

ALBRT: Cellular Composition Prediction in Routine Histology Images

Muhammad Dawood¹, Kim Branson², Nasir M. Rajpoot¹, Fayyaz ul Amir Afsar Minhas¹

¹Department of Computer Science University of Warwick, Coventry, UK

² Artificial Intelligence & Machine Learning, GlaxoSmithKline

{Muhammad.Dawood, N.M.Rajpoot, Fayyaz.Minhas}@warwick.ac.uk, kim.m.branson@gsk.com

Abstract

Cellular composition prediction, i.e., predicting the presence and counts of different types of cells in the tumor microenvironment from a digitized image of a Hematoxylin and Eosin (H&E) stained tissue section can be used for various tasks in computational pathology such as the analysis of cellular topology and interactions, subtype prediction, survival analysis, etc. In this work, we propose an image-based cellular composition predictor (ALBRT) which can accurately predict the presence and counts of different types of cells in a given image patch. ALBRT, by its contrastive-learning inspired design, learns a compact and rotation-invariant feature representation that is then used for cellular composition prediction of different cell types. It offers significant improvement over existing state-of-the-art approaches for cell classification and counting. The patch-level feature representation learned by ALBRT is transferrable for cellular composition analysis over novel datasets and can also be utilized for downstream prediction tasks in CPath as well. The code and the inference web-server for the proposed method are available at the URL: <https://github.com/engrodawood/ALBRT>.

1. Introduction

Predicting the presence of different types of cells and quantifying their counts in digitized images of Hematoxylin and Eosin (H&E) stained tissue slides can be useful for various downstream prediction tasks in computational pathology (CPath) such as survival prediction [24, 31, 16], prognosis [34, 25], recurrence prediction [14], gene expression and biological process analysis [13, 35]. Cellular composition can give insights into cellular diversity in the tumor microenvironment (TME) [9] as well as tissue organization in the whole slide image (WSI) which can be helpful in therapeutic decision making [23]. For instance, a high proportion of stromal tumor-infiltrating lymphocytes (sTILs) are associated with good prognosis [18, 7, 8], whereas excessive tumor-associated macrophages infiltration signifies

poor prognosis in breast, bladder, and cervix carcinomas [4, 21]. Similarly, tumors with co-presence of immune cells and tumor cells in the TME can be considered potential candidates for immunotherapy [2, 9]. Despite its usefulness, it is not feasible for pathologists to perform such a quantitative analysis on whole slide images and across multiple cases due to the amount of effort and time required for this purpose. As a consequence, computational approaches for patch-level cellular composition analysis in WSIs are needed. However, computational prediction of the presence and quantification of different cell types in a WSI patch of H&E stained tissue sections can be challenging due to inter-class phenotype similarities and intra-class variations. Moreover, presence of overlapping cells and diffused background complicates the prediction problem further.

Existing methods for cellular composition prediction can be divided into two major classes: segmentation-based, and regression-based methods. Segmentation-based methods first segment different type of cells in a patch and then obtain cellular counts through post-processing. For example, watershed transform has been used for red and white blood cell segmentation and counting from microscopy images [27]. Similar methods have been used for single cell segmentation in microscopy images of budding yeast colonies using immersion simulation based self-organizing transform [32]. Segmentation-based methods require accurate cellular boundaries annotation, or in some cases, scribbled annotations [20, 17] during training which can be laborious, error-prone and time consuming. Moreover, state-of-the-art cell segmentation methods such as HoVer-Net [12], fail to learn a compact patch-level representation from images which can be useful for downstream prediction tasks in CPath.

To overcome the pixel-level annotation bottleneck, several regression-based methods have been proposed for cellular composition prediction. Regression-based methods use patch-level cell-type counts as target labels for cellular composition prediction. For instance, recently a convolutional neural network (CNN) has been proposed for counting cells in microscopic images [19]. Similarly, Cohen et al. proposed a deep neural network called Count-ception, which

uses overlapping stride for cell counting [28]. Beside this, density maps [22], and manifold learning [33] have been used for counting embryonic, and bacterial cells from light microscopy images. Recently, a manifold regularized network was proposed for counting and localizing cells in histology and microscopy images [36]. However, these methods are restricted to predict the counts of a single cell type having uniform shape.

In this paper, we propose a multi-headed CNN-based method that uses patch and sub-patch level cellular count during training, and predicts the presence and counts of different type of cells in a patch. The main contributions of the proposed work are listed below:

- The proposed method harnesses patch and sub-patch level cellular count information during training to predict the presence and counts of different type of cells in a patch.
- The proposed method explicitly models rotation invariance in its design through a contrastive learning inspired learning mechanism.
- We have assessed the generalization performance of the proposed method on an independent set and demonstrated the transferability of the learned feature representation for cellular composition prediction over a dataset with different cell types and annotation structure in comparison to the dataset used for training ALBRT.
- The proposed method outperforms state of the art segmentation-based method for cellular composition prediction.
- We demonstrate that the learned cellular representation is compact, rotation invariant, and can capture inter- and intra-class variations. This representation can be useful for other downstream prediction tasks in CPath.
- The source code and dataset used in this work along with a webserver for patch-level cellular composition analysis are available at the URL: <https://github.com/engrodawood/ALBRT>.

2. Materials and Methods

The workflow of the proposed approach for patch level cellular composition prediction is shown in Figure 1. The proposed model takes a patch of size 256×256 pixels from H&E WSI as input and predicts its cellular composition, i.e., counts of different type of cells in the input patch. The counts of different type of cells predicted by the model are subsequently used for detecting presence or absence of a particular cell type in a patch. Below we provide details about

the datasets used in the study, preprocessing, and the proposed model for cellular composition prediction.

2.1. Datasets and Preprocessing

We used two publicly available datasets: PanNuke [10, 11] and NuCLS [3]. For model training and validation, PanNuke breast cancer data of four major cell types (Neoplastic, Inflammatory, Connective and Epithelial cells) available in three pre-defined folds was used. The dataset consisting of 2, 351 patches of size 256×256 pixels at a spatial resolution of 0.25 microns-per-pixel (MPP). For each patch, nuclear segmentation masks outlining cellular boundaries of different type of cells in a patch are available. In total, there are 50, 063 annotated nuclei (19, 900 neoplastic, 6, 160 inflammatory, 10, 266 connective and 13, 737 epithelial cells). We compute patch and sub-patch level cellular counts from the nuclear segmentation masks and used those as target labels for cellular composition prediction.

For assessing model generalization, we used NuCLS dataset [3] consisting of 1, 744 field of views (FOVs) partitioned into 5 pre-defined train and validation folds. FOVs from the corrected single-rater NuCLS dataset are of varying sizes with a spatial resolution of 0.20 MPP. To match the resolution of these patches with PanNuke, we extracted a fixed patch of size 320×320 pixels from each FOV and resized it to 256×256 pixels. Patches smaller than 320×320 pixels are excluded. The dataset provides bounding box annotation of various types of nuclei in a patch. The annotated nuclei were grouped into 3 super classes (Tumor, Stromal, and sTILs) and 6 sub-classes (Tumor, Mitotic Figure, Stromal, Macrophages, Lymphocytes, and plasma cells). The pre-processed dataset consists of 33, 205 annotated nuclei, which include 12, 581 Tumor, 145 Mitotic Figure, 5, 949 Stromal, 756 Macrophages, 9, 850 Lymphocytes, and 3, 924 plasma cells. We calculated patch-level cellular counts of sub-class and super-class cell type by counting cells that lie in a fixed 320×320 window and used the counts as target labels for model performance evaluation.

For both PanNuke and NuCLS datasets, cells with the majority of their annotated nuclei outside the image patch are excluded from the target count. In addition to cellular counts for each patch, patch-level binary labels indicating the presence or absence of different cell types are also constructed, i.e., if the count of a specific cell type in a patch is greater than zero, the patch would be labelled positive for that cell type, otherwise, negative. These binary cell type labels are used to analyze the sensitivity of various predictors to correctly detect the presence of different types of cells in a given patch.

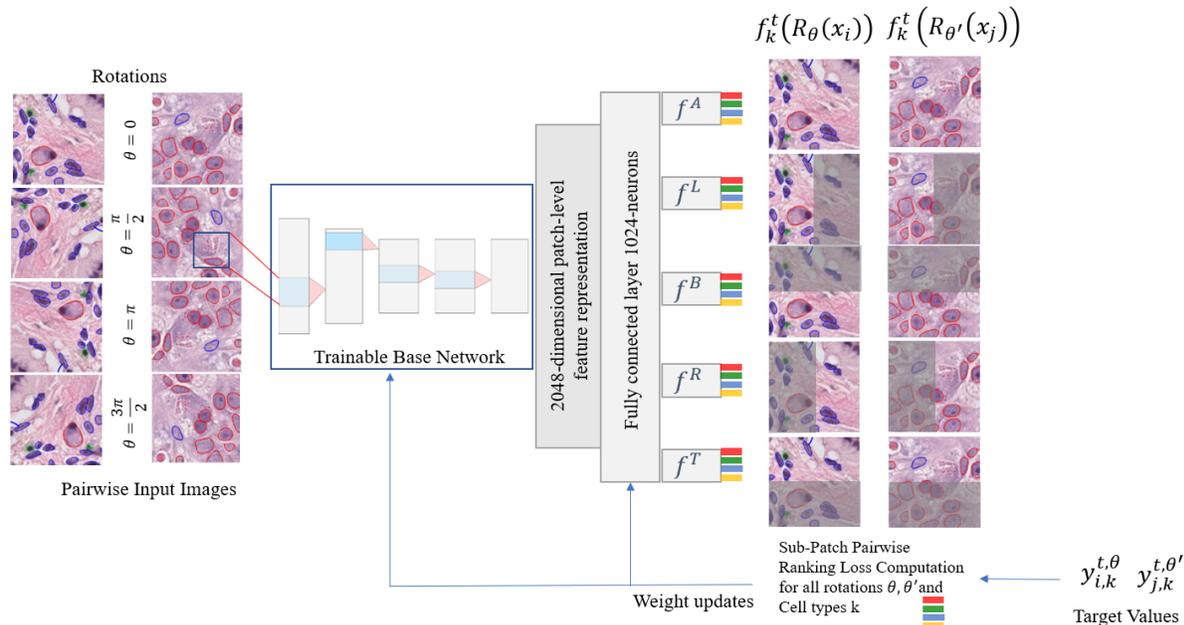


Figure 1. Workflow of the proposed approach for cellular composition prediction. The model accepts original images along with their rotated variants and predicts cellular counts in different region of the input image. The model has 5-multi-output branches. The first branch (A) predicts the counts of different type of cells in the whole of the input patch, whereas the remaining branches (L,B,R,T) predict the cellular counts in the left, bottom, right, and top halves of the input patch, respectively.

2.2. Proposed Model for Cellular Composition Prediction

To model cellular composition prediction as a learning problem, consider a training dataset $D = \{(x_i, y_i) | i = 1 \dots N\}$ consisting of N patches and their associated cellular counts. Here, x_i represents an RGB image patch and $y_i \in \mathcal{R}^K$ represents the corresponding vector of cellular counts of K different cell types, e.g., neoplastic, inflammatory, connective, and epithelial cells. The objective of the learning task is to train a predictor with a set of learnable parameters W such that the output of the predictor function $z_i = f(x_i; W)$ for K cell types matches their true counts for test images. The predictor function was subsequently also used as a classification score for detecting the presence or absence of a particular cell type k in a patch.

We propose a mutli-headed convolutional neural network that uses patch and sub-patch level cellular counts for image based cellular composition prediction, as shown in Figure 1. Our hypothesis is that learning sub-patch level features and invariances can lead to precise prediction at the patch level. The proposed model builds on this hypothesis and consists of five multi-output branches with each branch predicting cellular composition within a specific region of the input image patch. The first branch (A) predicts the counts of different type of cells in the whole of the input patch, whereas the remaining branches (L,B,R,T) predict cellular counts in the left, bottom, right, and top halves of the input patch, respec-

tively. During training, patch and sub-patch level counts of different type of cells in each image are compared with other images and also with their rotated (90° , 180° , 270°) variants through a pairwise ranking-based loss. The proposed model uses Xception network [5] pretrained on ImageNet [29] as a base convolutional network and the flattened output of its convolutional layers is passed onto a fully connected layer consisting of 1024 neurons followed by the five branches for predicting cellular counts in different regions of the input patch. The key motivation behind selecting the Xception network as a base network is that it uses depthwise separable convolution. Since raw patches contain cells of different sizes, depthwise separable convolution with its multiple kernels is expected to capture cell size diversity.

In order to formulate the learning problem, let $f_k^t(R_\theta(x_i))$ be the predicted cell counts of branch $t \in \{A, L, B, R, T\}$ for cell type k ($k=1, \dots, K$) for an input patch x_i after it has been rotated by an angle $\theta \in \{0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}\}$. Let $y_{i,k}^{t,\theta}$ be the corresponding target counts. We use pairwise ranking loss function $l_k()$ as in Equation 1, which compares the predicted cellular composition of images in corresponding branches with their true counts across multiple rotations and sub-patches. The overall loss function $L(f; D)$ in Equation 2 performs pairwise comparison of all patches in the training dataset across all cell types in a batch. For efficient implementation, the computation of the loss function is vectorized at the batch level.

$$l_k(f; (x_i, y_{ik}), (x_j, y_{jk})) = \sum_{t \in \{A, L, B, R, T\}} \sum_{\theta} \sum_{\theta'} \begin{cases} \max(0, 1 - (y_{ik}^{t\theta} - y_{jk}^{t\theta'})(f_k^t(R_{\theta}(x_i)) - f_k^t(R_{\theta'}(x_j)))) & \text{if } y_{ik}^{t\theta} \neq y_{jk}^{t\theta'} \\ (f_k^t(R_{\theta}(x_i)) - f_k^t(R_{\theta'}(x_j)))^2 & \text{otherwise} \end{cases} \quad (1)$$

$$L(f; D) = \sum_{k=1}^K \sum_{i=1}^N \sum_{j=1}^N l_k(f; (x_i, y_{ik}), (x_j, y_{jk})) \quad (2)$$

Minimization of this loss function penalizes the number of mis-ranked examples and attempts to rank patches and sub-patches correctly based on their relative cellular counts. In line with the concept of contrastive learning, it enables the network to learn rotational invariance and sub-patch level translation invariance. Due to the use of a pairwise ranking loss, the predicted cellular counts of the network will be relative and are converted into absolute counts by simple curve-fitting over predicted counts using the ground truth training data.

2.3. Model training and evaluation

We evaluated the performance of the proposed model using 3-fold cross-validation with pre-defined folds in the PanNuke dataset, i.e., data of one-fold was held out for testing and training was performed on the remaining two folds. We assessed the model’s regression performance on the test fold using mean absolute error (MAE) and Spearman correlation (SCC) as performance metrics. In order to analyze the model sensitivity in detecting presence of a particular cell type, the predicted count of that cell type was used as prediction score and compared with the corresponding actual binary label (presence or absence) for that patch using the Area under the Receiver Operating Characteristic Curve (AUROC). In each cross validation run, the proposed model was trained for 15 epochs using adaptive momentum-based optimizer [15] with a batch size of 32 and initial learning rate of 0.001. However, the learning rate by reduced by a factor of 10 after five successive epochs. The experiments were performed using Keras [6] and TensorFlow [1] deep learning libraries.

2.4. Baseline Method and Comparison

We compared the performance of proposed method with the state-of-the-art nuclear segmentation and classification method HoVer-Net [12]. HoVer-Net is a multi-headed CNN that uses set of residual units as backbone feature extractor followed by three branches for cells segmentation and classification. We assessed the performance of HoVer-Net on PanNuke dataset and used their results as baseline for the proposed method.

2.5. Transferability of Cellular Feature Representation

In order to assess whether the 2048-dimensional patch-level feature representation (see Figure 1) learned by the proposed model using the PanNuke dataset can be employed for predicting cellular counts for a different dataset with novel cell types and annotation structure, we used the NuCLS dataset which contains cellular count information for multiple cell-types divided into super and sub-classes. Since cell types in NuCLS dataset are different from PanNuke, we finetuned only the fully connected layers of ALBRT main branch (A) on super-class and sub-class cellular counts. Furthermore, to assess the ability of the learned feature representation to inherently model rotational invariance, no additional augmentations were used. As a baseline for comparison, we used the Xception model pretrained over ImageNet [29] with both mean-squared error and pairwise ranking losses without any rotational augmentations.

3. Results and discussion

3.1. Quantitative results

Table 1 shows the performance of the proposed approach for cellular composition prediction in terms of mean MAE and spearman correlation coefficient. From the table, it can be seen that the counts of all cell types listed in the table are predicted with an average MAE of less than 1.78, and a Spearman correlation higher than 0.77. Moreover, for counting epithelial and inflammatory cells, the MAE is around 1.08, while for neoplastic, and connective cell the MAE is 1.78 and 1.71, respectively. Similarly, neoplastic, inflammatory, connective and epithelial cells are predicted with a higher Spearman correlation of 0.94, 0.77, 0.85 and 0.93, respectively. Figure 2 shows model predictions for a set of patches. The figure visualizes input patches with cellular boundaries outlined using different colors along with their true and predicted cellular counts. From the figure, it can be seen that the difference between the true counts and predicted counts is minimal for all test images.

3.2. Comparison with HoVer-Net

Table 1 shows HoVer-Net prediction performance on PanNuke dataset in terms of mean MAE and Spearman corre-

Table 1. HoVer-Net and ALBRT cellular composition prediction and classification performance comparison for different types of cells present in PanNuke dataset. We report mean absolute error (MAE), Spearman Correlation coefficient (SCC) and AUROC. The numbers in parenthesis show the standard deviation across test folds.

Method	Metric	Neoplastic	Inflammatory	Connective	Epithelial
HoVer-Net	MAE	2.34 (0.17)	1.16 (0.06)	1.88 (0.04)	1.55 (0.25)
	SCC	0.89 (0.02)	0.72 (0.00)	0.79 (0.03)	0.84 (0.48)
	AUROC	0.93 (0.01)	0.82 (0.00)	0.91 (0.00)	0.97 (0.01)
ALBRT	MAE	1.78 (0.09)	1.08 (0.03)	1.71 (0.04)	1.08 (0.09)
	SCC	0.94 (0.01)	0.77 (0.01)	0.85 (0.01)	0.93 (0.02)
	AUROC	0.98 (0.00)	0.88 (0.01)	0.94 (0.00)	0.99 (0.00)

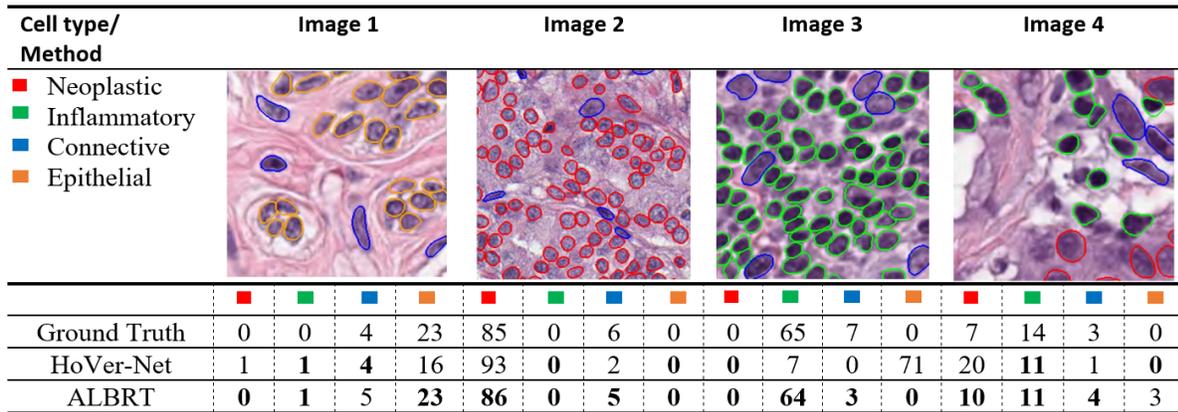


Figure 2. Comparison of ALBRT and HoVer-Net prediction with ground truth for a set of test images. Images are overlaid with annotated cellular boundaries. Below each image true cellular counts, HoVer-Net and ALBRT predicted cellular counts for different types of cells are shown.

lation coefficient. From the table, it can be seen that for all cell types the MAE of HoVer-Net is higher in comparison to ALBRT. Similarly, the Spearman correlation of ALBRT for all cell types is also higher than HoVer-Net. It is important to mention here that ALBRT uses only cellular count information and does not need precise cellular boundaries annotation during training which can be difficult and time consuming to obtain.

3.3. Saliency Map Analysis

We computed gradient-based saliency maps to visualize regions of an input patch that contribute to cellular composition prediction. We passed set of test images as input to a pretrained model and cellular counts of each cell type as a target and extracted the saliency maps using Grad-CAM [30] as shown in Figure 3. From the figure, it can be seen that for each cell type the extracted saliency map correlates with input patch regions where a specific type of cell is present. This clearly shows that the representation learned by ALBRT truly captures the information that contributes to cellular composition.

3.4. Cell Type presence prediction

Figure 4 shows the AUROC comparison of HoVer-Net and ALBRT. From the figure, it can be seen that the AUROC of ALBRT is significantly higher compared to HoVer-Net. This clearly shows that the proposed approach is capable of detecting the presence of a particular type of cell in a patch and can be used for both detecting the presence of different types of cells as well as counting cells of each type in the patch.

3.5. ALBRT captures interclass and intraclass cellular variability

An interesting feature of ALBRT is that it captures intra-class phenotype similarities and inter-class phenotype variations. To illustrate this, we used the 1024-dimensional output of the final fully connected layer in the network for a given patch as a latent representation and fit an unsupervised Uniform Manifold Approximation and Projection (UMAP) [26] as shown in Figure 2. From the figure, it can be seen that the learned latent representation is capable of classifying cells into different classes. Additionally, it can also cluster patches based on their cell count similarity in a given class.

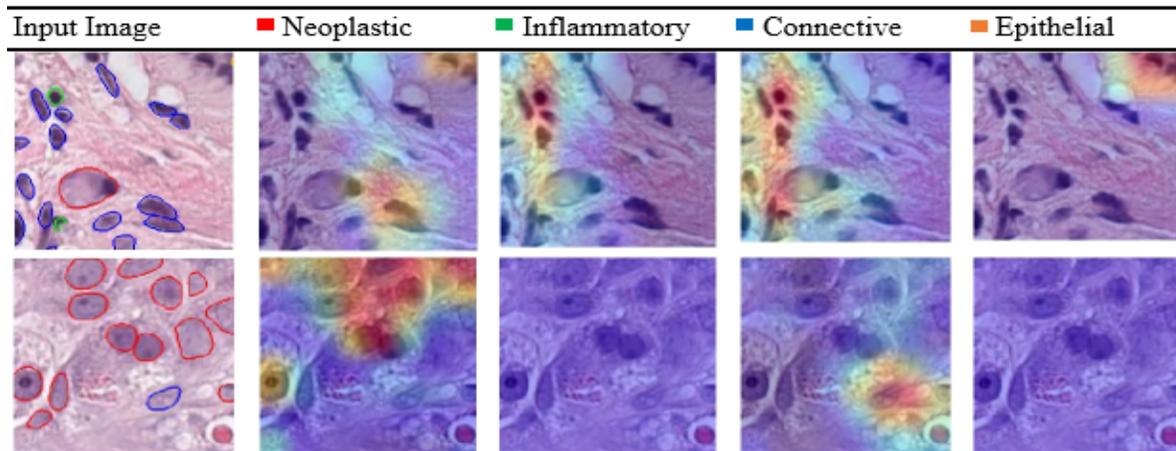


Figure 3. Saliency maps for two text patches. The first column shows the input images overlaid with annotated cellular boundaries in different colors. The other four columns show the saliency map of each cell type obtained using Grad-CAM

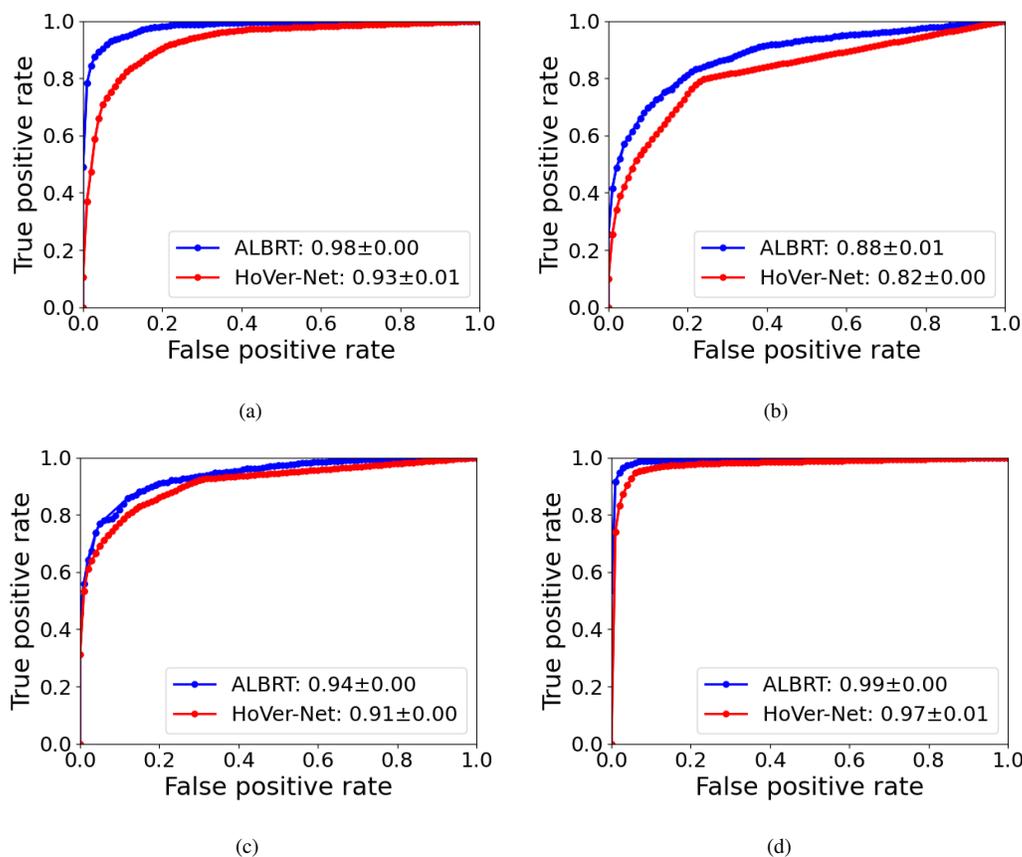


Figure 4. AUROC Comparisons of HoVer-Net and ALBRT for different types of cell present in PanNuke dataset. (a) Neoplastic (b) Inflammatory (c) Connective (d) Epithelial

This clearly shows that the latent representation learned by ALBRT truly contribute to cellular composition and cell type classification. Moreover, the representation can be used for other downstream classification and regression problems in CPath.

3.6. Transferability of ALBRT’s Cellular Feature Representation

Tables 2 and 3 show ALBRT’s generalization performance comparison with ALBRT using ImageNet weights for base network, and Xception network pretrained on ImageNet for

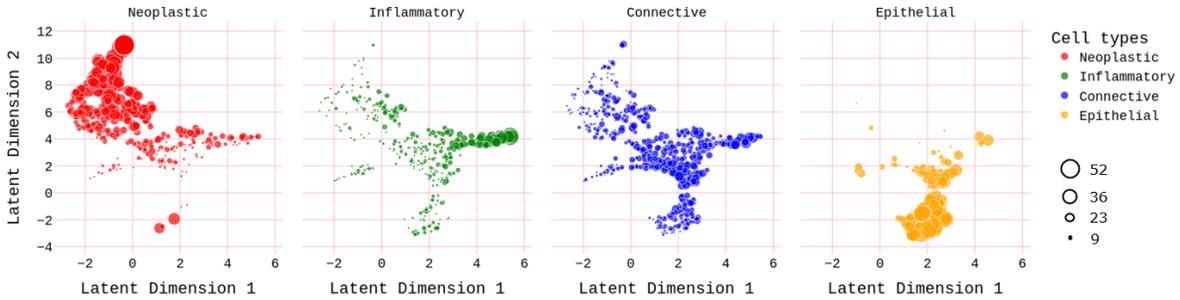


Figure 5. UMAP plots: each blob represents the latent representation of a given input image. The plot on the left shows the latent representation of Neoplastic cells, followed by inflammatory cells, and then connective cells and the final one shows representation of Epithelial cells.

Table 2. Results of cellular composition prediction using an Xception model with different patch-level feature representations over NuCLS dataset super classes in terms of mean absolute error (MAE), Spearman Correlation coefficient (SCC), and AUROC. The numbers in parenthesis show the standard deviation across test folds. The patch-level feature representation produced by the Xception network trained under the proposed ALBRT framework over the PanNuke dataset exhibits better prediction performance over the NuCLS dataset in comparison to baseline methods indicating transferability of the learned feature representation.

Method	Metric	Tumor	Stromal	sTILs
ImageNet based features with MSE Loss (Baseline)	MAE	5.17 (0.47)	3.95 (0.43)	5.37 (0.34)
	SCC	0.66 (0.02)	0.50 (0.04)	0.69 (0.07)
	AUROC	0.84 (0.02)	0.76 (0.04)	0.81 (0.04)
ImageNet based features with Pairwise-Ranking Loss	MAE	4.95 (0.72)	3.65 (0.48)	4.38 (0.71)
	SCC	0.64 (0.03)	0.51 (0.02)	0.75 (0.05)
	AUROC	0.84 (0.03)	0.77 (0.03)	0.84 (0.03)
ALBRT features with Pairwise-Ranking Loss	MAE	3.51 (0.65)	3.15 (0.37)	3.54 (0.79)
	SCC	0.79 (0.05)	0.64 (0.05)	0.85 (0.05)
	AUROC	0.91 (0.03)	0.83 (0.04)	0.92 (0.03)

super-class, and sub-class cell types. From the tables, it can be seen that for both super-class and sub-class cell types, ALBRT-based feature representation outperforms the baseline Xception model (with ImageNet pre-training and MSE or pairwise-ranking losses) in all performance metrics. For instance, ALBRT-based feature representation predicts the cellular counts of Tumor and sTILs with a mean Spearman correlation of 0.79, and 0.85, whereas ImageNet-based features give a correlation of 0.64, and 0.75. This shows that the feature representation learned by ALBRT is transferable for cellular composition prediction on other datasets. Furthermore, ALBRT using pretrained ImageNet weights for base network outperforms Xception network that uses MSE loss for finetuning. This shows that the loss function used for model training is a good choice for cellular composition prediction compared to MSE.

4. Conclusions and Future work

In this work, we have proposed an image-based cellular composition predictor that accurately predicts the presence

and counts of different types of cells in a given patch. Overcoming the shortcoming of segmentation-based cellular composition prediction approaches of using precise pixel-level cellular boundaries annotation, we proposed a ranking-based model that learns compact, and rotation-invariant representation from patches using patch and sub-patch level cellular counts information during training. We have assessed the transferability of the learned representation using the proposed method on an external dataset with novel cell types and annotation structure. Our analysis show that the learned representation is transferrable to other dataset and can be used for various downstream prediction tasks in CPath. In future, we plan to do more detailed comparative analysis on any other cellular composition datasets as well as extending the approach for downstream predictive tasks such as survival analysis.

Acknowledgments

MD would like to acknowledge the PhD studentship support from GlaxoSmithKline. FM and NR supported by the

Table 3. Results of cellular composition prediction with different patch-level feature representations over NuCLS dataset sub-classes in terms of mean absolute error (MAE), Spearman Correlation coefficient (SCC), and AUROC. The numbers in parenthesis show the standard deviation across test folds. The patch-level feature representation produced by the proposed ALBRT framework trained over the PanNuke dataset exhibits better prediction performance for the NuCLS dataset in comparison to baseline methods indicating transferability of the learned feature representation to novel datasets.

Method	Metric	Tumor	Mitotic	Stromal	Macrophage	Lymphocyte	Plasma cells
ImageNet based features with MSE Loss (Baseline)	MAE	5.05 (0.46)	0.26 (0.01)	3.39 (0.45)	1.17 (0.18)	5.43 (0.87)	5.06 (0.74)
	SCC	0.67 (0.02)	0.09 (0.07)	0.54 (0.06)	0.22 (0.05)	0.58 (0.07)	0.21 (0.12)
	AUROC	0.84 (0.03)	0.63 (0.11)	0.77 (0.05)	0.71 (0.06)	0.77 (0.04)	0.63 (0.08)
ImageNet based features with Pairwise-Ranking Loss	MAE	4.90 (0.76)	0.16 (0.09)	3.14 (0.57)	0.88 (0.17)	4.83 (1.3)	3.67 (1.25)
	SCC	0.64 (0.03)	0.02 (0.04)	0.55 (0.02)	0.16 (0.02)	0.62 (0.07)	0.28 (0.08)
	AUROC	0.85 (0.03)	0.54 (0.07)	0.78 (0.02)	0.65 (0.03)	0.79 (0.03)	0.68 (0.05)
ALBRT features with Pairwise-Ranking Loss	MAE	3.47 (0.64)	0.15 (0.08)	2.74 (0.48)	0.78 (0.18)	3.81 (1.11)	3.02 (1.05)
	SCC	0.79 (0.05)	0.16 (0.04)	0.67 (0.04)	0.33 (0.07)	0.74 (0.07)	0.51 (0.03)
	AUROC	0.91 (0.04)	0.70 (0.04)	0.84 (0.03)	0.80 (0.06)	0.88 (0.04)	0.83 (0.02)

PathLAKE digital pathology consortium which is funded from the Data to Early Diagnosis and Precision Medicine strand of the government’s Industrial Strategy Challenge Fund, managed and delivered by UK Research and Innovation (UKRI).

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org. 4
- [2] Khalid AbdulJabbar, Shan E Ahmed Raza, Rachel Rosenthal, Mariam Jamal-Hanjani, Selvaraju Veeriah, Ayse Akarca, Tom Lund, David A Moore, Roberto Salgado, Maise Al Bakir, et al. Geospatial immune variability illuminates differential evolution of lung adenocarcinoma. *Nature Medicine*, 26(7):1054–1062, 2020. 1
- [3] Mohamed Amgad, Lamees A Atteya, Hagar Hussein, Kareem Hosny Mohammed, Ehab Hafiz, Maha AT Elsebaie, Ahmed M Alhusseiny, Mohamed Atef AlMoslemany, Abdelmagid M Elmatboly, Philip A Pappalardo, et al. Nucls: A scalable crowdsourcing, deep learning approach and dataset for nucleus classification, localization and segmentation. *arXiv preprint arXiv:2102.09099*, 2021. 2
- [4] L Bingle, NJ Brown, and Claire E Lewis. The role of tumour-associated macrophages in tumour progression: implications for new anticancer therapies. *The Journal of Pathology: A Journal of the Pathological Society of Great Britain and Ireland*, 196(3):254–265, 2002. 1
- [5] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017. 3
- [6] Francois Chollet et al. Keras, 2015. 4
- [7] Carsten Denkert, Gunter von Minckwitz, Silvia Darb-Esfahani, Bianca Lederer, Barbara I Heppner, Karsten E Weber, Jan Budczies, Jens Huober, Frederick Klauschen, Jenny Furlanetto, et al. Tumour-infiltrating lymphocytes and prognosis in different subtypes of breast cancer: a pooled analysis of 3771 patients treated with neoadjuvant therapy. *The lancet oncology*, 19(1):40–50, 2018. 1
- [8] MV Dieci, P Conte, G Bisagni, AA Brandes, A Frassoldati, L Cavanna, A Musolino, F Giotta, A Rimanti, O Garrone, et al. Association of tumor-infiltrating lymphocytes with distant disease-free survival in the shorter randomized adjuvant trial for patients with early her2+ breast cancer. *Annals of Oncology*, 30(3):418–423, 2019. 1
- [9] Filippo Galli, Jesus Vera Aguilera, Belinda Palermo, Svetomir N Markovic, Paola Nisticò, and Alberto Signore. Relevance of immune cell and tumor microenvironment imaging in the new era of immunotherapy. *Journal of Experimental & Clinical Cancer Research*, 39:1–21, 2020. 1
- [10] Jevgenij Gamper, Navid Alemi Koohbanani, Ksenija Benet, Ali Khuram, and Nasir Rajpoot. Pannuke: an open pan-cancer histology dataset for nuclei instance segmentation and classification. In *European Congress on Digital Pathology*, pages 11–19. Springer, 2019. 2
- [11] Jevgenij Gamper, Navid Alemi Koohbanani, Simon Graham, Mostafa Jahanifar, Syed Ali Khuram, Ayesha Azam, Katherine Hewitt, and Nasir Rajpoot. Pannuke dataset extension, in-

- sights and baselines. *arXiv preprint arXiv:2003.10778*, 2020. [2](#)
- [12] Simon Graham, Quoc Dang Vu, Shan E Ahmed Raza, Ayesha Azam, Yee Wah Tsang, Jin Tae Kwak, and Nasir Rajpoot. Hover-net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images. *Medical Image Analysis*, 58:101563, 2019. [1](#), [4](#)
- [13] Bryan He, Ludvig Bergenstr hle, Linnea Stenbeck, Abubakar Abid, Alma Andersson,  ke Borg, Jonas Maaskola, Joakim Lundeberg, and James Zou. Integrating spatial gene expression and breast tumour morphology via deep learning. *Nature biomedical engineering*, 4(8):827–834, 2020. [1](#)
- [14] Meng-Yao Ji, Lei Yuan, Xiao-Da Jiang, Zhi Zeng, Na Zhan, Ping-Xiao Huang, Cheng Lu, and Wei-Guo Dong. Nuclear shape, architecture and orientation features from h&e images are able to predict recurrence in node-negative gastric adenocarcinoma. *Journal of translational medicine*, 17(1):1–12, 2019. [1](#)
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [4](#)
- [16] Sung Min Ko, Janghee Lee, Soong June Bae, Su Jung Baik, Junghwan Ji, Dooreh Kim, Sung Gwe Ahn, and Joon Jeong. Body mass index and absolute lymphocyte count predict disease-free survival in korean breast cancer patients. *British journal of cancer*, pages 1–7, 2021. [1](#)
- [17] Navid Alemi Koohbanani, Mostafa Jahanifar, Neda Zamani Tajadin, and Nasir Rajpoot. Nuclick: A deep learning framework for interactive segmentation of microscopic images. *Medical Image Analysis*, 65:101771, 2020. [1](#)
- [18] Hajime Kuroda, Tsengelmaa Jamiyan, Rin Yamaguchi, Akinari Kakumoto, Akihito Abe, Oi Harada, and Atsuko Masunaga. Tumor-infiltrating b cells and t cells correlate with postoperative prognosis in triple-negative carcinoma of the breast. *BMC cancer*, 21(1):1–10, 2021. [1](#)
- [19] Falko Lavitt, Demi R Rijlaarsdam, Dennet van der Linden, Ewelina Weglarz-Tomczak, and Jakub M Tomczak. Deep learning and transfer learning for automatic cell counting in microscope images of human cancer cell lines. *Applied Sciences*, 11(11):4912, 2021. [1](#)
- [20] Hyeonsoo Lee and Won-Ki Jeong. Scribble2label: Scribble-supervised cell segmentation via self-generating pseudo-labels with consistency. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 14–23. Springer, 2020. [1](#)
- [21] Russell D Leek, Claire E Lewis, Ruth Whitehouse, Michael Greenall, Jane Clarke, and Adrian L Harris. Association of macrophage infiltration with angiogenesis and prognosis in invasive breast carcinoma. *Cancer research*, 56(20):4625–4629, 1996. [1](#)
- [22] Victor Lempitsky and Andrew Zisserman. Learning to count objects in images. *Advances in neural information processing systems*, 23:1324–1332, 2010. [2](#)
- [23] Yuxin Lin, Jianxin Xu, and Huiyin Lan. Tumor-associated macrophages in tumor metastasis: biological roles and clinical therapeutic applications. *Journal of hematology & oncology*, 12(1):1–16, 2019. [1](#)
- [24] Cheng Lu, David Romo-Bucheli, Xiangxue Wang, Andrew Janowczyk, Shridar Ganesan, Hannah Gilmore, David Rimm, and Anant Madabhushi. Nuclear shape and orientation features from h&e images predict survival in early-stage estrogen receptor-positive breast cancers. *Laboratory investigation*, 98(11):1438–1448, 2018. [1](#)
- [25] Alberto Mantovani, Federica Marchesi, Alberto Malesci, Luigi Laghi, and Paola Allavena. Tumour-associated macrophages as treatment targets in oncology. *Nature reviews Clinical oncology*, 14(7):399–416, 2017. [1](#)
- [26] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018. [5](#)
- [27] Ana Carolina Borges Monteiro, Yuzo Iano, and Reinaldo Padilha Fran a. Detecting and counting of blood cells using watershed transform: an improved methodology. In *Brazilian Technology Symposium*, pages 301–310. Springer, 2017. [1](#)
- [28] Joseph Paul Cohen, Genevieve Boucher, Craig A Glastonbury, Henry Z Lo, and Yoshua Bengio. Count-ception: Counting by fully convolutional redundant counting. In *Proceedings of the IEEE International conference on computer vision workshops*, pages 18–26, 2017. [2](#)
- [29] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. [3](#), [4](#)
- [30] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. [5](#)
- [31] Muhammad Shaban, Syed Ali Khurram, Muhammad Moazam Fraz, Najah Alsubaie, Iqra Masood, Sajid Mushtaq, Mariam Hassan, Asif Loya, and Nasir M Rajpoot. A novel digital score for abundance of tumour infiltrating lymphocytes predicts disease free survival in oral squamous cell carcinoma. *Scientific reports*, 9(1):1–13, 2019. [1](#)
- [32] Quanli Wang. A novel immersion simulation based self-organizing transform with application to single-cell segmentation from microscopy images, 2016. [1](#)
- [33] Yi Wang, Yuexian Zou, and Wenwu Wang. Manifold-based visual object counting. *IEEE Transactions on Image Processing*, 27(7):3248–3263, 2018. [2](#)
- [34] Haihua Yuan, Chengcheng Qian, Renjie Cai, Wenying Zhang, Jiongyi Wang, Xiaohua Hu, YANJIE ZHANG, Bin Jiang, and Feng Liu. Neutrophil-lymphocyte ratio and circulating tumor cells counts predict prognosis in gastrointestinal cancer patients. *Frontiers in Oncology*, 11:2586, 2021. [1](#)
- [35] Xiaohui Zhan, Jun Cheng, Zhi Huang, Zhi Han, Bryan Helm, Xiaowen Liu, Jie Zhang, Tian-Fu Wang, Dong Ni, and Kun Huang. Correlation analysis of histopathology and proteogenomics data for breast cancer. *Molecular & Cellular Proteomics*, 18(8):S37–S51, 2019. [1](#)

- [36] Yuxin Zheng, Zhao Chen, Yanfei Zuo, Xiaosong Guan, Zhaoyu Wang, and Xiao Mu. Manifold-regularized regression network: A novel end-to-end method for cell counting and localization. In *Proceedings of the 2020 the 4th International Conference on Innovation in Artificial Intelligence*, pages 121–124, 2020. [2](#)