Multi-resolution characterization of molecular taxonomies in bulk and single-cell transcriptomics data

Eric R. Reed[1,2] and Stefano Monti [1,2,3,*]

[1]Section of Computational Biomedicine, Boston University School of Medicine, Boston, MA 02118, USA,
[2]Bioinformatics Program, College of Engineering, Boston University, Boston, MA 02118, USA and [3]Department of Biostatistics, Boston University School of Public Health, Boston, MA 02118, USA

ABSTRACT

As high-throughput genomics assays become more efficient and cost effective, their utilization has become standard in large-scale biomedical projects. These studies are often explorative, in that relationships between samples are not explicitly defined *a priori*, but rather emerge from data-driven discovery and annotation of molecular subtypes, thereby informing hypotheses and independent evaluation. Here, we present *K2Taxonomer*, a novel unsupervised recursive partitioning algorithm and associated R package that utilize ensemble learning to identify robust subgroups in a 'taxonomy-like' structure. *K2Taxonomer* was devised to accommodate different data paradigms, and is suitable for the analysis of both bulk and single-cell transcriptomics, and other '-omics', data. For each of these data types, we demonstrate the power of *K2Taxonomer* to discover known relationships in both simulated and human tissue data. We conclude with a practical application on breast cancer tumor infiltrating lymphocyte (TIL) single-cell profiles, in which we identified co-expression of translational machinery genes as a dominant transcriptional program shared by T cells subtypes, associated with better prognosis in breast cancer tissue bulk expression data.

GRAPHICAL ABSTRACT
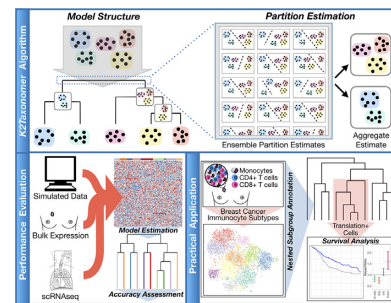
INTRODUCTION

As high-throughput transcriptomic assays become more efficient and cost-effective, they are being routinely integrated into large-scale biomedical projects (1–4). Bulk gene expression profiling by RNA sequencing (RNAseq) has been widely adopted in multiple high-throughput genomics studies, the paramount example being The Cancer Genome Atlas (TCGA) data commons, which currently include 10,558 bulk RNA sequencing (RNAseq) profiles across 33 cancer types (https://portal.gdc.cancer.gov/). Furthermore, since its first published application in 2009 (5), the size of single-cell RNA sequencing (scRNAseq) studies has exploded, such that it is now commonplace for studies to generate tens of thousands of profiles (6). As the scale of these studies and the associated datasets increases, so does their utility as a resource from which biological information can be extracted through the application of machine learning approaches. Common deliverables of these types of analysis include the discovery and characterization of molecular subtypes, which are prevalent in both bulk and single-cell gene expression studies. For example, TCGA bulk expression data have been utilized to characterize subtypes of numerous cancers (7), including but not limited to: breast (8), colorectal (9), liver (10) and bladder cancer (11,12). Similarly, the characterization of molecular subtypes is a stan-

*To whom correspondence should be addressed. Tel: +1 617 358 7087; Email: smonti@bu.edu

dard component of the scRNAseq data analysis workflow, insofar as estimation and annotation of subpopulations of cells is one of the primary goals of the assay (13).

The general framework for subtype characterization can be summarized in two steps: (i) estimation of data-driven groups of observations via application of an unsupervised learning procedure, followed by (ii) annotation of each group based on the identification of distinct patterns of gene expression relative to other groups. While most approaches focus on discovering a 'flat' set of non-overlapping groups or subtypes, in this manuscript we present an alternative approach, devised to emphasize 'taxonomy-like' hierarchical relationships between observations to discover nested subgroups.

Whereas a wide range of unsupervised learning algorithms is available for the analysis of bulk gene expression data, the considerable sparsity of scRNAseq data has motivated the development of novel methods specifically tailored to the analysis of this type of sparse, high-dimensional data. Popular software packages, such as *Seurat* (14) and *Scran* (15), generate 'flat' clusters, in which a finite set of mutually exclusive cell types is estimated. In so doing, they fail to capture the 'taxonomy-like', hierarchical structure that may exist among subgroups of observations at multiple levels of resolution, driven by transcriptional signatures based on different factors, including but not limited to: shared lineage, cell state, pathway activity, or morphological origin. Complementary methods exist to model such relationships, such as *Neighborhood Joining* (16,17) and more recent single-cell *trajectory inference* approaches (18), which estimate 'pseudo-temporal' states of individual cells indicative of developmental progression. Given the stringent interpretation of such models, their suitability depends on the assumption that the measured similarity between neighbors of cell profiles arises from a distinct continuous progression of molecular activity. However, the relative similarity between cell profiles may be confounded by numerous factors, including: cell cycle, spatial patterning, cell stress, and batch effects (19). To overcome these shortcomings, one recent method, *partition-based graph abstraction* (*PAGA*) (20), was devised to model complex topologies by estimating a graph of 'high-confidence' connections between labeled cell types based on their shared nearest neighbors. This method has the advantage of being able to first identify disconnected subgraphs from which to model separate trajectories. Even so, a ubiquitous attribute of *trajectory inference* approaches, including *PAGA*, is that all distances between cell profiles are computed based on a single set of features, generated by selection filtering and/or dimensionality reduction (21), thus precluding the discovery of nested structures defined by distinct transcriptional programs shared by relatively few cells.

Hierarchical clustering (HC) algorithms at face value address the need for a multi-resolution representation of the relationship among observations, and while originally adopted for the analysis of bulk gene expression data (22), numerous packages have also been developed for scRNAseq analysis, such as *pcaReduce* (23), *ascend* (24) and *BackSPIN* (25). However, since the number of possible subgroupings increases with the number of observations, robustly identifying such relationships can be challenging. As

a result, tree-cutting methods are often applied, ultimately yielding a flat set of non-overlapping clusters. Furthermore, as with *trajectory inference*, the bottom-up nature of HC's sample aggregation procedure forces the use of the same set of genes/features to drive the agglomeration at all levels of the hierarchy. Finally, HC methods do not support the clustering of *groups* of samples, or group of cell profiles representing cell types in scRNAseq, nor the identification of an interpretable taxonomic hierarchy over these groups or cell types.

Here, we introduce *K2Taxonomer*, a novel taxonomy discovery approach and associated R package for the estimation and in-silico characterization of hierarchical subgroup structures in both bulk and single-cell data. An important feature of the approach is that it can analyze both individual samples as well as sample groups such as, but not limited to, those corresponding to scRNAseq cell types. The package employs a recursive partitioning algorithm, which utilizes repeated perturbations of the data at each partition to estimate ensemble-based $K = 2$ subgroups. For scRNAseq analysis, *K2Taxonomer* utilizes the *constrained k-means algorithm* (26) to estimate partitions of the data at the cell type level, while preserving the influence of each individual cell profile. A defining feature of the method is that each recursive split of the input data is based on a distinct set of features selected to be most discriminatory within the subset of samples member of the current hierarchy branch. This makes the approach quite distinct from both standard clustering algorithms and trajectory inference algorithms, and particularly apt to discover nested taxonomies. *K2Taxonomer* thus fills a methodological gap and provides a rigorous way to further resolve biological insights from clustered scRNA-seq data. In addition, the package includes functionalities to comprehensively characterize and statistically test each subgroup based on their estimated stability, gene expression profiles, and a priori phenotypic annotation of individual profiles. Importantly, all results are aggregated into an automatically generated interactive portal to assist in parsing the results.

In this manuscript, we assess the performance of *K2Taxonomer* for partitioning both bulk gene expression and scRNAseq data, using both simulated and publicly available data sets, and we compare it to agglomerative clustering procedures. For bulk gene expression data, performance is assessed in terms of unsupervised sorting of breast cancer subtypes and established genotypic markers, using breast cancer patient tumor tissue data from the Molecular Taxonomy of Breast Cancer International Consortium (METABRIC) (27) and the TCGA compendia. For scRNAseq data, performance is assessed in terms of recapitulation of established relationships between 28 annotated cell types of the airway of healthy subjects (28). We conclude with a case study where we perform a *K2Taxonomer*-based analysis of breast cancer tumor infiltrating lymphocytes (TILs) profiled by scRNAseq (29). Our analysis significantly expands upon previously published results and identifies a phenotypically diverse subgroup of CD4 and CD8 cells, characterized by constitutive up-regulation of a subset of translation machinery genes. We further show that high expression of these genes in breast cancer tissue bulk expression is associated with better survival, supporting re-

cent findings on the role of the translation machinery assembly in T cell activation (30,31), and demonstrate that this coordinated expression of the translation machinery is pervasive among T cell subpopulations to such an extent that the expression levels of these genes in bulk measurements of tumor tissue is predictive of the degree of immune infiltration. The complete suite of analysis results is accessible through an automatically generated and publicly accessible portal (https://montilab.bu.edu/k2BRCAtcell/).

While we focus on the analysis of transcriptomics data, we emphasize that our approach is applicable to other 'omics' data, such as those generated by high-throughput proteomics and metabolomics assays. Moreover, applications of *K2Taxonomer* are not limited to cancer-centric studies. For example, we applied an earlier prototype of *K2Taxonomer* to the analysis of a toxicogenomic study on the effects of environmental exposures on adipocyte activity, and the tool proved to be instrumental to the identification of chemical subgroups and the pathways contributing to either their deleterious or beneficial effects on energy homeostasis (32).

MATERIALS AND METHODS

K2Taxonomer algorithm overview

K2Taxonomer implements a recursive partitioning algorithm that takes as input either a set of individual observations or a set of sample groups and returns a top-down hierarchical taxonomy of those samples or groups (Figure 1). Here we summarize the *K2Taxonomer* algorithm. A detailed description of the method is provided in the following three subsections.

To achieve robust model estimation, each partition is defined based on the aggregation of repeated partition estimations from distinct perturbations of the original set. Each of these partition estimates is created in three steps. First, a perturbation-specific data set is generated by bootstrapping features, i.e. sampling features from the original data set with replacement. Next, this perturbation-specific data set undergoes variability-based feature selection filtering. Finally, a $K = 2$ clustering algorithm is run, producing a perturbation-specific partition estimate. These three steps are repeated, generating a set of perturbation-specific partition estimates, which are aggregated into a cosine similarity matrix. The aggregate partition is then estimated based on a $K = 2$ tree cut following hierarchical clustering of the transformed cosine similarity matrix into a distance matrix, calculated as 1 - cosine similarity, with a user-specified agglomeration method.

The *K2Taxonomer* package supports both *observation-level* and *group-level* analysis modalities, depending on whether the analysis end-points are single samples or groups of samples, respectively. In the group-level modality, the algorithm is applied to data sets where observations have a priori-assigned group labels, and the objective is to identify intermediate relationships between these groups. This functionality was specifically incorporated to enable partitioning and annotations of cell types estimated by scR-NAseq clustering algorithms, but it is applicable to any data set with group-level labels.

For further customization of analyses, the *K2Taxonomer* R package permits the use of user-specified functions for performing perturbation-specific partition estimates.

K2Taxonomer feature filtering

A distinguishing property of *K2Taxonomer* when compared to other methods, such as traditional agglomerative hierarchical cluster or trajectory inference (21) is the manner in which feature selection is implemented. Even in large studies of high-throughput data sets, the number of features is typically much larger than the number of observations. This generally requires filtering the data set prior to modeling in order to reduce variance and computational expense of model fit. One way to do this is through feature selection, in which features suspected to contain more information about the relationship between observations are chosen for down-stream analysis. For unsupervised learning, relative information estimation is commonly calculated via variability-based metrics. Assuming the amount of noise is consistent across features, these metrics will capture the relative magnitude of the signal of individual features. Two common choices are standard deviation (SD) and median absolute deviation (MAD), of which the former is more statistically efficient with a small sample size and the latter is more robust to outliers (33). Implementation of these feature selection techniques prior to modeling may be problematic when learning hierarchical models. The magnitude of variability-based metrics is influenced by the frequency of observations for which the signal-to-noise ratio is higher, such that the subset of features is more likely to capture broader relationships between larger sets of observations and less likely to capture relationships between smaller sets of observations. This can obscure important relationships within smaller sets of observations, as when evaluating a sub-group of samples in a hierarchical procedure. In addition, an appropriate choice of the number features to use for modeling is difficult to determine a priori and may be obscured by many factors, including: the number of subgroups, number of observations belonging to each subgroup, and the number of features distinguishing individual subgroups.

To overcome these challenges, *K2Taxonomer* produces a model fit for each partition independently, such that feature selection is only performed within the subgroup of observations being evaluated at a given step. In particular, at each recursive step, the objective of partition estimation is to split the data based only on the dominant relationship between two subgroups. Since the selected features need only capture one relationship, a much smaller subset of features will be sufficient to discover this partition. By default, *K2Taxonomer* uses the square root of the total number of features, which is used in a related albeit supervised learning method, random forests (34). In doing so, the percentage of filtered features is dependent on the total number of features. For example, if the data set consists of 1,000 or 10,000 features, *K2Taxonomer* will estimate partitions using 3.2% or 1.0% of the total number of features, respectively. The appropriateness of using the square root of the total number of features against fixed percentages is later assessed with simulation-based testing.
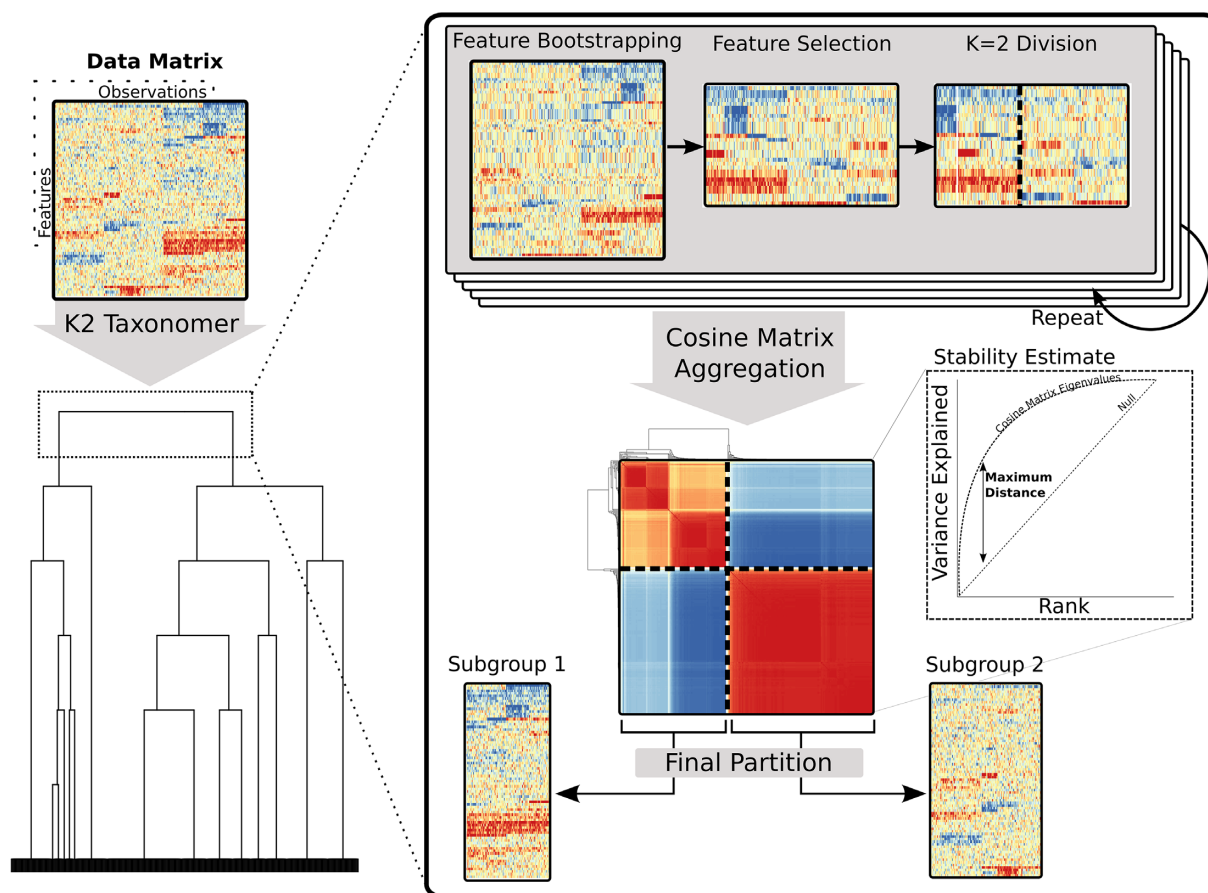
Figure 1. Schematic of the *K2Taxonomer* recursive partitioning algorithm. For each partition, *K2Taxonomer* generates an ensemble of $K = 2$ estimates from the feature bootstrapped data followed by variability-based feature selection. This ensemble is aggregated to a cosine matrix followed by hierarchical clustering and tree cutting. A stability estimate, indicative of the consistency of $K = 2$ estimates, is calculated based on an eigendecomposition of the cosine matrix. See supplementary methods for a more thorough description of the elements of this procedure.

The *K2Taxonomer* package includes options to perform both SD and MAD based feature selection. In the case of group-level analysis, *K2Taxonomer* can perform F-statistic based feature selection based on the ratio of between-group to within-group variability implemented by the *limma* R package (35).

K2Taxonomer data partitioning

To estimate each partition, *K2Taxonomer*, performs feature-level bootstrap aggregation, similar to that of consensus clustering (36). More specifically, each data partition represents the aggregation of a set of partitions estimated from perturbations of the original data set in which features have been sampled with replacement. Feature selection and $K = 2$ clustering are independently performed within each perturbation-specific data set. The final partition estimate is calculated by aggregating the set of perturbation-specific partitions into a *cosine similarity* matrix (defined below), which further undergoes hierarchical clustering, followed by a $K = 2$ tree cut.

K2Taxonomer package implements separate clustering methods tailored to analysis of either observation-level and group-level data input. For observation-level data, the

perturbation-specific partitions are estimated via hierarchical clustering of the Euclidean distance matrix, followed by a $K = 2$ tree cut. By default, Ward's agglomerative method is performed at this step because it has been shown to generally perform well compared to other hierarchical methods (22). For group-level data, perturbation-specific partitions are estimated via constrained K-means clustering (26). This algorithm performs semi-supervised clustering, in which group-level information is included as a pairwise 'must-link' constraint, preserving relationships between observations from the same group.

To assess the robustness of the partitioning of the aggregated results, hereby referred to as partition stability, as well as to facilitate interpretability, a cosine similarity matrix is computed, with each pairwise cosine similarity measurement functionally equivalent to the Pearson correlation of standardized variables.

Let an 'item' denote a single observation or group, depending on whether observation- or group-level analysis is being performed, respectively. The cosine similarity of two items is a measure proportional to the number of times across perturbation iterations that the two items are assigned to the same group in the perturbation-specific dichotomous partitions. It takes its maximum/minimum

value when the two items are always/never assigned to the same group.

If we represent with '−1' and '1' the assignments of an item to one or the other group in a dichotomous partition, we can then represent, and compare, the complete set of assignments of any two items across p perturbation-specific partitions as the vectors

$$X_i = (x_{i1}, ..., x_{ip}|x_i \in \{-1, 1\})$$
$$X_j = (x_{j1}, ..., x_{jp}|x_j \in \{-1, 1\}),$$

where X_i and X_j represent the ith and jth item, respectively.

We can then define the cosine similarity of X_i and X_j as

$$CS(X_i, X_j) = \frac{X_i \cdot X_j}{||X_i|| ||X_j||}.$$

where $X_i \cdot X_j$ represents the dot product and $||X_i|| ||X_j||$ represents the product of the Euclidean norms of X_i and X_j. Next, we prove the equivalence of the cosine similarity and Pearson correlation of two assignment vectors. The cosine similarity can be rewritten as follows:

$$CS(X_i, X_j) = \frac{X_i \cdot X_j}{||X_i|| ||X_j||} = \frac{(p - Z) - Z}{p}$$
$$= 1 - \frac{2Z}{p}.$$

In the above derivation, since X_i and X_j only take values in $\{-1,1\}$, their dot product, $X \cdot Y$, is equal to the difference between the number of iterations, Z, the two items are assigned to the same group, and the number of iterations, $p - Z$, the two items are assigned to different groups. Furthermore, the product of the Euclidean norms of X_i and X_j, $||X_i|| ||X_j||$, is equal to p.

Similarly, taking advantage of the relationship between Pearson correlation of standardized variables, $r()$, and Euclidean distance, $d()$, we have

$$r(X_i, X_j) = 1 - \frac{d^2(X_i, X_j)}{2p} = 1 - \frac{(X_i - X_j)^2}{2p}$$
$$= 1 - \frac{4Z}{2p} = 1 - \frac{2Z}{p},$$

where we used the fact that the squared Euclidean distance of X_i and X_j, $d^2(X_i, X_j)$, is equal to $(X_i - X_j)^2$. Furthermore, for $X_i - X_j$, the difference between mismatched adjacent elements is 2 and the difference between matched adjacent elements is 0. Therefore, $(X_i - X_j)^2 = 4Z$.

The function, $1 - \frac{2Z}{p}$, is the Hamann similarity index (37). Consistent with Pearson correlation, the range of possible values for the Hamann similarity is between 1 and −1; these extremes occur if Z is equal to p and 0, respectively, indicating that X_i and X_j are either identical or fully dissimilar. Furthermore, if elements of X_i and X_j share 50% of their matching assignments, then $Z = \frac{p}{2}$ and the Hamann similarity is equal to 0, indicating a lack of a relationship between their perturbation-specific partition estimates. This is not true for the related phi similarity (38), another correlation metric for dichotomous variables, the calculation of which includes adjustment for the marginal distribution of X_i and X_j. In this case, the marginal distribution of X_i and X_j is irrelevant because the elements of X_i and X_j are only meaningful in relation to their matching assignments.

K2Taxonomer **partition stability**

To assess the robustness of partition estimates, indicating the consistency of the perturbation-specific partition results, we developed a partition stability metric, which is calculated using the eigen-decomposition of the matrix of pairwise cosine similarities, Q, of dimension, N, the number of items. The eigen-decomposition of Q satisfies

$$QU = diag(\Lambda)U,$$

where U is the matrix of eigenvectors corresponding with Λ, the vector of rank-ordered eigenvalues

$$\Lambda = (\lambda_1, ..., \lambda_k, ..., \lambda_N | \lambda_k > \lambda_{k+1}).$$

Each eigenvalue is proportional to the 'variance explained' by each eigenvector, such that the cumulative sum of variance explained by the first k eigenvectors, v_k, is given by

$$v_k = \frac{\sum_{l=1}^{k} \lambda_l}{N}.$$

In this context, the variance explained by eigenvectors captures the consistency with which pairs of items received the same or different assignments across perturbation-specific partitions. Therefore, we can summarize this consistency by evaluating the difference between the variance explained by the eigenvectors of the estimated cosine matrix and the variance explained by these eigenvectors if there was no consistency across perturbation-specific partition assignments. We denote this deviation as the partition stability, PS, calculated as the maximum difference between v_k and $\frac{k}{N}$, the null value corresponding to all items being linearly independent,

$$PS = \max_k \left(v_k - \frac{k}{N} \right).$$

The possible values for the partition stability range between $1 - \frac{1}{N}$ and 0, with the former representing the case in which $\lambda_1 = N$, where every perturbation-specific partition is identical, i.e., all values of the cosine matrix are either −1 or 1. Conversely, a partition stability of 0 represents the case when the perturbation-specific partition assignments are random, i.e., all values of the cosine matrix are close to 0. The maximum value for a given partition is dependent on the number of items in the partition, approaching 1 when N is large, and equal to 0.5 when $N = 2$. Using the *K2Taxonomer* package, partition stability can be used to set stopping criteria for creating new partitions, thereby serving as a way to control the number of terminal subgroups without prior knowledge.

Finally, partition stability is used as a heuristic for calculating branch heights in dendrogram creation of the *K2Taxonomer* output. For a series of m partitions resulting in a given partition, z_m, the branch height, h_m, is

calculated as

$$h_m = \log(N_m) + \sum_{l=1}^{m} \log(PS_l) + c,$$

where c is a constant added to ensure that the minimum height for a node is equal to 1.

K2Taxonomer R package functionalities

In addition to running the recursive partitioning algorithm, the *K2Taxonomer* R package provides functionalities for comprehensive annotation of the estimated subgroups, via subgroup-level statistical analyses, including: differential analysis, gene set enrichment analysis, and phenotypic variable testing. Differential analysis of gene expression is carried out using the *limma* R package, which is well-suited to the analysis of normally distributed data such as microarray gene expression, as well as log-transformed and normalized RNAseq data (35). Gene set enrichment analysis is carried out on a set of user-provided gene sets and implemented in two ways: over-representation analysis based on a hypergeometric test, and differential analysis of single-sample gene set projections scores based on the *GSVA* R package (39). Finally, phenotypic variable testing is carried out on user-provided variables labeling individual observations or groups, supporting both continuous and categorical variables. Testing of association between continuous variable and taxonomy subgroups can be performed based on the parametric Student's *t*-test or the nonparametric Wilcoxon rank-sum test, while categorical testing is carried out using Fisher's exact test. All subgroup-level statistical analyses are corrected for multiple hypothesis testing based on the FDR procedure (40). The full set of results are compiled into an interactive-web portal for exploration and visualization. Differential analysis comparisons are carried out at the partition-level, i.e., comparing only the two subgroups at a particular node. However, the web portal includes functionality for performing post-hoc differential analysis of any combination of user-selected subgroups.

Statistical analysis

The implementation of *K2Taxonomer* for this manuscript was run with *R (v3.6.0)*, *limma (v3.42.2)*, and *GSVA (v1.34.0)*. All *P*-values reported in this manuscript are two-sided.

Simulation-based performance assessment

The performance of *K2Taxonomer* was assessed in comparison to Ward's agglomerative method for recapitulating induced hierarchical structure in simulated data. See supplementary methods for a comprehensive description of the strategy implemented for data generation and performance assessment for observation- and group-level analyses.

Performance assessment using breast cancer primary tumor bulk gene expression data

K2Taxonomer was evaluated in its observation-level modality for its ability to recover the Pam50 subtypes, as well ER-,

PR-, and HER2-status, and the aggregate three-gene genotype of ER-, PR-, and HER2-status, in the METABRIC and TCGA breast cancer datasets, independently (41,42). Distribution of these variables, as well as additional clinical variables for METABRIC and TCGA data sets are summarized in Supplementary Tables S1 and S2, respectively. See supplementary methods for a detailed description of these two data sets, as well as methods for acquisition and preprocessing. *K2Taxonomer* was also compared to two agglomerative clustering algorithms, Ward's and average. These specific methods were chosen because they have been previously shown to outperform other common agglomerative methods (22,43). Given the sensitivity of hierarchical clustering to the level of feature filtering, analyses included individual runs on four filtered data subsets of the total number of features: 100%, 25%, 10% and 5%, while *K2Taxonomer* was only run on the full set (i.e. 100%) of features. This should be kept in mind when comparing performances, since the best-performing pre-filtering level is not known a priori, and it is in general dataset dependent. For every pre-filtering level, the median absolute deviation (MAD) score was used for feature selection, and Euclidean distance was used to estimate observation-level distance. Performance was assessed as the entropy of each of the phenotypes (e.g., PAM50 labels), induced by the inferred sample sub-grouping, with lower entropy indicating 'purer' subgroups, hence better performance (44). The different methods were evaluated and compared by the relative decrease in entropy as the number of mutually exclusive clusters, *K*, increased from 2 to 8 based on tree cuts of the dendrograms produced by each model.

Healthy airway tissue scRNAseq gene expression analysis

K2Taxonomer was applied in its group-level modality to partition scRNAseq gene expression profiles of 28 estimated cell types from a publicly available data set of airway tissue of healthy patients (28), was evaluated against the known relationships among the included cell types, and was compared to the partitioning obtained by two agglomerative methods, Ward's and average, as well as *PAGA*, a graph-estimation and trajectory inference algorithm (20), and *ascend*, a fully unsupervised scRNAseq hierarchical clustering algorithm (24). See supplementary methods for a detailed description of this data set, as well as methods for acquisition and preprocessing. For agglomerative methods, cell type-level data processing and feature selection were performed consistent with the results of group-level analysis of simulated data (see supplementary methods). As with agglomerative methods, *PAGA* was run on F-statistic-based pre-filtered feature sets at different percentages of the full 18,417 genes: 100%, 25%, 10% and 5%. Due to limitations of computational resources, *ascend* analyses excluded those based on 100% of genes. Moreover, *PAGA* and *ascend* were each applied using three sets of the top principal components: 10, 15 and 20, such that a total of 12 and 9 unique *PAGA* and *ascend* models were estimated, respectively.

Analyses with *PAGA* were carried out using the *scanpy* (v1.6.0) python package. Neighborhood graphs were estimated from the component-based dimensionality reduced data sets using 'pp.neighbors()'. Finally, the *PAGA* algo-

rithm was run, setting the 'group' argument of 'tl.paga()' to the cell type labels.

Analyses with *ascend* were carried out using the *ascend* (v0.99.69) R package. The *ascend*-based hierarchical models were estimated from the component-based dimensionality reduced data sets using the 'runCORE()' function. Given that, unlike *K2Taxonomer* and *PAGA, ascend* can only be carried out as a fully unsupervised analysis without group-level input, concordance of these models with those of other methods was evaluated based on purity-based assignment of cell type labels to individual subgroups of each *ascend* model. Cell types were assigned to individual partitions based on two criteria: first, if at least 75% of the profiles of a given cell type were assigned to the subgroup and, second, if the total number of cell profiles across individual cell types that met the first criteria made up at least 75% of the total cell profiles in the subgroup.

Breast cancer immune cell scRNAseq gene expression analysis

Publicly available scRNAseq gene expression of raw counts from immunocytes of two TNBC patients was obtained from *GEO*, accession number GSE110938 (29). The data was processed in accordance with the original manuscript (29), recapitulating the reported 5,759 individual cells, 4,844 and 915 from either sample, with 15 623 genes passing QC criteria, selection of 1675 highly variable genes, and 10 latent variables estimated by *ZINB-WaVE (v1.8.0)* (45). To enable exploration of the data at finer resolution, clustering of the latent variables with *Seurat (v1.3.4)* was modified by setting the 'resolution' argument of 'FindClusters()' to 1.1, rather than the default, 0.8 (46). This resulted in 13 estimated cell clusters. Of the 10 cell clusters reported in the original manuscript, two cell clusters, 'CD4+ FOXP3+' and 'CD4+ IL7R+', were further split into three and two individual clusters, respectively (Supplementary Table S3).

K2Taxonomer was applied in its group-level modality to partition these 13 estimated cell subtypes based on the normalized count matrix estimated by *ZINB-WaVE*. According to the developers, *ZINB-WaVE*, normalized count estimates are not recommended for differential analysis (45), hence differential analysis was performed based on drop-out imputed and batch-corrected normalized counts estimated using the *bayNorm (v1.4.14)* R package (47). Pathway-level analysis was carried out using Reactome gene sets downloaded from *mSigDB (v7.0)* (48). Signatures of up-regulated genes were derived from each subgroup based on their FDR corrected *P*-value (FDR < 1e−10) and minimum subgroup-specific expression, (mean[log$_2$ counts] > 0.5), then restricted to a maximum of 50 genes.

To investigate the concordance between *K2Taxonomer* estimated model and that estimated by a fully unsupervised scRNAseq hierarchical clustering method, we ran *ascend* on the same *ZINB-WaVE* reduced dimension data set that was used as input for *Seurat*. Subgroup-level cell cluster assignments in the resulting *ascend* estimated model was carried out using the same criteria employed for the healthy airway tissue scRNAseq analysis.

To validate the clinical relevance of signatures of TILs derived by *K2Taxonomer*, we performed survival analysis based on gene signature projection scores, as well as on selected genes in the METABRIC breast cancer primary tumor gene expression data set. Gene set projection was carried out using *GSVA* (39). Multivariate Survival analysis was performed using Cox proportional hazards tests. All models included age and Pam50 subtype as covariates. To account for possible confounding effects of inflammation and proliferation, we generated separate patient-level activity scores for each, using a gene set projections of published signatures of deleterious breast cancer inflammation markers (49) and breast cancer proliferation (50) (Supplementary Table S4).

RESULTS

We report here our extensive evaluation of K2Taxonomer on both simulated and real data. We evaluated the method's performance both on data where the analysis end-points were single samples (*observation-level*), and groups of samples (*group-level*). The latter corresponds to scenarios where the goal is to define a taxonomy over sample groups, such as cell types in single-cell experiments, or chemical perturbations profiled in multiple replicates (51).

K2Taxonomer discovers hierarchical taxonomies on simulated data

We first evaluated *K2Taxonomer*'s capability to recapitulate hierarchical relationships induced in simulated data, as measured by the Baker's gamma coefficient estimate of similarity between the structure of two dendrograms, where the structure of each dendrogram is quantified by a matrix based on the number of partitions separating each pair of leaves (52). As a term of reference, we compared *K2Taxonomer*'s performance to Ward's agglomerative method.

For observation-level analysis, *K2Taxonomer* demonstrated robust performance for moderate levels of background noise, e.g., standard deviations equal to 0.5 and 1.0, regardless of the proportion of features with signal or the number of terminal clusters (Figure 2a, Supplementary Table S5). For higher levels of background noise, e.g., standard deviations equal to 2.0 and 3.0, the performance of *K2Taxonomer* was more dependent on either parameter, performing better with more features with signal and fewer terminal clusters. *K2Taxonomer* significantly outperformed Ward's method in 221 out of the 400 combinations of parameters tested (FDR < 0.05), while Ward's performed better for 17 combinations (Figure 2A, Supplementary Table S5). Furthermore, the differences between Baker's gamma coefficients for the 221 results significantly in favor of *K2Taxonomer* were generally larger than the 17 results significantly in favor of Ward's method, with median differences of 0.14 and 0.03, respectively. In general, *K2Taxonomer* significantly outperformed Ward's method when the background noise, the number of terminal groups and percent features with signal increased.

Remarkably, for group-level analysis *K2Taxonomer* outperformed Ward's method for all 400 combinations of variables tested (Figure 2b, Supplementary Table S6).
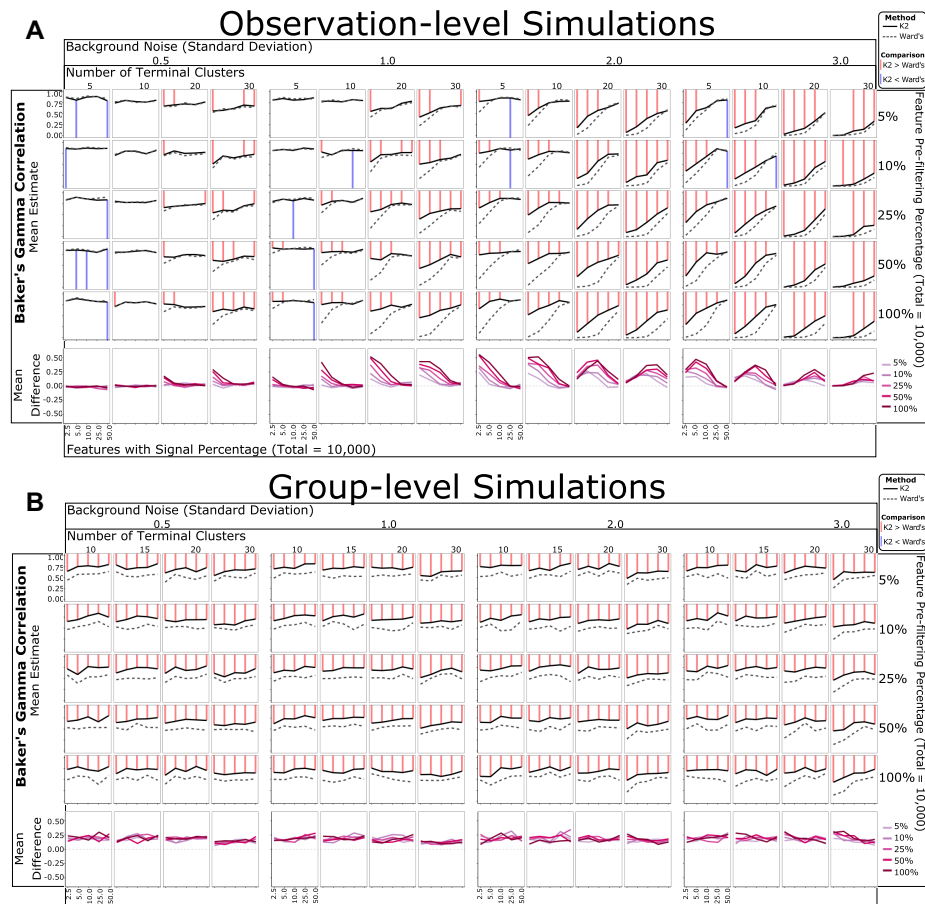
Figure 2. Simulation-based performance assessment of *K2Taxonomer* and Ward's agglomerative method. Mean Baker's gamma correlation estimates measuring the similarity of either *K2Taxonomer* and Ward's agglomerative method estimates to the true hierarchy from which the simulated data was generated. Each combination of parameters was simulated 25 times. The red and blue lines are indicative of statistically significant differences between the correlation estimates (FDR < 0.05) based on a Wilcoxon signed-rank test. (**A**) Observation-level analyses with 300 observations and 10,000 features. (**B**) Group-level analyses with 1000 observations and 10,000 features.

Using the square root of the total number of features as the partition-specific feature filtering parameter for running *K2Taxonomer* demonstrated stable performance. When compared to selecting a fixed percentage of the total number of features (Supplementary Figure S1, Supplementary Table S7), the square root outperformed larger percentages when the number of features was large, and outperformed smaller percentages when the number of features was small.

K2Taxonomer **accurately sorts breast cancer subtypes without pre-filtering of features**

We evaluated *K2Taxonomer*'s ability to sort Pam50 subtypes, ER-status, PR-status, and HER2-status from bulk gene expression data from METABRIC and the TCGA BRCA bulk gene expression data, separately. A fourth variable, defined by the Cartesian product of ER-status, PR-status and HER2-status was also assessed. Performance was assessed in terms of the decrease in entropy as the number of cluster estimates, K, increased from 2 to 8 (Figure 3A and B). We also compared *K2Taxonomer*'s performance to two agglomerative clustering methods, Ward's and average. Since standard hierarchical clustering is sensitive to the level

of feature filtering, the comparison was repeated for multiple pre-filtering levels.

In general, *K2Taxonomer* accurately segregated the known subtypes and phenotypes, performing as well or better than either method (Figure 3B, Supplementary Tables S8 and S9). When applied to the METABRIC data, *K2Taxonomer* analysis yielded the lowest entropy score compared to all other methods for $K = 3$ and higher with few exceptions. Other methods produced similar entropy measurements at selected higher levels of K. For example, Ward's method resulted in similar entropy scores for Pam50 subtypes and HER2-status at $K = 4$ and $K = 5$, respectively, but for different pre-filtering levels, 5% and 100%, respectively. When applied to the TCGA BRCA data, the difference in performance was less pronounced. *K2Taxonomer* resulted in the lowest entropy score for Pam50 scores, genotype, ER- and PR-status for $K = 4$. Ward's method at 5% pre-filtering level produced the smallest entropy score for HER2-status for $K = 4$.

It should be emphasized that the pre-filtering level to be used with hierarchical clustering is not known a priori, and it would thus preclude us in practice from selecting the level yielding the best results shown in the above comparison.
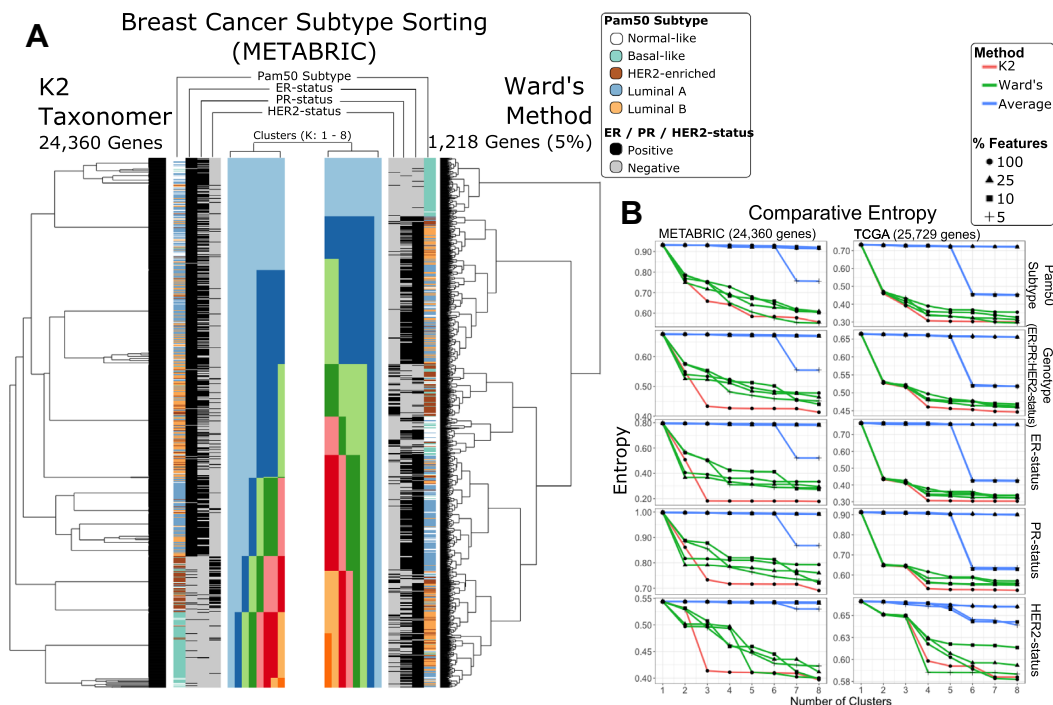
Figure 3. Breast cancer subtyping performance assessment of bulk gene expression data. Comparison of sorting of breast cancer Pam50 subtypes and genotypes (ER-, PR- and ER-status) for two bulk gene expression data sets, METABRIC and TCGA. An aggregate, three gene genotype status was also included by combining the individual genotypes. Performance was assessed based on reduction of entropy as the number cluster estimate increased based on tree cutting. *K2Taxonomer* was only run on the full set of features, while either agglomerative method, average and Ward's, were run on three additional subsets of the data. (**A**) Illustration of the results generated by *K2Taxonomer* and Ward's method for the METABRIC dataset. These results reflect Ward's method run on 5% of the total number of features, which demonstrated the best performance among agglomerative methods. (**B**) Entropy measurements for each method as *K* increased across the METABRIC and TCGA data sets.

K2Taxonomer accurately identifies and organizes subgroups of shared progenitors and epithelial cells from healthy airway scRNAseq cell clusters

To assess the capability of *K2Taxonomer* to recapitulate biologically relevant subgroupings of cell types estimated from scRNAseq data, we ran group-level analysis using 29 cell types estimations assigned to 77,969 cells of airway tissue from 35 samples across 10 healthy subjects and multiple locations (Figure 4A, B) (28). In addition to agglomerative methods, we evaluated *K2Taxonomer's* performance against that of partition-based graph abstraction (PAGA) (20), which has been shown to outperform similar methods, especially for analyzing large-scale scRNAseq data sets (18).

K2Taxonomer was remarkably accurate in capturing the higher-order organization of the 28 cell types. The first partition separated all epithelial cell subtypes from non-epithelial cell types (Figure 4B). Further partitioning of the 17 epithelial cell subtypes yielded five subgroups, characterized by shared morphology, labeled as 'Ciliated', 'Basal', 'Submucosal', 'Brush-like' and 'AE' (Alveolar Epithelium). The 'Ciliated' subgroup was comprised of differentiated multiciliated cells and their precursor, deuterosomal cells (28). The 'Basal' subgroup was comprised of epithelial cell progenitors, basal and cycling basal cells, as well as epithelial cell intermediary, suprabasal cells (53). The 'Brush-like' subgroup was comprised of three rare airway epithelial cell types: brush, ionocyte, and pulmonary neuroen-

docrine cells (PNECs), as well as their likely shared 'brush-like' precursor cells (54), as suggested in the original publication of these data (28). Further partitioning of the 11 non-epithelial cell types yielded two main subgroups characterized by shared progenitor cells: immune (55) and mesenchymal stem cells (56), with the immune cell subgroup also including endothelial cells. Further partitions of the immune cell subgroup included progeny of monoblasts: monocytes, dendritic cells, and macrophages (57), followed by endothelial cells separated from all non-monoblast progeny immune cells subtypes in the adjacent subgroup.

In contrast, agglomerative hierarchical clustering of these cell clusters, even if evaluated at multiple *F*-statistic-based pre-filtering levels, yielded significantly different results poorly reflective of the known taxonomic cell type organization (Figure 4C, Supplementary Figure S2). While the mesenchymal stem cell subgroup, comprised of fibroblasts, smooth muscle, and pericytes, was identified by Ward's method, and while there were other instances of concordant subgroups, none of these consisted of more than two cell types.

PAGA-based trajectory analysis (20) of the cell types performed better than agglomerative clustering, and demonstrated both improvements and drawbacks compared to *K2Taxonomer*. (Figure 4D, Supplementary Figure S3). Figure 4D shows what we selected to be the 'best' PAGA result for runs on a grid of F-statistic-based pre-filtering levels (25% of genes) and principal component-based dimensionality reduction size (15 principal components).
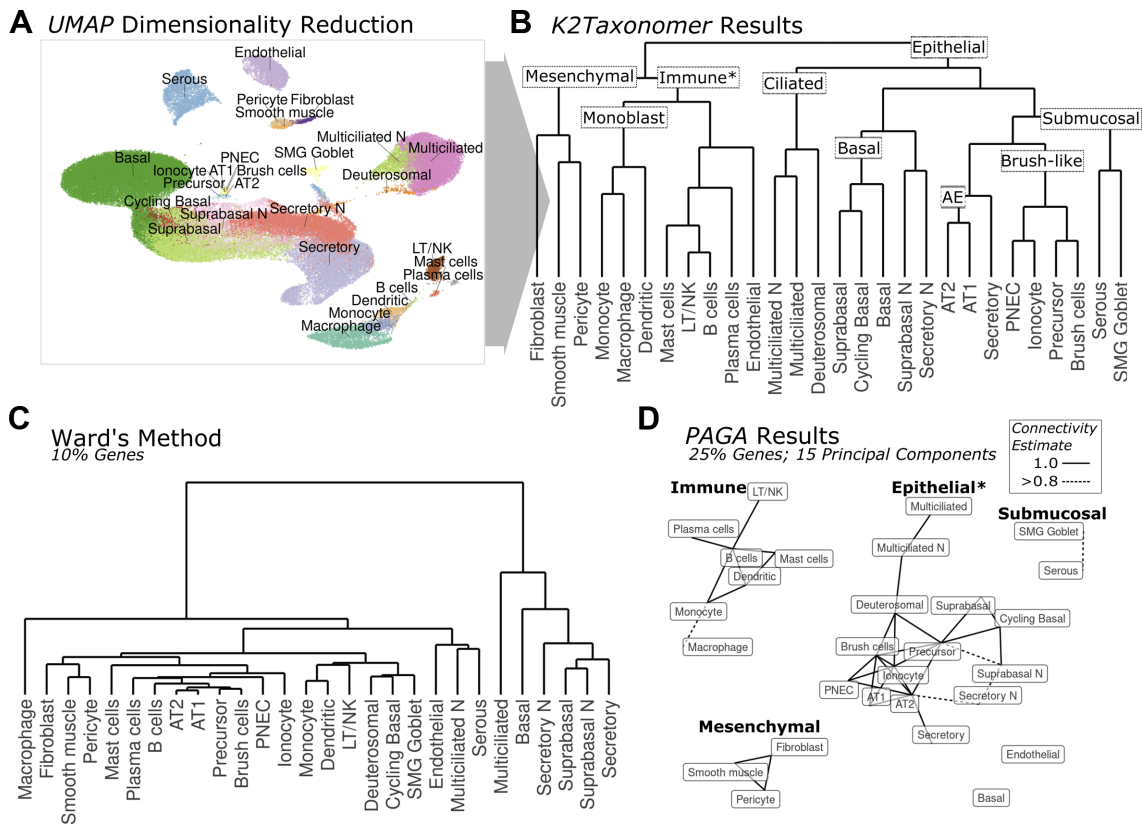
Figure 4. Subgrouping of healthy airway cell types from scRNAseq data. (**A**) tSNE dimensionality reduction of healthy airway scRNAseq data with labels for 28 cell types annotated by (29). Note cell types labelled ending in '*N*' indicate those which only included cells from nasal samples. (**B**) *K2Taxonomer* results with nine annotated subgroups. The '*' in the 'Immune' subgroup label indicates the impurity of this subgroup caused by the presence of endothelial cells. (**C**) Ward's agglomerative clustering results for a selected analysis performed on 10% of the total number of genes. The results for both Ward's method and Average method run on additional gene subsets: 100%, 25%, 10% and 5%, are shown in Supplementary Figure S2. (**D**) *PAGA* graph-based trajectory results for a selected analysis performed on 15 principal components estimated from 25% of the total number of genes with four annotated disconnected subgroups. The '*' in the 'Epithelial' subgroup label indicates the incompleteness of this subgraph caused by the absence of basal cells and submucosal cells. Edges indicate *PAGA* connectivity estimates >0.8. The results for analyses run on additional numbers of principal components: 10, 15 and 20, as well as additional gene subsets: 100%, 25%, 10% and 5%, are shown in Supplementary Figure S3.

In this case, *PAGA* recapitulated four subgroups identified by *K2Taxonomer* as disconnected subgraphs: 'Immune', 'Epithelial', 'Submucosal', and 'Mesenchymal'. Unlike *K2Taxonomer*, this model accurately segregated endothelial cells as its own disconnected vertex. However, all *PAGA* models failed to include basal cells in the epithelial cell subgraph (Figure 4D, Supplementary Figure S3). Furthermore, PAGA segregated the histologically separated submucosal cells lines: SMG goblet and serous, from the other epithelial cell lines. The most apparent difference between the *K2Taxonomer* and *PAGA* results is the high connectedness of the *PAGA* subgraphs, especially considering that these are all 'near perfect-confidence' connections. As a result, subgroup relationships within these subgraphs are difficult to distinguish, and approaches to estimate tree-like graphs, such as minimum spanning tree algorithms, yield multiple equivalent solutions. Finally, the *PAGA* results varied based on the feature pre-filtering level and the number of principal components, most notably in the estimations of connections between immune cells and epithelial cells (Supplementary Figure S3).

While models generated by *ascend*-based hierarchical clustering were also able to recapitulate some of the subgroups, we did not observe any improvements compared to *K2Taxonomer* or *PAGA* (Supplementary Figure S4). Whereas the first *K2Taxonomer* partition segregated all epithelial cell types, each of the *ascend* models separated multiciliated cells from all other cell types at the first split. Moreover, in six out of nine of these models, the multiciliated cells did not subgroup with deuterosomal cells. Like *PAGA*, the structure and subgroups identified by individual *ascend* models varied widely based on feature pre-filtering level and number of principal components, such that a 'best' model wasn't discernable. For example, the *ascend* model run with 25% of genes and 10 principal components was the only model capable to identify a purely 'Immune' cell subgroup, but did not include the 'Basal' or 'Monoblast' subgroups identified by some of the other models. In terms of interpretability the *ascend* models presented additional challenges. As expected, individual profiles with the same cell type label were scattered throughout the *ascend* models, such that a purity-based heuristic approach was used to assign cell type labels at individual partitions. As a result, in most cases individual cell type labels 'dropped-out' of models before segregating in cell type-specific subgroups. Moreover, this assignment procedure required considerable *ad hoc* manipulation of the objects generated by the *ascend* R package.

K2Taxonomer identifies subgroups of TILs characterized by differential regulation of TNF signaling, translation and mitotic activity from BRCA tumor scRNAseq cell clusters

We performed *K2Taxonomer* analysis on scRNAseq data of 13 TIL cell clusters reflecting further subdivision of the 10 cell types reported in the original study (29). The higher resolution was achieved by reproducing the reported methods (29) with the exception of selecting a higher resolution parameter when performing clustering with *Seurat* (14) (Supplementary Table S3). The full set of results is also available through the interactive portal at https://montilab.bu.edu/k2BRCAtcell.

The results of *K2Taxonomer* partitioning and annotation of breast cancer TIL cell clusters estimated from scRNAseq data is summarized in Figure 5a. Biologically informative subgroups, characterized by strongly significant differential expression of gene expression and sample-level pathway enrichment are highlighted and labeled within each boxed sub-dendrogram. The full set of differential results for genes and pathways across all partitions are reported in Supplementary Tables S10 and S11, respectively. Three distinct multi-cell subgroups emerged, labeled as: 'Trm All', 'CD4+ CCL5-' and 'Translation+', characterized by consistent up-regulation of PD-1 signaling (Reactome PD-1 signaling, FDR = 1.1e−241), translation (Reactome eukaryotic translation initiation, FDR = 5.6e-137), and TNF signaling (Reactome TNFS bind their physiological receptors, FDR ∼ 0.00), respectively (Figure 5B). 'Trm All' *and* 'Treg' subgroups each included a mitotic cell subgroup characterized by high cell cycle activity (Figure 5B). Furthermore, the 'CD4+ CCL5−' subgroup, comprised of the 'CD4+ CXCL13+' cell cluster and 'Treg' subgroup, is characterized by consistent down-regulation of CCL5 (FDR ∼ 0.00) and up-regulation of TNFRSF4 (FDR ∼ 0.00) (Figure 5C, D). Furthermore, additional up-regulation of TNFRSF4 (FDR = 1.1e−7) and RGS1 (FDR = 4.3e−55) distinguish non-mitotic 'Treg' subgroups (Figure 5C). Gene-level markers of the 'Translation+' subgroup included numerous ribosomal proteins, epitomized by up-regulation of RPS27 (FDR = 9.6e−246) (Figure 5C, D).

We compared the subgroup discovery of this model to that of *ascend* (24) (Supplementary Figure S5). The models differed substantially. Notably, unlike *K2Taxonomer*, *ascend*-based models did not segregate monocytes from T cell subtypes in the first partition and segregated mitotic Treg cells, *CD4+ FOXP3+ (3)*, from the two other Treg cell subtypes in the second partition. Moreover, the *ascend*-based model segregated individual profiles of *CD4+ Trm* cells, such that this cell type 'dropped-out' out of the model early in partitioning.

Confounding effects of inflammation and proliferation on the association between tumor infiltrating cell activity and patient survival

To assess the clinical relevance of *K2Taxonomer* annotation of single-cell immune cell subgroups, we performed survival analysis, via Cox proportional hazards testing, modeling the relationship between *K2Taxonomer* subgroup gene signature scores and patient survival in the METABRIC breast cancer bulk gene expression data set.

For these models, we examined two possible sources of confounding factors. First, inflammation has a well-described paradoxical role in breast cancer progression (58), such that the content of different subpopulations of lymphocytes has been associated with both better and worse prognosis (59). Given the physiological similarities between different lymphocyte subtypes (60), we hypothesized that expression patterns associated with tumor-promoting inflammation could mask those associated with tumor-suppressing TILs subsets. Second, we hypothesized that the signatures of the two mitotic T cell subgroups were similar enough to signatures of proliferative activity in non-immune tumor cells to result in a spurious association between T cell mitosis and worse prognosis. To assess and correct for these confounding effects, multivariate survival models were run without and with the inclusion of inflammation and proliferation scores as individual covariates. These patient-level scores were estimated by projecting published signatures of 'bad' inflammation (49) and proliferation (50), each of which had been previously reported to be associated with poor prognosis in breast cancer.

The results of each of these analyses are summarized in Figure 5A. The full set of survival results for unadjusted and adjusted models, including the genes belonging to each subgroup signature are reported in Supplementary Table S12. Controlling for inflammation and proliferation scores increased the overall significance of the association between subgroup-driven signatures of TILs and improved survival (hazard ratio < 1, FDR < 0.05). Furthermore, signatures of two cell subgroups, 'CD8+ mit. Trm' and 'Treg mit.', characterized by increased cell cycle activity (Figure 5B), were associated with worse patient survival in models ignoring inflammation and proliferation scores, but were subsequently statistically insignificant in models including these covariates, likely reflecting the effect of confounding by proliferation activity (Figure 5A). This is further illustrated in Figure 5E, which shows the 95% confidence intervals of hazard ratios of 'marginal' inflammation and proliferation models (left-most), as well as the confidence intervals of hazard ratios of select subgroups of cell subtypes, unadjusted and adjusted for inflammation and proliferation. Controlling for inflammation and proliferation allowed us to disentangle the contribution to survival of different components. For example, in the 'CD8+ mit. Trm' subgroup, we observed that the 'CD8+ mit. Trm' signature score was highly associated with worse patient survival in the unadjusted model, but the association became insignificant in the full model adjusted for proliferation and inflammation. On the other hand, there were instances where the hazard ratio achieved or improved significance (i.e., patient survival was significantly better) only after controlling for inflammation and proliferation in the full adjusted model, as observed in the 'Trm All' subgroup and, to a lesser extent, in the 'Translation+' subgroup (Figure 5E).

High expression of TNFRSF4, a marker for Treg cell activity is associated with worse survival when adjusting for CCL5 expression

TNFRSF4 and CCL5 were found to be the top two markers constitutively up- and down-regulated, respectively, within Treg subgroups, with TNFRSF4 the top marker further dis-
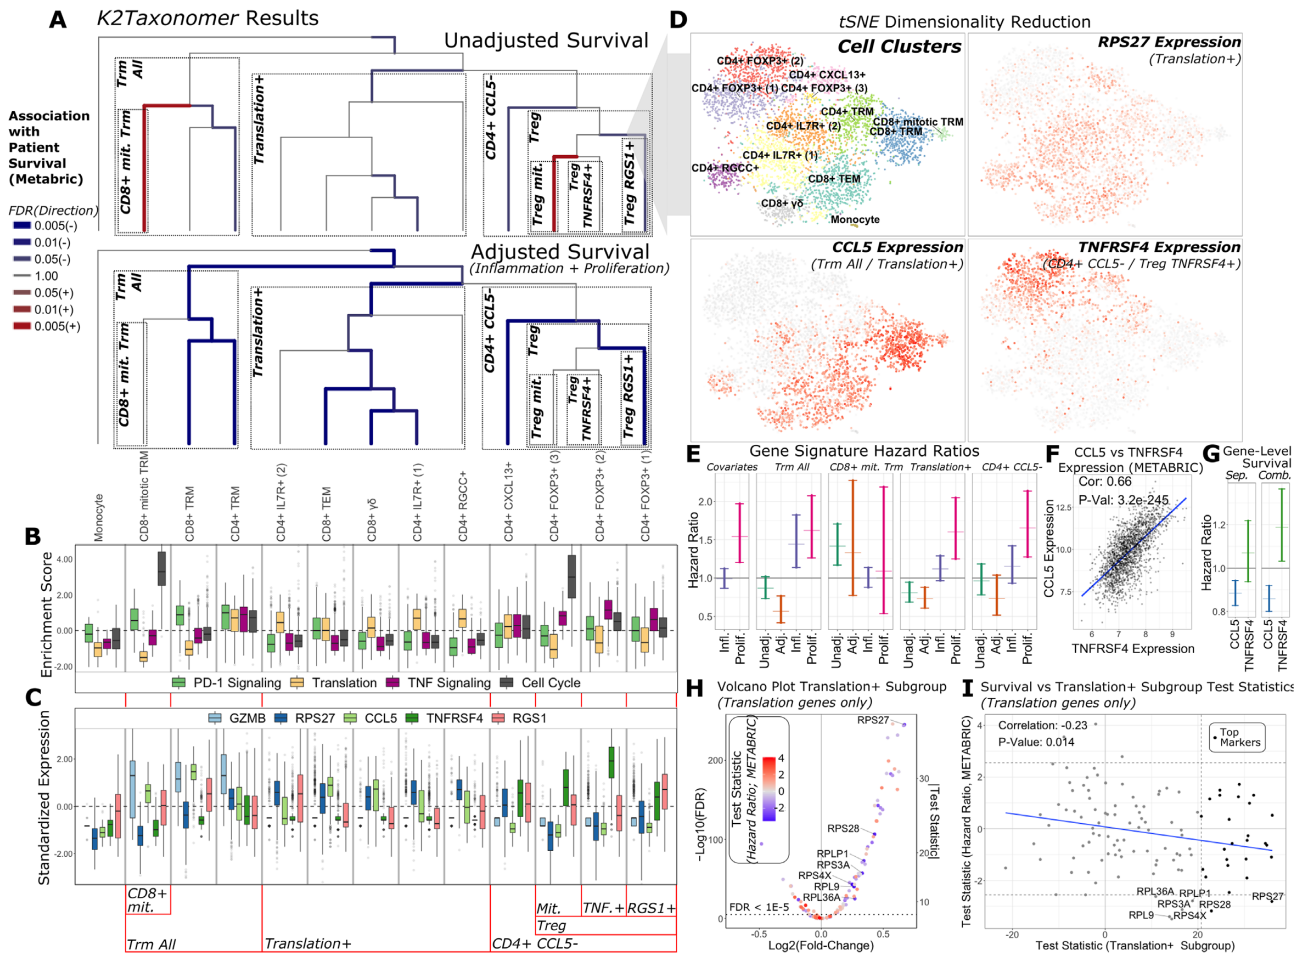
Figure 5. *K2Taxonomer* annotation of scRNA-seq clustering of breast cancer immune cell data and in-silico validation via patient survival on METABRIC breast cancer bulk gene expression data set. (**A**) *K2Taxonomer* annotation of 13 cell subtypes of breast cancer immune cell populations. Cell type labels are in accordance with the original publication (29). Color and thickness of each edge indicate direction and strength, respectively, of the association between the projected signature of up-regulated genes of each subgroup and patient survival in METABRIC breast cancer cohort via Cox proportional hazards testing. The top and bottom dendrograms show the results without and with adjustments of covariates for inflammation and proliferation. Blue and red are indicative of hazard ratio <1 and hazard ratio >1, respectively. All models included age and PAM50 subtype as covariates. (**B**) Boxplots of gene set projection scores of selected REACTOME pathways, enriched in subgroups of immune cells. These pathways include: PD-1 Signaling, enriched in the *Trm All* subgroup, Translation, enriched in the *Translation+* subgroup, TNF Signaling, enriched in the *CD4+ CCL5-* and *Treg TNFRSF4+* subgroups, and Cell Cycle, enriched in the *CD8+ mit. Trm and Treg mit. Subgroups*. The center line, hinges, and whiskers indicate the median, interquartile range, and extreme values truncated at 1.5 * the interquartile range, respectively. (**C**) Boxplots of markers constitutively regulated in selected *K2Taxonomer* subgroups. GZMB is upregulated in the *Trm All* subgroup. CCL5 and TNFRSF4 are up- and down-regulated, respectively, in the *CD4+ CCL5−* subgroup. TNFRSF4 is further up-regulated in the *Treg TNFRSF4+ subgroup,* while *RGS1+* is up-regulated in the *Treg RGS1+* subgroup. Finally, RPS27 is up-regulated in the *Translation+* subgroup. The center line, hinges and whiskers indicate the median, interquartile range and extreme values truncated at 1.5 * the interquartile range, respectively. (**D**) tSNE dimensionality reduction of the single-cell breast cancer immune cell data, indicating the cell subtype label assignment of every cell, as well as Z-scored expression of selected genes from C. (**E**) 95% confidence intervals of hazard ratios from Cox proportional hazards testing of gene set projections of cellular subgroups on the METABRIC data set. *Covariates* shows the results of the survival model of sample-level inflammation and proliferation scores without a *K2Taxonomer* derived signature. Every other model shows the confidence interval of the subgroup-specific model without and with adjusting for inflammation and proliferation score, as well as the confidence intervals of inflammation and proliferation in the full model. All models included age and Pam50 breast cancer subtype as covariates. (**F**) Comparison of the expression of CCL5 and TNFRSF4 expression in the METABRIC dataset. (**G**) 95% confidence intervals of hazard ratios from Cox proportional hazards testing of gene-level expression of CCL5 and TNFRSF4, modelled separately, *Sep.,* and combined in a single model, *Comb.* These models also included age, Pam50 breast cancer subtype, as well as sample-level inflammation and proliferation score as covariates. (**H**) Volcano plot of differential expression analysis of the *Translation+* subgroup in scRNAseq data of individual genes in the REACTOME eukaryotic translation initiation gene set. An alternative coding of the *y-axis* indicating the absolute value of the test statistic is shown on the right side of the plot. The colors indicates the association of each gene with survival in the METABRIC data set. Genes significantly associated with better survival (hazard ratios < 1, FDR < 0.1) are labelled. (**I**) Comparison of the association between survival and expression of the REACTOME eukaryotic translation initiation gene set (y-axis) and the test statistics indicating up-regulation in the *Translation+* subgroup (x-axis) in the METABRIC data set. Genes that were included as top markers of the *Translation+* subgroup are highlighted. Genes significantly associated with better survival (hazard ratios < 1, FDR < 0.1) are labelled. The blue line indicates the linear fit of these two variables.

criminating between the two non-mitotic Treg subgroups, *Treg TNFRSF4+* and *Treg RGS1+* (Figure 5A, C, D). Furthermore, their expression was highly correlated in the METABRIC data set (rho = 0.66, *P*-value = 3.2e−245) (Figure 5F), supporting a pattern of co-expression within TIL microenvironments. To assess whether TNFRSF4 and CCL5 expression levels could serve as markers for immuno-suppressive activity of Treg cells, we performed survival analysis of each gene modeled separately and in a combined model (Figure 5G). When modeled separately, the expression of TNFRSF4 is not associated with patient survival (*P*-value = 0.32), while CCL5 is associated with better patient prognosis (*P*-value = 1.94E−4). However, in the combined model both genes are associated with patient survival, with TNFRSF4 associated with worse patient survival (*P*-value = 0.015).

Taken together these results indicate that, in bulk gene expression data, markers of Treg cell activity are highly correlated with markers of overall tumor immune infiltration, confounding associations between expression of these markers and patient survival.

Up-regulation of specific translation genes characterizes a subgroup of TILs and is associated with better survival prognosis, independent of inflammation activity

The 'Translation+' subgroup was a notable instance where the subgroup-specific signature projection was associated with better patient survival, regardless of adjustment for inflammation and proliferation (Figure 5A, E). To assess the extent to which up-regulation of translation-specific genes in this subgroup associated with better patient prognosis, we ran separate survival analysis for each of the 112 genes from the Reactome Eukaryotic Translation Initiation gene set, which were shared between the single-cell BRCA gene set and METABRIC data set. Of the 112 genes, 61 were up-regulated in the 'Translation+' subgroup (FDR < 1E−5), including 26 genes within the top 50 marker 'Translation+' subgroup signatures (Figure 5H, I). The full set of survival results for these 112 genes are shown in Supplementary Table S13. The test statistics derived from single gene Cox proportional hazards models were negatively correlated with the corresponding genes' test statistics of their up-regulation in the 'Translation+' subgroup (rho = −0.23, *P*-value = 0.014) (Figure 5I). Furthermore, seven of the 112 genes were associated with better patient survival (FDR < 0.1). All of these genes were significantly up-regulated in the 'Translation+' subgroup. Of these seven genes, RPL36A had the minimum 'Translation+' subgroup-associated test statistic (FDR = 1.34e−25) and two (RPS28 and RPS27) were members of the top 50 markers, comprising the 'Translation+' subgroup signature. RPS27 was the top translation gene associated with the 'Translation+' subgroup (FDR = 9.6e−246).

In summary, we employed *K2Taxonomer* to characterize the up-regulation of translational machinery as the dominant transcriptional program shared by a diverse subgroup of TILs. These findings informed additional analyses, which demonstrated that association between expression of translational machinery genes and better patient survival tracked with their over-expression in a subgroup of TILs. Taken together these findings suggest that the up-regulation of specific translational machinery genes is widespread across TILs, serving as a predictor of the level of immune infiltration in breast cancer tissue from bulk gene expression data.

DISCUSSION

In this manuscript, we presented extensive assessment and practical applications of *K2Taxonomer*, a novel unsupervised recursive partitioning algorithm for taxonomy discovery in both bulk and single-cell high-throughput transcriptomic profiles. An important distinctive feature of the algorithm is that each partition is estimated based on a feature set selected to be most discriminatory within that partition, thus permitting the use of large sets of features to be used as input, without pre-filtering or dimensionality reduction approaches. Additionally, to minimize generalization error, each partition is based on an ensemble (61) of partition estimates from repeated perturbations of the data. The adoption of an ensemble approach also makes it possible to compute a stability measure for each partition, which can be used to assess the robustness of each partition, as well as a stopping criterion for limiting the number of subgroup estimates. Finally, to facilitate comprehensive exploration of the results, as well as to share results for independent interrogation, the *K2Taxonomer* R package includes functionalities to automatically generate interactive web-portals. One of these portals was utilized extensively to annotate subgroups as part of our analysis of breast cancer TILs in scRNAseq data and is publicly available (https://montilab.bu.edu/k2BRCAtcell/).

Our extensive evaluation and benchmarking of the methods on simulated data showed *K2Taxonomer*'s high accuracy, its superior performance when compared to representative other methods, and its capability to (re)discover known nested taxonomies. As we have shown in its multiple applications, *K2Taxonomer* may be applied in a fully unsupervised mode to partition individual-level data, or it can take group-level labels as input to estimate inter-group relationships among the known groups. In the group-level settings, *K2Taxonomer* is not directly comparable to fully unsupervised 'flat' or hierarchical clustering methods. In our simulations, the comparison to Ward agglomerative clustering was attained by averaging the profiles within each simulated 'group' and by then clustering these 'average profiles.' Side-by-side comparison with any of the single cell algorithms (23–25) on simulated data was thus deemed not necessary, since their main difference from standard agglomerative clustering is in their handling of the highly sparse single cell profiles, a need that is eliminated when dealing with the afore mentioned 'average profiles.'

In the group-level analysis, partition estimates are based on the *constrained K-means algorithm* (26), which estimates clusters at the level of known group labels. This approach is perfectly suited to the downstream analysis of scRNAseq data, following the estimation of mutually exclusive cell types using scRNAseq clustering methods such as *Seurat* (14) or *Scran* (15). It thus provides a unique tool for the analysis of ever-growing single cell data repositories such as the Human Cell Atlas (62), the Human Tumor Atlas Network (HTAN) (62) and the Human BioMolecular Atlas

Program (HuBMAP) (63) among others. By preserving the single observation information within each group, and by thus being able to tailor the feature set to each of the groups, we expect our approach to outperform methods in which group-level information is summarized into single statistical measures. This conclusion is supported by our simulation analysis, where *K2Taxonomer* was shown to significantly outperform Ward's agglomerative method based on group-level test statistics. Even when adopted for observation-level analysis, where inference was performed on the full set of individual observations, *K2Taxonomer* was still shown to significantly outperform standard agglomerative methods, on both simulated and real data, although not to as large an extent.

In our analysis of healthy airway cell types' annotation (28), we employed *K2Taxonomer* to (re)discover subgroups of cell types characterized by shared lineage. Remarkably, our analysis accurately recapitulated the known taxonomic structure relating the different cell types to an extent not matched by the other methods evaluated. This example illustrates a prototypical use of the tool: in those cases where a data set and its associated cell type estimations are publicly available, *K2Taxonomer* facilitates their immediate repurposing for additional insight and discovery. This is a defining advantage over fully unsupervised scRNAseq hierarchical clustering methods, such as *ascend* (24), which cannot make use of prior labeling of the cell profiles for subgroup discovery. Subgroup annotation from *ascend* analysis of these data was complicated by the fact that cell profiles belonging to the same cell type label were split across multiple partitions, and often to an extent that it made it challenging to assign a representative cell type label to individual subgroups. Moreover, even for annotated subgroups, none of the *ascend* models were able to recapitulate taxonomic relationships between cell types to the same level as the *K2Taxonomer* model. This was also true when comparing the *K2Taxonomer* and *ascend* models generated from the breast cancer TILs scRNAseq data set.

It is important to emphasize that in many data sets continuous lineage trajectories are non-existent or obscured by phenotype-driven inter-group transcriptional relationships. While *K2Taxonomer* cannot identify precursor relationships between cell types, the strategy of pairing recursive partitioning with local feature selection allows the discernment of relative relationships between groups rather than 'all-or-nothing' connections as in graph-based trajectory models. The advantages of this strategy are exemplified by the analysis of the healthy airway, in which the data included samples of multiple individuals and airway locations. *PAGA* analysis of these data produced highly-connected graphs with no decipherable trajectories or subgroups beyond that of disconnected subgraphs. On the other hand, *K2Taxonomer* recapitulated clear lineage-driven hierarchies of airway cell subgroups. While the subgrouping of endothelial cells as a subset of immune cells does not reflect the expected hierarchy, *K2Taxonomer* generated this model without parameterization. The 'best' *PAGA* model had the unfair advantage of having been chosen from a set of distinct models generated from combinations of feature selection and dimensionality reduction pa-

rameters. In practice, the a priori choice of the optimal values of these parameters is challenging.

Our extensive analysis of single-cell data from breast cancer TILs showcased the incorporation of *K2Taxonomer* in an advanced in-silico study that yielded significant novel insights. In contrast to bulk gene expression, which captures average expression across all cells, identifying dominant transcriptional programs driving phenotypic similarities between subgroups of cell populations offers additional insights to deconvolute the cellular microenvironment of these samples beyond their individual transcriptional signatures. Molecular convergence of cells of disparate lineages is exemplified by subpopulations of CD8+ and CD4+ T cells, each of which exists in various functional states as naïve, effector, and memory subpopulations (64). Importantly, our *K2Taxonomer*-based analysis showed that concordant subpopulations of CD8+ and CD4+ T cells share transcriptional signatures that may outweigh those arising from their shared lineage. For example, both CD8+ Trm and CD4+ Trm cells have been reported to express surface molecules, CD69 and CD11a (65). Concordantly, CD8+ Trm and CD4+ Trm cells were segregated into a common subgroup by K2T, demonstrating the relative dominance of their shared transcriptional activity. Projection of the expression signature of Trm cell subgroups was associated with better survival in the METABRIC data set. Past studies focusing on CD8+ Trm cell markers have reported similar findings (29,66).

Unlike Trm cells, the presence of immune-suppressing Treg cells in the microenvironment has been associated with poor prognosis in breast cancer (67–69). After identifying TNFRSF4 as heterogeneously expressed across the Treg cell subgroup, we showed that TNFRSF4 expression was associated with worse patient survival in the METABRIC data set when adjusted for CCL5 expression, which was downregulated among all Treg cells. This supports previous findings that TNFRSF4, also known as OX40, is a marker of high Treg cell immunosuppressive activity (70,71). The high level of co-expression of TNFRSF4 and CCL5 in the METABRIC data set suggests that either gene is associated with immune infiltration in breast cancer tumors. Additionally, this provides a resolution as to why projections of the signature of the Treg cell subgroup were associated with better patient survival, while the signature of the Treg cell subset characterized by high TNFRSF4 expression, was not. These results are consistent with previous studies establishing the ratio between Treg and CD8+ T cell abundance as a prognostic marker of breast cancer that reflects immune inhibitory function of Treg cells (72). Moreover, these results strongly suggest a prognostic value for markers that capture the degree of immune inhibitory activity of Treg cell populations.

Finally, *K2Taxonomer* identified a diverse subgroup of breast cancer TILs characterized by consistent upregulation of translational genes. Increased ribosomal biogenesis has been previously implicated in increased tumorigenesis (73–76), but has only recently been implicated in T cell activation (30) and expansion (31). Unlike the majority of other subgroups, the signature of the T cell subgroup overexpressing translational machinery genes was associ-

ated with better patient survival in METABRIC patients regardless of adjustments for inflammation (49) and proliferation (50) signatures. Furthermore, the association of the expression of specific translational genes with better patient survival was significantly correlated to their overexpression in this T cell subgroup. These results suggest that overexpression of these T cell-specific translational genes is not masked by tumor-specific gene expression and is therefore indicative of CD4+ and CD8+ T cell tumor infiltration.

In summary, *K2Taxonomer* demonstrated a remarkable ability to discover biologically relevant taxonomies when applied to the analysis of both bulk gene expression and scRNAseq data and to outperform standard agglomerative methods. In multiple practical applications, we showcased the versatility of *K2Taxonomer* to analyze scRNAseq data toward the characterization of genes and pathways distinguishing specific subgroups, thereby generating hypotheses that were then in-silico validated in independent bulk gene expression data. As noted, while we here focused on the analysis of transcriptomics data, the proposed approach is equally applicable to other bulk and single-cell 'omics' data, such as those generated by high-throughput proteomics and metabolomics assays.

DATA AVAILABILITY

K2Taxonomer is publicly available as an R package on GitHub, https://github.com/montilab/K2Taxonomer. Results are available through the interactive portal at https://montilab.bu.edu/k2BRCAtcell.

The data that support these findings are publicly available and were accessed from several repositories. The bulk microarray gene expression and subject data from the METABRIC breast cancer study is available through the cBioPortal, https://www.cbioportal.org/study/summary?id=brca_metabric. The bulk RNAseq gene expression and subject data from the TCGA BRCA project is available through the Genomic Data Commons, https://gdc.cancer.gov/access-data/. The single-cell RNAseq gene expression and subject data of healthy airway tissue is available through the UCSC Cell Browser page, https://www.genomique.eu/cellbrowser/HCA/?ds=HCA_airway_epithelium. Finally, the single-cell RNAseq gene expression and subject data of breast cancer immune cells is available in the Gene Expression Omnibus, GSE110938.

SUPPLEMENTARY DATA

Supplementary Data are available at NAR Online.

REFERENCES

1. Lonsdale,J., Thomas,J., Salvatore,M., Phillips,R., Lo,E., Shad,S., Hasz,R., Walters,G., Garcia,F., Young,N. *et al.* (2013) The genotype-tissue expression (GTEx) project. *Nat. Genet.*, **45**, 580–585.
2. Gao,G.F., Parker,J.S., Reynolds,S.M., Silva,T.C., Wang,L.-B., Zhou,W., Akbani,R., Bailey,M., Balu,S., Berman,B.P. *et al.* (2019) Before and after: comparison of legacy and harmonized TCGA genomic data commons' data. *Cell Syst.*, **9**, 24–34.
3. Barretina,J., Caponigro,G., Stransky,N., Venkatesan,K., Margolin,A.A., Kim,S., Wilson,C.J., Lehár,J., Kryukov,G.V., Sonkin,D. *et al.* (2012) The Cancer Cell Line Encyclopedia enables predictive modelling of anticancer drug sensitivity. *Nature*, **483**, 603–607.
4. Taliun,D., Harris,D.N., Kessler,M.D., Carlson,J., Szpiech,Z.A., Torres,R., Taliun,S.A.G., Corvelo,A., Gogarten,S.M., Kang,H.M. *et al.* (2021) Sequencing of 53,831 diverse genomes from the NHLBI TOPMed Program. *Nature*, **590**, 290–299.
5. Tang,F., Barbacioru,C., Wang,Y., Nordman,E., Lee,C., Xu,N., Wang,X., Bodeau,J., Tuch,B.B., Siddiqui,A. *et al.* (2009) mRNA-Seq whole-transcriptome analysis of a single cell. *Nat. Methods*, **6**, 377–382.
6. Svensson,V., Vento-Tormo,R. and Teichmann,S.A. (2018) Exponential scaling of single-cell RNA-seq in the past decade. *Nat. Protoc.*, **13**, 599–604.
7. Hoadley,K.A., Yau,C., Hinoue,T., Wolf,D.M., Lazar,A.J., Drill,E., Shen,R., Taylor,A.M., Cherniack,A.D., Thorsson,V. *et al.* (2018) Cell-of-origin patterns dominate the molecular classification of 10,000 tumors from 33 types of cancer. *Cell*, **173**, 291–304.
8. Raj-Kumar,P.-K., Liu,J., Hooke,J.A., Kovatich,A.J., Kvecher,L., Shriver,C.D. and Hu,H. (2019) PCA-PAM50 improves consistency between breast cancer intrinsic and clinical subtyping reclassifying a subset of luminal A tumors as luminal B. *Sci. Rep.*, **9**, 7956.
9. Guinney,J., Dienstmann,R., Wang,X., de Reyniès,A., Schlicker,A., Soneson,C., Marisa,L., Roepman,P., Nyamundanda,G., Angelino,P. *et al.* (2015) The consensus molecular subtypes of colorectal cancer. *Nat. Med.*, **21**, 1350–1356.
10. Ma,X., Gu,J., Wang,K., Zhang,X., Bai,J., Zhang,J., Liu,C., Qiu,Q. and Qu,K. (2019) Identification of a molecular subtyping system associated with the prognosis of Asian hepatocellular carcinoma patients receiving liver resection. *Sci. Rep.*, **9**, 7073.
11. Wang,W.-H., Xie,T.-Y., Xie,G.-L., Ren,Z.-L. and Li,J.-M. (2018) An integrated approach for identifying molecular subtypes in human colon cancer using gene expression data. *Genes*, **9**, 397.
12. Aine,M., Eriksson,P., Liedberg,F., Sjödahl,G. and Höglund,M. (2015) Biological determinants of bladder cancer gene expression subtypes. *Sci. Rep.*, **5**, 10957.
13. Duò,A., Robinson,M.D. and Soneson,C. (2018) A systematic performance evaluation of clustering methods for single-cell RNA-seq data. *F1000Res*, **7**, 1141.
14. Stuart,T., Butler,A., Hoffman,P., Hafemeister,C., Papalexi,E., Mauck,W.M., Hao,Y., Stoeckius,M., Smibert,P. and Satija,R. (2019) Comprehensive integration of single-cell data. *Cell*, **177**, 1888–1902.
15. Lun,A.T.L., McCarthy,D.J. and Marioni,J.C. (2016) A step-by-step workflow for low-level analysis of single-cell RNA-seq data with Bioconductor. *F1000Res*, **5**, 2122.
16. Saitou,N. and Nei,M. (1987) The neighbor-joining method: a new method for reconstructing phylogenetic trees. *Mol. Biol. Evol.*, **4**, 406–425.
17. Hughes,A.L. and Friedman,R. (2009) A phylogenetic approach to gene expression data: evidence for the evolutionary origin of mammalian leukocyte phenotypes. *Evol. Dev.*, **11**, 382–390.
18. Saelens,W., Cannoodt,R., Todorov,H. and Saeys,Y. (2019) A comparison of single-cell trajectory inference methods. *Nat. Biotechnol.*, **37**, 547–554.
19. Tritschler,S., Büttner,M., Fischer,D.S., Lange,M., Bergen,V., Lickert,H. and Theis,F.J. (2019) Concepts and limitations for learning

developmental trajectories from single cell genomics. *Development*, **146**, dev170506.

20. Wolf,F.A., Hamey,F.K., Plass,M., Solana,J., Dahlin,J.S., Göttgens,B., Rajewsky,N., Simon,L. and Theis,F.J. (2019) PAGA: graph abstraction reconciles clustering with trajectory inference through a topology preserving map of single cells. *Genome Biol.*, **20**, 59.

21. Lähnemann,D., Köster,J., Szczurek,E., McCarthy,D.J., Hicks,S.C., Robinson,M.D., Vallejos,C.A., Campbell,K.R., Beerenwinkel,N., Mahfouz,A. *et al.* (2020) Eleven grand challenges in single-cell data science. *Genome Biol.*, **21**, 31.

22. Hossen,B. (2015) Methods for evaluating agglomerative hierarchical clustering for gene expression data: a comparative study. *CBB*, **3**, 88.

23. Zurauskienė,J. and Yau,C. (2016) pcaReduce: hierarchical clustering of single cell transcriptional profiles. *BMC Bioinformatics*, **17**, 140.

24. Senabouth,A., Lukowski,S.W., Hernandez,J.A., Andersen,S.B., Mei,X., Nguyen,Q.H. and Powell,J.E. (2019) ascend: R package for analysis of single-cell RNA-seq data. *GigaScience*, **8**, giz087.

25. Zeisel,A., Munoz-Manchado,A.B., Codeluppi,S., Lonnerberg,P., La Manno,G., Jureus,A., Marques,S., Munguba,H., He,L., Betsholtz,C. *et al.* (2015) Cell types in the mouse cortex and hippocampus revealed by single-cell RNA-seq. *Science*, **347**, 1138–1142.

26. Wagstaff,K., Cardie,C., Rogers,S. and Schrödl,S. (2001) Constrained K-means clustering with background knowledge. In: *ICML*. pp. 577–584.

27. METABRIC Group, Curtis,C., Shah,S.P., Chin,S.-F., Turashvili,G., Rueda,O.M., Dunning,M.J., Speed,D., Lynch,A.G., Samarajiwa,S. *et al.* (2012) The genomic and transcriptomic architecture of 2,000 breast tumours reveals novel subgroups. *Nature*, **486**, 346–352.

28. Deprez,M., Zaragosi,L.-E., Truchi,M., Becavin,C., Ruiz García,S., Arguel,M.-J., Plaisant,M., Magnone,V., Lebrigand,K., Abelanet,S. *et al.* (2020) A single-cell atlas of the human healthy airways. *Am. J. Respir. Crit. Care Med.*, **202**, 1636–1645.

29. Savas,P., Virassamy,B., Ye,C., Salim,A., Mintoff,C.P., Caramia,F., Salgado,R., Byrne,D.J., Teo,Z.L., Dushyanthen,S. *et al.* (2018) Single-cell profiling of breast cancer T cells reveals a tissue-resident memory subset associated with improved prognosis. *Nat. Med.*, **24**, 986–993.

30. Ricciardi,S., Manfrini,N., Alfieri,R., Calamita,P., Crosti,M.C., Gallo,S., Müller,R., Pagani,M., Abrignani,S. and Biffo,S. (2018) The translational machinery of human CD4+ T cells is poised for activation and controls the switch from quiescence to metabolic remodeling. *Cell Metab.*, **28**, 895–906.

31. Araki,K., Morita,M., Bederman,A.G., Konieczny,B.T., Kissick,H.T., Sonenberg,N. and Ahmed,R. (2017) Translation is actively regulated during the differentiation of CD8+ effector T cells. *Nat. Immunol.*, **18**, 1046–1057.

32. Kim,S., Reed,E., Monti,S. and Schlezinger,J. (2021) A Data-Driven Transcriptional Taxonomy of Adipogenic Chemicals to Identify White and Brite Adipogens. bioRxiv doi: https://doi.org/10.1101/519629, 29 April 2021, preprint: not peer reviewed.

33. Rousseeuw,P.J. and Croux,C. (1993) Alternatives to the median absolute deviation. *J. Am. Statist. Assoc.*, **88**, 1273–1283.

34. Breiman,L. (2001) Random Forests. *Mach. Learn.*, **45**, 5–32.

35. Ritchie,M.E., Phipson,B., Wu,D., Hu,Y., Law,C.W., Shi,W. and Smyth,G.K. (2015) limma powers differential expression analyses for RNA-sequencing and microarray studies. *Nucleic Acids Res.*, **43**, e47.

36. Monti,S., Tamayo,P., Mesirov,J. and Golub,T. (2003) Consensus clustering: a resampling-based method for class discovery and visualization of gene expression microarray data. *Machine Learning*, **52**, 91–118.

37. Hamann,U. (1961) Merkmalsbestand und Verwandtschaftsbeziehungen der Farinose. Ein Betrag zum System der Monokotyledonen. *Willdenowia*, **2**, 639–768.

38. Yule,G.U. (1912) On the methods of measuring the association between two variables. *J. R. Statistic. Soc.*, **75**, 576–642.

39. Hänzelmann,S., Castelo,R. and Guinney,J. (2013) GSVA: gene set variation analysis for microarray and RNA-Seq data. *BMC Bioinformatics*, **14**, 7.

40. Benjamini,Y. and Hochberg,Y. (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Stat. Soc.: Ser. B (Methodological)*, **57**, 289–300.

41. The Cancer Genome Atlas Network (2012) Comprehensive molecular portraits of human breast tumours. *Nature*, **490**, 61–70.

42. Pereira,B., Chin,S.-F., Rueda,O.M., Vollan,H.-K.M., Provenzano,E., Bardwell,H.A., Pugh,M., Jones,L., Russell,R., Sammut,S.-J. *et al.* (2016) The somatic mutation profiles of 2,433 breast cancers refine their genomic and transcriptomic landscapes. *Nat. Commun.*, **7**, 11479.

43. Erman,N., Korosec,A. and Suklan,J. (2015) Performance of selected agglomerative hierarchical clustering methods. *IIASS*, **8**, 180–204.

44. Kim,H. and Park,H. (2007) Sparse non-negative matrix factorizations via alternating non-negativity-constrained least squares for microarray data analysis. *Bioinformatics*, **23**, 1495–1502.

45. Risso,D., Perraudeau,F., Gribkova,S., Dudoit,S. and Vert,J.-P. (2018) A general and flexible method for signal extraction from single-cell RNA-seq data. *Nat. Commun.*, **9**, 284.

46. Stuart,T., Butler,A., Hoffman,P., Hafemeister,C., Papalexi,E., Mauck,W.M., Hao,Y., Stoeckius,M., Smibert,P. and Satija,R. (2019) Comprehensive Integration of Single-Cell Data. *Cell*, **177**, 1888–1902.

47. Tang,W., Bertaux,F., Thomas,P., Stefanelli,C., Saint,M., Marguerat,S. and Shahrezaei,V. (2019) bayNorm: Bayesian gene expression recovery, imputation and normalization for single-cell RNA-sequencing data. *Bioinformatics*, **36**, 1174–1181.

48. Liberzon,A., Subramanian,A., Pinchback,R., Thorvaldsdottir,H., Tamayo,P. and Mesirov,J.P. (2011) Molecular signatures database (MSigDB) 3.0. *Bioinformatics*, **27**, 1739–1740.

49. Winslow,S., Leandersson,K., Edsjö,A. and Larsson,C. (2015) Prognostic stromal gene signatures in breast cancer. *Breast Cancer Res.*, **17**, 23.

50. Venet,D., Dumont,J.E. and Detours,V. (2011) Most random gene expression signatures are significantly associated with breast cancer outcome. *PLoS Comput. Biol.*, **7**, e1002240.

51. Igarashi,Y., Nakatsu,N., Yamashita,T., Ono,A., Ohno,Y., Urushidani,T. and Yamada,H. (2015) Open TG-GATEs: a large-scale toxicogenomics database. *Nucleic Acids Res.*, **43**, D921–D927.

52. Baker,F.B. (1974) Stability of two hierarchical grouping techniques case 1: sensitivity to data errors. *J. Am. Statist. Assoc.*, **69**, 440.

53. Mori,M., Mahoney,J.E., Stupnikov,M.R., Paez-Cortez,J.R., Szymaniak,A.D., Varelas,X., Herrick,D.B., Schwob,J., Zhang,H. and Cardoso,W.V. (2015) Notch3-Jagged signaling controls the pool of undifferentiated airway progenitors. *Development*, **142**, 258–267.

54. Goldfarbmuren,K.C., Jackson,N.D., Sajuthi,S.P., Dyjack,N., Li,K.S., Rios,C.L., Plender,E.G., Montgomery,M.T., Everman,J.L., Bratcher,P.E. *et al.* (2020) Dissecting the cellular specificity of smoking effects and reconstructing lineages in the human airway epithelium. *Nat. Commun.*, **11**, 2485.

55. Torang,A., Gupta,P. and Klinke,D.J. (2019) An elastic-net logistic regression approach to generate classifiers and gene signatures for types of immune cells and T helper cell subsets. *BMC Bioinformatics*, **20**, 433.

56. Pittenger,M.F., Discher,D.E., Péault,B.M., Phinney,D.G., Hare,J.M. and Caplan,A.I. (2019) Mesenchymal stem cell perspective: cell biology to clinical progress. *npj Regen Med*, **4**, 22.

57. Goasguen,J.E., Bennett,J.M., Bain,B.J., Vallespi,T., Brunning,R., Mufti,G.J. and for the International Working Group on Morphology of Myelodysplastic Syndrome (IWGM-MDS) (2009) Morphological evaluation of monocytes and their precursors. *Haematologica*, **94**, 994–997.

58. DeNardo,D.G. and Coussens,L.M. (2007) Inflammation and breast cancer. Balancing immune response: crosstalk between adaptive and innate immune cells during breast cancer progression. *Breast Cancer Res.*, **9**, 212.

59. Zhang,S.-C., Hu,Z.-Q., Long,J.-H., Zhu,G.-M., Wang,Y., Jia,Y., Zhou,J., Ouyang,Y. and Zeng,Z. (2019) Clinical implications of tumor-infiltrating immune cells in breast cancer. *J. Cancer*, **10**, 6175–6184.

60. Fang,P., Li,X., Dai,J., Cole,L., Camacho,J.A., Zhang,Y., Ji,Y., Wang,J., Yang,X.-F. and Wang,H. (2018) Immune cell subset differentiation and tissue inflammation. *J. Hematol. Oncol.*, **11**, 97.

61. Christiansen,B. (2018) Ensemble averaging and the curse of dimensionality. *J. Climate*, **31**, 1587–1596.

62. Regev,A., Teichmann,S.A., Lander,E.S., Amit,I., Benoist,C., Birney,E., Bodenmiller,B., Campbell,P., Carninci,P., Clatworthy,M. *et al.* (2017) The Human Cell Atlas. *eLife*, **6**, e27041.

63. HuBMAP Consortium (2019) The human body at cellular resolution: the NIH Human Biomolecular Atlas Program. *Nature*, **574**, 187–192.

64. Golubovskaya,V. and Wu,L. (2016) Different subsets of t cells, memory, effector functions, and car-t immunotherapy. *Cancers*, **8**, 36.

65. Teijaro,J.R., Turner,D., Pham,Q., Wherry,E.J., Lefrançois,L. and Farber,D.L. (2011) Cutting edge: tissue-retentive lung memory cd4 t cells mediate optimal protection to respiratory virus infection. *J.I.*, **187**, 5510–5514.

66. Wang,Z.-Q., Milne,K., Derocher,H., Webb,J.R., Nelson,B.H. and Watson,P.H. (2016) Cd103 and intratumoral immune response in breast cancer. *Clin. Cancer Res.*, **22**, 6290–6297.

67. Dziobek,K., Biedka,M., Nowikiewicz,T., Szymankiewicz,M., Łukaszewska,E. and Dutsch-Wicherek,M. (2018) Analysis of Treg cell population in patients with breast cancer with respect to progesterone receptor status. *Wo*, **22**, 236–239.

68. Merlo,A., Casalini,P., Carcangiu,M.L., Malventano,C., Triulzi,T., Mènard,S., Tagliabue,E. and Balsari,A. (2009) Foxp3 expression and overall survival in breast cancer. *JCO*, **27**, 1746–1752.

69. Bates,G.J., Fox,S.B., Han,C., Leek,R.D., Garcia,J.F., Harris,A.L. and Banham,A.H. (2006) Quantification of regulatory t cells enables the identification of high-risk breast cancer patients and those at risk of late relapse. *JCO*, **24**, 5373–5380.

70. Piconese,S., Timperi,E. and Barnaba,V. (2014) 'Hardcore' OX40 [+] immunosuppressive regulatory T cells in hepatic cirrhosis and cancer. *OncoImmunology*, **3**, e29257.

71. Aspeslagh,S., Postel-Vinay,S., Rusakiewicz,S., Soria,J.-C., Zitvogel,L. and Marabelle,A. (2016) Rationale for anti-OX40 cancer immunotherapy. *Eur. J. Cancer*, **52**, 50–66.

72. Liu,F., Lang,R., Zhao,J., Zhang,X., Pringle,G.A., Fan,Y., Yin,D., Gu,F., Yao,Z. and Fu,L. (2011) CD8+ cytotoxic T cell and FOXP3+ regulatory T cell infiltration in relation to breast cancer survival and molecular subtypes. *Breast Cancer Res. Treat.*, **130**, 645–655.

73. Barna,M., Pusic,A., Zollo,O., Costa,M., Kondrashov,N., Rego,E., Rao,P.H. and Ruggero,D. (2008) Suppression of Myc oncogenic activity by ribosomal protein haploinsufficiency. *Nature*, **456**, 971–975.

74. Mendillo,M.L., Santagata,S., Koeva,M., Bell,G.W., Hu,R., Tamimi,R.M., Fraenkel,E., Ince,T.A., Whitesell,L. and Lindquist,S. (2012) Hsf1 drives a transcriptional program distinct from heat shock to support highly malignant human cancers. *Cell*, **150**, 549–562.

75. Santagata,S., Mendillo,M.L., Tang,Y. -c., Subramanian,A., Perley,C.C., Roche,S.P., Wong,B., Narayan,R., Kwon,H., Koeva,M. *et al.* (2013) Tight coordination of protein translation and hsf1 activation supports the anabolic malignant state. *Science*, **341**, 1238303–1238303.

76. Pelletier,J., Thomas,G. and Volarević,S. (2018) Ribosome biogenesis in cancer: new players and therapeutic avenues. *Nat. Rev. Cancer*, **18**, 51–63.