CelLink: Integrate single-cell multi-omics data with few linked features and imbalanced cell populations

Xin Luo[1], Yuanhao Huang[1], Yicheng Tao[1], Fan Feng[2], Alexander Hopkirk[2], Thomas S.R. Bate[2], Diane C. Saunders[2], Peter Orchard[1], Catherine Robertson[1], Shristi Shrestha[3], Jean-Philippe Cartailler[3], Stephen C. J. Parker[1], Marcela Brissova[2], Jie Liu[1]

[1]Department of Computational Medicine and Bioinformatics, University of Michigan, Ann Arbor, MI, USA
[2]Division of Diabetes, Endocrinology and Metabolism, Department of Medicine, Vanderbilt University Medical Center, Nashville, TN, USA
[3]Center for Stem Cell Biology, Vanderbilt University, Nashville, Tennessee, USA

Abstract

Single-cell multi-omics technologies capture complementary molecular layers, enabling a comprehensive view of cellular states and functions. However, integrating these data types poses significant challenges when their features are weakly linked and cell population sizes are imbalanced. Currently, no method efficiently addresses these two issues simultaneously. Therefore, we developed CelLink, a novel single-cell multi-omics data integration method designed to overcome these challenges. CelLink normalizes and smooths feature profiles to align scales across datasets and integrates them through a multi-phase pipeline that iteratively employs the optimal transport algorithm. It dynamically refines cell-cell correspondences, identifying and excluding cells that cannot be reliably matched, thus avoiding performance degradation caused by erroneous imputations. This approach effectively adapts to weak feature linkage and imbalanced cell populations between datasets. Benchmarking CelLink on scRNA-seq and spatial proteomics datasets, as well as paired CITE-seq data, demonstrates its superior performance across various evaluation metrics, including data mixing, cell manifold structure preservation, and feature imputation accuracy. Compared to state-of-the-art methods, CelLink significantly outperforms others in imbalanced cell populations while consistently achieving better performance for balanced datasets. Moreover, CelLink uniquely enables cell subtype annotation, correction of mislabelled cells, and spatial transcriptomic analyses by imputing transcriptomic profiles for spatial proteomics data. CelLink sets a new milestone for multi-omics data integration. Its great ability to impute paired single-cell multi-omics profiles positions it as a pivotal tool for building single-cell multi-modal foundation models and advancing spatial cellular biology.

Introduction

The integration of single-cell multi-omics data provides complementary information from various molecular layers for individual cells. This approach enhances our understanding of cellular functions, states, and interactions by capturing diverse biological signals such as epigenomic, transcriptomic, and proteomic data, each offering unique insights into cellular processes. Despite its promise, the integration of multi-omics data still faces two challenges. The first challenge is handling proteomics data with its rapid expansion, which exhibits weak feature overlap and distinct systematic biases compared to other data. For example, the relationship between mRNA and protein expression levels can show a weak correlation due to multiple factors, including variations in their respective degradation rates and the influence of post-transcriptional regulatory mechanisms [11]. The second challenge is integrating multi-omics data with significant discrepancies in cell population sizes across different datasets. For example, for the HPAP024 donor from the Human Pancreas Analysis Program [16], the scRNA-seq data contains 54 beta cells comprising 4.26% of the total while the CODEX data includes 35,140 beta cells representing 14.49% of its total.

Traditionally, methods such as Seurat V3 [28] and LIGER [29] require significant overlap in features between datasets, such as scRNA-seq with scATAC-seq or with spatial transcriptomics. These methods are not suitable when the datasets have weak feature correspondence. To address this, MMD-MA [22] aligns the cell-cell similarity (manifold) structures between the two data types after projecting them into the kernel spaces. SCOT v2 [9] aligns unpaired multi-omics data by comparing the distance between samples rather than comparing the samples themselves across modalities via Gromov-Wasserstein optimal transport. MaxFuse [8] progressively updates the cell-cell matching between the two data types through fuzzy smoothing. Nonetheless, it is common for single-cell multi-omics datasets from the same tissue to have significantly different cell compositions. None of these methods adapt to datasets with imbalanced cell population sizes.

To overcome these challenges, we developed CelLink, a novel method to integrate single-cell multi-omics data. CelLink adeptly handles weak feature linkage through normalization and smoothing techniques, effectively neutralizing systematic biases introduced by different biotechnologies. Furthermore, it employs a dynamic cell-cell correspondence framework, which is iteratively updated to manage the imbalanced cell populations in datasets, resulting in excellent performance in case of discrepant cell populations. Meanwhile, CelLink also demonstrates a slight superiority in scenarios with balanced populations. The benchmarking results show that CelLink consistently outperforms other state-of-the-art integration methods across a broad spectrum of datasets under various evaluation metrics. Moreover, CelLink is the only method that automatically identifies cells that cannot be aligned between the multi-omics data. This avoids the issue where including such cells would degrade the alignment performance, as the imputed feature values for these cells would inherently be problematic. Additionally, CelLink enables various novel downstream analyses after mutually imputing the feature profiles between the single-cell transcriptomics and spatial proteomics data, including cell subtyping, exploring spatial gene expression patterns, and analyzing spatial cell-cell communications.

Results

Cell-cell alignment between modalities via iterative optimal transport

The inputs of CelLink are the expression matrices of the two modalities X and Y (Fig. 1A). First, to reduce the discrepancies of the two data modalities, we normalize them on a common numeric scale. Next, to mitigate the inherent biases arising from different biotechnological platforms, we perform cell-wise smoothing from k-NN graphs based on the feature profile for each modality. Then, we define "linked features", which refer to corresponding entities in X and Y, such as the GCG coding gene in scRNA-seq and the GCG protein marker in CODEX data. We denote the profiles of these linked features in X and Y as X_l and Y_l, respectively. After preprocessing the datasets, we input the normalized and smoothed linked features $\tilde{X}_l$ and $\tilde{Y}_l$ into the integration pipeline of CelLink. The integration pipeline offers two choices. If cell type or cell state is provided, it starts with Balanced Optimal Transport (BOT), continues with iterative unbalanced Optimal Transport (UOT) until convergence, and finally transfers the feature profile from one modality to another. Otherwise, it just performs a one-time interoperable UOT, followed by feature transferring.

In the first phase, BOT is performed to measure the cellular correspondence between modalities (Fig. 1B). To align cells fairly, we determine the mass of all cells in OT as 1.0. The cost matrix of BOT, defined as D_{xy}, is constructed based on the pairwise dissimilarity between cells in X and Y. See the Methods section for more details. Given that BOT requires the total mass of cells equal across modalities, we split each data into multiple partitions, ensuring that the number of cells in each pair of partitions is close. The outcome of BOT is a transport map, which establishes the initial cell-cell correspondence between modalities in the iterative procedure, denoted as U_0. When the cell-type or cell-state populations of the two data are significantly imbalanced, some cells will be mapped to erroneous cell types or cell states. This misalignment occurs because the total mass attributed to a particular cell type or cell state does not match between the two datasets. To solve this problem, we filter out all matched cells and retain all unmatched cells whose predicted cell types are erroneous, followed by correcting their alignments in the following phase. The predicted cell type is the one with the highest weight from the other

modality calculated via U_0.

In the second phase, we conduct iterative UOT for the unmatched cells in one modality to all cells in another (Fig. 1C). UOT does not require the total mass of cells to be equal between the two datasets, which largely addresses the challenge of disparate cell population sizes. However, the optimization of the optimal transport algorithm faces a trade-off between efficiency and global convergence issues. This results in incorrect matches for certain cells, whose feature distributions overlap with other similar cell types due to the weak feature linkage between datasets. To resolve this issue, we iteratively remove the correctly matched cells in each round and realign the remaining unmatched cells. This strategy gradually reduces the matching space for the remaining cells thereby improving their matching. The iteration continues until the predicted cell types for unmatched cells stabilize, indicating the convergence of UOT. This stability suggests that the remaining cells cannot be aligned with cells in the other modality. The cellular correspondence matrices in each iteration are assumed as $U_1, U_2, ..., U_t$, where t denotes the iteration times. The remaining unmatched cells will be excluded in the subsequent step.

In the third phase, we reciprocally transfer the feature profile from one modality to another. The outputs are the imputed feature profiles of all modalities along with their cell-cell correspondence matrices (Fig. 1D). The series of cellular correspondence matrices $U_0, U_1, ..., U_t$ is consolidated into a unified matrix U based on cell IDs, wherein the cell correspondence vector for each initially unmatched cell in U_0 is substituted with its counterpart in U_f. Here, f represents the index of the correspondence matrices that correct the unmatched cell. The feature transfer for one data is achieved by multiplying its consolidated correspondence matrix U with another data. After feature transfer, X and Y are jointly embedded into a low-dimensional space for further evaluation.

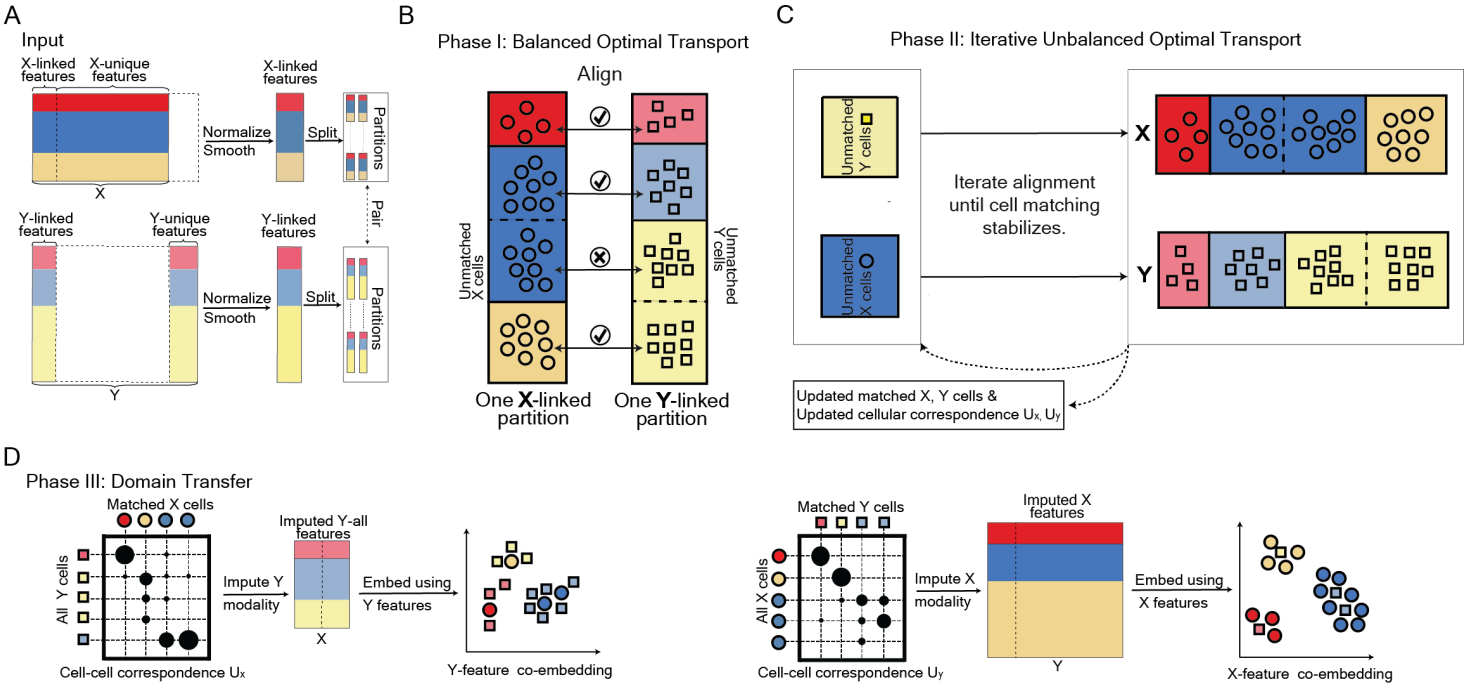

Figure 1. The pipeline of CelLink. A. The input is composed of the all-feature and linked-feature of X and Y modalities. CelLink utilizes the all-feature modality to smooth and denoise the linked-feature modality by the nearest neighbor (NN) graph and low-rank approximation. The smoothed linked modality X_{ss} and Y_{ss} are used to perform the integration. **B.** The first phase is to map cells using balanced optimal transport. CelLink predicts the cell type for each cell by the weight from the cell-cell transport map. The matched cells will be filtered while the unmatched ones will be re-aligned in the next phase. The transport map will be stored as the cell-cell correspondence matrix. **C.** The second phase is to iteratively align the unmatched cells using unbalanced optimal transport. The iterative alignment is performed separately for each modality. In each iteration, unmatched cells are identified via the cell-cell transport map and re-aligned in the next run. This transport map is then preserved as the corrected cell-cell correspondence matrix. The alignment stops when the predicted cell type for all unmatched cells does not change, indicating that they cannot be aligned and the model reaches convergence. **D.** In the third phase, all cell-cell correspondence matrices are consolidated to establish the final cell-cell correspondence framework. CelLink mutually imputes the feature profiles between modalities. The goodness of the alignment is evaluated after jointly embedding the imputed modalities. The output of CelLink is the final correspondence matrix U_x, U_y and the imputed modalities.

Benchmarking CelLink in the integration of scRNA-seq and CODEX data

We benchmarked CelLink across three scRNA-seq and CODEX datasets. One of these datasets features a balanced cell population of 10,000 randomly selected cells from tonsil tissue [18], with 46 protein markers from CODEX data. The other two datasets from HPAP (Home Pancreas Analysis Program) donors exhibit imbalanced cell populations from pancreas tissue [16]. The first HPAP dataset includes 1,340 cells from scRNA-seq and 12,030 cells from CODEX data with 24 protein markers, while the second one includes 1,743 cells from scRNA-seq and 1,516 cells from CODEX data with 24 protein markers as well. The cell-type composition of these three datasets is shown in Supplementary Table 4. We benchmarked CelLink with four baseline methods: CelLink (non-iterative), MaxFuse [8], Seurat v3 [28], and SCOT v2 [9]. The non-iterative version of CelLink represents the case when no cell type or cell state is provided, prompting CelLink to perform the one-time interoperable UOT [27] in the second phase.

A key objective of multi-omics integration is to precisely align the feature profiles from one modality to another. Given the absence of paired cells in the scRNA-seq and CODEX data, we assessed the integration performance based on the quality of cell-type clustering and data-type mixing from five views.

Firstly, the UMAP plots, generated by embedding the original and imputed data into the shared latent space, intuitively showed that CelLink consistently mixed data modalities and separated cell types well with all kinds of feature profiles (Fig. 2A, Supp Fig. 1-3). Another method MaxFuse, although successfully integrated the dataset with balanced cell populations regarding the protein profile, failed to separate the cell types in datasets with imbalanced cell populations (Fig. 2A, Supp Fig. 1B, Supp Fig. 2B). Next, the preservation of local cohesion of data modalities and cell types was evaluated by the Graph Connectivity Score (GCS). This result indicated that CelLink perfectly conserved the local connection within cell types and data types using all kinds of feature profiles compared with other methods in all datasets (Fig. 2B, Supp Fig. 4A-4B). Subsequently, we assessed the preservation of global differences between data types and cell types using the Average Silhouette Width (ASW). CelLink significantly outperformed other methods in all datasets (Fig. 2C). Additionally, cell-type matching accuracy (CTMA) and the count of precisely matched cells (CPMC) are also considered critical metrics to determine whether a method can succeed in aligning the maximum number of cells. From the result, CelLink far surpassed other methods in the imbalanced datasets (Fig. 2D). The heatmaps revealed that CelLink is the only method capable of aligning different cell types with high purity in the presence of imbalanced cell population (Fig. 2E, Supp. Fig. 5A). Although CelLink naturally exhibits higher cell-type matching accuracy owing to its use of cell-type labeling, it does not promote cells of the same type to align. Instead, it just allows unmatched cells to realign until their matched cell types no longer change.

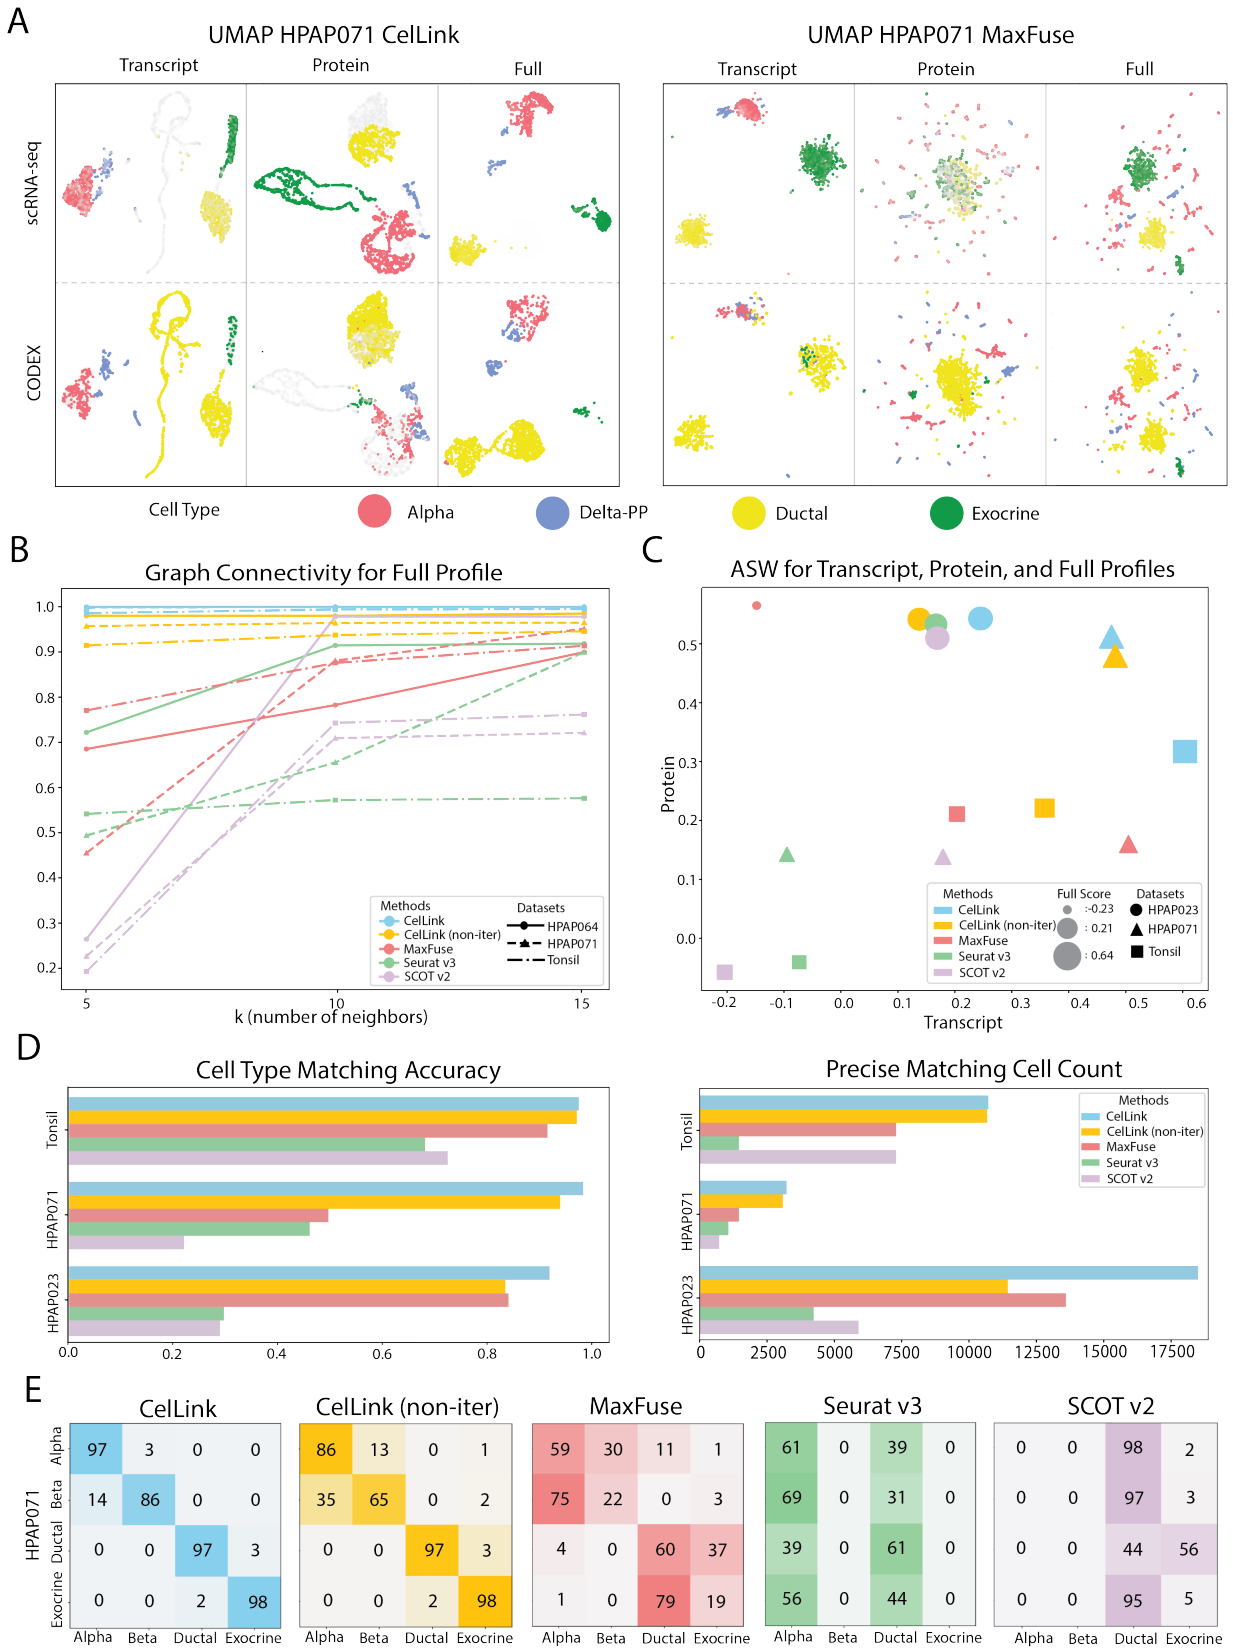

Figure 2. Systematic benchmarking of CelLink on scRNA-seq and CODEX datasets. A. The UMAP plots exhibit the intuitive integration results from CelLink and MaxFuse on HPAP071 donors with the transcript, protein, and full (transcript & protein) profiles. For the transcript profile, scRNA-seq data refers to the original profile while CODEX refers to the the imputed one. For the protein profile, scRNA-seq data refers to the imputed profile while CODEX refers to the original one. **B.** The line plot displays the graph connectivity scores for CelLink and four other baselines across three scRNA-seq and CODEX datasets. Each line indicates the trend of the score when the edges of the graph are built with

the increasing number of neighboring cells. The color of the line represents the method. The line type and point shape represent the dataset. **C.** The scatterplot displays the average silhouette width for CelLink and four other baselines across three datasets using the transcript, protein, and full profiles. The size of the point corresponds to the ASW of the full profile. **D.** Left: The horizontal barplot exhibits the cell-type matching accuracy for CelLink and four other baselines across three datasets. Right: The horizontal barplot exhibits the number of cells that are precisely matched across three datasets. **E.** The heatmaps display the detailed cell-type matching result for CelLink and four other baselines on the HPAP071 dataset.

Additional benchmarking CelLink in the integration of CITE-seq data

Since scRNA-seq and CODEX data do not contain paired cells, We additionally benchmarked CelLink on four CITE-seq datasets, which provide ground-truth transcriptomic and proteomic profiles for each cell. The first two datasets, derived from PBMC and BMC data [12], were sub-sampled to 10,000 cells each. The third dataset was generated by splitting the PBMC CITE-seq data into imbalanced groups: one group with 3,200 cells and another with 6,800 cells. Similarly, the fourth dataset was created by splitting the BMC CITE-seq data into imbalanced groups: one group with 8,000 cells and another with 2,000 cells. Detailed cell-type distributions of the third and fourth datasets are provided in Supplementary Table 2. For these two datasets, we used the transcriptomic profile from the first group and the proteomic profile from the second group for integration, while their corresponding alternative feature profiles served as the masked ground truth for evaluating the integration performance.

From the UMAP plots, CelLink achieved excellent mixing of data modalities and separation of cell types using the transcript, protein, and full feature profiles in all of the datasets (Fig. 3A-B, Supp Fig. 6A-B). In terms of the count of precisely matched cells and cell-type matching accuracy, CelLink consistently outperformed all other baseline methods in all datasets (Fig. 3C).

With the ground-truth transcript and protein profiles for each cell, we induced three additional metrics to evaluate the integration performance at the single-cell level. The first metric, Fraction Of Samples Closer Than True Match (FOSCTTM), assesses how closely the true matching cell is aligned relative to other cells post-integration. CelLink achieved the best performance with both transcript and protein profiles in the PBMC and BMC CITE-seq datasets (Fig. 3D). The second metric, RMSE of Feature Imputation, calculates the average RMSE between the aligned and original feature profiles for all cells. CelLink consistently showed the lowest RMSE score across all datasets using both transcript and protein profiles (Fig. 3E). The third metric, Correlation of Feature Imputation, computes the average Pearson Correlation Coefficient between the aligned and original feature profiles. CelLink achieved the highest Correlation across all datasets using both transcript and protein profiles (Fig. 3E).

In conclusion, CelLink effectively recovers masked feature profiles after integration, demonstrating high precision in imputing both transcriptomic and proteomic features.

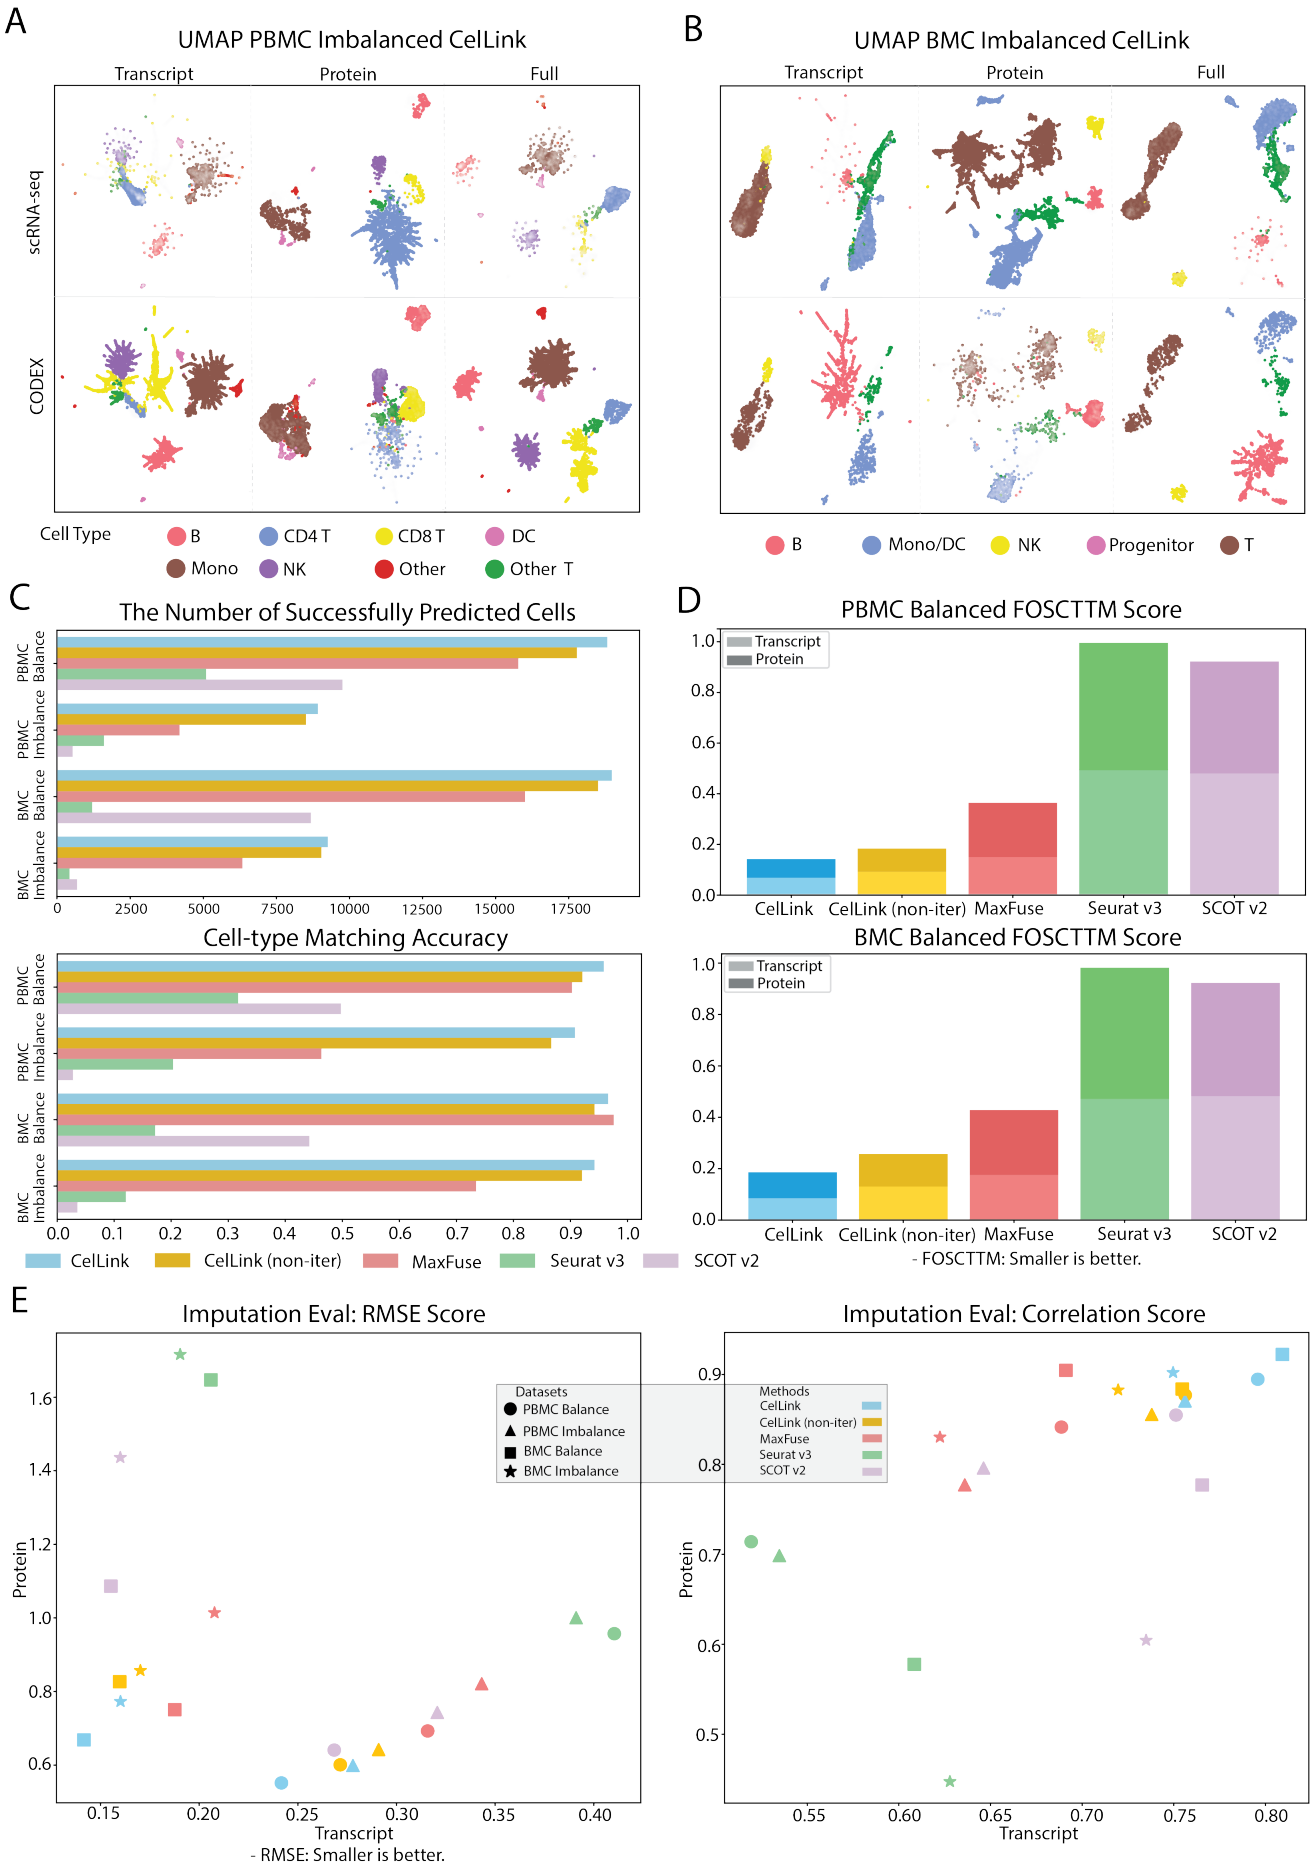
A
UMAP PBMC Imbalanced CelLink
Transcript
Protein
Full
scRNA-seq
CODEX
Cell Type
B
CD4 T
CD8 T
DC
Mono
NK
Other
Other T
B
UMAP BMC Imbalanced CelLink
Transcript
Protein
Full
B
Mono/DC
NK
Progenitor
T
C
The Number of Successfully Predicted Cells
PBMC Balance
PBMC Imbalance
BMC Balance
BMC Imbalance
0
2500
5000
7500
10000
12500
15000
17500
Cell-type Matching Accuracy
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1.0
CelLink
CelLink (non-iter)
MaxFuse
Seurat v3
SCOT v2
D
PBMC Balanced FOSCTTM Score
Transcript
Protein
CelLink
CelLink (non-iter)
MaxFuse
Seurat v3
SCOT v2
BMC Balanced FOSCTTM Score
- FOSCTTM: Smaller is better.
E
Imputation Eval: RMSE Score
Datasets
PBMC Balance
PBMC Imbalance
BMC Balance
BMC Imbalance
Protein
Transcript
- RMSE: Smaller is better.
Imputation Eval: Correlation Score
Methods
CelLink
CelLink (non-iter)
MaxFuse
Seurat v3
SCOT v2
Protein
Transcript

Figure 3. Systematic benchmarking of CelLink on CITE-seq datasets. A. The UMAP plots intuitively exhibited the integration results of CelLink on the imbalanced PBMC data with the transcript, protein, and full (transcript & protein) profiles. **B.**The UMAP plots intuitively exhibited the integration results of CelLink on the imbalanced BMC data with the transcript, protein, and full (transcript & protein) profiles. **C.** Top: The horizontal barplot displayed the number of successfully predicted cells for the four methods on four CITE-seq datasets. Bottom: The horizontal barplot displayed the cell-type matching accuracy for the four methods on four CITE-seq datasets. **D.** Top: The stacked barplot exhibited the FOSCTTM for the four methods on the PBMC CITE-seq dataset using the transcript and protein profile. Bottom: The stacked barplot exhibited the FOSCTTM for the four methods on the BMC CITE-seq dataset using the transcript and protein profile. **E.** Left: The scatterplot manifested the RMSE score of the transcript and protein profiles for the four methods on all CITE-seq datasets. Right: The scatterplot manifested the Correlation score of the transcript and protein profiles for the four methods on all CITE-seq datasets.

CelLink enables cell subtyping and identifies mislabeled cells.

While numerous cell-type annotation methods have been developed, efficient approaches for cell subtype annotation and correction are still lacking. CelLink not only enables annotating the subtypes of cells but also identifies the mislabelled cells. We applied CelLink on this task using scRNA-seq and CODEX data from Bandyopadhyay, Shovik, et al. [1] of bone marrow tissue (Fig. 4A). We randomly subsampled 10,000 cells from scRNA-seq and CODEX data according to the provided cell lineage and subtype annotations (Supplementary table 3). With cell lineage, CelLink successfully integrated the scRNA-seq and CODEX data, as cells originating from the same lineage are grouped in the embedding space (Fig. 4B).

Meanwhile, different cell subtypes were separated well in the result (Fig. 4C). Then, we predicted the cell subtypes in the CODEX data using the cell subtype annotation from the scRNA-seq data, and vice versa. Specifically, the subtype of each cell was predicted by the subtype with the highest weight in the inferred cell-cell correspondence matrix. The performance of subtype prediction was benchmarked with multiple baseline methods, including MaxFuse, Seurat v3, SCOT v2, and non-iterative UOT [8, 28, 9, 27]. CelLink significantly outperformed all other methods regarding the subtype prediction accuracy rate and the count of precisely matched cells (Fig. 4E, Supp Fig. 7A-B). Notably, CelLink achieved a high subtyping accuracy of 87.2% for the CODEX data, confirming its effectiveness in cell subtyping. However, a few subtypes were still not predicted perfectly, including CD4+ T-Cell, GMP, Adipo-MSC, and Macrophages in scRNA-seq and CD8+ T-Cell, Thy1+ MSC in CODEX.

To investigate the reason, we examined the expression of marker genes in the mispredicted cells. Among the 526 mispredicted CD4+ T cells, only 106 expressed the CD4 gene while lacking expression of both CD8A and CD8B genes (Fig. 4F). This suggests that about 80% of the CD4+ T cells were incorrectly predicted as CD8+ T cells due to problematic annotation. Similarly, 267 out of 358 mispredicted Adipo-MSCs express Thy1, while 187 out of 315 mispredicted Thy1+ MSCs did not express Thy1, leading to cross-misprediction between Adipo-MSCs and Thy1+ MSCs (Supp. Fig. 7C). GMP cells were also primarily mispredicted for the same reason, with only 120 out of 662 mispredicted GMP cells expressed both CD34 and CD38 (Supp. Fig. 7C). Some Macrophage cells were mispredicted not only because of the wrong annotation but also due to their overlapping distribution with Monocyte cells regarding both gene and protein expressions (Fig. 4G). Although CD14 is the marker of Monocyte cells, and CD68 and CD163 are markers of Macrophage cells, they can still express each other's markers at lower levels.

In summary, the misprediction of these subtypes is mainly due to two reasons, namely incorrect cell-subtype annotation in scRNA-seq, and the non-distinguishable expression level of the gene expression and protein marker intensities between specific subtypes. CelLink is the first method that enables accurate prediction of cell subtypes by label transferring from one dataset to another and effectively identifying mislabeled cells.

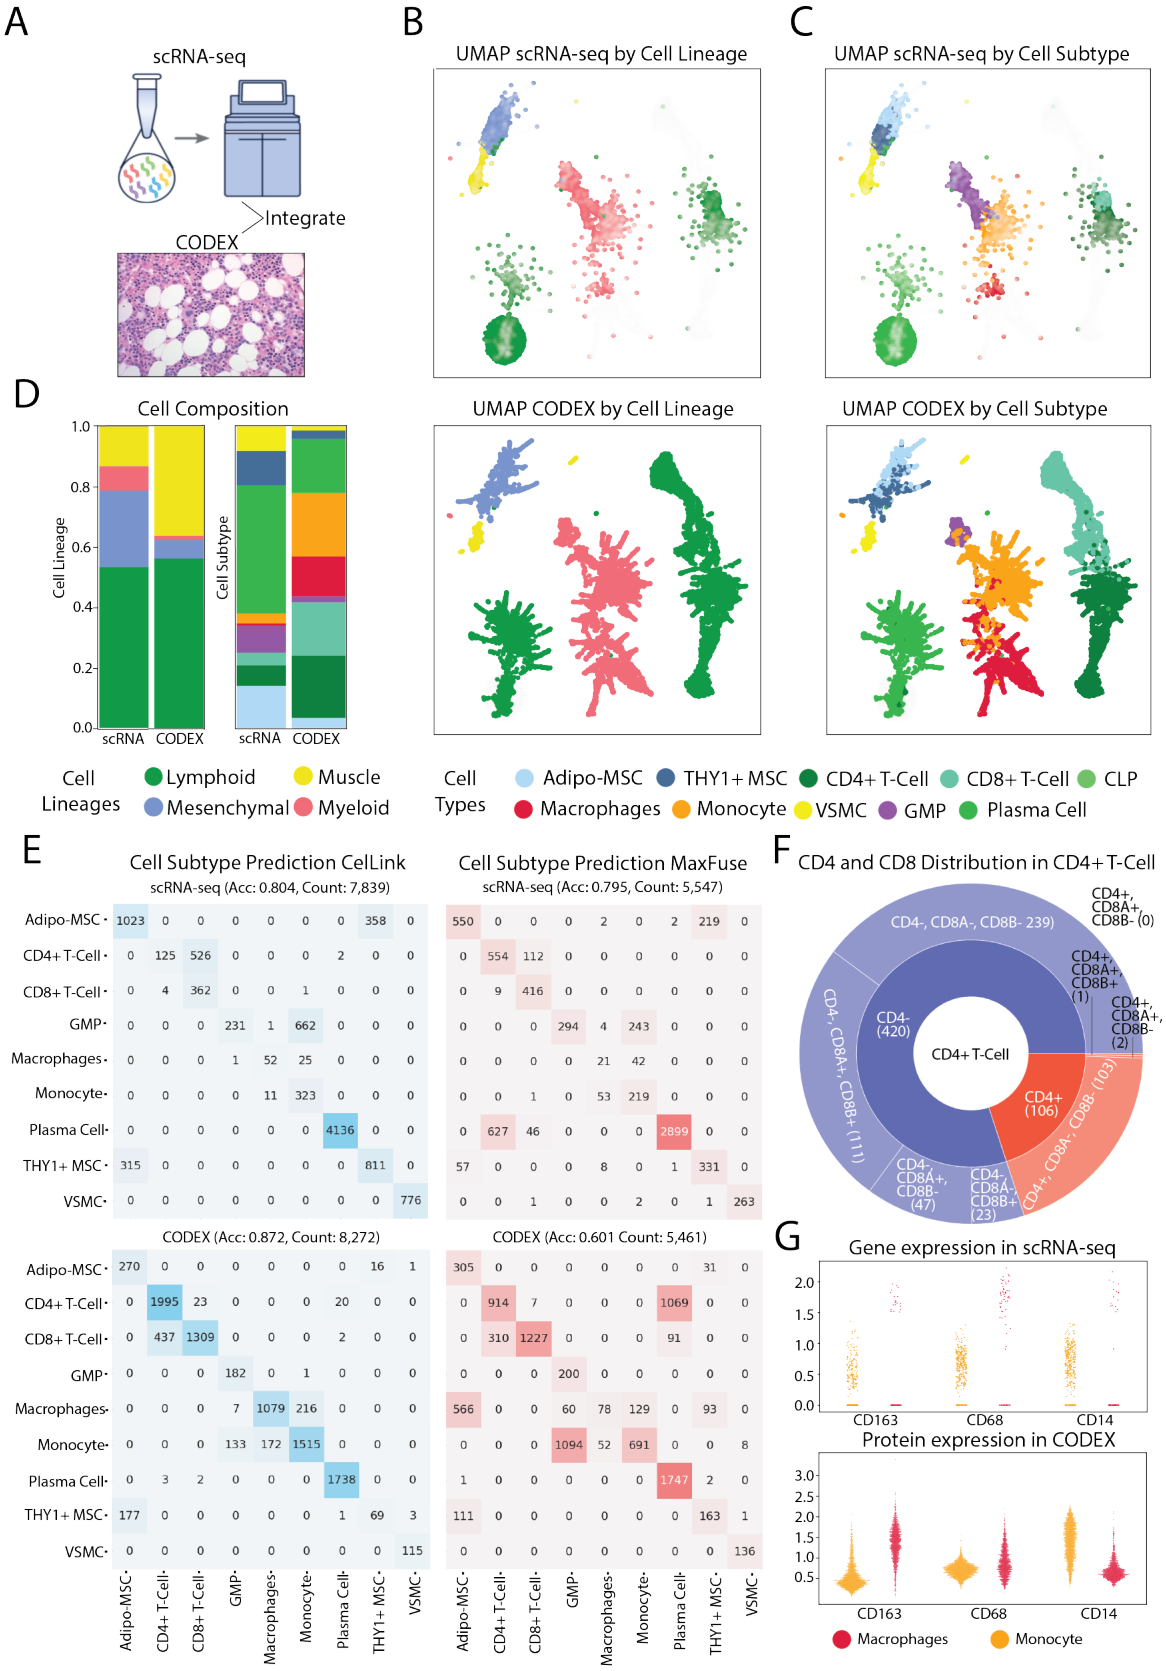

Figure 4. CelLink enables annotation and correction of cell subtypes. A. Overview of integrating scRNA-seq and CODEX data by Bandyopadhyay, Shovik, et al. from bone marrow tissue [1]. **B.** The UMAP plots depicted the distribution of cells annotated by the cell lineage after integration in the embedding space. The four colored cell lineages are "Lymphoid", "Muscle", "Mesenchymal", and "Myeloid". **C.** The UMAP plots depicted the distribution of cells annotated by the cell subtype after integration in the embedding space. The subtypes of Lymphoid cells are "CD4+ T-Cell", "CD8+ T-Cell", "CLP (Common Lymphoid Progenitor)", and "Plasma Cell". The subtypes of Mesenchymal cells are "Adipo-MSC" and "THY1+

MSC". The subtypes of Myeloid cells are "Macrophages" and "Monocyte". The subtype of the muscle cell is the "VSMC (Vascular Smooth Muscle Cells)". **D.** The stacked plots depicted the imbalanced cell composition between scRNA-seq and CODEX data grouped by cell lineage and cell subtypes, respectively. **E.** The heatmaps exhibited the prediction performance of cell subtypes by CelLink compared with MaxFuse. Acc: The cell-subtype prediction accuracy. Count: The total number of cells whose subtypes are predicted precisely across all subtypes. **F.** The sunburst chart displayed the distribution of CD4, CD8A, and CD8B coding genes among the mispredicted CD4+ T Cells. Each partition described the number of cells with a distinct combination of CD4 and CD8 marker expressions. **G.** The strip plots displayed the distribution of CD163, CD68, and CD14 markers within Macrophage and Monocyte cells. The upper plot represented the scRNA-seq data, while the lower plot corresponded to the CODEX data.

CelLink reveals spatial gene expression patterns and cellular interactions.

Recently, there has been a fast-growing number of spatial proteomics datasets, such as CODEX, MIBI-TOF, IMC, and CosMx SMI [2, 19, 6, 13], which measure 30-100 protein markers for single cells along with their spatial coordinates. However, unlike spatial transcriptomics, these proteomic datasets are limited in some spatial analyses, such as spatial cell-cell communications, due to the small number of measured markers. CelLink addresses this limitation by imputing gene expression profiles for spatial proteomics data.

We applied CelLink to infer the spatial cell-cell communication in human pancreas tissue and validated the results by GLKB. We utilized a scRNA-seq and CODEX data from an AAB (auto-antibody positive) HPAP024 donor [16], with the scRNA-seq data comprising 1,266 cells, and the CODEX data comprising 12,351 cells (Supplementary table 4). To infer spatial cellular communication, we experimented with three steps. Firstly, CelLink integrated these two datasets and imputed 10,000 HVGs for the CODEX data. (Fig. 5A). To validate the imputation, we compared the spatial patterns of cell-type specific protein markers with their corresponding coding genes and found them all to be consistent (Supp Fig. 8A-D). Cells in the CODEX data are annotated with four cell types, and each cell type exhibited unique spatial patterns (Fig. 5B). Secondly, We used the Squidpy package and Moran's I index to identify significant ligand-receptor interactions between cell-type pairs, based on their expression levels and spatial autocorrelation (see Methods). We summarized the results for several key ligand-receptor pairs in the pancreatic tissue (Fig. 5C). For example, beta cells act as the sender for the INS-INSR ligand-receptor pair and the IGF signaling pathway. Alpha and beta cells function as the sender and receiver for the GCG-GLP1R ligand-receptor pair. Thirdly, to validate these findings, we employed RAG (Retrieved Augmented Generation) of GLKB (Genomic Literature Knowledge Base) to search for supporting evidence from existing literature [15]. The retrieved answers from GLKB proved that all of the inferred cell-type level ligand-receptor interactions are correct (Fig. 5D).

To investigate the activities in different regions and microenvironments, we conducted cell neighborhood analysis for the CODEX data, identifying four distinct cell neighborhoods. Each cell neighborhood is dominated by one primary cell type mixed with others (Fig. 5E). The result indicates that alpha and beta cells are closely located in the spatial region, corresponding to human islets. In contrast, a subset of ductal cells is positioned near exocrine cells, consistent with the exocrine region of the tissue. Further, we detected the spatial distribution of ligand-receptor pairs regarding cell neighborhoods. For example, in the MDK-ITGA6 ligand-receptor pair, the sender cells are located in alpha and ductal-enriched neighborhoods, while the receptor cells are found in ductal and exocrine-enriched neighborhoods (Fig. 5F, Supp Fig. 8E). This result suggests a coordinated function among various cell neighborhoods, promoting the migration and adhesion of immune cells across the islet in response to the tissue damage and inflammation [3].

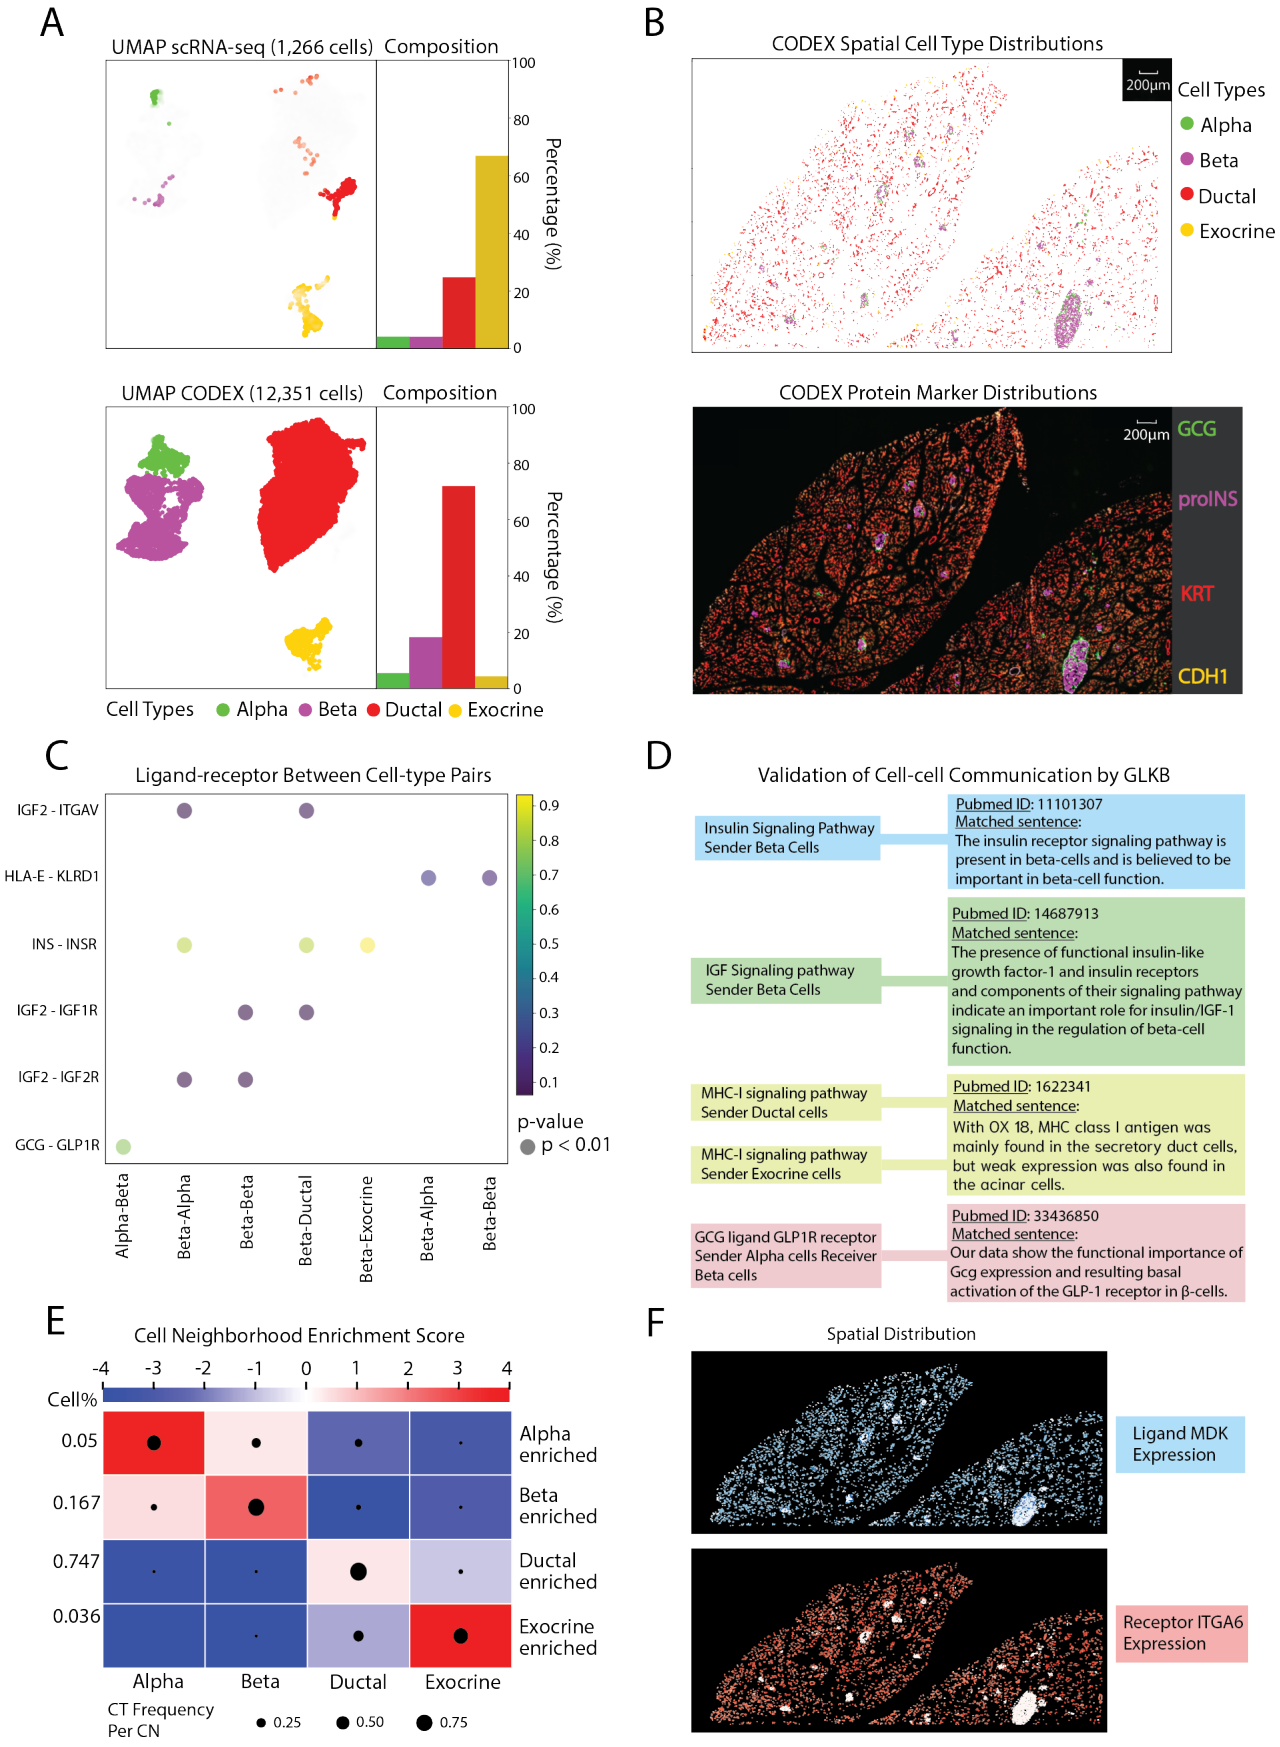

Figure 5. CelLink infers spatial cellular interactions in human pancreas tissue. A. Left: The UMAP plot depicts the cell distribution from scRNA-seq and CODEX data colored by cell type after integration in the embedding space. Right: The barplot depicts the imbalanced cell-type composition between the 1,266 cells from scRNA-seq data and 12,351 cells from

CODEX data. **B.** Upper: The spatial distribution of the cell types from CODEX data after cell segmentation within the sub-region of the head of the human pancreas tissue. Lower: The spatial distribution of the cell-type specific protein markers from the CODEX image within the sub-region of the head of the human pancreas tissue. **C.** The bubble plot demonstrates several significant and pivotal ligand-receptor interactions between cell-type pairs in the pancreas tissue. **D.** The figure illustrates the validation of the inferred cell-type resolution ligand-receptor interactions and signaling pathways in Figure C by RAG of GLKB. The left side displays the query content and the right side shows the retrieved answers from academic papers in PubMed. **E.** The heatmap depicts the identified cell neighborhoods among the selected region in the CODEX data. The size of the point in the heatmap represents the percentage of a cell type within a specific cell neighborhood type. **F.** The spatial distribution of Ligand-receptor pair MDK-ITGA6 in the selected region.

Discussion

Single-cell multi-omics integration is a pivotal case of machine learning domain adaptation in bioinformatics. Traditionally domain adaptation problems are solved by sample alignment or feature alignment between the source and target data. In the case of single-cell multiomics integration, sample alignment methods typically look for cell-to-cell matching between the two datasets. The cell-to-cell matching is either one-to-one like MaxFuse and Seurat V3 [8, 5], or many-to-many like UniPort and SCOT v2 [4, 9]. However, feature alignment methods transform features from different domains into a common latent space, which demands sufficient samples and related features to effectively capture the underlying patterns of the aligned distributions, like LiGER and scGLUE [29, 5]. These methods face challenges in integrating multi-omics data with weak feature linkage and imbalanced cell populations. Therefore, CelLink is developed based on optimal transport, a classic sample alignment method.

The key issue of imbalanced cell populations and weak feature linkage lies not only in the cell counts but also in the exacerbated imbalance among cell types or cell states, comprising various sub-distributions of the two datasets. Previous optimal transport-based methods do not perform well on this problem because they do not address the varying proportions and complexities of these sub-distributions within the overall distribution [10, 9, 4]. Although Unbalanced Optimal Transport (UOT) incorporates the imbalanced mass between the source and target distributions [27], it still hardly achieves the globally optimal transport map when different sub-distributions in the source and target data partially overlap in cases of weak feature linkage, such as between transcript and protein data.

Motivated by this, we developed CelLink based on two key processes. The first process, similar to MaxFuse, solves the issue of weakly linked features by smoothing the feature profile using neighboring cells and normalizing the feature expression values to enhance the comparability. The second process addresses imbalanced cell populations by performing iterative optimal transport. In this process, cell type or cell state is utilized as the label of sub-distribution to refine the alignment for cells that overlap with other sub-distributions. In this way, CelLink automatically filters out noise cells that cannot be aligned. To reduce the computation cost, CelLink also provides functions to split the data into multiple partitions in low sample sizes.

We compared CelLink with four other baselines using multiple datasets with various omics profiles. We demonstrated that CelLink not only significantly outperformed other methods in scenarios with imbalanced cell populations, but also showed slight improvements in datasets with balanced cell populations. Even without cell-type annotation, the non-iterative pipeline of CelLink still achieved excellent integration performance. Given CelLink's strong feature imputation capabilities, it largely helps build multi-modality foundation models by generating paired single-cell multi-omics datasets.

Methods

The Pipeline of CelLink

Model Annotation

Assume the dimension of X as $n_x * n_g$, and the dimension of Y as $n_y * n_p$. Assume the number of linked gene-protein features as n_l. Then the linked modality X_l and Y_l are $n_x * n_l$ and $n_y * n_l$. Assume the number of cells that can be potentially aligned as n_c. Assume $N = max(n_x, n_y)$. The annotations of other variables mentioned in the pipeline are shown in the following table.

Variable Name	Meaning
X, Y	The all-feature modalities X, Y
n_x, n_y	The number of cells in modality X, Y
n_{mx}, n_{my}	The number of matched cells in modality X, Y
n_g, n_p	The number of features in modality X, Y
n_l	The number of linked features between X and Y
$S_x \in (-1,1)^{n_x * n_x}$, $S_y \in (-1,1)^{n_y * n_y}$	The cell-cell similarity matrices of X, Y
$G_x \in \{-1,1\}^{n_x * n_x}$, $G_y \in \{-1,1\}^{n_y * n_y}$	The K-nearest-neighbor graphs of X, Y
K_x, K_y	The diagonal matrices of the number of neighbor cells for each cell in X, Y
X_l, Y_l	The linked-feature modalities of X, Y
$\tilde{X}_l, \tilde{Y}_l$	The smoothed linked-feature modalities $\tilde{X}_l, \tilde{Y}_l$
$D_{xy} \in (0,2)^{n_x * n_y}$	The distance matrix between X_{ss} and Y_{ss}
U_x, U_y	The final cell-cell correspondence matrix between X and Y

Table 1: Annotations of the variables and their meanings

Model Input and Preprocessing

The input of CelLink is the all-feature modality and linked-feature modality X, Y. To better differentiate cells of different cell types, we smooth gene expressions and protein marker intensities for each cell based on its neighbor cells which have the highest feature similarity to it. The smoothing process is as follows: (i) Calculate the cell-cell similarity matrices of all-feature X, Y as $S_x(n_x * n_x)$ and $S_y(n_y * n_y)$. The default metric to calculate similarity is the Pearson Correlation Coefficient. (ii) Build the K-nearest-neighbor graph for X and Y based on S_x and S_y, named G_x and G_y, with nodes as cells and edges as the correlation coefficient between cells. (iii) Smooth the linked-feature profile of each cell from its nearest neighbors $\tilde{X}_l = w * X_l + (1-w) * K_x^{-1} * G_x * X_l, \tilde{Y}_l = w * Y_l + (1-w) * K_y^{-1} * G_y * Y_l$, where $K_x = diag(k_1^x, k_2^x, ..., k_{n_x}^x), K_y = diag(k_1^y, k_2^y, ..., k_{n_y}^y)$, and k_i^x, k_j^y represent the number of neighboring cells for cell i and j.

Data Partition To better fit the model and save the computational resource, we provide an optional function to split data into multiple partitions so that cell counts between modalities are down-sampled to be closed. For example, we split a CODEX data comprising 120,000 cells into 10 partitions. Each partition will integrate with a scRNA-seq with 10,000 cells.

Each pair of batching of $\tilde{X}_l$ and $\tilde{Y}_l$ are the inputs of the phase I.

Phase I: Balanced Optimal Transport

(i) Calculate the cell-pairwise distance matrix $D_{xy}(n_x * n_y)$ from X_{ss} and Y_{ss}. The distance between cell i from modality X and cell j from modality Y is $1 - corr(X_{ss,i}, Y_{ss,j})$.

(ii) Determine D_{xy} as the cost matrix, infer the original cell-cell correspondence matrix $U_0(n_x * n_y)$ between X_{ss} and Y_{ss} by balanced optimal transport. We assume the mass of each cell as one. Thus, the mass transported from cells in X to Y is the correspondence. The OT problem is expressed as

$$U^* = argmin_{U \in R_+^{n_x * n_y}} (\sum_{i=1}^{n_x} \sum_{j=1}^{n_y} U_{i,j} * D_{i,j} + \rho * \Omega(\gamma)) \tag{1}$$

subject to (1) $U 1_{n_y} = 1_{n_x}$ (2) $U^T 1_{n_x} = 1_{n_y}$ (3) $U_{i,j} \geq 0$ for all i, j.

The entropy regularization term $\Omega(\gamma) = \sum_{i=1}^{n_x} \sum_{j=1}^{n_y} \gamma_{i,j} log(\gamma_{i,j})$ is incorporated by the Sinkhorn algorithm to adjust the sparsity of transport map and accelerate the convergence of the problem [20].

(iii) With the transport map U, we predict cell types for each cell by the maximum of weight the cell type from the other modality. The cell-type matched cells are filtered out while the unmatched ones are retained and will be re-aligned in the second phase. The transport map, also known as the original cell-cell correspondence, is stored as U_0.

Phase II: Iterative Unbalanced Optimal Transport

Iterative unbalanced OT is executed distinctly for unmatched cells within each modality with all cells from the other. Take unmatched cells in modality X as an example. Assume the number of retained unmatched X cells in the kth iteration is n_{xk}.

The procedure is:
(i) Calculate the cell-pairwise distance matrix between n_{xk} unmatched cells in X and all cells in Y. The distance metric is the same as phase I.
(ii) Determine the distance matrix as the transport cost. Since our goal is to find the optimal mapping for unmatched cells in X, we allow cells from X to reach their nearest destination in Y by adjusting the mass of Y cells. Thus, we still assume the initial mass of each cell as one, but only enforce that the mass of cells in X remains unchanged while achieving the minimal transport cost through unbalanced optimal transport (UOT). The UOT problem is expressed as

$$U^* = argmin_{U \in R_+^{n_{xk} * n_y}} (\sum_{i=1}^{n_{xk}} \sum_{j=1}^{n_y} U_{i,j} * D_{i,j} + \rho * \Omega(\gamma) + \rho_x * KL(U1_{n_y}, 1_{n_{xk}}) + \rho_y * KL(U^T 1_{n_{xk}}, 1_{n_y})) \tag{2}$$

where the entropy regularization term $\Omega(\gamma) = \sum_{i=1}^{n_{xk}} \sum_{j=1}^{n_y} \gamma_{i,j} log(\gamma_{i,j})$. ρ, ρ_x, ρ_y are penalty coefficients of the transport map sparsity, the mass change of X after transport, and the mass change of Y after transport, respectively.
(iii) With the transport map U_k (k denotes the current iteration index), we predict cell types for each cell by the maximum of weight the cell type from the other modality. The cell-type matched cells are filtered out while the unmatched ones are retained and will be re-aligned in the second phase. The transport map is stored as U_k.
(iv) Repeat (i), (ii), (iii) until the predicted cell types for unmatched cells do not change, indicating the convergence of unbalanced OT and the left cells cannot be aligned.
The optimization of both balanced and unbalanced optimal transport (OT) mentioned above is executed using the Sinkhorn algorithm. This algorithm efficiently produces a relatively dense transport map (cell-cell correspondence matrix). For users requiring sparse results, albeit with much higher computational complexity, we offer an alternative option: skip Phase I and implement iterative UOT with L-BFGS-B optimization [7] in Phase II.

Phase III: Domain Transfer and Output
We individually impute feature profiles of the other modality for matched cells in X and Y. The procedure to impute Y modality for X is as follows:
(i) Assume the UOT loop runs t times. The series of cellular correspondence matrices $U_0, U_1, ..., U_t$ is consolidated into a unified matrix U_x based on cell IDs. For every cell in U_0 that does not align correctly, its cell correspondence vector is replaced with the corresponding vector from U_f. The subscript f stands for the index of the matrix among the series that correctly aligns the previously unmatched cell.
(ii) Assume the number of matched cells in X is n_{mx}, and the dimension of U is $n_{mx} * n_y$. The Y-feature modality for X is imputed as $X_{imputed} = U * Y$.
(iii) To further evaluate the goodness of the integration, we jointly embed $X_{imputed}$ $(n_{mx} * n_p)$ and $Y(n_y * n_p)$ into a low-dimensional space and visualize it by UMAP plot.
(iv) The procedure to impute X modality for Y is in the same way, with the final cell-cell correspondence matrix and the X-modality feature for Y denoted as $U_y(n_{my} * n_x)$ and $Y_{imputed}$ $(n_{my} * n_g)$. The output of CelLink is U_x, U_y and $X_{imputed}, Y_{imputed}$.

Systematic benchmarking of various single-cell multi-omics datasets
We compared CelLink with four other baseline methods, including non-iterative CelLink, MaxFuse, Seurat V3, and SCOT v2. Non-iterative, CelLink, MaxFuse, and SCOT v2 were implemented in Python, and Seurat V3 was implemented in R. We pre-processed all benchmarking datasets in the same way for all methods and used the default hyperparameters suggested by each method. The pre-processing step includes excluding low-quality cells and extracting highly variable genes and proteins for integration. We benchmarked CelLink on seven single-cell multi-omics datasets: three scRNA-seq and CODEX datasets, and four CITE-seq datasets. The evaluation metrics are shown below. We utilized metrics (1), (2), (5), (6) for the datasets with unpaired cells, and (3), (4), (5), (6) for the datasets with paired cells.

Evaluation Metrics
(1) Average Silhouette Width (ASW): This metric evaluates the global integration performance by measuring the goodness of clustering. For each cell in modality X, it compares its distance to all other cells of the same cell type with the distance to all other cells of the nearest cell type in modality Y in the embedding space. Cells in modality Y are compared with modality X in the same way. ASW is calculated as:

$$s(i) = \frac{b(i) - a(i)}{max(b(i), a(i))} \tag{3}$$

$$ASW = \frac{\sum_{i=1}^{n_x + n_y} s(i)}{n_x + n_y} \tag{4}$$

where for cell i, $a(i)$ is its average distance to all cells in the same cell-type cluster, $b(i)$ is its smallest average distance to all cells among other cell types and $s(i)$ is its Silhouette width.

(2) Graph Connectivity Score (GCS): This metric evaluates the local integration performance by measuring the cohesion of data modalities and cell types in the embedding space. It calculates the proportion of the largest connected component within a cell-cell network to the total number of cells. GCS is calculated as:

$$GCS = \frac{1}{M}\sum_{j=1}^{M}\frac{LCC_j}{N_j} \tag{5}$$

, where M denotes the number of cell types, LCC_j represents the number of cells in the largest connected component, and N_j, is the total number of cells in cell type j. The largest connected component is the maximum number of cells that are linked in the $k-$NN graph built from the real and aligned transcriptomic or proteomic expression profile of cells from all modalities.

(3) Feature Imputation accuracy (FIA): This metric evaluates the integration performance by calculating two scores (sub-metrics): (i) the Pearson Correlation coefficient and (ii) the RMSE score between the imputed feature and true feature profiles for each cell in both modalities. In this case, assume $n_x = n_y = n$. They are calculated as:

$$(i)FIA_{corr} = \frac{1}{2n}\sum_{i=1}^{n}(corr(f_{xi}, \hat{f_{xi}}) + corr(f_{yi}, \hat{f_{yi}})) \tag{6}$$

$$(ii)FIA_{RMSE} = \sqrt{\frac{1}{2n}\sum_{i=1}^{n}((f_{xi} - \hat{f_{xi}})^T(f_{xi} - \hat{f_{xi}}) + (f_{yi} - \hat{f_{yi}})^T(f_{yi} - \hat{f_{yi}}))} \tag{7}$$

where $f_{xi}, \hat{f_{xi}}, f_{yi}, \hat{f_{yi}}$ represent the true and imputed feature profile for cell i in modality X and Y, and $corr(f_{xi}, \hat{f_{xi}})$ denotes the Pearson Correlation coefficient between them.

(4) Fraction Of Samples Closer Than True Match (FOSCTTM): This metric evaluates the integration performance by determining how many cells from an alternate modality are closer to the target cell than its ground-truth paired cell in the alternate modality. In this case, assume $n_x = n_y = n$. FOSCTTM is calculated as:

$$FOSCTTM = \frac{1}{2n}(\sum_{i=1}^{n}\frac{n_x^{(i)}}{n} + \sum_{i=1}^{n}\frac{n_y^{(i)}}{n}) \tag{8}$$

where for each i, $n_x^{(i)} = \#\{j \mid d(x_i, y_j) < d(x_i, y_i)\}$ with d denoting the Euclidean distance between the embedding of cells.

(5) Cell-type matching accuracy (CTMA): This metric evaluates the integration performance by checking the percentage of cells whose matched cell types are equal to their original cell types. We matched the cell type for each cell in one modality by the cell type that has the maximum correspondence score from the other modality. For a given cell i, its correspondence score for cell type t is calculated as $Score_{i,t} = \sum_{jint} U(i,j)$. Then, CTMA is calculated as:

$$CTMA = \frac{n_{mx} + n_{my}}{n_x + n_y} \tag{9}$$

(6) Count of precisely matched cells (CPMC): This metric evaluates the integration performance by calculating the number of cells whose matched cell types are equal to their original cell types. It is a supplementary metric to (5), judging whether a method can succeed in aligning the maximum number of cells. CPMC is calculated as:

$$CPMC = n_{mx} + n_{my} \tag{10}$$

where the meanings of the parameters n_x, n_y, n_{mx}, n_{my} are consistent with those provided in the model annotation table. Since cell-type information is utilized as a criterion to determine if each cell is successfully aligned in each iteration, it is natural that CelLink significantly outperforms other methods. Thus, we consider this metric as a supplementary evaluation.

Cell subtyping

Given the cell subtypes of one data, CelLink performs cell subtype annotations for another data by the inferred cell-cell correspondence matrix. For example, assume scRNA-seq as X and CODEX as Y. If cell subtypes of X are provided, subtypes of cells in Y are annotated by the maximum subtype correspondence score in the matrix U_y. For a given cell i, the subtype correspondence score for subtype t is calculated as $Score_{i,t} = \sum_{j \in t} U_y(i,j)$. If cell subtypes of Y are provided, subtypes of cells in X can be annotated similarly.

Spatial Cell-cell Communication Analysis

Spatial cellular communication is inferred by assessing ligand-receptor (LR) interaction scores between specific cell-type pairs relative to random pairs and calculating their spatial autocorrelation scores. The former assessment is conducted using the Squidpy package [25], which employs permutation tests to generate random cell-type pairs. The second assessment calculates the Moran's I index score of a LR Pair between a specific cell-type pair:

$$I = \frac{N}{W} * \frac{\sum_{i=1}^{N} \sum_{j=1}^{N} w_{ij}(x_i - \bar{x})(x_j - \bar{x})}{\sum_{i=1}^{N} (x_i - \bar{x})^2} \tag{11}$$

Here N denotes the number of cells in the cell-type pair, w_{ij} denotes the spatial weight between cell i and j, and $\bar{x}$ denotes the mean expression level of ligand and receptor across all cells. The calculation is performed by the "esda" package.

LR pairs were considered significant if they showed statistical significance in both assessments and had the cell-type level expression of ligand and receptor above 0.1 to reduce noise.

HPAP scRNA-seq and CODEX datasets analysis

scRNA-seq: The team conducted a comprehensive single-cell RNA sequencing analysis starting with basic quality control a Nextflow-based snRNA-Seq workflow [26] for data processing and downstream analysis using Seurat 4 [12] and other tools described here. STARsolo [17] was used to align the 10X Chromium short read data to the human GRCh38 (primary assembly) genome and mapped to GENCODE v42 annotations. The initial quality filtering was fairly standard, keeping cells with more than 500 features and less than 15% mitochondrial reads while ensuring genes were detected in at least three cells. Identifying real cells versus empty droplets was approached systematically by testing three different methods: EmptyDrops [23] with a strict 1% FDR, a knee-point threshold method, and finally a Dropkick-based threshold approach. The team ultimately selected the Dropkick method [14], applying sample-specific UMI thresholds. Ambient RNA contamination was handled through a two-step SoupX-based process [30]. After testing both automated and manual approaches, we settled on a hybrid method: first using automatic contamination estimation, followed by a manual adjustment with an additional 0.2 contamination fraction. Doublet detection relied on scDoubletFinder [24] with a sample-based increase of the doublet rate parameter up to a value of 0.3, which on average identified 16% of cells as potential doublets, which were then removed. A manual inspection was performed to remove cell clusters single-donor derived. Integration relied on Harmony [21] and included the following covariates: HPAP ID, technical platforms, and tissue source. For clustering, we used 2,000 most variable genes, using 15 principal components and a Louvain algorithm with 0.3 resolution. The final dataset includes key clinical metadata such as age, disease type, and T1D-related groupings. CODEX: HPAP OCT tissue blocks were sectioned at 8 µm and prepared for CODEX imaging according to the PhenoCycler User Manual (Akoya Biosciences). PhenoImager (Akoya Biosciences) acquired image tiles and performed tile stitching, deconvolution, background subtraction, and image registration. Processed images were then segmented into cell objects using HALO AI (Indica Labs) trained on manually annotated DAPI images. Cells were exported as tables with associated cell marker intensity data. These cell tables were further processed by Pearson residual transformation followed by Principal Component Analysis (PCA). The resulting principal components were used as input for Uniform Manifold Approximation and Projection (UMAP), reducing the number of available dimensions to two. These two dimensions were used in Leiden clustering, which resulted in a list of clusters and Z-score normalized average intensities for each cell marker. These clusters and cell marker intensities were visualized as a heatmap, which was then inspected for cell marker intensity signatures that could identify the type of cell corresponding to that cluster. Clusters that contained conflicting cell marker data were sub-clustered to identify less common cell populations. Cells were then assigned cell types according to their cluster membership.

Pre-process of the scRNA-seq and CODEX datasets

We pre-processed the three pancreas and one bone marrow scRNA-seq and CODEX datasets in the same way. For the scRNA-seq data, we normalized the gene expressions by log transformation. We selected the most 10,000 highly variable genes and next filtered the non-static ones whose standard deviations were above 0.5. For the CODEX data, we normalized the protein marker intensities by log transformation and scaled them. We filtered the non-static protein markers whose standard deviations were above 0.5 as well.

References

[1] Shovik Bandyopadhyay, Michael P Duffy, Kyung Jin Ahn, Jonathan H Sussman, Minxing Pang, David Smith, Gwendolyn Duncan, Iris Zhang, Jeffrey Huang, Yulieh Lin, et al. Mapping the cellular biogeography of human bone marrow niches using single-cell transcriptomics and proteomic imaging. *Cell*, 187(12):3120–3140, 2024.

[2] Sarah Black, Darci Phillips, John W Hickey, Julia Kennedy-Darling, Vishal G Venkataraaman, Nikolay Samusik, Yury Goltsev, Christian M Schürch, and Garry P Nolan. Codex multiplexed tissue imaging with dna-conjugated antibodies. *Nature protocols*, 16(8):3802–3835, 2021.

[3] Warren A Campbell, Amanda Fritsch-Kelleher, Isabella Palazzo, Thanh Hoang, Seth Blackshaw, and Andy J Fischer. Midkine is neuroprotective and influences glial reactivity and the formation of müller glia-derived progenitor cells in chick and mouse retinas. *Glia*, 69(6):1515–1539, 2021.

[4] Kai Cao, Qiyu Gong, Yiguang Hong, and Lin Wan. uniport: a unified computational framework for single-cell data integration with optimal transport. *bioRxiv*, 2022. doi: 10.1101/2022.02.14.480323.

[5] Zhi-Jie Cao and Ge Gao. Multi-omics single-cell data integration and regulatory inference with graph-linked embedding. *Nature Biotechnology*, 40(10):1458–1466, 2022.

[6] Qing Chang, Olga I Ornatsky, Iram Siddiqui, Alexander Loboda, Vladimir I Baranov, and David W Hedley. Imaging mass cytometry. *Cytometry part A*, 91(2):160–169, 2017.

[7] Laetitia Chapel, Rémi Flamary, Haoran Wu, Cédric Févotte, and Gilles Gasso. Unbalanced optimal transport through non-negative penalized linear regression. *Advances in Neural Information Processing Systems*, 34:23270–23282, 2021.

[8] Shuxiao Chen, Bokai Zhu, Sijia Huang, John W Hickey, Kevin Z Lin, Michael Snyder, William J Greenleaf, Garry P Nolan, Nancy R Zhang, and Zongming Ma. Integration of spatial and single-cell data across modalities with weakly linked features. *Nature Biotechnology*, pages 1–11, 2023.

[9] Pinar Demetci, Rebecca Santorella, Björn Sandstede, and Ritambhara Singh. Unsupervised integration of single-cell multi-omics datasets with disparities in cell-type representation. *bioRxiv*. doi: 10.1101/2021.11.09.467903.

[10] Pinar Demetci, Rebecca Santorella, Björn Sandstede, William Stafford Noble, and Ritambhara Singh. Scot: single-cell multi-omics alignment with optimal transport. *Journal of computational biology*, 29(1):3–18, 2022.

[11] Saad Haider and Ranadip Pal. Integrated analysis of transcriptomic and proteomic data. *Current genomics*, 14(2):91–110, 2013.

[12] Yuhan Hao, Stephanie Hao, Erica Andersen-Nissen, William M Mauck, Shiwei Zheng, Andrew Butler, Maddie J Lee, Aaron J Wilk, Charlotte Darby, Michael Zager, et al. Integrated analysis of multimodal single-cell data. *Cell*, 184(13):3573–3587, 2021.

[13] Shanshan He, Ruchir Bhatt, Carl Brown, Emily A Brown, Derek L Buhr, Kan Chantranuvatana, Patrick Danaher, Dwayne Dunaway, Ryan G Garrison, Gary Geiss, et al. High-plex imaging of rna and proteins at subcellular resolution in fixed tissue by spatial molecular imaging. *Nature Biotechnology*, 40(12):1794–1806, 2022.

[14] Cody N Heiser, Victoria M Wang, Bob Chen, Jacob J Hughey, and Ken S Lau. Automated quality control and cell identification of droplet-based single-cell data using dropkick. *Genome Research*, 31(10):1742–1752, 2021.

[15] Yuanhao Huang, Zhaowei Han, Xin Luo, Xuteng Luo, Yijia Gao, Meiqi Zhao, Feitong Tang, Yiqun Wang, Jiyu Chen, Chengfan Li, et al. Building a literature knowledge base towards transparent biomedical ai. *bioRxiv*, pages 2024–09, 2024.

[16] Klaus H Kaestner, Alvin C Powers, Ali Naji, HPAP Consortium, and Mark A Atkinson. Nih initiative to improve understanding of the pancreas, islet, and autoimmunity in type 1 diabetes: the human pancreas analysis program (hpap). *Diabetes*, 68(7): 1394–1402, 2019.

[17] Benjamin Kaminow, Dinar Yunusov, and Alexander Dobin. Starsolo: accurate, fast and versatile mapping/quantification of single-cell and single-nucleus rna-seq data. *Biorxiv*, pages 2021–05, 2021.

[18] Julia Kennedy-Darling, Salil S Bhate, John W Hickey, Sarah Black, Graham L Barlow, Gustavo Vazquez, Vishal G Venkataraaman, Nikolay Samusik, Yury Goltsev, Christian M Schürch, et al. Highly multiplexed tissue imaging using repeated oligonucleotide exchange reaction. *European Journal of Immunology*, 51(5):1262–1277, 2021.

[19] Leeat Keren, Marc Bosse, Steve Thompson, Tyler Risom, Kausalia Vijayaragavan, Erin McCaffrey, Diana Marquez, Roshan Angoshtari, Noah F Greenwald, Harris Fienberg, et al. Mibi-tof: A multiplexed imaging platform relates cellular phenotypes and tissue structure. *Science advances*, 5(10):eaax5851, 2019.

[20] Philip A Knight. The sinkhorn–knopp algorithm: convergence and applications. *SIAM Journal on Matrix Analysis and Applications*, 30(1):261–275, 2008.

[21] Ilya Korsunsky, Nghia Millard, Jean Fan, Kamil Slowikowski, Fan Zhang, Kevin Wei, Yuriy Baglaenko, Michael Brenner, Po-ru Loh, and Soumya Raychaudhuri. Fast, sensitive and accurate integration of single-cell data with harmony. *Nature methods*, 16 (12):1289–1296, 2019.

[22] Jie Liu, Yuanhao Huang, Ritambhara Singh, Jean-Philippe Vert, and William Stafford Noble. Jointly embedding multiple single-cell omics measurements. In *Algorithms in bioinformatics:... International Workshop, WABI..., proceedings. WABI (Workshop)*, volume 143. NIH Public Access, 2019.

[23] Aaron TL Lun, Samantha Riesenfeld, Tallulah Andrews, The Phuong Dao, Tomas Gomes, Participants in the 1st Human Cell Atlas Jamboree, and John C Marioni. Emptydrops: distinguishing cells from empty droplets in droplet-based single-cell rna sequencing data. *Genome biology*, 20:1–9, 2019.

[24] Christopher S McGinnis, Lyndsay M Murrow, and Zev J Gartner. Doubletfinder: doublet detection in single-cell rna sequencing data using artificial nearest neighbors. *Cell systems*, 8(4):329–337, 2019.

[25] Giovanni Palla, Hannah Spitzer, Michal Klein, David Fischer, Anna Christina Schaar, Louis Benedikt Kuemmerle, Sergei Rybakov, Ignacio L Ibarra, Olle Holmberg, Isaac Virshup, et al. Squidpy: a scalable framework for spatial omics analysis. *Nature methods*, 19(2):171–178, 2022.

[26] ParkerLab. snRNAseq-NextFlow: Nextflow-based snRNA-Seq workflow. https://github.com/ParkerLab/snRNAseq-NextFlow, 2024. Repository continuously updated. [Online; last accessed 2024-04].

[27] Khiem Pham, Khang Le, Nhat Ho, Tung Pham, and Hung Bui. On unbalanced optimal transport: An analysis of sinkhorn algorithm. In *International Conference on Machine Learning*, pages 7673–7682. PMLR, 2020.

[28] Tim Stuart, Andrew Butler, Paul Hoffman, Christoph Hafemeister, Efthymia Papalexi, William M Mauck, Yuhan Hao, Marlon Stoeckius, Peter Smibert, and Rahul Satija. Comprehensive integration of single-cell data. *cell*, 177(7):1888–1902, 2019.

[29] Joshua D Welch, Velina Kozareva, Ashley Ferreira, Charles Vanderburg, Carly Martin, and Evan Z Macosko. Single-cell multi-omic integration compares and contrasts features of brain cell identity. *Cell*, 177(7):1873–1887, 2019.

[30] Matthew D Young and Sam Behjati. Soupx removes ambient rna contamination from droplet-based single-cell rna sequencing data. *Gigascience*, 9(12):giaa151, 2020.