cell polarization state prediction. Each method's interaction matrices were used to extract a set of features consisting of seven statistical features (sum, max column/row sums, diagonal sum, Frobenius norm, max/min elements) as input to the classifiers.




Figure S9: *Per-cell SLDS real-world data components.* **a**, Switched activation patterns of transition matrices for first 5 cells. **b**, Per-cell transition matrices for first 5 cells.




Figure S10: *Components identified by the GNN baseline on Real World Data.* **a**, Loading matrices from GNN trained on the real-world NK-tumor data. **b**, Example reconstruction for first two features, h ground truth, \hat{h} prediction results. **c**, Approximated transition matrices from loadings and latent states for the first 8 time points.