
Uncertainty Quantification for Atlas-Level Cell Type Transfer

Jan Engelmann^{1*} Leon Hetzel^{1 2*} Giovanni Palla^{1 3*} Lisa Sikkema^{1 3} Malte Luecken¹ Fabian Theis^{1 2 3}

Abstract

Single-cell reference atlases are large-scale, cell-level maps that capture cellular heterogeneity within an organ using single cell genomics. Given their size and cellular diversity, these atlases serve as high-quality training data for the transfer of cell type labels to new datasets. Such label transfer, however, must be robust to domain shifts in gene expression due to measurement technique, lab specifics and more general batch effects. This requires methods that provide uncertainty estimates on the cell type predictions to ensure correct interpretation. Here, for the first time, we introduce uncertainty quantification methods for cell type classification on single-cell reference atlases. We benchmark four model classes and show that currently used models lack calibration, robustness, and actionable uncertainty scores. Furthermore, we demonstrate how models that quantify uncertainty are better suited to detect unseen cell types in the setting of atlas-level cell type transfer.

1. Introduction

Single-cell genomics enables the characterization of cellular states at unprecedented scale and resolution. This has led to the generation of diverse single-cell RNA sequencing (scRNA-seq) datasets that capture cellular heterogeneity at the tissue and organ level. To combine multiple studies into a single representation of a human organ, consortia such as the Human Cell Atlas (Regev et al., 2017) are leading the generation of *integrated scRNA-seq atlases*. Such atlases are built by integrating multiple datasets across samples, individuals, and technologies, and are deemed to comprehensively capture cellular variation in the given tissue or organ. Recently, the first such comprehensive representation of the human lung was introduced in the Human Lung Cell

Atlas (HLCA) (Sikkema et al., 2022).

Next to their potential for discovery, integrated atlases can rapidly accelerate the analysis of new datasets in a harmonised way. Most notably, this is achieved by transferring cell type labels from the reference atlas to new, unannotated data from the same tissue. This task relies on cell type classifiers which are trained on the latent representation and expert cell labels of the integrated atlas. In order to analyze a new dataset, one “projects” the query data to the reference’s latent space and applies the classifiers to infer cell type labels. The encoding of the query data is based on heavily regularized batch effect correction approaches to match the reference and is therefore not guaranteed to align perfectly (Lotfollahi et al., 2021). Note that the resulting misalignment naturally becomes greater when the projected data is comprised of samples from a different biological condition, e.g., disease or development, or contains unseen cell states.

The scenarios described above result in variable degrees of dataset “shifts”, complicating the transfer of cell type labels. In addition, it is paramount to distinguish shifted latent representations due to unresolved technical batch effects from shifts due to biological differences between the datasets, e.g., new cell types or disease-related cell states. Such a distinction can be accomplished with a classifier that provides a measure of uncertainty, which yields actionable insights for downstream tasks.

In machine learning research data shifts, and how classifiers can overcome them, have been widely studied in *uncertainty quantification* (Gal, 2016; Gawlikowski et al., 2021). Crucially, model predictions should be reliable, i.e. calibrated, as they are used by data analysts to draw key interpretations from the data which influence further downstream analyses. Cell type classifiers should maintain these properties across dataset shifts while remaining sensitive to subtle biologically driven shifts in gene expression. Finally, they should be transparent to the analyst, by providing an accurate measure of confidence on the prediction task.

In this work, we evaluate uncertainty quantification methods for atlas-level cell type transfer. We find that models that quantify uncertainty are more robust, better calibrated, and provide high quality uncertainty measures that enable them to identify unseen cell types in the projected data.

*Equal contribution ¹Institute of Computational Biology, Helmholtz Center Munich, Germany ²Department of Mathematics, TU Munich, Germany ³TUM School of Life Sciences Weihenstephan, Technical University of Munich, Germany. Correspondence to: Fabian Theis <fabian.theis@helmholtz-muenchen.de>.

2. Modelling and evaluation of uncertainty

2.1. Dataset and model choice

The HLCA dataset, on which all evaluations are performed, consists of a core of 14 datasets including 166 scRNA-seq healthy tissue samples of the human respiratory system from 107 individuals (Sikkema et al., 2022), see Figure 7 in appendix A for UMAPs. Batch effects between datasets were removed using scANVI with dataset as the batch covariate (Xu et al., 2021). All cells in the integrated atlas were re-annotated using iterative clustering, data-driven marker gene detection, and a consensus manual annotation from 6 lung experts, leading to 58 consensus cell type labels (Sikkema et al., 2022). The resulting training data consists of $\sim 580\,000$ cells where each cell is represented by a 30 dimensional embedding and has a unique cell type.

As we focus on the models’ ability to provide uncertainty estimates, we consider two baseline models as well as two state-of-the-art classifiers:

- *Weighted K-nearest neighbor* (WKNN) classifier which is used in the scArches method (Lotfollahi et al., 2021) and in the HLCA (Sikkema et al., 2022).
- *Random Forest* (RF) classifier from scikit-learn (Pedregosa et al., 2012).
- *Deep kernel learning with spectral normalized residual network classifier* (DKL) (van Amersfoort et al., 2021; Wilson et al., 2015)
- *Multi-input multi-output classifier* (MIMO) (Havasi et al., 2020), with 3 or 8 subnetworks.

Note that we evaluate the models’ uncertainty only on the latent embedding and do not have access to the network used for data integration. This makes methods like scANVI (Xu et al., 2021) for uncertainty prediction inapplicable.

Uncertainty estimates can be divided into *predictive*, *model*, and *data uncertainty* (Gal, 2016; Gawlikowski et al., 2021). Data uncertainty entails label and measurement noise that is inherent to the data, while model uncertainty measures the uncertainty of the model’s predictions due to parameter uncertainty. An application of model uncertainty is OOD detection, because it is higher for OOD cells than for ambiguous ID cells, which both have high predictive uncertainty. The uncertainties are calculated using:

$$\underbrace{\mathbb{H}[y|\mathbf{x}, \mathcal{D}]}_{\text{predictive}} = \underbrace{\mathbb{I}[y, \boldsymbol{\omega}|\mathbf{x}, \mathcal{D}]}_{\text{model}} + \underbrace{\mathbb{E}_{p(\boldsymbol{\omega}|\mathcal{D})}[\mathbb{H}[y|\mathbf{x}, \boldsymbol{\omega}]]}_{\text{data}},$$

where \mathbb{H} is the entropy and \mathbb{I} the mutual information, (\mathbf{x}, y) the embedding and cell type of the cell, \mathcal{D} the training data and $\mathbb{E}_{p(\boldsymbol{\omega}|\mathcal{D})}$ the expectation over the posterior (Smith &

Gal, 2018). The uncertainties reported throughout this work have been scaled to $[0, 1]$.

In the context of label transfer from reference datasets, data uncertainty would apply to transitioning cell states which cannot be associated clearly with a single cell type. In contrast, model uncertainty highlights how familiar the model is with cells similar to the provided query cell, that is, in an ideal scenario, *the model knows when it does not know*.

DKL and MIMO belong to the class of single forward pass uncertainty estimation models which we chose for their computational efficiency and performance. Other model classes that perform uncertainty estimation include: Bayesian Neural Networks (Blundell et al., 2015), Bayesian Dropout (Gal & Ghahramani, 2015), or Ensemble models (Lakshminarayanan et al., 2016), but they are not further considered here.

We performed hyperparameter optimisation for all model classes across ~ 100 distinct parameter combinations. Furthermore, all reported metrics are computed from model fits across three test sets evaluated on three different train-validation splits, resulting in 9 unique train-validation-test splits.

2.2. Evaluation setup

As we are interested in both the evaluation of cell type transfer and the ability to provide meaningful uncertainty scores, the evaluation is divided in two parts.

In section 3.1, we evaluate all models in terms of their balanced accuracy, F1-score and expected calibration error (ECE) (Guo et al., 2017). The leave-out set is a dataset integrated in the HLCA but not used during classifier training. Next to the predictive performance, we analyse the calibration of the classifiers. Calibration considers the relation between the frequency of correct predictions and the binned prediction probability and, therefore, indicates how reliable a classifier is. A well calibrated classifier will be correct 7 out of 10 times when providing a confidence score of 0.7. For calibration evaluation we report both ECE and calibration curves.

In section 3.2, we test the models’ ability to quantify uncertainty for several out-of-distribution (OOD) scenarios and their ability to discriminate between in-distribution training data (ID) and OOD data. To evaluate the models’ uncertainty quality, we exclude three cell types during training and assess their performance through the predictive and model uncertainty. We evaluate OOD data discrimination using the Wasserstein distance (Krishnan & Tickoo, 2020) as well as the area under the precision-recall curve (AUPR).

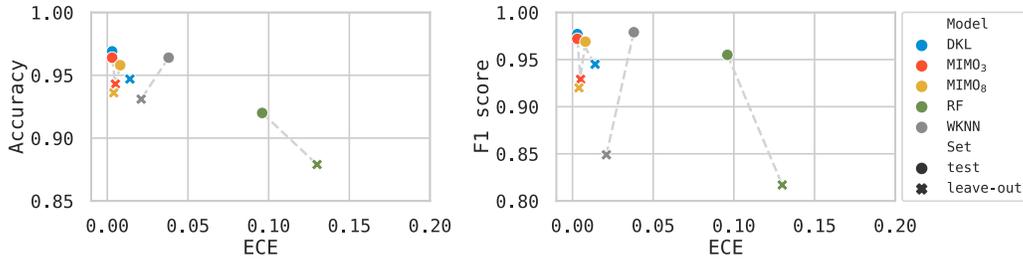


Figure 1. Accuracy (\uparrow) against ECE (\downarrow) (left) and F1 score (\uparrow) against ECE (\downarrow) (right) across all model classes. Each point represents the mean of three runs of nested cross-validation. The standard deviation is not reported as it is $< .002$ and does not affect the interpretation of the results.

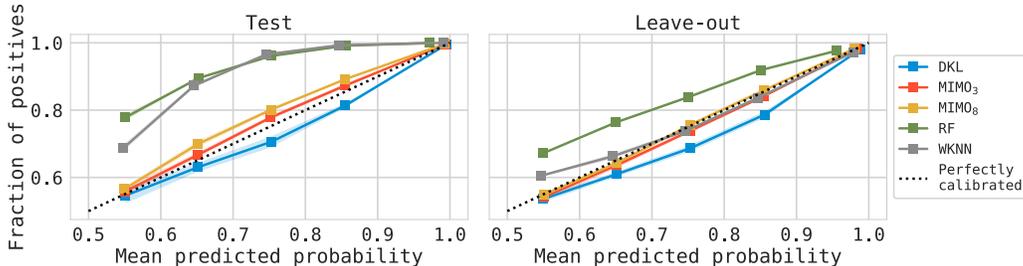


Figure 2. Fraction of positive predictions against the binned mean predicted probability across five bins for all five model classes. Each point represents the mean of three runs of nested cross-validation. The shaded area is the standard deviation.

3. Results

3.1. Label transfer and calibration

We first evaluate the models on their ability to classify cell types on the integrated HLCA dataset using the most fine-grained cell type annotation (58 cell types in total). In a typical use case a new, leave-out dataset introduces a data-shift due to differences in experimental design, sample handling, and cell type composition. We use a random train-test split in this analysis to serve as upper bounds on the performance.

We find that all models have comparable accuracy and F1-scores (Figure 1) on the test set. In contrast, the baseline models perform worse on the leave-out dataset, while DKL and MIMO are able to maintain their good performance. These results show that DKL and MIMO are superior in cell type classification to RF and WKNN.

The same is also reflected in the models’ calibration, where the ECE for MIMO and DKL is consistently lower (better) compared to the baseline models. The calibration curves in Figure 2 provide a more detailed view on calibration by showing the models’ calibration at different confidence levels. It can be seen that WKNN and RF are further away from perfect calibration in the test set than DKL and MIMO. Interestingly, the decline in performance on the leave-out dataset leads to an improved calibration of WKNN, making it comparable to models that quantify uncertainty for the leave-out scenario. Nevertheless, this calibration plot,

together with the ECE, should be seen in light of model accuracy, since a model should be robust not only in terms of calibration, but also accuracy under data shifts. We report exact scores in Tables 2,3 in the appendix A.

Overall, this analysis indicates that models that quantify uncertainty, such as DKL and MIMO, are better suited for the cell type transfer task. They achieve comparable performance in terms of accuracy for both test and leave-out set, while also showing superior calibration compared to the WKNN and RF models.

3.2. Uncertainty quantification on unseen cell types

To evaluate the models’ ability to detect unseen cell types, we left out three cell types—B cells, Mast cells, and Ionocytes—from the training data and compute the predictive uncertainty for all classes seen during training, as well as the left out class. A classifier trained on K classes should result in high entropy of the output for the unseen, $K + 1$ class.

Figure 3 reports predictive uncertainty for B cells for all model types (see Figure 4 in appendix A for additional cell types). DKL clearly identifies B cells as a new, unseen cell type which is shown through the high uncertainty score for the respective observations. In contrast, predictive uncertainty for WKNN is diffused and does not highlight B cells. Consequently, WKNN’s uncertainty estimates can be identified as less reliable which may cause incorrect cell type transfer. We quantitatively evaluate the densities of the

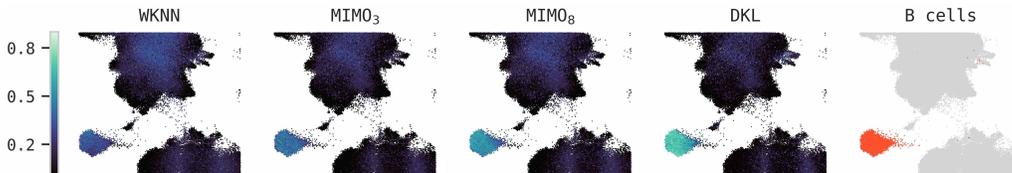


Figure 3. Predictive uncertainty on a subset of the HLCA UMAP with B cells left out during classifier training. It can be observed that the MIMO and DKL models predict higher uncertainty for the B cells. RF results are not reported due to poor classification accuracy.

Table 1. Area under the precision-recall curve (AUPR) for classification between OOD (unseen cell type) and ID (validation set sampled from reference atlas), using different models and uncertainty types. DKL and MIMO outperform WKNN which is the baseline used in the HLCA. Softmax confidence uses the probability of the predicted class as a measure of uncertainty.

MODEL	UNCERTAINTY TYPE	AUPR (POSITIVE CLASS: LEFT OUT CELLTYPE)		
		B CELLS	IONOCYTES	MAST CELLS
DKL	PREDICTIVE MODEL	0.929 ± 0.026	0.940 ± 0.017	0.948 ± 0.009
	SOFTMAX CONFIDENCE	0.824 ± 0.051	0.717 ± 0.202	0.793 ± 0.011
MIMO ₈	PREDICTIVE MODEL	0.911 ± 0.009	0.677 ± 0.108	0.943 ± 0.024
	SOFTMAX CONFIDENCE	0.928 ± 0.005	0.823 ± 0.058	0.964 ± 0.010
WKNN	PREDICTIVE MODEL	0.605 ± 0.020	0.236 ± 0.006	0.825 ± 0.013
	SOFTMAX CONFIDENCE	0.510 ± 0.014	0.121 ± 0.008	0.707 ± 0.016

predictive uncertainty for seen and unseen cell types via the Wasserstein distance, cf. (Krishnan & Tickoo, 2020), see Figure 5 in appendix A, from which the same conclusions can be drawn.

Finally, we formulate the task of detecting a new cell type as a binary classification problem and quantify the performance through the AUPR. The results are shown in Table 1. Each entry is the AUPR of a model when the given uncertainty type is used for classification. MIMO₈ performs best using model uncertainty for all left out cell types. DKL on the other hand has higher AUPR using predictive uncertainty. Since model uncertainty should be better suited for OOD detection, this suggests that MIMO₈ exhibits better disentanglement of model and data uncertainty. DKL’s issue to highlight OOD cells using model uncertainty can also be observed in Figure 6 in appendix A. Nonetheless, DKL has higher AUPR than MIMO₈ for B cells and Ionocytes using predictive uncertainty. When using predictive uncertainty for DKL and model uncertainty for MIMO, both models outperform the WKNN baseline across cell types.

These results indicate that models that quantify uncertainty are superior in detecting unseen cell types, since they produce higher uncertainty on the left out classes. This is a desirable property of the models as it can be used by analysts to discover new cell states which are not present in the integrated scRNA-seq reference atlas.

This analysis requires further investigation and should be carried out on a real OOD scenario, i.e., where a new dataset is projected onto the reference and also excluded from the classifier’s optimisation. Nonetheless, one has to appreciate

that neither of the baseline models, RF and WKNN, provides such a measure on model uncertainty, and, therefore, an evaluation of models’ prediction confidence is not possible. Put differently, it is not possible to state *the model doesn’t know when it doesn’t know*. In contrast, models like DKL and MIMO, that provide measures on both data and model uncertainty, should be preferred for building and extending integrated scRNA-seq reference atlases.

4. Conclusion

In this work, we presented uncertainty quantification on cell type transfer for single-cell reference atlases. We evaluated the considered models on a variety of metrics, including performance, calibration, and uncertainty quantification. We demonstrate that the baseline methods are not well calibrated and lack desirable robustness. In contrast, models that quantify uncertainty provide better confidence and more sensible uncertainty scores, which provide actionable results for downstream interpretations. In fact, such uncertainty measures can be used by analysts to uncover new cell types and cell states not present in the reference integrated atlas, as well as retain a notion on the “confidence” of the prediction, so that more careful conclusions can be drawn. In the future, we seek to extend this benchmark by including real-world query-to-reference mapping as well provide a more extensive evaluation on model and data uncertainty under dataset shifts.

References

- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural networks, 2015. URL <https://arxiv.org/abs/1505.05424>.
- Gal, Y. Uncertainty in deep learning. 2016.
- Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning, 2015. URL <https://arxiv.org/abs/1506.02142>.
- Gawlikowski, J., Tassi, C. R. N., Ali, M., Lee, J., Humt, M., Feng, J., Kruspe, A. M., Triebel, R., Jung, P., Roscher, R., Shahzad, M., Yang, W., Bamler, R., and Zhu, X. X. A survey of uncertainty in deep neural networks. *CoRR*, abs/2107.03342, 2021. URL <https://arxiv.org/abs/2107.03342>.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. In *International Conference on Machine Learning*, pp. 1321–1330. PMLR, 2017.
- Havasi, M., Jenatton, R., Fort, S., Liu, J. Z., Snoek, J., Lakshminarayanan, B., Dai, A. M., and Tran, D. Training independent subnetworks for robust prediction. October 2020.
- Krishnan, R. and Tickoo, O. Improving model calibration with accuracy versus uncertainty optimization. December 2020.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles, 2016. URL <https://arxiv.org/abs/1612.01474>.
- Lotfollahi, M., Naghipourfar, M., Luecken, M. D., Khajavi, M., Büttner, M., Wagenstetter, M., Avsec, Ž., Gayoso, A., Yosef, N., Interlandi, M., Rybakov, S., Misharin, A. V., and Theis, F. J. Mapping single-cell data to reference atlases by transfer learning. *Nat. Biotechnol.*, August 2021.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., VanderPlas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in python. *CoRR*, abs/1201.0490, 2012. URL <http://arxiv.org/abs/1201.0490>.
- Regev, A., Teichmann, S. A., Lander, E. S., Amit, I., Benoist, C., Birney, E., Bodenmiller, B., Campbell, P., Carninci, P., Clatworthy, M., Clevers, H., Deplancke, B., Dunham, I., Eberwine, J., Eils, R., Enard, W., Farmer, A., Fugger, L., Göttgens, B., Hacohen, N., Haniffa, M., Hemberg, M., Kim, S., Klenerman, P., Kriegstein, A., Lein, E., Linnarsson, S., Lundberg, E., Lundeberg, J., Majumder, P., Marioni, J. C., Merad, M., Mhlanga, M., Nawijn, M., Netea, M., Nolan, G., Pe’er, D., Phillipakis, A., Ponting, C. P., Quake, S., Reik, W., Rozenblatt-Rosen, O., Sanes, J., Satija, R., Schumacher, T. N., Shalek, A., Shapiro, E., Sharma, P., Shin, J. W., Stegle, O., Stratton, M., Stubbington, M. J. T., Theis, F. J., Uhlen, M., van Oudenaarden, A., Wagner, A., Watt, F., Weissman, J., Wold, B., Xavier, R., Yosef, N., and Human Cell Atlas Meeting Participants. The human cell atlas. *Elife*, 6, December 2017.
- Sikkema, L., Strobl, D., Zappia, L., Madisson, E., Markov, N., Zaragosi, L., Ansari, M., Arguel, M., Apperloo, L., Bécavin, C., Berg, M., Chichelnitskiy, E., Chung, M., Collin, A., Gay, A., Hooshyar Kashani, B., Jain, M., Kapellos, T., Kole, T., Mayr, C., von Papen, M., Peter, L., Ramírez-Suástegui, C., Schmiering, J., Taylor, C., Walzthoeni, T., Xu, C., Bui, L., de Donno, C., Dony, L., Guo, M., Gutierrez, A., Heumos, L., Huang, N., Ibarra, I., Jackson, N., Kadur Lakshminarasimha Murthy, P., Lotfollahi, M., Tabib, T., Talavera-Lopez, C., Travaglini, K., Wilbrey-Clark, A., Worlock, K., Yoshida, M., Consortium, L. B. N., Desai, T., Eickelberg, O., Falk, C., Kaminski, N., Krasnow, M., Lafyatis, R., Nikolčić, M., Powell, J., Rajagopal, J., Rozenblatt-Rosen, O., Seibold, M., Sheppard, D., Shepherd, D., Teichmann, S., Tsankov, A., Whitsett, J., Xu, Y., Banovich, N., Barbry, P., Duong, T., Meyer, K., Kropski, J., Pe’er, D., Schiller, H., Tata, P., Schultze, J., Misharin, A., Nawijn, M., Luecken, M., and Theis, F. An integrated cell atlas of the human lung in health and disease. *bioRxiv*, 2022. doi: 10.1101/2022.03.10.483747. URL <https://www.biorxiv.org/content/early/2022/03/11/2022.03.10.483747>.
- Smith, L. and Gal, Y. Understanding Measures of Uncertainty for Adversarial Example Detection, March 2018. URL <http://arxiv.org/abs/1803.08533>.
- van Amersfoort, J., Smith, L., Jesson, A., Key, O., and Gal, Y. Improving deterministic uncertainty estimation in deep learning for classification and regression. *CoRR*, abs/2102.11409, 2021. URL <https://arxiv.org/abs/2102.11409>.
- Wilson, A. G., Hu, Z., Salakhutdinov, R., and Xing, E. P. Deep kernel learning. November 2015.
- Xu, C., Lopez, R., Mehlman, E., Regier, J., Jordan, M. I., and Yosef, N. Probabilistic harmonization and annotation of single-cell transcriptomics data with deep generative models. *Mol. Syst. Biol.*, 17(1):e9620, January 2021.

A. Appendix

Table 2. Test set classification accuracy, F1 score and ECE

	ACCURACY (\uparrow)	F1 SCORE (\uparrow)	ECE (\downarrow)
DKL	0.969	0.977	0.003
MIMO ₃	0.964	0.972	0.003
MIMO ₈	0.958	0.969	0.008
RF	0.920	0.955	0.096
WKNN	0.964	0.979	0.038

Table 3. Leave-out set classification accuracy, F1 score and ECE

	ACCURACY (\uparrow)	F1 SCORE (\uparrow)	ECE (\downarrow)
DKL	0.947	0.945	0.014
MIMO ₃	0.943	0.929	0.005
MIMO ₈	0.936	0.920	0.004
RF	0.879	0.817	0.130
WKNN	0.931	0.849	0.021

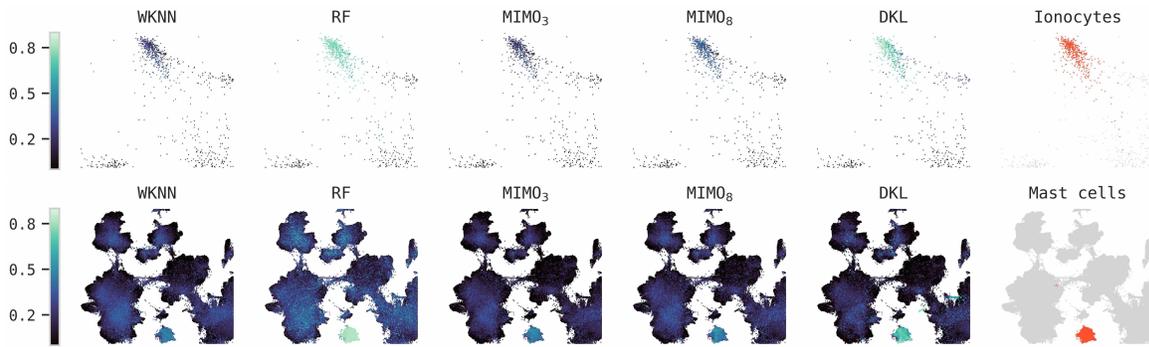


Figure 4. Predictive uncertainty for all models on left out cell types (rightmost panels).

Uncertainty Quantification for Atlas-Level Cell Type Transfer

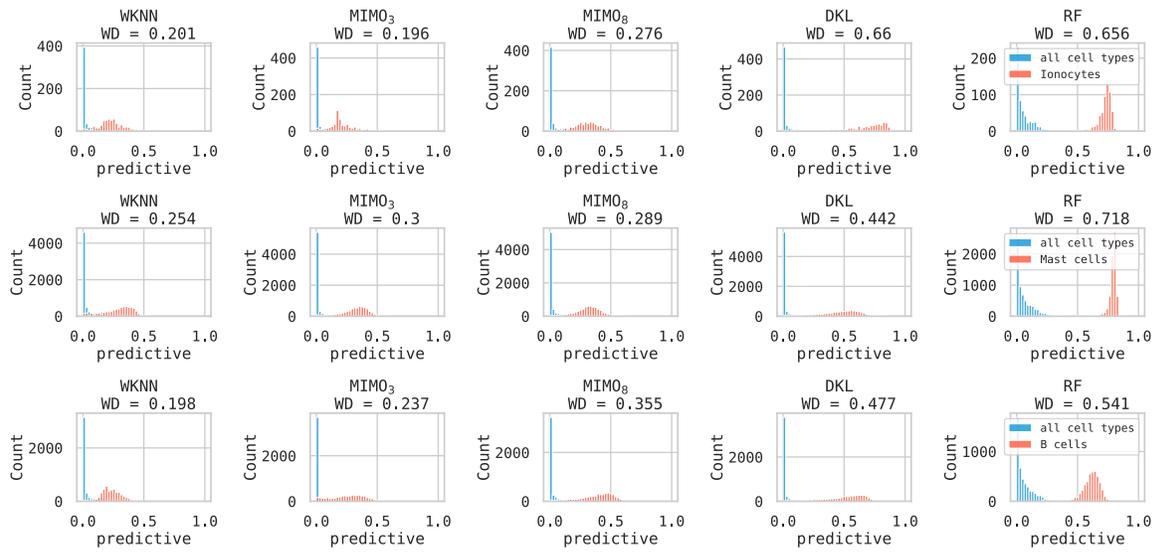


Figure 5. Distribution of predictive uncertainty for the cell types the model was trained on (blue) and the left out cell type (orange). The computed Wasserstein distance, WD, (\uparrow) between both distributions is shown in the title. This quantifies the models' ability to distinguish the OOD cell type from the ID training data. All models outperform the WKNN baseline. While RF shows the largest WD, using this model cannot be recommended due to poor performance in classification accuracy.

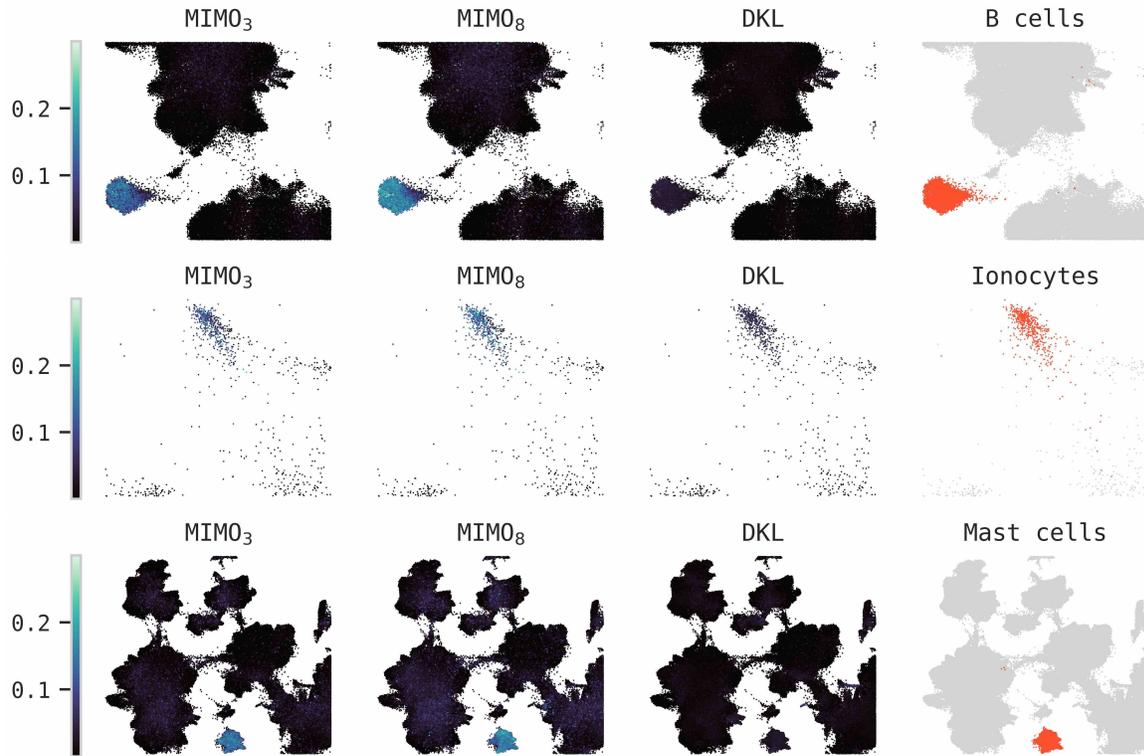
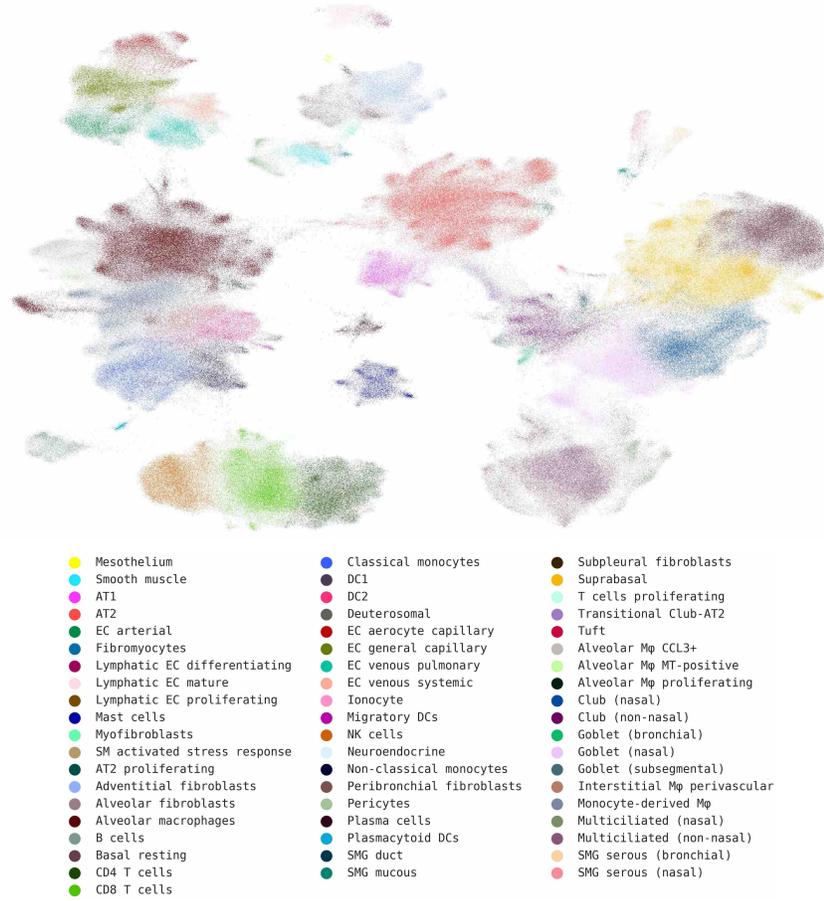


Figure 6. Model uncertainty for MIMO and DKL on left out cell types (rightmost panels).

Finest cell type annotation level



Leave-out dataset Jain Misharin 2021

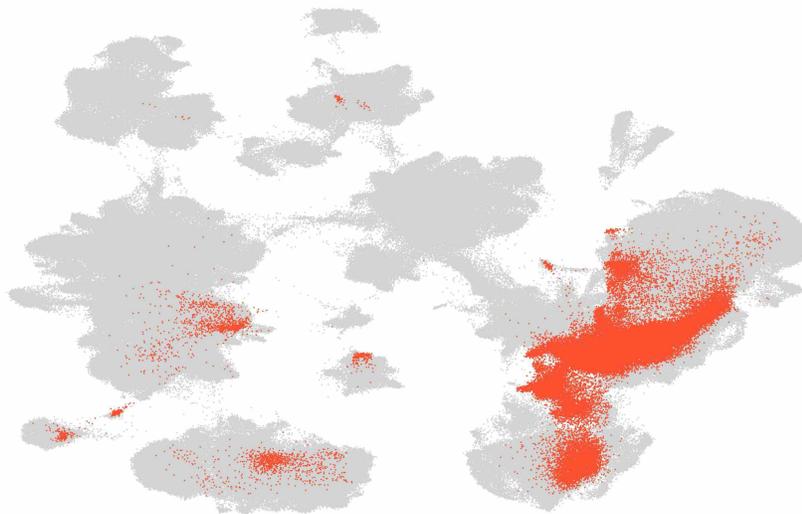


Figure 7. Human Lung Cell Atlas UMAPs colored by cell types (top) and leave-out dataset (bottom).