2022 44th Annual International Conference of
the IEEE Engineering in Medicine & Biology Society (EMBC)
Scottish Event Campus, Glasgow, UK, July 11-15, 2022

2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC) | 978-1-7281-2782-8/22/$31.00 ©2022 IEEE | DOI: 10.1109/EMBC48229.2022.9871710

Multi-Scale Graphical Representation of Cell Environment

Helen Theissen[1], Tapabrata Chakraborty[1], Stefano Malacrino[1],
Daniel Royston[2,3] and Jens Rittscher[1,2]

***Abstract*— We present a multi-scale graphical network that can capture the relevant representations of individual cell morphology, topological structure of cell communities in a tissue image, as well as whole slide level attributes. This helps to effectively merge the disease relevant cell morphology to the overall topological context within the sample, within one unified deep framework. From the explainability point of view, instead of empirical design, the graphs are designed with biomedical considerations in mind in order to have translational validity. We also provide a clinically interpretable visualisation of the cells and their micro- and macro-environment by leveraging label noise reduction. We demonstrate the efficacy of our methodology on myeloproliferative neoplasms (MPN), a haematopoietic stem cell disorder as an exemplar test case. The proposed method achieves an encouraging performance in the robust separation of different MPN subtypes in this exciting new dataset as part of this work.**

I. Introduction

Characterising cells and the spatial architecture of cell networks in tissue samples is an open problem in the biomedical imaging community. Automated representations in this area can help in novel biomarker detection, drug development, clinical decision making, etc [1]. While progress based on deep learning and graphical models has been made, a number of challenges still need to be addressed before they can be used to guide human pathologists such as the following.

First, the representations of cells and the cellular network need to capture disease relevant features. Recent developments in tissue and cell characterisation based on graph neural networks have shown how graphs can combine local node features with structural relation of its neighbourhood. However, it remains unclear how the graphs themselves should be constructed and how distance thresholds are chosen [2] [3]. Here, we present a way to construct cell graphs based on biological foundations. Second, most works on cell graph representations are constrained to handcrafted features of nuclei and methods leveraging convolutional neural networks stay confined to a 2D or 3D data structure. While the latter have been successfully applied for applications such as cell classification [4] [5], few groups have tried to combine deep features of the entire cell with a topological representation of cell networks. To address this, we aim to combine a graph neural network model with a deep feature extractor trained for the task at hand at the cell level to increase the disease relevance of the features. Third, cell graphs have been exploited for disease prediction, but they rarely support interpretability through visualisation. Hence, we combine decision-making support on the slide level with a disease relevant spatial visualisation of the cell network.

Pati *et al.* designed a hierarchical cell-to-tissue graph neural network to represent a biopsy [6]. Like Jaume *et al.* [2] they used a set of handcrafted features extracted from nuclei to represent nodes and their corresponding features, and hence lacked the generalisation of task agnosticism and the data-driven ethos of learned features.

The contributions of this paper are twofold. First, we show that incorporating the immediate neighbourhood of a cell can reduces the spatial noise in the cell-level feature space. Instead of using nuclei features as most other works [7] [2] we detect and segment complete cells for feature extraction. Second, we merge low-level deep features extracted from cell images with the topological context of the cell population in a whole slide image. In addition to a prediction to distinguish disease subtypes, our approach provides a spatial visualisation of immediate clinical utility.

We demonstrate the potential of the proposed method in the context of the diagnosis of myeloproliferative neoplasms. These haematopoetic stem cell disorders (blood cancers) are driven by changes in the megakaryocyte (MK) cell population in the bone marrow, thus making them an optimal exemplar case. This is a new dataset introduced as part of this project.

II. Multi-scale Graphical Representation

We propose a multi-scale graph network to capture disease relevant representation of a cell population in a whole slide image consisting of: 1) a feature extractor at the cell level that provides disease relevant cell representations 2) a cellular neighbourhood graph yielding a local representation of the morphological and topological structure 3) a community detection mechanism to identify relevant subgraphs and 4) a slide-level classification.

A. Cell Features

We obtain a d-dimensional cell representation $x(v) \in \mathbb{R}^d$ by using a supervised setup to distinguish between MK cells stemming from reactive and diseased biopsies. The cells were extracted and segmented using the method described

[1]Institute of Biomedical Engineering (IBME) and the Big Data Institute (BDI), Dept. of Engineering Science, University of Oxford, Oxford, UK
[2]NIHR Oxford Biomedical Research Centre, Oxford University Hospitals NHS Foundation Trust, Oxford, Oxfordshire, UK
[3]Department of Cellular Pathology, John Radcliffe Hospital, Oxford University NHS Foundation Trust, Oxford, United Kingdom

*HT is funded by the Engineering and Physical Sciences Research Council (EPSRC) and Medical Research Council (MRC), grant number EP/L016052/1. TC is supported by the UKRI Innovate UK DART Porgramme and the Oxford CRUK Cancer Centre.

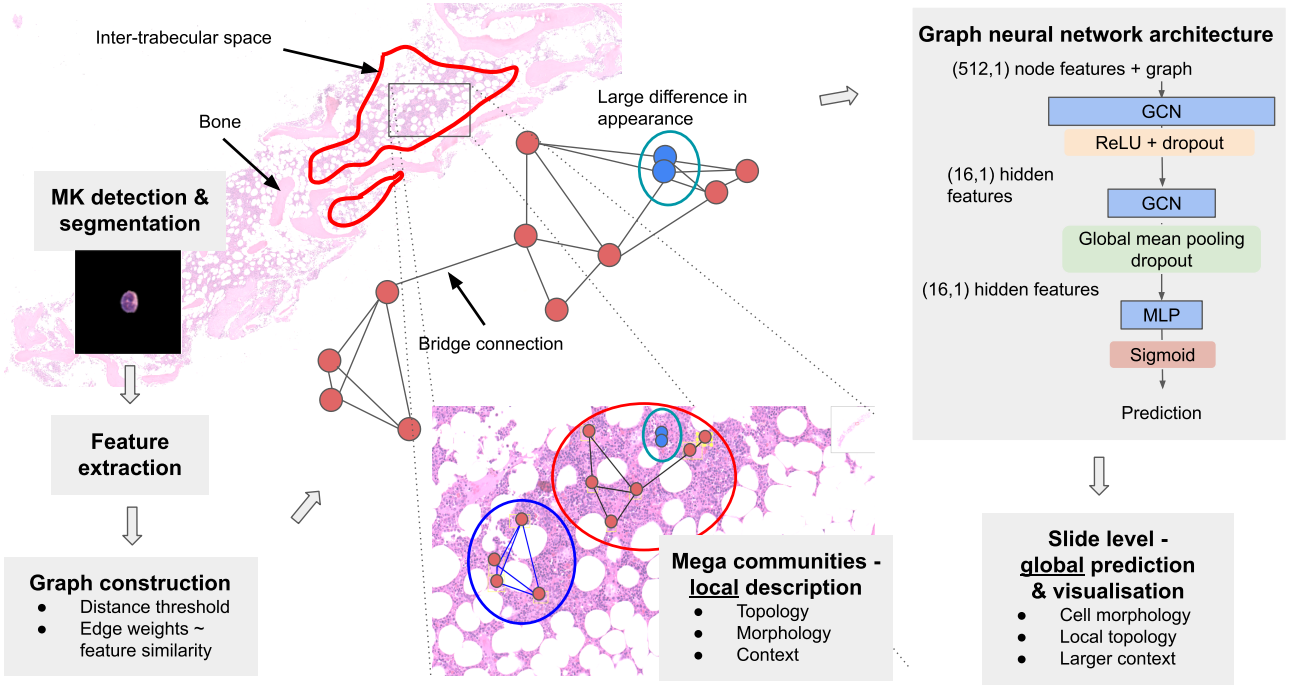

Fig. 1: The proposed multi-scale graph representation. to be read from the top left in an anti-clockwise manner. First on the left, the cells are detected and deep features are extracted to form the base level representation of the immediate cell neighbourhood. Second at bottom right, we have community detection to identify larger patterns in the cell population, over regions of the sample. The third stage on top right gives a yet wider representation providing disease predictions on the whole slide.

in [8] and the resulting 384×384 pixel images were pre-process using channel-wise histogram equalisation to account for non-uniform staining within and across the slides. The feature extractor itself is comprised of a ResNet18 block pretrained on ImageNet [9] followed by classification to increase the disease relevance of the features.

B. Cell neighbourhood

The cell neighbourhood graph is defined as an undirected graph $G := (V, E)$ with $|E|$ edges and $|V|$ nodes composed of the MK cells within a radius R of the target cell. Therefore, the graphs may have a varying number of nodes or cells. Every node v is represented by a corresponding d-dimensional feature vector $x(v) \in \mathbb{R}^d$ extracted from the previous cell classification stage. The adjacency matrix $A \in \mathbb{R}^{|V|\times|V|}$ is symmetric and the entry corresponding to nodes u and v is $A_{u,v} = 1$ if edge $e_{u,v} \in E$. Every edge is weighted by a Gaussian threshold kernel W defined as

$$W_{u,v} = \begin{cases} exp^{-\frac{[dist(u,v)]^2}{2\varsigma^2}} & \text{if } dist(u,v) \leq R \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where the neighbourhood radius $R = 768$ pixels at $40\times$ magnification is determined based on the observation that there should not be a connection between cells that are distant or that are located in separate inter-trabecular spaces (tissue portions separated by thick strands of bone). Unlike other models [3] [10], we do not use a k-nearest-neighbour approach because this would potentially force the graph into an edge topology that has no biological background. In fact, highly connected graphs might be a disease relevant feature. In case of single nodes we omit the cells from the graph representation. Our graphical model consists of two graph convolutional layers and an average pooling layer to aggregate the node level features into a cell neighbourhood representation. The task is predicting the slide-level label in order to assess the boost in performance when merging topology and morphological features. As the cell morphology seen on the slide is just a 2D slice of an actual 3D cell it is by nature noisy data. By classifying the cell neighbourhood representation we use the smoothing effect of the message passing in graph neural networks to our advantage. As demonstrated in [11], message passing leads to a low-pass effect on the input signals of the node features. Assuming that the cell features carry disease relevant information, this leads to a better class separation for the disease prediction. By implementing two graph convolutional layers before the global pooling layer we obtain a node-level representation which includes information that has been passed down from the 2-hop neighbourhood of the target cell within radius R.

C. Cell Communities

Our aim is to obtain a more refined spatial representation of the MK population in a BMT compared to simply obtaining slide-level average features purely based on the topology of the cells. Whereas the cellular neighbourhood captures the immediate local neighbourhood of the cell,

TABLE I: Bone marrow trephine dataset

	Total	Reactive	ET	PV	MF
Slides	136	43	48	19	26
Cells	53866	7332	17724	12486	16324

the communities should represent the larger spatial structure taking into account morpho-topological similarity of neighbouring cells. The graphs are built as described above. Therefore, each slide contains a graph potentially made up of several disconnected components. Assuming that cells with similar embeddings have stronger links between them, we create a weighted graph with edge weights proportional to the cosine similarity in the feature space. This space is based on the features extracted after the pooling layer of the cell neighbourhood model. Using a community detection algorithm we are identify groups of cells which are morpho-topologically more closely related. To obtain a quantitative description of each community we compute the following properties. The graph density d, defined as

$$d = \frac{2|E|}{|V|(|V|-1)}, \tag{2}$$

The clustering coefficient c, defined as

$$c = \frac{1}{|V|}\sum_{v \in V} c_v, \text{ with } c_v = \frac{2T_v}{deg(v)(deg(v)-1)}, \tag{3}$$

for nodes with $deg(v) > 1$, otherwise $c_v = 0$. $deg(v)$ and T_v denote the degree of node v and the number of triangles through the node, respectively. The spatial node (cell) density d_S, defined as

$$d_S = \frac{2|V|}{D^2}, \text{ with } D = \frac{1}{|V|^2}\sum_{u,v \in V} dist(u,v). \tag{4}$$

where D is the spatial graph diameter. Finally, we compute the number of nodes and the average cell prediction P_{avg}s.

III. Experiments

A. Bone Marrow Trephines from the Oxford Archive

The dataset consists of 136 bone marrow trephines including 48 ET, 19 PV, 26 MF and 43 reactive samples acquired from OUH NHS Foundation Trust archive. The MPN patients were diagnosed according to the latest WHO classification scheme from 2016. The reactive control group was selected based on patients showing neither signs of a bone marrow malignancy nor another myeloid disorder. This dataset contains 53,866 MK cells of which more than 50 % were manually validated by a trained human expert. (Ref. Table I)

B. Cell Feature Extraction

The cell feature extractor was fine-tuned using 5-fold cross-validation. For all experiments we use the same cross-validation setup. Weighted random sampling was employed in each batch to counter data imbalance. We used data augmentation strategies such as colour perturbation, noise, rotation and flips during training to avoid overfitting. Binary cross entropy loss was minimised during training by the ADAM optimiser using a learning rate of 0.0001 and the batch size of 64. The model was trained for a maximum of 100 epochs using the accuracy on the validation set as a stopping criterion. The feature extractor itself is comprised of a ResNet18 block pretrained on ImageNet [9] followed by a classification layer to increase the disease relevance of the features.

TABLE II: Mean classification performance on the test sets

Task	Level	accuracy	recall	precision
R/MPN	cell	0.872	0.932	0.921
ET/PV	cell	0.726	0.721	0.517
ET/PMF	cell	0.793	0.775	0.756
PV/PMF	cell	0.734	0.749	0.812
R/MPN	neighbourhood	0.938	0.973	0.957
ET/PV	neighbourhood	0.935	0.924	0.933
ET/PMF	neighbourhood	0.925	0.943	0.892
PV/PMF	neighbourhood	0.896	0.895	0.94

C. Cellular Neighbourhood

The cell level features are fed into the cell neighbourhood graph as node features. We used weighted random sampling to account for dataset imbalance. However, we only retrieve a subset of 40 % because closely clustered cells are more likely to be sampled as neighbouring cells. The models are implemented using pytorch-geometric [12]. The ADAM optimiser is used for parameter updating based on binary cross entropy loss minimisation for a maximum of 50 epochs. The batch size and the learning rate are set to 16 and 0.00001, respectively.

D. Cell Communities and Whole Slide Prediction

We apply community detection on each of the weighted subgraphs obtained by employing the aforementioned distance constraint based on the greedy modularity algorithm implemented in *networkX* [13]. From the resulting cell communities we can extract commonly used graph features.

IV. Results

A. Cell Feature Extraction & Neighbourhood Denoising

Table 2 shows the performance for the cell classification based on ResNet18 and the cell neighbourhood. The former yields a mean prediction accuracy of 87.2 % on the unseen test set for the task reactive vs MPN. The same classification task based on the cell neighbourhood model yields a mean prediction accuracy of 93.8 %. The improvement is even larger for the distinction between disease subtypes.

B. Community Detection and Visualisation

Figure 2a shows the subgraphs overlayed by cell level predictions and 2b shows the community graphs overlayed by the neighbourhood level predictions. The latter clearly shows how the neighbourhood model makes the predictions spatially more consistent resulting in a clinically useful visualisation. For example MK populations of differing disease states can be easily identified as indicated by the dashed circles. The two sets of community features correspond to the communities enclosed by a dashed circle.

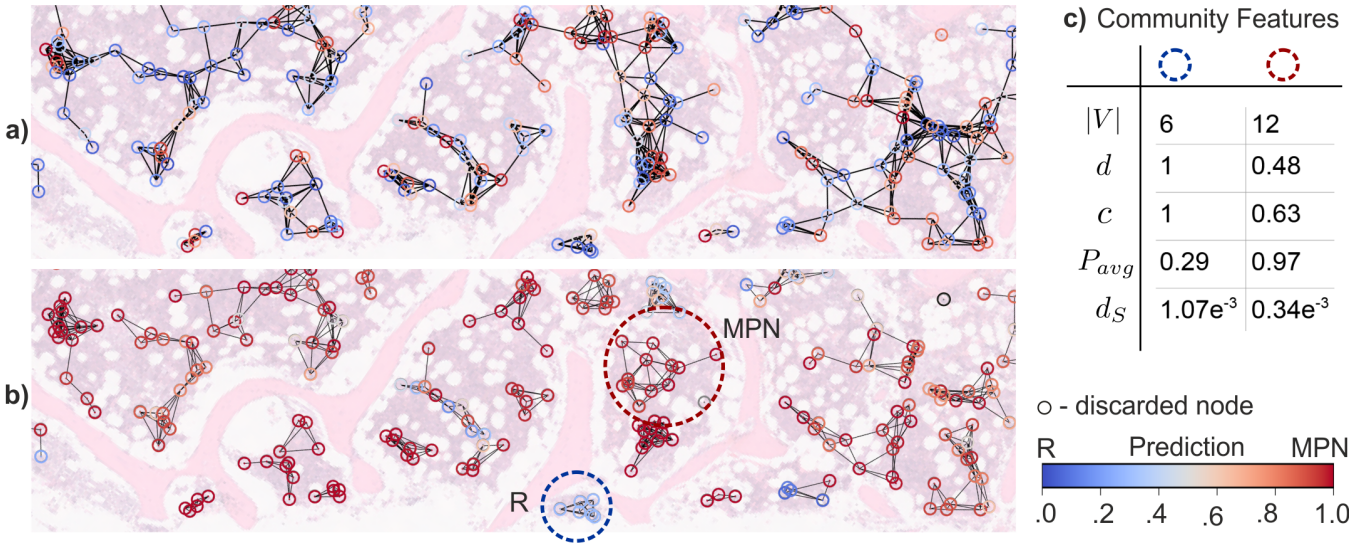

Fig. 2: Exemplar prediction overlay of a heterogeneous ET case. The location of the megakaryocyte cells are represented by circles where the colour denotes the prediction outcome of the binary classification between myeloproliferative neoplasms (MPN) and reactive. The colour bar to the right has 0 corresponding to reactive (R) and 1 to MPN. The edges of the initial constructed graphs are indicated by lines between the cells. Part a) shows the cell predictions obtained by the cell classifier whereas b) indicates the spatially more consistent predictions from the cell neighbourhood model. The latter only depicts intra-community edges. Table (c) to the right has a numerical description of the circled cell communities in part b) based on the following features $|V|$, d, c, P_{avg} and d_S denoting number of cells, graph density, clustering coefficient, average neighbourhood prediction and spatial density. Correlating the marked graphs in (b) with the tabulated values in (c), we can see that the MPN graph, covering a larger area than the R graph, contains more cells and has a lower density, which makes intuitive sense.

TABLE III: Slide level classification performance on the test sets

Task	precision	recall	F-score
R/MPN	0.966	0.925	0.945
ET/PV	0.944	0.895	0.876
ET/PMF	0.963	1.00	0.981
PV/PMF	0.867	1.00	0.929

C. Whole Slide Prediction

Based on majority voting of neighbourhood prediction we can deduce a more robust slide-level prediction compared with a slide-level prediction purely based on cell classification for distinguishing between the MPN disease subtypes, essential thrombocythaemia (ET), polycythaemia vera (PV) and primary myelofibrosis (PMF). Table 3 depicts the performance for the binary classification cases for reactive vs. MPN, ET vs. PV, ET vs. PMF and PV vs. PMF. For all of these combinations, our method achieves high performance in precision, recall and F-score.

V. Conclusion

We have presented a graphical network that is able to represent the cell population in a tissue image at different scales. We do this by collating individual morphological cell features with the immediate cellular neighbourhood as well as with the larger topology in the sample. We encapsulated this multi-scale representation in a data-driven fashion, resulting in a more biomedically interpretable visualisation of samples. Finally, the topological structure of the cell population can enhance the performance of disease prediction. For our exemplar bone marrow trephine dataset, this method enables the distinction between disease subtypes.

References

[1] P. Lang *et al.*, "Cellular imaging in drug discovery," *Nature Reviews Drug Discovery*, 5(1): 343-356, 2006.

[2] G. Jaume *et al.*, "Towards explainable graph representations in digital pathology," *arXiv preprint arXiv:2007.00311*, 2020.

[3] M. Sureka *et al.*, "Visualization for histopathology images using graph convolutional neural networks," in *2020 IEEE 20th International Conference on Bioinformatics and Bioengineering (BIBE)*. IEEE, 331–335, 2020.

[4] M. Ferlaino *et al.*, "Towards deep cellular phenotyping in placental histology," *arXiv preprint arXiv:1804.03270*, 2018.

[5] H. Theissen *et al.*, "Learning cellular phenotypes through supervision," in *43rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 2021.

[6] P. Pati *et al.*, "Hact-net: A hierarchical cell-to-tissue graph neural network for histopathological image classification," in *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Graphs in Biomedical Image Analysis*. Springer, 208-219,2020.

[7] D. Anand, S. Gadiya, and A. Sethi, "Histographs: graphs in histopathology," in *Medical Imaging 2020: Digital Pathology*. International Society for Optics and Photonics, (11320)113200O, 2020.

[8] K. Sirinukunwattana *et al.*, "Artificial intelligence–based morphological fingerprinting of megakaryocytes: a new tool for assessing disease in mpn patients," *Blood advances*, 4(14): 3284-3294, 2020.

[9] K. He *et al.*, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778, 2016.

[10] Y. Zhou *et al.*, "Cgc-net: Cell graph convolutional network for grading of colorectal cancer histology images," in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 0–0, 2019.

[11] N. Hoang and T. Maehara, "Revisiting graph neural networks: All we have is low-pass filters," *arXiv preprint arXiv:1905.09550*, vol. 2, 2019.

[12] M. Fey and J. E. Lenssen, "Fast graph representation learning with pytorch geometric," *arXiv preprint arXiv:1903.02428*, 2019.

[13] A. Hagberg, P. Swart, and D. S Chult, "Exploring network structure, dynamics, and function using networkx," Los Alamos National Lab.(LANL), Los Alamos, NM (United States), Tech. Rep., 2008.