

Predicting Immunotherapy Response with Graph-Derived and Baseline Donor-Level Features in Pan-Cancer Single-Cell Tumor Data

CIS 5528: Predictive Modeling in Biomedicine

Grant Merrigan

grant.merrigan@temple.edu

PSM Bioinformatics

Tie Wen

tie.wen@temple.edu

Temple University

Khadija Wharton

khadija.wharton@temple.edu

Temple University

Calvin Kang

calvinkang@temple.edu

Temple University

Abstract

Predicting response to immune checkpoint inhibitors (ICI) remains a major challenge, partly due to intra-tumor heterogeneity among malignant cells. Most approaches rely on averaged gene expression and overlook the structural organization of tumor cell states. Here, we introduce a graph-based framework that quantifies malignant-cell heterogeneity from single-cell RNA sequencing data using graph diffusion entropy (GDE). For each patient, malignant cells are represented as a similarity graph in latent transcriptional space, and entropy over diffusion processes captures both the diversity and organization of tumor states. Applied to a large pan-cancer single-cell ICI dataset, raw GDE is strongly confounded by graph size; normalization reduces this effect but yields only modest, non-significant differences between favorable and unfavorable outcomes. In contrast, cluster-level analysis reveals greater dispersion of entropy across malignant subpopulations in unfavorable tumors, indicating more uneven tumor-state organization. A baseline XGBoost model using clinical, tumor, and microenvironmental features achieves moderate predictive performance, suggesting these features remain dominant predictors. Overall, donor-level GDE alone provides limited predictive signal, but graph-derived measures capturing variability across tumor subpopulations offer a complementary perspective on heterogeneity. These findings highlight both the promise and limitations of graph-based representations of single-cell data for predicting immunotherapy response.

1 Introduction

Immune checkpoint inhibitors (ICI) have improved cancer survival rates across multiple malignancies, with some patients achieving durable, long-term responses (Gondal, Cieslik, and Chinnaiyan 2025). Although approximately 57% of cancer patients in the United States are eligible for ICI therapies, only about 20% ultimately respond (Haslam, Olivier, and Prasad 2025). Even among responders, outcomes vary substantially, with some patients achieving durable remission, while others develop acquired resistance or relapse (Haslam, Olivier, and Prasad 2025; Heo et al. 2021). In addition, ICI therapies are often costly and can be associated with immune-related adverse events that resemble autoimmune disorders (Choi and Lee 2020). These challenges highlight the need for improved methods to predict patient response prior to treatment.

A key driver of this variability may be intra-tumor heterogeneity, defined by the presence of diverse molecular and transcriptional profiles within malignant cells of a single tumor (Pribluda, Cruz, and Jackson 2015; Ma et al. 2019). The tumor microenvironment consists of stromal, immune, and malignant cell populations; importantly, malignant cells are not uniform, but instead exist in multiple distinct transcriptional states (Tirosh et al. 2016). Although tumors are often categorized by histological subtype, individual malignant cells can differ substantially in gene expression programs, proliferative potential, metabolic activity, and immune-related signaling pathways (Andreatta, Garnica, and Carmona 2025; Guruprasad et al. 2021). These differences can influence how tumor cells interact with the immune system and respond to therapeutic pressure. As a result, tumors composed of more diverse or adaptable malignant states may be more likely to evade immune surveillance and develop resistance to therapy (Gondal, Cieslik, and Chinnaiyan 2025; Marusyk, Almendro, and Polyak 2012; S. K. Kim and Cho 2022). Quantifying this heterogeneity in a patient-specific and structurally informative manner is therefore essential for understanding treatment response and improving outcome prediction.

Most current approaches summarize tumors using average gene expression levels, which may fail to capture how malignant cell states are structured and interconnected within a tumor (Dagogo-Jack and Shaw 2018). Advances in single-cell RNA sequencing, however, have enabled direct measurement of cell-to-cell transcriptional variation and have become essential for studying tumor heterogeneity. These developments highlight the need for analytical methods that represent relationships between cells, rather than methods that summarize gene expression at the aggregate level (Huang et al. 2023). Single-cell data can, therefore, be represented as graphs, where individual cells are modeled as nodes connected by transcriptional similarity (Hetzel, Günemann, and Theis 2021). Within this framework, the structure of malignant cell states can be characterized by modeling how signals diffuse across the graph, enabling the derivation of patient-level measures of structural heterogeneity within tumor cell populations.

We implement this approach by constructing patient-specific graphs of malignant cells based on their gene expression profiles and apply it to a large, publicly available dataset of single-cell tumor profiles collected before and after treatment (Gondal, Cieslik, and Chinnaiyan 2025). We compute a graph diffusion entropy (GDE) metric for each patient, which summarizes both the diversity and organization of malignant cell states by measuring how signals spread across the tumor-state graph. Importantly, this GDE score serves as an engineered feature representing tumor heterogeneity at the patient level. We first assess whether GDE is associated with clinical outcomes using statistical testing and then evaluate its contribution to predictive performance when incorporated into XGBoost, alongside established clinical and tumor features. The primary contribution of this work lies in the construction of a graph-derived feature, enabling evaluation of whether tumor structure provides additional predictive signal beyond conventional features.

2 Hypothesis

Pre-treatment tumors exhibiting higher malignant-cell heterogeneity, as quantified by a graph-based diffusion entropy measure, are associated with unfavorable clinical outcomes. Specifically, graph diffusion entropy computed from patient-specific malignant cell expression profiles will show a statistically significant association with patient outcomes and will provide greater discriminative power than a baseline model trained on standard gene-expression features alone.

3 Methods

This study analyzes a large, publicly available single-cell RNA-sequencing dataset of tumor samples from pan-cancer patients treated with immune checkpoint blockade (ICB), obtained from Zenodo (record 14511579). The dataset integrates scRNA-seq profiles across multiple cancer types and clinical studies, with harmonized annotations for treatment timing, cell type, malignant status, and patient-level clinical outcomes following immunotherapy.

Combined outcome category	Number of patients
Unfavourable	81
Favourable	41
UT	23
n/a	22

Table 1: Clinical outcome label counts in the dataset. “UT” denotes another patient outcome present in the dataset; “n/a” indicates missing or unassigned outcome. Patients with an outcome of “UT” or “n/a” were excluded, resulting in a final analytical cohort of 122 patients.

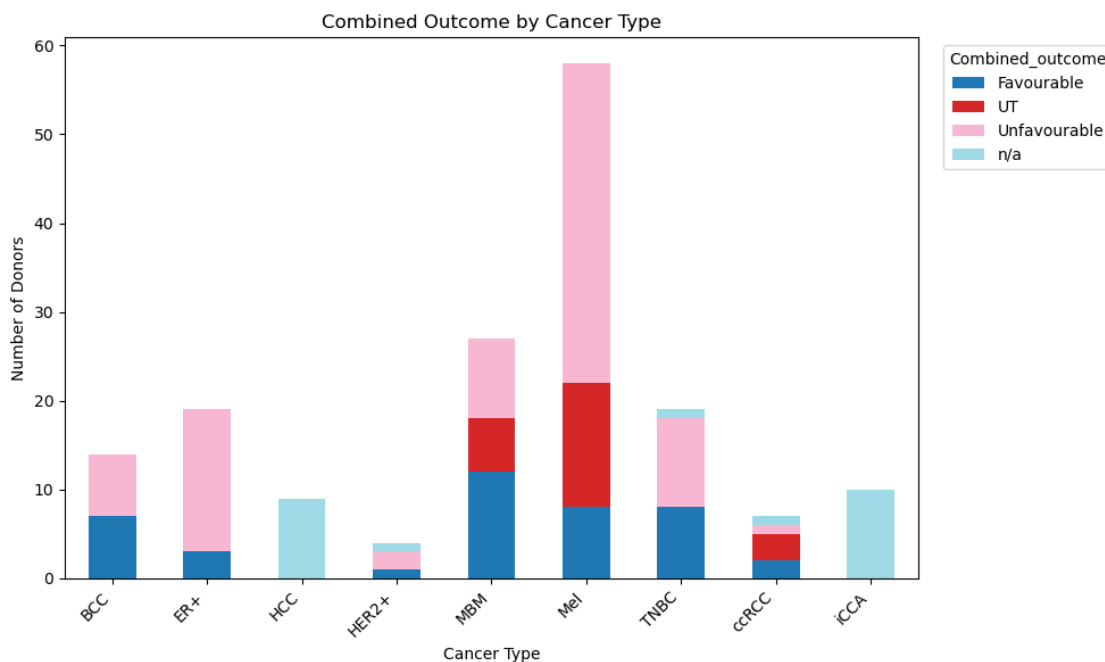


Figure 1: Stacked bar chart showing combined patient outcomes for each cancer type in the integrated single-cell ICB dataset.

3.1 Malignant Cell Identification via Copy Number Variation

To restrict downstream analyses to tumor-intrinsic transcriptional heterogeneity, malignant cells were identified using inferred copy number variation (CNV) profiles derived from single-cell RNA sequencing data. Large-scale chromosomal gains and losses are a hallmark of cancer and are frequently detectable at the transcriptomic level due to coordinated shifts in gene expression across contiguous genomic regions.

CNV inference was performed using the Python library `infercnvpy`, which operates directly on `AnnData` objects within the Scanpy ecosystem. The workflow proceeded as follows:

1. **Gene–chromosome mapping:** A gene transfer format (GTF) file corresponding to the appropriate genome build (e.g., GRCh38) was parsed to extract chromosomal coordinates for each gene. These coordinates allowed genes in the expression matrix to be ordered along chromosomes, enabling detection of regional expression shifts rather than isolated gene-level differences.
2. **Reference cell specification:** Putative non-malignant cell populations, such as immune or stromal cells identified through canonical marker expression and clustering, were designated as reference cells. These reference populations provided a baseline diploid expression profile against which epithelial or tumor-enriched clusters were compared.
3. **CNV signal computation:** Using `infercnvpy.tl.infercnv`, gene expression values were smoothed across neighboring genomic windows along each chromosome. For each cell, deviations from the reference baseline were computed, producing a per-cell CNV profile that reflected inferred chromosomal amplifications or deletions.
4. **Malignancy classification:** Cells exhibiting large-scale, coherent chromosomal shifts, rather than localized stochastic variation, were classified as malignant. This classification was further validated by examining consistency within epithelial clusters and by visualizing CNV heatmaps across chromosomes.

Computationally, this step was implemented in Python using Scanpy and `infercnvpy` within a reproducible workflow. After standard preprocessing, including normalization, log-transformation, and highly variable gene selection, CNV inference was performed separately for each patient to avoid cross-sample artifacts. The resulting CNV scores or classifications were stored in the `AnnData.obs` metadata and used to subset malignant cells for all subsequent graph construction and heterogeneity analyses.

This approach ensured that the downstream graph diffusion entropy metric quantified heterogeneity intrinsic to tumor cells rather than variation introduced by immune infiltration or stromal composition.

3.2 Preprocessing and Latent Representation

Gene expression values were normalized and log-transformed using standard single-cell preprocessing workflows. Highly variable genes were selected within the malignant compartment to reduce technical noise and emphasize biologically meaningful variation. Principal component analysis (PCA) was then applied to the tumor-restricted expression matrix, and the leading principal components were retained as a low-dimensional representation of malignant transcriptional state.

To contextualize the diversity of the cohort, Figure 2 summarizes the distribution of patients (donors) and cancer types included in the integrated dataset.

This latent space provides a compact description of tumor cell variation while preserving the major axes of transcriptional diversity needed for graph-based analysis.

3.3 Patient-Specific Malignant-Cell Similarity Graph Construction

For each patient, malignant cells were represented as nodes in a transcriptional-state graph constructed in PCA space. A k -nearest-neighbor (kNN) graph was built using Euclidean distance,

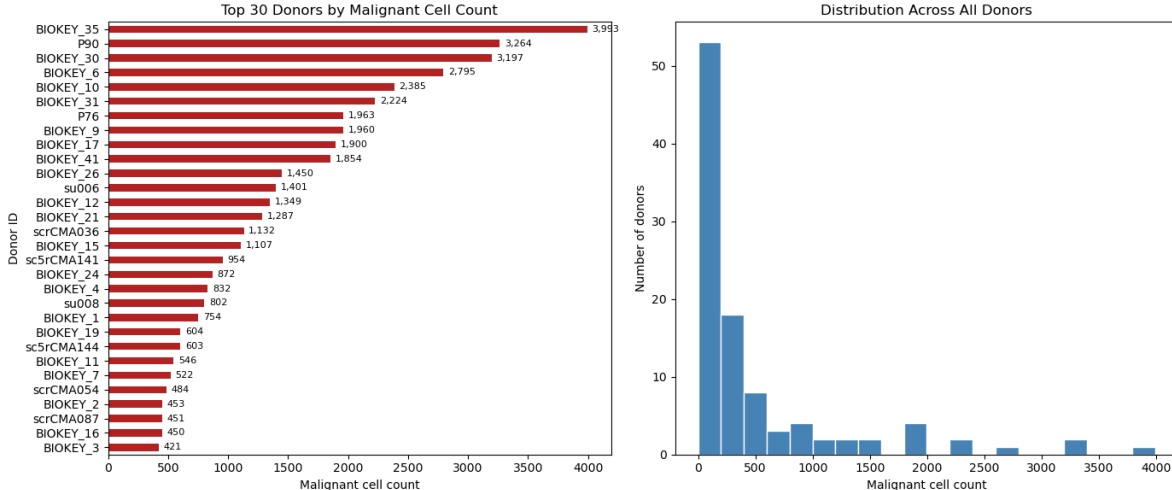


Figure 2: **Left:** Top 30 donors ranked by malignant cell count. **Right:** Distribution of malignant cell counts across all donors. The wide, right-skewed variation in malignant cell numbers highlights substantial inter-patient variation in graph size, motivating patient-specific graph construction and evaluation of diffusion-based heterogeneity metrics.

with k chosen to balance sensitivity to local structure and robustness to noise. Edge weights were assigned using a Gaussian kernel, such that transcriptionally similar cells were more strongly connected.

This graph representation treated each tumor as a structured system of related cell states, rather than a collection of independent observations. The resulting topology reflected how densely or sparsely malignant states were organized within a patient’s tumor.

3.4 Graph Convolutional Refinement of Graphs

After constructing each graph using a k -nearest-neighbor (kNN) strategy in latent space, the initial graph reflected local similarity between nodes based primarily on Euclidean distance. Although this structure captured immediate neighborhood relationships, it did not explicitly model higher-order dependencies or the broader topological organization of the graph.

To incorporate contextual information from neighboring nodes and higher-order graph-topological structure, we introduced a Graph Convolutional Network (GCN) refinement module to refine node embeddings prior to diffusion-based entropy computation in both cell-level and cluster-level graph representations.

Let $W \in \mathbb{R}^{n \times n}$ denote the weighted adjacency matrix of the graph and $X \in \mathbb{R}^{n \times d}$ represent the initial node feature matrix in the latent space, where n is the number of nodes and d is the dimensionality of the latent representation. For cell-level graphs, each node corresponded to an individual malignant cell; for cluster-level graphs, each node corresponded to a Leiden cluster represented by its aggregated latent features. The GCN performed message passing across connected nodes and produced updated node representations according to:

$$X' = \sigma \left(D^{-1/2} W D^{-1/2} X \Theta \right),$$

where D is the degree matrix defined as $D_{ii} = \sum_j W_{ij}$, Θ denotes learnable parameters, and $\sigma(\cdot)$ is a nonlinear activation function.

This symmetric normalization ensured stable propagation of information across the graph and prevented dominance by highly connected nodes.

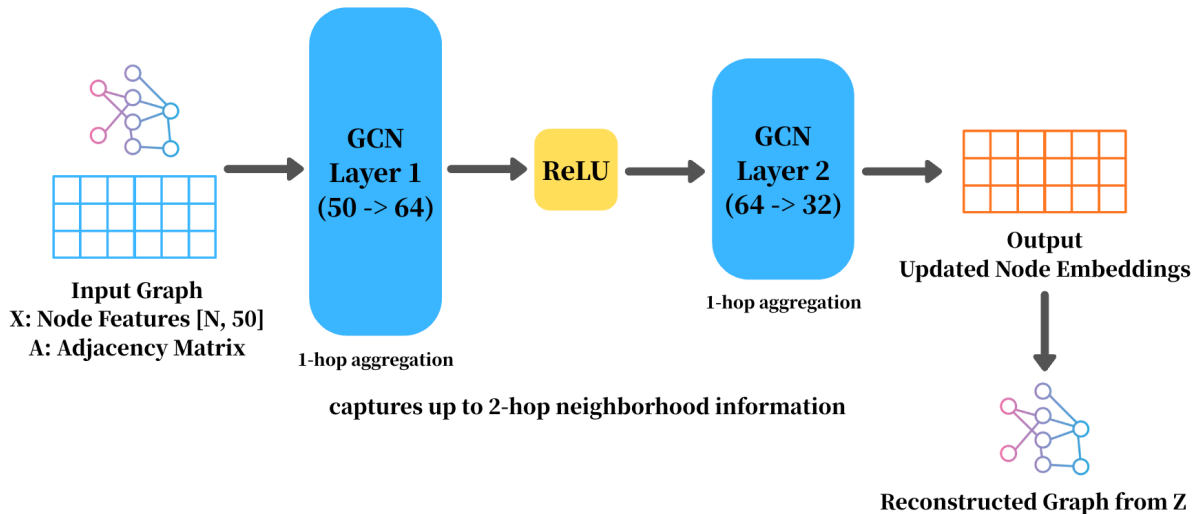


Figure 3: Architecture of the two-layer GCN for fixed-graph representation learning.

In this study, the GCN was used as a self-supervised representation learning module. A two-layer GCN architecture was adopted so that node embeddings could incorporate information from an approximately 2-hop neighborhood through stacked message passing, while maintaining limited depth to reduce the risk of overfitting and oversmoothing given the scale of the patient-specific graphs. Training was performed using a self-supervised reconstruction-based objective. In each iteration, a random subset of node features was masked or perturbed, and the GCN was trained to reconstruct the masked information from the remaining graph context via neighborhood aggregation. By repeating this masking-and-reconstruction procedure with different random masks across iterations, all nodes were progressively exposed to the reconstruction objective, enabling the learnable parameters Θ to capture a shared neighborhood-aware refinement rule for updating node embeddings across patient-specific graphs.

The refinement step allowed each node to integrate information from its local neighborhood, effectively smoothing similarity patterns over the graph while preserving structural heterogeneity. The updated embeddings X' provided a more robust representation of the global organization of the graph and were subsequently used to reconstruct the weighted graph for downstream diffusion modeling.

Importantly, this representation learning step did not replace the diffusion process; rather, it enhanced the structural fidelity of the graph before entropy-based quantification.

3.5 Graph-Derived Donor-Level Features

Unless otherwise stated, graph-derived features were computed from the GCN-refined malignant-cell graphs reconstructed from the updated node embeddings. To characterize malignant-cell organization beyond baseline clinical, tumor, and tumor-microenvironment features, we extracted donor-level graph features from each patient-specific graph. These features were grouped into two broad categories: GDE-related diffusion features and other graph-derived structural features. GDE-related features described how random-walk signals spread across the malignant-cell graph, while other graph-derived features summarized the static geometry, dominant structural patterns,

and cluster organization of malignant-cell states.

3.5.1 GDE-Related Diffusion Features

Tumor heterogeneity was quantified using graph diffusion entropy (GDE) computed at the patient level. For each patient, malignant cells were represented as nodes in a tumor-state graph, where edges connected transcriptionally similar cells. These relationships were summarized in a weighted adjacency matrix W .

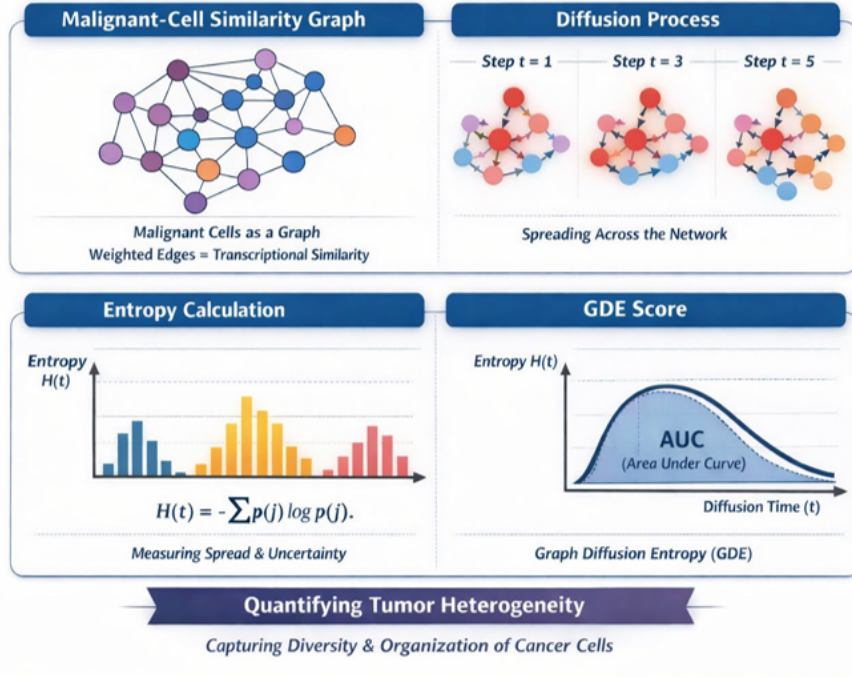


Figure 4: Schematic overview of the graph diffusion entropy (GDE) calculation used to summarize malignant-cell heterogeneity at the patient level.

To model how transcriptional state can spread across this network, the adjacency matrix was normalized to form a transition matrix,

$$P = D^{-1}W, \quad (1)$$

where D is a diagonal matrix containing the total connection strength of each cell. This normalization ensures that each row of P represents the relative likelihood of moving from one cell to its neighbors.

Diffusion was simulated by repeatedly applying this transition matrix. Starting from each malignant cell, the resulting distribution after t steps was given by

$$p_i^{(t)} = e_i P^t, \quad (2)$$

where e_i is a one-hot row vector, and i is the index of the starting cell. This describes how widely the influence of a given cell spreads across the tumor-state space.

At each diffusion step, entropy was used to summarize how spread out this distribution was:

$$H(t) = - \sum_j p^{(t)}(j) \log p^{(t)}(j). \quad (3)$$

Here, $p^{(t)}(j)$ denotes the probability of the diffusion process occupying node j after t steps, and $H(t)$ quantifies the degree of uncertainty or dispersion of this distribution across the graph. Higher entropy values indicate broader exploration of transcriptionally distinct malignant-cell states.

Rather than relying on a single diffusion time, entropy was computed across a range of diffusion steps $t \in \{t_1, \dots, t_T\}$ to capture heterogeneity at multiple scales. Entropy values were first averaged across starting cells at each diffusion step, producing an entropy trajectory $H(t)$ for each patient. These trajectories were then summarized using the area under the entropy–time curve (AUC), yielding a single patient-level GDE score that reflects both the magnitude and persistence of diffusion-driven dispersion across the tumor-state graph.

In addition to cell-level GDE, we also computed Leiden-based cluster-graph diffusion features. In this analysis, Leiden clustering was applied to the malignant-cell graph to identify graph communities corresponding to malignant-cell subpopulations. Each Leiden community was then represented as a cluster-level node, and diffusion entropy was computed on the resulting cluster graph. This allowed us to evaluate whether diffusion behavior differed not only across individual malignant cells, but also across broader malignant-cell subpopulations.

Conceptually, GDE captures not only how many different malignant-cell states are present, but also how those states are organized and connected. In the context of immune checkpoint blockade, higher GDE values may reflect tumors with greater intrinsic transcriptional diversity and plasticity, potentially influencing immune recognition, immune evasion, or adaptive resistance mechanisms prior to treatment.

3.5.2 Other Graph-Derived Structural Features

In addition to GDE-related diffusion measures, we extracted non-diffusion graph-derived features to summarize the static organization of malignant-cell states. These features were divided into static graph features and k-means-based cluster-structure features.

Static graph features. Static graph features describe the overall geometry and structural dominance of the malignant-cell graph without relying on random-walk diffusion. These features include summaries of cell-state dispersion, such as centroid-distance and pairwise-distance statistics, as well as graph spectral features such as the leading eigenvalue ratio. Together, these features capture whether malignant-cell states are compact, broadly dispersed, or dominated by a strong major structural pattern.

K-means cluster-structure features. To further summarize subpopulation-level organization, we extracted k-means-based cluster-structure features from the malignant-cell latent space. These features describe the number of clusters, cluster-size balance, dominance of the largest cluster, fraction of small clusters, and separation between clusters. Unlike the Leiden-based cluster graph used for diffusion analysis, these k-means features were used as donor-level structural summaries rather than for GDE computation.

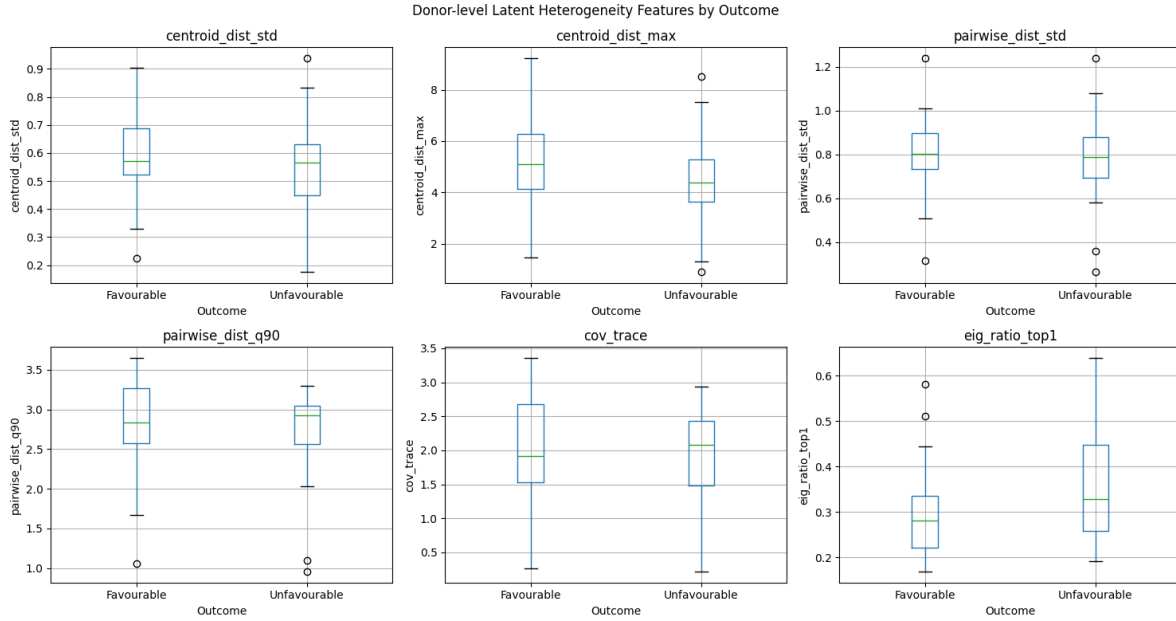


Figure 5: Static graph features used to summarize the overall geometry and structural organization of malignant-cell states.

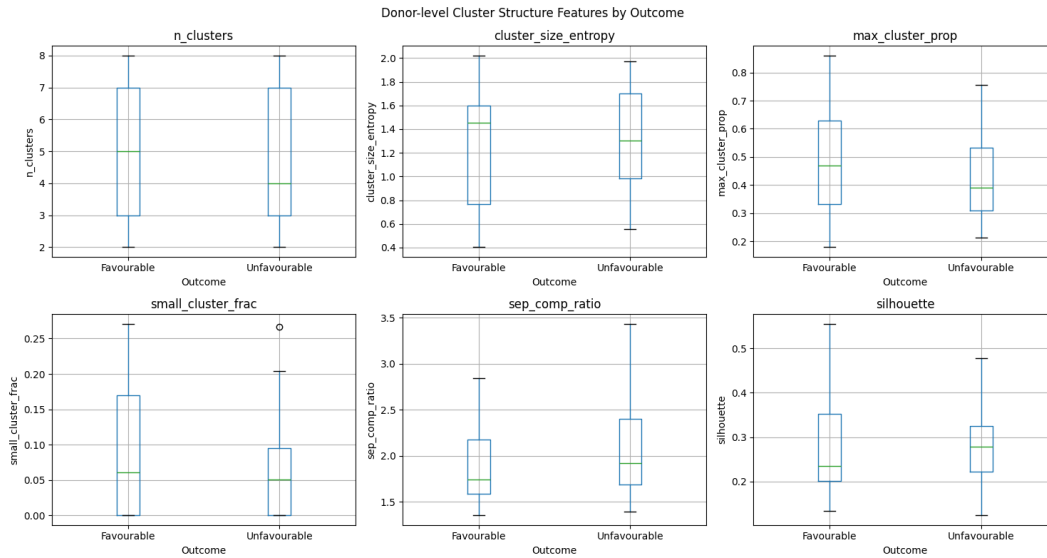


Figure 6: K-means-based cluster-structure features used to summarize malignant-cell subpopulation organization.

Together, these features capture both fine-grained and coarse-grained patterns of malignant-cell organization. Following topological and biological interpretation, three graph-derived features with strong representational relevance and group-level discriminative potential were prioritized for detailed analysis based on their biological interpretability and observed group-level separation. The top-1 eigenvalue characterizes the extent to which the malignant-cell graph is dominated by a single major structural pattern, with higher values in the unfavourable group suggesting a

more concentrated tumor-state organization. Normalized random-walk entropy AUC captures the degree to which malignant-cell states can mix and spread across the graph through diffusion dynamics. The separation-compactness ratio measures the distinctness of cell-state communities by comparing between-cluster separation with within-cluster compactness, where higher values in the unfavourable group indicate stronger boundaries between malignant-cell subpopulations.

Overall, unfavourable tumors appear to exhibit a more dominant, less mixed, and more clearly separated malignant-cell state organization. These graph-derived features may provide complementary structural signals for downstream predictive tasks, and the observed group-level differences suggest that they could improve outcome modeling and classification when integrated with baseline clinical and tumor-microenvironment features.

3.6 Baseline Model Evaluation

To assess the additional predictive signal provided by the graph-based approach, a baseline model was first trained using non-graph covariate features. These covariates were grouped into the following categories:

- **Clinical:** age, sex, stage, and line of therapy/method of action.
- **Tumor:** malignant fraction as a purity proxy, TMB, MSI, and PD-L1 status when available.
- **TME:** fine-grained cell-type composition, including CD8 T cells, regulatory T cells, myeloid, stromal, and endothelial cell populations, as well as key immune-state proportions.
- **Technical/study:** study ID, tissue site, sampling timepoint, sequencing platform or chemistry, and per-sample QC/depth metrics.

3.7 Association with Clinical Outcomes

The association between graph diffusion entropy and clinical outcome was assessed by splitting the cohort into favorable and unfavorable groups. A two-sided Mann–Whitney U test with a significance threshold of $p < 0.05$ was used to determine whether GDE-related scores differed between groups.

The predictive value of graph-derived features was then evaluated by incorporating them as additional training features in predictive models. Classical machine learning models, including Random Forest, XGBoost, and Logistic Regression, were considered to assess whether graph-based heterogeneity provided predictive signal beyond conventional non-graph-based representations. Model performance was assessed using cross-validation to limit overfitting and information leakage.

4 Results

4.1 Baseline Feature Model

The baseline model was trained on 176 donor/timepoint samples (116 unfavorable, 60 favorable). On the held-out test set, the XGBoost baseline achieved an accuracy of 0.694, macro-F1 of 0.630, ROC-AUC of 0.674, and average precision of 0.613. Performance was stronger for the unfavorable class (precision 0.74, recall 0.83) than for the favorable class (precision 0.56, recall 0.42), indicating that the model captured some outcome signal but remained biased toward the majority class.

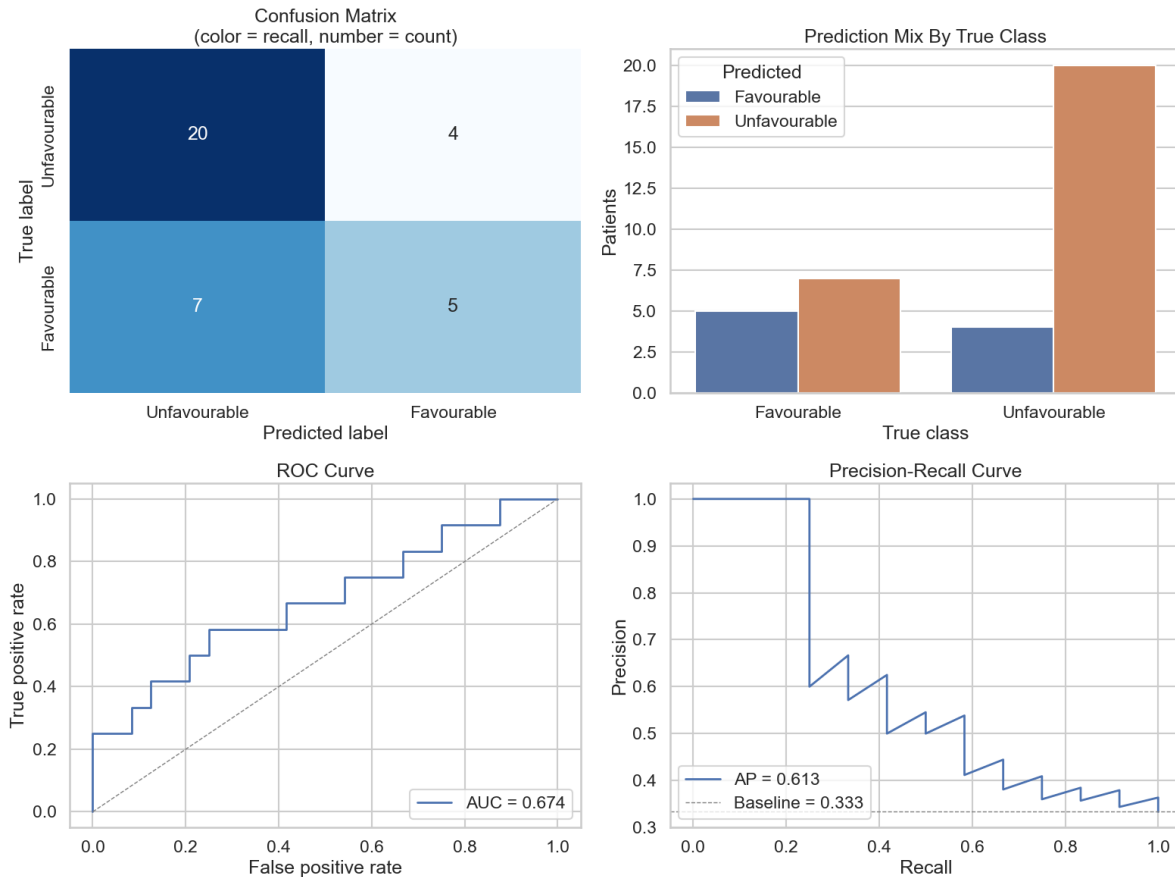


Figure 7: Held-out performance of the baseline XGBoost model trained on donor-level engineered features. The figure shows the confusion matrix, class prediction mix, ROC curve, and precision-recall curve from the BaselineModel notebook.

Feature importance was dominated by a mixture of fine-grained cell-type composition variables, immune-state signatures, and study-specific covariates. Highly ranked features included plasma-cell and naive-T-cell fractions, all-cell exhaustion signal, malignant-cell cytotoxic signal, pDC abundance, low-end RNA QC summaries, and treatment/study indicators. These rankings suggest that the baseline model relies heavily on tumor microenvironment composition and cohort-specific context rather than a single dominant biological marker.

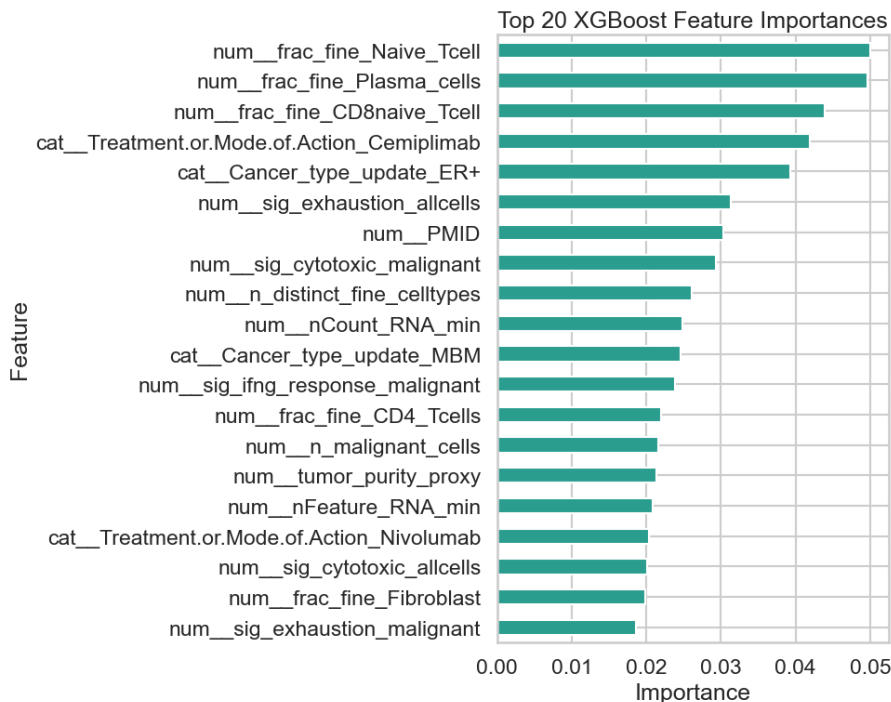


Figure 8: Top 20 XGBoost feature importances from the baseline model. The most influential features are a mixture of fine-grained cell-type composition, immune-state signatures, and study/treatment covariates.

4.2 Pre-treatment Graph Diffusion Entropy

For the graph-based analysis, restricting to pre-treatment samples with at least 50 malignant cells and valid favorable/unfavorable labels produced 50 donors (33 unfavorable, 17 favorable). Raw diffusion entropy AUC showed only a modest difference between outcomes (mean 97.14 in favorable versus 94.37 in unfavorable tumors), but it was strongly correlated with graph size (Pearson $r = 0.881$, Spearman $\rho = 0.943$), indicating that donor-to-donor variation in malignant-cell count substantially inflates the unnormalized entropy score.

To address this, we normalized donor-level GDE by $\log(n_{\text{malignant cells}})$. After normalization, favorable tumors had slightly higher scores (mean 16.06, median 15.75) than unfavorable tumors (mean 15.54, median 15.56), but the difference did not reach significance (Mann–Whitney $U = 363.0$, $p = 0.093$, rank-biserial correlation = 0.294). Thus, the primary donor-level GDE metric showed a directional trend but not a statistically significant separation between outcome groups in this cohort.

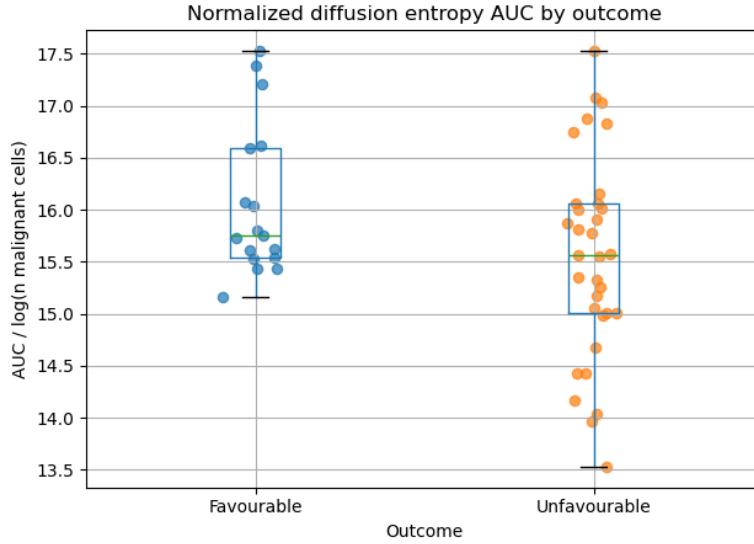


Figure 9: Normalized donor-level graph diffusion entropy by clinical outcome, computed as entropy AUC divided by $\log(n_{\text{malignant cells}})$. Favorable tumors show a mild upward shift, but the group difference is not statistically significant.

The shape of the normalized entropy trajectory was also outcome-dependent. Favorable tumors maintained slightly higher normalized entropy across most walk steps, suggesting broader exploration of malignant-state space during diffusion, although the separation remained small in magnitude.




Figure 10: Outcome-stratified normalized diffusion entropy curves averaged across donors. The favorable group remains modestly above the unfavorable group over most walk steps.

4.3 Leiden Cluster-Graph Analysis

All 50 donors also yielded valid Leiden cluster graphs. The normalized cluster-graph GDE score again showed only a small mean difference between groups (16.79 for favorable versus 16.43 for unfavorable tumors), implying that simple donor-level entropy on the reduced cluster graph does not strongly separate clinical outcomes on its own.

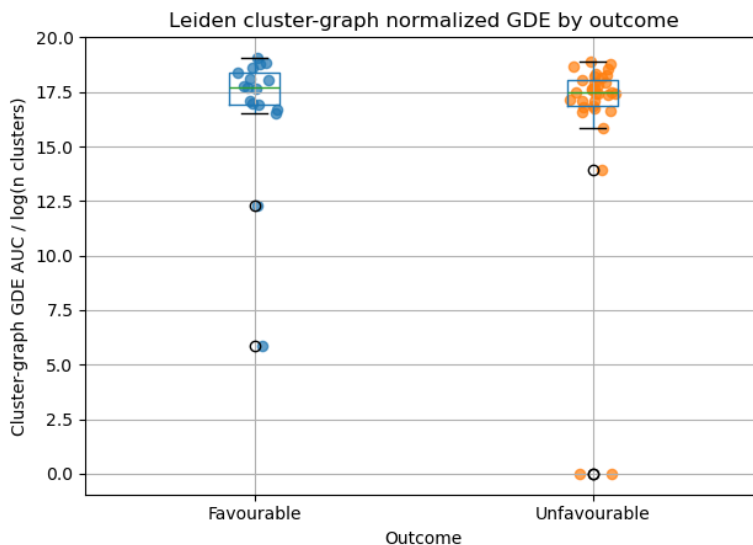


Figure 11: Normalized Leiden cluster-graph diffusion entropy by outcome. Aggregating malignant cells into transcriptional clusters preserves the overall trend but does not materially increase group separation.

However, cluster-level entropy dispersion was more informative than the donor-level mean. After excluding one extreme favorable outlier (BIOKEY_28), the variance of cluster-level entropy AUC values was higher in unfavorable tumors (mean 0.206) than in favorable tumors (mean 0.108). Similarly, the unfavorable group had larger mean IQR and MAD of cluster entropy values, indicating that unfavorable tumors tend to contain a wider spread of low-entropy and high-entropy malignant clusters. In contrast, the fraction of high-entropy clusters alone was nearly identical between groups (0.147 in unfavorable versus 0.134 in favorable tumors), suggesting that heterogeneity is better captured by the spread of cluster entropies than by a simple threshold count.




Figure 12: Variance of cluster-level entropy AUC values by outcome after removing one extreme favorable outlier. Unfavorable tumors show greater within-donor dispersion of cluster entropies, suggesting more uneven malignant-state organization.

Representative Leiden cluster graphs in UMAP space illustrate this structural variability. Favorable and unfavorable tumors can have similar numbers of clusters while differing in how evenly those clusters are connected and distributed, motivating the use of diffusion-based summaries rather than simple graph size alone.

Three favourable and three unfavourable Leiden cluster graphs in UMAP space

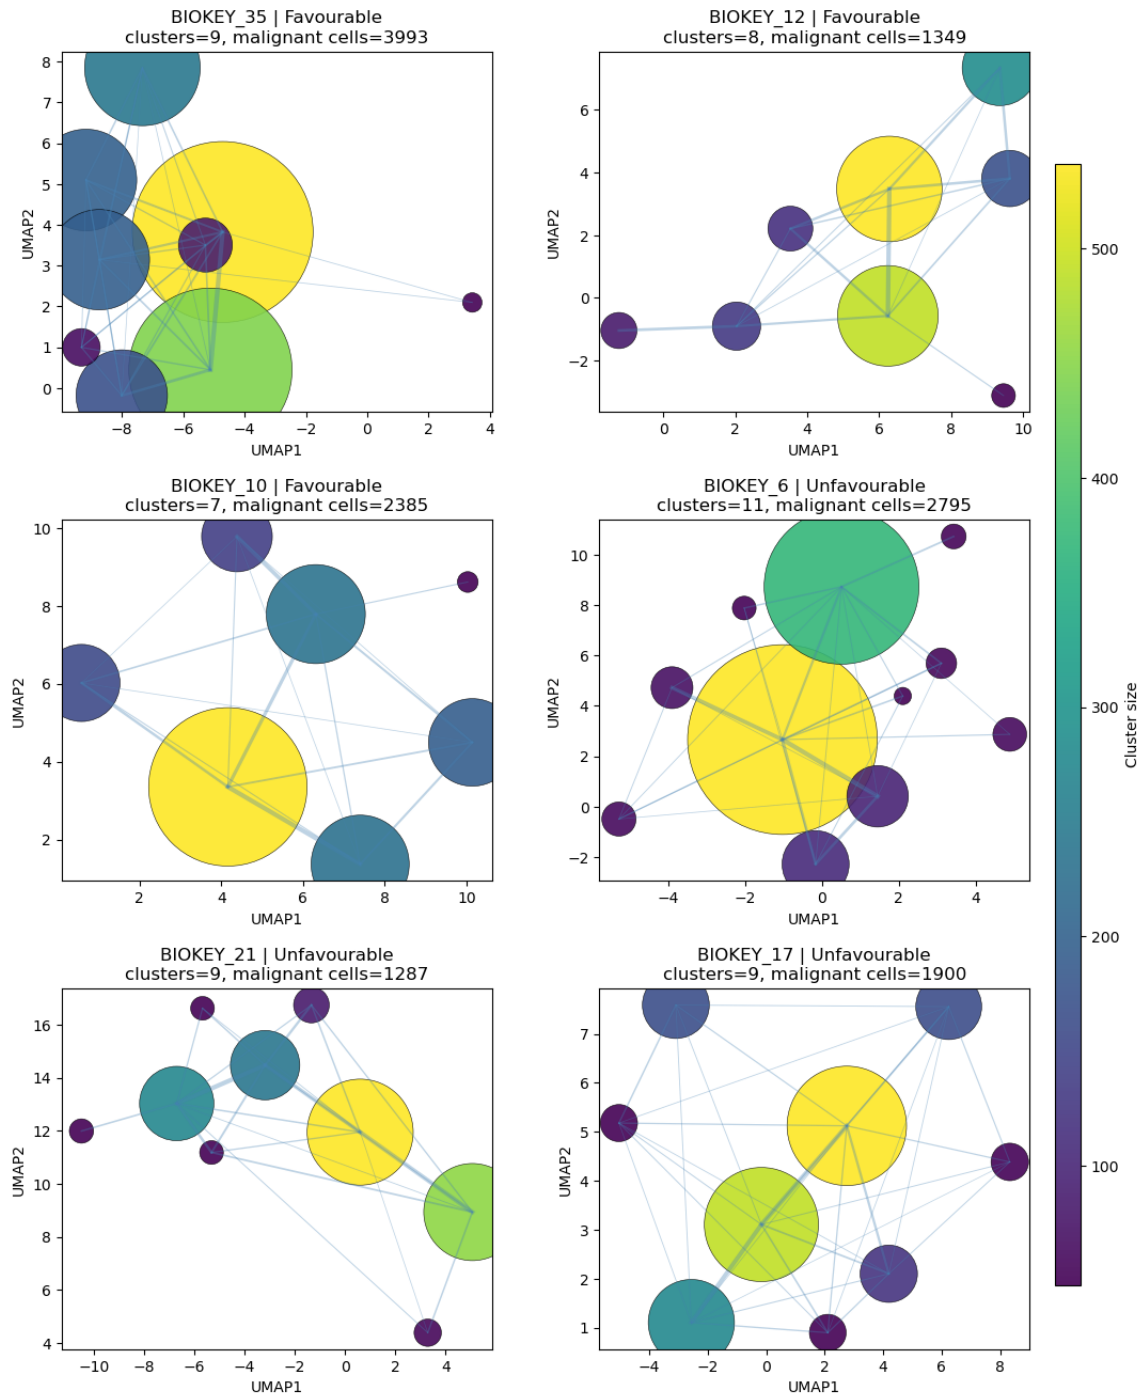


Figure 13: Examples of favorable and unfavorable malignant-cell Leiden cluster graphs in UMAP space. Nodes correspond to cluster centroids and edge width reflects normalized between-cluster connectivity.

Taken together, these results show that a standard donor-level baseline model has moderate predictive utility, while graph-derived donor-level entropy alone provides only weak outcome sep-

aration once graph size is controlled. The most promising signal from the graph analysis appears at the level of *within-donor cluster entropy dispersion*, where unfavorable tumors exhibit a broader spread of malignant cluster behaviors even when the mean donor-level entropy is similar.

References

- Andreatta, Massimo, Josep Garnica, and Santiago J. Carmona (2025). “Identification of malignant cells in single-cell transcriptomics data”. In: *Communications Biology* 8.1, p. 1264. DOI: 10.1038/s42003-025-08695-4. URL: <https://doi.org/10.1038/s42003-025-08695-4>.
- Choi, J. and S. Y. Lee (2020). “Clinical Characteristics and Treatment of Immune-Related Adverse Events of Immune Checkpoint Inhibitors”. In: *Immune Network* 20.1, e9. DOI: 10.4110/in.2020.20.e9. URL: <https://doi.org/10.4110/in.2020.20.e9>.
- Dagogo-Jack, I. and A. T. Shaw (Feb. 2018). “Tumour heterogeneity and resistance to cancer therapies”. In: *Nature Reviews Clinical Oncology* 15.2. Epub 2017-11-08, pp. 81–94. DOI: 10.1038/nrclinonc.2017.166.
- Gondal, M. N., M. Cieslik, and A. M. Chinnaiyan (2025). “Integrated cancer cell-specific single-cell RNA-seq datasets of immune checkpoint blockade-treated patients”. In: *Scientific Data* 12.1, p. 139. DOI: 10.1038/s41597-025-04381-6. URL: <https://doi.org/10.1038/s41597-025-04381-6>.
- Guruprasad, Puneeth et al. (2021). “The current landscape of single-cell transcriptomics for cancer immunotherapy”. In: *Journal of Experimental Medicine* 218.1, e20201574. DOI: 10.1084/jem.20201574. URL: <https://doi.org/10.1084/jem.20201574>.
- Haslam, Alyson, Timothée Olivier, and Vinay Prasad (2025). “How many people in the US are eligible for and respond to checkpoint inhibitors: An empirical analysis”. In: *International Journal of Cancer* 156.12, pp. 2352–2359. DOI: 10.1002/ijc.35347. URL: <https://doi.org/10.1002/ijc.35347>.
- Heo, Ja Yoon et al. (2021). “Clinical pattern of failure after a durable response to immune checkpoint inhibitors in non-small cell lung cancer patients”. In: *Scientific Reports* 11.1, p. 2514. DOI: 10.1038/s41598-021-81666-x. URL: <https://doi.org/10.1038/s41598-021-81666-x>.
- Hetzl, Leon, David S. Fischer and Stephan Günemann, and Fabian J. Theis (2021). “Graph representation learning for single-cell biology”. In: *Current Opinion in Systems Biology* 28, p. 100347.
- Huang, D. et al. (2023). “Advances in single-cell RNA sequencing and its applications in cancer research”. In: *Journal of Hematology and Oncology* 16.11, p. 98. DOI: 10.1186/s13045-023-01494-6. URL: <https://doi.org/10.1186/s13045-023-01494-6>.
- Kim, S. K. and S. W. Cho (2022). “The Evasion Mechanisms of Cancer Immunity and Drug Intervention in the Tumor Microenvironment”. In: *Frontiers in Pharmacology* 13, p. 868695. DOI: 10.3389/fphar.2022.868695. URL: <https://doi.org/10.3389/fphar.2022.868695>.
- Ma, Lichun et al. (2019). “Tumor Cell Biodiversity Drives Microenvironmental Reprogramming in Liver Cancer”. In: *Cancer Cell* 36.4, 418–430.e6. DOI: 10.1016/j.ccell.2019.08.007. URL: <https://doi.org/10.1016/j.ccell.2019.08.007>.
- Marusyk, Andriy, Vanessa Almendro, and Kornelia Polyak (2012). “Intra-tumour heterogeneity: a looking glass for cancer?” In: *Nature Reviews Cancer* 12.5, pp. 323–334. DOI: 10.1038/nrc3261. URL: <https://doi.org/10.1038/nrc3261>.
- Pribluda, Ariel, Cecile C. de la Cruz, and Erica L. Jackson (2015). “Intratumoral Heterogeneity: From Diversity Comes Resistance”. In: *Clinical Cancer Research* 21.13, pp. 2916–2923. DOI: 10.1158/1078-0432.CCR-14-1213. URL: <https://doi.org/10.1158/1078-0432.CCR-14-1213>.

Tirosh, Itay et al. (2016). “Dissecting the multicellular ecosystem of metastatic melanoma by single-cell RNA-seq”. In: *Science* 352.6282, pp. 189–196. DOI: 10.1126/science.aad0501. URL: <https://doi.org/10.1126/science.aad0501>.