

Feature selection for preserving biological trajectories in single-cell data

Jolene S. Ranek^{1,2}, Wayne Stallaert³, Justin Milner⁴, Natalie Stanley^{*2,5}, and Jeremy E. Purvis^{*1,2}

¹Department of Genetics, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

²Computational Medicine Program, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

³Department of Computational and Systems Biology, University of Pittsburgh, Pittsburgh, PA, USA

⁴Department of Microbiology and Immunology, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

⁵Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

May 9, 2023

Abstract

Single-cell technologies can readily measure the expression of thousands of molecular features from individual cells undergoing dynamic biological processes, such as cellular differentiation, immune response, and disease progression. While examining cells along a computationally ordered pseudotime offers the potential to study how subtle changes in gene or protein expression impact cell fate decision-making, identifying characteristic features that drive continuous biological processes remains difficult to detect from unenriched and noisy single-cell data. Given that all profiled sources of feature variation contribute to the cell-to-cell distances that define an inferred cellular trajectory, including confounding sources of biological variation (e.g. cell cycle or metabolic state) or noisy and irrelevant features (e.g. measurements with low signal-to-noise ratio) can mask the underlying trajectory of study and hinder inference. Here, we present DELVE (dynamic selection of locally covarying features), an unsupervised feature selection method for identifying a representative subset of dynamically-expressed molecular features that recapitulates cellular trajectories. In contrast to previous work, DELVE uses a bottom-up approach to mitigate the effect of unwanted sources of variation confounding inference, and instead models cell states from dynamic feature modules that constitute core regulatory complexes. Using simulations, single-cell RNA sequencing data, and iterative immunofluorescence imaging data in the context of the cell cycle and cellular differentiation, we demonstrate that DELVE selects features that more accurately characterize cell populations and improve the recovery of cell type transitions. This feature selection framework provides an alternative approach for improving trajectory inference and uncovering co-variation amongst features along a biological trajectory. DELVE is implemented as an open-source python package and is publicly available at: <https://github.com/jranek/delve>.

1 Introduction

2 High-throughput single-cell technologies, such as flow and mass cytometry [1, 2, 3], single-cell RNA sequencing [4, 5, 6, 7], and imaging-based profiling techniques [8, 9, 10, 11] have transformed our ability to study how cell 3 populations respond and dynamically change during processes like development [12, 13, 14, 15] and immune 4 response [16, 17, 18]. By profiling many features (e.g. proteins or genes) for many thousands of cells from a 5 biological sample, these technologies provide high-dimensional snapshot measurements that can be used to gain 6 fundamental insights into the molecular mechanisms that govern phenotypic changes.

7 Trajectory inference methods [19] have been developed to model dynamic biological processes from snapshot 8 single-cell data. By assuming cells are asynchronously changing over time such that a profiled biological sample 9 from a single experimental time point describes a range of the underlying dynamic process, computational trajectory 10 inference approaches have leveraged minimum spanning tree approaches [20, 21, 22], curve-fitting [23, 24], graph- 11 based techniques [25, 26, 27], probabilistic approaches [28, 29, 30], or optimal transport [31, 32] to order cells 12 based on their similarities in feature expression. Once a trajectory model is fit, regression [33, 34, 35] can be 13

*Correspondence to Natalie Stanley, email: natalies@cs.unc.edu and Jeremy Purvis, email: purvisj@email.unc.edu

14 performed along estimated pseudotime (e.g. distance through the inferred trajectory from a start cell) to identify
15 specific cell state changes associated with differentiation or disease trajectories. Moreover, these inferred cellular
16 trajectories have the potential to elucidate higher-order gene interactions [36], gene regulatory networks [37],
17 predict cell fate probabilities [29], or find shared mechanisms of expression dynamics across disease conditions or
18 species [38, 39].

19 While trajectory analysis has proven useful in the context of single-cell biology, the identification of characteristic
20 genes or proteins that drive continuous biological processes relies on having inferred accurate cellular trajectories,
21 which can be challenging, especially when trajectory inference is performed on the original full unenriched dataset.
22 Single-cell data are noisy measurements that suffer from limitations in detection sensitivity, where dropout [40],
23 low signal-to-noise, or sample degradation [41] can result in spurious signals that can overwhelm true biological
24 differences. Furthermore, all profiled sources of feature variation contribute to the cell-to-cell distances that
25 define the inferred cellular trajectory; thus, including confounding sources of biological variation (e.g. cell cycle,
26 metabolic state) or irrelevant features (e.g. extracted imaging measurements that contain low signal-to-noise ratio)
27 can distort or mask the intended trajectory of study [42, 43]. With the accumulation of large-scale single-cell data
28 and multi-modal measurements [44], appropriate filtering of noisy, information-poor, or irrelevant features can
29 serve as a crucial and necessary step for cell type identification, inference of dynamic phenotypes, and identification
30 of putative driver features (e.g. genes, proteins).

31 Feature selection methods [45] are a class of supervised or unsupervised approaches that can remove redundant
32 or information-poor features prior to performing trajectory inference, and therefore, they have great potential
33 for improving the interpretation of downstream analysis, while easing the computational burden by reducing
34 dataset dimension. In the supervised-learning regime, classification-based [46] or information-theoretic approaches
35 [47, 48] have been used to evaluate features according to their discriminative power or association with cell types.
36 Despite having great power to detect biologically-relevant features, these methods rely on expensive or laborious
37 manual annotations (e.g. cell types) which are often unavailable [49] thus precluding them from use. In the
38 unsupervised-learning regime, computational approaches often aim to identify relevant features based on intrinsic
39 properties of the complete dataset; however, these methods have some limitations with respect to retaining features
40 that are useful for defining cellular trajectories. Namely, although unsupervised variance-based approaches [50, 51],
41 which effectively sample features based on their overall variation across cells, have been extensively used to identify
42 features that define cell types without the need for ground truth annotations, (1) they can be overwhelmed by
43 noisy or irrelevant features that dominate data variance, and (2) are insensitive to lineage-specific features (e.g.
44 transcription factors) that have a small variance and gradual progression of expression. Alternatively, unsupervised
45 similarity-based [52, 53, 25] or subspace-learning [54, 55] feature selection methods evaluate features according
46 to their association with a cell-similarity graph defined by all features or the underlying structure of the data (e.g.
47 pairwise similarities defined by uniform manifold approximation and projection (UMAP) [56], eigenvectors of
48 the graph Laplacian matrix [57]). While these approaches have the potential to detect smoothly varying genes
49 or proteins that define cellular transitions, they rely on the cell-similarity graph from the full dataset and can
50 fail to identify relevant features when the number of noisy features outweighs the number of informative ones
51 [58, 59].

52 To address these limitations, we developed DELVE (dynamic selection of locally covarying features), an unsupervised
53 feature selection method for identifying a representative subset of molecular features that robustly recapitulate
54 cellular trajectories. In contrast to previous work [55, 52, 50, 54, 53, 25], DELVE uses a bottom-up approach to
55 mitigate the effect of unwanted sources of variation confounding feature selection and trajectory inference, and
56 instead models cell states from dynamic feature modules that constitute core regulatory complexes. Features are
57 then ranked for selection according to their association with the underlying cell trajectory graph using data diffusion
58 techniques. We demonstrate the power of our approach for improving inference of cellular trajectories through
59 achieving an increased sensitivity to detect diverse and dynamically expressed features that better delineate cell
60 types and cell type transitions from single-cell RNA sequencing and protein immunofluorescence imaging data.
61 Overall, this feature selection framework provides an alternative approach for uncovering co-variation amongst
62 features along a biological trajectory.

63 Results

64 Overview of the DELVE algorithm

65 We propose DELVE, an unsupervised feature selection framework for modeling dynamic cell state transitions
66 using graph neighborhoods (Fig. 1). Our approach extends previous unsupervised similarity-based [52, 53, 25]

67 or subspace-learning feature selection [55] methods by computing the dependence of each gene on the cellular
 68 trajectory graph structure using a two-step approach. Inspired by the molecular events that occur during differ-
 69 entiation, where the coordinated spatio-temporal expression of key regulatory genes govern lineage specification
 70 [60, 61, 62, 63], we reasoned that we can approximate cell state transitions by identifying groups of features that
 71 are temporally co-expressed or co-regulated along the underlying dynamic process.

72 In step one, DELVE identifies groups of features that are temporally co-expressed by clustering features according
 73 to their average pairwise change in expression across prototypical cellular neighborhoods (Figure 1 step 1). As has
 74 been done previously [64, 65], we model cell states using a weighted k -nearest neighbor (k -NN) affinity graph,
 75 where nodes represent cells and edges describe the transcriptomic or proteomic similarity amongst cells according
 76 to all profiled features. Here, DELVE leverages a distribution-focused sketching method [66] to effectively sample
 77 cellular neighborhoods across all cell types. This sampling approach has three main advantages: (1) cellular
 78 neighborhoods are more reflective of the distribution of cell states, (2) redundant cell states are removed, and
 79 (3) fewer cellular neighborhoods are required to estimate feature dynamics resulting in increased scalability.
 80 Following clustering, each DELVE module contains a set of features with similar local changes in co-variation
 81 across prototypical cell states along the cellular trajectory. Feature-wise permutation testing is then used to assess
 82 the significance of dynamic expression variation across grouped features as compared to random assignment. By
 83 identifying and excluding modules of features that have static, random, or noisy patterns of expression variation,
 84 this approach effectively mitigates the effect of unwanted sources of variation confounding feature ranking and
 85 selection, and subsequent trajectory inference.

86 In step two, DELVE approximates the underlying cellular trajectory by constructing an affinity graph between cells,
 87 where cell similarity is now redefined according to a core set of dynamically expressed regulators. All profiled
 88 features are then ranked according to their association with the underlying cellular trajectory graph using graph
 89 signal processing techniques [67, 68] (Fig. 1 step 2). More concretely, a graph signal is any function that has a
 90 real defined value on all of the nodes. In this context, we consider all features as graph signals and rank them
 91 according to their total variation in expression along the cellular trajectory graph using the Laplacian Score [52].
 92 Intuitively, DELVE retains features that are considered to be globally smooth, or have similar expression values
 93 amongst similar cells along the approximate cellular trajectory graph. In contrast, DELVE excludes features that
 94 have a high total variation in signal, or expression values that are rapidly oscillating amongst neighboring cells, as
 95 these features are likely noisy or not involved in the underlying dynamic process that was seeded. The output of
 96 DELVE is a ranked set of features that best preserve the local trajectory structure. For a more detailed description
 97 on the problem formulation, the mathematical foundations behind feature ranking, and the impact of nuisance
 98 features on trajectory inference, see DELVE in the Methods section.

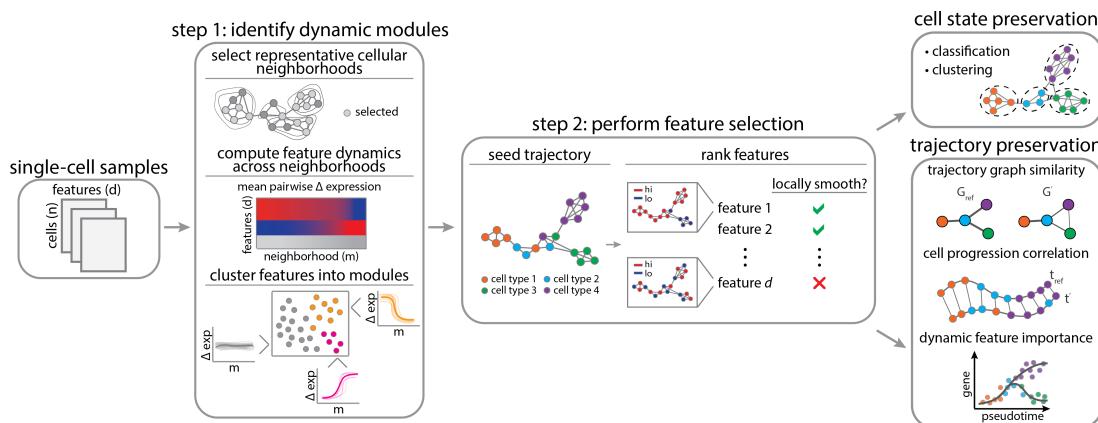


Figure 1: Schematic overview of the DELVE pipeline. Feature selection is performed in a two-step process. In step 1, DELVE clusters features according to their expression dynamics along local representative cellular neighborhoods defined by a weighted k -nearest neighbor affinity graph. Neighborhoods are sampled using a distribution-focused sketching algorithm that preserves cell-type frequencies and spectral properties of the original dataset [66]. A permutation test with a variance-based test statistic is used to determine if a set of features are (1) dynamically changing (dynamic) or (2) exhibiting random patterns of variation (static). In step 2, dynamic modules are used to seed or initialize an approximate cell trajectory graph and the trajectory is refined by ranking and selecting features that best preserve the local structure using the Laplacian Score [52]. In this study, we compare DELVE to the alternative unsupervised feature selection approaches on how well selected features preserve cell type and cell type transitions according to several metrics.

99 **DELVE outperforms existing feature selection methods on representing cellular trajectories 100 in the presence of single-cell RNA sequencing noise**

101 Although feature selection is a common preprocessing step in single-cell analysis [69] with the potential to reveal
102 cell-type transitions that would have been masked in the original high-dimensional feature space [42], there has
103 been no systematic evaluation of feature selection method performance on identifying biologically-relevant features
104 for trajectory analysis in single-cell data, especially in the context of noisy data that contain biological or technical
105 challenges (e.g. low total mRNA count, low signal-to-noise ratio, or dropout). In this study, we compared DELVE
106 to eleven other feature selection approaches and evaluated methods on their ability to select features that represent
107 cell types and cell type transitions using simulated data where the ground truth was known. In the sections below,
108 we will describe an overview of the feature selection methods considered and outline the simulation design and
109 evaluation criteria in more detail. We will then provide qualitative and quantitative assessments of how noise
110 impacts feature selection method performance and subsequent inference of cellular trajectories.

111 **Overview of feature selection methods**

112 We performed a systematic evaluation of twelve feature selection methods for preserving cellular trajectories in noisy
113 single-cell data. Feature selection methods were grouped into five general categories prior to evaluation: supervised,
114 similarity, subspace-learning, variance, and baseline approaches. For more details on the feature selection methods
115 implemented and hyperparameters, see Benchmarked feature selection methods and Supplementary Table 1.

116 **Supervised approaches:** To illustrate the performance of *ground-truth* feature selection that could be obtained
117 through supervised learning on expert annotated cell labels, we performed Random Forest classification. Random
118 Forest classification [46] is a supervised ensemble learning algorithm that uses an ensemble of decision trees to
119 partition the feature space such that all of the cells with the same cell type label are grouped together. Here, each
120 decision or split of a tree was chosen by minimizing the Gini Impurity score [70]. This approach was included to
121 provide context for unsupervised feature selection method performance.

122 **Similarity approaches:** We considered four similarity-based approaches as unsupervised feature selection methods
123 that rank features according to their association with a cell similarity graph defined by all profiled features (e.g.
124 Laplacian Score, neighborhood variance, hotspot) or dynamically-expressed features (e.g. DELVE). First, the
125 Laplacian Score (LS) [52] is an unsupervised locality-preserving feature selection method that ranks and selects
126 features according to (1) the total variation in feature expression across neighboring cells using a cell similarity graph
127 defined by all features and (2) a feature's global variance. Next, neighborhood variance [25] is an unsupervised
128 feature selection method that selects features with gradual changes in expression for building biological trajectories.
129 Here, features are selected if their variance in expression across local cellular neighborhoods is smaller than their
130 global variance. Hotspot [53] performs unsupervised feature selection through a local autocorrelation test statistic
131 that measures the association of a gene's expression with a cell similarity graph defined by all features. Lastly,
132 DELVE (dynamic selection of locally covarying features) is an unsupervised feature selection method that ranks
133 features according to their association with the underlying cellular trajectory graph. First, features are clustered into
134 modules according to changes in expression across local representative cellular neighborhoods. Next, modules of
135 features with dynamic expression patterns (denoted as *dynamic seed*) are used to construct an approximate cellular
136 trajectory graph. Features are then ranked according to their association with the approximate cell trajectory graph
137 using the Laplacian Score [52].

138 **Subspace learning approaches:** We considered two subspace-learning feature selection methods as unsupervised
139 methods that rank features according to how well they preserve the overall cluster structure (e.g. MCFS) or
140 manifold structure (e.g. SCMER) of the data. First, multi-cluster feature selection (MCFS) [55] is an unsupervised
141 feature selection method that selects features that best preserve the multi-cluster structure of data by solving an
142 L1 regularized least squares regression problem on the spectral embedding defined by all profiled features. The
143 optimization is solved using the least angles regression algorithm [71]. Next, single-cell manifold-preserving feature
144 selection (SCMER) [54] is an unsupervised feature selection method that selects a subset of features that best
145 preserves the pairwise similarity matrix between cells defined in uniform manifold approximation and projection
146 [56] based on all profiled features. To do so, it uses elastic net regression to find a sparse solution that minimizes
147 the KL divergence between a pairwise similarity matrix between cells defined by all features and one defined using
148 only the selected features.

149 **Variance approaches:** We considered two variance-based feature selection approaches (e.g. highly variable genes
150 [50], max variance) as unsupervised methods that use global expression variance as a metric for ranking feature
151 importance. First, highly variable gene selection (HVG) [50] is an unsupervised feature selection method that
152 selects features according to a normalized dispersion measure. Here, features are binned based on their average
153 expression. Within each bin, genes are then z-score normalized to identify features that have a large variance, yet a
154 similar mean expression. Next, max variance is an unsupervised feature selection method that ranks and selects
155 features that have a large global variance in expression.

156 **Baseline approaches:** We considered three baseline strategies (e.g. all, random, dynamic seed) that provide
157 context for the overall performance of feature selection. First, all features illustrates the performance when feature
158 selection is not performed and all features are included for analysis. Second, random features represents the
159 performance quality when a random subset of features are sampled. Lastly, dynamic seed features indicate the
160 performance from dynamically-expressed features identified in step 1 of the DELVE algorithm prior to feature
161 ranking and selection.

162 Single-cell RNA sequencing simulation design

163 To validate our approach and benchmark feature selection methods on representing cellular trajectories, we
164 simulated 90 single-cell RNA sequencing datasets with 1500 cells and 500 genes using Splatter. Splatter [72]
165 simulates single-cell RNA sequencing data with various trajectory structures (e.g. linear, bifurcation, tree) using
166 a gamma-Poisson hierarchical model. Importantly, this approach provides ground truth reference information
167 (e.g. cell type annotations, differentially expressed genes per cell type and trajectory, and a latent vector that
168 describes an individual cell's progression through the trajectory) that we can use to robustly assess feature selection
169 method performance, as well as quantitatively evaluate the limitations of feature selection strategies for trajectory
170 analysis. Moreover, to comprehensively evaluate feature selection methods under common biological and technical
171 challenges associated with single-cell RNA sequencing data, we added relevant sources of single-cell noise to the
172 simulated data. First, we simulated low signal-to-noise ratio by enforcing a mean-variance relationship amongst
173 genes; this ensures that lowly expressed genes are more variable than highly expressed genes. Next, we modified
174 the total number of profiled mRNA transcripts, or library size. This has been shown previously to vary amongst
175 cells within a single-cell experiment and can influence both the detection of differentially expressed genes [73],
176 as well as impact the reproducibility of the inferred lower-dimensional embedding [74]. Lastly, we simulated the
177 inefficient capture of mRNA molecules, or dropout, by undersampling gene expression from a binomial distribution;
178 this increases the amount of sparsity present within the data. For more details on the splatter simulation, see
179 *Splatter simulation*. For each simulated trajectory, we performed feature selection according to all described feature
180 selection strategies, and considered the top 100 ranked features for downstream analysis and evaluation.

181 Qualitative assessment of feature selection method performance

182 Prior to evaluating feature selection method performance quantitatively, we began our analysis with a qualitative
183 assessment of the importance of feature selection for representing cellular trajectories when the data contain
184 irrelevant or noisy genes. First, we visually compared the cellular trajectories generated from a feature selection
185 strategy with PHATE (potential of heat-diffusion for affinity-based transition embedding). PHATE [75] is a
186 nonlinear dimensionality reduction method that has been shown to effectively learn and represent the geometry of
187 complex continuous and branched biological trajectories. As an illustrative example, Fig. 2a shows the PHATE
188 embeddings for simulated linear differentiation trajectories generated from four feature selection approaches (all,
189 DELVE, Laplacian Score (LS), and random) when subjected to a decrease in the signal-to-noise ratio. Here,
190 we simulated a reduction in the signal-to-noise ratio and stochastic gene expression by modifying the biological
191 coefficient of variation (BCV) parameter within Splatter [72]. This scaling factor controls the mean-variance
192 relationship between genes, where lowly expressed genes are more variable than highly expressed genes (See
193 *Splatter simulation*). Under low noise conditions where the data contained a high signal-to-noise ratio, we observed
194 that excluding irrelevant features with DELVE or the Laplacian Score (LS) produced a much smoother, denoised
195 visualization of the linear trajectory, where cells were more tightly clustered according to cell type. This was
196 compared to the more diffuse presentation of cell states obtained based on all genes. We then examined how
197 noise influences the quality of selected features from a feature selection strategy. As the signal-to-noise ratio
198 decreased (high, medium, low), we observed that the linear trajectory became increasingly harder to distinguish,
199 whereby including both irrelevant and noisy genes often masked the underlying trajectory structure (Fig. 2a all
200 genes, medium to low signal-to-noise ratio). Furthermore, we found that unsupervised similarity-based or subspace
201 learning feature selection methods that initially define a cell similarity graph according to all irrelevant, noisy, and

202 informative genes often selected genes that produced noisier embeddings as the amount of noise increased (e.g. Fig. 2a LS: medium signal-to-noise ratio), as compared to DELVE (e.g. Fig. 2a DELVE medium signal-to-noise ratio). 203 We reason that this is due to spurious similarities amongst cells, reduced clusterability, and increased diffusion times. 204 These qualitative observations were consistent across different noise conditions (e.g. decreased signal-to-noise, 205 decreased library size, increased dropout) and trajectory types (e.g. linear, bifurcation, tree) (See Supplementary 206 Figs. 1 - 9). Although a qualitative comparison, this example illustrates how including irrelevant or noisy genes can 207 define spurious similarities amongst cells, which can (1) influence a feature selection method ability to identify 208 biologically-relevant genes and (2) impact the overall quality of an inferred lower dimensional embedding following 209 selection. Given that many trajectory inference methods use lower dimensional representations in order to infer 210 a cell's progression through a differentiation trajectory, it is crucial to remove information-poor features prior to 211 performing trajectory inference in order to obtain high quality embeddings, clustering assignments, or cellular 212 orderings that are reproducible for both qualitative interpretation and downstream trajectory analysis. 213

214 Quantitative assessment of feature selection method performance

215 We next quantitatively examined how biological or technical challenges associated with single-cell RNA sequencing 216 data may influence a feature selection method's ability to detect the particular genes that define cell types or cell type 217 transitions. To do so, we systematically benchmarked the 12 described feature selection strategies on their capacity 218 to preserve trajectories according to three sets of quantitative comparisons. Method performance was assessed by 219 evaluating if selected genes from an approach were (1) differentially expressed within a cell type or along a lineage, 220 (2) could be used to classify cell types, and (3) could accurately estimate individual cell progression through the 221 cellular trajectory. Fig. 2b-d shows feature selection method performance for simulated linear differentiation 222 trajectories when subjected to the technical challenge of having a reduction in the signal-to-noise ratio.

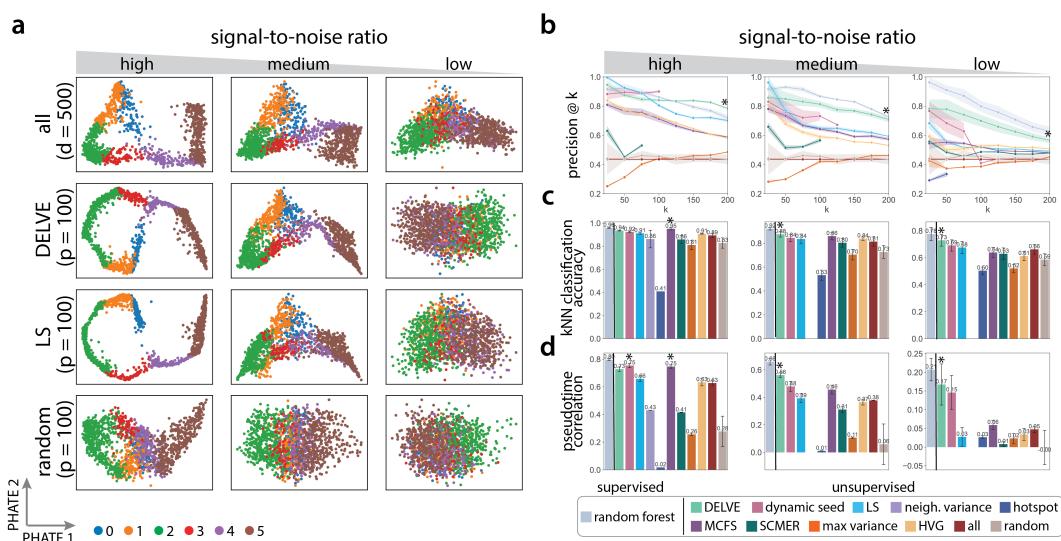


Figure 2: Comparison of feature selection methods on preserving linear trajectories when subjected to a reduction in the signal-to-noise ratio. (a) Example PHATE [75] visualizations of simulated linear trajectories using four feature selection approaches (all features, DELVE, Laplacian Score (LS) [52], and random selection) when subjected to a reduction in the signal-to-noise ratio (high, medium, low). Here, we simulated a reduction in the signal-to-noise ratio and stochastic gene expression by modifying the biological coefficient of variation (bcv) parameter within Splatter [72] that controls the mean-variance relationship between genes, where lowly expressed genes are more variable than highly expressed genes (high: bcv = 0.1, medium: bcv = 0.25, low: bcv = 0.5). d indicates the total number of genes ($d = 500$) and p indicates the number of selected genes following feature selection ($p = 100$). (b-d) Performance of twelve different feature selection methods: random forest [46], DELVE, dynamic seed features, LS [52], neighborhood variance [25], hotspot [53], multi-cluster feature selection (MCFS) [55], single-cell manifold preserving feature selection (SCMER) [54], max variance, highly variable gene selection (HVG) [50], all features, random features. Following feature selection, trajectory preservation was quantitatively assessed according to several metrics: (b) the precision of differentially expressed genes at k selected genes, (c) k -NN classification accuracy, and (d) pseudotime correlation to the ground truth cell progression across 10 random trails. Error bars/bands represent the standard deviation. * indicates the method with the highest median score. For further details across other trajectory types and noise conditions, see Supplementary Figs. 1 - 9.

223 First, we assessed the biological relevancy of selected genes, as well as the overall recovery of relevant genes
224 as the signal-to-noise ratio decreased by computing a precision score. Precision@k is a metric that defines the
225 proportion of selected genes (k) that are known to be differentially-expressed within a cell type or along a lineage
226 (See Precision@k). Overall, we found that DELVE achieved the highest precision@k score between selected genes
227 and the ground truth, validating that our approach was able to select genes that are differentially expressed and
228 was the strongest in defining cell types and cell type transitions (See Fig. 2b). Importantly, DELVE's ability to
229 recover informative genes was robust to the number of genes selected (k) and to the amount of noise present in the
230 data. In contrast, variance-based, similarity-based, or subspace-learning approaches exhibited comparatively worse
231 recovery of cell type and lineage-specific differentially expressed genes.

232 Given that a key application of single-cell profiling technologies is the ability to identify cell types or cell states
233 that are predictive of sample disease status, responsiveness to drug therapy, or are correlated with patient clinical
234 outcomes [76, 77, 78, 79, 65], we then assessed whether selected genes from a feature selection strategy can
235 correctly classify cells according to cell type along the underlying cellular trajectory; this is a crucial and necessary
236 step of trajectory analysis. Therefore, we trained a k -nearest neighbor (k -NN) classifier on the selected feature
237 set (see *k-nearest neighbor classification*) and compared the predictions to the ground truth cell type annotations
238 by computing a cell type classification accuracy score. Across all simulated trajectories, we found that DELVE
239 selected genes that often achieved the highest median k -NN classification accuracy score (high signal-to-noise ratio:
240 0.937, medium signal-to-noise ratio: 0.882, low signal-to-noise ratio: 0.734) and produced k -NN graphs that were
241 more faithful to the underlying biology (See Fig. 2c). Moreover, we observed a few results that were consistent
242 with the qualitative interpretations. First, removing irrelevant genes with DELVE, LS, or MCFS achieved higher
243 k -NN classification accuracy scores (e.g. high signal-to-noise ratio; DELVE = 0.937, LS = 0.915, and MCFS
244 = 0.955, respectively) than was achieved by retaining all genes (all = 0.900). Next, DELVE outperformed the
245 Laplacian Score, suggesting that using a bottom-up framework and excluding nuisance features prior to performing
246 ranking and selection is crucial for recovering cell-type specific genes that would have been missed if the cell
247 similarity graph was initially defined based on all genes. Lastly, when comparing the percent change in performance
248 as the amount of noise corruption increased (e.g. high signal-to-noise ratio to medium signal-to-noise ratio) for
249 linear trajectories, we found that DELVE often achieved the highest average classification accuracy score (0.905)
250 and lowest percent decrease in performance (-6.398%), indicating that DELVE was the most robust unsupervised
251 feature selection method to noise corruption (See Supplementary Fig. 10a). In contrast, the existing unsupervised
252 similarity-based or subspace learning feature selection methods that achieved high to moderate average k -NN
253 classification accuracy scores (e.g. MCFS = 0.905, LS = 0.874) had larger decreases in performance (e.g. MCFS =
254 -9.673%, LS = -8.390%) as the amount of noise increased. This further highlights the limitations of current feature
255 selection methods on identifying cell type-specific genes from noisy single-cell omics data.

256 Lastly, when undergoing dynamic biological processes such as differentiation, cells exhibit a continuum of cell
257 states with fate transitions marked by linear and nonlinear changes in gene expression [80, 81, 82]. Therefore, we
258 evaluated how well feature selection methods could identify genes that define complex differentiation trajectories
259 and correctly order cells along the cellular trajectory in the presence of noise. To infer cellular trajectories and
260 to estimate cell progression, we used the diffusion pseudotime algorithm [83] on the selected gene set from each
261 feature selection strategy, as this approach has been shown previously [19] to perform reasonably well for inference
262 of simple or branched trajectory types (See Trajectory inference and analysis). Method performance was then
263 assessed by computing the Spearman rank correlation between estimated pseudotime and the ground truth cell
264 progression. We found that DELVE approaches more accurately inferred cellular trajectories and achieved the
265 highest median pseudotime correlation to the ground truth measurements, as compared to alternative methods or
266 all features (See Fig. 2d). Furthermore, similar to the percent change in classification performance, we found
267 that DELVE was the most robust unsupervised feature selection method in estimating cell progression, as it often
268 achieved the highest average pseudotime correlation (0.645) and lowest percent decrease in performance (-22.761%)
269 as the amount of noise increased (See Supplementary Fig. 10b high to medium signal-to-noise ratio). In contrast, the
270 alternative methods incorrectly estimated cellular progression and achieved lower average pseudotime correlation
271 scores (e.g. MCFS = 0.602, LS = 0.526) and higher decreases in performance as the signal-to-noise ratio decreased
272 (MCFS = -38.884%, LS = -40.208%).

273 We performed this same systematic evaluation across a range of trajectory types (e.g. linear, bifurcation, tree)
274 and biological or technical challenges associated with single-cell data (See Supplementary Figs. 1 - 12). Fig. 3
275 displays the overall ranked method performance of feature selection methods on preserving cellular trajectories
276 when subjected to different sources of single-cell noise (pink: decreased signal-to-noise ratio, green: decreased
277 library size, and blue: increased dropout). Ranked aggregate scores were computed by averaging results across
278 all datasets within a condition; therefore, this metric quantifies how well a feature selection strategy can recover
279 genes that define cell types or cell type transitions underlying a cellular trajectory when subjected to that biological

280 or technical challenge (See Aggregate scores). Across all conditions, we found that DELVE often achieved an
 281 increased recovery of differentially expressed genes, higher cell type classification accuracy, higher correlation
 282 of estimated cell progression, and lower percent change in performance in noisy data. While feature selection
 283 method performance varied across biological or technical challenges, we found that the Laplacian score (LS) and
 284 multi-cluster feature selection (MCFS) performed reasonably well under low amounts of noise corruption and
 285 are often the second and third ranked unsupervised methods. Altogether, this simulation study demonstrates that
 286 DELVE more accurately recapitulates cellular dynamics and can be used to effectively interrogate cell identity and
 287 lineage-specific gene expression dynamics from noisy single-cell data.

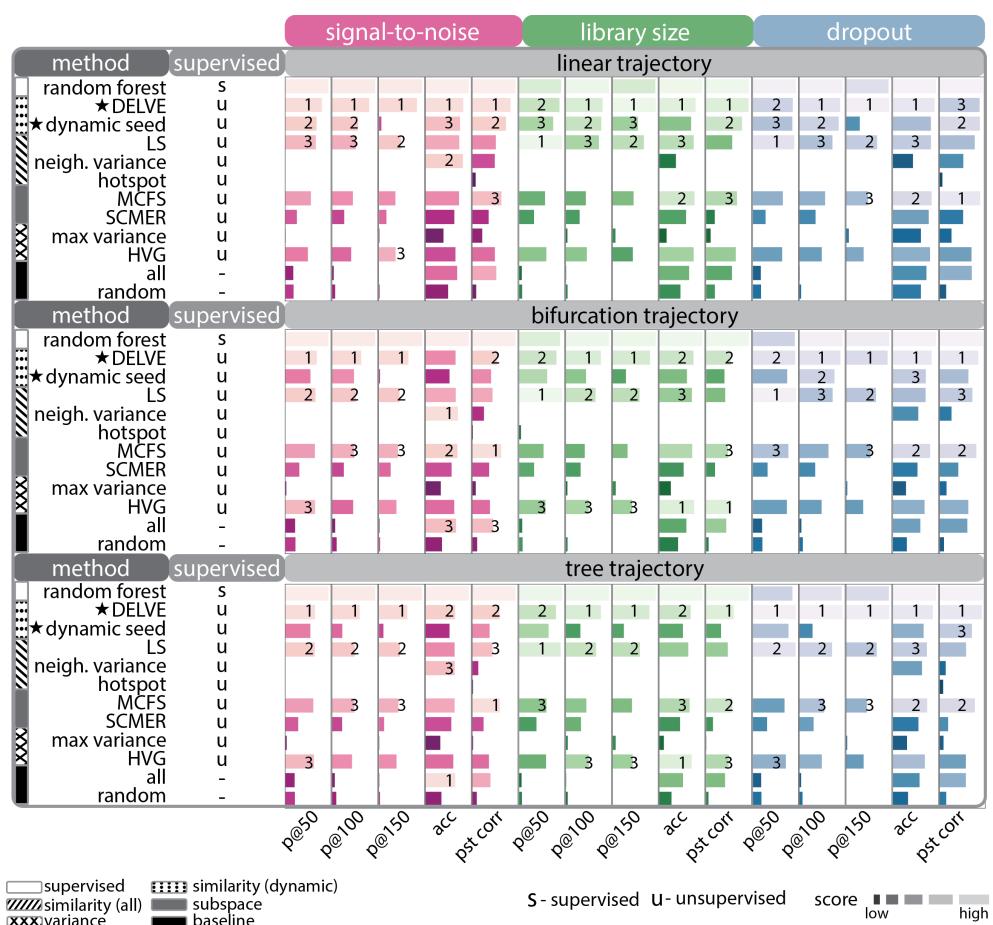


Figure 3: DELVE outperforms existing feature selection methods on representing trajectories in the presence of single-cell RNA sequencing noise. Feature selection methods were ranked by averaging their overall performance across datasets from different trajectory types (e.g. linear, bifurcation, tree) when subjected to noise corruption (e.g. decreased signal-to-noise ratio, decreased library size, and increased dropout). Several metrics were used to quantify trajectory preservation, including, precision of dynamically-expressed genes with 50 selected genes ($p@50$), precision at 100 selected genes ($p@100$), precision at 150 selected genes ($p@150$), k -NN classification accuracy of cell type labels (acc), and pseudotime correlation (pst). Here, higher-ranked methods are indicated by a longer lighter bar, and the star illustrates our approach (DELVE) as well as the performance from dynamic seed features of step 1 of the algorithm. DELVE often achieves the highest precision of lineage-specific differentially expressed genes, highest classification accuracy, and highest pseudotime correlation across noise conditions and trajectory types. Of note, random forest was included as a baseline representation to illustrate feature selection method performance when trained on ground truth cell type annotations; however, it was not ranked, as this study is focused on unsupervised feature selection method performance on trajectory preservation.

288 Revealing molecular trajectories of proliferation and cell cycle arrest

289 Recent advances in spatial single-cell profiling technologies [8, 9, 84, 10, 11, 85, 86, 87, 88] have enabled the
 290 simultaneous measurement of transcriptomic or proteomic signatures of cells, while also retaining additional
 291 imaging or array-derived features that describe the spatial positioning or morphological properties of cells. These

292 spatial single-cell modalities have provided fundamental insights into mammalian organogenesis [88, 89] and
293 complex immune responses linked to disease progression [90, 91]. By leveraging imaging data to define cell-to-
294 cell similarity, DELVE can identify smoothly varying spatial features that are strongly associated with cellular
295 progression, such as changes in cell morphology or protein localization.

296 To demonstrate this, we applied DELVE to an integrated live cell imaging and protein iterative indirect immunofluo-
297 rescence imaging (4i) dataset consisting of 2759 human retinal pigmented epithelial cells (RPE) undergoing the cell
298 cycle (See *RPE analysis*). In a recent study [92], we performed time-lapse imaging on an asynchronous population
299 of non-transformed RPE cells expressing a PCNA-mTurquoise2 reporter to record the cell cycle phase (G0/G1, S,
300 G2, M) and age (time since last mitosis) of each cell. We then fixed the cells and profiled them with 4i to obtain
301 measurements of 48 core cell cycle effectors. The resultant dataset consisted of 241 imaging-derived features
302 describing the expression and localization of different protein markers (e.g. nucleus, cytoplasm, perinuclear region –
303 denoted as ring), as well as morphological measurements from the images (e.g. size and shape of the nucleus).
304 Given that time-lapse imaging was performed prior to cell fixation, this dataset provides the unique opportunity to
305 rigorously evaluate feature selection methods on a real biological system (cell cycle) with technical challenges (e.g.
306 many features with low signal-to-noise ratio, autofluorescence, sample degradation).

307 We first tested whether DELVE can identify a set of dynamically-expressed cell cycle-specific features to construct
308 an approximate cellular trajectory graph for feature selection. Overall, we found that DELVE successfully identified
309 dynamically-expressed seed features ($p = 13$) that are known to be associated with cell cycle proliferation (e.g.
310 increase in DNA content and area of the nucleus) and captured key mechanisms previously shown to drive cell cycle
311 progression (Fig. 4a right), including molecular events that regulate the G1/S and G2/M transitions. For example,
312 the G1/S transition is governed by the phosphorylation of RB by cyclin:CDK complexes (e.g. cyclinA/CDK2 and
313 cyclinE/CDK2), which control the expression of E2F transcription factors that regulate S phase genes [93]. We
314 also observed an increase in expression of Skp2, which reduces p27-mediated inhibition of E2F1 target genes
315 [94, 95]. In addition, our approach identified S phase events that are known to be associated with DNA replication,
316 including an accumulation of PCNA foci at sites of active replication [96] and a DNA damage marker, pH2AX,
317 which becomes phosphorylated in response to double-stranded DNA breaks in areas of stalled replication [97, 98].
318 Lastly, we observed an increase in expression of cyclin B localized to different regions of the cell, which is a
319 primary regulator of G2/M transition alongside CDK1 [99, 100]. Of note, phosphorylation of RB also controls cell
320 cycle re-entry and is an important biomarker that is often used for distinguishing proliferating from arrested cells
321 [101, 102]. Furthermore, by ordering the average pairwise change in expression of features across ground truth
322 phase annotations, we observed that DELVE dynamically-expressed seed features exhibited non-random patterns
323 of expression variation that gradually increased throughout the canonical phases of the cell cycle (Fig. 4a), and
324 were amongst the top ranked features that were biologically predictive of cell cycle phase and age measurements
325 using a random forest classification and regression framework, respectively (See *Random forest*, Figure 4a right,
326 Supplementary Fig. 13). Collectively, these results illustrate that the dynamic feature module identified by DELVE
327 represents a minimum cell cycle feature set (Fig. 4b dynamic seed) that precisely distinguishes individual cells
328 according to their cell cycle progression status and can be used to construct an approximate cellular trajectory for
329 ranking feature importance.

330 We then comprehensively evaluated feature selection methods on their ability to retain imaging-derived features
331 that define cell cycle phases and resolve proliferation and arrest cell cycle trajectories. We reasoned that cells
332 in similar stages of the cell cycle (as defined by the cell cycle reporter) should have similar cell cycle signatures
333 (4i features) and should be located near one another in a low dimensional projection. Fig. 4b shows the PHATE
334 embeddings from each feature selection strategy. Using the DELVE feature set, we obtained a continuous PHATE
335 trajectory structure that successfully captured the smooth progression of cells through the canonical phases of
336 the cell cycle, where cells were tightly grouped together according to ground truth cell cycle phase annotations
337 (Fig. 4b). Moreover, we observed that the two DELVE approaches (i.e. DELVE and dynamic seed), in addition to
338 hotspot and HVG selection, produced qualitatively similar denoised lower-dimensional visualizations comparable
339 to the supervised random forest approach that was trained on ground truth cell cycle phase annotations. In contrast,
340 similarity-based approaches such as Laplacian score and neighborhood variance, which define a cell similarity
341 graph according to all features, showed more diffuse presentations of cell states. Variance-based (max variance) or
342 subspace-learning approaches (SCMER, MCFS) produced qualitatively similar embeddings to that produced using
343 all features.

344 To quantitatively assess if selected features from a feature selection strategy were biologically predictive of cell
345 cycle phases, we performed three complementary analyses. We first focused on the task of cell state classification,
346 where our goal was to learn the ground truth cell cycle phase annotations from the selected feature set. To do
347 so, we trained a support vector machine (SVM) classifier and compared the accuracy of predictions to their

348 ground truth phase annotations (See *Support Vector Machine*). We performed nested 10-fold cross validation to
349 obtain a distribution of predictions for each method. Overall, we found that DELVE achieved the highest median
350 classification accuracy (DELVE = 0.960) obtaining a similar performance to the random forest classifier trained
351 on cell cycle phase annotations (random forest = 0.957), and outperforming existing unsupervised approaches
352 (e.g. hotspot = 0.935, max variance = 0.902, HVG = 0.889, MCFS = 0.870, SCMER = 0.797, LS = 0.770),
353 as well as all features (0.946), suggesting that selected features with DELVE were more biologically predictive
354 of cell cycle phases (Fig. 4d). We next aimed to assess how well a feature selection method could identify and
355 rank cell cycle phase-specific features according to their representative power. To test this, we trained a random
356 forest classifier on the ground truth phase annotations using nested 10-fold cross validation (See *Random forest*).
357 We then compared the average ranked feature importance scores from the random forest to the selected features
358 from a feature selection strategy using the precision@k metric. Strikingly, we found that DELVE achieved the
359 highest median precision@k score (DELVE $p@30 = 0.800$) and appropriately ranked features according to their
360 discriminative power of cell cycle phases despite being a completely unsupervised approach (Fig. 4d). This was
361 followed by hotspot with a precision@k score of (hotspot $p@30 = 0.633$) and highly variable gene selection
362 (HVG $p@30 = 0.500$). In contrast, the Laplacian Score and max variance obtained low precision scores ($p@30 =$
363 0.367 and 0.333 respectively), whereas neighborhood variance and subspace-learning feature selection methods
364 MCFS and SCMER were unable to identify cell cycle phase-specific features from noisy 4i data and exhibited
365 precision scores near random ($p@30 = 0.267$, 0.267, and 0.233, respectively). Lastly, we assessed if selected
366 image-derived features could be used for downstream analysis tasks like unsupervised cell population discovery.
367 To do so, we clustered cells using the KMeans++ algorithm [103] on the selected feature set and compared the
368 predicted labels to the ground truth annotations using a normalized mutual information (NMI) score over 25 random
369 initializations (See *Unsupervised clustering*). We found that hotspot, DELVE, and dynamic seed features were
370 better able to cluster cells according to cell cycle phases and achieved considerably higher median NMI scores
371 (0.615, 0.599, 0.543, respectively), as compared to retaining all features (0.155) (Fig. 4d). Moreover, we found
372 that clustering performance was similar to that of the random forest trained on cell cycle phase annotations (0.626).
373 In contrast, variance-based approaches achieved moderate NMI clustering scores (HVG: 0.421, max variance:
374 0.361) and alternative similarity-based and subspace learning approaches obtained low median NMI scores (~ 0.2)
375 and were unable to cluster cells into biologically-cohesive cell populations. Of note, many trajectory inference
376 methods use clusters when fitting trajectory models [24, 104, 105, 23], thus accurate cell-to-cluster assignments
377 following feature selection is crucial for both cell type annotation and discovery, as well as for accurate downstream
378 trajectory analysis interpretation. Collectively, these results highlight that feature selection with DELVE identifies
379 imaging-derived features from noisy protein immunofluorescence imaging data that are more biologically predictive
380 of cell cycle phases.

381 We then focused on the much harder task of predicting an individual cell's progression through the cell cycle. A
382 central challenge in trajectory inference is the destructive nature of single-cell technologies, where only a static
383 snapshot of cell states is profiled. To move toward a quantitative evaluation of cell cycle trajectory reconstruction
384 following feature selection, we leveraged the ground truth age measurements determined from time-lapse imaging
385 of the RPE-PCNA reporter cell line. We first evaluated whether selected features could be used to accurately
386 predict cell cycle age by training a support vector machine (SVM) regression framework using nested ten fold
387 cross validation (See *Support Vector Machine*). Method performance was subsequently assessed by computing
388 the mean squared error (MSE) between the predictions and the ground truth age measurements. Overall, we
389 found that DELVE achieved the lowest median MSE (3.261), outperforming both supervised (random forest
390 = 3.296) and unsupervised approaches (e.g. second best performer hotspot = 3.654) suggesting that selected
391 features more accurately estimate the time following mitosis (Fig. 4c). Crucially, this highlights DELVE's ability
392 to learn new biologically-relevant features that might be missed when performing a supervised or unsupervised
393 approach. Lastly, we assessed whether selected imaging features could be used to accurately infer proliferation
394 and arrest cell cycle trajectories using common trajectory inference approaches (Fig. 4d). Briefly, we constructed
395 predicted cell cycle trajectories using the diffusion pseudotime algorithm [83] under each feature selection strategy
396 (See *Trajectory inference and analysis*). Cells were separated into proliferation or arrest lineages according to
397 their average expression of pRB, and cellular progression was estimated using ten random root cells that had the
398 youngest age. Feature selection method performance on trajectory inference was then quantitatively assessed by
399 computing the Spearman rank correlation between estimated pseudotime and the ground truth age measurements.
400 We found that DELVE achieved the highest median correlation of estimated pseudotime to the ground truth age
401 measurements (proliferation: 0.656, arrest: 0.405) as compared to alternative methods (second best performer
402 hotspot; proliferation: 0.632, arrest: 0.333) or all features (proliferation: 0.330, arrest: 0.135), indicating that our
403 approach was better able to resolve both proliferation and cell cycle arrest trajectories where other approaches
404 failed (Fig. 4d). Of note, DELVE was robust to the choice in hyperparameters and obtained reproducible results
405 across a range of hyperparameter choices (See Supplementary Fig. 14).

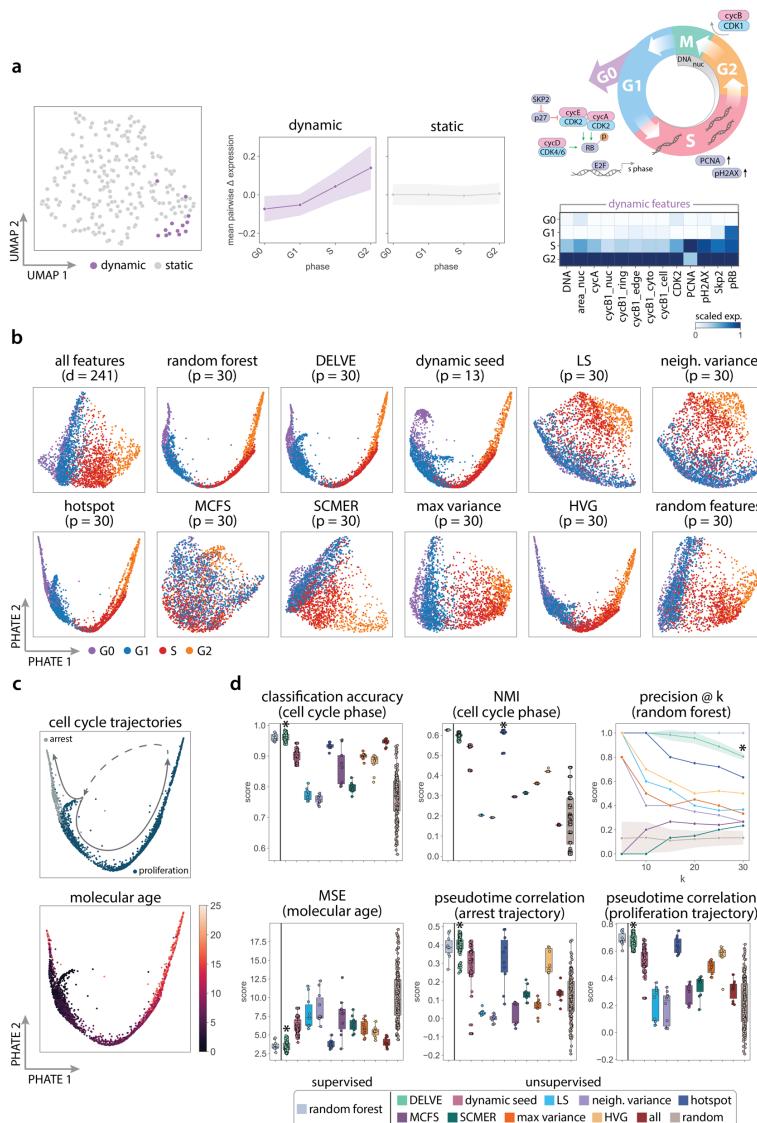


Figure 4: DELVE recovered signatures of proliferation and arrest in noisy protein immunofluorescence imaging data. (a) DELVE identified one dynamic module consisting of 13 seed features that represented a minimum cell cycle. (a left) UMAP visualization of image-derived features where each point indicates a dynamic or static feature identified by the model. (a middle) The average pairwise change in expression of features within DELVE modules ordered across ground truth cell cycle phase annotations. (a right top) Simplified signaling schematic of the cell cycle highlighting the role of DELVE dynamic seed features within cell cycle progression. (a right bottom) Heatmap of the standardized average expression of dynamic seed features across cell cycle phases. (b) Feature selection was performed to select the top ($p = 30$) ranked features from the original ($d = 241$) feature set according to a feature selection strategy. PHATE visualizations illustrating the overall quality of low-dimensional retinal pigmented epithelial cell cycle trajectories following feature selection. Cells were labeled according to ground truth cell cycle phase annotations from time-lapse imaging. Each panel represents a different feature selection strategy. (c) PHATE visualizations following DELVE feature selection, where cells were labeled according to cell cycle trajectory (top) or ground truth age measurements (bottom). (d) Performance of twelve feature selection methods on representing the cell cycle according to several metrics, including classification accuracy between predicted and ground truth phase annotations using a support vector machine classifier on selected features, normalized mutual information (NMI) between predicted and ground truth phase labels to indicate clustering performance, precision of phase-specific features determined by a random forest classifier trained on cell cycle phase annotations, mean-squared error between predicted and ground truth molecular age measurements using a support vector machine regression framework on selected features, and the correlation between estimated pseudotime to the ground truth molecular age measurements following trajectory inference on selected features. Error bands represent the standard deviation. * indicates the method with the highest median score. DELVE achieved the highest classification accuracy, highest p@k score, and high NMI clustering score indicating robust prediction of cell cycle phase. Furthermore, DELVE achieved the lowest mean squared error and highest correlation between estimated pseudotime and ground truth age measurements indicating robust prediction of cell cycle transitions.

406 As a secondary validation, we applied DELVE to 9 pancreatic adenocarcinoma (PDAC) cell lines (e.g. BxPC3,
407 CFPAC, MiaPaCa, HPAC, Pa01C, Pa02C, PANC1, UM53) profiled with 4i (See *PDAC analysis*) and performed a
408 similar evaluation of cell cycle phase and phase transition preservation (See Supplementary Figs. 16 - 24). Across
409 all cell lines and metrics, we found that DELVE approaches and hotspot considerably outperformed alternative
410 methods on recovering the cell cycle from noisy 4i data and often achieved the highest classification accuracy
411 scores, clustering scores, and the highest correlation of cellular progression along proliferative and arrested cell
412 cycle trajectories (See Supplementary Fig. 15). Notably, DELVE was particularly useful in resolving cell cycle
413 trajectories from the PDAC cell lines that had numerous imaging measurements with low signal-to-noise ratio (e.g.
414 CFPAC, MiaPaCa, PANC1, and UM53), whereas the alternative strategies were unable to resolve cell cycle phases
415 and achieved scores near random (See Supplementary Figs. 17, 19, 23, 24).

416 **Identifying molecular drivers of CD8+ T cell effector and memory formation**

417 To demonstrate the utility of our approach in a complex differentiation setting consisting of heterogeneous cell
418 subtypes and shared and distinct molecular pathways, we applied DELVE to a single-cell RNA sequencing time
419 series dataset consisting of 29,893 mouse splenic CD8+ T cells responding to acute viral infection [106]. Here,
420 CD8+ T cells were profiled over 12 time points following infection with the Armstrong strain of lymphocytic
421 choriomeningitis virus (LCMV): Naive, d3-, d4-, d5-, d6-, d7-, d10-, d14-, d21-, d32-, d60-, and d90- post-infection
422 (See *CD8 T cell differentiation analysis*). During an immune response to acute viral infection, naive CD8+ T
423 cells undergo a rapid activation and proliferation phase, giving rise to effector cells that can serve in a cytotoxic
424 role to mediate immediate host defense, followed by a contraction phase giving rise to self-renewing memory
425 cells that provide long-lasting protection and are maintained by antigen-dependent homeostatic proliferation
426 [107, 108, 109]. Despite numerous studies detailing the molecular mechanisms of CD8+ T cell effector and
427 memory fate specification, the molecular mechanisms driving activation, fate commitment, or T cell dysfunction
428 continue to remain unclear due to the complex intra- and inter-temporal heterogeneity of the CD8+ T cell response
429 during infection. Therefore, we applied DELVE to the CD8+ T cell dataset to resolve the differentiation trajectory
430 and investigate transcriptional changes that are involved in effector and memory formation during acute viral
431 infection with LCMV.

432 Following unsupervised seed selection, we found that DELVE successfully identified three gene modules consti-
433 tuting core regulatory complexes involved in CD8+ T cell viral response and had dynamic expression patterns
434 that varied across experimental time following viral infection (Fig. 5a-c). Namely, dynamic module 0 contained
435 genes involved in early activation and interferon response (e.g. Ly6a, Bst2, Ifi27l2a) [110, 111], and proliferation
436 (e.g. Cenpa, Cenpf, Ccnb2, Ube2c, Top2a, Tubb4b, Birc5, Cks2, Cks1b, Nusap1, Hmgb2, Rrm2, H2afx, Pclaf,
437 Stmn1) [112]. Dynamic module 1 contained genes involved in effector formation, including interferon- γ cytotoxic
438 molecules, such as perforin/granzyme pathway (e.g. Gzma, Gzmk), integrins (e.g. Itgb1), killer cell lectin-like
439 receptor family (e.g. Klr1, Klr1d, Klrk1, Klr1c), chemokine receptors (e.g. Cxcr3, Cxcr6, Ccr2), and a canonical
440 transcription factor involved in terminal effector formation (e.g. Id2) [113, 114, 115, 116]. Lastly, dynamic module
441 2 contained genes involved in long-term memory formation (e.g. Sell, Bcl2, Il7r, Ltb) [117, 118, 119, 120]. To
442 quantitatively examine if genes within a dynamic module were meaningfully associated with one another, or had
443 experimental evidence of co-regulation, we constructed gene association networks using experimentally-derived
444 association scores from the STRING database [121]. Here, a permutation test was performed to assess the statistical
445 significance of the observed experimental association amongst genes within a DELVE module as compared to
446 random gene assignment (See Protein-protein interaction networks). Notably, across all three dynamic modules,
447 DELVE identified groups of genes that had statistically significant experimental evidence of co-regulation (p -value
448 = 0.001), where DELVE networks had a larger average degree of experimentally-derived edges than the null
449 distribution (Fig. 5b: dynamic modules). Degree centrality is a simple measurement of the number of edges (e.g.
450 experimentally derived associations between genes) connected to a node (e.g. gene); therefore, in this context,
451 networks with a high average degree may contain complexes of genes that are essential for regulating a biological
452 process. In contrast, genes identified by DELVE that exhibited random or noisy patterns of expression variation
453 (static module) had little to no evidence of co-regulation (p -value = 1.0) and achieved a much lower average degree
454 than networks defined by random gene assignment (Fig. 5b).

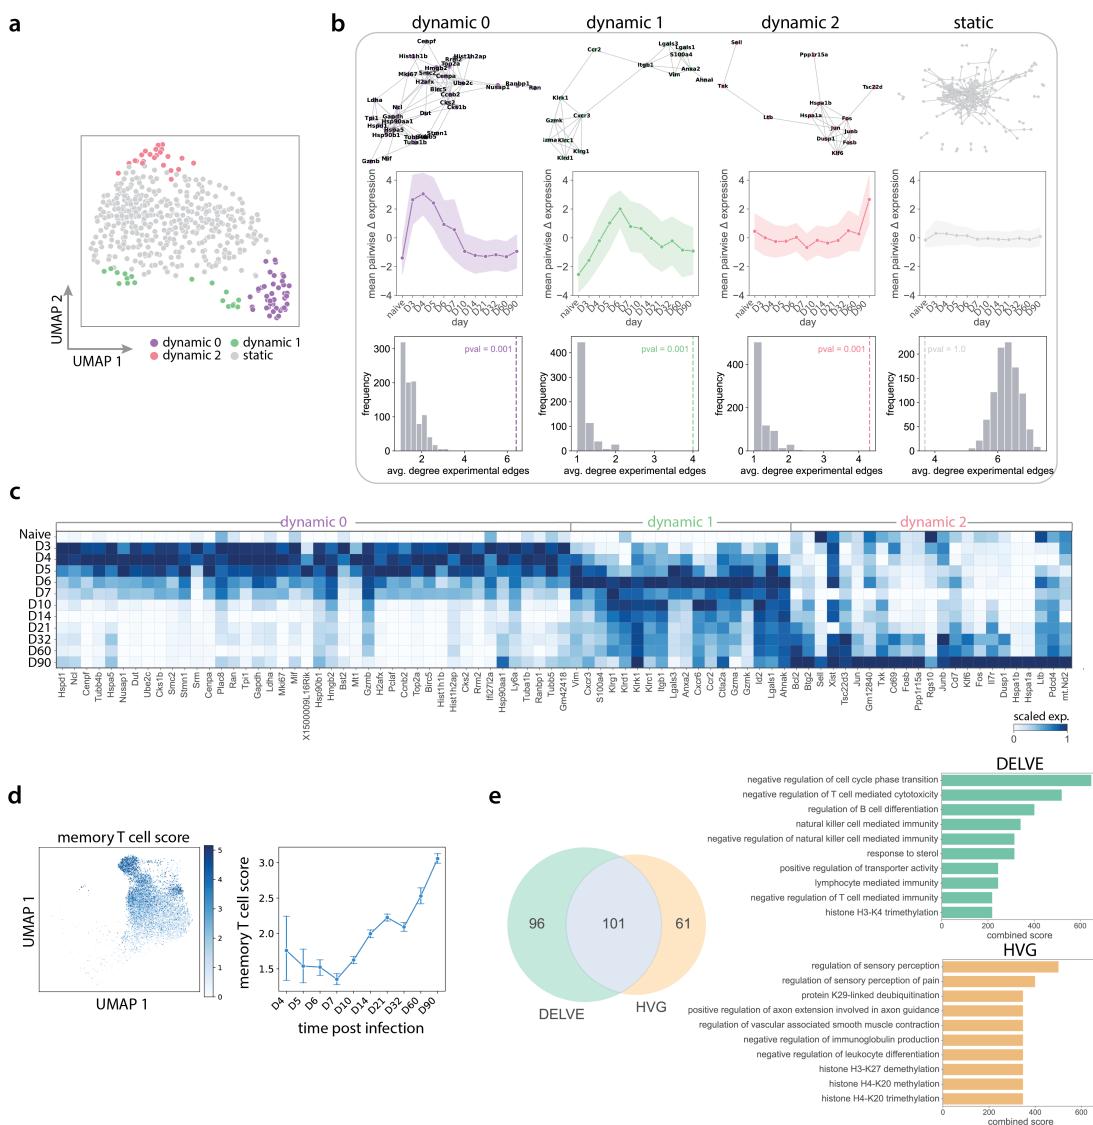


Figure 5: DELVE identified molecular drivers of CD8+ T cell effector and memory formation. (a) DELVE identified three dynamic modules representing cell cycle and early activation (dynamic 0), effector formation and cytokine signaling (dynamic 1), and long-term memory formation (dynamic 2) during CD8+ T cell differentiation response to viral infection with lymphocytic choriomeningitis virus (LCMV). UMAP visualization of ($d = 500$) genes where each point indicates a dynamic or static gene identified by the model. (b) A permutation test was performed using experimentally-derived association scores from the STRING protein-protein interaction database [121] to assess whether genes within DELVE dynamic modules had experimental evidence of co-regulation as compared to random assignment. (b top) STRING association networks, where nodes represent genes from a DELVE module and edges represent experimental evidence of association. (b middle) Average pairwise change in expression amongst genes within a module ordered by time following infection. (b bottom) Histograms showing the distribution of the average degree of experimentally-derived edges of gene networks from $R = 1000$ random permutations. The dotted line indicates the observed average degree from genes within a DELVE module. p -values were computed using a one-sided permutation test. (c) Heatmap visualization of the standardized average expression of dynamically-expressed genes identified by DELVE ordered across time following infection. (d) Trajectory inference was performed along the memory lineage using the diffusion pseudotime algorithm [83], where cell similarity was determined by DELVE selected genes or highly variable genes (HVG). UMAP visualization of memory T cell scores according to the average expression of known memory markers (Bcl-2, Sell, IL7r). Line plot indicates the onset of expression and cellular commitment to the memory lineage following infection. (e) Genes from the full dataset were regressed along estimated pseudotime using a generalized additive model to determine lineage-specific significant genes. The venn diagram illustrates the quantification and overlap of memory lineage-specific genes across feature selection strategies. The barplots show the top ten gene ontology terms associated with the temporally-expressed gene lists specific to each feature selection strategy.

455 Next, we examined if DELVE could be used to improve the identification of genes associated with long-term
456 CD8+ T cell memory formation following trajectory inference. To do so, we first prioritized cells along the
457 memory lineage by computing a memory T cell score according to the average expression of three known memory
458 markers (Bcl2, Sell, and Il17r) (See *CD8 T cell differentiation analysis*, Fig. 5d). We then reconstructed the memory
459 CD8+ T cell differentiation trajectory from middle-late stage cellular commitment using the diffusion pseudotime
460 algorithm [83] on the top 250 ranked genes following DELVE feature selection or highly variable gene selection.
461 Therefore, in this context, cell ordering was reflective of the differences in cell state along the memory lineage
462 according to (1) dynamically-expressed genes that had experimental evidence of co-regulation or (2) variance-based
463 selection. Lastly, we performed a regression analysis for each gene ($d = 500$) in the original dataset along
464 estimated pseudotime using a negative binomial generalized additive model (GAM). Genes were considered to be
465 differentially expressed along the memory lineage if they had a q -value < 0.05 following Benjamini-Hochberg
466 false discovery rate correction [122] (See Trajectory inference and analysis). Overall, we found that ordering
467 cells according to similarities in selected gene expression using DELVE was more reflective of long-term memory
468 formation and achieved an increased recovery of memory lineage-specific genes, as directly compared to the
469 standard approach of highly variable gene selection (Fig. 5e). To determine the biological relevance of these
470 memory lineage-specific genes, we performed gene set enrichment analysis on the temporally-expressed genes
471 specific to each feature selection strategy using EnrichR [123]. Here, DELVE obtained higher combined scores
472 and identified more terms involved in immune regulation and memory CD8+ T cell formation, including negative
473 regulation of cell cycle phase transition, negative regulation of T cell mediated cytotoxicity, lymphocyte mediated
474 immunity, and negative regulation of T cell mediated immunity (Fig. 5e).

475 Discussion

476 Computational trajectory inference methods have transformed our ability to study the continuum of cellular states
477 associated with dynamic phenotypes; however, current approaches for reconstructing cellular trajectories can be
478 hindered by biological or technical noise inherent to single-cell data [42, 43]. To mitigate the effect of unwanted
479 sources of variation confounding trajectory inference, we designed a bottom-up unsupervised feature selection
480 method that ranks and selects features that best approximate cell state transitions from dynamic feature modules that
481 constitute core regulatory complexes. The key innovation of this work is the ability to parse temporally co-expressed
482 features from noisy information-poor features prior to performing feature selection; in doing so, DELVE constructs
483 cell similarity graphs that are more reflective of cell state progression for ranking feature importance.

484 In this study, we benchmarked twelve feature selection methods [46, 52, 25, 53, 55, 54, 50] on their ability to
485 identify biologically relevant features for trajectory analysis from simulated single-cell RNA sequencing data where
486 the ground truth was known. We found that DELVE achieved the highest recovery of differentially expressed
487 genes within a cell type or along a cellular lineage, highest cell type classification accuracy, and most accurately
488 estimated individual cell progression across a variety of trajectory topologies and biological or technical challenges.
489 Furthermore, through a series of qualitative and quantitative comparisons, we illustrated how noise (e.g. stochasticity,
490 sparsity, low library size) and information-poor features create spurious similarities amongst cells and considerably
491 impact the performance of existing unsupervised similarity-based or subspace learning-based feature selection
492 methods on identifying biologically-relevant features.

493 Next, we applied DELVE to a variety of biological contexts and demonstrated improved recovery of cellular
494 trajectories over existing unsupervised feature selection strategies. In the context of studying the cell cycle from
495 protein imaging-derived features [92], we illustrated how DELVE identified molecular features that were strongly
496 associated with cell cycle progression and were more biologically predictive of cell cycle phase and age, as
497 compared to the alternative unsupervised feature selection methods. Importantly, DELVE often achieved similar or
498 better performance to the supervised Random Forest classification approach without the need for training on ground
499 truth cell cycle labels. Lastly, in the context of studying heterogeneous CD8+ T cell response to viral infection from
500 single-cell RNA sequencing data [106], we showed how DELVE identified gene complexes that had experimental
501 evidence of co-regulation and were strongly associated with CD8+ T cell differentiation. Furthermore, we showed
502 how performing feature selection with DELVE prior to performing trajectory inference improved the identification
503 and resolution of gene programs associated with long-term memory formation that would have been missed by the
504 standard unsupervised feature selection approach.

505 This study highlights how DELVE can be used to improve inference of cellular trajectories in the context of noisy
506 single-cell omics data; however, it is important to note that feature selection can greatly bias the interpretation of the
507 underlying cellular trajectory [42], thus careful consideration should be made when performing feature selection
508 for trajectory analysis. Furthermore, we provided an unsupervised framework for ranking features according to

509 their association with temporally co-expressed genes, although we note that DELVE can be improved by using a
 510 set of previously established regulators (See Step 2: feature ranking). Future work could focus on extending this
 511 framework for applications such as (1) deconvolving cellular trajectories using biological system-specific seed
 512 graphs or (2) studying complex biological systems such as organoid models or spatial microenvironments.

513 Methods

514 DELVE

515 DELVE identifies a subset of dynamically-changing features that preserve the local structure of the underlying
 516 cellular trajectory. In this section, we will (1) describe computational methods for the identification and ranking
 517 of features that have non-random patterns of dynamic variation, (2) explain DELVE's relation to previous work,
 518 and (3) provide context for the mathematical foundations behind discarding information-poor features prior to
 519 performing trajectory inference.

520 Problem formulation

521 Let $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$ denote a single-cell dataset, where $\mathbf{x}_i \in \mathbb{R}^d$ represents the vector of d measured features (e.g.
 522 genes or proteins) measured in cell i . We assume that the data have an inherent trajectory structure, or biologically-
 523 meaningful ordering, that can be directly inferred by a limited subset of p features where $p \ll d$. Therefore, our
 524 goal is to identify this limited set of p features from the original high-dimensional feature set that best approximate
 525 the transitions of cells through each stage of the underlying dynamic process.

526 Step 1: dynamic seed selection

527 **Graph construction:** Our approach DELVE extends previous similarity-based [52, 25, 53] or subspace-learning
 528 [55] feature selection methods by computing the dependence of each gene on the underlying cellular trajectory.
 529 In step 1, DELVE models cell states using a weighted k -nearest neighbor affinity graph of cells ($k = 10$), where
 530 nodes represent cells and edges describe the transcriptomic or proteomic similarity amongst cells according to the d
 531 profiled features encoded in \mathbf{X} . More specifically, let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ denote a between-cell affinity graph, where \mathcal{V}
 532 represents the cells and the edges, \mathcal{E} , are weighted according to a Gaussian kernel as,

$$w_{ij} = \begin{cases} \exp\left(-\frac{\|\mathbf{x}_{v_i} - \mathbf{x}_{v_j}\|^2}{2\sigma_i^2}\right), & \text{if } v_j \in \mathcal{N}_i \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

533 Here, \mathbf{W} is a $n \times n$ between-cell similarity matrix, where cells v_i and v_j are connected with an edge with edge
 534 weight w_{ij} if the cell v_j is within the set of v_i 's neighbors, as denoted by notation \mathcal{N}_i . Moreover, σ_i , specific for
 535 a particular cell i , represents the Gaussian kernel bandwidth parameter that controls the decay of cell similarity
 536 edge weights. We chose a bandwidth parameter as the distance to the 3rd nearest neighbor as this has been shown
 537 previously in refs. [53] and [124] to provide reasonable decay in similarity weights.

538 **Identification of feature modules:** To identify groups of features with similar co-expression variation, DELVE
 539 clusters features according to changes in expression across prototypical cell neighborhoods. First, cellular neighbor-
 540 hoods are defined according to the average expression of each set of k nearest neighbors (\mathcal{N}_i) as, $\mathbf{Z} = \{\mathbf{z}_i \in \mathbb{R}^d\}_{i=1}^n$,
 541 where each $\mathbf{z}_i = \frac{1}{k} \sum_{\mathcal{N}_i} \mathbf{x}_i$ represents the center of the k nearest neighbors for cell i across all measured features.
 542 Next, DELVE leverages Kernel Herding sketching [66] to effectively sample m representative cell neighborhoods,
 543 or rows, from the per-cell neighbor averaged feature matrix, \mathbf{Z} , as $\tilde{\mathbf{Z}} = \{\tilde{\mathbf{z}}_i \in \mathbb{R}^d\}_{i=1}^m$. This sampling approach
 544 ensures that cellular neighborhoods are more reflective of the original distribution of cell states, while removing
 545 redundant cell states to aid in the scalability of estimating expression dynamics. DELVE then computes the average
 546 pairwise change in expression of features across representative cellular neighborhoods, Δ , as,

$$\Delta = \frac{1}{m-1} \sum_{i=1}^m (\tilde{\mathbf{Z}} - \mathbf{j}_m \tilde{\mathbf{z}}_i^T), \quad (2)$$

547 where \mathbf{j}_m is a column vector of ones with length m , such that $\mathbf{j}_m \in \mathbb{R}^m$. Lastly, features are clustered according to
 548 the transpose of their average pairwise change in expression across the representative cellular neighborhoods, Δ^T ,

549 using the KMeans++ algorithm [103]. In this context, each DELVE module contains a set of features with similar
 550 local changes in co-variation across cell states along the cellular trajectory.

551 **Dynamic expression variation permutation testing:** To assess the significance of dynamic expression variation
 552 across grouped features within a DELVE module, we perform a permutation test as follows. Let $\bar{S}_c^2(P_c)$ denote the
 553 average sample variance of the average pairwise change in expression across m cell neighborhoods for the set of p
 554 features (a set of features denoted as P_c) within a DELVE cluster c as,

$$\bar{S}_c^2(P_c) = \frac{1}{|P_c|} \sum_{p=1}^{|P_c|} \sum_{i=1}^m \frac{(\Delta_{i,p} - \bar{\Delta}_p)^2}{m-1}. \quad (3)$$

555 Moreover, let R_q denote a set of randomly selected features sampled without replacement from the full feature
 556 space d , such that $|P_c| = |R_q|$, and $\tilde{S}_c^2(R_q)$ denote the average sample variance of randomly selected feature sets
 557 averaged across t random permutations as,

$$\tilde{S}_c^2(R_q) = \frac{1}{t} \sum_{q=1}^t \bar{S}_c^2(R_q). \quad (4)$$

558 Here, DELVE considers a module of features as being *dynamically-expressed* if the average sample variance of the
 559 change in expression of the set of features within a DELVE cluster (or specifically feature set P_c), is greater than
 560 random assignment, R_q , across randomly permuted trials as,

$$\bar{S}_c(P_c) > \tilde{S}_c(R_q). \quad (5)$$

561 In doing so, this approach is able to identify and exclude modules of features that have static, random, or noisy
 562 patterns of expression variation, while retaining dynamically expressed features for ranking feature importance. Of
 563 note, given that KMeans++ clustering is used to initially assign features to a group, feature-to-cluster assignments
 564 can tend to vary due to algorithm stochasticity. Therefore, to reduce the variability and find a core set of features
 565 that are consistently dynamically-expressed, this process is repeated across ten random clustering initializations
 566 and the set of dynamically-expressed features are defined as the intersection across runs.

567 Step 2: feature ranking

568 Following dynamic seed selection, in step two, DELVE ranks features according to their association with the
 569 underlying cellular trajectory graph. First, DELVE approximates the underlying cellular trajectory by constructing
 570 a between-cell affinity graph, where the nodes represent the cells and edges are now re-weighted according to a
 571 Gaussian kernel between all cells based on the core subset of dynamically expressed regulators from step 1, such
 572 that $\tilde{\mathbf{X}} = \{\tilde{\mathbf{x}}_i \in \mathbb{R}^p\}$ where $p \ll d$ as,

$$\tilde{w}_{ij} = \begin{cases} \exp\left(-\frac{\|\tilde{\mathbf{x}}_{v_i} - \tilde{\mathbf{x}}_{v_j}\|^2}{2\sigma_i^2}\right), & \text{if } v_j \in \mathcal{N}_i \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

573 Here, $\tilde{\mathbf{W}}$ is a $n \times n$ between-cell similarity matrix, where cells v_i and v_j are connected with an edge with edge
 574 weight \tilde{w}_{ij} if the cell is within the set of v_i 's neighbors, denoted as \mathcal{N}_i . Moreover, as previously mentioned,
 575 σ_i represents the Gaussian kernel bandwidth parameter for a particular cell i as the distance to the 3rd nearest
 576 neighbor.

577 Features are then ranked according to their association with the underlying cellular trajectory graph using graph
 578 signal processing techniques [52, 67, 68]. A graph signal f is any function that has a real defined value on all of the
 579 nodes, such that $f \in \mathbb{R}^n$ and f_i gives the signal at the i th node. Intuitively, we can consider all features as graph
 580 signals and rank them according to their variation in expression along the approximate cell trajectory graph to see if
 581 they should be included or excluded from downstream analysis. Let \mathbf{L} denote the unnormalized graph Laplacian,
 582 with $\mathbf{L} = \mathbf{D} - \tilde{\mathbf{W}}$, where \mathbf{D} is a diagonal degree matrix with each element as $d_{ii} = \sum_j \tilde{w}_{ij}$. The local variation
 583 in the expression of feature signal f can then be defined as the weighted sum of differences in signals around a
 584 particular cell i as,

$$(\mathbf{L}\mathbf{f})(i) = \sum_j \tilde{\mathbf{w}}_{ij} (\mathbf{f}(i) - \mathbf{f}(j)). \quad (7)$$

585 This metric effectively measures the similarity in expression of a particular node's graph signal, denoted by the
 586 feature vector, \mathbf{f} , around its k nearest neighbors. By summing the local variation in expression across all neighbors
 587 along the cellular trajectory, we can define the total variation in expression of feature graph signal \mathbf{f} as,

$$\mathbf{f}^T \mathbf{L} \mathbf{f} = \sum_{ij} \tilde{\mathbf{w}}_{ij} (\mathbf{f}(i) - \mathbf{f}(j))^2. \quad (8)$$

588 Otherwise known as the Laplacian quadratic form, in this context, the total variation represents the global smoothness
 589 of the particular graph signal encoded in \mathbf{f} (e.g. expression of a particular gene or protein) along the approximate
 590 cellular trajectory graph. Intuitively, DELVE ultimately retains features that have a low total variation in expression,
 591 or have similar expression values amongst similar cells along the approximate cellular trajectory graph. In contrast,
 592 DELVE excludes features that have a high total variation in expression, or those which have expression values that
 593 are rapidly oscillating amongst neighboring cells, as these features are likely noisy or not involved in the underlying
 594 dynamic process that was initially seeded.

595 In this work, we ranked features according to their association with the cell-to-cell affinity graph defined by
 596 dynamically expressed features from DELVE dynamic modules using the Laplacian score [52]. This measure takes
 597 into account both the total variation in expression, as well as the overall global variance. For each of the original d
 598 measured features, or *graph signals* encoded in \mathbf{f} with $\mathbf{f} \in \mathbb{R}^n$, the Laplacian score L_f is computed as,

$$L_f = \frac{\tilde{\mathbf{f}}^T \mathbf{L} \tilde{\mathbf{f}}}{\tilde{\mathbf{f}}^T \mathbf{D} \tilde{\mathbf{f}}}. \quad (9)$$

599 Here, \mathbf{L} represents the unnormalized graph Laplacian, such that $\mathbf{L} = \mathbf{D} - \tilde{\mathbf{W}}$, \mathbf{D} is a diagonal degree matrix with
 600 the i th element of the diagonal d_{ii} as $d_{ii} = \sum_j \tilde{\mathbf{w}}_{ij}$, $\tilde{\mathbf{f}}$ represents the mean centered expression of feature \mathbf{f} as
 601 $\tilde{\mathbf{f}} = \mathbf{f} - \frac{\mathbf{f}^T \mathbf{D} \mathbf{1}}{\mathbf{1}^T \mathbf{D} \mathbf{1}}$, and $\mathbf{1} = [1, \dots, 1]^T$. By sorting features in ascending order according to their Laplacian score,
 602 DELVE effectively ranks features that best preserve the local trajectory structure (e.g. an ideal numerator has
 603 a small local variation in expression along neighboring cells), as well as best preserve cell types (e.g. an ideal
 604 denominator has large variance in expression for discriminative power).

605 Benchmarked feature selection methods

606 In this section, we describe the twelve feature selection methods evaluated for representing biological trajectories.
 607 For more details on implementation and hyperparameters, see Supplementary Table 1.

608 **Random forest** To quantitatively compare feature selection approaches on preserving biologically relevant genes
 609 or proteins, we aimed to implement an approach that would leverage ground truth cell type labels to determine
 610 feature importance. Random forest classification [46] is a supervised ensemble learning algorithm that uses an
 611 ensemble of decision trees to partition the feature space such that all of the samples (cells) with the same class
 612 (cell type labels) are grouped together. Each decision or split of a tree was chosen by minimizing the Gini impurity
 613 score as,

$$G(m) = \sum_{i=1}^C p_{mi}(1 - p_{mi}). \quad (10)$$

614 Here, p_{mi} is the proportion of cells that belong to class i for a feature node m , and C is the total number of classes
 615 (e.g. cell types). We performed random forest classification using the sklearn v0.23.2 package in python. Nested
 616 10-fold cross-validation was performed using stratified random sampling to assign cells to either a training or
 617 test set. The number of trees was tuned over a grid search within each fold prior to training the model. Feature
 618 importance scores were subsequently determined by the average Gini importance across folds.

Max variance Max variance is an unsupervised feature selection approach that uses sample variance as a criterion for retaining discriminative features, where \tilde{S}_f^2 represents the sample variance for feature $\mathbf{f} \in \mathbb{R}^n$ as,

$$S_f^2 = \frac{1}{n-1} \sum_{i=1}^n (f_i - \bar{f})^2, \quad (11)$$

619 where f_i indicates the expression value of feature \mathbf{f} in cell i . We performed max variance feature selection by
620 sorting features in descending order according to their variance score and selecting the top p maximally varying
621 features.

622 **Neighborhood variance** Neighborhood variance [25] is an unsupervised feature selection approach that uses
623 a local neighborhood variance metric to select gradually-changing features for building biological trajectories.
624 Namely, the neighborhood variance metric \tilde{S}_f^2 quantifies how much feature f varies across neighboring cells
625 as,

$$\tilde{S}_f^2 = \frac{1}{nk_c - 1} \sum_{i=1}^n \sum_{j=1}^{k_c} (f_i - f_{\mathcal{N}_{(i,j)}})^2. \quad (12)$$

626 Here, f_i represents the expression value of feature f for cell i , $\mathcal{N}_{(i,j)}$ indicates the j nearest neighbor of cell i ,
627 and k_c is the minimum number of k -nearest neighbors required to form a fully connected graph. Features were
628 subsequently selected if they had a smaller neighborhood variance \tilde{S}_f^2 than global variance S_f^2 ,

$$\frac{S_f^2}{\tilde{S}_f^2} > 1. \quad (13)$$

629 **Highly variable genes** Highly variable gene selection [50] is an unsupervised feature selection approach that
630 selects features according to a normalized dispersion measure. First, features are binned based on their average
631 expression. Within each bin, genes are then z-score normalized to identify features that have a large variance, yet a
632 similar mean expression. We selected the top p features using the highly variable genes function in Scanpy v1.8.1
633 (flavor = Seurat, bins = 20, n_top_genes = p).

634 **Laplacian score** Laplacian score (LS) [52] is a locality-preserving unsupervised feature selection method that
635 ranks features according to (1) how well a feature's expression is consistent across neighboring cells defined by a
636 between-cell similarity graph define by all profiled features and (2) the feature's global variance. First, a weighted
637 k -nearest neighbor affinity graph of cells ($k = 10$) is constructed according to pairwise Euclidean distances between
638 cells based on all features, \mathbf{X} . More specifically, let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} represents the cells and edges, \mathcal{E} , are
639 weighted using a Gaussian as follows. Specifically, edge weights between cells i and j can be defined as,

$$w_{ij} = \begin{cases} \exp\left(-\frac{\|\mathbf{x}_{v_i} - \mathbf{x}_{v_j}\|^2}{2\sigma_i^2}\right), & \text{if } v_j \in \mathcal{N}_i \\ 0, & \text{otherwise.} \end{cases} \quad (14)$$

640 Here \mathbf{W} is a $n \times n$ between-cell similarity matrix, where cells v_i and v_j are connected with an edge with edge
641 weight w_{ij} if the cell v_j is within the set of v_i 's neighbors, \mathcal{N}_i . Moreover, as previously described, σ_i represents
642 the bandwidth parameter for cell i defined as the distance to the 3rd nearest neighbor. For each feature \mathbf{f} , where
643 $\mathbf{f} \in \mathbb{R}^n$ represents the value of the feature across all n cells, we compute the Laplacian score, L_f as,

$$L_f = \frac{\tilde{\mathbf{f}}^T \mathbf{L} \tilde{\mathbf{f}}}{\tilde{\mathbf{f}}^T \mathbf{D} \tilde{\mathbf{f}}}. \quad (15)$$

644 Here, \mathbf{L} represents the unnormalized graph Laplacian, with $\mathbf{L} = \mathbf{D} - \mathbf{W}$, \mathbf{D} is a diagonal degree matrix with the i th
645 element of the diagonal d_{ii} as $d_{ii} = \sum_j w_{ij}$, $\tilde{\mathbf{f}}$ represents the mean centered expression of feature \mathbf{f} as $\tilde{\mathbf{f}} = \mathbf{f} - \frac{\mathbf{f}^T \mathbf{D} \mathbf{1}}{\mathbf{1}^T \mathbf{D} \mathbf{1}}$,
646 and $\mathbf{1} = [1, \dots, 1]^T$. We performed feature selection by sorting features in ascending order according to their
647 Laplacian score and selecting the top p features.

648 **MCFS** Multi-cluster feature selection (MCFS) [55] is an unsupervised feature selection method that selects for
 649 features that best preserve the multi-cluster structure of data by solving an L1 regularized least squares regression
 650 problem on the spectral embedding. Similar to the Laplacian score, first k -nearest neighbor affinity graph of cells
 651 ($k = 10$) is computed to encode the similarities in feature expression between cells i and j using a Gaussian kernel
 652 as,

$$w_{ij} = \begin{cases} \exp\left(-\frac{\|\mathbf{x}_{v_i} - \mathbf{x}_{v_j}\|^2}{2\sigma_i^2}\right), & \text{if } v_j \in \mathcal{N}_i \\ 0, & \text{otherwise.} \end{cases} \quad (16)$$

653 Similar to previous formulations above, \mathbf{W} is an $n \times n$ between cell affinity matrix, where a pair of cells v_i and v_j
 654 are connected with an edge with weight w_{ij} if cell v_j is within the set of v_i 's neighbors, \mathcal{N}_i . Further, σ_i represents
 655 the kernel bandwidth parameter chosen to be the distance to the third nearest neighbor from cell i . Next, to represent
 656 the intrinsic dimensionality of the data, the spectral embedding [125] is computed through eigendecomposition of
 657 the unnormalized graph Laplacian \mathbf{L} , where $\mathbf{L} = \mathbf{D} - \mathbf{W}$ as,

$$\mathbf{Ly} = \lambda \mathbf{Dy}. \quad (17)$$

658 Here, $\mathbf{Y} = \{\mathbf{y}_k\}_{k=1}^K$ are the eigenvectors corresponding to the K smallest eigenvalues, \mathbf{W} is a symmetric affinity
 659 matrix encoding cell similarity weights, and \mathbf{D} represents a diagonal degree matrix with each element as $d_{ii} =$
 660 $\sum_j w_{ij}$. Given that eigenvectors of the graph Laplacian represent frequency harmonics [68] and low frequency
 661 eigenvectors are considered to capture the informative structure of the data, MCFS computes the importance of
 662 each feature along each intrinsic dimension \mathbf{y}_k by finding a relevant subset of features by minimizing the error
 663 using an L1 norm penalty as,

$$\min_{\mathbf{a}_k} \|\mathbf{y}_k - \mathbf{X}^T \mathbf{a}_k\|^2 \quad \text{s.t.} \quad \|\mathbf{a}_k\|_1 \leq \gamma. \quad (18)$$

664 Here, the non-zero coefficients, \mathbf{a}_k , indicate the most relevant features for distinguishing clusters from the embedding
 665 space, \mathbf{y}_k and γ controls the sparsity and ensures the least relevant coefficients are shrunk to zero. The optimization
 666 is solved using the least angles regression algorithm [71], where for every feature, the MCFS score is defined
 667 as,

$$\text{MCFS}(j) = \max_k \|a_{k,j}\|. \quad (19)$$

668 Here, j and k index feature and eigenvector, respectively. We performed multi-cluster feature selection with the
 669 number of eigenvectors K chosen to be the number of ground truth cell types present within the data, as this is the
 670 traditional convention in spectral clustering [57] and the number of nonzero coefficients was set to the number of
 671 selected features, p .

672 **SCMER** Single-cell manifold-preserving feature selection (SCMER) [54] selects a subset of p features that
 673 represent the embedding structure of the data by learning a sparse weight vector \mathbf{w} by formulating an elastic net
 674 regression problem that minimizes the KL divergence between a cell similarity matrix defined by all features and
 675 one defined by a reduced subset of features. More specifically, let \mathbf{P} denote a between-cell pairwise similarity
 676 matrix defined in UMAP [56] computed with the full data matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ and \mathbf{Q} denote a between-cell pairwise
 677 similarity matrix defined in UMAP computed with the dataset following feature selection $\mathbf{Y} \in \mathbb{R}^{n \times p}$, where
 678 $\mathbf{Y} = \mathbf{Xw}$ and $p \ll d$. Here, elastic net regression is used to find a sparse and robust solution of \mathbf{w} that minimizes
 679 the KL divergence as,

$$\text{KL}(\mathbf{P} \parallel \mathbf{Q}) = \sum_i \sum_j p_{ij} \log \frac{p_{ij}}{q_{ij}}. \quad (20)$$

680 Features with non-zero weights in \mathbf{w} are considered useful for preserving the embedding structure and selected
 681 for downstream analysis. We performed SCMER feature selection using the scmer v.0.1.0a3 package in python
 682 by constructing a k -nearest neighbor graph ($k = 10$) according to pairwise Euclidean distances of cells based
 683 on their first 50 principal components and using the default regression penalty weight parameters (lasso =
 684 $3.87e - 4$, ridge = 0).

685 **Hotspot** Hotspot [53] is an unsupervised gene module identification approach that performs feature selection
686 through a test statistic that measures the association of a gene's expression with the between-cell similarity graph
687 defined based on the full feature matrix, \mathbf{X} . More specifically, first, a k -nearest neighbor cell affinity graph ($k = 10$)
688 is defined based on pairwise Euclidean distances between all pairs of cells using a Gaussian kernel as,

$$w_{ij} = \begin{cases} \exp\left(-\frac{\|\mathbf{x}_{v_i} - \mathbf{x}_{v_{ij}}\|^2}{\sigma_i^2}\right), & \text{if } v_j \in \mathcal{N}_i \\ 0, & \text{otherwise.} \end{cases} \quad (21)$$

689 Here, cells v_i and v_j are connected with an edge with edge weight w_{ij} if the cell v_j is within the set of v_i 's
690 neighbors such that $\sum_j w_{ij} = 1$ for each cell and σ_i represents the bandwidth parameter for cell i defined as the
691 distance to the $\frac{k}{3}$ neighbor. For a given feature $\mathbf{f} \in \mathbb{R}^n$, representing expression across all n cells where f_i is the
692 mean-centered and standardized expression of feature \mathbf{f} in cell i according to a null distribution model of gene
693 expression, the local autocorrelation test statistic representing the dependence of each gene on the graph structure is
694 defined as,

$$H_f = \sum_i \sum_{j \neq i} w_{ij} f_i f_j. \quad (22)$$

695 Hotspot was implemented using the hotspot v1.1.1 package in python, where we selected the top p features by
696 sorting features in ascending order according to their significance with respect to a null model defined by a negative
697 binomial distribution with the mean dependent on library size.

698 **All features** To consider a baseline representation without feature selection, we evaluated performance using all
699 features from each dataset following quality control preprocessing.

700 **Random features** As a second baseline strategy, we simply selected a subset of random features without
701 replacement. Results were computed across twenty random initializations for each dataset.

702 **DELVE** DELVE was run as previously described. Here, we constructed a weighted k -nearest neighbor affinity
703 graph of cells ($k = 10$), and 1000 neighborhoods were sketched to identify dynamic seed feature clusters ($c = 3$
704 for the simulated dataset, $c = 5$ for the RPE cell cycle dataset, $c = 5$ for the CD8 T cell differentiation dataset,
705 and $c = 10$ for PDAC cell cycle datasets). Results were computed across twenty random initializations for each
706 dataset.

707 **Datasets**

708 We evaluated feature selection methods based on how well retained features could adequately recover biological
709 trajectories under various noise conditions, biological contexts, and single-cell technologies.

710 **Splatter simulation** Splatter [72] is a single-cell RNA sequencing simulation software that generates count data
711 using a gamma-Poisson hierarchical model with modifications to alter the mean-variance relationship amongst
712 genes, the library size, or sparsity. We used splatter to simulate a total of 90 ground truth datasets with different
713 trajectory structures (e.g. linear, bifurcation, and tree topologies). First, we estimated simulation parameters
714 by fitting the model to a real single-cell RNA sequencing dataset consisting of human pluripotent stem cells
715 differentiating into mesoderm progenitors [126]. We then used the estimated parameters (`mean_rate` = 0.0173,
716 `mean_shape` = 0.54, `lib_loc` = 12.6, `lib_scale` = 0.423, `out_prob` = 0.000342, `out_fac_loc` = 0.1,
717 `out_fac_scale` = 0.4, `bcv` = 0.1, `bcv_df` = 90.2, `dropout` = `None`) to simulate a diverse set of ground
718 truth reference trajectory datasets with the splatter paths function (python wrapper `scprep SplatSimulate v1.1.0` of
719 `splatter v1.18.2`). Here, a reference trajectory structure (e.g. bifurcation) was used to simulate linear and nonlinear
720 changes in the mean expression of genes along each step of the specified differentiation path. We simulated ten
721 differentiation datasets (1500 cells, 500 genes, 6 clusters) for each trajectory type (linear, bifurcation, tree) by
722 modifying (1) the probability of a cell belonging to a cluster by randomly sampling from a Dirichlet distribution
723 with six categories and a uniform concentration of one and (2) the path skew by randomly sampling from a beta
724 distribution ($\alpha = 10$, $\beta = 10$). The output of each simulation is a ground truth reference consisting of cell-to-cluster
725 membership, differentially expressed genes per cluster or path, as well as a latent `step` vector that indicates the

726 progression of each cell within a cluster. Lastly, we modified the step vector to be monotonically increasing across
727 clusters within the specified differentiation path to obtain a reference pseudotime measurement.

728 To estimate how well feature selection methods can identify genes that represent cell populations and are differenti-
729 ally expressed along a differentiation path in noisy single-cell RNA sequencing data, we added relevant sources of
730 biological and technical noise to the reference datasets.

731 1. *Biological Coefficient of Variation (BCV)*: To simulate the effect of stochastic gene expression, we modified
732 the biological coefficient of variation parameter within splatter ($BCV = 0.1, 0.25, 0.5$). This scaling factor
733 controls the mean-variance relationship between genes, where lowly expressed genes are more variable than
734 highly expressed genes, following a γ distribution.

735 2. *Library size*: The total number of profiled mRNA transcripts per cell, or library size, can vary between cells
736 within a single-cell RNA sequencing experiment and can influence the detection of differentially expressed
737 genes [73], as well as impact the reproducibility of the lower-dimensional representation of the data [74].
738 To simulate the effect of differences in sequencing depth, we proportionally adjusted the gene means for
739 each cell by modifying the location parameter ($1.lib_loc = 12, 11, 10$) of the log-normal distribution within
740 splatter that estimates the library size scaling factors.

741 3. *Technical dropout*: Single-cell RNA sequencing data contain a large proportion of zeros, where only a small
742 fraction of total transcripts are detected due to capture inefficiency and amplification noise [127]. To simulate
743 the inefficient capture of mRNA molecules and account for the trend that lowly expressed genes are more
744 likely to be affected by dropout, we undersampled mRNA counts by sampling from a binomial distribution
745 with the scale parameter or dropout rate proportional to the mean expression of each gene as previously
746 described in ref. [128] as,

$$r_i = \exp(-\lambda\mu_i^2). \quad (23)$$

747 Here, μ_i represents the log mean expression of gene i , and λ is a hyperparameter that controls the magnitude
748 of dropout ($\lambda = 0, 0.05, 0.1$).

749 In our subsequent feature selection method analyses, we selected the top $p = 100$ features under each feature
750 selection approach.

751 **RPE analysis** The retinal pigmented epithelial (RPE) dataset [92] is an iterative indirect immunofluorescence
752 imaging (4i) dataset consisting of RPE cells undergoing the cell cycle. Here, time-lapse imaging was performed
753 on an asynchronous population of non-transformed human retinal pigmented epithelial cells expressing a PCNA-
754 mTurquoise2 reporter in order to record the cell cycle phase (G0/G1, S, G2, M) and molecular age (time since last
755 mitosis) of each cell. Following time-lapse imaging, cells were fixed and 48 core cell cycle effectors were profiled
756 using 4i [8]. For preprocessing, we min-max normalized the data and performed batch effect correction on the
757 replicates using ComBat [129]. Lastly, to refine phase annotations and distinguish G0 from G1 cells, we selected
758 cycling G1 cells according to the bimodal distribution of $\frac{pRB}{RB}$ expression as described in ref. [92]. Of note, cells
759 were excluded if they did not have ground truth phase or age annotations. The resultant dataset consisted of 2759
760 cells \times 241 features describing the expression and localization of different protein markers. For our subsequent
761 analysis, we selected the top $p = 30$ features for each feature selection approach.

762 **PDAC analysis** The pancreatic ductal adenocarcinoma (PDAC) dataset is an iterative indirect immunofluorescence
763 imaging dataset consisting of 9 human PDAC cell lines: BxPC3, CFPAC, HPAC, MiaPaCa, Pa01C, Pa02C, Pa16C,
764 PANC1, UM53. For each cell line (e.g. BxPC3) under control conditions, we min-max normalized the data. Cell
765 cycle phases (G0, G1, S, G2, M) were annotated *a priori* based on manual gating cells according to the abundance
766 of core cell cycle markers. Phospho-RB (pRB) was used to distinguish proliferative cells (G1/S/G2/M, high pRB)
767 from arrested cells (G0, low pRB). DNA content, E2F1, cyclin A (cycA), and phospho-p21 (p-p21) were used to
768 distinguish G1 (DNA content = 2C, low cycA), S (DNA content = 2-4C, high E2F1), G2 (DNA content = 4C, high
769 cycA), and M (DNA content = 4C, high p-p21). For our subsequent analysis, we selected the top $p = 30$ features
770 for each feature selection approach.

771 **CD8 T cell differentiation analysis** The CD8 T cell differentiation dataset [106] is a single-cell RNA sequencing
772 dataset consisting of mouse splenic CD8+ T cells profiled over 12-time points ($d = \text{day}$) following infection
773 with the Armstrong strain of the lymphocytic choriomeningitis virus: Naive, d3-, d4-, d5-, d6-, d7-, d10-, d14-,
774 d21-, d32-, d60-, d90- post-infection. Spleen single-cell RNA sequencing data were accessed from the Gene
775 Expression Omnibus using the accession code GSE131847 and concatenated into a single matrix. The dataset was
776 subsequently quality control filtered according to the distribution of molecular counts. To remove dead or dying

771 cells, we filtered cells that had more than twenty percent of their total reads mapped to mitochondrial transcripts.
772 Genes that were observed in less than three cells or had less than 400 counts were also removed. Following cell
773 and gene filtering, the data were transcripts-per-million normalized, log+1 transformed, and variance filtered using
774 highly variable gene selection, such that the resulting dataset consisted of 29893 cells \times 500 genes (See *Highly*
775 *variable genes*). Lastly, to obtain lineage labels for trajectory analysis, cells were scored and gated according to
776 their average expression of known memory markers (Bcl2, Sell, Il7r) using the score_genes function in Scanpy
777 v1.8.1. When evaluating feature selection methods, we selected the top $p = 250$ features for each feature selection
778 approach.

779 Evaluation

780 Classification and regression

781 ***k*-nearest neighbor classification** To quantitatively compare feature selection methods on retaining features that
782 are representative of cell types, we aimed to implement an approach that would assess the quality of the graph
783 structure. k -nearest neighbors classification is a supervised learning algorithm that classifies data based on labels
784 of the k -most similar cells according to their gene or protein expression, where the output of this algorithm is a
785 set of labels for every cell. We performed k -nearest neighbors classification to predict cell type labels from the
786 simulated single-cell RNA sequencing datasets as follows. First, 3-fold cross-validation was performed using
787 stratified random sampling to assign cells to either a training or a test set. Stratified random sampling was chosen
788 to mitigate the effect of cell type class imbalance. Within each fold, feature selection was then performed on the
789 training data to identify the top $p = 250$ relevant features according to a feature selection strategy. Next, a k -nearest
790 neighbor classifier ($k = 3$) was fit on the feature selected training data to predict the cell type labels of the feature
791 selected test query points. Here, labels are predicted as the mode of the cell type labels from the closest training
792 data points according to Euclidean distance. Classification performance was subsequently assessed according to the
793 median classification accuracy with respect to the ground truth cell type labels across folds.

794 **Support Vector Machine** The Support Vector Machines (SVM) [130] is a supervised learning algorithm that con-
795 structs hyperplanes in the high-dimensional feature space to separate classes. We implemented SVM classification
796 or regression using the sklearn v0.23.4 package in python. SVM classification was used to predict cell cycle phase
797 labels for both RPE and PDAC 4i datasets, whereas regression was used to predict age measurements from time
798 lapse imaging for the RPE dataset. Here, Nested 10-fold cross-validation was performed using stratified random
799 sampling to assign cells to either a training set or a test set. Within each fold, feature selection was performed to
800 identify the p most relevant features according to a feature selection strategy. SVM hyperparameters were then
801 tuned over a grid search and phase labels were subsequently predicted from the test data according to those p
802 features. Classification performance was assessed according to the median classification accuracy with respect to
803 the ground truth cell type labels across folds. Regression performance was assessed according to the average mean
804 squared error with respect to ground truth age measurements across folds.

805 Precision@k

806 To evaluate the biological relevance of selected features from each method, we computed precision@k ($p@k$) as
807 the proportion of top k selected features that were considered to be biologically relevant according to a ground
808 truth reference as,

$$p@k = \frac{|\mathcal{F}_{s,k} \cap \mathcal{F}_r|}{|\mathcal{F}_{s,k}|}, \quad (24)$$

809 where $\mathcal{F}_{s,k}$ indicates the set of selected features at threshold k , where $\mathcal{F}_{s,k} \subset \mathcal{F}_s$, and \mathcal{F}_r indicates the set of
810 reference features. Reference features were defined as either (1) the ground truth differentially expressed features
811 within a cluster or along a differentiation path from the single-cell RNA sequencing simulation study (see *Splatter*
812 *simulation*) or (2) the features determined to be useful for classifying cells according to cell cycle phase using
813 a random forest classifier trained on ground truth phase annotations from time-lapse imaging for the protein
814 immunofluorescence imaging datasets (See *Random forest, RPE analysis, PDAC analysis*).

815 Unsupervised clustering

816 To evaluate feature selection method performance on retaining features that are informative for identifying canonical
817 cell types, we performed unsupervised clustering on the data defined by the top p ranked features from a feature

818 selection strategy. More specifically, for each feature selection approach, clustering was performed on the selected
819 data using either the KMeans++ algorithm [103] with the number of centroids set as the same number of
820 ground truth cell cycle phase labels for the protein immunofluorescence imaging datasets (RPE: $c = 4$, PDAC:
821 $c = 5$).

822 To assess the accuracy of clustering assignments, we quantified a normalized mutual information (NMI) score
823 between the predicted cluster labels and the ground truth cell type labels. Normalized mutual information [131] is a
824 clustering quality metric that measures the amount of shared information between two cell-to-cluster partitions (\mathbf{u}
825 and \mathbf{v} , such that the i th entry u_i gives the cluster assignment of cell i) as,

$$\text{NMI} = \frac{2I(\mathbf{u}; \mathbf{v})}{H(\mathbf{u})H(\mathbf{v})}, \quad (25)$$

826 where, $I(\mathbf{u}; \mathbf{v})$ measures the mutual information between ground truth cell type labels \mathbf{u} and cluster labels \mathbf{v} , and
827 $H(\mathbf{u})$ or $H(\mathbf{v})$ indicates the Shannon entropy or the amount of uncertainty for a given set of labels. Here, a score of
828 1 indicates that clustering on the selected features perfectly recovers the ground truth cell type labels. KMeans++
829 clustering was implemented using the KMeans function in sklearn v0.23.4.

830 Protein-protein interaction networks

831 In this work, we aimed to test whether features within DELVE dynamic clusters had experimental evidence
832 of co-regulation as compared to random assignment. The STRING (search tool for the retrieval of interacting
833 genes/proteins) database [121] is a relational database that computes protein association scores according to
834 information derived from several *evidence* channels, including computational predictions (e.g. neighborhood,
835 fusion, co-occurrence), co-expression, experimental assays, pathway databases, and literature text mining. To assess
836 the significance of protein interactions amongst features within a DELVE cluster, we performed a permutation test
837 with a test statistic derived from STRING association scores using experimental evidence as follows.

838 Let $\mathcal{G}_p = (\mathcal{N}_p, \mathcal{E}_p)$ denote a graph of p proteins from a DELVE cluster comprising the nodes \mathcal{N}_p , and \mathcal{E}_p denote the
839 set of edges, where edge weights encode the association scores of experimentally-derived protein-protein interaction
840 evidence from the STRING database. Moreover, let $\mathcal{G}_r = (\mathcal{N}_r, \mathcal{E}_r)$ denote a graph of r proteins randomly sampled
841 without replacement from the full feature space d such that $r = p$ comprising the nodes \mathcal{N}_r , and \mathcal{E}_r denote the
842 set of edges encoding the experimentally-derived association scores between those r proteins from the STRING
843 database. We compute the permutation p -value as described previously in ref. [132] as,

$$p\text{-value} = \frac{N + 1}{R + 1}. \quad (26)$$

844 Here N indicates the number of times that $T_r \geq T_{\text{obs}}$ out of R random permutations ($R = 1000$), where T_r is
845 the average degree of a STRING association network from randomly permuted features as $T_r = \frac{|\mathcal{N}_r|}{|\mathcal{E}_r|}$, and T_{obs}
846 is the average degree of a STRING association network from the features identified within a DELVE cluster as
847 $T_{\text{obs}} = \frac{|\mathcal{N}_p|}{|\mathcal{E}_p|}$. Of note, networks with higher degree are more connected, and thus show greater experimental
848 evidence of protein-protein interactions. Experimental evidence-based association scores were obtained from
849 the STRING database using the stringdb v0.1.5 package in python and networks were generated using networkx
850 v2.5.1.

851 Trajectory inference and analysis

852 To evaluate how well feature selection methods could identify features that recapitulate the underlying cellular
853 trajectory and can be used for trajectory analysis, we computed three metrics to assess trajectory preservation at
854 different stages of inference: accuracy of the inferred trajectory graph, correlation of estimated pseudotime to
855 the ground truth cell progression measurements, and the significance of dynamic features identified following
856 regression analysis.

857 To obtain predicted trajectories, we performed trajectory inference using the diffusion pseudotime algorithm [83]
858 based on 20 diffusion map components generated from a k -nearest neighbor graph ($k = 10$), where edge weights
859 were determined by pairwise Euclidean distances between cells according to selected feature expression. Inference
860 was performed for the following lineages – simulated trajectories: all cells, arrested trajectory: cells with G0 phase
861 annotation, proliferative trajectory: cells with G1, S, G2, or M phase annotation, CD8 T cell memory lineage: cells
862 following day 7 of infection with a memory score. Moreover, for each feature selection approach, we estimated

863 pseudotime using ten random root cells according to *a priori* biological knowledge. Root cells were chosen as
864 either (1) cells with the smallest ground truth pseudotime annotation for the simulated datasets, (2) cells with the
865 youngest molecular age for 4i cell cycle datasets, or (3) cells from the day 7 population along the memory lineage
866 for the CD8 T cell differentiation dataset. Feature selection trajectory performance was subsequently assessed as
867 follows.

868 1. *Trajectory graph similarity*: Partition-based graph abstraction (PAGA) [104] performs trajectory inference
869 by constructing a coarse grained trajectory graph of the data. First cell populations are determined either
870 through unsupervised clustering, graph partitioning, or *a priori* experimental annotations. Next, a statistical
871 measure of edge connectivity is computed between cell populations to estimate the confidence of a cell
872 population transition. To assess if feature selection methods retain features that represent coarse cell type
873 transitions, we compared predicted PAGA trajectory graphs to ground truth cell cycle reference trajectories
874 curated from the literature [92]. First, PAGA connectivity was estimated between ground truth cell cycle
875 phase groups using the k -nearest neighbor graph ($k = 10$) based on pairwise Euclidean distances between
876 cells according to selected feature expression. We then computed the Jaccard distance between predicted and
877 reference trajectories as,

$$d_j(\mathbf{W}_p, \mathbf{W}_r) = 1 - \frac{|\mathbf{W}_p \cap \mathbf{W}_r|}{|\mathbf{W}_p \cup \mathbf{W}_r|}. \quad (27)$$

878 \mathbf{W}_p indicates the predicted cell type transition adjacency matrix, where each entry $W_{p,ij}$ represents the
879 connectivity strength between cell populations i and j from PAGA and \mathbf{W}_r indicates the reference trajectory
880 adjacency matrix with entries encoding ground truth cell type transitions curated from the literature. Here, a
881 lower Jaccard distance indicates that predicted trajectories better capture known cellular transitions.

882 2. *Pseudotime correlation*: To evaluate if feature selection methods retain features that accurately represent a
883 cell's progression through a biological trajectory, we computed the Spearman rank correlation coefficient
884 between estimated pseudotime following feature selection and ground truth cell progression annotations
885 (e.g. the ground truth pseudotime labels generated from simulations, time-lapse imaging molecular age
886 measurements).

887 3. *Regression analysis*: To identify genes associated with the inferred CD8+ T cell differentiation trajectory
888 following feature selection, we performed regression analysis for each gene ($d = 500$) along estimated
889 pseudotime using a negative binomial generalized additive model (GAM). Genes were considered to be
890 differentially expressed along the memory lineage if they had a q -value < 0.05 following Benjamini-
891 Hochberg false discovery rate correction [122].

892 4. *Gene Ontology*: To identify the biological relevance of the differentially expressed genes along inferred
893 CD8+ T cell differentiation trajectories specific to each feature selection strategy, we performed gene set
894 enrichment analysis on the set difference of significant genes from either highly variable gene selection or
895 DELVE feature selection using Enrichr [123]. Here, we considered the mouse gene sets from GO Biological
896 Process 2021.

897 Diffusion pseudotime was implemented using the dpt function in Scanpy v1.8.1, PAGA was implemented using
898 the paga function in scanpy v1.8.1, GAM regression was implemented using the statsmodels v0.12.2 package in
899 python, and gene set enrichment analysis was performed using the enrichr function in gseapy v1.0.4 package in
900 python.

901 PHATE visualizations

902 To qualitatively compare lower dimensional representations from each feature selection strategy, we performed non-
903 linear dimensionality reduction using PHATE (potential of heat-diffusion for affinity-based transition embedding)
904 [75] as this approach performs reasonably well for representing complex continuous biological trajectories. PHATE
905 was implemented using the phate v1.0.7 package in python. Here, we used the same set of hyperparameters across
906 all feature selection strategies ($knn = 30$, $t = 10$, $decay = 40$).

907 Aggregate scores

908 To rank feature selection methods on preserving biological trajectories in the presence of single-cell noise, we
909 computed rank aggregate scores by taking the mean of scaled method scores across simulated single-cell RNA
910 sequencing datasets from a trajectory type and noise condition (e.g. linear trajectory, dropout noise). More

911 specifically, we first defined an overall method score per dataset as the median of each metric. Method scores were
912 subsequently min-max scaled to ensure datasets were equally weighted prior to computing the average.

913 Data and code availability

914 The raw publicly available single-cell datasets used in this study are available in the Zenodo repository <https://doi.org/10.5281/zenodo.4525425> for the RPE cell cycle dataset [133], the Zenodo repository <https://doi.org/10.5281/zenodo.7860332> for the PDAC cell cycle datasets [134], and the Gene Expression Omnibus (GEO) under the accession code GSE131847 for the CD8 T cell differentiation dataset [135]. The preprocessed datasets are available in the Zenodo repository <https://doi.org/10.5281/zenodo.7883604> [136]. DELVE is implemented as an open-source python package and is publicly available at <https://github.com/jranek/delve>. Source code including all functions for benchmarking feature selection methods including preprocessing, feature selection, evaluation, and plotting are publicly available at: https://github.com/jranek/delve_benchmark.

923 Funding

924 This work was supported by the National Institutes of Health, F31-HL156433 (JSR), 5T32-GM067553 (JSR),
925 R01-GM138834 (JEP), NSF CAREER Award 1845796 (JEP), and NSF Award 2242980 (JEP).

926 Authors' contributions

927 JSR and WS conceptualized the study. JSR, NS, and JEP designed the method and computational analyses. JSR
928 performed the data preprocessing, benchmarking, evaluation, and analysis. WS provided input for the RPE and
929 PDAC cell cycle study. JM provided input for the CD8 T cell study. JSR wrote the manuscript with help from all
930 authors. All authors read and approved of the final manuscript.

931 Acknowledgements

932 We would like to thank Tarek Zikry, Garrett Sessions, and Alec Plotkin for their insightful discussions related to
933 this work.

934 References

- 935 [1] Matthew H Spitzer and Garry P Nolan. Mass cytometry: Single cells, many features. *Cell*, 165(4):780–791, May 2016.
- 936 [2] N Baumgarth and M Roederer. A practical approach to multicolor flow cytometry for immunophenotyping. *J. Immunol. Methods*, 243(1-2):77–97, September 2000.
- 937 [3] Dmitry R Bandura, Vladimir I Baranov, Olga I Ornatsky, Alexei Antonov, Robert Kinach, Xudong Lou, Serguei Pavlov, Sergey Vorobiev, John E Dick, and Scott D Tanner. Mass cytometry: technique for real time
938 single cell multitarget immunoassay based on inductively coupled plasma time-of-flight mass spectrometry. *Anal. Chem.*, 81(16):6813–6822, August 2009.
- 939 [4] Grace X Y Zheng, Jessica M Terry, Phillip Belgrader, Paul Ryvkin, Zachary W Bent, Ryan Wilson, Solongo B
940 Ziraldo, Tobias D Wheeler, Geoff P McDermott, Junjie Zhu, Mark T Gregory, Joe Shuga, Luz Montesclaros,
941 Jason G Underwood, Donald A Masquelier, Stefanie Y Nishimura, Michael Schnall-Levin, Paul W Wyatt,
942 Christopher M Hindson, Rajiv Bharadwaj, Alexander Wong, Kevin D Ness, Lan W Beppu, H Joachim
943 Deeg, Christopher McFarland, Keith R Loeb, William J Valente, Nolan G Ericson, Emily A Stevens,
944 Jerald P Radich, Tarjei S Mikkelsen, Benjamin J Hindson, and Jason H Bielas. Massively parallel digital
945 transcriptional profiling of single cells. *Nat. Commun.*, 8:14049, January 2017.
- 946 [5] Evan Z Macosko, Anindita Basu, Rahul Satija, James Nemesh, Karthik Shekhar, Melissa Goldman, Itay
947 Tirosh, Allison R Bialas, Nolan Kamitaki, Emily M Martersteck, John J Trombetta, David A Weitz, Joshua R
948 Sanes, Alex K Shalek, Aviv Regev, and Steven A McCarroll. Highly parallel genome-wide expression
949 profiling of individual cells using nanoliter droplets. *Cell*, 161(5):1202–1214, May 2015.
- 950 [6] Simone Picelli, Omid R Faridani, Asa K Björklund, Gösta Winberg, Sven Sagasser, and Rickard Sandberg.
951 Full-length RNA-seq from single cells using smart-seq2. *Nat. Protoc.*, 9(1):171–181, January 2014.

956 [7] Rapolas Zilionis, Juozas Nainys, Adrian Veres, Virginia Savova, David Zemmour, Allon M Klein, and Linas
957 Mazutis. Single-cell barcoding and sequencing using droplet microfluidics. *Nat. Protoc.*, 12(1):44–73,
958 January 2017.

959 [8] Gabriele Gut, Markus D Herrmann, and Lucas Pelkmans. Multiplexed protein maps link subcellular
960 organization to cellular states. *Science*, 361(6401), August 2018.

961 [9] Leeat Keren, Marc Bosse, Steve Thompson, Tyler Risom, Kausalia Vijayaragavan, Erin McCaffrey, Diana
962 Marquez, Roshan Angoshtari, Noah F Greenwald, Harris Fienberg, Jennifer Wang, Neeraja Kambham, David
963 Kirkwood, Garry Nolan, Thomas J Montine, Stephen J Galli, Robert West, Sean C Bendall, and Michael
964 Angelo. MIBI-TOF: A multiplexed imaging platform relates cellular phenotypes and tissue structure. *Sci
965 Adv*, 5(10):eaax5851, October 2019.

966 [10] Charlotte Giesen, Hao A O Wang, Denis Schapiro, Nevena Zivanovic, Andrea Jacobs, Bodo Hattendorf,
967 Peter J Schüffler, Daniel Grolimund, Joachim M Buhmann, Simone Brandt, Zsuzsanna Varga, Peter J Wild,
968 Detlef Günther, and Bernd Bodenmiller. Highly multiplexed imaging of tumor tissues with subcellular
969 resolution by mass cytometry. *Nat. Methods*, 11(4):417–422, April 2014.

970 [11] Yury Goltsev, Nikolay Samusik, Julia Kennedy-Darling, Salil Bhate, Matthew Hale, Gustavo Vazquez, Sarah
971 Black, and Garry P Nolan. Deep profiling of mouse splenic architecture with CODEX multiplexed imaging.
972 *Cell*, 174(4):968–981.e15, August 2018.

973 [12] Junyue Cao, Malte Spielmann, Xiaojie Qiu, Xingfan Huang, Daniel M Ibrahim, Andrew J Hill, Fan Zhang,
974 Stefan Mundlos, Lena Christiansen, Frank J Steemers, Cole Trapnell, and Jay Shendure. The single-cell
975 transcriptional landscape of mammalian organogenesis. *Nature*, 566(7745):496–502, February 2019.

976 [13] Jeffrey A Farrell, Yiqun Wang, Samantha J Riesenfeld, Karthik Shekhar, Aviv Regev, and Alexander F
977 Schier. Single-cell reconstruction of developmental trajectories during zebrafish embryogenesis. *Science*,
978 360(6392), June 2018.

979 [14] David Fawkner-Corbett, Agne Antanaviciute, Kaushal Parikh, Marta Jagielowicz, Ana Sousa Gerós, Tarun
980 Gupta, Neil Ashley, Doran Khamis, Darren Fowler, Edward Morrissey, Chris Cunningham, Paul R V
981 Johnson, Hashem Koohy, and Alison Simmons. Spatiotemporal analysis of human intestinal development at
982 single-cell resolution. *Cell*, 184(3):810–826.e23, February 2021.

983 [15] Emily M Holloway, Michael Czerwinski, Yu-Hwai Tsai, Joshua H Wu, Angeline Wu, Charlie J Childs,
984 Katherine D Walton, Caden W Sweet, Qianhui Yu, Ian Glass, Barbara Treutlein, J Gray Camp, and Jason R
985 Spence. Mapping development of the human intestinal niche at Single-Cell resolution. *Cell Stem Cell*,
986 28(3):568–580.e4, March 2021.

987 [16] Max Kaufmann, Hayley Evans, Anna-Lena Schaupp, Jan Broder Engler, Gurman Kaur, Anne Willing,
988 Nina Kursawe, Charlotte Schubert, Kathrine E Attfield, Lars Fugger, and Manuel A Friese. Identifying
989 CNS-colonizing T cells as potential therapeutic targets to prevent progression of multiple sclerosis. *Med (N
990 Y)*, 2(3):296–312.e8, March 2021.

991 [17] Tyler Risom, David R Glass, Inna Averbukh, Candace C Liu, Alex Baranski, Adam Kagel, Erin F McCaffrey,
992 Noah F Greenwald, Belén Rivero-Gutiérrez, Siri H Strand, Sushama Varma, Alex Kong, Leeat Keren,
993 Sucheta Srivastava, Chunfang Zhu, Zumana Khair, Deborah J Veis, Katherine Deschryver, Sujay Venkam,
994 Carlo Maley, E Shelley Hwang, Jeffrey R Marks, Sean C Bendall, Graham A Colditz, Robert B West, and
995 Michael Angelo. Transition to invasive breast cancer is associated with progressive changes in the structure
996 and composition of tumor stroma. *Cell*, 185(2):299–310.e18, January 2022.

997 [18] Prabhu S Arunachalam, Florian Wimmers, Chris Ka Pun Mok, Ranawaka A P M Perera, Madeleine Scott,
998 Thomas Hagan, Natalia Sigal, Yupeng Feng, Laurel Bristow, Owen Tak-Yin Tsang, Dhananjay Wagh, John
999 Coller, Kathryn L Pellegrini, Dmitri Kazmin, Ghina Alaaeddine, Wai Shing Leung, Jacky Man Chun Chan,
1000 Thomas Shiu Hong Chik, Chris Yau Chung Choi, Christopher Huerta, Michele Paine McCullough, Huibin
1001 Lv, Evan Anderson, Srilatha Edupuganti, Amit A Upadhyay, Steve E Bosinger, Holden Terry Maecker,
1002 Purvesh Khatri, Nadine Roush, Malik Peiris, and Bali Pulendran. Systems biological assessment of
1003 immunity to mild versus severe COVID-19 infection in humans. *Science*, 369(6508):1210–1220, September
1004 2020.

1005 [19] Wouter Saelens, Robrecht Cannoodt, Helena Todorov, and Yvan Saeys. A comparison of single-cell trajectory
1006 inference methods. *Nat. Biotechnol.*, 37(5):547–554, May 2019.

1007 [20] Cole Trapnell, Davide Cacchiarelli, Jonna Grimsby, Prapti Pokharel, Shuqiang Li, Michael Morse, Niall J
1008 Lennon, Kenneth J Livak, Tarjei S Mikkelsen, and John L Rinn. The dynamics and regulators of cell fate
1009 decisions are revealed by pseudotemporal ordering of single cells. *Nat. Biotechnol.*, 32(4):381–386, April
1010 2014.

1011 [21] Zhicheng Ji and Hongkai Ji. TSCAN: Pseudo-time reconstruction and evaluation in single-cell RNA-seq
1012 analysis. *Nucleic Acids Res.*, 44(13):e117, July 2016.

1013 [22] Jaehoon Shin, Daniel A Berg, Yunhua Zhu, Joseph Y Shin, Juan Song, Michael A Bonaguidi, Grigori
1014 Enikolopov, David W Nauen, Kimberly M Christian, Guo-Li Ming, and Hongjun Song. Single-Cell RNA-
1015 Seq with waterfall reveals molecular cascades underlying adult neurogenesis. *Cell Stem Cell*, 17(3):360–372,
1016 September 2015.

1017 [23] Robrecht Cannoodt, Wouter Saelens, Dorine Sichien, Simon Tavernier, and Yvan Saeys. SCORPIUS
1018 improves trajectory inference and identifies novel modules in dendritic cell development. *bioRxiv*, October
1019 2016.

1020 [24] Kelly Street, Davide Risso, Russell B Fletcher, Diya Das, John Ngai, Nir Yosef, Elizabeth Purdom, and
1021 Sandrine Dudoit. Slingshot: cell lineage and pseudotime inference for single-cell transcriptomics. *BMC
1022 Genomics*, 19(1):477, June 2018.

1023 [25] Joshua D Welch, Alexander J Hartemink, and Jan F Prins. SLICER: inferring branched, nonlinear cellular
1024 trajectories from single cell RNA-seq data. *Genome Biol.*, 17(1):106, May 2016.

1025 [26] F Alexander Wolf, Fiona K Hamey, Mireya Plass, Jordi Solana, Joakim S Dahlin, Berthold Göttgens,
1026 Nikolaus Rajewsky, Lukas Simon, and Fabian J Theis. PAGA: graph abstraction reconciles clustering with
1027 trajectory inference through a topology preserving map of single cells. *Genome Biol.*, 20(1):59, March 2019.

1028 [27] Junyue Cao, Malte Spielmann, Xiaojie Qiu, Xingfan Huang, Daniel M Ibrahim, Andrew J Hill, Fan Zhang,
1029 Stefan Mundlos, Lena Christiansen, Frank J Steemers, Cole Trapnell, and Jay Shendure. The single-cell
1030 transcriptional landscape of mammalian organogenesis. *Nature*, 566(7745):496–502, February 2019.

1031 [28] Laleh Haghverdi, Maren Büttner, F Alexander Wolf, Florian Buettner, and Fabian J Theis. Diffusion
1032 pseudotime robustly reconstructs lineage branching. *Nat. Methods*, 13(10):845–848, October 2016.

1033 [29] Manu Setty, Vaidotas Kiseliolas, Jacob Levine, Adam Gayoso, Linas Mazutis, and Dana Pe'er. Character-
1034 ization of cell fate probabilities in single-cell data with palantir. *Nat. Biotechnol.*, 37(4):451–460, April
1035 2019.

1036 [30] Shobana V Stassen, Gwinky G K Yip, Kenneth K Y Wong, Joshua W K Ho, and Kevin K Tsia. Generalized
1037 and scalable trajectory inference in single-cell omics data with VIA. *Nat. Commun.*, 12(1):5528, September
1038 2021.

1039 [31] Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon, Joshua
1040 Gould, Siyan Liu, Stacie Lin, Peter Berube, Lia Lee, Jenny Chen, Justin Brumbaugh, Philippe Rigolet,
1041 Konrad Hochedlinger, Rudolf Jaenisch, Aviv Regev, and Eric S Lander. Optimal-Transport analysis of Single-
1042 Cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943.e22,
1043 February 2019.

1044 [32] Alexander Tong, Jessie Huang, Guy Wolf, David van Dijk, and Smita Krishnaswamy. TrajectoryNet: A
1045 dynamic optimal transport network for modeling cellular dynamics. *Proc Mach Learn Res*, 119:9526–9536,
1046 July 2020.

1047 [33] Koen Van den Berge, Hector Roux de Bézieux, Kelly Street, Wouter Saelens, Robrecht Cannoodt, Yvan
1048 Saeys, Sandrine Dudoit, and Lieven Clement. Trajectory-based differential expression analysis for single-cell
1049 sequencing data. *Nat. Commun.*, 11(1):1201, March 2020.

1050 [34] Wenpin Hou, Zhicheng Ji, Zeyu Chen, E John Wherry, Stephanie C Hicks, and Hongkai Ji. A statistical
1051 framework for differential pseudotime analysis with multiple single-cell RNA-seq samples. *bioRxiv*, page
1052 2021.07.10.451910, July 2021.

1053 [35] Kieran R Campbell and Christopher Yau. Uncovering pseudotemporal trajectories with covariates from
1054 single cell and bulk expression data. *Nat. Commun.*, 9(1):2442, June 2018.

1055 [36] Shila Ghazanfar, Yingxin Lin, Xianbin Su, David Ming Lin, Ellis Patrick, Ze-Guang Han, John C Marioni,
1056 and Jean Yee Hwa Yang. Investigating higher-order interactions in single-cell data with scHOT. *Nat.*
1057 *Methods*, 17(8):799–806, August 2020.

1058 [37] Atul Deshpande, Li-Fang Chu, Ron Stewart, and Anthony Gitter. Network inference with granger causality
1059 ensembles on single-cell transcriptomics. *Cell Rep.*, 38(6):110333, February 2022.

1060 [38] Ayelet Alpert, Lindsay S Moore, Tania Dubovik, and Shai S Shen-Orr. Alignment of single-cell trajectories
1061 to compare cellular expression dynamics. *Nat. Methods*, 15(4):267–270, April 2018.

1062 [39] Reiichi Sugihara, Yuki Kato, Tomoya Mori, and Yukio Kawahara. Alignment of single-cell trajectory trees
1063 with CAPITAL. *Nat. Commun.*, 13(1):5972, October 2022.

1064 [40] Stephanie C Hicks, F William Townes, Mingxiang Teng, and Rafael A Irizarry. Missing data and technical
1065 variability in single-cell RNA-sequencing experiments. *Biostatistics*, 19(4):562–578, October 2018.

1066 [41] John W Hickey, Elizabeth K Neumann, Andrea J Radtke, Jeannie M Camarillo, Rebecca T Beuschel,
1067 Alexandre Albanese, Elizabeth McDonough, Julia Hatler, Anne E Wiblin, Jeremy Fisher, Josh Croteau,
1068 Eliza C Small, Anup Sood, Richard M Caprioli, R Michael Angelo, Garry P Nolan, Kwanghun Chung,
1069 Stephen M Hewitt, Ronald N Germain, Jeffrey M Spraggins, Emma Lundberg, Michael P Snyder, Neil L
1070 Kelleher, and Sinem K Saka. Spatial mapping of protein composition and tissue organization: a primer for
1071 multiplexed antibody-based imaging. *Nat. Methods*, 19(3):284–295, March 2022.

1072 [42] Mohammed Charrouf, Marcel J T Reinders, and Ahmed Mahfouz. Untangling biological factors influencing
1073 trajectory inference from single cell data. *NAR Genom Bioinform*, 2(3):lqaa053, September 2020.

1074 [43] Sophie Tritschler, Maren Büttner, David S Fischer, Marius Lange, Volker Bergen, Heiko Lickert, and
1075 Fabian J Theis. Concepts and limitations for learning developmental trajectories from single cell genomics.
1076 *Development*, 146(12), June 2019.

1077 [44] Chenxu Zhu, Sebastian Preissl, and Bing Ren. Single-cell multimodal omics: the power of many. *Nat.*
1078 *Methods*, 17(1):11–14, January 2020.

1079 [45] Pengyi Yang, Hao Huang, and Chunlei Liu. Feature selection revisited in the single-cell era. *Genome Biol.*,
1080 22(1):321, December 2021.

1081 [46] Leo Breiman. Random forests. *Mach. Learn.*, 45(1):5–32, October 2001.

1082 [47] Hanchuan Peng, Fuhui Long, and Chris Ding. Feature selection based on mutual information: criteria of max-
1083 dependency, max-relevance, and min-redundancy. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(8):1226–1238,
1084 August 2005.

1085 [48] Pablo A Estévez, Michel Tesmer, Claudio A Perez, and Jacek M Zurada. Normalized mutual information
1086 feature selection. *IEEE Trans. Neural Netw.*, 20(2):189–201, February 2009.

1087 [49] Thomas Liechti, Lukas M Weber, Thomas M Ashhurst, Natalie Stanley, Martin Prlic, Sofie Van Gassen, and
1088 Florian Mair. An updated guide for the perplexed: cytometry in the high-dimensional era. *Nat. Immunol.*,
1089 22(10):1190–1197, October 2021.

1090 [50] Rahul Satija, Jeffrey A Farrell, David Gennert, Alexander F Schier, and Aviv Regev. Spatial reconstruction
1091 of single-cell gene expression data. *Nat. Biotechnol.*, 33(5):495–502, May 2015.

1092 [51] Tim Stuart, Andrew Butler, Paul Hoffman, Christoph Hafemeister, Efthymia Papalexi, William M Mauck, 3rd,
1093 Yuhan Hao, Marlon Stoeckius, Peter Smibert, and Rahul Satija. Comprehensive integration of Single-Cell
1094 data. *Cell*, 177(7):1888–1902.e21, June 2019.

1095 [52] Xiaofei He, Deng Cai, and Partha Niyogi. Laplacian score for feature selection. In *Advances in Neural*
1096 *Information Processing Systems*, volume 18. MIT Press, 2005.

1097 [53] David DeTomaso and Nir Yosef. Hotspot identifies informative gene modules across modalities of single-cell
1098 genomics. *Cell Syst.*, 12(5):446–456.e9, May 2021.

1099 [54] Shaoheng Liang, Vakul Mohanty, Jinzhuang Dou, Qi Miao, Yuefan Huang, Muharrem Müftüoğlu, Li Ding,
1100 Weiyi Peng, and Ken Chen. Single-cell manifold-preserving feature selection for detecting rare cell
1101 populations. *Nature Computational Science*, 1(5):374–384, May 2021.

1102 [55] Deng Cai, Chiyuan Zhang, and Xiaofei He. Unsupervised feature selection for multi-cluster data. In
1103 *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*,
1104 KDD '10, pages 333–342, New York, NY, USA, July 2010. Association for Computing Machinery.

1105 [56] Leland McInnes, John Healy, and James Melville. UMAP: Uniform manifold approximation and projection
1106 for dimension reduction. *arXiv*, February 2018.

1107 [57] Andrew Y Ng, Michael I Jordan, and Yair Weiss. On spectral clustering: analysis and an algorithm. In
1108 *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and*
1109 *Synthetic*, NIPS'01, pages 849–856, Cambridge, MA, USA, January 2001. MIT Press.

1110 [58] Ofir Lindenbaum, Uri Shaham, Jonathan Svirsky, Erez Peterfreund, and Yuval Kluger. Differentiable
1111 unsupervised feature selection based on a gated laplacian. *arXiv*, July 2020.

1112 [59] Uri Shaham, Ofir Lindenbaum, Jonathan Svirsky, and Yuval Kluger. Deep unsupervised feature selection by
1113 discarding nuisance and correlated features. *Neural Netw.*, 152:34–43, August 2022.

1114 [60] Sebastian J Arnold and Elizabeth J Robertson. Making a commitment: cell lineage allocation and axis
1115 patterning in the early mouse embryo. *Nat. Rev. Mol. Cell Biol.*, 10(2):91–103, February 2009.

1116 [61] Norbert Perrimon, Chrysoula Pitsouli, and Ben-Zion Shilo. Signaling mechanisms controlling cell fate and
1117 embryonic patterning. *Cold Spring Harb. Perspect. Biol.*, 4(8):a005975, August 2012.

1118 [62] George Britton, Idse Heemskerk, Rachel Hodge, Amina A Qutub, and Aryeh Warmflash. A novel self-
1119 organizing embryonic stem cell system reveals signaling logic underlying the patterning of human ectoderm.
1120 *Development*, 146(20), October 2019.

1121 [63] Panteleimon Rompolas, Kailin R Mesa, Kyogo Kawaguchi, Sangbum Park, David Gonzalez, Samara
1122 Brown, Jonathan Boucher, Allon M Klein, and Valentina Greco. Spatiotemporal coordination of stem cell
1123 commitment during epidermal homeostasis. *Science*, 352(6292):1471–1474, June 2016.

1124 [64] Jacob H Levine, Erin F Simonds, Sean C Bendall, Kara L Davis, El-Ad D Amir, Michelle D Tadmor, Oren
1125 Litvin, Harris G Fienberg, Astraea Jager, Eli R Zunder, Rachel Finck, Amanda L Gedman, Ina Radtke,
1126 James R Downing, Dana Pe'er, and Garry P Nolan. Data-Driven phenotypic dissection of AML reveals
1127 progenitor-like cells that correlate with prognosis. *Cell*, 162(1):184–197, July 2015.

1128 [65] Emma Dann, Neil C Henderson, Sarah A Teichmann, Michael D Morgan, and John C Marioni. Differential
1129 abundance testing on single-cell data using k-nearest neighbor graphs. *Nat. Biotechnol.*, September 2021.

1130 [66] Vishal Athreya Baskaran, Jolene Ranek, Siyuan Shan, Natalie Stanley, and Junier B Oliva. Distribution-
1131 based sketching of single-cell samples. In *Proceedings of the 13th ACM International Conference on*
1132 *Bioinformatics, Computational Biology and Health Informatics*, number Article 26 in BCB '22, pages 1–10,
1133 New York, NY, USA, August 2022. Association for Computing Machinery.

1134 [67] Xiaowen Dong, Dorina Thanou, Laura Toni, Michael Bronstein, and Pascal Frossard. Graph signal processing
1135 for machine learning: A review and new perspectives. *IEEE Signal Process. Mag.*, 37(6):117–127, November
1136 2020.

1137 [68] David I Shuman, Sunil K Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. The emerging
1138 field of signal processing on graphs: Extending High-Dimensional data analysis to networks and other
1139 irregular domains. *arXiv*, October 2012.

1140 [69] Malte D Luecken and Fabian J Theis. Current best practices in single-cell RNA-seq analysis: a tutorial. *Mol.*
1141 *Syst. Biol.*, 15(6):e8746, June 2019.

1142 [70] Hemant Ishwaran. The effect of splitting on random forests. *Mach. Learn.*, 99(1):75–118, April 2015.

1143 [71] Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. Least angle regression. *aos*, 32(2):407–
1144 499, April 2004.

1145 [72] Luke Zappia, Belinda Phipson, and Alicia Oshlack. Splatter: simulation of single-cell RNA sequencing data.
1146 *Genome Biol.*, 18(1):174, September 2017.

1147 [73] Christoph Hafemeister and Rahul Satija. Normalization and variance stabilization of single-cell RNA-seq
1148 data using regularized negative binomial regression. *Genome Biol.*, 20(1):296, December 2019.

1149 [74] Valentine Svensson, Eduardo da Veiga Beltrame, and Lior Pachter. Quantifying the tradeoff between
1150 sequencing depth and cell number in single-cell RNA-seq. *bioRxiv*, page 762773, September 2019.

1151 [75] Kevin R Moon, David van Dijk, Zheng Wang, Scott Gigante, Daniel B Burkhardt, William S Chen, Kristina
1152 Yim, Antonia van den Elzen, Matthew J Hirn, Ronald R Coifman, Natalia B Ivanova, Guy Wolf, and Smita
1153 Krishnaswamy. Visualizing structure and transitions in high-dimensional biological data. *Nat. Biotechnol.*,
1154 37(12):1482–1492, December 2019.

1155 [76] Siyuan Shan, Vishal Athreya Baskaran, Haidong Yi, Jolene Ranek, Natalie Stanley, and Junier B Oliva.
1156 Transparent single-cell set classification with kernel mean embeddings. In *Proceedings of the 13th ACM*
1157 *International Conference on Bioinformatics, Computational Biology and Health Informatics*, number Article
1158 25 in BCB '22, pages 1–10, New York, NY, USA, August 2022. Association for Computing Machinery.

1159 [77] Jolene S Ranek, Natalie Stanley, and Jeremy E Purvis. Integrating temporal single-cell gene expression
1160 modalities for trajectory inference and disease prediction. *Genome Biol.*, 23(1):1–32, September 2022.

1161 [78] Robert V Bruggner, Bernd Bodenmiller, David L Dill, Robert J Tibshirani, and Garry P Nolan. Auto-
1162 mated identification of stratifying signatures in cellular subpopulations. *Proc. Natl. Acad. Sci. U. S. A.*,
1163 111(26):E2770–7, July 2014.

1164 [79] Daniel B Burkhardt, Jay S Stanley, 3rd, Alexander Tong, Ana Luisa Perdigoto, Scott A Gigante, Kevan C
1165 Herold, Guy Wolf, Antonio J Giraldez, David van Dijk, and Smita Krishnaswamy. Quantifying the effect of
1166 experimental perturbations at single-cell resolution. *Nat. Biotechnol.*, 39(5):619–629, May 2021.

1167 [80] Gabriel Torregrosa and Jordi Garcia-Ojalvo. Mechanistic models of cell-fate transitions from single-cell
1168 data. *Current Opinion in Systems Biology*, 26:79–86, June 2021.

1169 [81] Peijie Zhou, Shuxiong Wang, Tiejun Li, and Qing Nie. Dissecting transition cells from single-cell transcrip-
1170 tome data through multiscale stochastic dynamics. *Nat. Commun.*, 12(1):5609, September 2021.

1171 [82] Michael J Casey, Patrick S Stumpf, and Ben D MacArthur. Theory of cell fate. *Wiley Interdiscip. Rev. Syst.*
1172 *Biol. Med.*, 12(2):e1471, March 2020.

1173 [83] Laleh Haghverdi, Maren Büttner, F Alexander Wolf, Florian Buettnner, and Fabian J Theis. Diffusion
1174 pseudotime robustly reconstructs lineage branching. *Nat. Methods*, 13(10):845–848, October 2016.

1175 [84] Michael Angelo, Sean C Bendall, Rachel Finck, Matthew B Hale, Chuck Hitzman, Alexander D Borowsky,
1176 Richard M Levenson, John B Lowe, Scot D Liu, Shuchun Zhao, Yasodha Natkunam, and Garry P Nolan.
1177 Multiplexed ion beam imaging of human breast tumors. *Nat. Med.*, 20(4):436–442, April 2014.

1178 [85] Robert R Stickels, Evan Murray, Pawan Kumar, Jilong Li, Jamie L Marshall, Daniela J Di Bella, Paola
1179 Arlotta, Evan Z Macosko, and Fei Chen. Highly sensitive spatial transcriptomics at near-cellular resolution
1180 with Slide-seqV2. *Nat. Biotechnol.*, 39(3):313–319, March 2021.

1181 [86] Yang Liu, Mingyu Yang, Yanxiang Deng, Graham Su, Archibald Enninful, Cindy C Guo, Toma Tebaldi,
1182 Di Zhang, Dongjoo Kim, Zhiliang Bai, Eileen Norris, Alisia Pan, Jiatong Li, Yang Xiao, Stephanie Halene,
1183 and Rong Fan. High-Spatial-Resolution Multi-Omics sequencing via deterministic barcoding in tissue. *Cell*,
1184 183(6):1665–1681.e18, December 2020.

1185 [87] Chenglong Xia, Jean Fan, George Emanuel, Junjie Hao, and Xiaowei Zhuang. Spatial transcriptome profiling
1186 by MERFISH reveals subcellular RNA compartmentalization and cell cycle-dependent gene expression.
1187 *Proc. Natl. Acad. Sci. U. S. A.*, 116(39):19490–19499, September 2019.

1188 [88] Ao Chen, Sha Liao, Mengnan Cheng, Kailong Ma, Liang Wu, Yiwei Lai, Xiaojie Qiu, Jin Yang, Jiangshan
1189 Xu, Shijie Hao, Xin Wang, Huifang Lu, Xi Chen, Xing Liu, Xin Huang, Zhao Li, Yan Hong, Yujia Jiang,
1190 Jian Peng, Shuai Liu, Mengzhe Shen, Chuanyu Liu, Quanshui Li, Yue Yuan, Xiaoyu Wei, Huiwen Zheng,
1191 Weimin Feng, Zhifeng Wang, Yang Liu, Zhaojun Wang, Yunzhi Yang, Haitao Xiang, Lei Han, Baoming Qin,
1192 Pengcheng Guo, Guangyao Lai, Pura Muñoz-Cánores, Patrick H Maxwell, Jean Paul Thiery, Qing-Feng
1193 Wu, Fuxiang Zhao, Bichao Chen, Mei Li, Xi Dai, Shuai Wang, Haoyan Kuang, Junhou Hui, Liqun Wang,
1194 Ji-Feng Fei, Ou Wang, Xiaofeng Wei, Haorong Lu, Bo Wang, Shiping Liu, Ying Gu, Ming Ni, Wenwei
1195 Zhang, Feng Mu, Ye Yin, Huanming Yang, Michael Lisby, Richard J Cornall, Jan Mulder, Mathias Uhlén,
1196 Miguel A Esteban, Yuxiang Li, Longqi Liu, Xun Xu, and Jian Wang. Spatiotemporal transcriptomic atlas of
1197 mouse organogenesis using DNA nanoball-patterned arrays. *Cell*, 185(10):1777–1792.e21, May 2022.

1198 [89] T Lohoff, S Ghazanfar, A Missarova, N Koulena, N Pierson, J A Griffiths, E S Bardot, C-H L Eng, R C V
1199 Tyser, R Argelaguet, C Guibentif, S Srinivas, J Briscoe, B D Simons, A-K Hadjantonakis, B Göttgens,
1200 W Reik, J Nichols, L Cai, and J C Marioni. Integration of spatial and single-cell transcriptomic data
elucidates mouse organogenesis. *Nat. Biotechnol.*, 40(1):74–85, January 2022.

1202 [90] Tyler Risom, David R Glass, Inna Averbukh, Candace C Liu, Alex Baranski, Adam Kagel, Erin F McCaffrey,
1203 Noah F Greenwald, Belén Rivero-Gutiérrez, Siri H Strand, Sushama Varma, Alex Kong, Leeat Keren,
1204 Sucheta Srivastava, Chunfang Zhu, Zumana Khair, Deborah J Veis, Katherine Deschryver, Sujay Vennam,
1205 Carlo Maley, E Shelley Hwang, Jeffrey R Marks, Sean C Bendall, Graham A Colditz, Robert B West, and
1206 Michael Angelo. Transition to invasive breast cancer is associated with progressive changes in the structure
1207 and composition of tumor stroma. *Cell*, 185(2):299–310.e18, January 2022.

1208 [91] Erin F McCaffrey, Michele Donato, Leeat Keren, Zhenghao Chen, Alea Delmastro, Megan B Fitzpatrick,
1209 Sanjana Gupta, Noah F Greenwald, Alex Baranski, William Graf, Rashmi Kumar, Marc Bosse, Chris-
1210 tine Camacho Fullaway, Pratista K Ramdial, Erna Forgó, Vladimir Jovic, David Van Valen, Smriti Mehra,
1211 Shabaana A Khader, Sean C Bendall, Matt van de Rijn, Daniel Kalman, Deepak Kaushal, Robert L Hunter,
1212 Niaz Banaei, Adrie J C Steyn, Purvesh Khatri, and Michael Angelo. The immunoregulatory landscape of
1213 human tuberculosis granulomas. *Nat. Immunol.*, 23(2):318–329, February 2022.

1214 [92] Wayne Stallaert, Katarzyna M Kedziora, Colin D Taylor, Tarek M Zikry, Jolene S Ranek, Holly K Sobon,
1215 Sovanny R Taylor, Catherine L Young, Jeanette G Cook, and Jeremy E Purvis. The structure of the human
1216 cell cycle. *Cell Syst.*, 13(1):103, January 2022.

1217 [93] Samuel Hume, Grigory L Dianov, and Kristijan Ramadan. A unified model for the G1/S cell cycle transition.
1218 *Nucleic Acids Res.*, 48(22):12483–12501, December 2020.

1219 [94] A C Carrano, E Eytan, A Hershko, and M Pagano. SKP2 is required for ubiquitin-mediated degradation of
1220 the CDK inhibitor p27. *Nat. Cell Biol.*, 1(4):193–199, August 1999.

1221 [95] L Zhang and C Wang. F-box protein skp2: a novel transcriptional target of E2F. *Oncogene*, 25(18):2615–
1222 2627, April 2006.

1223 [96] Jeroen Essers, Arjan F Theil, Céline Baldeyron, Wiggert A van Cappellen, Adriaan B Houtsmuller, Roland
1224 Kanaar, and Wim Vermeulen. Nuclear dynamics of PCNA in DNA replication and repair. *Mol. Cell. Biol.*,
1225 25(21):9350–9359, November 2005.

1226 [97] Simran Khurana and Philipp Oberdoerffer. Replication stress: A lifetime of epigenetic change. *Genes*,
1227 6(3):858–877, September 2015.

1228 [98] Bianca M Sirbu, Frank B Couch, Jordan T Feigerle, Srividya Bhaskara, Scott W Hiebert, and David
1229 Cortez. Analysis of protein dynamics at active, stalled, and collapsed replication forks. *Genes Dev.*,
1230 25(12):1320–1327, June 2011.

1231 [99] Arne Lindqvist, Verónica Rodríguez-Bravo, and René H Medema. The decision to enter mitosis: feedback
1232 and redundancy in the mitotic entry network. *J. Cell Biol.*, 185(2):193–202, April 2009.

1233 [100] Olivier Gavet and Jonathon Pines. Activation of cyclin B1-Cdk1 synchronizes events in the nucleus and the
1234 cytoplasm at mitosis. *J. Cell Biol.*, 189(2):247–259, April 2010.

1235 [101] Justin Moser, Iain Miller, Dylan Carter, and Sabrina L Spencer. Control of the restriction point by rb and
1236 p21. *Proc. Natl. Acad. Sci. U. S. A.*, 115(35):E8219–E8227, August 2018.

1237 [102] R A Weinberg. The retinoblastoma protein and cell cycle control. *Cell*, 81(3):323–330, May 1995.

1238 [103] David Arthur and Sergei Vassilvitskii. k-means++: the advantages of careful seeding. In *Proceedings of*
1239 *the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '07, pages 1027–1035, USA,
1240 January 2007. Society for Industrial and Applied Mathematics.

1241 [104] F Alexander Wolf, Fiona K Hamey, Mireya Plass, Jordi Solana, Joakim S Dahlin, Berthold Göttgens,
1242 Nikolaus Rajewsky, Lukas Simon, and Fabian J Theis. PAGA: graph abstraction reconciles clustering with
1243 trajectory inference through a topology preserving map of single cells. *Genome Biol.*, 20(1):59, March 2019.

1244 [105] Xiaojie Qiu, Qi Mao, Ying Tang, Li Wang, Raghav Chawla, Hannah A Pliner, and Cole Trapnell. Reversed
1245 graph embedding resolves complex single-cell trajectories. *Nat. Methods*, 14(10):979–982, October 2017.

1246 [106] Nadia S Kurd, Zhaoren He, Tiani L Louis, J Justin Milner, Kyla D Omilusik, Wenhao Jin, Matthew S Tsai,
1247 Christella E Widjaja, Jad N Kanbar, Jocelyn G Olvera, Tiffani Tysl, Lauren K Quezada, Brigid S Boland,
1248 Wendy J Huang, Cornelis Murre, Ananda W Goldrath, Gene W Yeo, and John T Chang. Early precursors
1249 and molecular determinants of tissue-resident memory CD8⁺ T lymphocytes revealed by single-cell RNA
1250 sequencing. *Sci Immunol*, 5(47), May 2020.

1251 [107] E John Wherry and Rafi Ahmed. Memory CD8 t-cell differentiation during viral infection. *J. Virol.*,
1252 78(11):5535–5545, June 2004.

1253 [108] Susan M Kaech, E John Wherry, and Rafti Ahmed. Effector and memory t-cell differentiation: implications
1254 for vaccine development. *Nat. Rev. Immunol.*, 2(4):251–262, April 2002.

1255 [109] E John Wherry, Volker Teichgräber, Todd C Becker, David Masopust, Susan M Kaech, Rustom Antia,
1256 Ulrich H von Andrian, and Rafi Ahmed. Lineage relationship and protective immunity of memory CD8 T
1257 cell subsets. *Nat. Immunol.*, 4(3):225–234, March 2003.

1258 [110] Amanda L Blasius, Emanuele Giurisato, Marina Celli, Robert D Schreiber, Andrey S Shaw, and Marco
1259 Colonna. Bone marrow stromal cell antigen 2 is a specific marker of type I IFN-producing cells in the naïve
1260 mouse, but a promiscuous cell surface antigen following IFN stimulation. *J. Immunol.*, 177(5):3260–3265,
1261 September 2006.

1262 [111] Mladen Jergović, Christopher P Coplen, Jennifer L Uhrlaub, David G Besselsen, Shu Cheng, Megan J
1263 Smithey, and Janko Nikolich-Žugich. Infection-induced type I interferons critically modulate the homeostasis
1264 and function of CD8+ naïve T cells. *Nat. Commun.*, 12(1):5303, September 2021.

1265 [112] Itay Tirosh, Benjamin Izar, Sanjay M Prakadan, Marc H Wadsworth, 2nd, Daniel Treacy, John J Trombetta,
1266 Asaf Rotem, Christopher Rodman, Christine Lian, George Murphy, Mohammad Fallahi-Sichani, Ken Dutton-
1267 Regester, Jia-Ren Lin, Ofir Cohen, Parin Shah, Diana Lu, Alex S Genshaft, Travis K Hughes, Carly G K
1268 Ziegler, Samuel W Kazer, Aleth Gaillard, Kellie E Kolb, Alexandra-Chloé Villani, Cory M Johannessen,
1269 Aleksandr Y Andreev, Eliezer M Van Allen, Monica Bertagnolli, Peter K Sorger, Ryan J Sullivan, Keith T
1270 Flaherty, Dennie T Frederick, Judit Jané-Valbuena, Charles H Yoon, Orit Rozenblatt-Rosen, Alex K Shalek,
1271 Aviv Regev, and Levi A Garraway. Dissecting the multicellular ecosystem of metastatic melanoma by
1272 single-cell RNA-seq. *Science*, 352(6282):189–196, April 2016.

1273 [113] Pratip K Chattopadhyay, Michael R Betts, David A Price, Emma Gostick, Helen Horton, Mario Roederer, and
1274 Stephen C De Rosa. The cytolytic enzymes granzyme a, granzyme b, and perforin: expression patterns, cell
1275 distribution, and their relationship to cell maturity and bright CD57 expression. *J. Leukoc. Biol.*, 85(1):88–97,
1276 January 2009.

1277 [114] Kyla D Omilusik, Marija S Nadjsombati, Laura A Shaw, Bingfei Yu, J Justin Milner, and Ananda W Goldrath.
1278 Sustained id2 regulation of E proteins is required for terminal differentiation of effector CD8+ T cells. *J.
1279 Exp. Med.*, 215(3):773–783, March 2018.

1280 [115] J Justin Milner, Hongtuyet Nguyen, Kyla Omilusik, Miguel Reina-Campos, Matthew Tsai, Clara Toma,
1281 Arnaud Delpoux, Brigid S Boland, Stephen M Hedrick, John T Chang, and Ananda W Goldrath. Delineation
1282 of a molecularly distinct terminally differentiated memory CD8 T cell population. *Proc. Natl. Acad. Sci. U.
1283 S. A.*, 117(41):25667–25678, October 2020.

1284 [116] Susan M Kaech and Weiguo Cui. Transcriptional control of effector and memory CD8+ T cell differentiation.
1285 *Nat. Rev. Immunol.*, 12(11):749–761, November 2012.

1286 [117] John T Harty and Vladimir P Badovinac. Shaping and reshaping CD8+ t-cell memory. *Nat. Rev. Immunol.*,
1287 8(2):107–119, February 2008.

1288 [118] J M Grayson, A J Zajac, J D Altman, and R Ahmed. Cutting edge: increased expression of bcl-2 in
1289 antigen-specific memory CD8+ T cells. *J. Immunol.*, 164(8):3950–3954, April 2000.

1290 [119] Susan M Kaech, Joyce T Tan, E John Wherry, Bogumila T Konieczny, Charles D Surh, and Rafi Ahmed.
1291 Selective expression of the interleukin 7 receptor identifies effector CD8 T cells that give rise to long-lived
1292 memory cells. *Nat. Immunol.*, 4(12):1191–1198, December 2003.

1293 [120] Vaibhav Upadhyay and Yang-Xin Fu. Lymphotoxin signalling in immune homeostasis and the control of
1294 microorganisms. *Nat. Rev. Immunol.*, 13(4):270–279, April 2013.

1295 [121] Damian Szklarczyk, Annika L Gable, Katerina C Nastou, David Lyon, Rebecca Kirsch, Sampo Pyysalo,
1296 Nadezhda T Doncheva, Marc Legeay, Tao Fang, Peer Bork, Lars J Jensen, and Christian von Mering.
1297 The STRING database in 2021: customizable protein-protein networks, and functional characterization of
1298 user-uploaded gene/measurement sets. *Nucleic Acids Res.*, 49(D1):D605–D612, January 2021.

1299 [122] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: A practical and powerful approach
1300 to multiple testing. *J. R. Stat. Soc.*, 57(1):289–300, January 1995.

1301 [123] Maxim V Kuleshov, Matthew R Jones, Andrew D Rouillard, Nicolas F Fernandez, Qiaonan Duan, Zichen
1302 Wang, Simon Koplev, Sherry L Jenkins, Kathleen M Jagodnik, Alexander Lachmann, Michael G McDermott,
1303 Caroline D Monteiro, Gregory W Gundersen, and Avi Ma'ayan. Enrichr: a comprehensive gene set
1304 enrichment analysis web server 2016 update. *Nucleic Acids Res.*, 44(W1):W90–7, July 2016.

1305 [124] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *J. Mach. Learn. Res.*,
1306 9(86):2579–2605, 2008.

1307 [125] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps and spectral techniques for embedding and
1308 clustering. In *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic*, NIPS'01, pages 585–591, Cambridge, MA, USA, January 2001. MIT Press.

1310 [126] Kyle M Loh, Angela Chen, Pang Wei Koh, Tianda Z Deng, Rahul Sinha, Jonathan M Tsai, Amira A Barkal,
1311 Kimberle Y Shen, Rajan Jain, Rachel M Morganti, Ng Shyh-Chang, Nathaniel B Fernhoff, Benson M George,
1312 Gerlinde Wernig, Rachel E A Salomon, Zhenghao Chen, Hannes Vogel, Jonathan A Epstein, Anshul Kundaje,
1313 William S Talbot, Philip A Beachy, Lay Teng Ang, and Irving L Weissman. Mapping the pairwise choices
1314 leading from pluripotency to human bone, heart, and other mesoderm cell types. *Cell*, 166(2):451–467, July
1315 2016.

1316 [127] Peter V Kharchenko, Lev Silberstein, and David T Scadden. Bayesian approach to single-cell differential
1317 expression analysis. *Nat. Methods*, 11(7):740–742, July 2014.

1318 [128] Emma Pierson and Christopher Yau. ZIFA: Dimensionality reduction for zero-inflated single-cell gene
1319 expression analysis. *Genome Biol.*, 16:241, November 2015.

1320 [129] W Evan Johnson, Cheng Li, and Ariel Rabinovic. Adjusting batch effects in microarray expression data
1321 using empirical bayes methods. *Biostatistics*, 8(1):118–127, January 2007.

1322 [130] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Mach. Learn.*, 20(3):273–297, September
1323 1995.

1324 [131] Hanneke van der Hoef and Matthijs J Warrens. Understanding information theoretic measures for comparing
1325 clusterings. *Behaviormetrika*, 46(2):353–370, October 2019.

1326 [132] Belinda Phipson and Gordon K Smyth. Permutation p-values should never be zero: calculating exact p-values
1327 when permutations are randomly drawn. *Stat. Appl. Genet. Mol. Biol.*, 9:Article39, October 2010.

1328 [133] Wayne Stallaert, Katarzyna M Kedziora, Colin D Taylor, Tarek M Zikry, Jolene S Ranek, Holly K Sobon,
1329 Sovanny R Taylor, Catherine L Young, Jeanette G Cook, and Jeremy E Purvis. The structure of the human
1330 cell cycle. Datasets. Zenodo Repository. <https://doi.org/10.5281/zenodo.4525425> (2022).

1331 [134] Wayne Stallaert, Bjoern Papke, Channing Der, and Jeremy E Purvis. Cell cycle heterogeneity in pancreatic
1332 ductal adenocarcinoma. Datasets. Zenodo Repository. <https://doi.org/10.5281/zenodo.7860332>
1333 (2023).

1334 [135] Nadia S Kurd, Zhaoren He, Tiani L Louis, J Justin Milner, Kyla D Omilusik, Wenhao Jin, Matthew S Tsai,
1335 Christella E Widjaja, Jad N Kanbar, Jocelyn G Olvera, Tiffani Tysl, Lauren K Quezada, Brigid S Boland,
1336 Wendy J Huang, Cornelis Murre, Ananda W Goldrath, Gene W Yeo, and John T Chang. Early precursors
1337 and molecular determinants of tissue-resident memory CD8⁺ T lymphocytes revealed by single-cell RNA
1338 sequencing. Datasets. Gene Expression Omnibus. <https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE131847> (2020).

1339 [136] Jolene Ranek, Wayne Stallaert, Justin Milner, Natalie Stanley, and Jeremy Purvis. Feature selection for
1340 preserving biological trajectories in single-cell data. Datasets. Zenodo Repository. <https://doi.org/10.5281/zenodo.7883604> (2023).

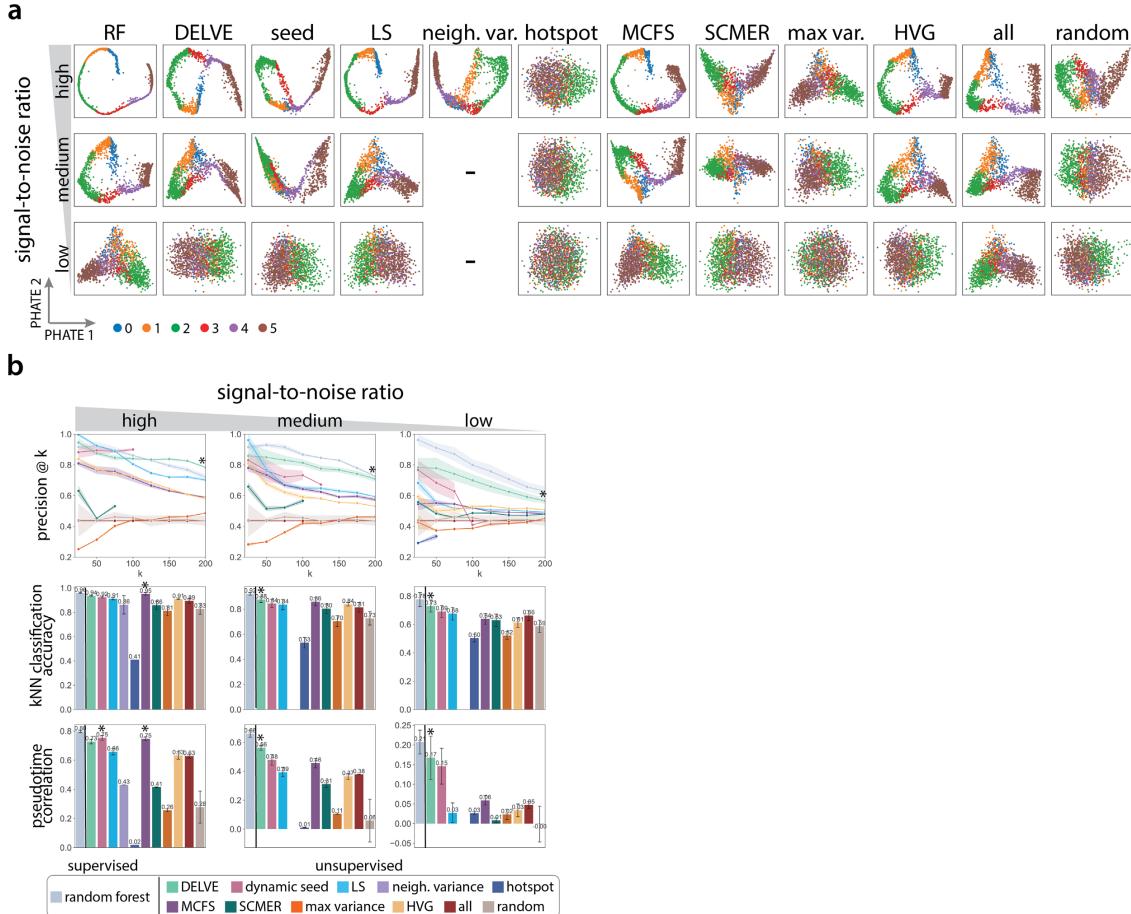
1343 **Supplementary Information**

1344 **Supplementary Tables**

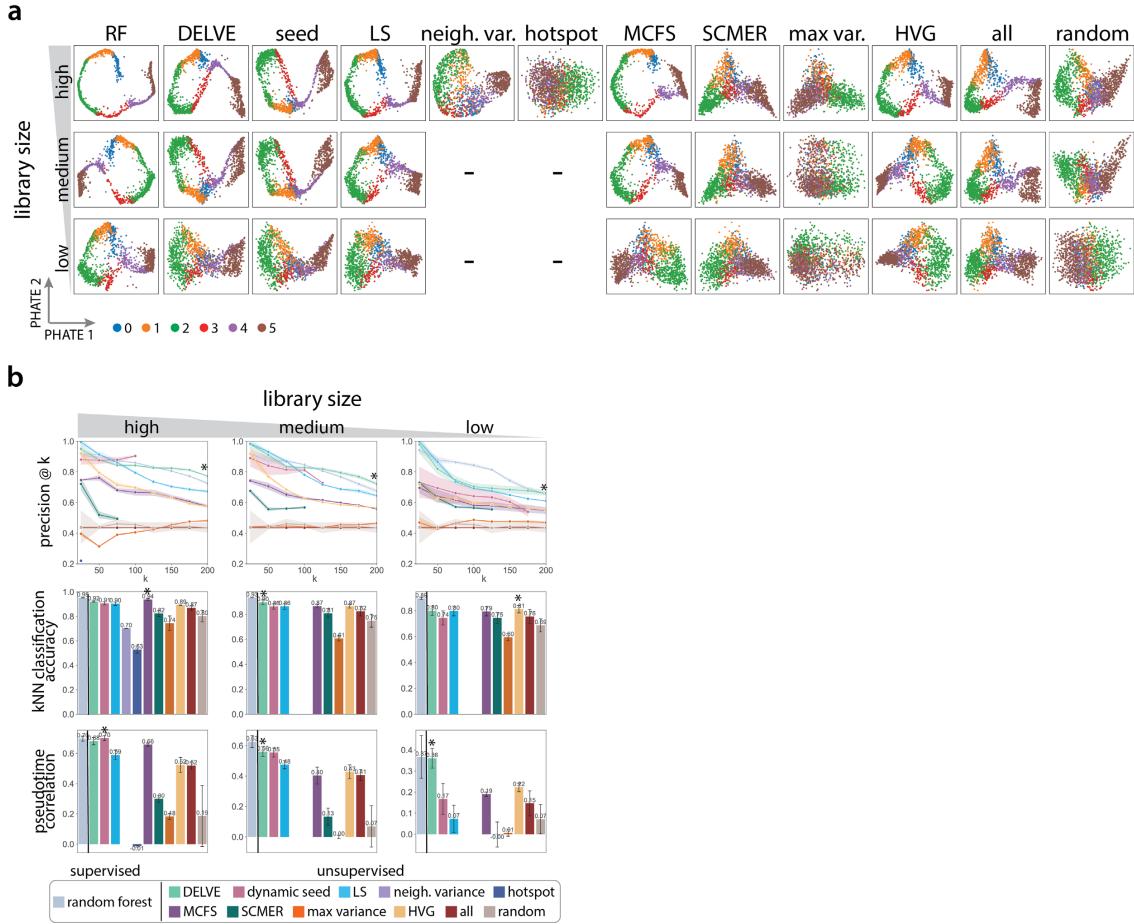
Supplementary Table 1: Implemented feature selection method parameters

Method name	Supervision	Parameter description	Parameters
Random Forest classifier	supervised	cell type labels number of trees	cell type labels n_estimators = 10, 100, 500
DELVE	unsupervised	number of nearest neighbors number of representative neighborhoods number of modules	$k = 10$ $m = 1000$ $c = 3, 5, 10$
Laplacian score	unsupervised	number of nearest neighbors kernel bandwidth parameter	$k = 10$ $\sigma_i = 3\text{rd-nearest neighbor distance}$
Neighborhood variance	unsupervised	NA	NA
Hotspot	unsupervised	number of nearest neighbors kernel bandwidth parameter number of principal components model	$k = 10$ $\sigma_i = k/3\text{-nearest neighbor distance}$ $n_pcs = 50$ danb
MCFS	unsupervised	number of nearest neighbors kernel bandwidth parameter number of nonzero coefficients number of eigenvectors	$k = 10$ $\sigma_i = 3\text{rd-nearest neighbor distance}$ p $c = \text{number of known cell types}$
SCMER	unsupervised	number of nearest neighbors number of principal components lasso regression parameter ridge regression parameter	$k = 10$ $n_pcs = 50$ $\text{lasso} = 3.87e - 4$ $\text{ridge} = 0$
Max variance	unsupervised	NA	NA
Highly variable genes	unsupervised	flavor bins n_top_genes	Seurat 20 p
All features	NA	NA	NA
Random features	NA	NA	NA

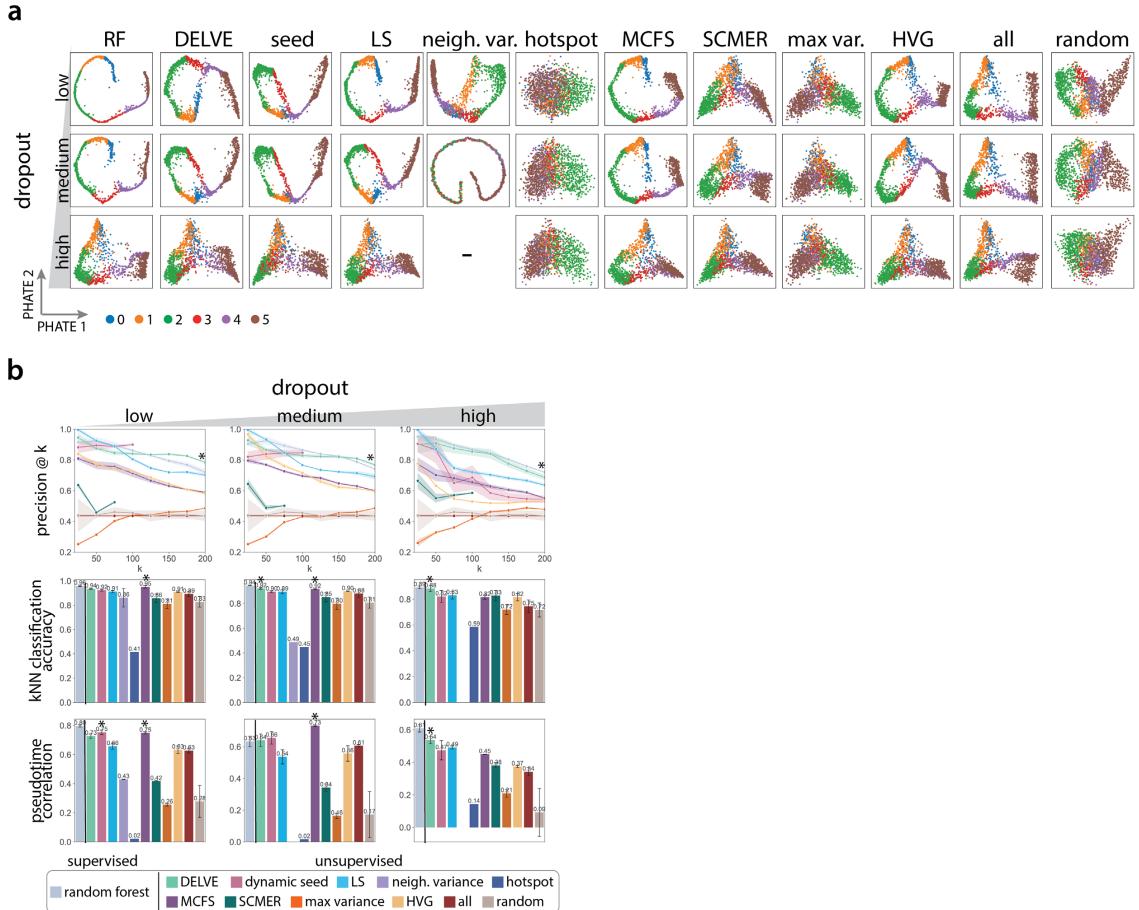
1345 **Supplementary Figures**



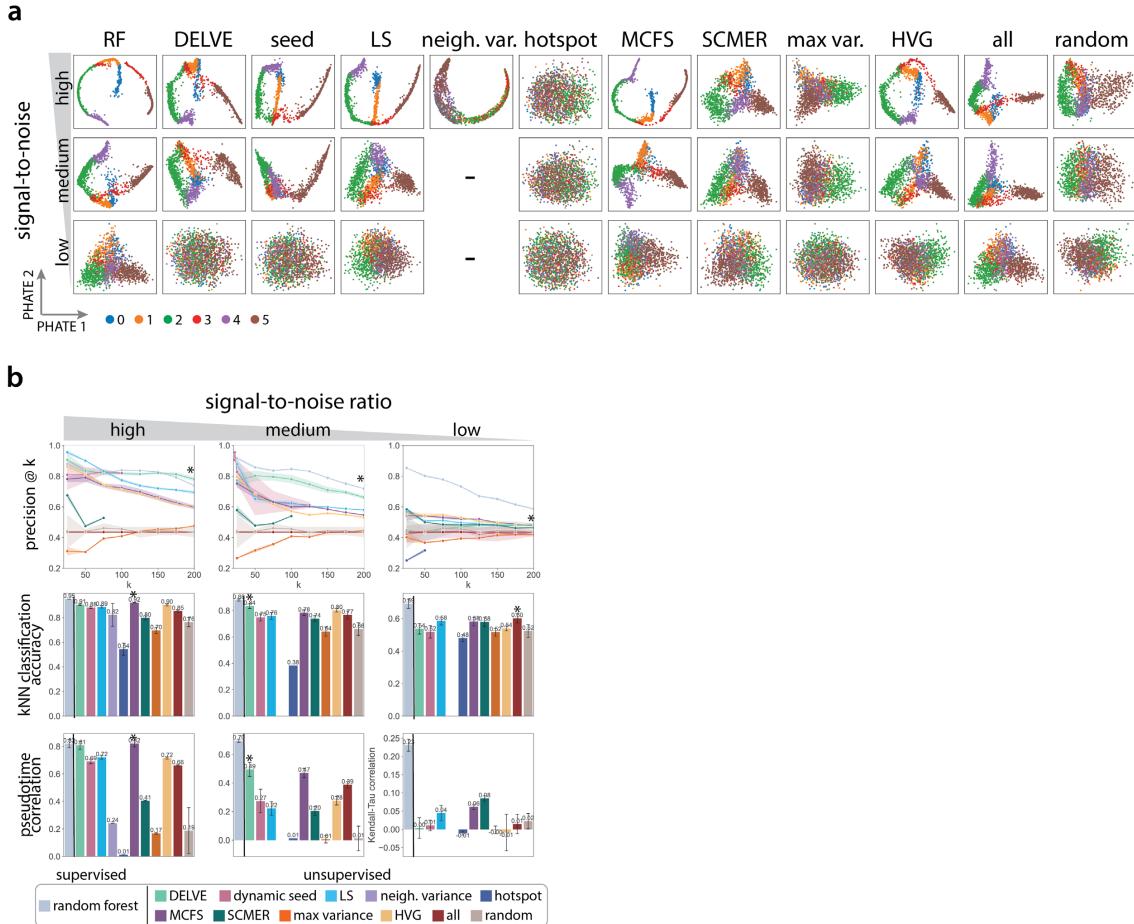
Supplementary Figure 1: Comparison of feature selection methods on preserving linear differentiation trajectories under a reduction in the signal-to-noise ratio. (a) Example PHATE [75] visualizations of simulated linear differentiation trajectories for twelve feature selection strategies when subjected to a reduction in the signal-to-noise ratio (high, medium, low). The signal-to-noise ratio was altered by modifying the biological coefficient of variation parameter within Splatter (high: BCV = 0.1, medium: BCV = 0.25, low: BCV = 0.5). This scaling factor controls the mean-variance relationship between genes, where lowly expressed genes are more variable than highly expressed genes. d indicates the total number of genes ($d = 500$) and p indicates the number of selected genes following feature selection ($p = 100$). (b) Performance of twelve different feature selection methods when subjected to a reduction in the signal-to-noise ratio. Following feature selection, trajectory preservation was quantitatively assessed according to several metrics: the precision of differentially expressed genes at k selected genes (top), k -NN classification accuracy (middle), and pseudotime correlation (bottom) across 10 random trials. Error bars/ bands represent the standard deviation. * indicates the method with the highest median score. - indicates that the method identified no features.



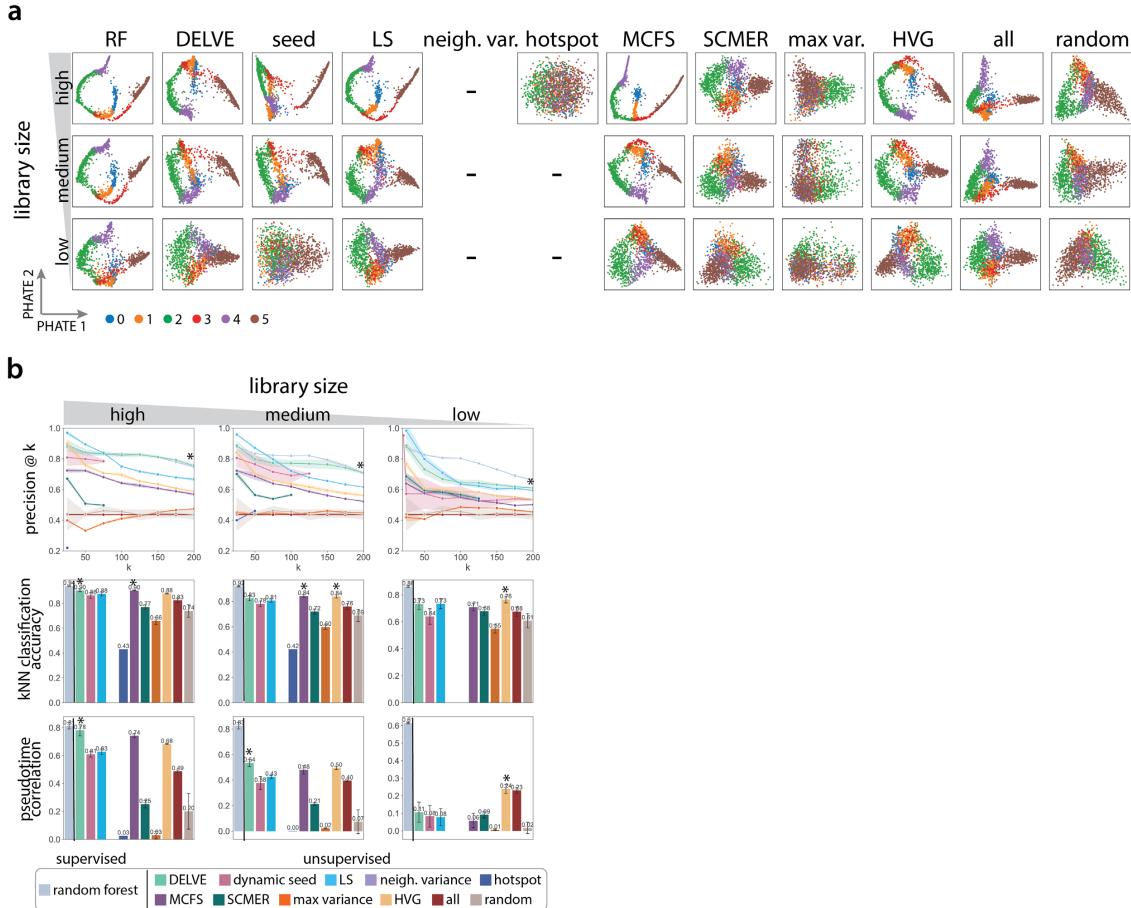
Supplementary Figure 2: Comparison of feature selection methods on preserving linear differentiation trajectories in the presence of library size noise. (a) Example PHATE [75] visualizations of simulated linear differentiation trajectories for twelve feature selection strategies when subjected to a reduction in the total mRNA count (high, medium, low). Library size was reduced by modifying the location parameter in the log-normal distribution in Splatter [72] that specifies library size scaling factors (high: location = 12, medium: location = 11, low: location = 10). d indicates the total number of genes ($d = 500$) and p indicates the number of selected genes following feature selection ($p = 100$). (b) Performance of twelve different feature selection methods when subjected to a reduction in total mRNA count. Following feature selection, trajectory preservation was quantitatively assessed according to several metrics: the precision of differentially expressed genes at k selected genes (top), k -NN classification accuracy (middle), and pseudotime correlation (bottom) across 10 random trials. Error bars/bands represent the standard deviation. * indicates the method with the highest median score. - indicates that the method identified no features.



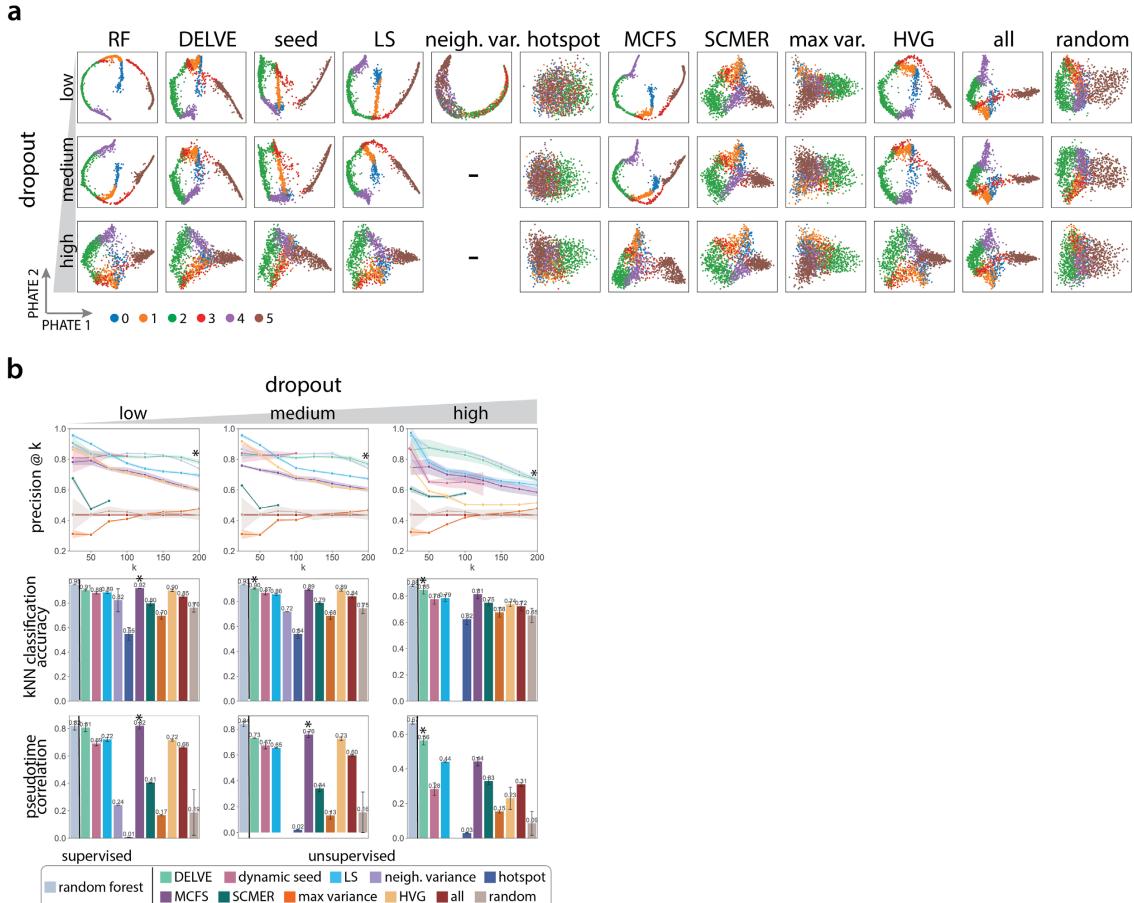
Supplementary Figure 3: Comparison of feature selection methods on preserving linear differentiation trajectories in the presence of dropout noise. (a) Example PHATE [75] visualizations of simulated linear differentiation trajectories for twelve feature selection strategies when subjected to an increase in the amount of dropout (low, medium, high). Technical dropout was simulated by undersampling mRNA counts by sampling from a binomial distribution with the scale parameter or dropout rate proportional to the mean expression of each gene (low: $\lambda = 0$, medium: $\lambda = 0.05$, high: $\lambda = 0.1$). d indicates the total number of genes ($d = 500$) and p indicates the number of selected genes following feature selection ($p = 100$). (b) Performance of twelve different feature selection methods when subjected to an increase in the amount of dropout noise. Following feature selection, trajectory preservation was quantitatively assessed according to several metrics: the precision of differentially expressed genes at k selected genes (top), k -NN classification accuracy (middle), and pseudotime correlation (bottom) across 10 random trials. Error bars/ bands represent the standard deviation. * indicates the method with the highest median score. - indicates that the method identified no features.



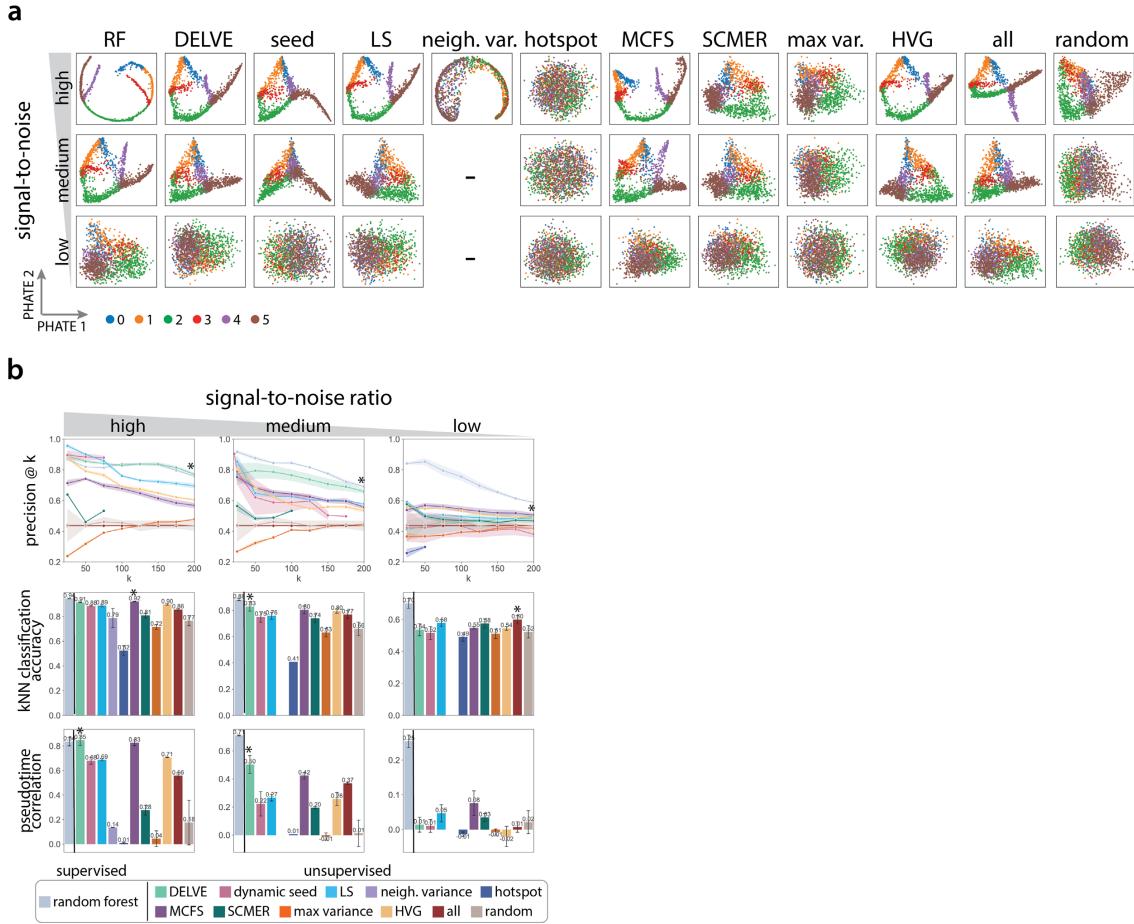
Supplementary Figure 4: Comparison of feature selection methods on preserving bifurcation differentiation trajectories under a reduction in the signal-to-noise ratio. (a) Example PHATE [75] visualizations of simulated bifurcation differentiation trajectories for twelve feature selection strategies when subjected to a reduction in the signal-to-noise ratio (high, medium, low). The signal-to-noise ratio was altered by modifying the biological coefficient of variation parameter within Splatter (high: BCV = 0.1, medium: BCV = 0.25, low: BCV = 0.5). This scaling factor controls the mean-variance relationship between genes, where lowly expressed genes are more variable than highly expressed genes. d indicates the total number of genes ($d = 500$) and p indicates the number of selected genes following feature selection ($p = 100$). (b) Performance of twelve different feature selection methods when subjected to a reduction in the signal-to-noise ratio. Following feature selection, trajectory preservation was quantitatively assessed according to several metrics: the precision of differentially expressed genes at k selected genes (top), k -NN classification accuracy (middle), and pseudotime correlation (bottom) across 10 random trials. Error bars/ bands represent the standard deviation. * indicates the method with the highest median score. - indicates that the method identified no features.



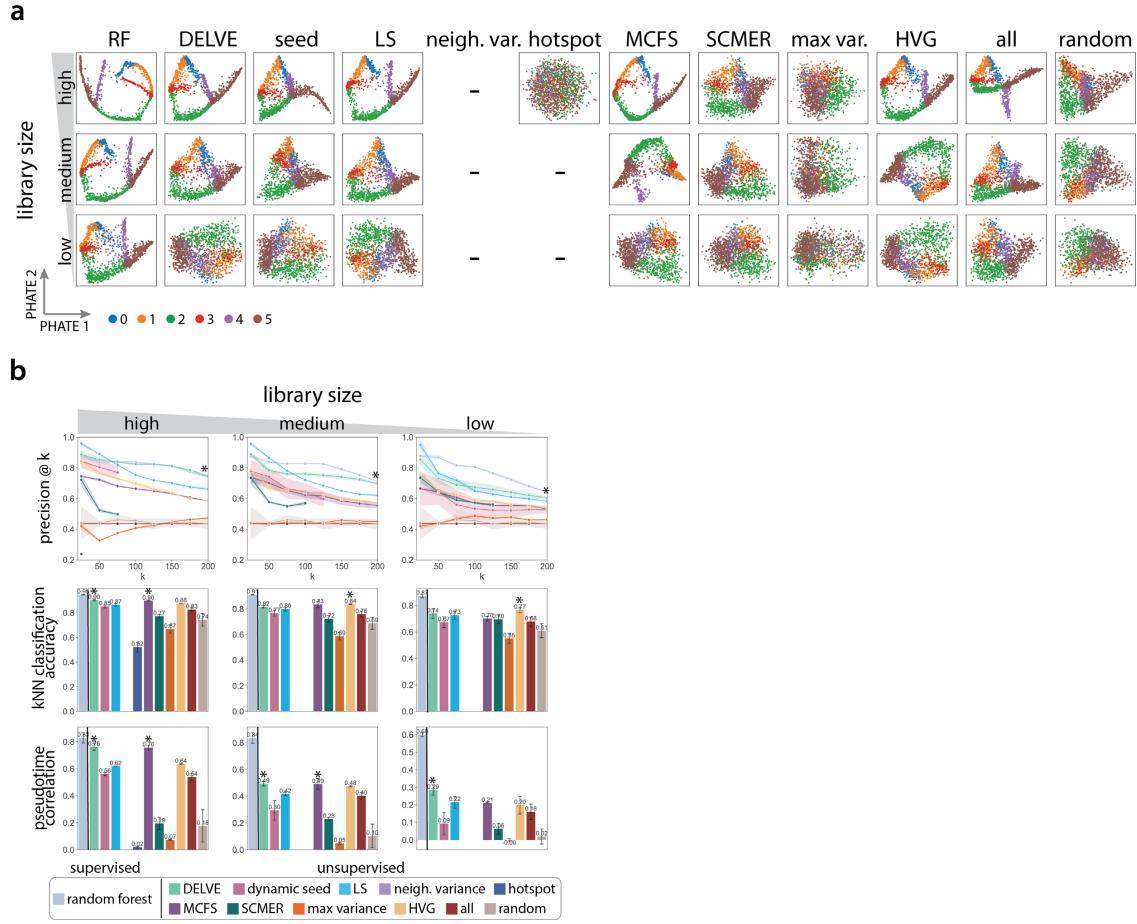
Supplementary Figure 5: Comparison of feature selection methods on preserving bifurcation differentiation trajectories in the presence of library size noise. (a) Example PHATE [75] visualizations of simulated bifurcation differentiation trajectories for twelve feature selection strategies when subjected to a reduction in the total mRNA count (high, medium, low). Library size was reduced by modifying the location parameter in the log-normal distribution in Splatter [72] that specifies library size scaling factors (high: location = 12, medium: location = 11, low: location = 10). d indicates the total number of genes ($d = 500$) and p indicates the number of selected genes following feature selection ($p = 100$). (b) Performance of twelve different feature selection methods when subjected to a reduction in total mRNA count. Following feature selection, trajectory preservation was quantitatively assessed according to several metrics: the precision of differentially expressed genes at k selected genes (top), k -NN classification accuracy (middle), and pseudotime correlation (bottom) across 10 random trials. Error bars/bands represent the standard deviation. * indicates the method with the highest median score. - indicates that the method identified no features.



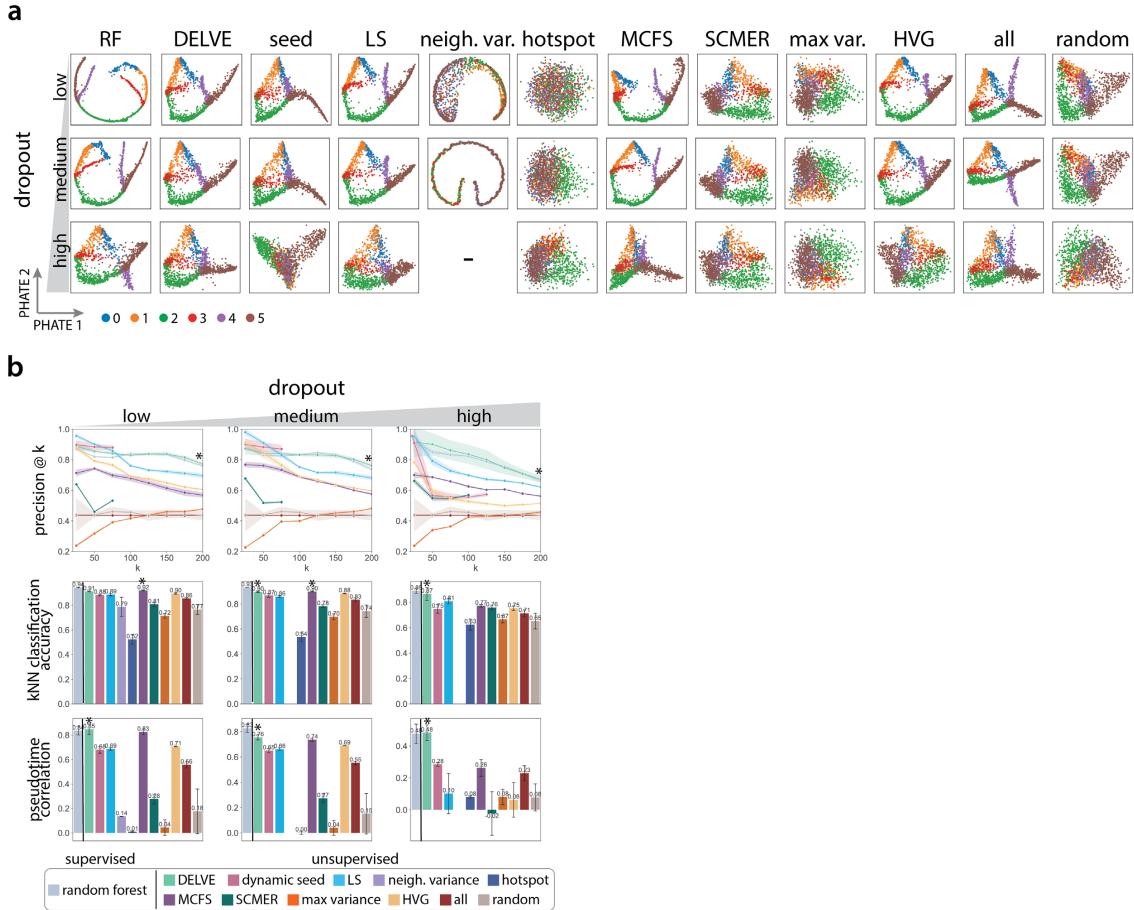
Supplementary Figure 6: Comparison of feature selection methods on preserving bifurcation differentiation trajectories in the presence of dropout noise. (a) Example PHATE [75] visualizations of simulated bifurcation differentiation trajectories for twelve feature selection strategies when subjected to an increase in the amount of dropout (low, medium, high). Technical dropout was simulated by undersampling mRNA counts by sampling from a binomial distribution with the scale parameter or dropout rate proportional to the mean expression of each gene (low: $\lambda = 0$, medium: $\lambda = 0.05$, high: $\lambda = 0.1$). d indicates the total number of genes ($d = 500$) and p indicates the number of selected genes following feature selection ($p = 100$). (b) Performance of twelve different feature selection methods when subjected to an increase in the amount of dropout noise. Following feature selection, trajectory preservation was quantitatively assessed according to several metrics: the precision of differentially expressed genes at k selected genes (top), k -NN classification accuracy (middle), and pseudotime correlation (bottom) across 10 random trials. Error bars/bands represent the standard deviation. * indicates the method with the highest median score. - indicates that the method identified no features.



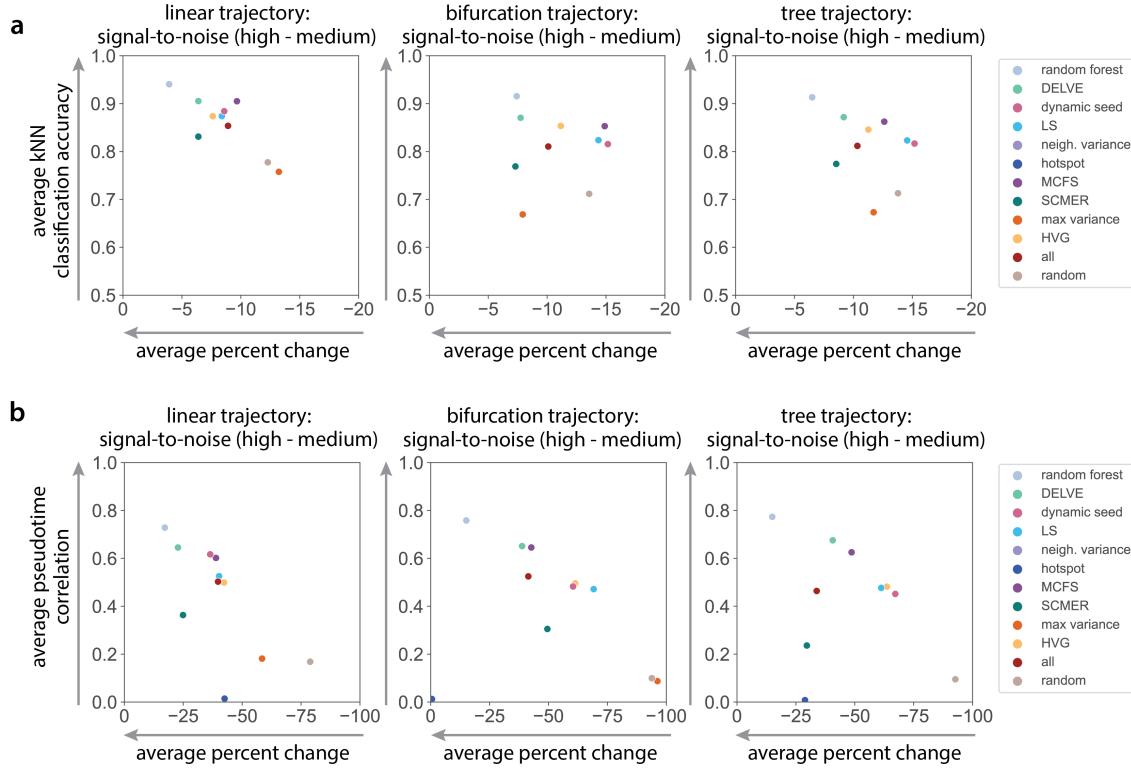
Supplementary Figure 7: Comparison of feature selection methods on preserving tree differentiation trajectories under a reduction in the signal-to-noise ratio. (a) Example PHATE [75] visualizations of simulated tree differentiation trajectories for twelve feature selection strategies when subjected to a reduction in the signal-to-noise ratio (high, medium, low). The signal-to-noise ratio was altered by modifying the biological coefficient of variation parameter within Splatter (high: BCV = 0.1, medium: BCV = 0.25, low: BCV = 0.5). This scaling factor controls the mean-variance relationship between genes, where lowly expressed genes are more variable than highly expressed genes. d indicates the total number of genes ($d = 500$) and p indicates the number of selected genes following feature selection ($p = 100$). (b) Performance of twelve different feature selection methods when subjected to a reduction in the signal-to-noise ratio. Following feature selection, trajectory preservation was quantitatively assessed according to several metrics: the precision of differentially expressed genes at k selected genes (top), k -NN classification accuracy (middle), and pseudotime correlation (bottom) across 10 random trials. Error bars/ bands represent the standard deviation. * indicates the method with the highest median score. - indicates that the method identified no features.



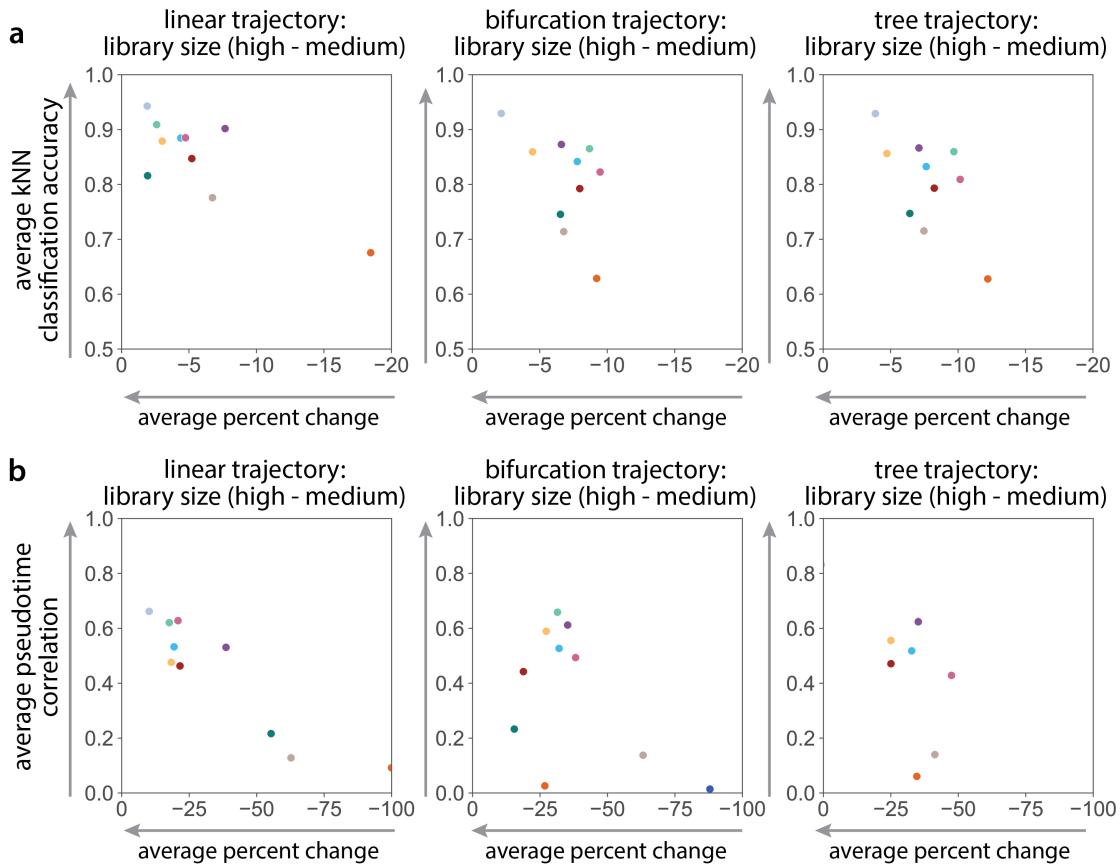
Supplementary Figure 8: Comparison of feature selection methods on preserving tree differentiation trajectories in the presence of library size noise. (a) Example PHATE [75] visualizations of simulated tree differentiation trajectories for twelve feature selection strategies when subjected to a reduction in the total mRNA count (high, medium, low). Library size was reduced by modifying the location parameter in the log-normal distribution in Splatter [72] that specifies library size scaling factors (high: location = 12, medium: location = 11, low: location = 10). d indicates the total number of genes ($d = 500$) and p indicates the number of selected genes following feature selection ($p = 100$). (b) Performance of twelve different feature selection methods when subjected to a reduction in total mRNA count. Following feature selection, trajectory preservation was quantitatively assessed according to several metrics: the precision of differentially expressed genes at k selected genes (top), k -NN classification accuracy (middle), and pseudotime correlation (bottom) across 10 random trials. Error bars/bands represent the standard deviation. * indicates the method with the highest median score. - indicates that the method identified no features.



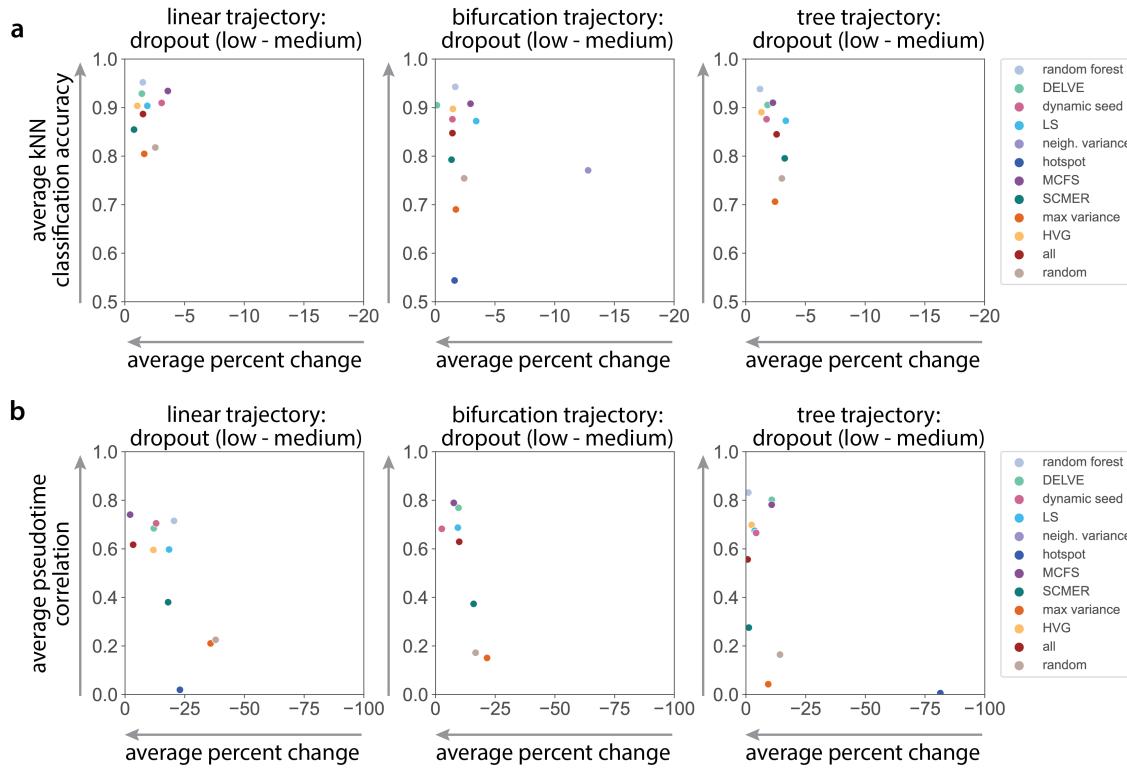
Supplementary Figure 9: Comparison of feature selection methods on preserving tree differentiation trajectories is the presence of dropout noise. (a) Example PHATE [75] visualizations of simulated tree differentiation trajectories for twelve feature selection strategies when subjected to an increase in the amount of dropout (low, medium, high). Technical dropout was simulated by undersampling mRNA counts by sampling from a binomial distribution with the scale parameter or dropout rate proportional to the mean expression of each gene (low: $\lambda = 0$, medium: $\lambda = 0.05$, low: $\lambda = 0.1$). d indicates the total number of genes ($d = 500$) and p indicates the number of selected genes following feature selection ($p = 100$). (b) Performance of twelve different feature selection methods when subjected to an increase in the amount of dropout noise. Following feature selection, trajectory preservation was quantitatively assessed according to several metrics: the precision of differentially expressed genes at k selected genes (top), k -NN classification accuracy (middle), and pseudotime correlation (bottom) across 10 random trials. Error bars/bands represent the standard deviation. * indicates the method with the highest median score. - indicates that the method identified no features.



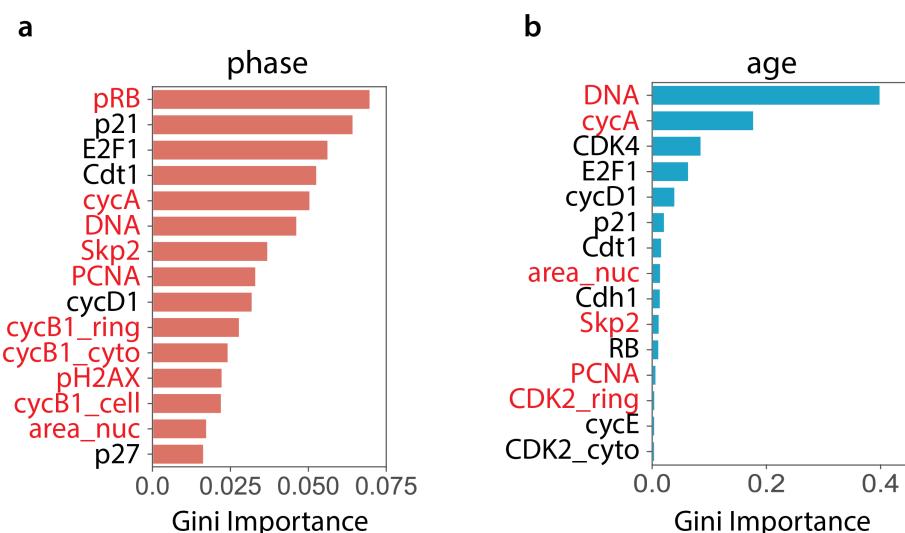
Supplementary Figure 10: Comparison of the robustness of feature selection methods on inferring differentiation trajectories under a reduction in the signal-to-noise ratio. (a) Average k -NN classification accuracy vs. average percent change in k -NN classification accuracy as the signal-to-noise ratio decreased (high to medium) and the mean-variance relationship amongst genes increased. The signal-to-noise ratio was altered by modifying the biological coefficient of variation parameter within Splatter [72] (high: BCV = 0.1, medium: BCV = 0.25). (b) Average pseudotime correlation vs. average percent change in pseudotime correlation as the signal-to-noise ratio decreased (high to medium).



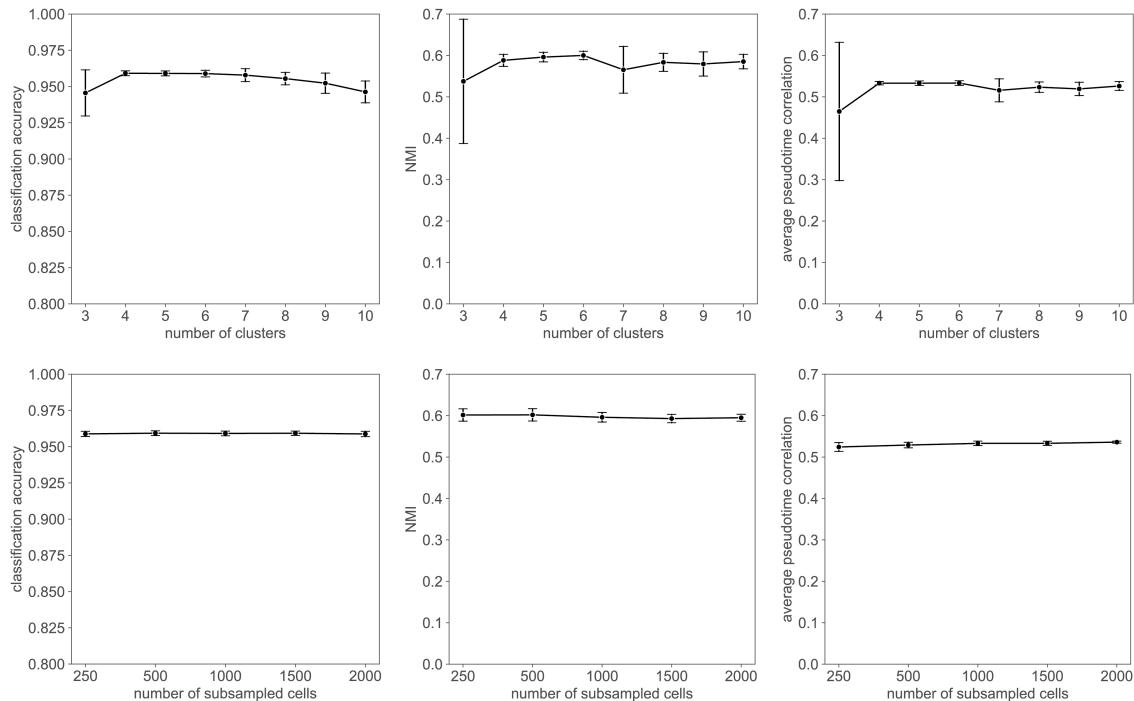
Supplementary Figure 11: Comparison of the robustness of feature selection methods on inferring differentiation trajectories under library size noise corruption. (a) Average k -NN classification accuracy vs. average percent change in k -NN classification accuracy as the total mRNA count decreased (high to medium). Library size was reduced by modifying the location parameter in the log-normal distribution in Splatter [72] that specifies library size scaling factors (high: location = 12, medium: location = 11). (b) Average pseudotime correlation vs. average percent change in pseudotime correlation as the total mRNA count decreased (high to medium).



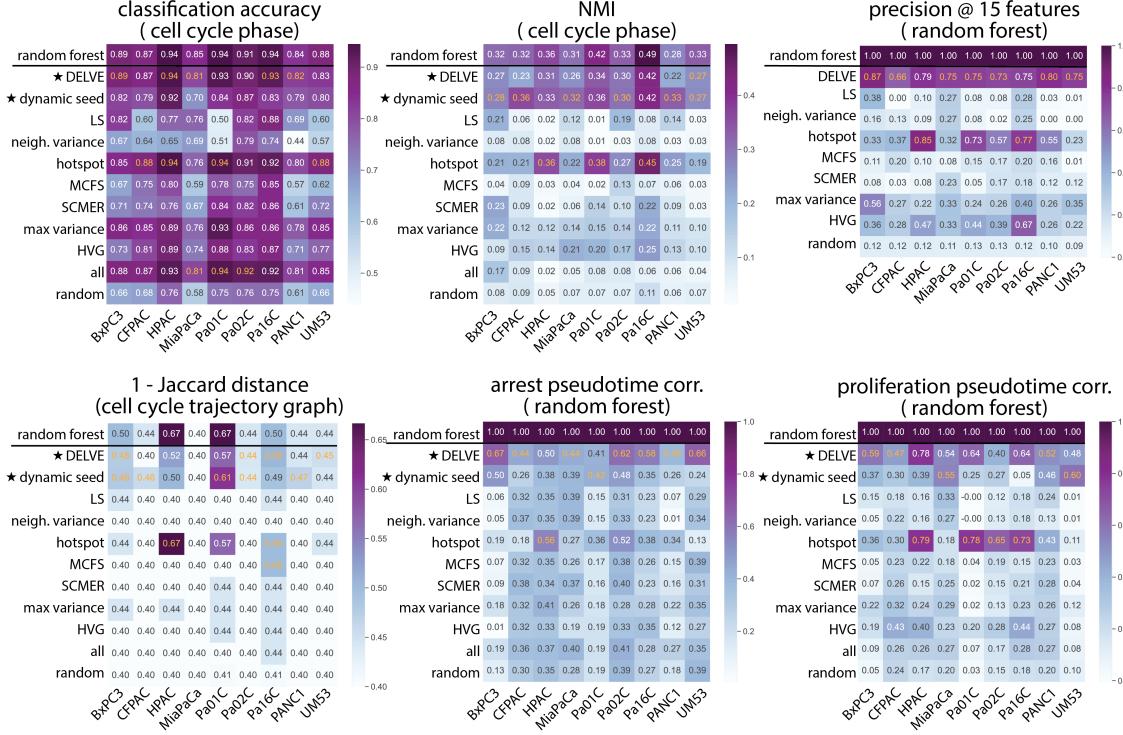
Supplementary Figure 12: Comparison of the robustness of feature selection methods on inferring differentiation trajectories under dropout noise corruption. (a) Average k -NN classification accuracy vs. average percent change in k -NN classification accuracy as the amount of dropout or sparsity increased (low to medium). Technical dropout was simulated by undersampling mRNA counts by sampling from a binomial distribution with the scale parameter or dropout rate proportional to the mean expression of each gene (low: $\lambda = 0$, medium: $\lambda = 0.05$). (b) Average pseudotime correlation vs. average percent change in pseudotime correlation as the amount of dropout or sparsity increased (low to medium).



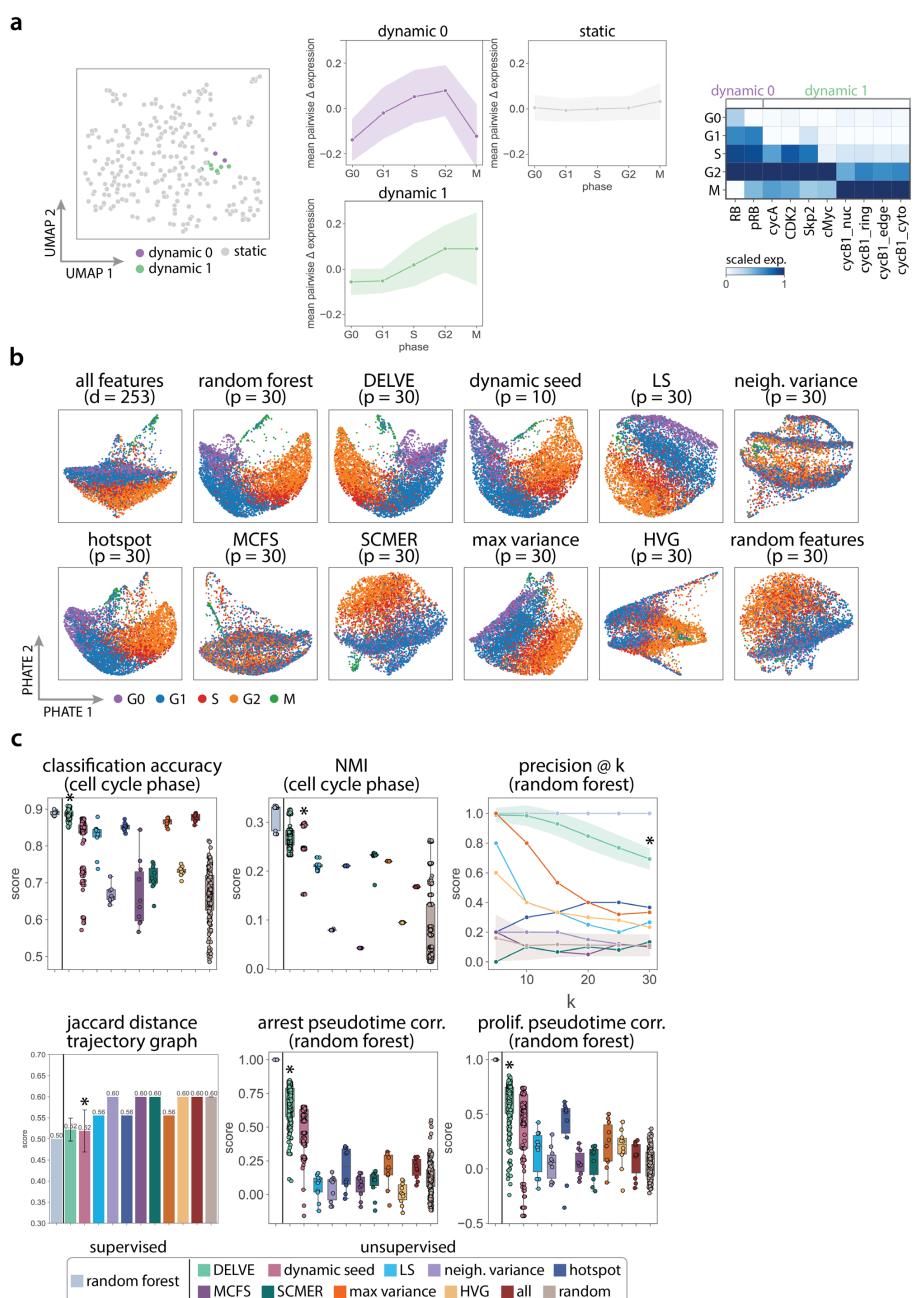
Supplementary Figure 13: Validation of DELVE seed selection on retinal pigmented epithelial (RPE) cell cycle dataset. Top ranked features identified by (a) random forest classifier trained on ground truth cell cycle phase annotations or (b) random forest regressor trained on ground truth cell cycle age measurements. Features highlighted in red were also identified by DELVE seed selection (See Figure 4a heatmap).



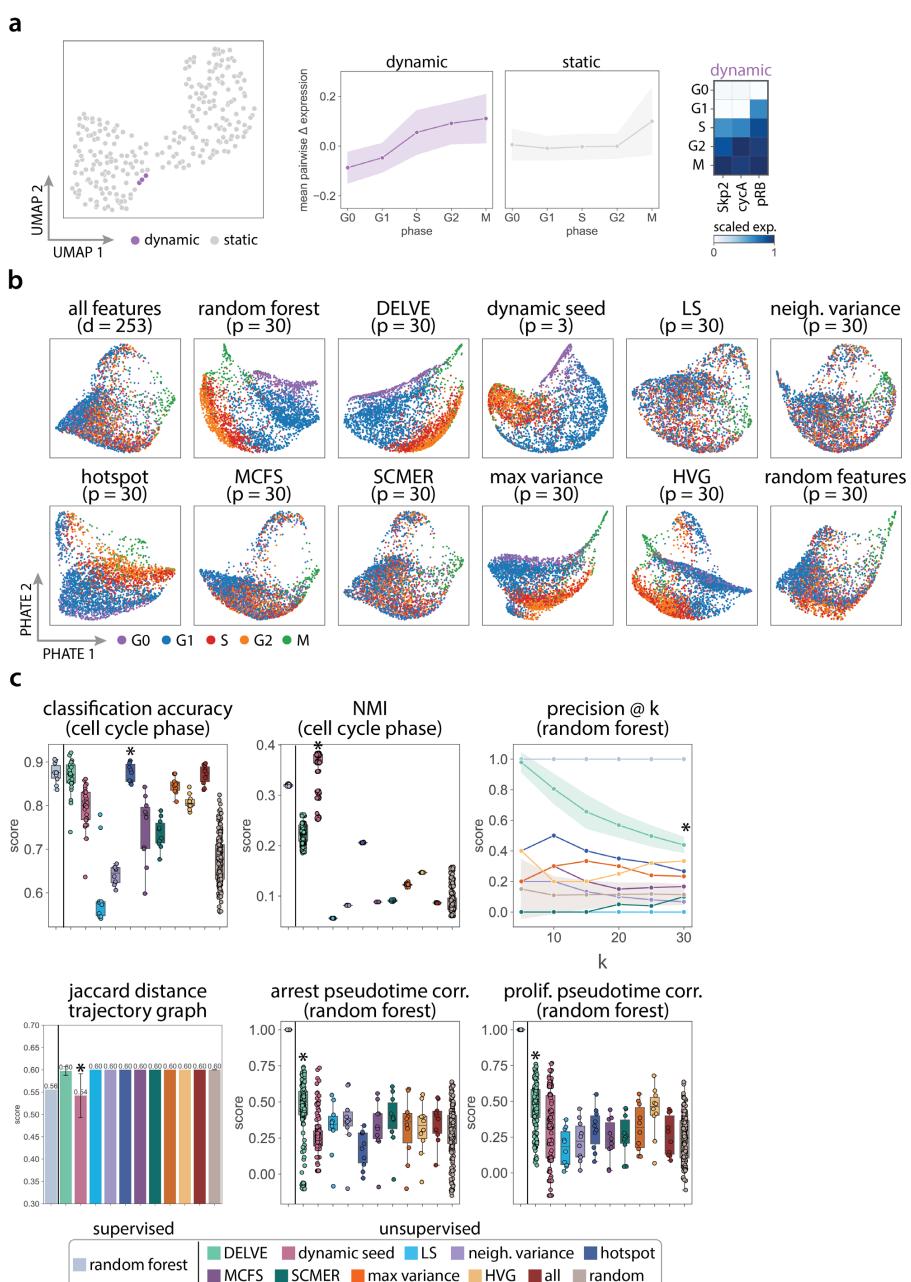
Supplementary Figure 14: DELVE is robust to changes in hyperparameters for the retinal pigmented epithelial (RPE) cell cycle dataset. DELVE achieves similar classification accuracy, normalized mutual information (NMI) clustering score, and pseudotime correlation across a range of (top) cluster sizes and (bottom) subsampled cells. Plots show results over 20 random trials. Error bars represent the standard deviation.



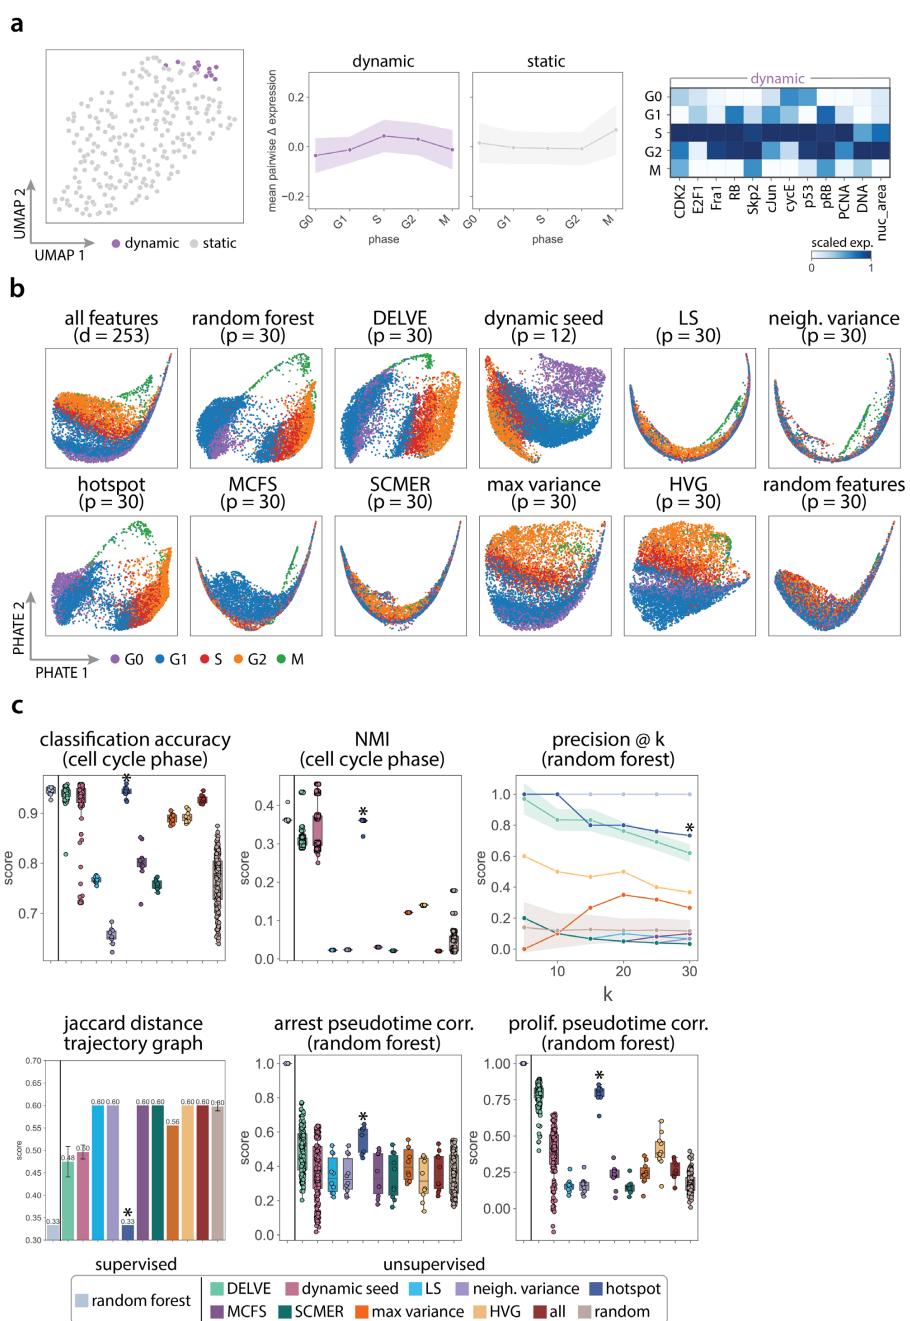
Supplementary Figure 15: Feature selection method performance on pancreatic adenocarcinoma (PDAC) cell cycle datasets. Performance of twelve feature selection methods on preserving cell cycle trajectories from 9 PDAC cancer cell lines (BxPC3, CFPAC, HPAC, MiaPaCa, Pa02C, Pa01C, Pa16C, PANC1, and UM53) profiled with protein immunofluorescence imaging. Following feature selection, cell cycle preservation was quantitatively assessed according to several metrics including: support vector machine classification accuracy to the ground truth phase annotations, normalized mutual information (NMI) clustering score to ground truth phase annotations, precision of cell cycle phase-specific imaging-derived features as measured by a random forest classifier trained on ground truth phase annotations, Jaccard distance between predicted cell cycle trajectory graphs and a ground truth reference cell cycle trajectory curated from the literature, and the Spearman rank correlation between estimated pseudotime and the ground truth as measured by a random forest classifier trained on ground truth phase annotations. Heatmaps show the average performance. Approaches with the highest average score are highlighted in yellow.



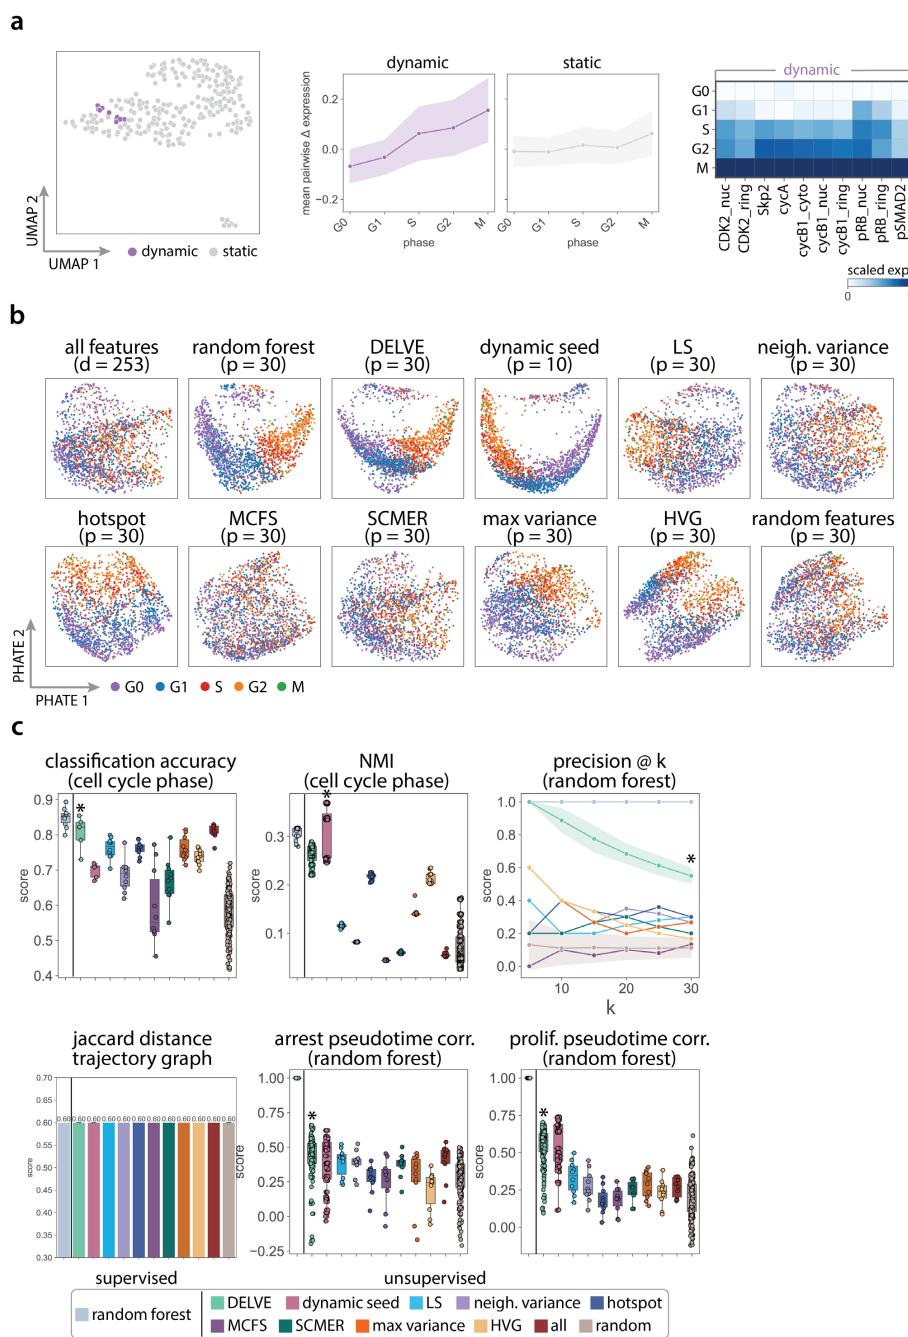
Supplementary Figure 16: DELVE recovers BxPC3 pancreatic adenocarcinoma cell cycle trajectories in protein immunofluorescence imaging data. BxPC3 cells were profiled with protein immunofluorescence imaging to measure 63 core cell cycle effectors resulting in a dataset with $d = 253$ imaging-derived features. (a) DELVE identified two modules of dynamic features representing a minimum cell cycle. (a left) UMAP visualization of image-derived features where each point indicates a dynamic or static feature identified by the model. (a middle) The average pairwise change in expression for features within a module ordered across ground truth cell cycle phase annotations. (a right) Heatmap illustrating the standardized average expression of dynamic seed features across cell cycle phases. (b) Feature selection was performed to select the top $p = 30$ ranked features. Example PHATE [75] visualizations of cell cycle trajectories for twelve feature selection approaches. (c) Quantitative assessment of twelve feature selection methods on preserving cell cycle phases and phase transitions according to several metrics including: support vector machine classification accuracy to the ground truth phase annotations, normalized mutual information (NMI) clustering score to ground truth phase annotations, precision of cell cycle phase-specific imaging-derived features as measured by a random forest classifier trained on ground truth phase annotations, Jaccard distance between predicted cell cycle trajectory graphs and a ground truth reference cell cycle trajectory curated from the literature, and the Spearman rank correlation between estimated pseudotime and the ground truth as measured by a random forest classifier trained on ground truth phase annotations. Boxplots show the results over 10 random trials, error bands represent the standard deviation, and * indicates the method with the highest median score.



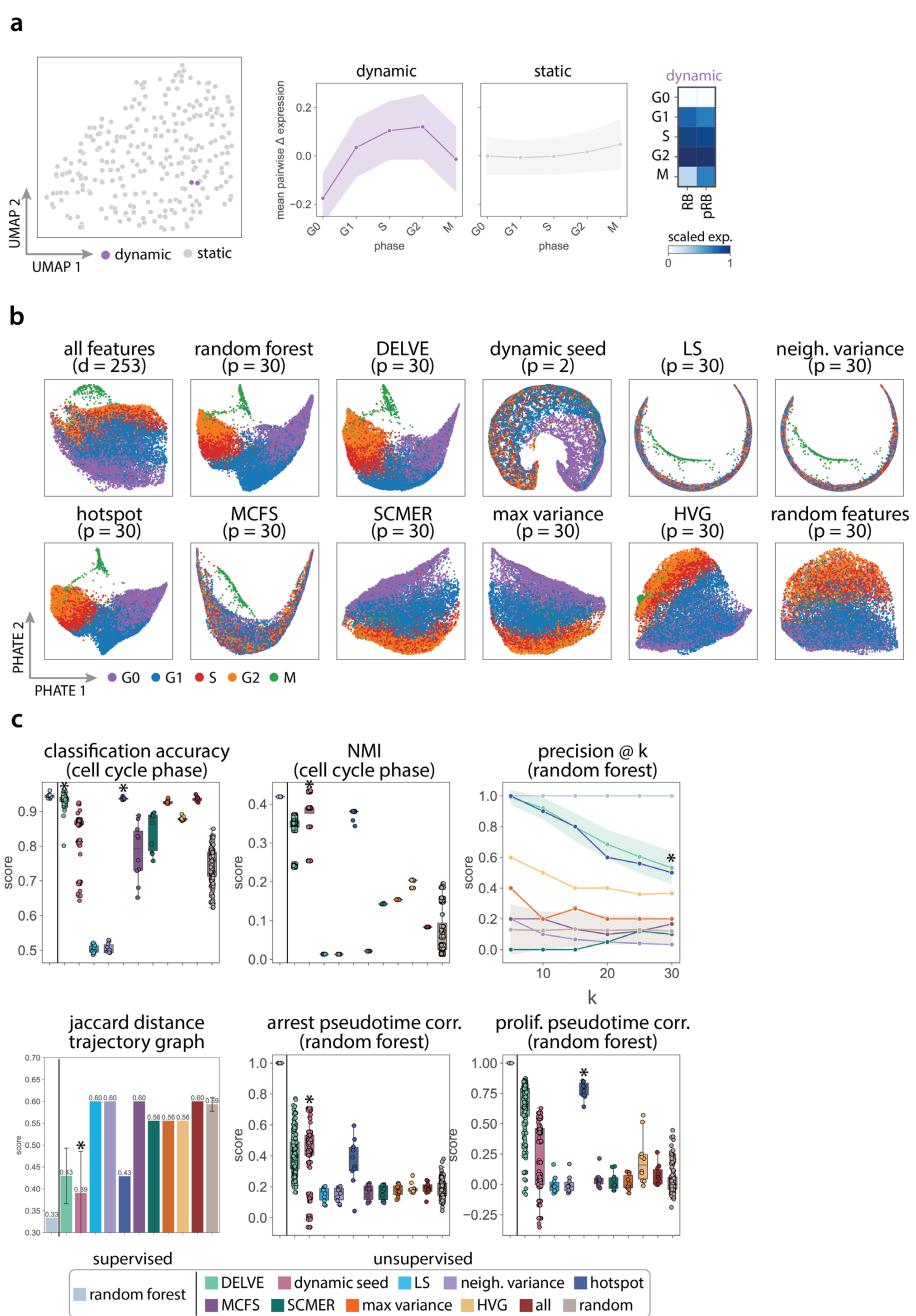
Supplementary Figure 17: DELVE recovers CFPAC pancreatic adenocarcinoma cell cycle trajectories in protein immunofluorescence imaging data. CFPAC cells were profiled with protein immunofluorescence imaging to measure 63 core cell cycle effectors resulting in a dataset with $d = 253$ imaging-derived features. (a) DELVE identified one module of dynamic features representing a minimum cell cycle. (a left) UMAP visualization of image-derived features where each point indicates a dynamic or static feature identified by the model. (a middle) The average pairwise change in expression for features within a module ordered across ground truth cell cycle phase annotations. (a right) Heatmap illustrating the standardized average expression of dynamic seed features across cell cycle phases. (b) Feature selection was performed to select the top $p = 30$ ranked features. Example PHATE [75] visualizations of cell cycle trajectories for twelve feature selection approaches. (c) Quantitative assessment of twelve feature selection methods on preserving cell cycle phases and phase transitions according to several metrics including: support vector machine classification accuracy to the ground truth phase annotations, normalized mutual information (NMI) clustering score to ground truth phase annotations, precision of cell cycle phase-specific imaging-derived features as measured by a random forest classifier trained on ground truth phase annotations, Jaccard distance between predicted cell cycle trajectory graphs and a ground truth reference cell cycle trajectory curated from the literature, and the Spearman rank correlation between estimated pseudotime and the ground truth as measured by a random forest classifier trained on ground truth phase annotations. Boxplots show the results over 10 random trials, error bands represent the standard deviation, and * indicates the method with the highest median score.



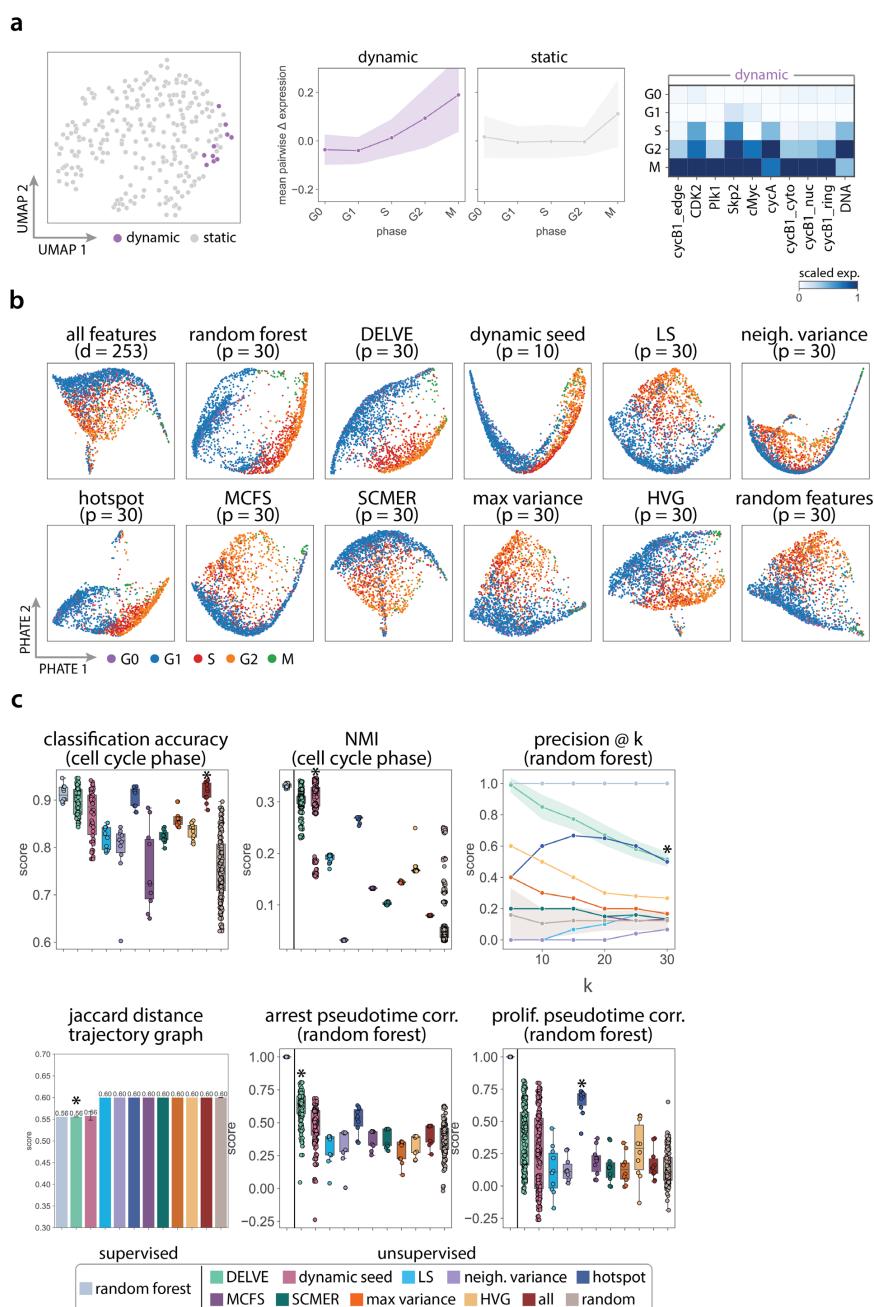
Supplementary Figure 18: DELVE recovers HPAC pancreatic adenocarcinoma cell cycle trajectories in protein immunofluorescence imaging data. HPAC cells were profiled with protein immunofluorescence imaging to measure 63 core cell cycle effectors resulting in a dataset with $d = 253$ imaging-derived features. (a) DELVE identified one module of dynamic features representing a minimum cell cycle. (a left) UMAP visualization of image-derived features where each point indicates a dynamic or static feature identified by the model. (a middle) The average pairwise change in expression for features within a module ordered across ground truth cell cycle phase annotations. (a right) Heatmap illustrating the standardized average expression of dynamic seed features across cell cycle phases. (b) Feature selection was performed to select the top $p = 30$ ranked features. Example PHATE [75] visualizations of cell cycle trajectories for twelve feature selection approaches. (c) Quantitative assessment of twelve feature selection methods on preserving cell cycle phases and phase transitions according to several metrics including: support vector machine classification accuracy to the ground truth phase annotations, normalized mutual information (NMI) clustering score to ground truth phase annotations, precision of cell cycle phase-specific imaging-derived features as measured by a random forest classifier trained on ground truth phase annotations, Jaccard distance between predicted cell cycle trajectory graphs and a ground truth reference cell cycle trajectory curated from the literature, and the Spearman rank correlation between estimated pseudotime and the ground truth as measured by a random forest classifier trained on ground truth phase annotations. Boxplots show the results over 10 random trials, error bands represent the standard deviation, and * indicates the method with the highest median score.



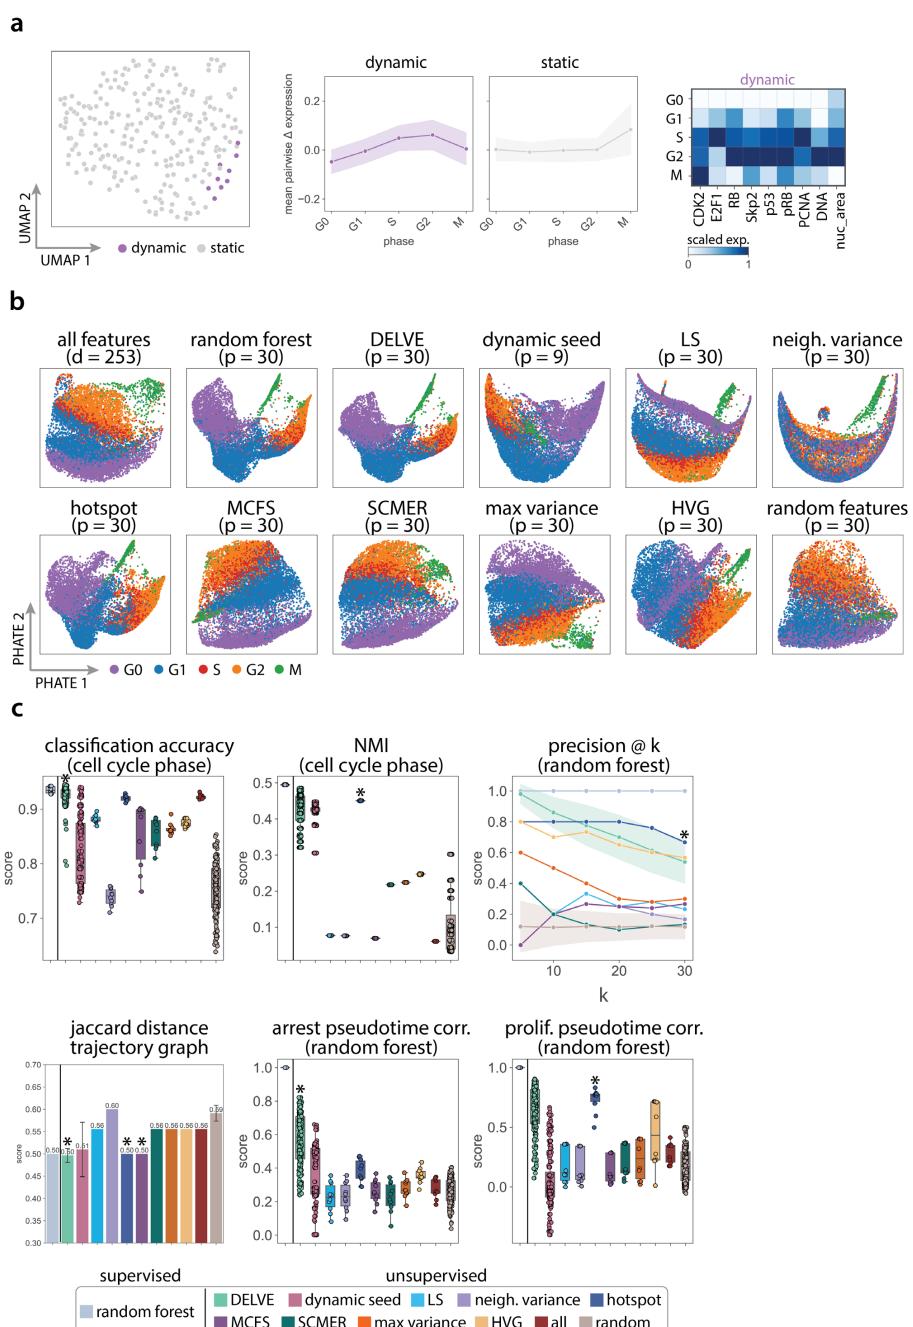
Supplementary Figure 19: DELVE recovers MiaPaCa pancreatic adenocarcinoma cell cycle trajectories in protein immunofluorescence imaging data. MiaPaCa cells were profiled with protein immunofluorescence imaging to measure 63 core cell cycle effectors resulting in a dataset with $d = 253$ imaging-derived features. (a) DELVE identified one module of dynamic features representing a minimum cell cycle. (a left) UMAP visualization of image-derived features where each point indicates a dynamic or static feature identified by the model. (a middle) The average pairwise change in expression for features within a module ordered across ground truth cell cycle phase annotations. (a right) Heatmap illustrating the standardized average expression of dynamic seed features across cell cycle phases. (b) Feature selection was performed to select the top $p = 30$ ranked features. Example PHATE [75] visualizations of cell cycle trajectories for twelve feature selection approaches. (c) Quantitative assessment of twelve feature selection methods on preserving cell cycle phases and phase transitions according to several metrics including: support vector machine classification accuracy to the ground truth phase annotations, normalized mutual information (NMI) clustering score to ground truth phase annotations, precision of cell cycle phase-specific imaging-derived features as measured by a random forest classifier trained on ground truth phase annotations, Jaccard distance between predicted cell cycle trajectory graphs and a ground truth reference cell cycle trajectory curated from the literature, and the Spearman rank correlation between estimated pseudotime and the ground truth as measured by a random forest classifier trained on ground truth phase annotations. Boxplots show the results over 10 random trials, error bands represent the standard deviation, and * indicates the method with the highest median score.



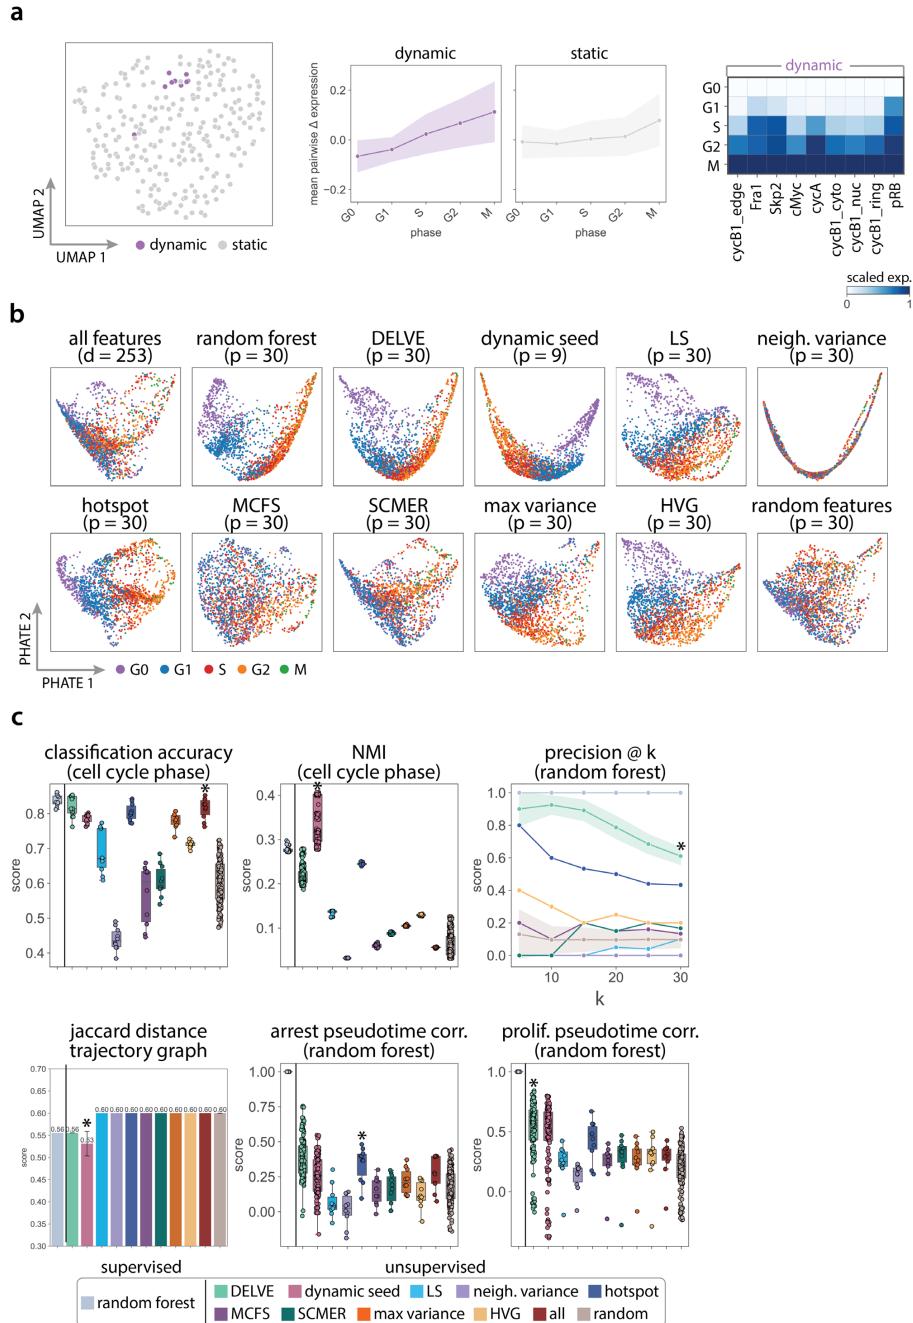
Supplementary Figure 20: DELVE recovers Pa01C pancreatic adenocarcinoma cell cycle trajectories in protein immunofluorescence imaging data. Pa01C cells were profiled with protein immunofluorescence imaging to measure 63 core cell cycle effectors resulting in a dataset with $d = 253$ imaging-derived features. (a) DELVE identified one module of dynamic features representing a minimum cell cycle. (a left) UMAP visualization of image-derived features where each point indicates a dynamic or static feature identified by the model. (a middle) The average pairwise change in expression for features within a module ordered across ground truth cell cycle phase annotations. (a right) Heatmap illustrating the standardized average expression of dynamic seed features across cell cycle phases. (b) Feature selection was performed to select the top $p = 30$ ranked features. Example PHATE [75] visualizations of cell cycle trajectories for twelve feature selection approaches. (c) Quantitative assessment of twelve feature selection methods on preserving cell cycle phases and phase transitions according to several metrics including: support vector machine classification accuracy to the ground truth phase annotations, normalized mutual information (NMI) clustering score to ground truth phase annotations, precision of cell cycle phase-specific imaging-derived features as measured by a random forest classifier trained on ground truth phase annotations, Jaccard distance between predicted cell cycle trajectory graphs and a ground truth reference cell cycle trajectory curated from the literature, and the Spearman rank correlation between estimated pseudotime and the ground truth as measured by a random forest classifier trained on ground truth phase annotations. Boxplots show the results over 10 random trials, error bands represent the standard deviation, and * indicates the method with the highest median score.



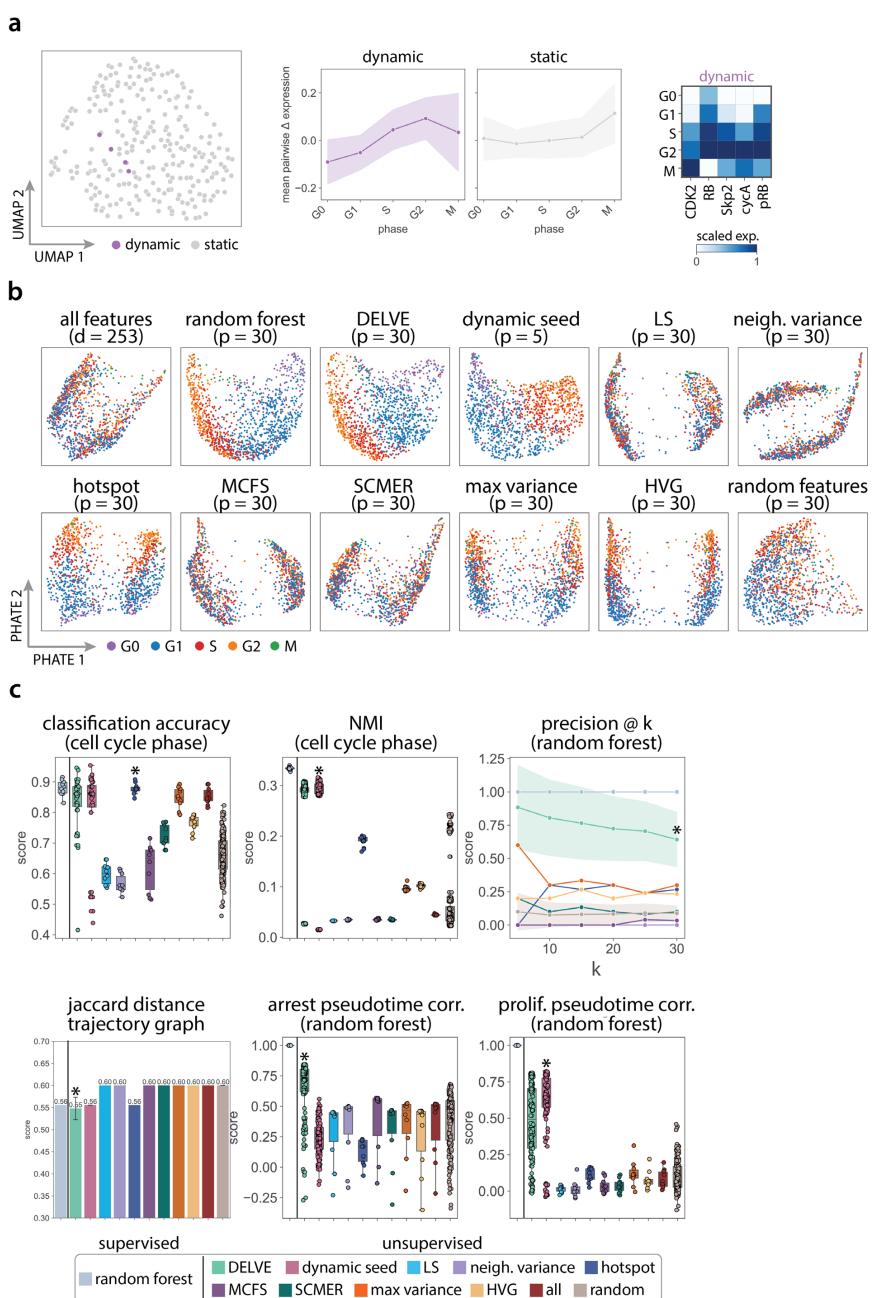
Supplementary Figure 21: DELVE recovers Pa02C pancreatic adenocarcinoma cell cycle trajectories in protein immunofluorescence imaging data. Pa02C cells were profiled with protein immunofluorescence imaging to measure 63 core cell cycle effectors resulting in a dataset with $d = 253$ imaging-derived features. (a) DELVE identified one module of dynamic features representing a minimum cell cycle. (a left) UMAP visualization of image-derived features where each point indicates a dynamic or static feature identified by the model. (a middle) The average pairwise change in expression for features within a module ordered across ground truth cell cycle phase annotations. (a right) Heatmap illustrating the standardized average expression of dynamic seed features across cell cycle phases. (b) Feature selection was performed to select the top $p = 30$ ranked features. Example PHATE [75] visualizations of cell cycle trajectories for twelve feature selection approaches. (c) Quantitative assessment of twelve feature selection methods on preserving cell cycle phases and phase transitions according to several metrics including: support vector machine classification accuracy to the ground truth phase annotations, normalized mutual information (NMI) clustering score to ground truth phase annotations, precision of cell cycle phase-specific imaging-derived features as measured by a random forest classifier trained on ground truth phase annotations, Jaccard distance between predicted cell cycle trajectory graphs and a ground truth reference cell cycle trajectory curated from the literature, and the Spearman rank correlation between estimated pseudotime and the ground truth as measured by a random forest classifier trained on ground truth phase annotations. Boxplots show the results over 10 random trials, error bands represent the standard deviation, and * indicates the method with the highest median score.



Supplementary Figure 22: DELVE recovers Pa16C pancreatic adenocarcinoma cell cycle trajectories in protein immunofluorescence imaging data. Pa16C cells were profiled with protein immunofluorescence imaging to measure 63 core cell cycle effectors resulting in a dataset with $d = 253$ imaging-derived features. (a) DELVE identified one module of dynamic features representing a minimum cell cycle. (a left) UMAP visualization of image-derived features where each point indicates a dynamic or static feature identified by the model. (a middle) The average pairwise change in expression for features within a module ordered across ground truth cell cycle phase annotations. (a right) Heatmap illustrating the standardized average expression of dynamic seed features across cell cycle phases. (b) Feature selection was performed to select the top $p = 30$ ranked features. Example PHATE [75] visualizations of cell cycle trajectories for twelve feature selection approaches. (c) Quantitative assessment of twelve feature selection methods on preserving cell cycle phases and phase transitions according to several metrics including: support vector machine classification accuracy to the ground truth phase annotations, normalized mutual information (NMI) clustering score to ground truth phase annotations, precision of cell cycle phase-specific imaging-derived features as measured by a random forest classifier trained on ground truth phase annotations, Jaccard distance between predicted cell cycle trajectory graphs and a ground truth reference cell cycle trajectory curated from the literature, and the Spearman rank correlation between estimated pseudotime and the ground truth as measured by a random forest classifier trained on ground truth phase annotations. Boxplots show the results over 10 random trials, error bands represent the standard deviation, and * indicates the method with the highest median score.



Supplementary Figure 23: DELVE recovers PANC1 pancreatic adenocarcinoma cell cycle trajectories in protein immunofluorescence imaging data. PANC1 cells were profiled with protein immunofluorescence imaging to measure 63 core cell cycle effectors resulting in a dataset with $d = 253$ imaging-derived features. (a) DELVE identified one module of dynamic features representing a minimum cell cycle. (a left) UMAP visualization of image-derived features where each point indicates a dynamic or static feature identified by the model. (a middle) The average pairwise change in expression for features within a module ordered across ground truth cell cycle phase annotations. (a right) Heatmap illustrating the standardized average expression of dynamic seed features across cell cycle phases. (b) Feature selection was performed to select the top $p = 30$ ranked features. Example PHATE [75] visualizations of cell cycle trajectories for twelve feature selection approaches. (c) Quantitative assessment of twelve feature selection methods on preserving cell cycle phases and phase transitions according to several metrics including: support vector machine classification accuracy to the ground truth phase annotations, normalized mutual information (NMI) clustering score to ground truth phase annotations, precision of cell cycle phase-specific imaging-derived features as measured by a random forest classifier trained on ground truth phase annotations, Jaccard distance between predicted cell cycle trajectory graphs and a ground truth reference cell cycle trajectory curated from the literature, and the Spearman rank correlation between estimated pseudotime and the ground truth as measured by a random forest classifier trained on ground truth phase annotations. Boxplots show the results over 10 random trials, error bands represent the standard deviation, and * indicates the method with the highest median score.



Supplementary Figure 24: DELVE recovers UM53 pancreatic adenocarcinoma cell cycle trajectories in protein immunofluorescence imaging data. UM53 cells were profiled with protein immunofluorescence imaging to measure 63 core cell cycle effectors resulting in a dataset with $d = 253$ imaging-derived features. (a) DELVE identified one module of dynamic features representing a minimum cell cycle. (a left) UMAP visualization of image-derived features where each point indicates a dynamic or static feature identified by the model. (a middle) The average pairwise change in expression for features within a module ordered across ground truth cell cycle phase annotations. (a right) Heatmap illustrating the standardized average expression of dynamic seed features across cell cycle phases. (b) Feature selection was performed to select the top $p = 30$ ranked features. Example PHATE [75] visualizations of cell cycle trajectories for twelve feature selection approaches. (c) Quantitative assessment of twelve feature selection methods on preserving cell cycle phases and phase transitions according to several metrics including: support vector machine classification accuracy to the ground truth phase annotations, normalized mutual information (NMI) clustering score to ground truth phase annotations, precision of cell cycle phase-specific imaging-derived features as measured by a random forest classifier trained on ground truth phase annotations, Jaccard distance between predicted cell cycle trajectory graphs and a ground truth reference cell cycle trajectory curated from the literature, and the Spearman rank correlation between estimated pseudotime and the ground truth as measured by a random forest classifier trained on ground truth phase annotations. Boxplots show the results over 10 random trials, error bands represent the standard deviation, and * indicates the method with the highest median score.