LIMEADE: LOCAL INTERPRETABLE MANIFOLD EXPLANATIONS FOR DIMENSION EVALUATIONS

Tarek M. Zikry & Genevera I. Allen
Department of Statistics
Irving Center for Cancer Dynamics
Center for Theoretical Neuroscience,
Zuckerman Mind Brain Behavior Institute
Columbia University
New York, NY 10027, USA
{tarek.zikry,genevera.allen}@columbia.edu

ABSTRACT

To visualize and analyze high-dimensional biological data, scientists often turn to manifold learning and dimensionality reduction techniques such as tSNE and UMAP. However, these methods (*1*) are non-projective, which means that new data cannot be projected on the manifold without refitting, and (*2*) lack the explainability to help practitioners understand which features drive manifold locations and neighborhoods. In practice, scientists often must turn to marginal distributions along a manifold or expert annotations to explain reduced dimension data. Here, we present Local Interpretable Manifold Explanations for Dimension Evaluations (LIMEADE), a surrogate model integrated with a dimensionality reduction method, similar to the LIME surrogate used in classification and regression models. We define LIMEADE as a group lasso-regularized multi-task regression problem that identifies sparse linear projections of the data aligning with local neighborhoods of the manifold space. When applied to single-cell proteomics data, LIMEADE effectively extracts biologically meaningful features, providing a more interpretable approach to feature selection and dimensionality reduction.

1 INTRODUCTION

Dimension reduction (DR), also known as manifold learning, has become commonplace for visualizing complex data. These unsupervised learning methods take a high dimensional dataset of size $n \times p$, and via a function f, map the data to a new set of components: $f : \mathbb{R}^{n \times p} \to \mathbb{R}^{n \times r}$, where $r \ll p$, with r typically set to 2 or 3 for ease of visualization. Each specific DR method defines a different function f, designed to capture meaningful information present in the data. Local methods seek to preserve the relationships between similar observations in the high-dimensional space when mapping to a new lower-dimensional space, whereas global methods focus on preserving structural patterns across the entire dataset (Van Der Maaten et al., 2009).

With increasingly complex biological datasets being collected, nonlinear DR methods such as t-SNE (Van der Maaten & Hinton, 2008), UMAP (McInnes et al., 2018), and PHATE (Moon et al., 2019), have gained popularity. These methods have proven effective in capturing lower-dimension structures across diverse fields, including cancer research (Zikry et al., 2024; Abdelmoula et al., 2016), neuroscience (Lee et al., 2021; He et al., 2021), and chemistry (Pouyet et al., 2018). However, unlike principal component analysis (PCA), these methods are *nonprojective*, meaning the function f is unknown; to project new data onto an existing manifold, whether it be from a replicate or additional treatment condition, the entire model needs to be refit at computational cost.

Furthermore, these methods are not inherently interpretable. Here, we define interpretable machine learning (IML) as in Allen et al. (2023), where the goal is to "generate human-understandable insights into data, the learned model, or the model output." Models such as tSNE, UMAP, and PHATE are black boxes. Without post-hoc metrics, users cannot precisely determine how and why a manifold is being constructed from their original data, a concern that is especially critical when these

methods inform conclusions in biomedicine (Patra et al., 2021). Past work shows how these methods are extremely sensitive to hyperparameter choices (Xia et al., 2024), random initializations (Kobak & Linderman, 2021) and most significantly, feature selection (Ranek et al., 2024).

A surrogate model is designed to mimic the predictions of a complex black-box model to enable insights into the underlying decision-making process Nóbrega & Marinho (2019). Local interpretable model-agnostic explanations (LIME) apply this principle to black box supervised classification models, converting the decision boundary into an optimization problem with regularization to find the best local linear approximation (Ribeiro et al., 2016). Further work applies this concept in DR (Bardos et al., 2022; Björklund et al., 2023; Bibal et al., 2020; Zeng et al., 2023; Collaris et al., 2023), where the decision boundary is determined by the placement of individual observations, instead of being drawn across classes. However, this past work does not adequately address how to constrain the optimization problem across the different dimensions of r; instead of leveraging information across components, many methods apply regularization by treating the dimensions independently or averaging their effects, rather than fully exploiting all available information.

Contribution Here, we present our method for Local Interpretable Manifold Explanations for Dimension Evaluations (LIMEADE). Unlike prior approaches, we leverage the full dimensionality of the manifold representation to train a surrogate model for interpretability. Instead of averaging across the dimensions of the reduced space Y, we formulate a multitask regression problem to learn an interpretable linear projection. To select covariates across the dimensions of r, we apply the Group Lasso, a multivariate extension of Lasso designed to select groups of features. We apply LIMEADE to a single-cell proteomics dataset modeling the cell cycle, where significant biological expert annotation and post-hoc feature importance was required to identify a small subset of features most relevant to the cell cycle. We demonstrate how LIMEADE can select the most relevant features and effectively recapitulate the underlying the cell cycle manifold.

2 METHOD

Notation and Model For LIMEADE, we take the input pair (X, Y), where $X_{n\times p}$ is our normalized original data, along with a reduced feature embedding after dimension reduction $Y_{n\times r}$, where $r \ll p$. The goal is to learn the f such that $Y = f(X)$. With projective methods such as PCA, this function is defined, and for a new data point x_{n+1}, we can easily project into $\mathbb{R}^r$ via $\hat{y}_{n+1} = f(x_{n+1})$. However, with nonprojective DR methods, f is unknown. Hence, we seek to learn a linear projection as a surrogate model of this relationship. Additionally, we apply the Group Lasso (Yuan & Lin, 2006) for regularization and feature sparsity *across* the dimensions of r. Unlike standard Lasso, which applies an ℓ_1 penalty to individual coefficients and promotes element-wise sparsity, the Group Lasso extends this framework by penalizing groups of coefficients together, which allows entire features to be either fully included or excluded across all dimensions of r. Specifically, we seek to learn $V_{p\times r}$, such that:

$$\hat{V} = \arg\min_{V} \left(\frac{1}{2}||Y - XV||_2^2 + \lambda \sum_{j=1}^{p} ||V_j||_2 \right) \tag{1}$$

where $j = 1, \ldots, J$ are individual features, grouping across dimensions of the reduced embedding. Equation 1 is the global formulation of LIMEADE; creating a single linear projection for the entire data. If there exist partitions such that $(X, Y) = [(X_{M_1}, Y_{M_1}), \ldots, (X_{M_k}, Y_{M_k})]$, then for each of the k partitions, we can form separate hyperplanes,

$$\hat{V}_{M_k} = \arg\min_{V} \left(\frac{1}{2}||Y_{M_k} - X_{M_k}V_k||_2^2 + \lambda \sum_{j=1}^{p} ||V_j||_2 \right) \tag{2}$$

Further formulations may include single-point interpretations, where a local neighborhood is defined using a radial basis function to weight nearby points, described in Section A.2, and previously seen in Ribeiro et al. (2016); Collaris et al. (2023); Bardos et al. (2022). We implement LIMEADE in R with the package `grpreg` using EBIC tuning (Chen & Chen, 2008) with $\gamma = 0.5$.

Real Data Application We evaluate the single-cell proteomics dataset described in Stallaert et al. (2022) as a motivating example for LIMEADE. This is an iterative indirect immunofluorescence imaging (4i) dataset collected in 8,850 asynchronous retinal pigment epithelium cells, where each of the 187 experimental markers collected is hypothesized to be an indicator of cell cycle progression. A primary goal of this study was to identify a small subset of these markers as key factors for cells progressing or exiting the cell cycle. Within the 187 markers, many are highly correlated, including measurements of the same protein in different parts of the cell. To this end, the authors first developed stepwise random forest regression models to predict cell cycle phase and ages, and took the top 40 features, measured via permutation importance, to feed into their DR method of choice—a 2-dimensional PHATE (Moon et al., 2019). Here, we seek to determine whether LIMEADE can recover a similar or more interpretable set of features to recapitulate the cell cycle. If so, it would provide a more direct and systematic approach to interpreting high-dimensional feature selection with regards to understanding a lower dimensional representation.

3 RESULTS

Simulation Study To test the ability of LIMEADE to globally recover significant features, we first generate data from a Gaussian mixture model with low-rank cluster centroids, following Zheng et al. (2024). Results can be seen in Table 3, where we assess the true positive rate (TPR) and false positive rate (FPR) of the nonzero features of $\hat{V}$. As we increase SNR, global LIMEADE does well in recovering true features and controlling false discovery at $n = 1000$ with both $p = 100$ and $p = 1000$. Next, we simulate differentially expressed blocks of features across clusters to test the feature recovery of partitioned LIMEADE, with results in Table 3. Again, we see extremely strong TPR and FPR results, with high detection of true features even at low SNR in the $p = 100$ setting, and extremely low false discovery rates. Full simulation details can be found in Appendix Section A.1. Global simulations across SNR of Y overlaid with cluster labels are visible in Figure 2, and representative plots of the simulated partitioned data are provided in Figure 3.

Table 1: Global Simulation

SNR	p=100		p=1000	
	TPR	**FPR**	**TPR**	**FPR**
0	0.302	0.306	0.004	0.005
1	0.442	0.264	0.008	0.005
2	0.662	0.066	0.135	0.004
3	0.724	0.016	0.389	0.001
4	0.782	0.031	0.488	0.000
5	0.864	0.063	0.570	0.000

Table 2: Partitioned Simulation

SNR	p=100		p=1000	
	TPR	**FPR**	**TPR**	**FPR**
0	0.212	0.210	0.001	0.002
0.25	0.232	0.203	0.003	0.002
0.5	0.576	0.169	0.005	0.002
0.75	0.846	0.083	0.015	0.002
1	0.909	0.042	0.169	0.000
1.25	0.968	0.000	0.636	0.000
1.5	0.997	0.000	0.849	0.000

Cell Cycle Manifold Projections We first seek to assess how well LIMEADE allows us to project new data onto the original manifold. We split our 4i dataset $X \in \mathbb{R}^{8850\times 187}$ and PHATE embedding $Y \in \mathbb{R}^{8850\times 2}$ into training and testing via an 80/20 split. We train first train a global LIMEADE on X_{train} and Y_{train} to get $\hat{V}$ via Equation 1. Then we project held-out data $Y_{test} = X_{test}\hat{V}$, with results visible in Figure 1A,B. As with other nonlinear DR methods, the axes of PHATE hold no intrinsic meaning and are left off the plots. After removing several extreme outliers, we observe that global LIMEADE performs reasonably well, accurately projecting new data onto the existing cell cycle manifold, making it sufficient for interpretation. Although some subtleties of the embedding are obfuscated, cells from the same cell cycle phase remain clustered together, and the phases appear in a canonical order, preserving the biological structure. Next, we apply partitioned LIMEADE over individual cell cycle phase clusters, as shown in Figures 1D and E. This approach further improves the accuracy of the projections, better capturing the nuanced, heterogeneous relationships within specific cell cycle stages.

Cell Cycle Interpretations When applying global LIMEADE to the cell cycle 4i data, the Group Lasso selects 24 nonzero features, instead of the 40 features selected to use for DR in Stallaert et al.

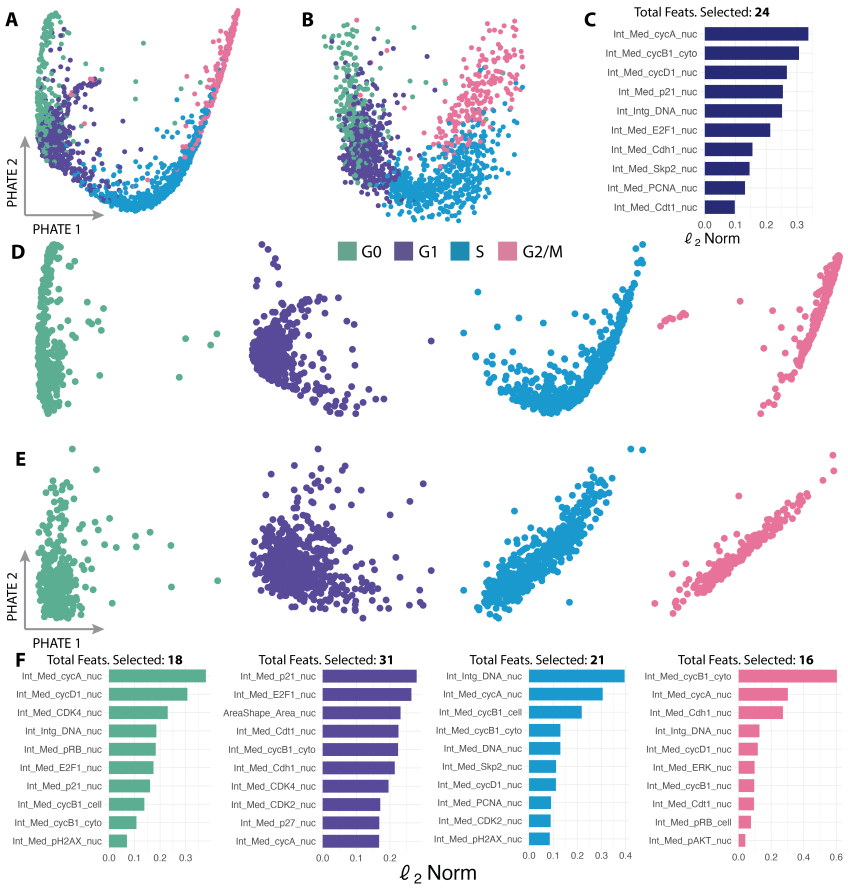

Figure 1: Global LIMEADE recapitulation of single-cell manifold from cell cycle 4i data. **A:** Original manifold Y_{test} generated using PHATE. **B.** Global LIMEADE linear projections $\hat{Y}_{test}$ learned from 4i features and PHATE embedding. **C.** Top 10 features from global LIMEADE. **D.** Original manifold Y_{test} for each cell cycle phase cluster. **E.** Partitioned LIMEADE linear projections. **F.** Top 10 features for partitioned LIMEADE by cell cycle phase.

(2022), after taking the ℓ_2 norm across columns of r in $\hat{V}$. We first examine the top 10 features selected from global LIMEADE, seen in Figure 1C. When compared to the top 10 features found to be important from Stallaert et al. (2022), 8 of the 10 features are the same. Of the original 40 features selected from the 4i study, 20 of them are selected from global LIMEADE, with 4 selected by LIMEADE not found in the original top 40. The full feature lists, along with overlaps, differences, and compared ranks, are provided in Table 3 in the Appendix.

As with standard Lasso, the Group Lasso will only select a single feature from a group of correlated features. The authors in Stallaert et al. (2022) and the original imaging technology paper Gut et al. (2018) note this correlation as part of the immunofluorescent imaging design. For example, the original model selected the cyclinB1 measurement for the ring, the cell, the cytoplasm, and a mean edge measurement. As noted specifically for cyclinB1, these are clearly correlated measurements due to the localization patterns of proteins (Porter et al., 2003), and hence possibly redundant information for the model. As expected, when using partitioned LIMEADE, different features are selected for different cell cycle clusters (Figure 1F). For example, partitioned LIMEADE, p21 is found to be the most important feature in G1, supporting known cell cycle dynamics (Barr et al., 2017). In G2/M, cytoplasmic cyclinB1 is found to be the most important feature, supported by how cyclinB1 is primarily cytoplasmic during interphase (Gong & Ferrell Jr, 2010). Moreover, each partitioned projection selects different numbers of features, with G1 at the most with 31 features

selected, compared to G2/M with only 16. To investigate G1 heterogeneity further, we carry out a locally-weighted LIMEADE in Section 3 on three selected cells along the G1 region of the manifold.

Conclusion Our robust simulation study and single-cell analysis demonstrate how LIMEADE deepens the understanding of black-box dimensionality reduction methods. With the novel utilization of the Group Lasso, LIMEADE can help researchers across domains identify key features driving low-dimensional embeddings along single observation, cluster-wide, or global scales, and improve interpretability for data-driven insights into complex high-dimensional datasets.

ACKNOWLEDGEMENTS

G.I.A acknowledges funding from NSF DMS-1554821.

REFERENCES

Walid M Abdelmoula, Benjamin Balluff, Sonja Englert, Jouke Dijkstra, Marcel JT Reinders, Axel Walch, Liam A McDonnell, and Boudewijn PF Lelieveldt. Data-driven identification of prognostic tumor subpopulations using spatially mapped t-sne of mass spectrometry imaging data. *Proceedings of the National Academy of Sciences*, 113(43):12244–12249, 2016.

Genevera I Allen, Luqin Gan, and Lili Zheng. Interpretable machine learning for discovery: Statistical challenges and opportunities. *Annual Review of Statistics and Its Application*, 11, 2023.

Avraam Bardos, Ioannis Mollas, Nick Bassiliades, and Grigorios Tsoumakas. Local explanation of dimensionality reduction. In *Proceedings of the 12th Hellenic Conference on Artificial Intelligence*, pp. 1–9, 2022.

Alexis R Barr, Samuel Cooper, Frank S Heldt, Francesca Butera, Henriette Stoy, Jörg Mansfeld, Béla Novák, and Chris Bakal. Dna damage during s-phase mediates the proliferation-quiescence decision in the subsequent g1 via p21 expression. *Nature communications*, 8(1):14728, 2017.

Adrien Bibal, Viet Minh Vu, Géraldin Nanfack, and Benoît Frénay. Explaining t-sne embeddings locally by adapting lime. In *28th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning: ESANN2020*, pp. 393–398. ESANN (i6doc. com), 2020.

Anton Björklund, Jarmo Mäkelä, and Kai Puolamäki. Slisemap: Supervised dimensionality reduction through local explanations. *Machine Learning*, 112(1):1–43, 2023.

Jiahua Chen and Zehua Chen. Extended bayesian information criteria for model selection with large model spaces. *Biometrika*, 95(3):759–771, 2008.

Dennis Collaris, Pratik Gajane, Joost Jorritsma, Jarke J van Wijk, and Mykola Pechenizkiy. Lemon: Alternative sampling for more faithful explanation through local surrogate models. In *International Symposium on Intelligent Data Analysis*, pp. 77–90. Springer, 2023.

Delquin Gong and James E Ferrell Jr. The roles of cyclin a2, b1, and b2 in early and late mitotic events. *Molecular biology of the cell*, 21(18):3149–3161, 2010.

Gabriele Gut, Markus D Herrmann, and Lucas Pelkmans. Multiplexed protein maps link subcellular organization to cellular states. *Science*, 361(6401):eaar7042, 2018.

Jing He, Michael Kleyman, Jianjiao Chen, Aydin Alikaya, Kathryn M Rothenhoefer, Bilge Esin Ozturk, Morgan Wirthlin, Andreea C Bostan, Kenneth Fish, Leah C Byrne, et al. Transcriptional and anatomical diversity of medium spiny neurons in the primate striatum. *Current Biology*, 31 (24):5473–5486, 2021.

Dmitry Kobak and George C Linderman. Initialization is critical for preserving global data structure in both t-sne and umap. *Nature biotechnology*, 39(2):156–157, 2021.

Eric Kenji Lee, Hymavathy Balasubramanian, Alexandra Tsolias, Stephanie Udochukwu Anakwe, Maria Medalla, Krishna V Shenoy, and Chandramouli Chandrasekaran. Non-linear dimensionality reduction on extracellular waveforms reveals cell type diversity in premotor cortex. *Elife*, 10: e67490, 2021.

Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

Kevin R Moon, David Van Dijk, Zheng Wang, Scott Gigante, Daniel B Burkhardt, William S Chen, Kristina Yim, Antonia van den Elzen, Matthew J Hirn, Ronald R Coifman, et al. Visualizing structure and transitions in high-dimensional biological data. *Nature biotechnology*, 37(12):1482–1492, 2019.

Caio Nóbrega and Leandro Marinho. Towards explaining recommendations through local surrogate models. In *Proceedings of the 34th ACM/SIGAPP symposium on applied computing*, pp. 1671–1678, 2019.

Sudhansu Shekhar Patra, GM Harshvardhan, Mahendra Kumar Gourisaria, Jnyana Ranjan Mohanty, and Subham Choudhury. Emerging healthcare problems in high-dimensional data and dimension reduction. *Advanced Prognostic Predictive Modelling in Healthcare Data Analytics*, pp. 25–49, 2021.

Lisa A Porter, Daniel J Donoghue, et al. Cyclin b1 and cdk1: nuclear localization and upstream regulators. *PROGRESS IN CELL CYCLE RESEARCH.*, 5:335–348, 2003.

Emeline Pouyet, Neda Rohani, Aggelos K Katsaggelos, Oliver Cossairt, and Marc Walton. Innovative data reduction and visualization strategy for hyperspectral imaging datasets using t-sne approach. *Pure and Applied Chemistry*, 90(3):493–506, 2018.

Jolene S Ranek, Wayne Stallaert, J Justin Milner, Margaret Redick, Samuel C Wolff, Adriana S Beltran, Natalie Stanley, and Jeremy E Purvis. Delve: feature selection for preserving biological trajectories in single-cell data. *Nature Communications*, 15(1):2765, 2024.

Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. " why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1135–1144, 2016.

Wayne Stallaert, Katarzyna M Kedziora, Colin D Taylor, Tarek M Zikry, Jolene S Ranek, Holly K Sobon, Sovanny R Taylor, Catherine L Young, Jeanette G Cook, and Jeremy E Purvis. The structure of the human cell cycle. *Cell systems*, 13(3):230–240, 2022.

Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.

Laurens Van Der Maaten, Eric O Postma, H Jaap Van Den Herik, et al. Dimensionality reduction: A comparative review. *Journal of machine learning research*, 10(66-71):13, 2009.

Lucy Xia, Christy Lee, and Jingyi Jessica Li. Statistical method scdeed for detecting dubious 2d single-cell embeddings and optimizing t-sne and umap hyperparameters. *Nature Communications*, 15(1):1753, 2024.

Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 68(1):49–67, 2006.

Xingchen Zeng, Haowen Zhou, Zhicong Li, Chenqi Zhang, Juncong Lin, Jiazhi Xia, Yanyi Yang, and Xiaoyan Kui. ihelp: Interactive hierarchical linear projections for interpreting non-linear projections. *Journal of Visualization*, 26(3):631–648, 2023.

Lili Zheng, Andersen Chang, and Genevera I Allen. Cluster quilting: Spectral clustering for patchwork learning. *arXiv preprint arXiv:2406.13833*, 2024.

Tarek M Zikry, Samuel C Wolff, Jolene S Ranek, Harris M Davis, Ander Naugle, Namit Luthra, Austin A Whitman, Katarzyna M Kedziora, Wayne Stallaert, Michael R Kosorok, et al. Cell cycle plasticity underlies fractional resistance to palbociclib in er+/her2- breast tumor cells. *Proceedings of the National Academy of Sciences*, 121(7):e2309261121, 2024.

A Appendix

A.1 Simulation Design

Global Simulation Following Zheng et al. (2024), we first randomly assign n points into K clusters via a one-hot encoded matrix $F^* \in \mathbb{R}^{n \times K}$. Next, we compare a rank r matrix $W \in \mathbb{R}^{p \times r}$, a linear indepedent matrix with entries randomly selected from $\{-SNR, 0, SNR\}$ Then, we generate the feature loadings matrix $Z \in \mathbb{R}^{p \times r}$, where

$$Z = \begin{bmatrix} Z_{\text{signal}} \in \mathbb{R}^{p_{\text{signal}} \times r} \\ Z_{\text{noise}} \in \mathbb{R}^{(p - p_{\text{signal}}) \times r} \end{bmatrix}.$$

The first p_{signal} rows, Z_{signal}, are drawn from a random orthogonal matrix, ensuring that the informative features capture structured variation across dimensions. The remaining $p - p_{\text{signal}}$ rows, Z_{noise},

are set to zero, ensuring that these features contain no meaningful signal in the low-dimensional representation. We then form out cluster centroid matrix $\Theta^* = WZ^\top \in \mathbb{R}^{K\times p}$, and data X are simulated from the multivariate normal distribution with mean $F^*\Theta^*$ and identity covariance matrix. Y is then generated from 2-dimensional UMAP with default parameters before applying LIMEADE. With $n = 1000$, $p_{signal} = 20$, $K = 3$, and $r = 2$, we increase SNR in increments of 0.5 and assess true positive rate (TPR) and false positive rate (FPR) for feature recovery as seen in Figure 3, averaged over 500 replicates.

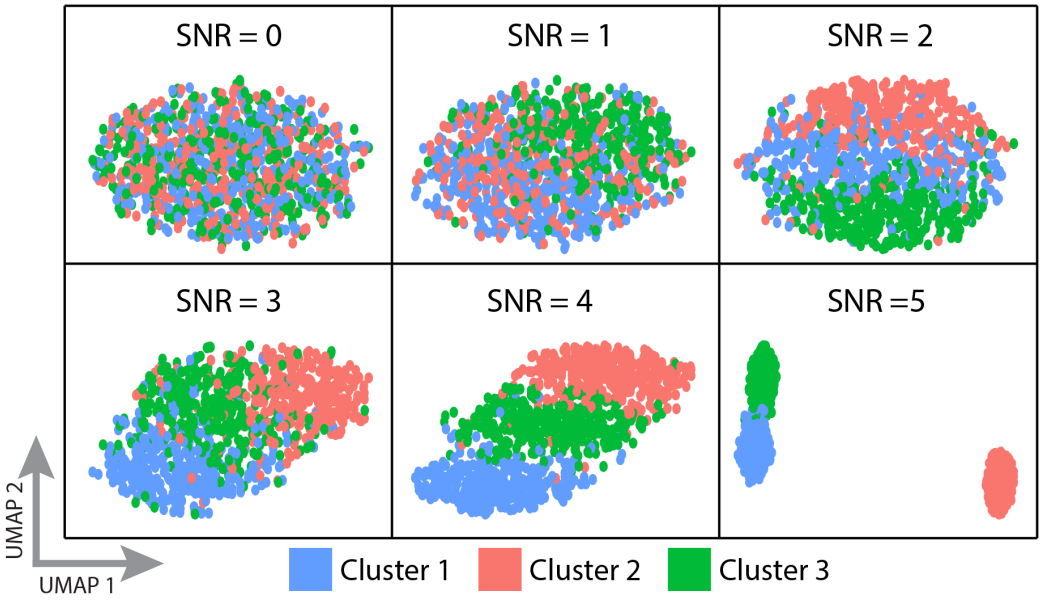

Figure 2: UMAP embedding of simulated low-rank cluster centroids across values of SNR for a single replicate.

Partitioned Simulation We simulate a dataset with 3 clusters, each having different subsets of $p_{signal} < p$ features from differentially expressed across the population. To represent differential expression, we generate a pseudotime time variable t sampled evenly from $(0, 2\pi)$ for each cluster. The signal features then follow a sinusoidal trajectory along the population, where for observation i and signal feature j:

$$X_{ij} = SNR \cdot \sin(t_i + j \cdot \pi / p_{\text{signal}}) + \epsilon_{ij}$$

with random Gaussian noise ϵ_{ij}, and SNR modulates the strength of the signal features over the noise. All other $p - p_{signal}$ noise features are generated from random Gaussian noise.

We choose each cluster matrix to have partially overlapping different signal feature blocks, each with 20 signal features. Cluster 1 has signal features $1:20$, cluster 2 with $11:30$, and cluster 3 with $21-40$. After generating these clusters into dataset X, we carry out UMAP on the entire data to get a reduced embedding Y. Representative visualizations are provided in Figure 3. Then, we partition the pairing (X, Y) into the separate clusters $[(X_1, Y_1), (X_2, Y_2), (X_3, Y_3)]$, and compute LIMEADE on each pairing separately to generate $[\hat{V}_1, \hat{V}_2, \hat{V}_3]$ and compute the true positive rate (TPR) and false positive rate (FPR) for feature recovery.

A.2 Local LIMEADE

We can additionally formulate a locally-weighted LIMEADE, similar to the original LIME formulation (Ribeiro et al., 2016).

$$\widehat{V} = \arg\min_V \left(\frac{1}{2} ||W^{1/2}(Y - XV)||_2^2 + \lambda \sum_{j=1}^{p} ||V_j||_2 \right) \tag{3}$$

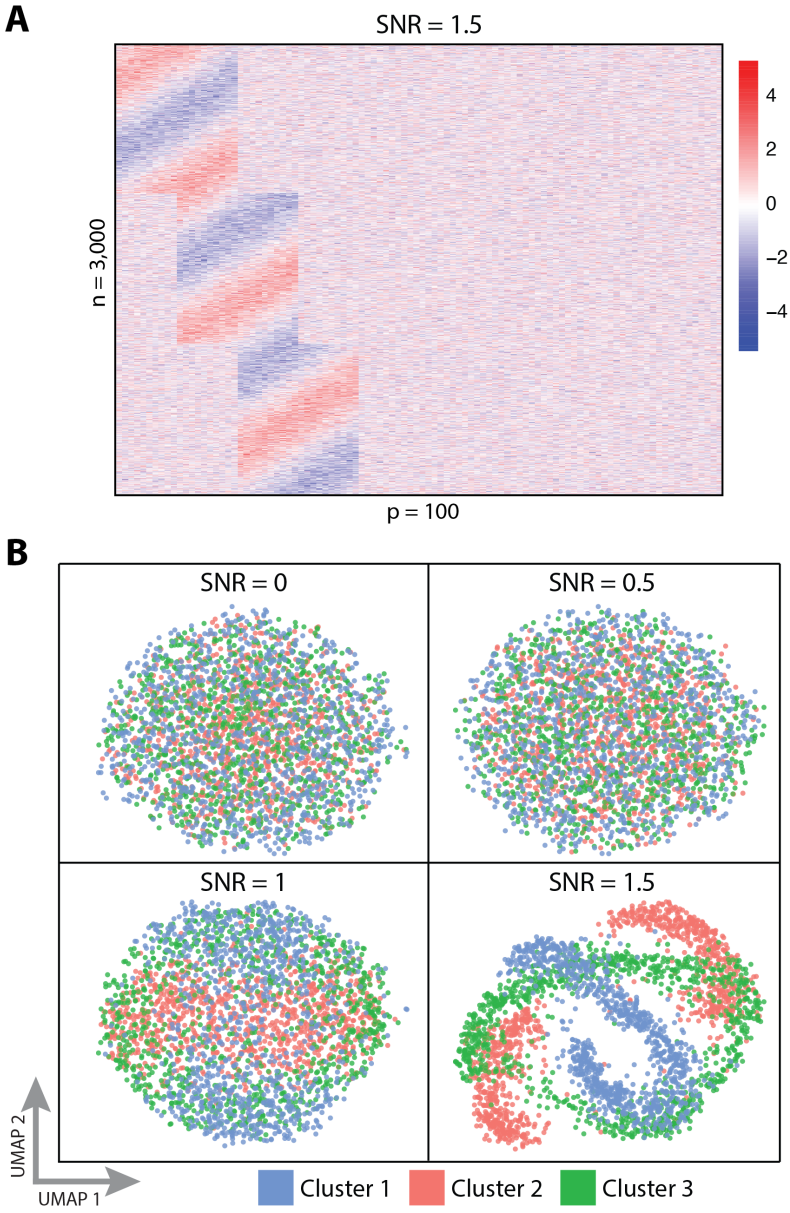

Figure 3: Simulated data with partitioned signal features. **A.** Representative data from a single replicate for simulated partitioned data described in Section A.1, with SNR of 1.5. **B.** UMAP embedding of clusters seen in **A**.

where, for row $x_i \in X$,

$$W_i = \text{diag}(w_1, w_2, \ldots, w_n), w_i = \exp\left(-\gamma \|x_i - X\|^2\right)$$

For new point x_{n+1}, we can first find the nearest neighbor in the original space: $x_j = \text{NearestNeighbor}(x_{n+1}, X)$, and take the associated $\hat{v}_j$, and then project based on that: $\hat{y}_{n+1} = x_{n+1}\hat{v}_j$. Existing literature focuses on variations of how to create these local neighborhoods through sampling methods (Bibal et al., 2020; Collaris et al., 2023), but with the cell cycle 4i data, we hand-select points of interest.

From the analysis of cell cycle phase partitions in Section 3, we see that the G1 cluster has selected the most features of any cell cycle phase. Therefore, we hypothesize that G1 may have the most

heterogenous sub-phase states. From In Figures 4A and B, we handpick three cells across G1 in the PHATE structure; one close to G0 cells, one close to G2/M cells, and one close to S phase cells. After carrying out locally-weighted LIMEADE, with $\gamma = 1/4p$, where p is the original number of features after preprocessing (176) from cell cycle 4i data X, we obtain three $\hat{V}$ estimates, one per point.

The top 10 features from each $\hat{V}$ are listed in Figure 4C. LIMEADE detects significant heterogeneity where features are found to be important in the respective cell placement on the manifold. The green point, nearest to S phase cells, has p21 protein as its top feature, as opposed to cyclinD1 for the cell near G2/M and cyclinA for the cell near G0.

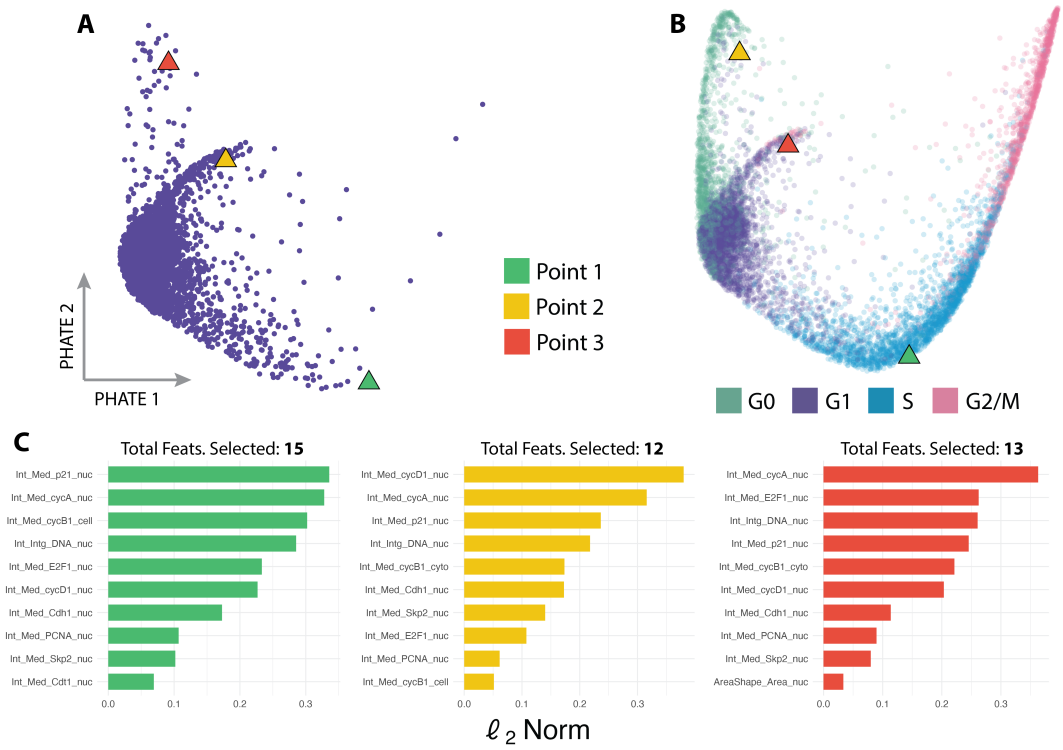

Figure 4: Local interpretations of heterogenous cells in G1. **A.** Selected points in different manifold locations in G1. **B.** Selected points in G1 placed along the population manifold to visualize proximity to different phases. **C.** LIMEADE feature interpretations for each of the three points along G1.

Feature	Global LIMEADE Rank	Original Rank
Int_Med_cycA_nuc	1	2
Int_Med_cycB1_cyto	2	16
Int_Med_cycD1_nuc	3	3
Int_Med_p21_nuc	4	4
Int_Intg_DNA_nuc	5	5
Int_Med_E2F1_nuc	6	1
Int_Med_Cdh1_nuc	7	9
Int_Med_Skp2_nuc	8	6
Int_Med_PCNA_nuc	9	11
Int_Med_Cdt1_nuc	10	7
Int_Med_p27_nuc	11	20
Int_Med_pRB_nuc	12	18
Int_Med_CDK4_nuc	13	27
AreaShape_Area_nuc	14	8
Int_Med_CDK2_nuc	15	26
Int_Med_pH2AX_nuc	16	12
Int_Med_PCNA_cyto	17	73
Int_Med_Skp2_cyto	18	81
Int_Med_DNA_nuc	19	28
Int_Med_cFos_nuc	20	29
Int_Med_cycE_nuc	21	10
Int_Med_E2F1_cell	22	48
Int_Med_p21_cell	23	211
Int_Med_CDK6_nuc	24	32
Int_Med_cMyc_nuc	–	13
Int_Med_RB_nuc	–	14
Int_Med_cycB1_cell	–	15
Int_Med_cycB1_ring	–	17
Int_Intg_p21_nuc	–	19
Int_Med_ERK_nuc	–	21
Int_Med_pp53_nuc	–	22
Int_Med_S6_nuc	–	23
Int_MeanEdge_cycB1_cell	–	24
Int_Med_pp21_nuc	–	25
AreaShape_Area_cell	–	30
Int_MeanEdge_Fra1_cell	–	31
Int_Med_p16_nuc	–	33
Int_Med_cycD1_cyto	–	34
Int_Med_GSK3b_nuc	–	35
Int_Med_p38_nuc	–	36
Int_Med_Bcl2_nuc	–	37
Int_Med_CDK2_cyto	–	38
Int_MeanEdge_Skp2_cell	–	39
Int_Med_cycA_cell	–	40

Table 3: Comparison of global LIMEADE feature ranking and original cell cycle 4i ranking. Red denotes features selected by global LIMEADE not included in the most important features for manifold visualization in Stallaert et al. (2022).