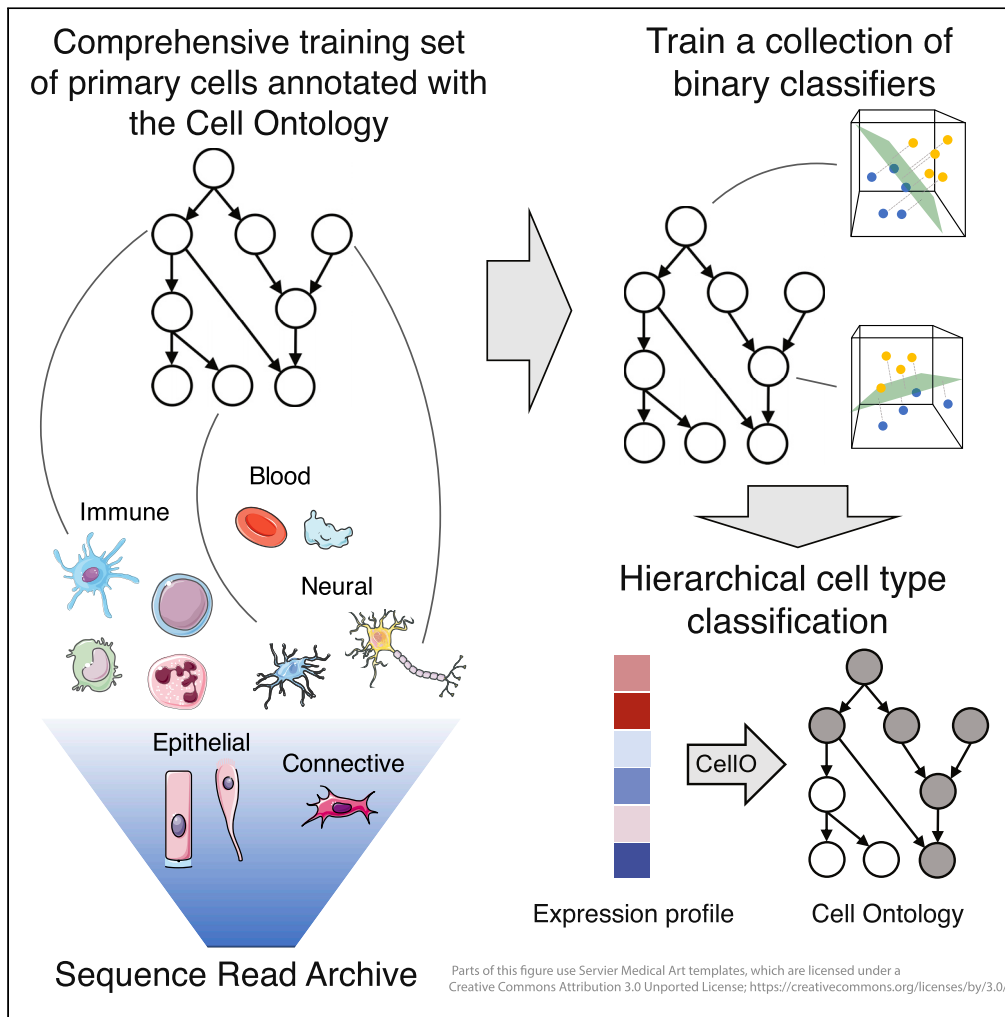


Article

CellO: comprehensive and hierarchical cell type classification of human cells with the Cell Ontology



Matthew N. Bernstein,
Zhongjie Ma,
Michael Gleicher,
Colin N. Dewey

colin.dewey@wisc.edu

HIGHLIGHTS

CellO hierarchically annotates single-cell RNA-seq data using the Cell Ontology

CellO is pre-trained on a comprehensive data set comprising diverse cell types

CellO achieves superior or comparable performance with existing methods

The CellO Viewer is a novel web application for exploring CellO's trained models



Article

Cello: comprehensive and hierarchical cell type classification of human cells with the Cell Ontology

Matthew N. Bernstein,¹ Zhongjie Ma,² Michael Gleicher,^{2,4} and Colin N. Dewey^{2,3,4,5,*}

SUMMARY

Cell type annotation is a fundamental task in the analysis of single-cell RNA-sequencing data. In this work, we present Cello, a machine learning-based tool for annotating human RNA-seq data with the Cell Ontology. Cello enables accurate and standardized cell type classification of cell clusters by considering the rich hierarchical structure of known cell types. Furthermore, Cello comes pre-trained on a comprehensive data set of human, healthy, untreated primary samples in the Sequence Read Archive. Cello's comprehensive training set enables it to run out of the box on diverse cell types and achieves competitive or even superior performance when compared to existing state-of-the-art methods. Lastly, Cello's linear models are easily interpreted, thereby enabling exploration of cell-type-specific expression signatures across the ontology. To this end, we also present the Cello Viewer: a web application for exploring Cello's models across the ontology.

INTRODUCTION

Cell type annotation is a fundamental task in the analysis of single-cell RNA-sequencing (scRNA-seq) data. Recently, a number of computational tools have been developed for automating the cell type annotation task. Nonetheless, many of these tools suffer from certain disadvantages that inhibit their use. First, many existing methods require the user to provide either a set of marker genes associated with each cell type (Zhang et al., 2019a; Pliner et al. 2019) or a suitable training data set with cells already annotated with cell type labels (Ma and Pellegrini 2020; Alquicira-Hernandez et al., 2019; Tan and Cahan 2019). Marker gene-based approaches are challenged by the fact that there is not a canonical set of marker genes for most cell types (Zhang et al., 2019b). Furthermore, finding an appropriate and labeled training set that contains all of the cell types in the target data set can be challenging, especially considering that existing approaches are sensitive to the chosen training set (Abdelaal et al., 2019).

Second, many existing methods use flat classification. Flat classification suffers from the possibility that predictions are logically inconsistent with the hierarchy of cell types. Specifically, for a given query, a flat classifier may output a probability for a cell type that is larger than the classifier's output for its parent cell type in the hierarchy (Obozinski et al., 2008). Such incoherent outputs reduce the scientific usefulness of the classifier. We assert that framing the cell type classification task as that of hierarchical classification against the Cell Ontology (Bard et al. 2005) poses a number of advantages over flat classification. The Cell Ontology provides a comprehensive hierarchy of animal cell types encoded as a directed acyclic graph (DAG). This DAG provides a rich source of prior knowledge to the cell type classification task that remains un-utilized in flat classification. In addition, if the algorithm is uncertain about which specific cell type the cell may be, the use of a hierarchy allows the algorithm to place a cell internally within the graph rather than at a leaf node. Thus, for cells whose specific cell types are absent from the training set, a classifier that uses a hierarchy is capable of providing more informative output than simply claiming that the cell is "uncertain" as is implemented by some flat classifiers such as ACTINN (Ma and Pellegrini 2020).

Finally, those methods that do perform hierarchical classification do not make use of the rich hierarchical relationships between known cell types encoded by the Cell Ontology. For example, CHETAH (de Kanter et al., 2019) classifies cells against a hierarchy; however, CHETAH infers this hierarchy from the data rather than utilizing the existing hierarchy encoded by the Cell Ontology. Garnett (Pliner et al. 2019) utilizes a

¹Morgridge Institute for Research, Madison, WI 53715, USA

²Department of Computer Sciences, University of Wisconsin - Madison, Madison, WI 53706, USA

³Department of Biostatistics and Medical Informatics, University of Wisconsin - Madison, Madison, WI 53792, USA

⁴Senior author

⁵Lead contact

*Correspondence:

colin.dewey@wisc.edu

<https://doi.org/10.1016/j.isci.2020.101913>



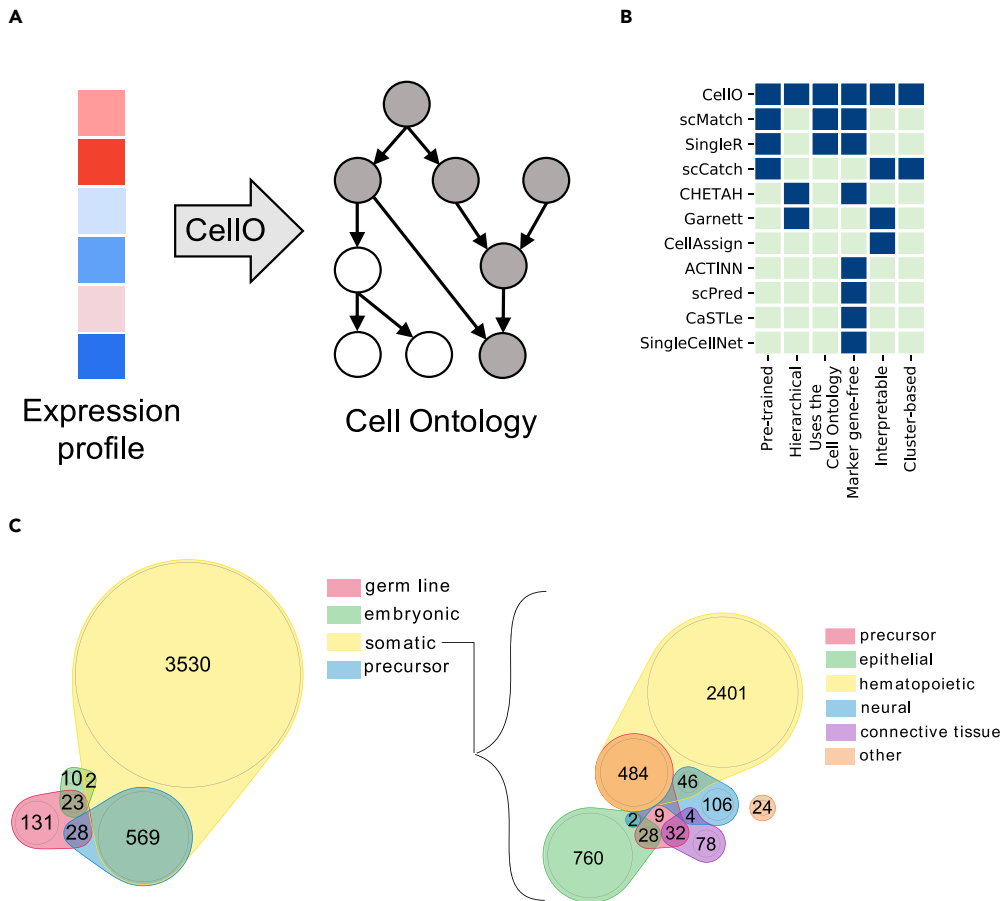


Figure 1. Overview of Cello

(A) A schematic overview of Cello's hierarchical classification approach. Cello performs hierarchical classification with the Cell Ontology. Given a gene expression profile, Cello annotates the cell with a set of cell types (gray nodes) that are consistent with the hierarchical structure of the Cell Ontology.

(B) We compare Cello to eight recent cell type annotation methods regarding the criteria we surmise are desirable in a cell type classification approach: whether the method (1) arrives pre-trained and can run out of the box, (2) incorporates a hierarchy of cell types, (3) specifically uses the Cell Ontology as its hierarchy, (4) requires cell-type-specific marker genes, (5) uses a model that can be interrogated to better understand how it arrived at its decision, and (6) whether the method operates on clusters or single cells. We compare Cello to scMatch (Hou et al., 2019), SingleR (Aran et al., 2019), scCatch (Shao et al., 2020), CHETAH (de Kanter et al., 2019), Garnett (Pliner et al., 2019), CellAssign (Zhang et al., 2019a), ACTINN (Ma and Pellegrini 2020), scPred (Alquicira-Hernandez et al., 2019), CaSTLe (Lieberman et al., 2018), and SingleCellNet (Tan and Cahan 2019). Cello meets more desirable criteria than existing methods.

(C) Euler diagrams of the cell types within the bulk RNA-seq expression profiles used to train Cello. This training set comprises most primary cell bulk RNA-seq samples within the SRA and consists of diverse cell types spanning various tissues, developmental stages, and stages of differentiation. These diagrams were created with nVenn (Pérez-Silva et al., 2018).

hierarchy of cell types; however, these hierarchies must be pre-specified by the user. Furthermore, Garnett requires that each cell within the hierarchy be associated with a set of marker genes. To the best of our knowledge, the only method that utilizes the graph structure of an ontology for the task at hand is URSA (Lee et al., 2013), which classifies gene expression profiles against the BRENDA Tissue Ontology (Gremse et al., 2011).

In this work, we present Cell Ontology-based Classification (Cello), a tool for annotating cells against the graph-structured Cell Ontology (Figure 1A). Cello is a discriminative, supervised machine learning approach for classifying clusters of cells in scRNA-seq data. Cello comes pre-trained on a comprehensive data set comprising the majority human primary samples in the Sequence Read Archive (SRA; Leinonen

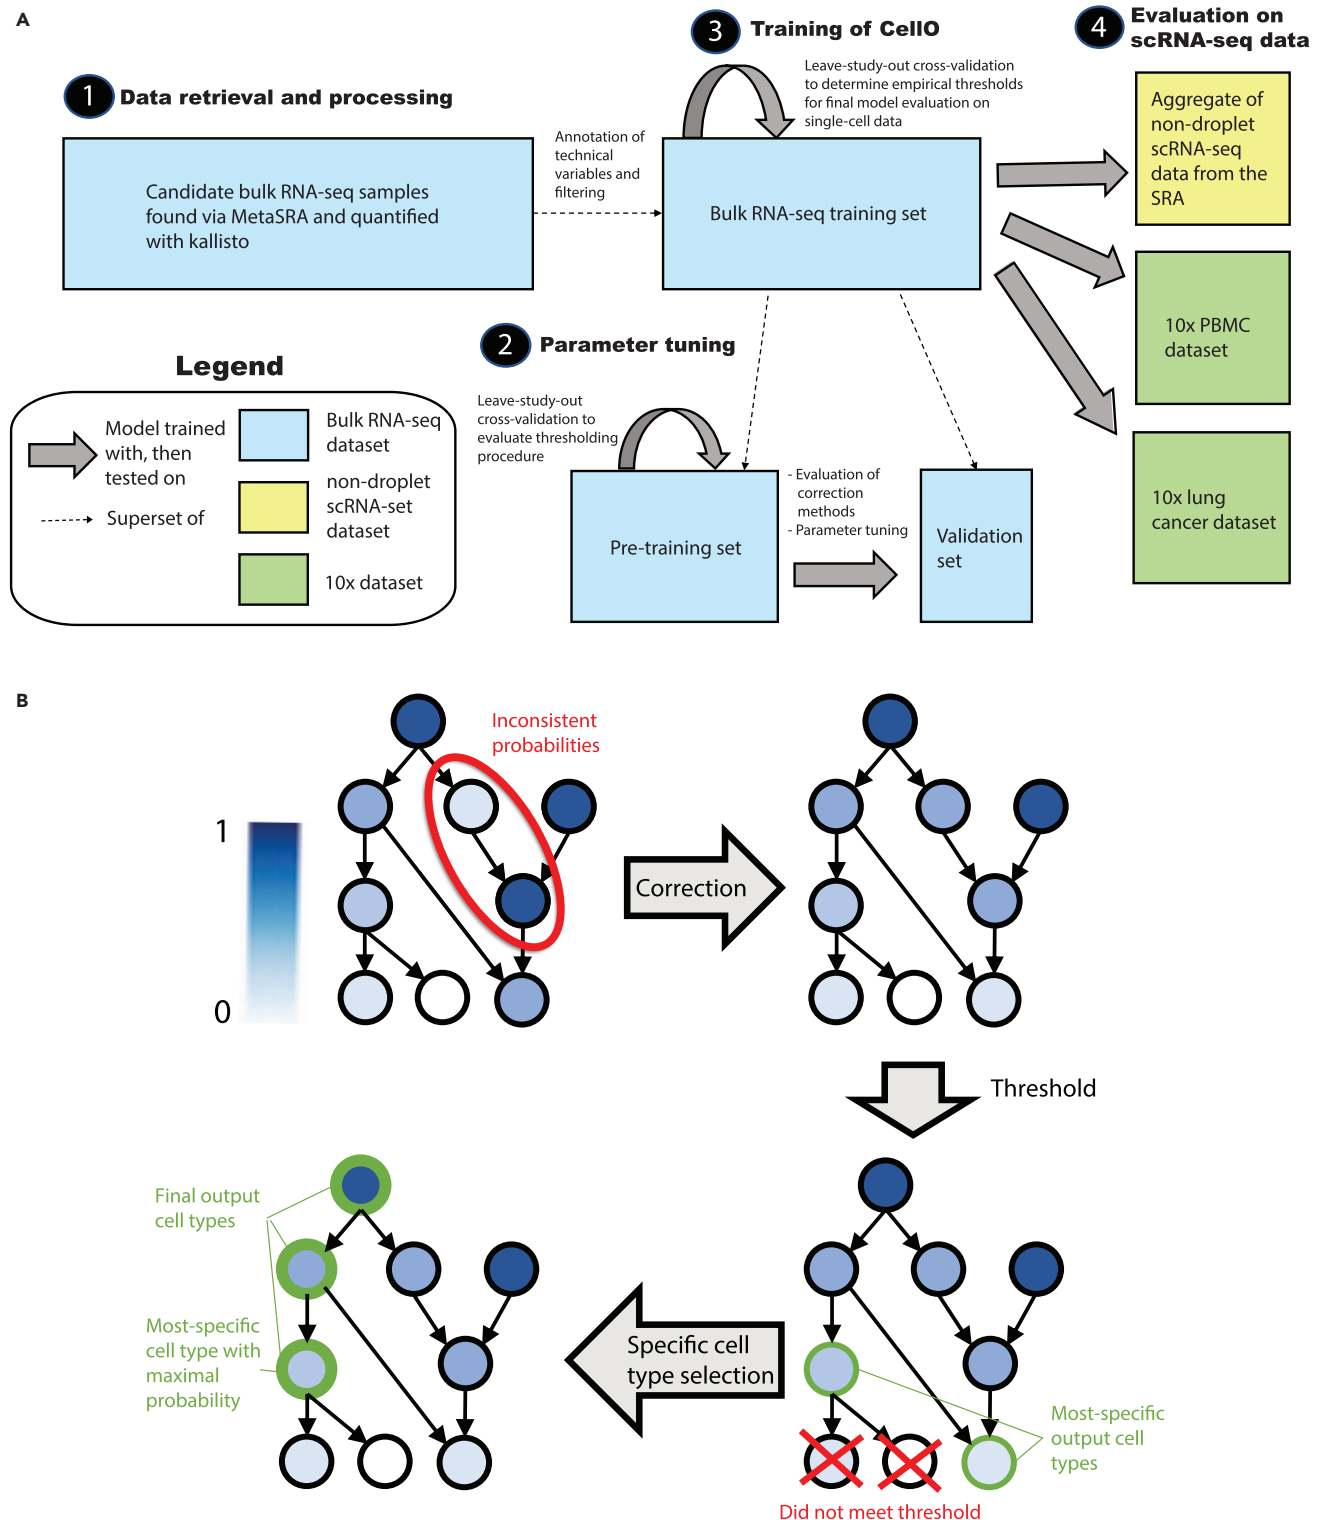


Figure 2. Overview of analyses and Cello's algorithm

(A) A schematic illustration of the data sets and analyses performed in this study. Initial candidate bulk RNA-seq samples were selected from the SRA via the MetaSRA, filtered for errors, and quantified using the kallisto algorithm (Bray et al., 2016), which resulted in a comprehensive bulk RNA-seq training set consisting of healthy, human primary cells. This training set was split into a pre-training and validation set for tuning the parameters of the binary classifiers, as well as for evaluating the graph correction methods (Transparent Methods). The full bulk RNA-seq data set was then used to train the final models that

Figure 2. Continued

were then evaluated on three sets of scRNA-seq data. The first set consisted of an aggregation of diverse non-droplet-based data sets from the SRA. The second data set consisted of FAC-sorted PBMCs from [Zheng et al. \(2017\)](#). The third set consisted of primary lung tumor cells from [Laughney et al. \(2020\)](#).

(B) A schematic illustration of CellO's classification procedure. First, for a given sample, the raw classifier probabilities are corrected with the cell ontology using IR (if CLR is used, this step is not necessary). We illustrate one edge of the graph whose incident nodes have probabilities that are logically inconsistent with the hierarchy and thus require correction because the child node has a higher probability than the parent. Once corrected, cell types whose raw probabilities meet their respective decision threshold are selected. Among these, the most specific cell types (i.e., lowest in the ontology) are examined and the cell type with the highest output probability is selected. CellO outputs this final selected cell type along with all ancestor terms.

[et al., 2011](#)) and therefore arrives ready to run on diverse scRNA-seq data sets. CellO offers a complementary approach to marker gene-based methods for scenarios in which the test set contains cell types with poorly characterized marker genes.

Lastly, CellO makes extensive use of linear models, which are particularly amenable to interpretation. To enable their interpretation, we present a web-based tool, the CellO Viewer, for exploring the cell type expression signals uncovered by the model (<https://uwgraphics.github.io/CellOViewer/>). We benchmarked CellO on a collection of diverse single-cell data sets and found CellO capable of accurately annotating data sets that existing state-of-the-art and ready-to-run (i.e. come pre-trained) annotation methods were unable to accurately annotate, thus highlighting CellO's ability to annotate diverse data sets out of the box. Through its use of the Cell Ontology, its comprehensive training set, and the interpretability of its models, CellO is a practical tool for scRNA-seq cell type annotation ([Figure 1B](#)).

RESULTS

A comprehensive curated RNA-seq data set of human primary cells

In order to capture robust cell type signals, we sought a data set of bulk RNA-seq samples comprising only healthy primary cells that originate from cells that have been isolated based on phenotypic characteristics downstream of gene expression itself (such as cell surface proteins). We thus avoid the circularity in using ground truth cell type labels determined by gene expression (via the expression of cell-type-specific marker genes) as are often provided in scRNA-seq data sets. We did not wish to include cells that underwent multiple passages, were diseased, or underwent other treatments, such as *in vitro* differentiation because these conditions alter gene expression. We therefore curated a data set from the SRA consisting of healthy, untreated, primary cells. To do so efficiently, we leveraged the annotations provided by the MetaSRA project ([Bernstein et al. 2017](#)), which includes sample-specific information including cell type, disease state, and sample type. We then manually curated the samples selected via the MetaSRA by both annotating technical variables and refining cell type annotations ([Transparent Methods](#)).

This curation effort resulted in a data set comprising 4,293 bulk RNA-seq samples from 264 studies. These samples were labeled with 310 cell type terms, of which 113 were the most specific cell types in our data set (i.e., no sample in our data was labeled with a descendant cell type term). These cell types were diverse, spanning multiple stages of development and differentiation ([Figure 1C](#)). We uniformly quantified and normalized (via log transcripts per million) gene expression from the raw RNA-seq data for these samples ([Figure 2A](#), [Transparent Methods](#)). To the best of our knowledge, this data set is the largest and most diverse set of bulk RNA-seq samples derived from only primary cells. Prior to this work, the most comprehensive bulk primary cell transcriptomic data set was compiled by ([Aran et al. 2017](#)), which contains data for 64 cell types from 6 studies. Although our data set consists of only RNA-seq data, this prior data set included samples assayed with several other technologies, such as microarrays. Another comprehensive set of primary cell expression data was collected by [Mabbott et al. \(2013\)](#), which contain primary cell data from 745 samples from 105 studies; however, these data are exclusively from microarrays.

Applications of hierarchical classification methods

We frame the cell type classification task as hierarchical classification against the Cell Ontology's DAG. The hierarchical classification task is inherently a multi-label classification task where each input sample (i.e., cell) is mapped to a "set" of output labels (i.e., cell types). Hierarchical classification extends multi-label classification by further requiring that the output labels are "consistent" with the labels' DAG. That is, for each label in a given output set of labels, the label's parent labels are also in the output set ([Figure 1A](#)). Moreover, when training a hierarchical classifier, samples that are annotated with specific labels (i.e., terms

that are lower in the DAG) can be aggregated in order to train the classifier to recognize more general, ancestral labels in the DAG.

We implemented two strategies for performing hierarchical classification against the Cell Ontology's DAG that both come packaged with CellO. First, we implemented cascaded logistic regression (CLR; [Obozinski et al., 2008](#)), which entails classifying a sample in a top-down fashion from the root of the ontology downward via a collection of binary classifiers. Specifically, each binary classifier is associated with a cell type and is trained to classify a sample conditioned on the sample belonging to all of the cell type's parents in the ontology.

Next, we implemented a collection of one-vs-rest binary classifiers for each cell type in the DAG. We will refer to this as the "independent classifiers" approach. This approach suffers from the possibility that the classifiers' outputs will be inconsistent with the hierarchical structure of the ontology. An inconsistency occurs when the output probability for a given cell type exceeds that of one of its parent cell types in the ontology ([Figure 2B](#)). We tested the use of independent logistic regression classifiers and found inconsistencies to be a frequent source of errors. Specifically, we performed leave-study-out cross-validation on the full set of bulk RNA-seq data and examined the consistency of all edges that were adjacent to at least one cell type whose classifier produced a non-negligible probability (>0.01) of the sample originating from that cell type. Of these edges, 12.1% were inconsistent ([Figure S1](#)). Nearly all samples ($>99\%$) contained at least one inconsistent edge and 34% of samples contained at least one severely inconsistent edge in which the child classifier's probability exceeded the parent classifier's probability by at least 0.25. We will use the term "correction" to refer to the task of reconciling the outputs of independent classifiers with a hierarchy ([Figure 2B](#)).

To date, the one correction method that has been applied to the task at hand is Bayesian network correction (BNC), which is implemented in the URSA tool ([Lee et al., 2013](#)). Therefore, as a baseline, we implemented a BNC algorithm following the description in [Lee et al. \(2013\) \(Transparent Methods\)](#). We also tested two correction methods that have yet to be applied to the cell type classification task: isotonic regression correction (IR; [Obozinski et al., 2008](#)) and a heuristic procedure called the True Path Rule (TPR; [Notaro et al., 2017](#)). IR uses a projection-based approach for correction that entails finding a set of consistent output cell type probabilities that minimize the sum of squared differences to the raw, and possibly inconsistent, classifier output probabilities. In contrast, TPR uses a heuristic procedure that involves a bottom-up pass through the ontology such that the outputs of children classifiers are averaged with the output of the parent classifier to allow information flow across the ontology graph.

To test these correction methods, we first partitioned the bulk RNA-seq data set into a pre-training and validation set ([Figure 2A](#); [Transparent Methods](#)). Using this validation set, we performed a grid search to find the optimal parameters for training each binary logistic regression classifier, and given the optimal set of parameters, compared how well the aforementioned correction methods either enhanced or degraded accuracy over the samples in the validation set. Overall, we find that IR and TPR output probabilities similar to those output by the independent classifiers in regards to both average precision scores across the cell types in the validation set ([Figure 3A](#)) and precision-recall curves when considering each sample-cell type pair as an independent prediction ([Figure 3B](#)). This indicates that IR and TPR do not degrade performance in comparison to independent classifiers. In contrast, we found that the BNC approach significantly degraded performance ([Figures 3A and 3B](#)). We note that these results are in line with work by [Obozinski et al \(2008\)](#), which demonstrates that IR outperforms BNC on the hierarchical protein function prediction task. Although both IR and TPR yielded similar results, we use IR as our correction method of choice due to its simplicity.

We also used this partition of the bulk RNA-seq samples to tune the parameters of the CLR algorithm ([Transparent Methods](#)). We found that after tuning, both IR and CLR achieved similar median F1-scores and median average precisions across cell types on the validation set ([Figure S13](#)), and therefore, both are included in the CellO software and evaluated throughout the remainder of this study.

Comparison to existing methods

We trained both CLR and IR on the full set of bulk RNA-seq samples in order to test their performance on single-cell data ([Figure 2A](#)). We note that training a single-cell classifier with bulk RNA-seq data may lead to




Figure 3. Reconciling the outputs of independent classifiers with a hierarchy

(A) Average precision scores across all cell types for the independent classifiers (Ind.), as well as for IR, TPR, and BNC on the validation set.

(B) Each paired sample and cell type prediction was considered independently. The set of all such predictions was ordered according to their prediction probability and the corresponding precision-recall curve was constructed for the independent classifiers, IR, TPR, and BNC.

models being poorly calibrated to the sparse single-cell expression profiles. To address this challenge, we first cluster single-cell data using the Leiden community detection algorithm (Traag et al. 2019) using the default resolution parameter of 1.0, as implemented in the Scanpy Python package (Wolf et al. 2018), and then compute each cluster's mean expression profile. The mean expression profiles are less sparse than those of the individual cells and thus better resemble the bulk RNA-seq data on which the algorithms were trained. Cello first classifies each cluster based on its mean expression profile and then assigns each cell to its cluster's assigned cell types.

We compiled a data set consisting of 7,366 healthy primary cells originating from non-droplet-based RNA-seq assays, such as SMART-Seq2 (Picelli et al., 2013) and MARS-seq (Jaitin et al., 2014), from the SRA in a manner similar to that used for compiling the bulk RNA-seq training data. This data set originated from 14 studies and were labeled with 125 cell type terms, of which 32 were most specific to the data. Of these cells, 4,936 were of cell types that were included in the bulk RNA-seq training set. This subset of cells originate from 12 studies and were labeled with 71 cell type terms of which 16 were most specific to the data set. We note that for many of the cells used in this analysis, the ground-truth cell types provided by the authors of the data were determined via *in silico* and/or manual approaches (e.g. via heuristic marker gene-based approaches), and thus, this analysis can be understood as an analysis of the “consistency” between the cell types as annotated by the authors and those annotated by the automated methods explored in this work.

We use the subset of 4,936 cells to evaluate IR, CLR, as well as a baseline one-nearest-neighbor (1NN) algorithm that simply returns the cell type labels of the most similar sample in the training set to the query expression profile using Pearson correlation as the similarity metric. We evaluated two aspects of these algorithms' classifications. First, we compute the average precision (a measure of area under the precision-recall curve) on each cell type's output probabilities. Second, for each cell, we evaluate a set of binary yes-no decisions for each cell type that result from thresholding the raw output probabilities and enforcing each cell to be annotated with only one most specific cell type (Figure 2B; Transparent Methods). We evaluate these binary decisions using precision, recall, and F1-score (harmonic mean of precision and recall). We modified the evaluation metrics to take into account samples that were labeled with a general cell type but not a specific cell type (e.g. T cell versus CD8+ T cell; Transparent Methods).

We found that IR and CLR outperformed 1NN according to F1-score, precision, and average precision (Figures 4A, 4B, and 4C). Specifically, IR, CLR, and 1NN produced median F1-scores of 0.81, 0.85, and 0.63, respectively, across all cell types. The cell types on which classification performance was poor were generally more specific and clustered within the hierarchy (Figure 4D). We note that Cello's average precision scores across cell types tend to be higher than its F1-scores (Figure 4D). This discrepancy indicates that Cello is doing well at discriminating among these cell types; however, the decision thresholds used by Cello to output hard classifications for some cell types may be non-optimal. We hypothesize that this is

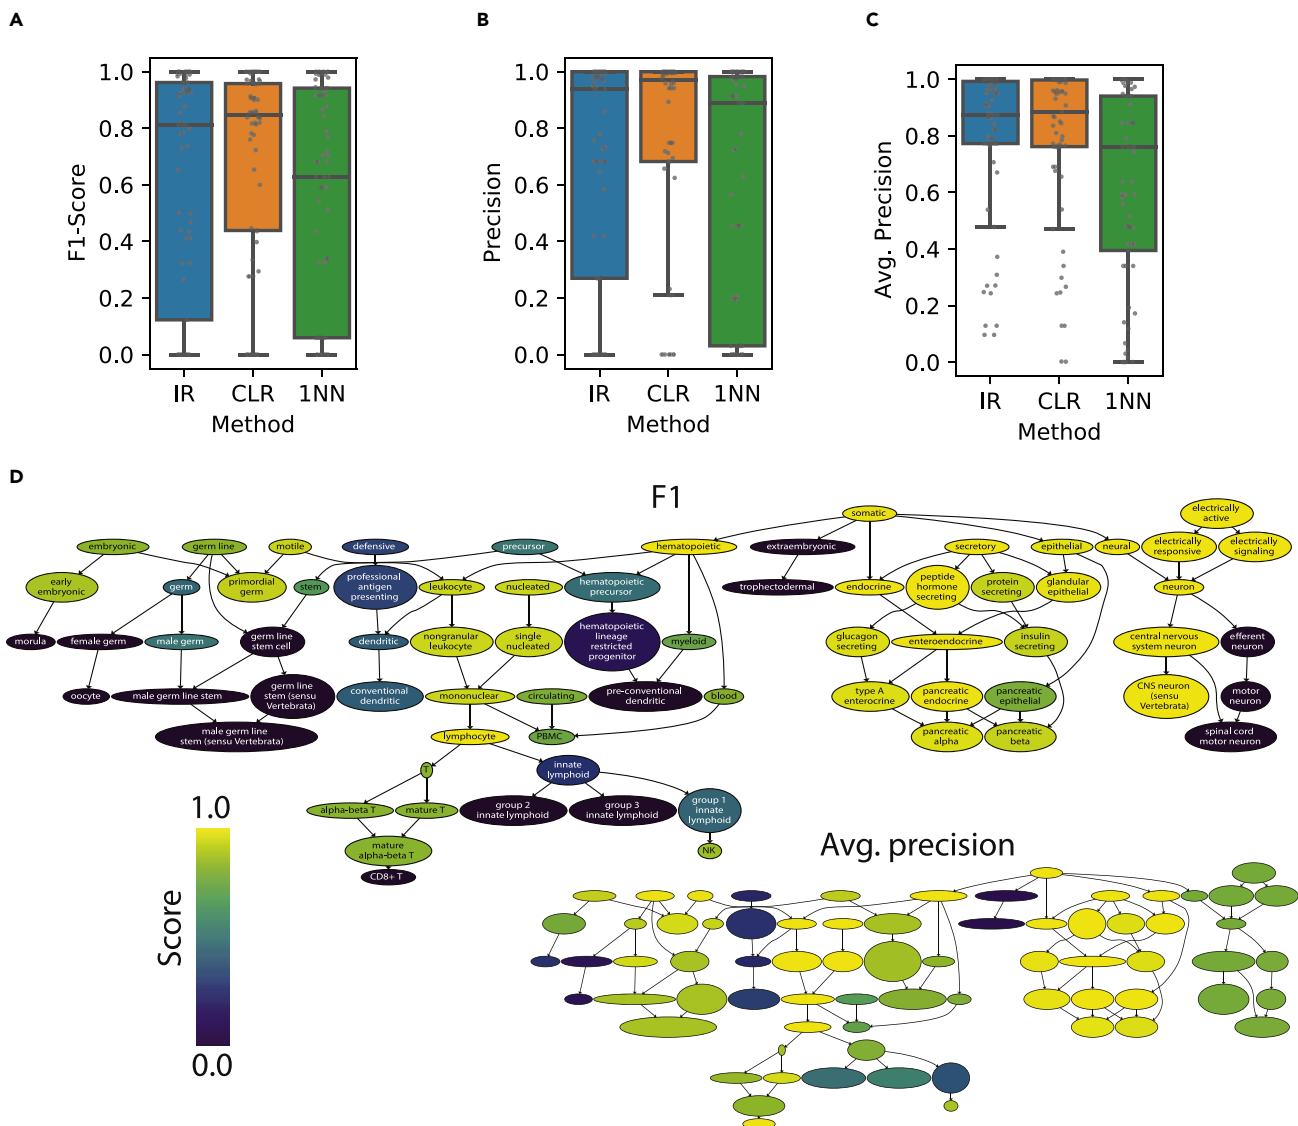


Figure 4. Results on non-droplet-based single-cell data

Cello’s performance on the 4,936 non-droplet-based cells considering only cells whose cell types are present in the bulk RNA-seq training set. We compare the distributions of (A) F1-score, (B) precision, and (C) average precision across all such cell types. (D) The subgraph spanning the non-droplet-based cells where each cell type is colored according to Cello’s (IR) F1-score (top) as well as by average precision (bottom).

due to some classifiers being poorly calibrated, especially for cell types lower in the ontology. This may be due to there being fewer “studies” in the training set generating these cell types, and thus, for a given cell type, Cello’s binary classifiers may be more prone to fitting the batch effects present in these fewer studies (Figure S16).

We also compared Cello to two existing methods, scMatch (Hou et al., 2019) and SingleR (Aran et al., 2019). scMatch and SingleR are most comparable to Cello because they come packaged with comprehensive reference data sets of human cells. Like Cello, these methods are designed to run out of the box on diverse single-cell data sets. scMatch comes packaged with a reference data set comprising data from the FANTOM5 project (Lizio et al., 2017). SingleR comes packaged with two comprehensive human reference data sets: a data set comprising data from the Blueprint (Fernández et al., 2016) and ENCODE (Sloan et al., 2016) projects and a reference set from the Human Primary Cell Atlas (Mabbott et al., 2013). We also built a reference set for SingleR from Cello’s training set in order to isolate methodological differences between




Figure 5. Comparison of Cello to existing approaches on non-droplet-based single-cell data

Evaluating Cello, SingleR, and scMatch on the non-droplet-based cells.

(A) The fraction of cell types in the single-cell test data set that are also present in each method's training set. IR and CLR are not shown separately because they share the same training set. We evaluate SingleR's built in reference sets from the Human Primary Cell Atlas (HPCA) and Blueprint + ENCODE (BE).

(B) The distribution of both F1-scores (left) and precisions (right) for only those cell types that are in each method's training set. We compare Cello to scMatch, SingleR with the Human Primary Cell Atlas (HPCA), and SingleR with the Blueprint + ENCODE reference (BE). Note each distribution evaluates different sets of cell types depending on the particular subset of cell types present in each method's training set.

SingleR and Cello. Unfortunately, scMatch does not support the creation of a custom reference set, and therefore, we were unable to perform this experiment with scMatch. To enable a comparison between scMatch, SingleR, and Cello, we project the outputs of scMatch and SingleR onto the Cell Ontology in order to evaluate scMatch and SingleR within the hierarchical classification framework. Specifically, for a given cell annotated by one of these methods with some cell type C, we also annotate the cell with all ancestors of C according to the Cell Ontology.

First, we note that Cello's training sets included over 50% more of the cell types in this test set compared to those of scMatch and SingleR. In fact, most cell types in this test set, such as pancreatic islet cell types, are absent from scMatch and SingleR's reference sets, thus indicating that a user would be required to supply their own reference set for annotating these cell types (Figure 5A). We thus evaluated each method on only cell types that exist in each respective method's prepackaged reference set and found that Cello outperformed existing approaches (Figure 5B). SingleR performed poorly with Cello's training set, which may be due to the high number of samples in Cello's training set, the high number genes (i.e. including non-coding genes), and the fact that Cello contains bulk RNA-seq samples at a higher level in the ontology (Figure S4A). We note that, in this analysis, for a given method, we may remove from the analysis a specific cell type that is absent from a given method's training set (e.g. stomach epithelial cell), but we would keep an ancestral cell type term (e.g. epithelial cell) if the method's training set contains a sample labeled with this ancestral term (e.g. a sample labeled as intestinal epithelial cell).

Next, we evaluated Cello on fluorescence-activated cell sorted (FAC-sorted) peripheral blood mononuclear cells (PBMC's) from Zheng et al. (2017) that were sequenced with Chromium 10x. We selected this data set because it is one of the few droplet-based data sets for which the cell type labels are determined phenotypically (via sorting) rather than computationally (via expression analysis). To reflect the size of typical single-cell data sets, we subsampled 2,000 cells from each of the ten sorted cell types and aggregated these cells together creating a data set consisting of 20,000 cells. We first compared IR and CLR to the 1NN baseline and again found that IR and CLR outperformed 1NN with respect to median F1-score, although the difference between IR/CLR and 1NN was smaller than in the comparison than on the aforementioned non-droplet-based scRNA-seq data set (Figure 6D). Specifically, IR, CLR, and 1NN produced median F1-scores of 0.97, 0.96, and 0.95, respectively, across all cell types.

Next, we compared Cello to scMatch and SingleR on this PBMC data set. In this analysis, we also ran SingleR using an immune-specific reference set of purified immune cells from Monaco et al. (2019). The cell types in this data set are better represented in scMatch and SingleR's respective reference sets, and

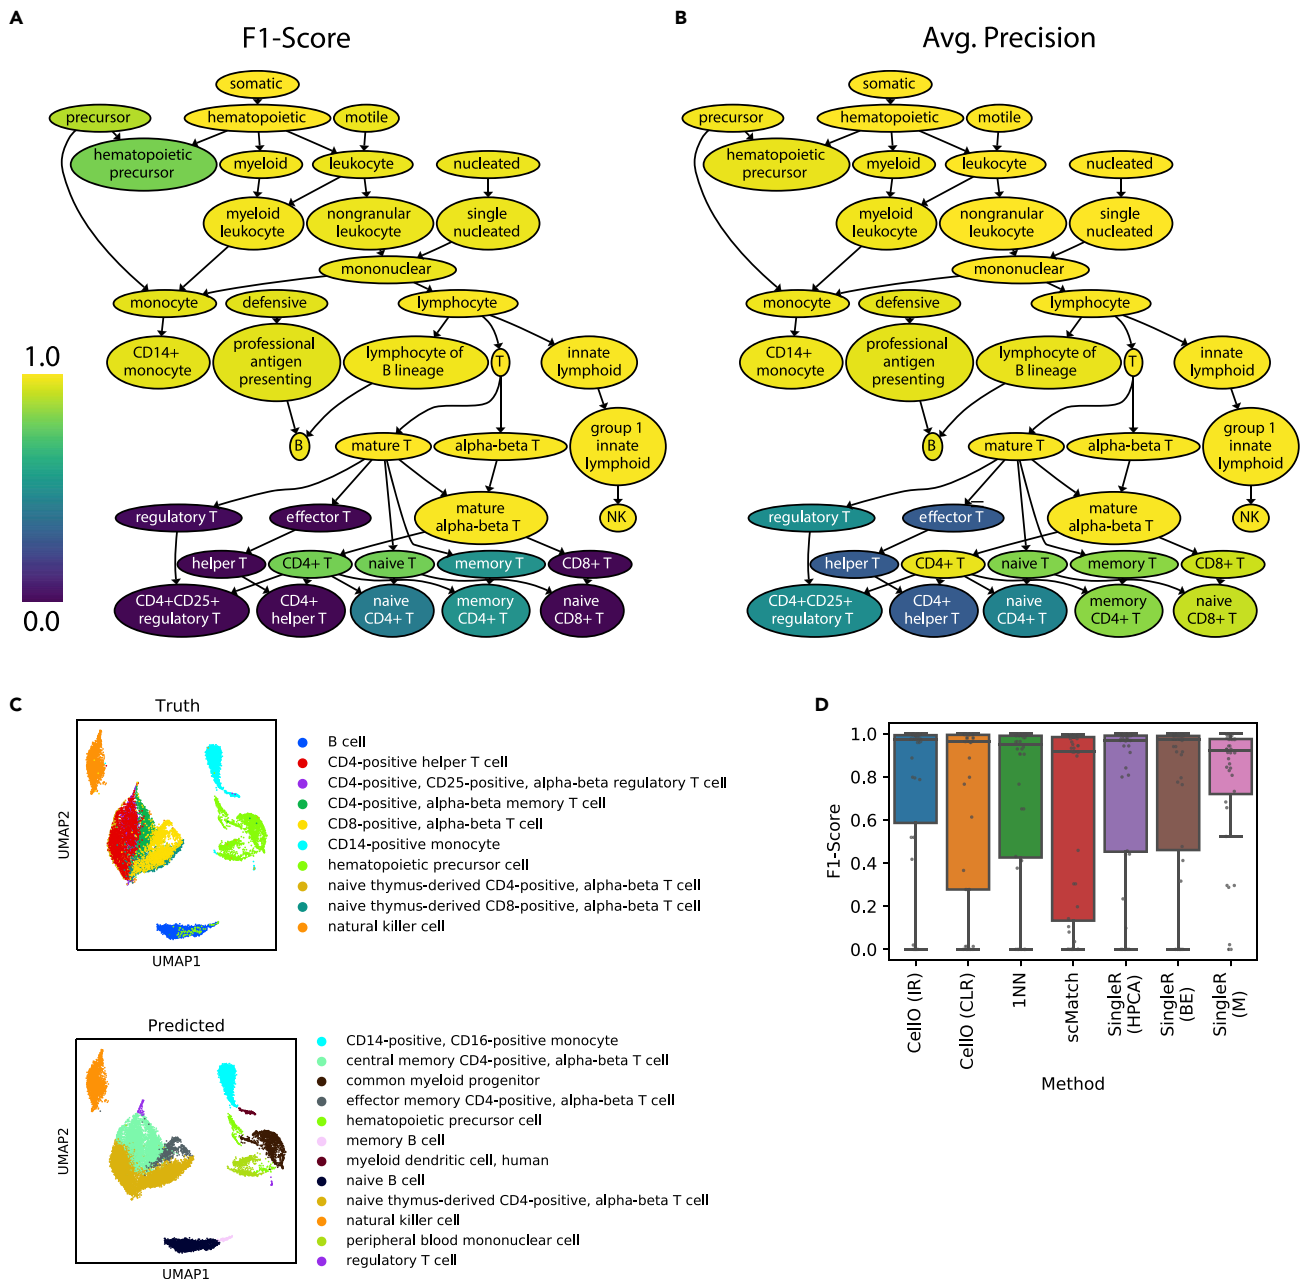


Figure 6. Results on 10x PBMC data

(A and B) (A) The subgraph of the Cell Ontology spanning the 10x PBMC data set from Zheng et al. (2017). Each cell type is colored according to Cello's (IR) F1-score, as well as (B) average precision.

(C) UMAP plots of the single-cell data set where cells are colored by their true cell type (top), as well as the most specific predicted cell type (i.e. lowest in the ontology) as output by Cello (bottom).

(D) Boxplots displaying the distribution of F1-scores across all cell types for IR, CLR, INN, scMatch, SingleR with the Human Primary Cell Atlas (HPCA), SingleR with the Blueprint + ENCODE reference (BE), and SingleR with the Monaco et al. reference (M).

thus, a comparison between Cello and these methods on this data better isolates performance differences between these methods. Like scMatch and SingleR, Cello struggled to accurately classify the T cell subtypes (Figures 6A–6C, S2C, S3, and S4B). Among the methods compared here, SingleR with the Monaco et al. reference set most accurately classified the T cell subtypes (Figure S3), though we note that this reference set is specialized for immune cell types, whereas the other reference sets, including Cello's, are more broad. We also note that Cello produced high average precision scores on most cell types including many

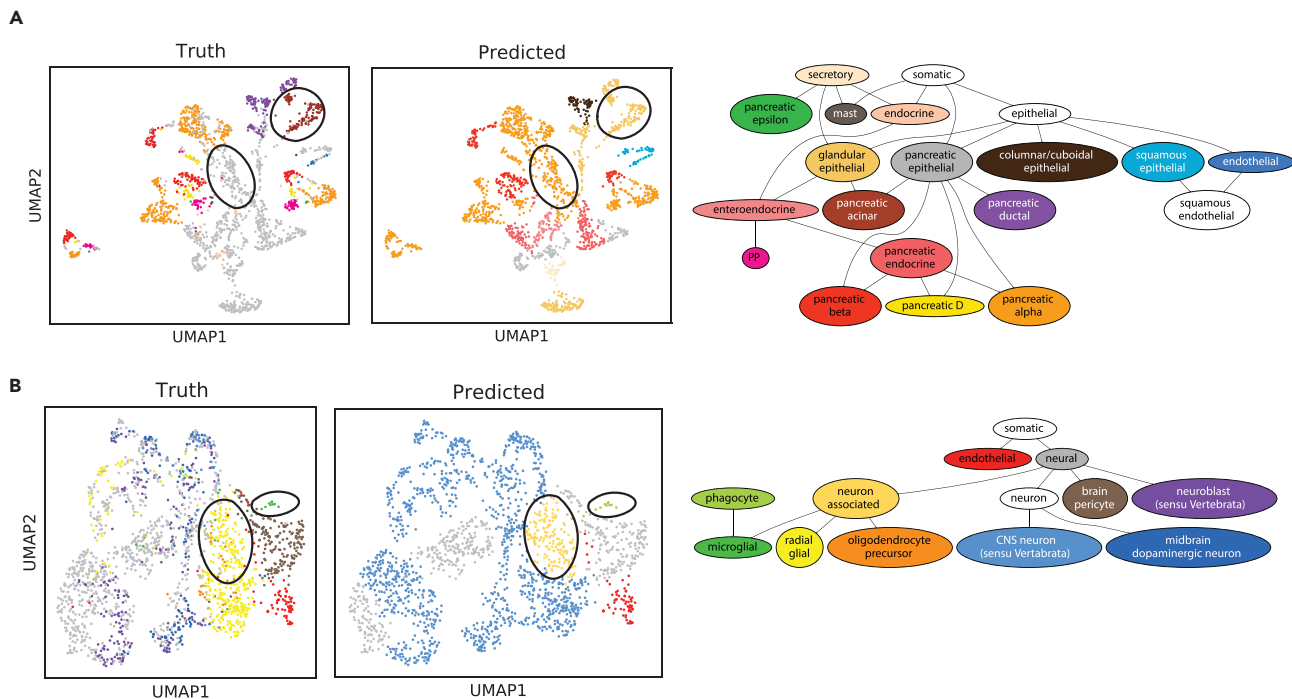


Figure 7. Examination of Cello's performance on difficult data sets

(A) UMAP plots of all healthy cells in [Segerstolpe et al. \(2016\)](#) including cells for which their specific cell types are not present in Cello's bulk RNA-seq training set. Cells are colored according to their true cell type (left) and (IR) predicted cell type (right). Highlighted are Cello's predictions made on pancreatic acinar cells (top right ovals), as well as a subset of uncharacterized pancreatic epithelial cells predicted as A cells (center ovals).

(B) UMAP plots of human, embryonic neural cells from [La Manno et al. \(2016\)](#). Cells are colored according to their true cell type (left) and predicted cell type (right). Highlighted are Cello's predictions made on both the microglial and glial cells and note that Cello annotates these cells using terms that are higher in the ontology's graph than their true terms.

of these T cell subtypes. Again, this indicates that Cello's classifiers have learned to discriminate between these cell types; however, the threshold for calling these cell types may be non-optimal.

Inspection of performance on challenging and diseased samples

We examined Cello's classifications on three challenging data sets. Two of these data sets comprised subsets of the 7,366 cells from non-droplet-based assays and are challenging because they contained cells for which their combination of cell types was absent from Cello's training set. First, we examined Cello's accuracy on 1,978 healthy pancreatic islet cells from [Segerstolpe et al. \(2016\)](#), which includes cell types that are absent from the bulk RNA-seq training set, specifically ductal cells, acinar cells, epsilon cells, and delta cells. We found that Cello was able to correctly annotate the acinar cells as glandular epithelial cells, which is an ancestral cell type to acinar cells in the Cell Ontology ([Figures 7A and S2A](#)). This highlights the advantage of classifying against a hierarchy in that it enables Cello to annotate cells with a term higher in the ontology DAG when it is unsure about a cell's more specific cell type. We also note that a number of cells were uncharacterized in the original study due to not meeting a stringent quality control filter. Cello annotated many of these cells as pancreatic A cells (a.k.a. pancreatic alpha cells), which is plausible owing to both their close position to annotated A cells according to UMAP, which is known to preserve some level of global structure in high dimensional data ([Becht et al., 2018](#)), as well as the fact that A cells were found to be the most abundant endocrine cell type in [Segerstolpe et al. \(2016\)](#) of those that passed their stringent quality control filtering.

We also further examined Cello's classification on 1,977 fetal neural cells from [La Manno et al. \(2016\)](#). Although the bulk RNA-seq training data contain samples of both embryonic cells and cells of various neural cell types, they do not contain any sample labeled as both neural cell and embryonic cell. Despite this discrepancy, Cello was able to annotate these cells with reasonable cell type labels ([Figures 7B and S2B](#)). We note that the microglial cells were annotated as phagocytes, which are an ancestral term to microglial

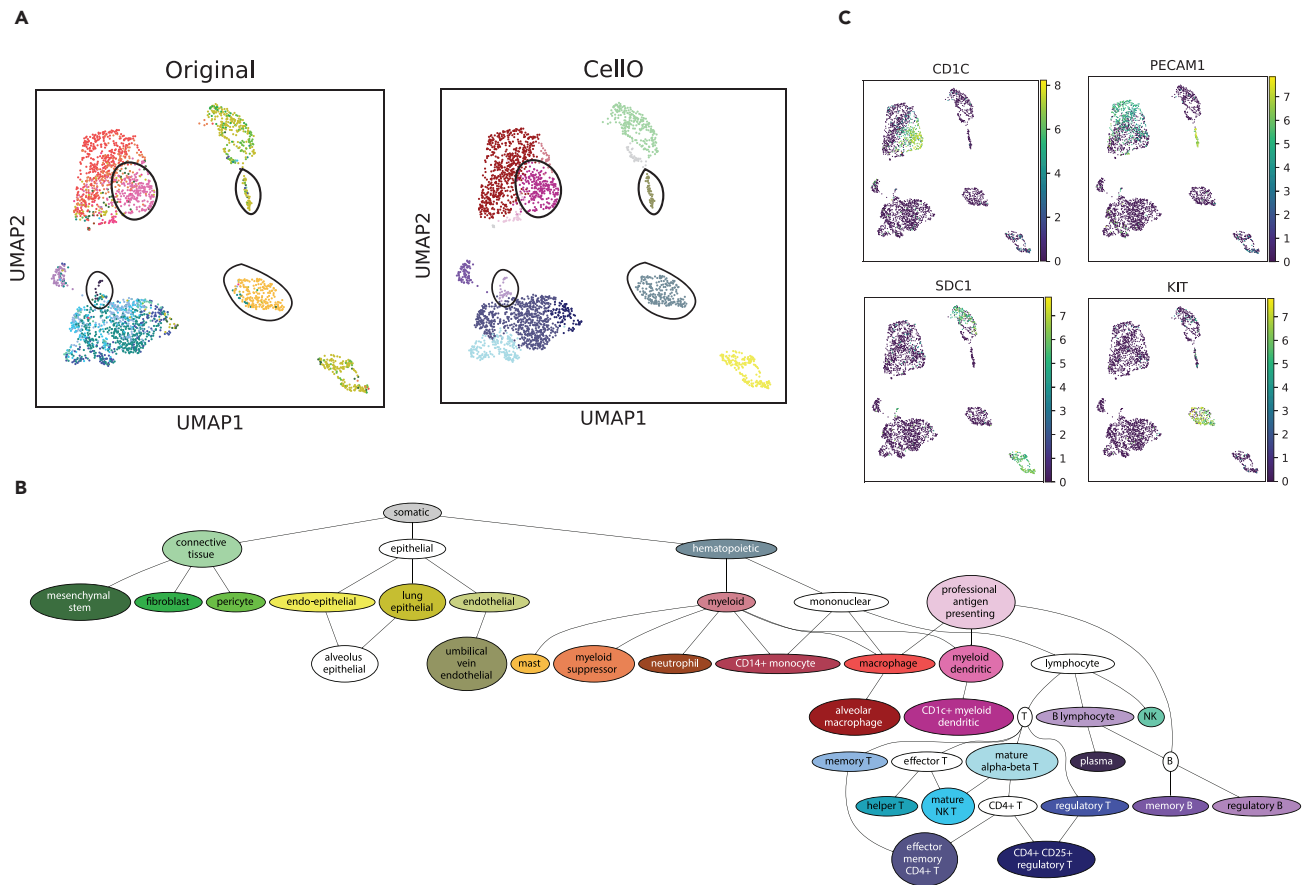


Figure 8. Examination of Cello on diseased cells

(A) UMAP plots of lung adenocarcinoma tumor LX675 from [Laughney et al. \(2020\)](#) colored by Cello's output using IR and a Leiden resolution parameter of 1.0 (left) and the original cell type labels provided by the authors. We highlight four subpopulations comprising putative CD1c + myeloid dendritic cells (top left), endothelial cells (top right), plasma cells (bottom left), and mast cells (bottom right).

(B) The legend for coloring cells in (A).

(C) UMAP plots of cells colored by their expression, in units $\log(\text{TPM}+1)$, of CD1c, a marker for CD1c+ myeloid dendritic cells, PECAM1, a marker for endothelial cells, SDC1, a marker for plasma cells, and KIT, a marker for mast cells.

cells in the Cell Ontology. Similarly, Cello annotated the glial cells as neuron associated cells. These examples again highlight Cello's ability to annotate cells with a term higher in the ontology DAG when it is unsure about a cell's more specific cell type.

Finally, we examined Cello's classification on eight lung adenocarcinoma tumor samples from [Laughney et al. \(2020\)](#) that were sequenced with 10x. This data set provides a good, yet challenging test set for Cello due to (a) the heterogeneity of constituent cell types, (b) the fact that it contains cell types absent from Cello's training set, and (c) these cells originate from diseased tissue, thereby providing insight into how Cello will perform on non-healthy samples. We compared Cello's output to the annotations provided by the authors, which were the result of a custom, *in silico* cell type annotation pipeline.

Overall, we found a high correspondence between the labels provided by Cello and the authors' annotation pipeline (Figures 8; Figures S5–S12). When the two methods differed, we attempted to determine which method is likely to be correct using manual inspection of known marker genes and in many of these discrepancies, Cello produced the correct cell type labels. For example, in tumor LX675 and LX682, Cello correctly annotates putative endothelial cells (Figures 8; Figure S5), based on expression of PECAM1 (Figure 8C) and CD34 (data not shown), whereas the authors of these data labeled these as epithelial cells. In tumor LX675, Cello labeled the myeloid dendritic cell population as CD1c+ myeloid dendritic cells, a more specific cell type than that provided by the authors (Figure 8). In tumor LX682, Cello correctly annotates the

putative fibroblasts as evidenced by their expression of FAP (Puré and Blomberg 2018) (Figure S5) and S100A4 (Strutz et al., 1995; data not shown). In tumor LX679, Cello appears to correctly annotate the plasmacytoid dendritic cell population as evidenced by the expression of IL3RA (Figure S6), a known marker for this cell type (Collin et al. 2013). We also note that for cells whose likely true cell types are absent from Cello's training set, Cello was able to label these cells using a correct but more general cell type. For example, in tumor LX675, mast cells were absent from Cello's training set, yet Cello accurately labeled these cells as hematopoietic cells (Figure 8). Similarly, when Cello was unsure of the labels for cells of cell types for which training data were sparse, such as for plasma cells and pericytes, Cello labeled these cells as more general cell type terms lymphocyte of B lineage and connective tissue cell, respectively (Figure 8). Lastly, we note some of the most clear-cut cases in which the authors' pipeline produced correct labels but Cello erred. For example, in tumor LX679, Cello labeled many of the epithelial cell types as prostate epithelial cells (Figure S6), which is clearly incorrect given that these are lung tumor samples. Other errors produced by Cello are due to rarer cell types within the sample clustering together with a more common, similar cell type. For example, in tumor LX682, the myeloid dendritic cells are labeled by Cello as macrophages due to these cells clustering together with the macrophage population (Figure S5).

Evaluation of robustness to clustering

In order to test Cello's robustness to cluster granularity when clustering single-cell data, we tested Cello on the Zheng et al. PBMC data set and Laughney et al. lung cancer data set using differing values for Leiden's resolution parameter. On the Zheng et al. data set, we tested Cello with five values for the resolution parameter and found both CLR and IR to perform similarly across these settings (Figure S15). In contrast to the Zheng et al. data set, the lack of robust ground-truth labels for the Laughney et al. data set made it difficult to perform a similar assessment. Instead, on this data set, we ran Cello using the default resolution of 1.0, as well as a higher resolution of 8.0, and then manually inspected the differences in the outputs produced by these two settings (Figures S5–S12). Overall, we found that the higher resolution led Cello to classify clusters with more granularity but at the cost of some errors. For example, on tumor LX679 and LX682, Cello was able to correctly classify the putative myeloid dendritic cells at the resolution of 8.0, whereas at the default resolution of 1.0, these cells were labeled incorrectly as alveolar macrophages (Figures S5 and S6). This is likely due to the fact that the dendritic cells were subsumed into the macrophage cluster at the lower resolution. However, in other instances, the higher resolution led to errors. These errors are likely due to the fact that at a more granular clustering resolution, each cluster's mean expression value remained sparse due to the averaging of fewer cells. For example, in tumor LX679, cells correctly labeled as respiratory epithelial cells at a resolution of 1.0 were incorrectly labeled as prostate epithelial cells at a resolution of 8.0 (Figure S6).

User-friendly software

We provide a Python package for running Cello, using either IR or CLR, on a user-provided gene expression matrix (<https://github.com/deweylab/Cello>). Cello reduces the burden of reformatting and preprocessing an input expression matrix by accepting a variety of input file formats, including comma or tab-separated text files and HDF5, and by accepting expression data in a variety of units including counts or transcripts per million (TPM). To address the scenario in which the input data set's genes do not match those expected by the pre-trained classifiers, we provide functionality for a user to re-train the models on the bulk RNA-seq training set with a custom gene set. On a personal laptop, training a new classifier took 31 minutes to train IR and 11 minutes to train CLR on the full set of 58,243 GENCODE genes. Training time is reduced when trained on a smaller set of genes (e.g., only protein-coding genes). We also note that the time required for Cello to perform classification is low because of the fact that it uses pre-trained logistic regression classifiers operating on cluster-averaged expression profiles. On the Zheng et al. (2017) data set, Cello took six minutes to run on a personal laptop (including time for clustering), whereas scMatch required six hours and nine minutes (run with five cores), and SingleR required between nine and 22 minutes depending on the reference set used. Finally, we note that the relative performance of IR and CLR varies across cell types. To guide a user on their selection of either IR or CLR, we provide the average precision values achieved by both methods on each cell type in the bulk RNA-seq validation set (Table S1, Bulk RNA-seq Validation Set Metrics). We also provide average precision values and F1-scores achieved by both methods on each cell type on the test set of 4,936 non-droplet-based single cells whose cell types were present in the bulk RNA-seq training set (Table S2, Single-cell RNA-seq Test Set Metrics).

Lastly, we note that due to Cello's comprehensive training set, which comprises cell types from many organs and tissues, some of Cello's errors are due to Cello annotating cells using a cell type that is unique to a tissue type that differs from the known tissue type of the target sample. For example, in some lung cancer tumors from Laughney et al., endothelial cells were classified as "endothelial cell of the umbilical vein" (Figures S5, S6, and S9), which is clearly incorrect given that these samples were taken from the lung. This error is likely due to the fact that the endothelial cells in Cello's training set largely originate from umbilical cord samples. Because these errors can be easily caught by the user, the Cello package enables a user to fix such errors by enabling the user to supply a blacklist of tissue types that do not pertain to the target sample. Cello then uses edges between the Cell Ontology and the Uberon ontology (which encodes anatomical entities; Mungall et al., 2012) to filter out cell types from Cello's output that are uniquely located in the blacklisted tissue types. For example, by blacklisting the Uberon term "umbilical vein", Cello correctly classifies the endothelial cells from Laughney et al.

Interpretability of models

Cello makes extensive use of linear models, which are particularly amenable to interpretation especially when the coefficients are sparse (Gleicher 2013). Although Cello's models are not regularized to be sparse (as in Gleicher 2013), we sparsify them by selecting the top ten genes per cell type according to the magnitude of the coefficients associated with each gene within each cell type's one-vs-rest binary classification model, which is used for Cello's IR classifier. To enable their interpretation, we present a web-based tool, the Cello Viewer, for exploring these discriminative genes uncovered by the models (<https://uwgraphics.github.io/CelloViewer/>). The tool supports two modes of operation: a "cell-centric" mode (Figure 9A) and a "gene-centric" mode (Figure 9B). In the cell-centric mode, the user can select cell types via a graphical display of the Cell Ontology in order to view and compare the most important genes for distinguishing those cell types. In the gene-centric view, the user can select genes and explore which cell types these genes are most important for distinguishing from the remaining cell types. The Cello Viewer uses an interactive display of the Cell Ontology's graph to enable the user to navigate between cell types across the ontology.

We found that across diverse cell types, many known cell-type-specific marker genes were recovered by the Cello models and are presented by the Cello Viewer. For example, CD3D, CD3E, and CD3G, which are canonical markers for T cells, were all present within the top ten genes ranked according to the magnitude of their coefficients within the binary logistic regression model used for distinguishing T cells from all other cell types. Similarly, CD4 and CD8 were present in the top genes for the CD4+ T cell and CD8+ T cell models, respectively (Figure 9). In a more complex example, the genes GCG, LOXL4, DPP4, GC, and FAP, known markers for pancreatic alpha cells, and INS, IAPP, and ADCYAP1, known markers for pancreatic beta cells (Segerstolpe et al. (2016)), all appear within the top ten genes for their respective cell types.

Interestingly, certain genes appear in the top ten coefficients for broad cell types but not more specific cell types, indicating that Cello is able to find signals specific to broad cell type categories. For example, DDX4 appeared in the top ten genes for distinguishing germ line cells but did not appear within the top ten genes for any of the more specific germ cell subtypes. DDX4 is known to be expressed in germ cells across both sexes (Hickford et al., 2011). Similarly, the gene NRG1 appeared in the top ten genes for distinguishing precursor cells but did not appear within the top ten genes for any of the more specific cell types that are descendants of precursor cells within the ontology. NRG1 is known to play a role in the development of a number of organ systems (Lemmens et al. 2007; Mei and Xiong 2008).

DISCUSSION

In this work, we explore the application of hierarchical classification algorithms to cell type classification with the Cell Ontology using a well-curated set of human primary cell RNA-seq samples. This data set may prove useful for future investigations of cell type expression patterns or for use in cell type deconvolution methods (Aran et al., 2019; Newman et al., 2015). We demonstrate that the trained classifiers perform well across cell types in diverse single-cell data sets and outperformed existing cell type annotation methods when trained on their comprehensive reference sets. We packaged these classifiers into an easy-to-run Python package called Cello.

In our exploration of methods for correcting the independent one-vs-rest classifiers, we found that discriminative methods outperformed the generative BNC approach implemented in URSA (Lee et al., 2013). We hypothesize that BNC suffers in comparison due to two causes. First, BNC's probabilistic model makes strong assumptions regarding the generative process of classifier scores and true cell type assignments.

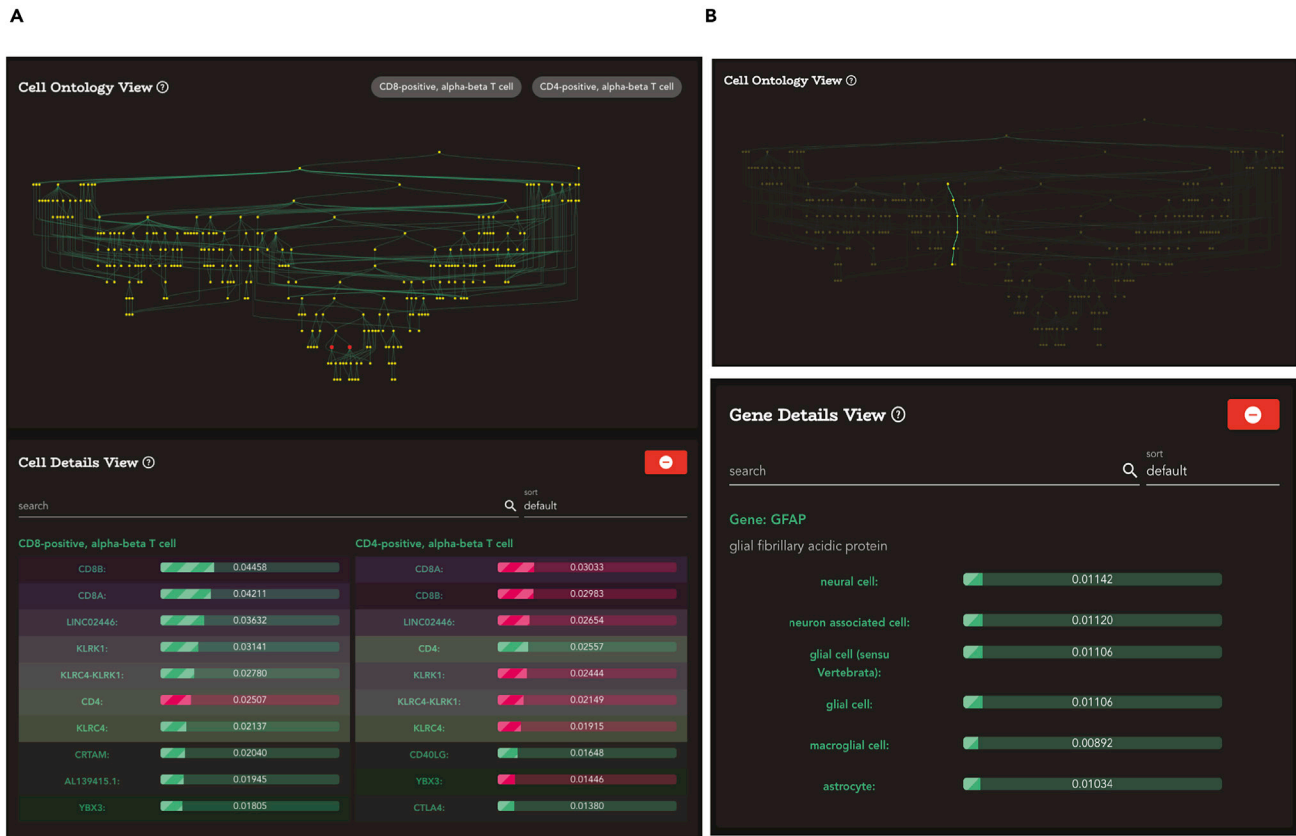


Figure 9. The Cello Viewer

Screenshots of the Cello Viewer web application for enabling the exploration of cell-type-specific expression signatures across the Cell Ontology. (A) Comparing the top ten genes between CD4+ T cells and CD8+ T cells (red nodes in the Graph View) ranked by the magnitude of their coefficients in their corresponding models. Genes that are shared between the two lists are highlighted with the same color. The Cello Viewer displays genes whose expressions are both positively correlated (green) and negatively correlated (red) with the selected cell types. (B) A screenshot of the gene-centric mode of the Cello Viewer with GFAP, an astrocyte marker, selected. For a given selected gene, the Cello Viewer will display the cell types within the DAG (top) and in list form (bottom) for which the selected gene appears within the top ten genes ranked by each model's coefficients.

Second, BNC requires estimating the conditional probability distribution of each classifier's output scores (i.e., distance from the decision boundary) conditioned on the true cell type labels, which may be difficult to estimate accurately given the limited quantity of training data available for each cell type.

By using linear models, Cello's trained parameters are easily interpreted as cell-type-specific signatures across the ontology. However, we note that since certain cell types undergo similar sorting and preparation procedures (e.g., fluorescence-activated cell sorting), it remains unclear to what extent these procedures affect gene expression and thus confound cell type. We sought to mitigate this effect by using data from a diversity of studies. We also note that the CLR algorithm may help to further mitigate this effect since the binary classifiers trained in this framework for each cell type condition on the sample belonging to the parent cell types. Thus, for a given cell type, if samples of its parent cell types were prepared through similar procedures, the learned model parameters for that cell type will better capture biological cell type signatures.

Limitations of the study

There are a number of avenues that require further investigation. First, future work will entail curating comprehensive training sets from the SRA for other species such as mouse. This work will rely partly on future inclusion of standardized mouse metadata in the MetaSRA.

Second, Cello is a cluster-based annotation method, and thus, its accuracy relies, in part, on the robustness of the clustering algorithm. If the clustering is too coarse, rare cell types may be missed. If clustering is too

fine, the algorithm may not be combining enough data to accurately annotate each cluster. Determining the optimal clustering in scRNA-seq data is a challenging, open problem that will require further investigation (Kiselev et al. 2019). Nonetheless, we demonstrated that Cello accurately classified a number of diverse data sets using Leiden's default parameter in the Scanpy package.

Third, we note that calibrating discriminative models trained on bulk RNA-seq data and applying them to single-cell data is challenging. In this work, we developed techniques for closing the gap between the performance of Cello when evaluated with average precision versus when evaluated with F1-score. The very high average precision scores across many cell types indicate that Cello is learning an accurate representation of these cell types and that with better calibration, Cello's accuracy when making binary yes-no decisions for each cell type could be improved. Future work will investigate alternative approaches to calibrating Cello's models in order to improve Cello's binary cell type decisions.

Fourth, the Cell Ontology encodes anatomical and functional relationships between cell types; however, there exist a number of other relationships between cell types that could be utilized to improve accuracy. Such examples include lineage-based relationships (i.e., one cell type derives from another cell type; Yuan et al., 2020) or evolutionary relationships between extant cell types and ancient cell types (Arendt et al., 2016; Liang et al., 2018). For example, the evolutionary relationships between cell types may be utilized to address inconsistencies in the independent classifiers approach that arise when certain cell types share a parent cell type via the currently encoded "is a" relationship but are purported to have divergent evolutionary origins.

Fifth and finally, we expect the performance of hierarchical classifiers to improve as both more data are collected and as the Cell Ontology is expanded. Most importantly, we expect the calibration of the classifiers to improve as more training data become available for each cell type. More training data will be collected both as data are continually added to the SRA and as improvements are made to the SRA's metadata, thereby allowing retrieval of previously undiscovered primary cell samples.

Resource availability

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the Lead Contact, Colin Dewey (colin.dewey@wisc.edu).

Data and code availability

A Python package for running Cello can be found on GitHub: <https://github.com/deweylab/Cello>. The data used in this work can be found on <https://doi.org/10.5281/zenodo.4289064>. The Cello Viewer can be accessed at <https://uwgraphics.github.io/CelloViewer/>. The code implementing the Cello Viewer can be found on GitHub: <https://github.com/uwgraphics/CelloViewer>. All code for performing the experiments in this work can be found on GitHub: <https://github.com/deweylab/cell-type-classification-paper>.

Materials availability

This study did not generate new unique reagents.

METHODS

All methods can be found in the accompanying [Transparent methods supplemental file](#).

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.isci.2020.101913>.

ACKNOWLEDGMENTS

M.N.B. thanks Gary H. Bernstein, John Steill, Ron Stewart, and Christina Kendzioriski for helpful conversations. This project has been made possible in part by grant U54 AI117924 from the National Institutes of Health and grant 2018-182626 from the Chan Zuckerberg Initiative DAF, an advised fund of Silicon Valley Community Foundation. M.N.B. acknowledges support of the Computation and Informatics in Biology and Medicine Training Program funded by grant: 5T15LM007359 from the National Institutes of Health. M.N.B.

also acknowledges support from the Morgridge Institute for Research. M.G. acknowledges support from the National Science Foundation award 1841349.

AUTHOR CONTRIBUTIONS

M.N.B. led the conceptual development of the ideas present in this article and implemented both the Cello software and experiments. C.N.D. supervised the conceptual development and implementation of the software and experiments. M.N.B. and C.N.D. wrote the manuscript. Z.M. implemented the Cello Viewer. M.G. supervised the development of the Cello Viewer. M.N.B., Z.M., M.G., and C.N.D. contributed to the conceptual design of the Cello Viewer.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: July 30, 2020

Revised: October 28, 2020

Accepted: December 2, 2020

Published: January 22, 2021

REFERENCES

- Abdelaal, T., Michielsen, L., Cats, D., Hoogduin, D., Mei, H., Marcel, J., Reinders, T., and Ahmed, M. (2019). A comparison of automatic cell identification methods for single-cell RNA sequencing data. *Genome Biol.* 20, 194.
- Alquicira-Hernandez, J., Anuja, S., Ji, H.P., Nguyen, Q., Joseph, E., and Powell. (2019). scPred: accurate supervised method for cell-type classification from single-cell RNA-seq data. *Genome Biol.* 20, 264.
- Aran, D., Hu, Z., and Butte, A.J. (2017). xCell: digitally portraying the tissue cellular heterogeneity landscape. *Genome Biol.* 18, 220.
- Aran, D., Looney, A.P., Liu, L., Wu, E., Fong, V., Hsu, A., Chak, S., Naikawadi, R.P., Wolters, P.J., Abate, A.R., et al. (2019). Reference-based analysis of lung single-cell sequencing reveals a transitional profibrotic macrophage. *Nat. Immunol.* 20, 163–172.
- Arendt, D., Musser, J.M., Baker, C.V.H., Bergman, A., Cepko, C., Erwin, D.H., Pavlicev, M., Schlosser, G., Widder, S., Laubichler, M.D., et al. (2016). The origin and evolution of cell types. *Nat. Rev. Genet.* 17, 744–757.
- Bard, J., Rhee, S.Y., and Ashburner, M. (2005). An ontology for cell types. *Genome Biol.* <https://doi.org/10.1186/gb-2005-6-2-r21>.
- Becht, E., McInnes, L., Healy, J., Charles-Antoine, D., Kwok, I.W.H., Ng, L.G., Ginhoux, F., and Newell, E.W. (2018). Dimensionality reduction for visualizing single-cell data using UMAP. *Nat. Biotechnol.* <https://doi.org/10.1038/nbt.4314>.
- Bernstein, M.N., Doan, A., and Dewey, C.N. (2017). MetaSRA: normalized human sample-specific metadata for the sequence Read archive. *Bioinformatics* 33, 2914–2923.
- Bray, N.L., Pimentel, H., Melsted, P., and Pachter, L. (2016). Near-optimal probabilistic RNA-seq quantification. *Nat. Biotechnol.* 34, 525–527.
- Collin, M., McGovern, N., and Haniffa, M. (2013). Human dendritic cell subsets. *Immunology* 140, 22–30.
- Fernández, J.M., de la Torre, V., Richardson, D., Royo, R., Puigròs, M., Moncunill, V., Fragkogianni, S., Clarke, L.; BLUEPRINT Consortium, and Flicek, P., et al. (2016). The blueprint data analysis portal. *Cell Syst.* 3, 491–495.e5.
- Gleicher, M. (2013). Explainers: expert explorations with crafted projections. *IEEE Trans. Vis. Comput. Graph.* 19, 2042–2051.
- Gremse, M., Chang, A., Ida, S., Grote, A., Scheer, M., Ebeling, C., and Schomburg, D. (2011). The BRENDA tissue ontology (BTO): the first all-integrating ontology of all organisms for enzyme sources. *Nucleic Acids Res.* 39, D507–D513.
- Hickford, D.E., Stephen, F., Pask, A.J., Shaw, G., and Renfree, M.B. (2011). DDX4 (VASA) is conserved in germ cell development in marsupials and monotremes. *Biol. Reprod.* 85, 733–743.
- Hou, R., Elena, D., Alistair, R., and Forrest, R. (2019). scMatch: a single-cell gene expression profile Annotation tool using reference datasets. *Bioinformatics* 35, 4688–4695.
- Jaitin, D.A., Kenigsberg, E., Keren-Shaul, H., Elefant, N., Paul, F., Zaretsky, I., Alexander, M., Cohen, N., Jung, S., Tanay, A., et al. (2014). Massively parallel single-cell RNA-seq for marker-free decomposition of tissues into cell types. *Science* 343, 776–779.
- de Kanter, J.K., Lijnzaad, P., Candelli, T., Margaritis, T., Frank, C., and Holstege, P. (2019). CHETAH: a selective, hierarchical cell type identification method for single-cell rna sequencing. *Nucleic Acids Res.* 47, e95.
- Kiselev, V.Y., Andrews, T.S., and Martin, H. (2019). Challenges in unsupervised clustering of single-cell RNA-seq data. *Nat. Rev.* 20, 273–282.
- Laughney, A.M., Hu, J., Nathaniel, R.C., Bakhroum, S.F., Setty, M., Lavallée, V.P., Xie, Y., Masilionis, I., Carr, A.J., Kottapalli, S., et al. (2020). Regenerative lineages and immune-mediated pruning in lung cancer metastasis. *Nat. Med.* 26, 259–269.
- Lee, Y.S., Krishnan, A., Zhu, Q., and Olga, G.T. (2013). Ontology-aware classification of tissue and cell-type signals in gene expression profiles across platforms and technologies. *Bioinformatics* 29, 3036–3044.
- Leinonen, R., Sugawara, H., Martin, S., and International Nucleotide Sequence Database Collaboration. (2011). The sequence Read archive. *Nucleic Acids Res.* 39, D19–D21.
- Lemmens, K., Doggen, K., and de Keulenaer, G.W. (2007). Role of neuregulin-1/ErbB signaling in cardiovascular physiology and disease: implications for therapy of heart failure. *Circulation* 116, 954–960.
- Liang, C., Musser, J.M., Cloutier, A., Prum, R.O., and Wagner, G.P. (2018). Pervasive correlated evolution in gene expression shapes cell and tissue type transcriptomes. *Genome Biol. Evol.* 10, 538–552.
- Lieberman, Y., Rokach, L., and Shay, T. (2018). CaSTLe - classification of single cells by transfer learning: harnessing the power of publicly available single cell RNA sequencing experiments to annotate new experiments. *PLoS One* 13, e0205499.
- Lizio, M., Harshbarger, J., Abugessaisa, I., Noguchi, S., Kondo, A., Severin, J., Mungall, C., Arenillas, D., Mathelier, A., Medvedeva, Y.A., et al. (2017). Update of the FANTOM web resource: high resolution transcriptome of diverse cell types in mammals. *Nucleic Acids Res.* 45, D737–D743.
- Ma, F., and Pellegrini, M. (2020). ACTINN: automated identification of cell types in single cell RNA sequencing. *Bioinformatics* 36, 533–538.

- Mabbott, N.A., Kenneth Baillie, J., Brown, H., Freeman, T.C., and Hume, D.A. (2013). An expression Atlas of human primary cells: inference of gene function from coexpression networks. *BMC Genomics* *14*, 632.
- La Manno, G., Gyllborg, D., Codeluppi, S., Nishimura, K., Salto, C., Amit, Z., Borm, L.E., Stott, S.R.W., Toledo, E.M., Villaescusa, J.C., et al. (2016). Molecular diversity of midbrain development in mouse, human, and stem cells. *Cell* *167*, 566–580.e19.
- Mei, L., and Xiong, W.C. (2008). Neuregulin 1 in neural development, synaptic plasticity and schizophrenia. *Nat. Rev. Neurosci.* *9*, 437–452.
- Monaco, G., Lee, B., Xu, W., Mustafah, S., Hwang, Y.Y., Carré, C., Burdin, N., Visan, L., Ceccarelli, M., Poidinger, M., et al. (2019). RNA-seq signatures normalized by mRNA abundance allow absolute deconvolution of human immune cell types. *Cell Rep.* *26*, 1627–1640.e7.
- Mungall, C.J., Torniai, C., Gkoutos, G.V., Lewis, S.E., and Haendel, M.A. (2012). Uberon, an integrative multi-species anatomy ontology. *Genome Biol.* *13*, R5.
- Newman, A.M., Liu, C.L., Green, M.R., Gentles, A.J., Feng, W., Xu, Y., Chuong, D.H., Maximilian, D., and Alizadeh, A.A. (2015). Robust enumeration of cell subsets from tissue expression profiles. *Nat. Methods* *12*, 453–457.
- Notaro, M., Schubach, M., Robinson, P.N., and Valentini, G. (2017). Prediction of human phenotype ontology terms by means of hierarchical ensemble methods. *BMC Bioinformatics* *18*, 449.
- Obozinski, G., Lanckriet, G., Grant, C., Jordan, M.I., and Noble, W.S. (2008). Consistent probabilistic outputs for protein function prediction. *Genome Biol.* *9*, S6.
- Pérez-Silva, J.G., Araujo-Voces, M., and Quesada, V. (2018). nVenn: generalized, quasi-proportional venn and euler diagrams. *Bioinformatics* *34*, 2322–2324.
- Picelli, S., Björklund, Å.K., Faridani, O.R., Sagasser, S., Winberg, G., and Sandberg, R. (2013). Smart-seq2 for sensitive full-length transcriptome profiling in single cells. *Nat. Methods* *10*, 1096–1098.
- Pliner, H.A., Shendure, J., and Cole, T. (2019). Supervised classification enables rapid annotation of cell atlases. *Nat. Methods* *16*, 983–986.
- Puré, E., and Blomberg, R. (2018). Pro-tumorigenic roles of fibroblast activation protein in cancer: back to the basics. *Oncogene* *37*, 4343–4357.
- Segerstolpe, Å., Palasantza, A., Eliasson, P., Andersson, E.M., Anne-Christine, A., Sun, X., Picelli, S., Sabirsh, A., Clausen, M., Bjursell, M.K., et al. (2016). Single-cell transcriptome profiling of human pancreatic islets in Health and type 2 diabetes. *Cell Metab.* *24*, 593–607.
- Shao, X., Liao, J., Lu, X., Xue, R., Ni, A., and Fan, X. (2020). scCATCH: automatic annotation on cell types of clusters from single-cell RNA sequencing data. *iScience* *23*, 100882.
- Sloan, C.A., Esther, T.C., Davidson, J.M., Venkat, S.M., Strattan, J.S., Hitz, B.C., Gabdank, I., Narayanan, A.K., Ho, M., Lee, B.T., et al. (2016). ENCODE data at the ENCODE portal. *Nucleic Acids Res.* *44*, D726–D732.
- Strutz, F., Okada, H., Lo, C.W., Danoff, T., Carone, R.L., Tomaszewski, J.E., and Neilson, E.G. (1995). Identification and characterization of a fibroblast marker: FSP1. *J. Cell Biol.* *130*, 393–405.
- Tan, Y., and Cahan, P. (2019). SingleCellNet: a computational tool to classify single cell RNA-seq data across platforms and across species. *Cell Syst.* *9*, 207–213.e2.
- Traag, V.A., Waltman, L., and van Eck, N.J. (2019). From louvain to leiden: guaranteeing well-connected communities. *Sci. Rep.* *9*, 5233.
- Wolf, F.A., Angerer, P., and Theis, F.J. (2018). SCANPY: large-scale single-cell gene expression data analysis. *Genome Biol.* *19*, 15.
- Yuan, M., Yang, X., Lin, J., Cao, X., Chen, F., Zhang, X., Li, Z., Zheng, G., Wang, X., Chen, X., et al. (2020). Alignment of cell lineage trees elucidates genetic programs for the development and evolution of cell types. *iScience* *23*, 101273.
- Zhang, A.W., O’Flanagan, C., Chavez, E.A., Lim, J.L.P., Ceglia, N., McPherson, A., Wiens, M., Walters, P., Chan, T., Hewitson, B., et al. (2019a). Probabilistic cell-type Assignment of single-cell RNA-seq for tumor microenvironment profiling. *Nat. Methods* *16*, 1007–1015.
- Zhang, X., Lan, Y., Xu, J., Quan, F., Zhao, E., Deng, C., Luo, T., Xu, L., Liao, G., Yan, M., et al. (2019b). CellMarker: a manually curated resource of cell markers in human and mouse. *Nucleic Acids Res.* *47*, D721–D728.
- Zheng, G.X., Terry, J.M., Belgrader, P., Paul, R., Bent, Z.W., Ryan, W., Ziraldo, S.B., Wheeler, T.D., McDermott, G.P., Zhu, J., et al. (2017). Massively parallel digital transcriptional profiling of single cells. *Nat. Commun.* *8*, 14049.


iScience, Volume 24

Supplemental Information

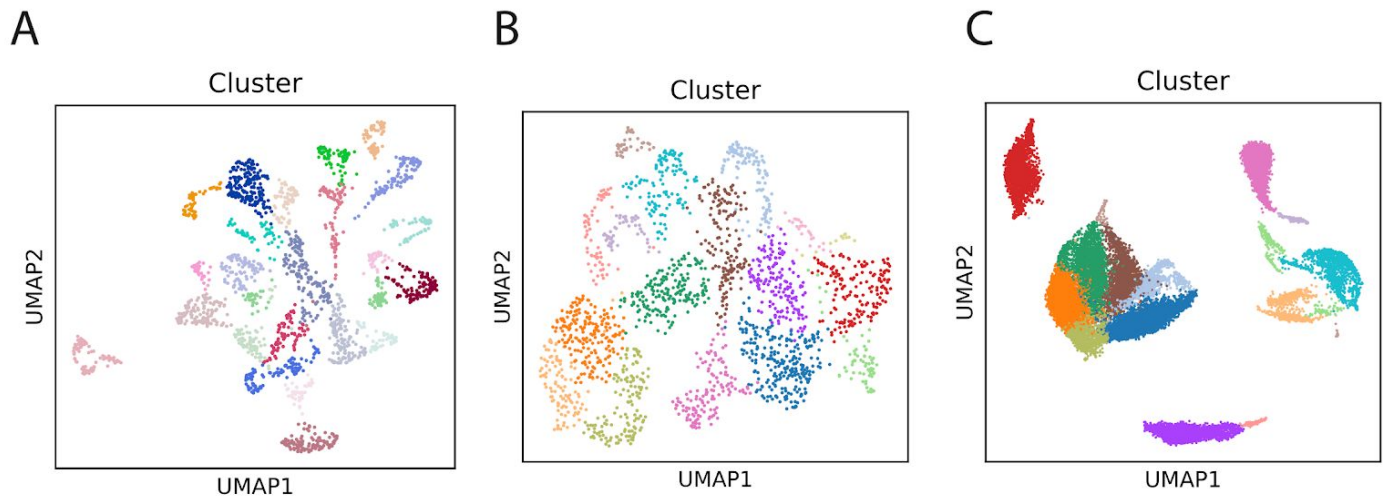
**CellO: comprehensive and hierarchical
cell type classification of human
cells with the Cell Ontology**

Matthew N. Bernstein, Zhongjie Ma, Michael Gleicher, and Colin N. Dewey

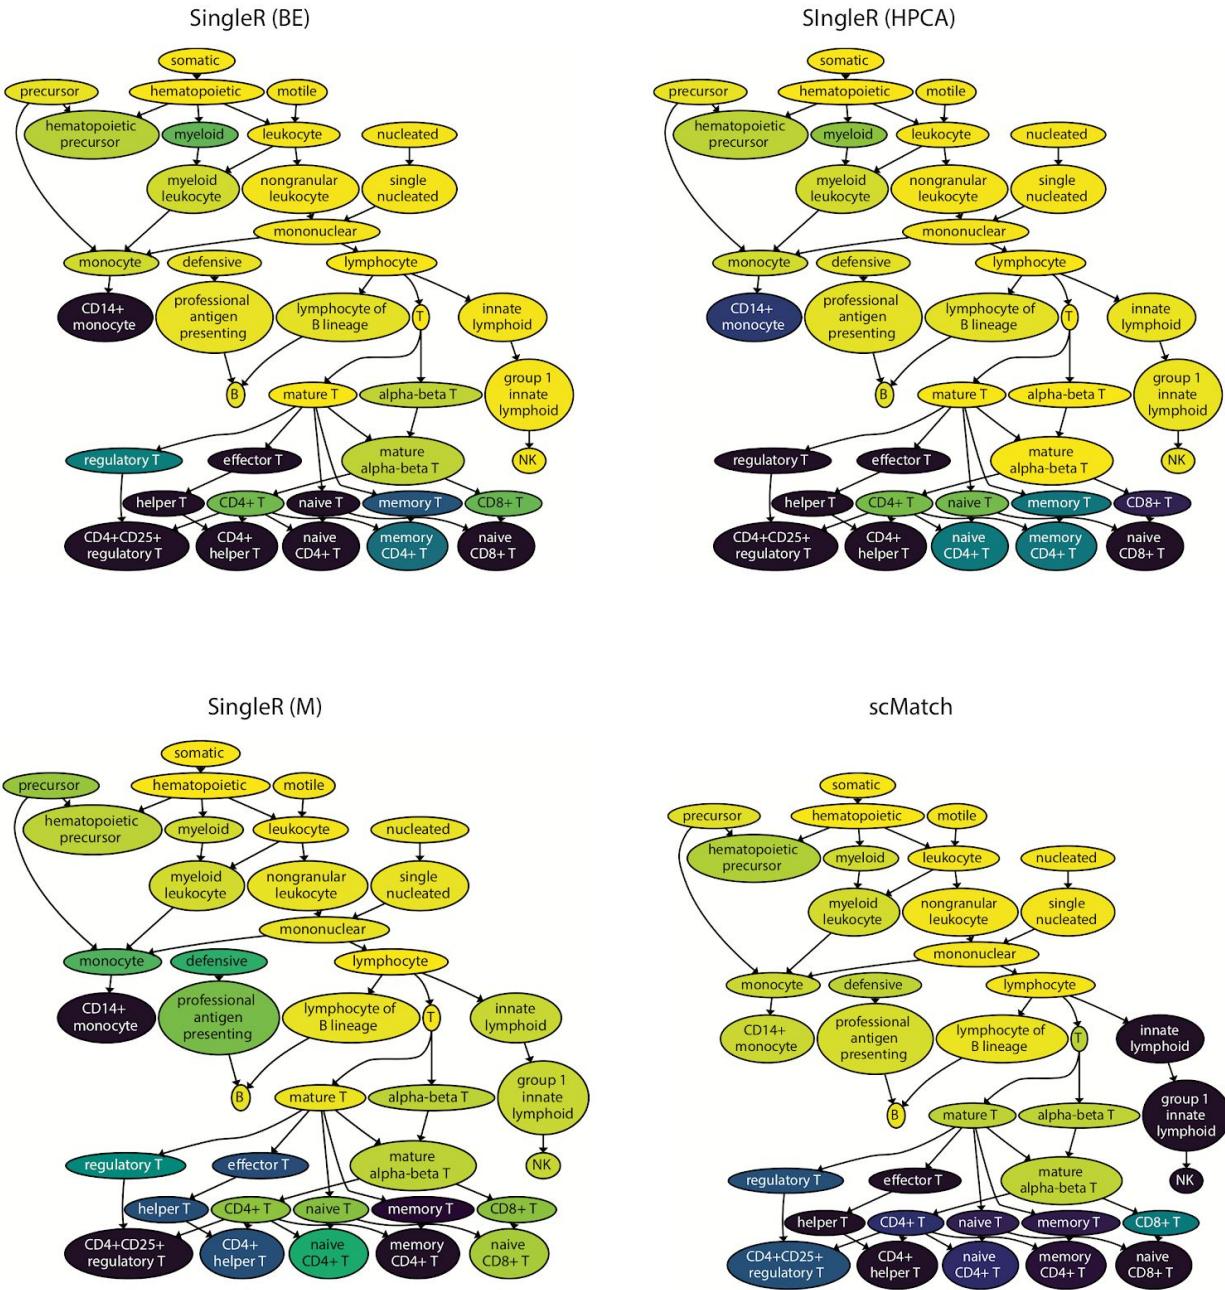
Supplemental Figures



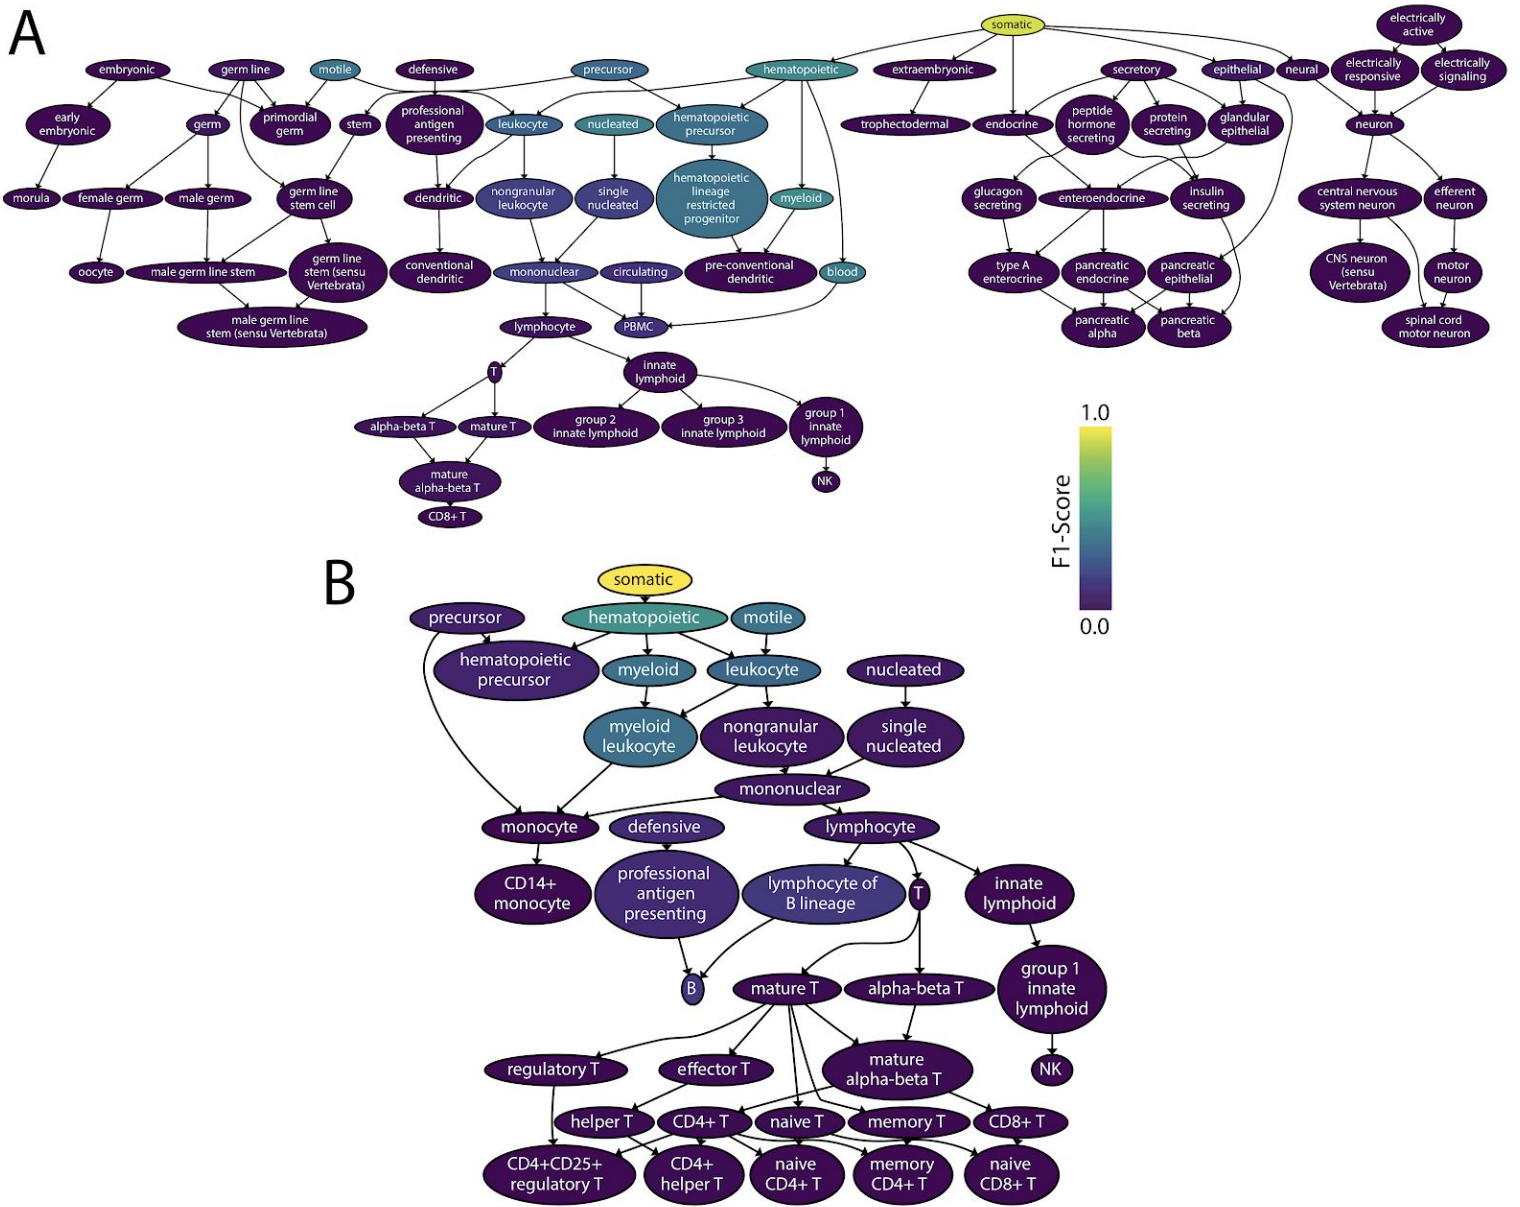
Supplemental Figure 1. **Distribution of edge inconsistencies (Related to Figure 3)**. The cumulative distribution function over the difference in probability between the parent and child classifiers for all edges for which either the parent or child classifier output a probability greater than 0.01.



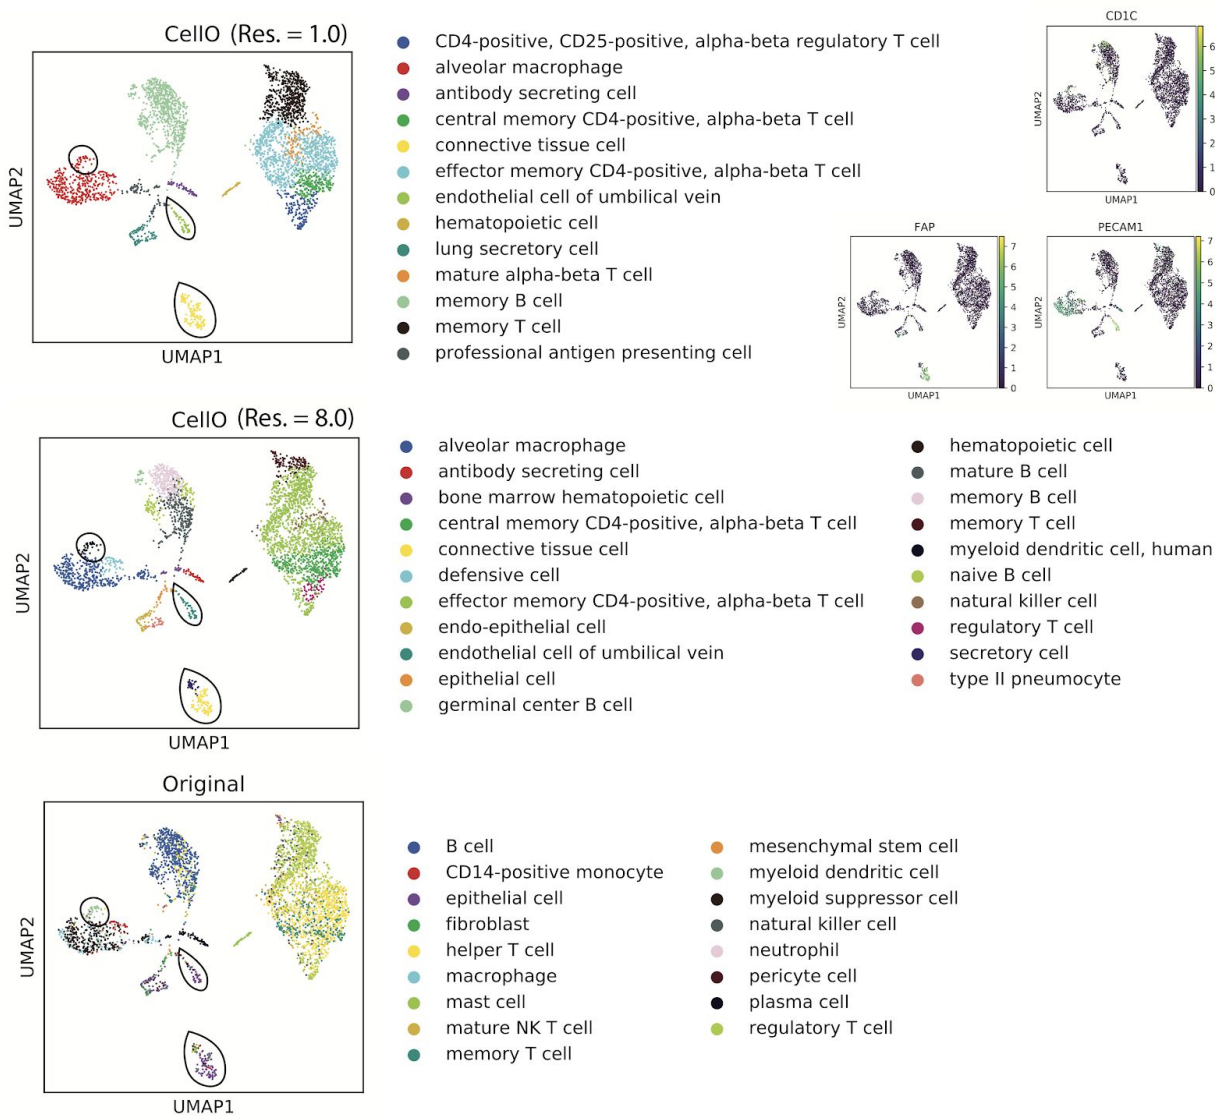
Supplemental Figure 2. **Cluster assignments for single-cell datasets evaluated in this work (Related to Figures 6C, 7A-B).** UMAP plots of the single-cell datasets examined this work where each cell is colored according to its Cello cluster assignment. These clusters were aggregated to compute a mean expression profile and on which Cello was run. Plots are displayed for (A) the Segerstolpe et al. (2016) dataset of healthy pancreatic cells, (B) the La Manno et al. (2016) dataset of fetal neural cells, and (C) the Zheng et al. (2017) dataset of PBMCs.



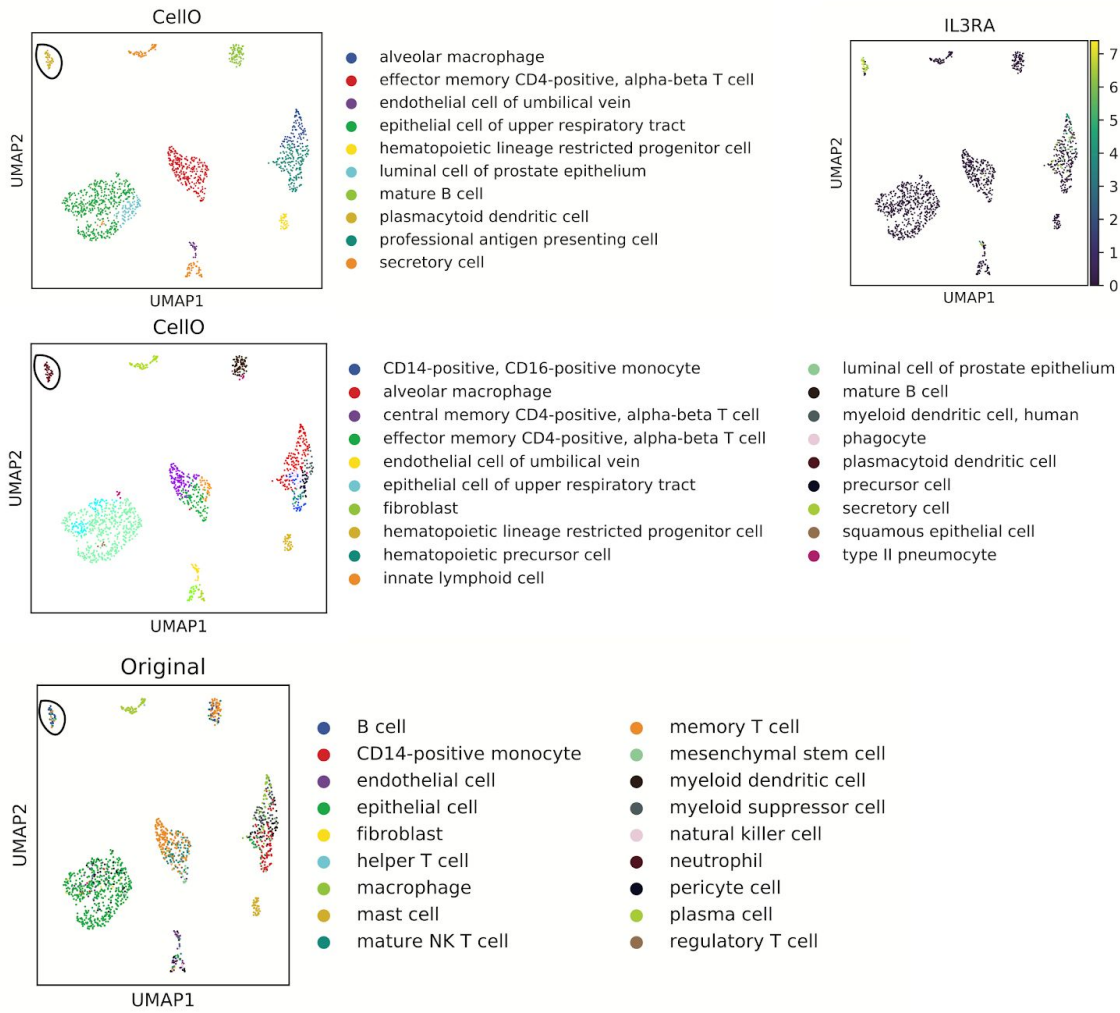
Supplemental Figure 3. **PBMC F1-scores produced by existing methods (Related to Figures 6A and 6D).** The subgraph spanning the PBMC's from Zheng et al. (2017) where each node is colored according to the F1-scores produced by SingleR with the Blueprint+Encode reference (BE), SingleR with the Human Primary Cell Atlas reference (HPCA), SingleR with the Monaco et al. (2019) reference (M), and scMatch.




Supplemental Figure 4. **PBMC F1-scores produced by SingleR using CellIO's training set (Related to Figures 4D and 6A).** (A) The subgraph spanning the non-droplet-based cells and (B) the subgraph spanning the PBMC's from Zheng et al. (2017) where each node is colored according to the F1-scores produced by SingleR with CellIO's reference.



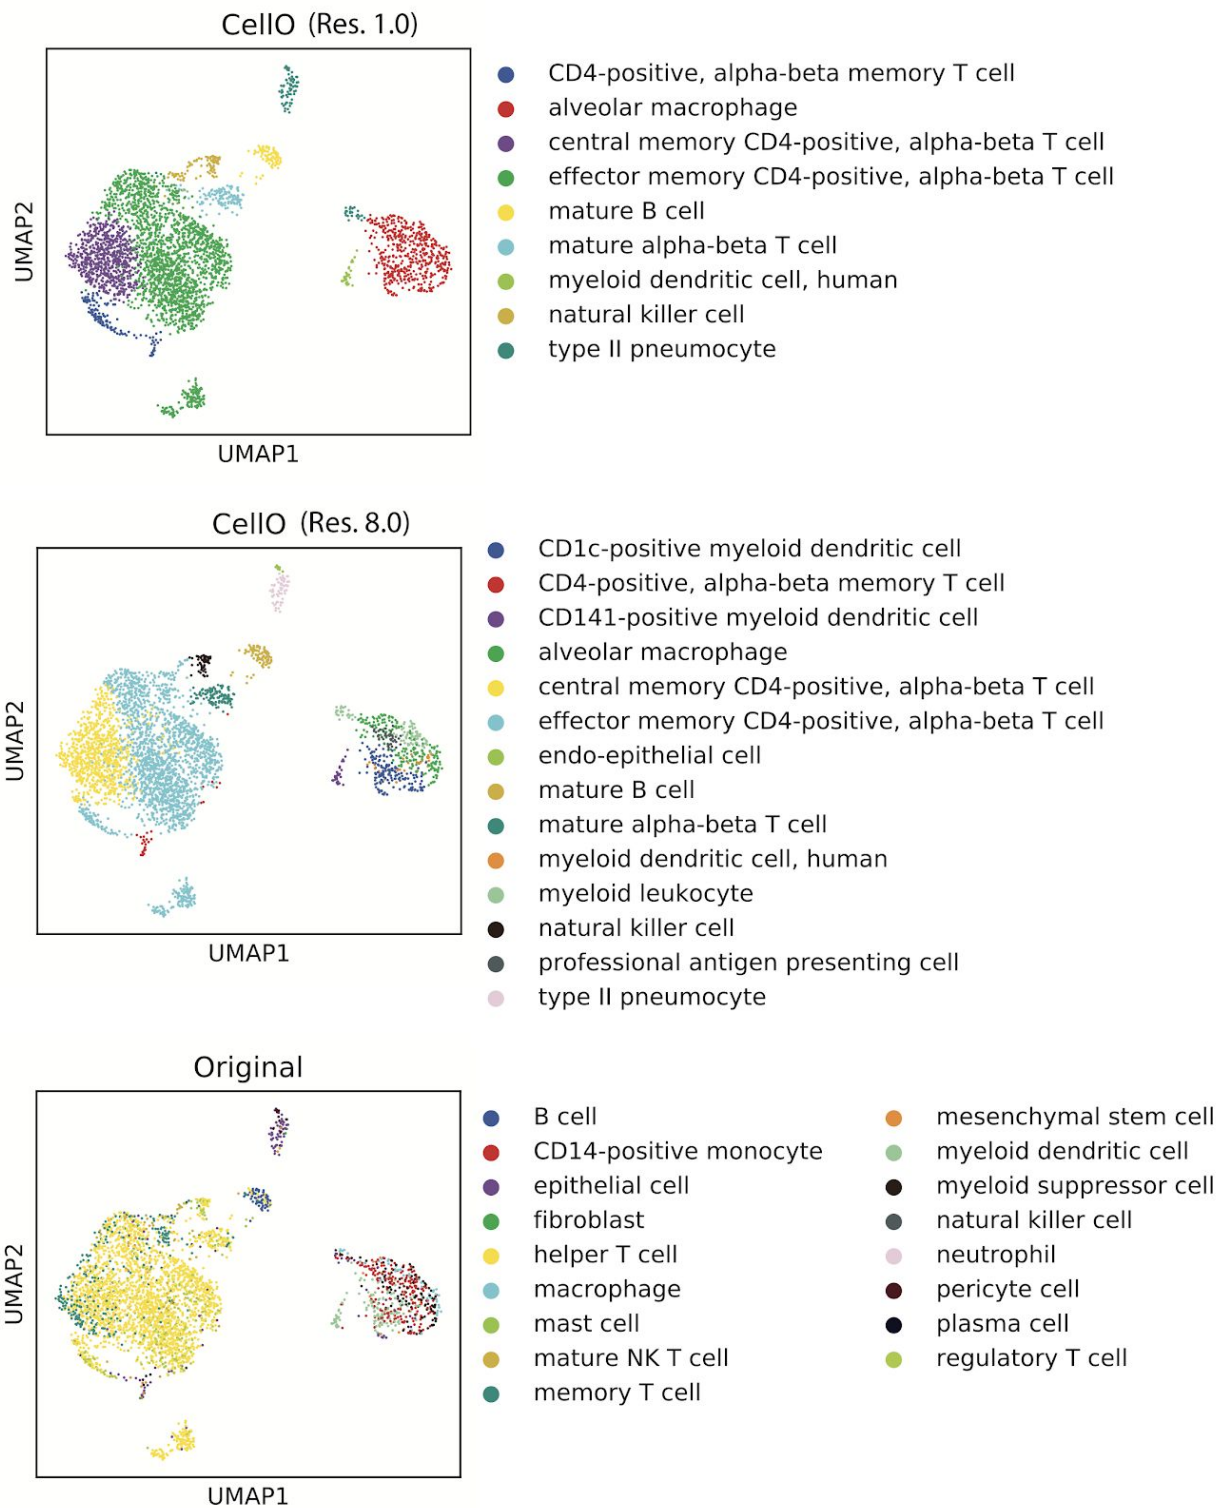
Supplemental Figure 5. **Examination of CellIO's Output for Tumor LX682 (Related to Figure 8).** UMAP plots of lung adenocarcinoma tumor LX682 from Laughney et al. colored by (top) CellIO's output using IR and a Leiden resolution parameter of 1.0 (middle) CellIO's output using IR and a Leiden resolution parameter of 8.0, and (bottom) the original cell type labels provided by the authors. We highlight three subpopulations of putative dendritic cells (top subpopulation), endothelial cells (middle subpopulation) and fibroblasts (bottom subpopulation). We verified these subpopulations using known marker genes for these cell types (top right).



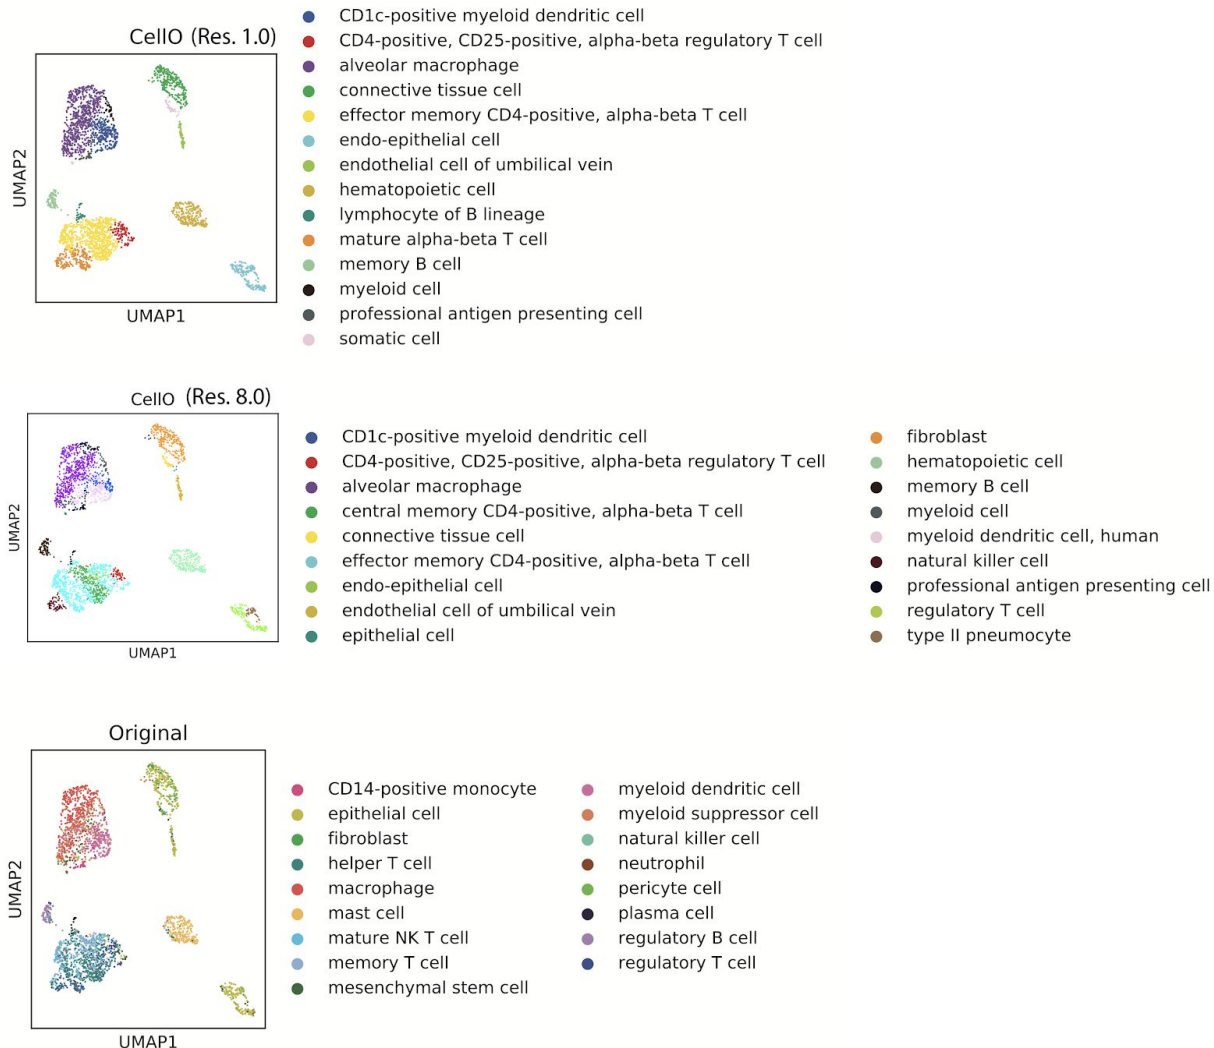
Supplemental Figure 6. **Examination of Cello's Output for Tumor LX679 (Related to Figure 8).** UMAP plots of lung adenocarcinoma tumor LX679 from Laughney et al. colored by (top) Cello's output using IR and a Leiden resolution parameter of 1.0 (middle) Cello's output using IR and a Leiden resolution parameter of 8.0, and (bottom) the original cell type labels provided by the authors. We highlight the plasmacytoid dendritic cell population and verify its identity using known marker genes for these cell types (top right).



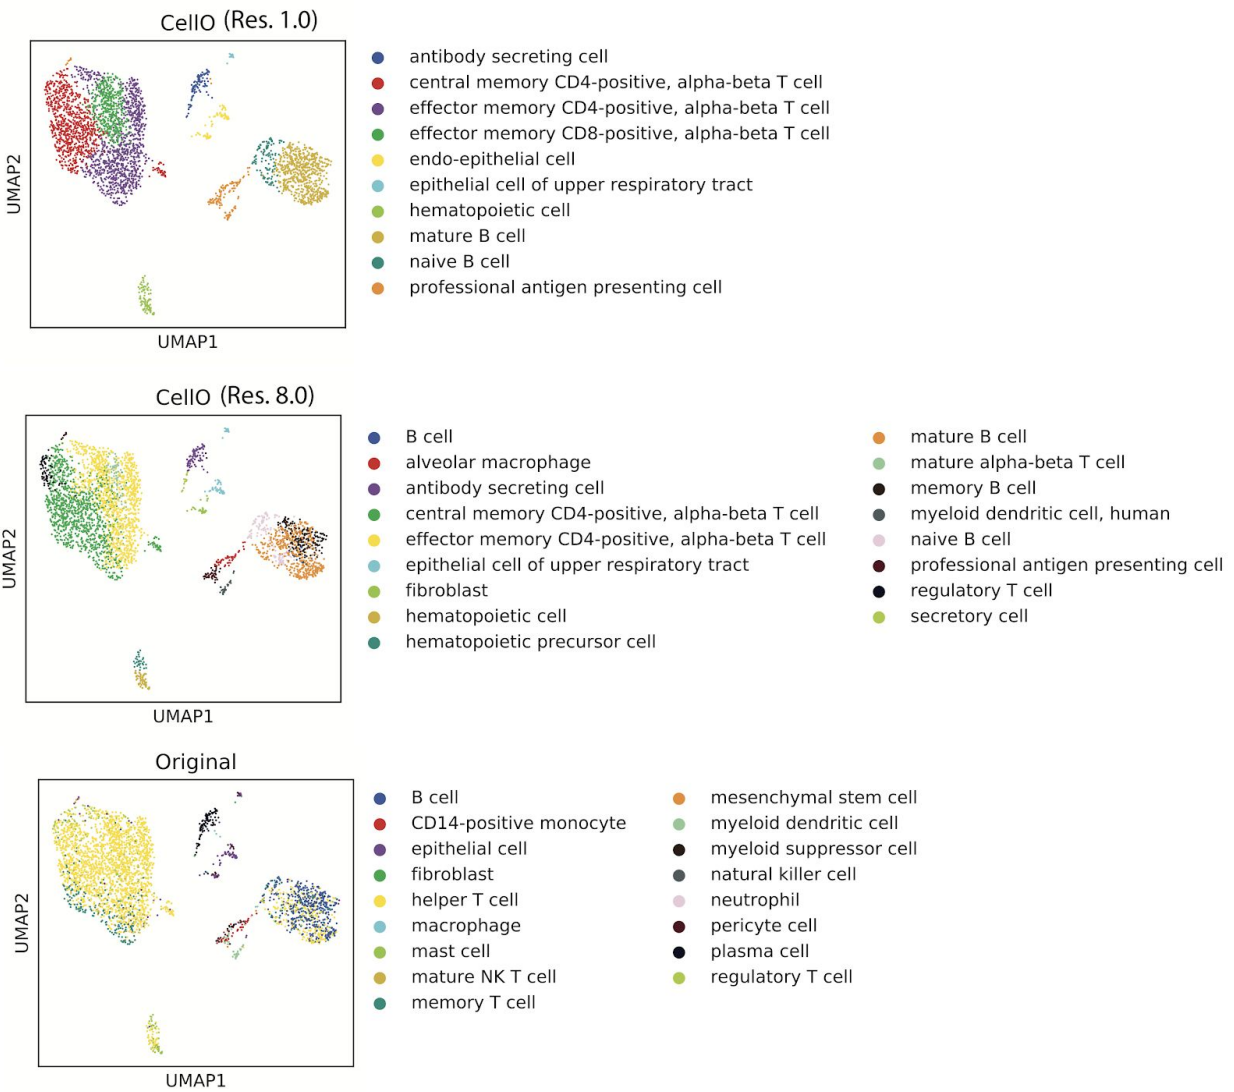
Supplemental Figure 7. **Examination of CellIO's Output for Tumor LX653 (Related to Figure 8)**. UMAP plots of lung adenocarcinoma tumor LX653 from Laughney et al. colored by (top) CellIO's output using IR and a Leiden resolution parameter of 1.0 (middle) CellIO's output using IR and a Leiden resolution parameter of 8.0, and (bottom) the original cell type labels provided by the authors.



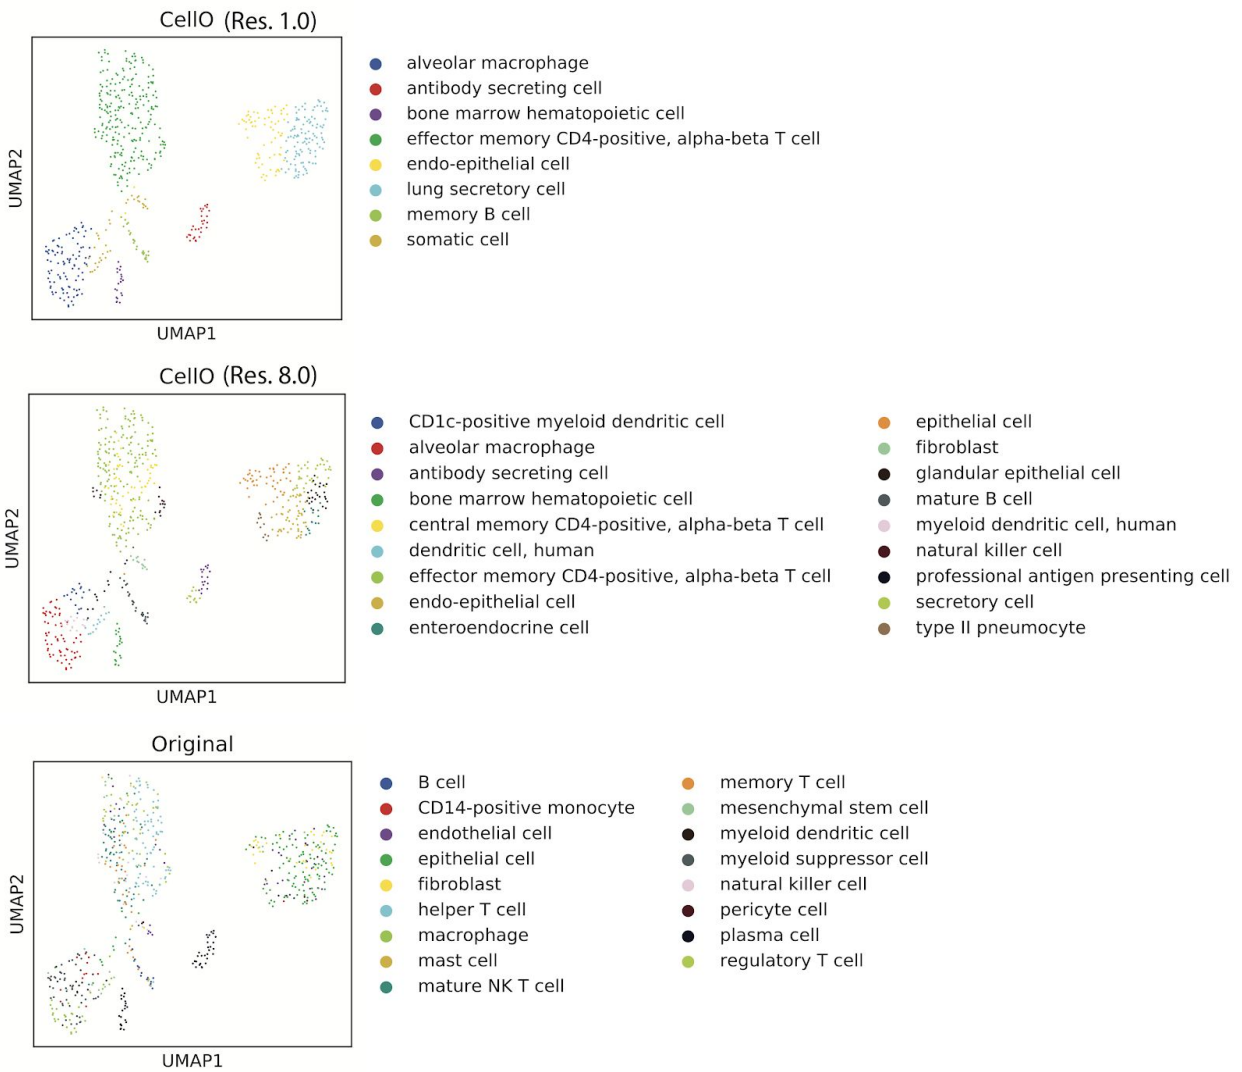
Supplemental Figure 8. **Examination of CellIO's Output for Tumor LX661 (Related to Figure 8).** UMAP plots of lung adenocarcinoma tumor LX661 from Laughney et al. colored by (top) CellIO's output using IR and a Leiden resolution parameter of 1.0 (middle) CellIO's output using IR and a Leiden resolution parameter of 8.0, and (bottom) the original cell type labels provided by the authors.



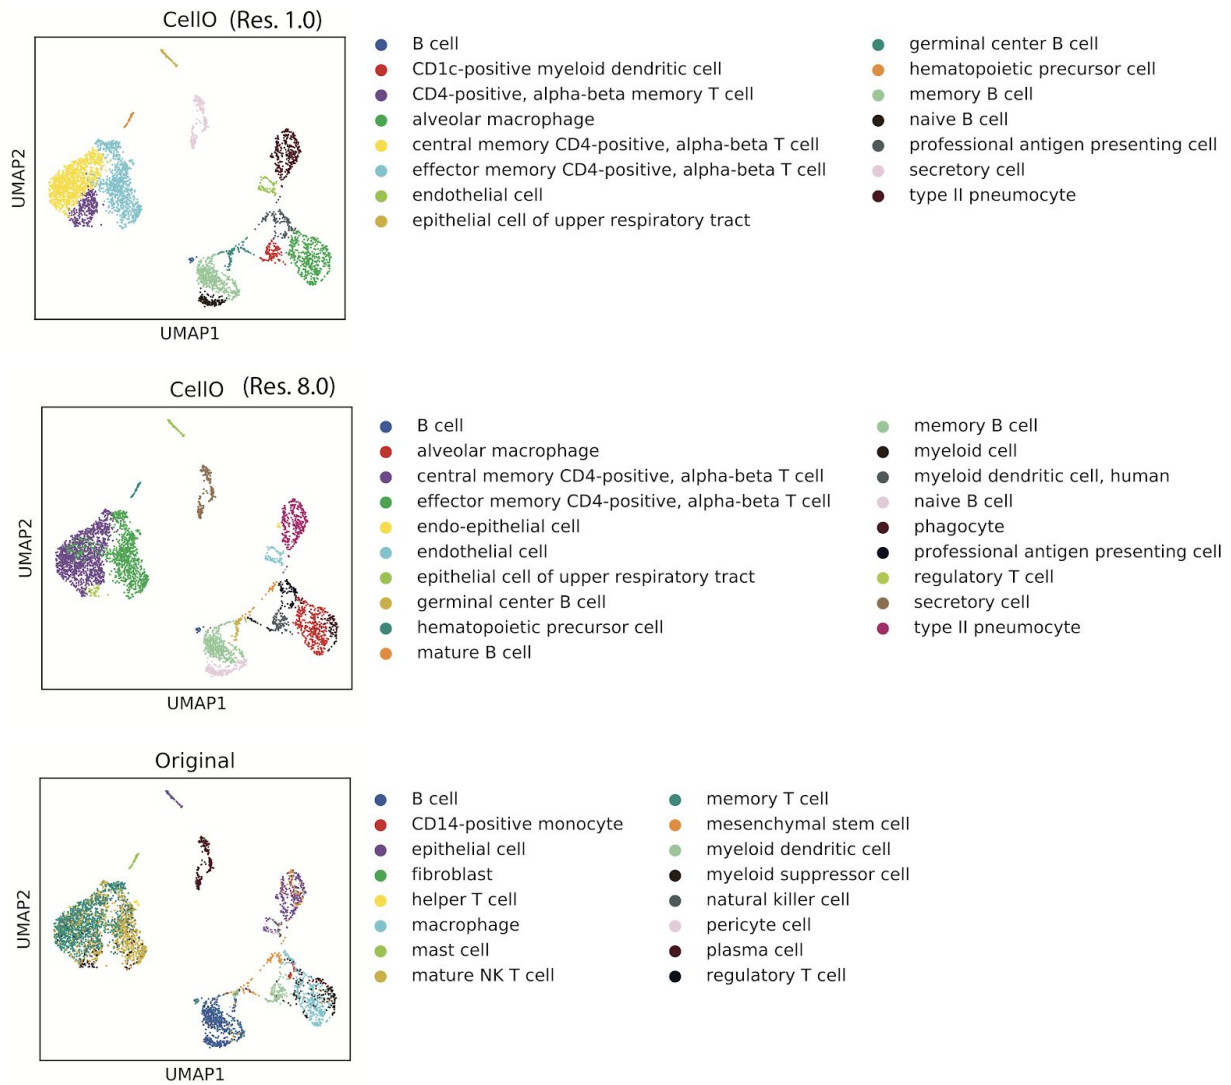
Supplemental Figure 9. **Examination of CellIO's Output for Tumor LX675 (Related to Figure 8).** UMAP plots of lung adenocarcinoma tumor LX675 from Laughney et al. colored by (top) CellIO's output using IR and a Leiden resolution parameter of 1.0 (middle) CellIO's output using IR and a Leiden resolution parameter of 8.0, and (bottom) the original cell type labels provided by the authors.



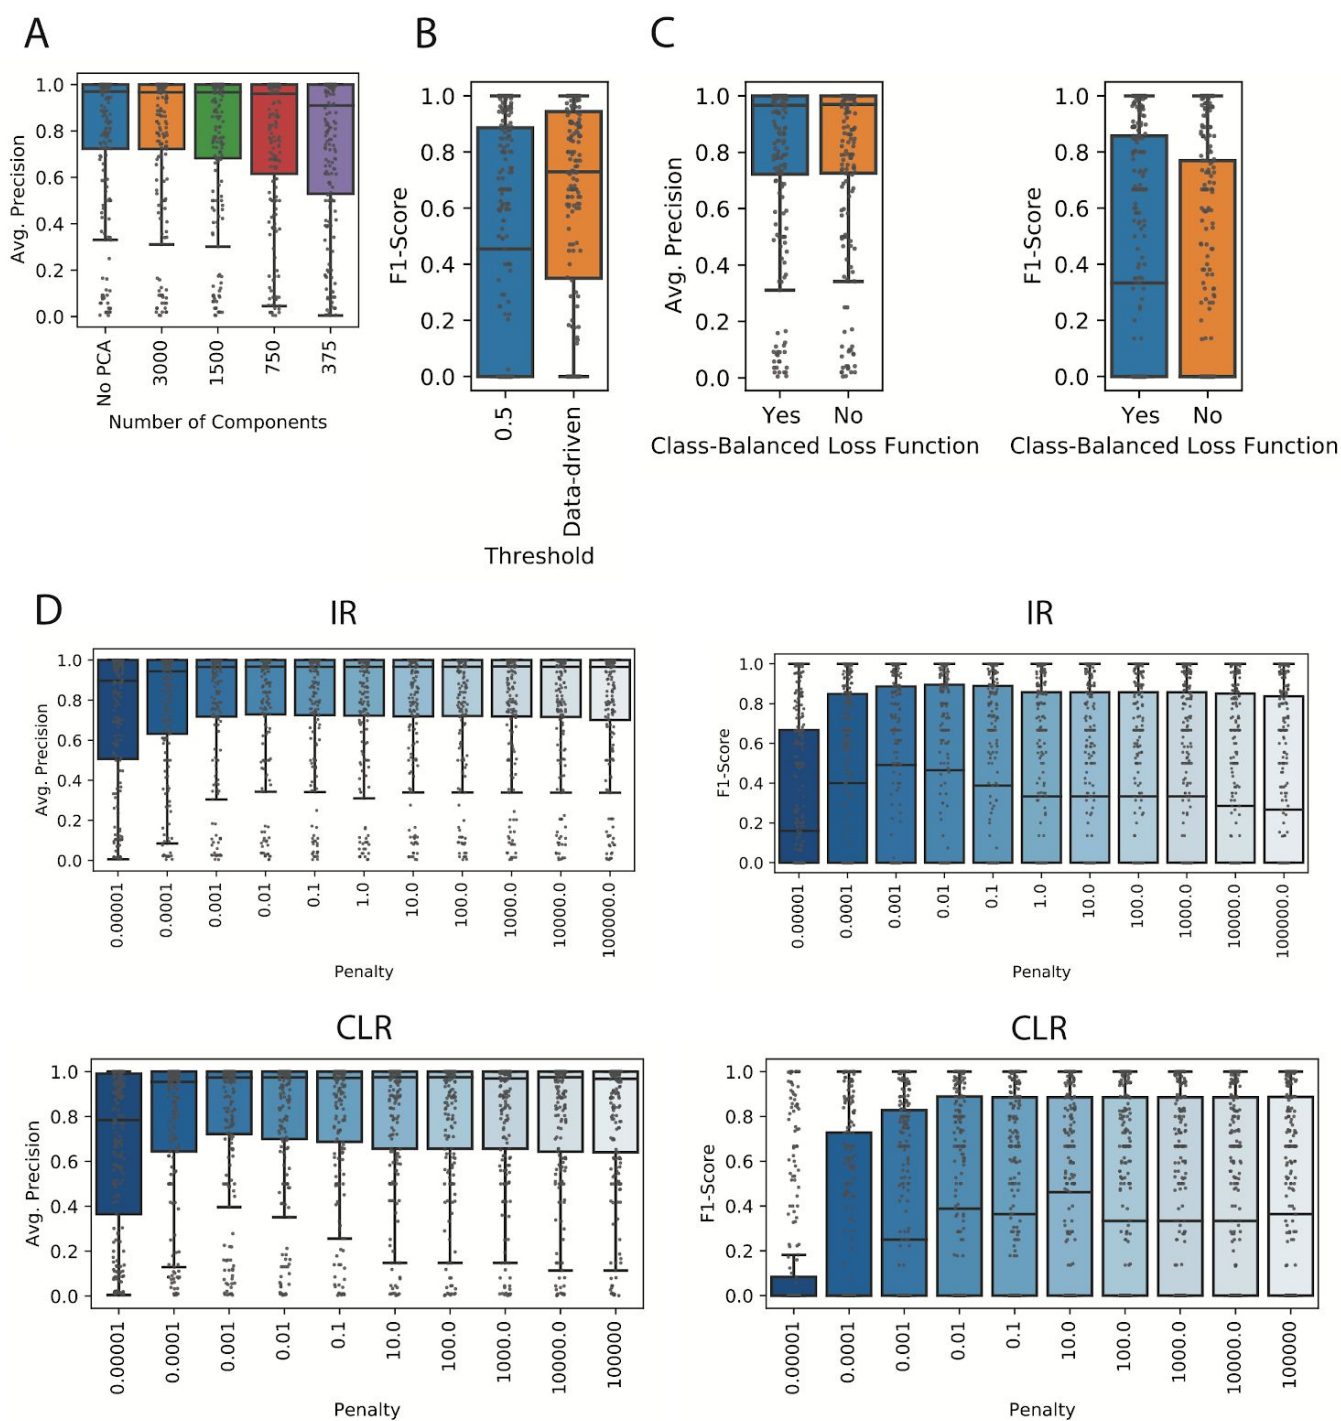
Supplemental Figure 10. **Examination of Cello's Output for Tumor LX676 (Related to Figure 8).** UMAP plots of lung adenocarcinoma tumor LX676 from Laughney et al. colored by (top) Cello's output using IR and a Leiden resolution parameter of 1.0 (middle) Cello's output using IR and a Leiden resolution parameter of 8.0, and (bottom) the original cell type labels provided by the authors.



Supplemental Figure 11. **Examination of Cello's Output for Tumor LX680 (Related to Figure 8).** UMAP plots of lung adenocarcinoma tumor LX680 from Laughney et al. colored by (top) Cello's output using IR and a Leiden resolution parameter of 1.0 (middle) Cello's output using IR and a Leiden resolution parameter of 8.0, and (bottom) the original cell type labels provided by the authors.

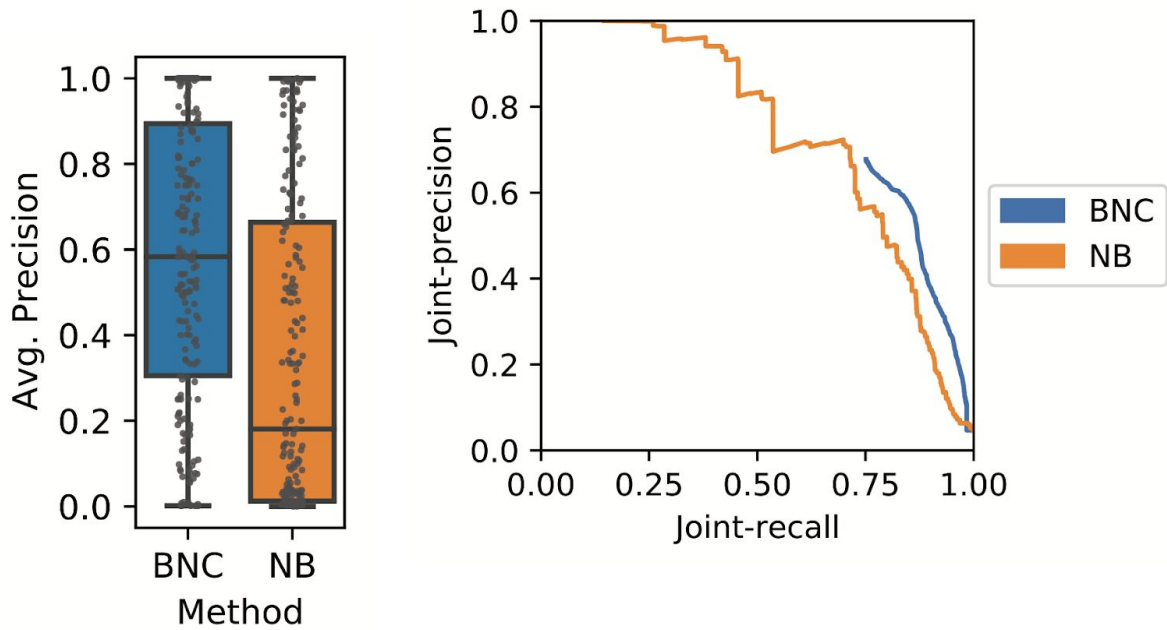


Supplemental Figure 12. **Examination of CellIO's Output for Tumor LX684 (Related to Figure 8).** UMAP plots of lung adenocarcinoma tumor LX684 from Laughney et al. colored by (top) CellIO's output using IR and a Leiden resolution parameter of 1.0 (middle) CellIO's output using IR and a Leiden resolution parameter of 8.0, and (bottom) the original cell type labels provided by the authors.

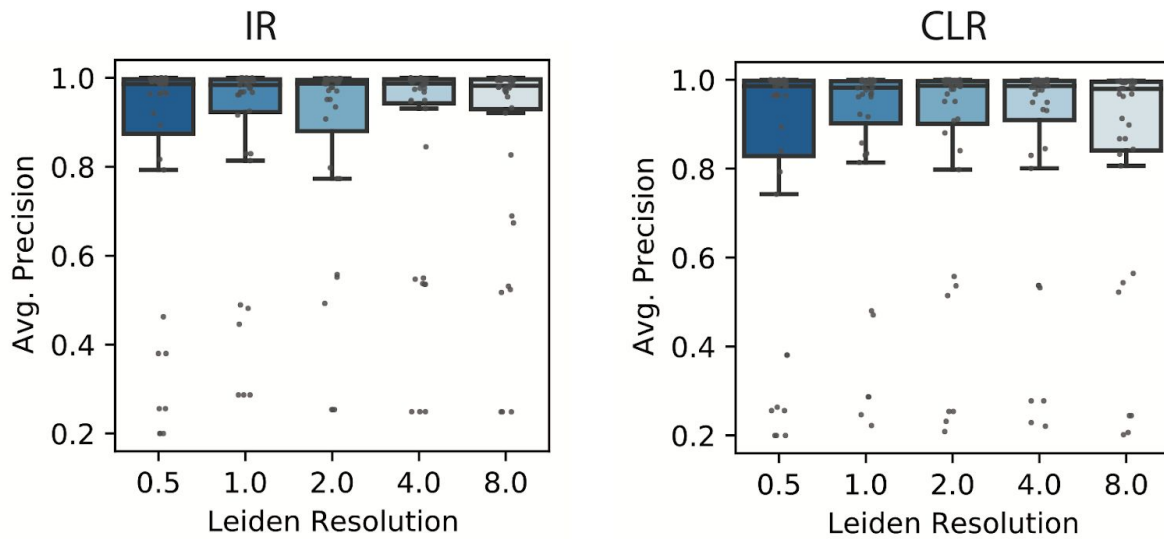


Supplemental Figure 13. **Parameter tuning (Related to Figure 2A).** (A) Average precision scores across cell types in the bulk RNA-seq validation set produced by the independent one-vs.-rest classifiers on data preprocessed with various numbers of principal components. (B) F1-scores produced by the independent classifiers when using a threshold of 0.5 for making a binary yes-no decision for all cell types versus a custom threshold for each cell type as empirically determined via a leave-study-out cross-validation experiment on the pre-training set of bulk RNA-seq samples. (C) Average precision (left) and F1-scores (right) across cell types in the bulk RNA-seq validation set produced by the independent one-vs.-rest classifiers trained either with or without the class-balanced loss-function. (D) Average


precision scores across cell types in the bulk RNA-seq validation set produced by the one-vs.-rest and CLR conditional classifiers when trained with various regularization strengths (i.e. penalties). These penalty values correspond to the inverse of the regularization strength and thus, smaller numbers indicate stronger regularization.



Supplemental Figure 14. **Comparing Bayesian network correction with naive Bayes (Related to Figure 3).** A comparison between BNC and a naive Bayes variant of BNC that considers each cell type independent of the graph-dependencies between cell types. We compared these approaches in terms of both the distribution of average precisions across cell types (left) and joint precision-recall curves constructed by considering each cell-cell type pair as an independent prediction.



Supplemental Figure 15. **Effect of clustering parameter on performance (Related to Figure 6).** We clustered the Zheng et al. (2017) PBMC dataset using various resolution parameters for the Leiden community-detection algorithm and then computed the average-precision across all cell types produced by IR (left) and CLR (right) on the mean expression profiles for the clusters generated under each parameter.



Supplemental Figure 16. **Effect of number of training studies on calibration quality (Related to Figure 2A).** We evaluated IR on the bulk RNA-seq validation set and for each cell type, we computed the F1-score (F1) using a decision-threshold of 0.5 and average precision (AP). We then plot the difference between F1 and AP against the logarithm of the number of studies in the bulk RNA-seq pre-training set that sequenced each cell type. A larger gap between F1 and AP (i.e. negative numbers with large magnitude) indicates poorer calibration. Error bands around the ordinary least-squares regression line indicate 0.95 confidence intervals estimated via bootstrapping.

Transparent Methods

Data curation

To create the training set of primary cell bulk RNA-seq samples, we first selected all samples labelled as a “primary cell” sample by the MetaSRA (v1.4). Thus, we followed the conservative definition for a primary cell sample by Bernstein et al. (2017), which requires that a sample has not undergone passaging beyond the first culture. We followed this selection with a manual curation of each sample’s technical variables by consulting sources of metadata that are not captured by the MetaSRA annotation process, such as fields in Gene Expression Omnibus (Clough and Barrett 2016) records and each study’s publication. In total, we annotated 27,097 samples (available to download at http://deweylab.biostat.wisc.edu/cell_type_classification/). We then removed all samples that were either incorrectly labelled as primary cell samples or had been experimentally treated. When found, we also corrected errors in the MetaSRA-provided Cell Ontology labels by both adding additional cell types that were missed by the MetaSRA as well as removing incorrect cell type labels. This curation effort resulted in the final set of 4,293 samples.

Data processing

We quantified the gene expression of all RNA-seq samples from the SRA (both bulk and non-droplet-based single-cell samples) with kallisto (v0.43.1) (Bray et al. 2016) using human genome release GRCh38 with GENCODE annotation version 27. We chose kallisto for gene expression quantification in order to prioritize processing speed on this large dataset, figuring that any small loss in accuracy (at the gene level) relative to a less approximate, but slower approach would not be significant for the cell type classification task. This produced estimated counts for 200,401 isoform-level genomic features. We summed the transcripts per million (TPM) values by gene to produce TPM’s for 58,243 gene-level features. The curated metadata and associated quantified samples are available to download at http://deweylab.biostat.wisc.edu/cell_type_classification.

Notation

In the following descriptions of the methods used in this work, we let $\mathbf{x} \in \mathbb{R}^G$ denote a gene expression profile, where G is the number of considered genes and \mathbf{x} is measured in units of $\log(\text{TPM}+1)$ where TPM are transcripts per million. We let n denote the number of samples, m denote the number of considered cell types, $y_{i,j} \in \{0, 1\}$ denote the cell type assignment for cell type $j \in [m]$ and sample $i \in [n]$.

Binary classification with logistic regression

We use L2-penalized logistic regression for each binary classifier as implemented by scikit-learn (Pedregosa et al. 2011) using the LIBLINEAR solver (Fan et al. 2008). To speed up the training of each binary classifier, we preprocessed the bulk RNA-seq training data using principal components analysis (PCA) as implemented in scikit-learn. Specifically for each cell type j , each classifier is trained by minimizing the following loss-function:

$$\min_{\beta_j, c_j} \frac{1}{2} \|\beta_j\|_2^2 + C \sum_{i=1}^n w_{i,j} \left[-y_{i,j} \log(h_{\beta_j, c_j}(\mathbf{x}_i)) - (1 - y_{i,j}) \log(1 - h_{\beta_j, c_j}(\mathbf{x}_i)) \right]$$

where

$$h_{\beta_j, c_j}(\mathbf{x}_i) := [1 + \exp((\mathbf{U}\mathbf{x}_i)^T \beta_j + c_j)]^{-1}$$

is the logistic function, $\beta_j \in \mathbb{R}^k$ are the model-coefficients for cell type j , c_j is the intercept for cell type j , C controls the strength of the regularization, n is the number of training samples, $w_{i,j}$ is a per-sample weight for handling class-imbalance in cell type j ’s model, and $\mathbf{U} \in \mathbb{R}^{k \times m}$ is the PCA loadings matrix.

We also note that most training sets are highly unbalanced. We found that the models were better calibrated when the loss-function weighted each training sample such that the positive and negative samples contribute equally to the loss function (via the $w_{i,j}$ weights above) as evidenced by the improved F1-scores when using a threshold of 0.5 for making discrete yes-no cell type decisions (**Figure S13C**). We implemented this class-balancing by setting the *class_weight* parameter to 'balanced' in scikit-learn's *LogisticRegression* class constructor.

Isotonic regression correction

We train a binary classifier for each cell type $j \in [m]$ to model $p(y_j | \mathbf{x})$ using logistic regression and a one-versus-rest training strategy. As proposed by Obozinski *et al.* (2008), these probabilities are then reconciled with the ontology graph using isotonic regression. Specifically, we output the set of probabilities:

$$p_1, \dots, p_m := \arg \min_{p'_1, \dots, p'_m} \sum_{i=1}^m (p'_j - \hat{p}_j)^2$$

subject to

$$\forall j \in [m], \forall k \in \text{Par}(j), p_j \leq p_k \text{ and } \forall j \in [m], 0 \leq p_j \leq 1$$

where $\forall j \in [m], \hat{p}_j := p(y_j = 1 | \mathbf{x})$ as output by each classifier and $\text{Par}(j)$ is the set of parent cell types for cell type j . Cello uses the quadprog Python package (<https://pypi.org/project/quadprog/>) for solving this quadratic optimization problem.

Bayesian network correction

A binary classifier is trained for each cell type and a one-versus-rest training strategy. The classifier outputs are then reconciled with the ontology graph using a Bayesian network as proposed by Lee *et al.* (2013). The true assignments for each cell type within a given sample, denoted y_1, \dots, y_m , are modelled as latent random variables, and the classifier outputs, denoted $f_1(\mathbf{x}), \dots, f_m(\mathbf{x})$ (signed distances to each decision boundary), are modelled as observed random variables in a Bayesian network. The final output probability for cell type j is then the marginal probability

$$p(y_j = 1 | f_1(\mathbf{x}), \dots, f_m(\mathbf{x}))$$

More specifically, for a given cell type j , we model the conditional distribution of the classifier's output $f_j(\mathbf{x})$ (distance to the decision boundary) conditional on its true label y_j as a discrete random variable constructed as follows. We partition the training data for cell type j into two folds ensuring that no study is split between folds while attempting to keep the sizes of the two folds as similar as possible. We then train on one fold and compute the classifier scores from the second fold (for each of the two folds).

Using all of these scores, we then compute a histogram where the bin sizes are determined using a second 2-fold cross-validation scheme. Specifically, we then test a number of bin sizes by first estimating a histogram density function using data in one fold and then computing the likelihood of the data in the second fold (performing this procedure for both folds). The histogram density function is given by

$$f(x) := \frac{1}{nh} \text{Count}(x)$$

where n is the total number of data points, h is the width of each bin, and $\text{Count}(x)$ is the number of data points sharing the same bin as x . We choose a bin size that maximizes the mean of the two data log likelihoods computed on each fold.

As described by Lee et al. (2013), the true cell type assignments $p(y_1, \dots, y_m)$ factor according to the ontology graph:

$$p(y_1, \dots, y_m) := \prod_{j=1}^m p(y_j \mid \text{Children}_j)$$

where

$$\text{Children}_j := \{y_k \mid k \text{ is a child of } j\}$$

are the assignments to the children cell types of cell type j in the ontology. These conditional distributions $p(y_j \mid \text{Children}_j)$ enforce consistency with the ontology by defining

$$p(y_j = 1 \mid \text{Children}_j) := 1.0$$

if $1 \in \text{Children}_j$. Otherwise,

$$p(y_j = 1 \mid \text{Children}_j) := \text{prior}_j$$

The prior_j values are computed from counts in the training data. Specifically, the prior for each leaf-node cell type c is simply the fraction of samples in the data set labelled as c . For each internal node c , the prior is computed as the fraction of all samples labelled as c , but not labelled as any child of c . A pseudocount of one was used in the calculation for all priors. Due to the size of the ontology, we perform approximate inference using Gibbs sampling rather than exact inference using the Lauritzen algorithm as was performed by Lee et al.

To test the impact of the graph-structured prior on this algorithm, we also tested a naive Bayes variant in which each cell type is predicted independently (i.e., without the graph-structured prior). We found a significant improvement in performance for the graph-structured BNC algorithm over the naive Bayes algorithm indicating that the graph structured prior is an important component of this algorithm (**Figure S14**).

True Path Rule

We train a binary classifier for each cell type $j \in [m]$ to model $p(y_j \mid \mathbf{x})$ using logistic regression and a one-versus-rest training strategy. As proposed by Notaro et al. (2017), this method involves two passes across the ontology: on a bottom-up pass, each cell type's output probability is averaged with the outputs of all child cell types classifiers for which the classifier makes a positive prediction according to a predefined threshold. More specifically, each cell type j 's output probability is set to

$$\bar{p}_j := \frac{1}{|C_j| + 1} \left(\hat{p}_j + \sum_{k \in C_j} \bar{p}_k \right)$$

where $\hat{p}_j := p(y_j = 1 \mid \mathbf{x})$ according to the classifier and

$$C_j := \{k \in \text{Children}(j) : \hat{p}_k > t\}$$

is the set of children of cell type j for which the classifier output a positive prediction according to a predefined threshold t . We used a threshold of $t = 0.5$. This bottom-up pass allows for sharing of information across the classifiers. In the top-down pass of the ontology, the output probabilities are set to ensure consistency with the ontology. This procedure works as follows: each node in the ontology is visited according to the topologically sorted order of nodes and for a given visited cell type j , its final probability is set to

$$p_j := \begin{cases} \bar{p}_j & \text{if } j \text{ is a root node} \\ \min_{k \in \text{Par}(j)} p_k & \text{if } \min_{k \in \text{Par}(j)} p_k < \bar{p}_j \\ \bar{p}_j & \text{otherwise} \end{cases}$$

where $\text{Par}(j)$ are the parent nodes of node j in the DAG.

Cascaded logistic regression

Classification is made in a top-down fashion starting from the root of the ontology downward as proposed by Obozinski et al. (2008). This is accomplished by training a logistic regression, binary classifier (although any binary classifier that outputs a prediction probability can be used) for each cell type $j \in [m]$ to model the distribution

$$q_j := p(y_j = 1 \mid \pi_j = 1, \mathbf{x})$$

where $\pi_j \in \{0, 1\}$ indicates whether the sample belongs to all of the parents of j in the ontology. In order to model these distributions, each cell type's negative training examples consist of those samples that are labeled with all parent cell types, but not the target cell type. Given these learned distributions, the probability that \mathbf{x} originates from cell type j is computed via

$$p(y_j = 1 \mid \mathbf{x}) = q_j \prod_{k \in A_j} q_k$$

where A_j denotes the ancestors of cell type j in the ontology's DAG.

One-nearest neighbors

Given a query gene expression profile \mathbf{x} , we return all cell type labels belonging to the training set expression profile

$$\arg \min_{\mathbf{x}' \in \mathbf{X}} 1 - \text{Corr}(\mathbf{x}, \mathbf{x}')$$

where $\text{Corr}(\mathbf{x}, \mathbf{x}')$ is the Pearson correlation of the expression values in \mathbf{x} and \mathbf{x}' as implemented in Python's SciPy package (Virtanen et al. 2020).

Quantification and statistical analysis

Partitioning bulk RNA-seq data into training and test sets

In order to find the optimal parameters and configurations for CellO, we partitioned the bulk RNA-seq training dataset into a pre-training set and validation set (**Figure 2A**). When creating this partition, we sought to satisfy a number of criteria that would enable unbiased estimation of performance across cell types. First, we required that no study be split between the pre-training and validation sets in order to ensure that a model is never tested on data from a study on which it was trained. This mitigates the possibility that the algorithm will provide an overly optimistic estimate of the generalization error when run on the validation set. Second, we sought an approximately 80/20 split of the data between the pre-training and validation sets. Third and finally, we sought for all cell types to be represented in both the pre-training and validation sets. We framed this partitioning task as an optimization problem where our four criteria were encoded in an objective function. Minimizing this objective function entails creating a partition that most closely meets the aforementioned four criteria. To optimize the objective function, we performed a simple hill-climbing procedure where we moved a study's data from the validation set to the pre-training set if such a move resulted in a partition that improved the objective function.

Parameter tuning

To choose the optimal number of principal components, we evaluated various numbers of principal components by training the collection of one-vs.-rest classifiers on the bulk RNA-seq pre-training set and evaluating on the bulk RNA-seq validation set. We found that using 3,000 principal components performed equally well to using the raw gene expression values as features. Further, we found performance degraded as the number of principal components decreased (**Figure S13A**). We then performed a parameter sweep of penalty-weight parameters (i.e. regularization-strength parameters) for

IR and CLR and chose the values that maximized the median average precisions for each method (**Figure S13D**).

Model calibration

We found that for many cell types, there were relatively few studies in the training set for which samples of these cell types originated and thus, we found that the models tended to overfit to the training set leading to poorly calibrated models. This was evidenced by the fact that the average-precision scores were high across most cell types; however, the F1-scores were much lower when using a default threshold of 0.5 for making binary cell type decisions (**Figure S13D**). The high average precision scores indicate that the positive and negative examples of each cell type are well separated in the ranking of samples when ranked according to the classifier output probability for the given cell type. However, the poor F1-scores indicate that the default threshold of 0.5 is non-optimal for separation.

To address this issue, we used a data-driven approach to find empirical thresholds according to a leave-study-out cross-validation experiment. Specifically, we performed leave-study-out cross-validation on the pre-training set of bulk RNA-seq samples and used these results to empirically choose a threshold for each cell type such that if a given threshold less than 0.5 led to a higher F1-score than the default threshold of 0.5, we select this empirical threshold. When applying these data-driven thresholds to the validation set, we observed a significant increase in the F1-scores across cell types under the independent one-vs.-rest classifiers approach (**Figure S13B**). Thus, when training the final CellIO models (both CLR and IR), we performed this same cross-validation procedure on the entire bulk RNA-seq training set to select data-driven thresholds.

We note that even with these selected thresholds, CellIO often output more than one specific cell type for a given sample (e.g. both natural killer cell and T cell), which may confuse a user of the tool. This phenomenon is a consequence of both mis-calibrated models (despite the data-driven thresholding procedure) and because CellIO performs multi-label hierarchical classification. To address this issue, whenever CellIO outputs more than one specific cell type, we select only the cell type with highest output probability along with all ancestor cell types. All F1-scores reported in the main text followed from this correction procedure and thus, are an apt measure of the practical utility of CellIO on real data (**Figure 2B**).

Model interpretation

We note that although the logistic regression model's coefficients $\beta_j \in \mathbb{R}^k$ (where j denotes the index for a given cell type) weight the k principal components rather than genes, each gene's contribution to the model's decision can be recovered by

$$\mathbf{w}_j := \mathbf{U}\beta_j$$

where $\mathbf{U} \in \mathbb{R}^{m \times k}$ is the PCA loading matrix and $\mathbf{W}_{j,g}$ describes the contribution of gene g to the model's decision. The full set of these vectors over all cell types can be explored within the CellIO Viewer web application.

Evaluation metrics

For a given input dataset, CellIO generates two sets of outputs: binary yes-no decisions for each cell type assignment as well as probability scores that quantify how likely each cell should be assigned to any given cell type. Both of these outputs can be represented as matrices. More specifically, given an input dataset, $\mathbf{X} \in \mathbb{R}^{n \times G}$ where n is the number of cells and G is the number of genes, the binary yes-no decisions can be represented as a matrix $B \in \{0, 1\}^{n \times m}$ where m is the number of cell types and $B_{i,j} = 1$ if the classifier predicts cell i to be of cell type j . The probability scores can also be represented as a matrix $S \in [0, 1]^{n \times m}$ where $S_{i,j}$ denotes the classifier's confidence that cell type i should be labelled as cell type j . Finally, the true cell type assignments can also be represented as a

matrix $T \in \{0, 1\}^{n \times m}$ where $T_{i,j} = 1$ if cell i is truly of cell type j .

We first define metrics for comparing B to T . For cell type j , we define the number of true positives (TP), false positives (FP), and false negatives (FN) as:

$$\begin{aligned} TP &:= |\{i \mid T_{i,j} = 1 \wedge B_{i,j} = 1\}| \\ FP &:= |\{i \mid T_{i,j} = 0 \wedge B_{i,j} = 1\}| \\ FN &:= |\{i \mid T_{i,j} = 1 \wedge B_{i,j} = 0\}| \end{aligned}$$

We then evaluate the classifier's performance on cell type j using precision, recall, and F1-score, which are defined as

$$\begin{aligned} \text{precision} &:= \frac{TP}{TP + FP} \\ \text{recall} &:= \frac{TP}{TP + FN} \\ \text{F1-score} &:= \frac{2 * \text{precision} * \text{recall}}{\text{precision} + \text{recall}} \end{aligned}$$

We note that for a given cell, the ground-truth assignments for cell types that are more-specific than the cell's most-specific ground-truth cell type are ambiguous since these samples could, in theory, truly be of a more specific cell type than they are labelled with (e.g., a cell labelled as a T cell could be a CD8+ T cell even if it isn't annotated as such). Thus, when computing the aforementioned metrics for a given cell type within these single cell datasets, we exclude those cells that are labelled most-specifically as an ancestor of the cell type. For example, for metrics calculated for CD8+ T cells, we would exclude from the calculations those cells that are most-specifically labelled as T cells.

References

- Clough, Emily, and Tanya Barrett. 2016. "The Gene Expression Omnibus Database." *Methods in Molecular Biology* 1418: 93–110.
- Fan, Rong-En, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. 2008. "LIBLINEAR: A Library for Large Linear Classification, Journal of Machine Learning Research." *Journal of Machine Learning Research: JMLR* 9: 1871–74.
- Pedregosa, Fabian, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, et al. 2011. "Scikit-Learn: Machine Learning in Python." *Journal of Machine Learning Research: JMLR* 12: 2825–30.
- Virtanen, Pauli, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, et al. 2020. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python." *Nature Methods* 17 (3): 261–72.