

HistoCartography: A Toolkit for Graph Analytics in Digital Pathology

Guillaume Jaume*

*IBM Research, Zurich – EPFL, Lausanne, *equal contribution*

GJA@ZURICH.IBM.COM

Pushpak Pati*

IBM Research, Zurich – ETH, Zurich

PUS@ZURICH.IBM.COM

Valentin Anklin

ETH, Zurich

ANKLINV@STUDENT.ETHZ.CH

Antonio Foncubierta

IBM Research, Zurich

FRA@ZURICH.IBM.COM

Maria Gabrani

IBM Research, Zurich

MGA@ZURICH.IBM.COM

Editors: M. Atzori, N. Burlutskiy, F. Ciompi, Z. Li, F. Minhas, H. Müller, T. Peng, N. Rajpoot, B. Torben-Nielsen, J. van der Laak, M. Veta, Y. Yuan, and I. Zlobec.

Abstract

Advances in entity-graph analysis of histopathology images have brought in a new paradigm to describe tissue composition, and learn the tissue structure-to-function relationship. Entity-graphs offer flexible and scalable representations to characterize tissue organization, while allowing the incorporation of prior pathological knowledge to further support model explainability. However, their analysis requires prerequisites for image-to-graph translation and knowledge of state-of-the-art algorithms applied to graph-structured data, which can potentially hinder their adoption. In this work, we aim to alleviate these issues by developing HISTOCARTOGRAPHY, a standardized python API with necessary preprocessing, machine learning and explainability tools to facilitate graph-analytics in computational pathology. Further, we have benchmarked the computational time and performance on multiple datasets across different imaging types and histopathology tasks to highlight the applicability of the API for building computational pathology workflows. HISTOCARTOGRAPHY is available at <https://github.com/histocartography/histocartography>.

Keywords: Graph Representation Learning, Computational Pathology, Python API

1. Introduction

Recent advancements in tissue-slide digitization have paved way for enhancing storage, sharing capabilities, and computer-aided inspection by leveraging Deep Learning (DL). Most DL approaches analyze tissue images in three steps, namely patch generation, patch-level feature extraction, and feature aggregation to produce image-level embeddings for downstream pathology tasks. However, they suffer from several limitations, (i) the trade-off between operational resolution and adequate context per-patch, (ii) the aggregation is often sub-optimal, (iii) comprehensive modeling of tissue composition is missing, and (iv) the lack of model transparency raises barriers to deployment in real life.

To circumvent these limitations, entity-graphs are proposed (Demir et al., 2004) where the nodes and edges of the graphs denote tissue entities and their interactions, respec-

tively. Entity-graphs, followed by Graph Neural Networks (GNNs)-based processing, have recently gained popularity in addressing various pathology tasks (Zhou et al., 2019; Chen et al., 2020; Pati et al., 2021; Anklin et al., 2021; Jaume et al., 2021). The entities can be biologically-defined, *e.g.*, nuclei, tissue regions, glands (Zhou et al., 2019; Pati et al., 2021; Anklin et al., 2021), or can be patches (Adnan et al., 2020; Aygunes et al., 2020). The entity-graphs enable to simultaneously capture local entity environment and global tissue composition. They can seamlessly scale to arbitrary tissue dimensions by incorporating arbitrary number of entities and interactions, thus offering an alternate to Multiple Instance Learning (MIL) (Campanella et al., 2019; Lu et al., 2021). The entity-graphs also enable to selectively operate on diagnostically relevant entities, instead of analyzing the entire tissue (Tellez et al., 2019b; Shaban et al., 2020). Furthermore, when the entities depict biological units, such as nuclei, glands etc., the analysis allows pathologists to directly comprehend and reason with the outcomes (Jaume et al., 2020, 2021). However, constructing an entity-graph based pathological workflow demands several prerequisites, such as entity detection, entity encoding, constructing the graph topology etc., alongside standard preprocessing, such as stain normalization, tissue detection etc. Additionally, the workflow requires to utilize the recent advancements in DL for processing graph-structured data. All these may inhibit the adoption of entity-graphs in computational pathology. In addition, the lack of a standardized framework with the aforementioned functionalities urge the researchers to reinvent the wheel, which is cumbersome, time-consuming, hampers reproducibility, and requires a wide range of technical acumen.

To overcome these constraints, we present HISTOCARTOGRAPHY, a novel open-source python library that facilitates graph-analytics in computational pathology. Specifically our contributions are: (i) a standardized, unit-tested python library that unifies a set of histology image manipulation tools, entity-graph builders, GNN models, and model explainability tools, (ii) a benchmark assessment of performance and scalability on classification and segmentation tasks in pathology, (iii) a comprehensive overview of graph representation and modeling in histology, and (iv) a review of extant libraries for histological image analysis.

2. Related Work

2.1 Graphs in Computational Pathology

Entity-graphs are proposed to realize the tissue composition-to-functionality relationship in terms of the phenotypical and structural characteristics of tissue. The entities can be nuclei (Demir et al., 2004; Zhou et al., 2019; Wang et al., 2020; Chen et al., 2020; Pati et al., 2021), tissue regions (Pati et al., 2021), patches (Anand et al., 2020; Adnan et al., 2020; Aygunes et al., 2020; Zhao et al., 2020; Li et al., 2018; Levy et al., 2021), etc. Typically nodes include handcrafted or DL features to characterize the entities, and the topology can depict the spatial or semantic relationship among the entities, *e.g.*, k-Nearest Neighbors (k-NN), region adjacency, or probabilistic models. The graphs can be processed using classic Machine Learning (ML) (Sharma et al., 2016, 2017) or GNNs to outperform state-of-the-art Convolutional Neural Network (CNN)-based approaches for several pathology tasks across multiple organs (García-Arteaga et al., 2017; Zhou et al., 2019; Zhao et al., 2020; Adnan et al., 2020; Pati et al., 2021; Studer et al., 2021; Anklin et al., 2021). Interestingly, when the graph-nodes depict biological entities, *e.g.*, nuclei, tissue regions, the entity-graphs combined with

feature attribution techniques can provide pathologist-friendly interpretations (Zhou et al., 2019; Jaume et al., 2020; Sureka et al., 2020) and explanations (Jaume et al., 2021), unlike pixelated blurry saliency maps. A detailed review of graphs in computational pathology is presented by Ahmedt-Aristizabal et al. (2021).

2.2 Extant Libraries in Computational Pathology

Several open-source libraries facilitate the development of computational pathology workflows. Most of them include helper functions to perform standard preprocessing and visualization. HISTOLAB (Arbitrio et al., 2020) includes Whole Slide Image (WSI)-level tissue detection and tile extraction modules. SYNTAX (Byfield et al., 2020) provides the same features with abstraction where modules can be stacked and run in a pre-defined pipeline. STAINTOOLS (Byfield et al., 2019) provides tools for stain normalization and augmentation. HISTOMICSTK (Beezley et al., 2021) enables to perform tissue detection, object detection and segmentation, image filtering, stain normalization and deconvolution, and handcrafted feature extraction. Further, HISTOMICSTK allows nuclei segmentation and classification using classical ML approaches. It also provides a User Interface (UI) to run containerized modules and pipelines. Though HISTOMICSTK includes valuable functionalities, it caters limited DL tools. Similarly, OPENSLIDE (Gilbert et al., 2020) provides a UI to read and visualize histology images that supports most of the WSI formats. Finally, QUPATH (Bankhead et al., 2021) offers a UI that allows to read, visualize and annotate WSIs. It also includes tools to perform stain normalization, nuclei and tissue detection, and implement basic ML models. However, QUPATH is not a python Application Programming Interface (API), which makes it difficult to integrate into existing workflow and DL frameworks, *e.g.*, PyTorch, Tensorflow. Most importantly, none of the frameworks provide graph-related helpers. With the advent of graph-techniques as a new paradigm for analyzing histology images, a standardized library is desired for reinforcing the development.

3. Histocartography: Graph Analytics Tool for Pathology

In this section, we highlight the core modules of HISTOCARTOGRAPHY, (1) *Preprocessing* module: a set of histology image processing tools and entity-graph builders, (2) *ML* module: helpers to learn from entity-graphs, (3) *Explainability* module: a set of GNN model interpretability tools. List of module-wise functionalities are summarized in Table 1. To facilitate integration and reduce boilerplate code, HISTOCARTOGRAPHY includes a pipeline runner which allows to pre-define pipeline steps along with loading and saving utilities.

3.1 Preprocessing Module

Stain normalization: Variation in Hematoxylin and Eosin (H&E) staining protocols for tissue specimens induces appearance variability that adversely impacts computational methods (Tellez et al., 2019a). To alleviate these variations, HISTOCARTOGRAPHY implements two popular normalization algorithms proposed by Macenko et al. (2009) and Vahadane et al. (2016), similar to STAINTOOLS and HISTOMICSTK, which supports both reference-based and reference-free normalization, *i.e.*, with manual stain vectors. Figure 1 highlights a sample normalization output using our API.

Table 1: Overview of HISTOCARTOGRAPHY functionalities, with the i/o, CPU and GPU compatibility, and availability in extant libraries for individual module. I, M, X, G, P and S denote an image (np.array (Harris et al., 2020)), a mask (np.array), features (torch.Tensor (Paszke et al., 2019)), a graph (DGLGraph (Wang et al., 2019)), predictions (torch.Tensor) and importance scores (torch.Tensor), respectively.

Module	Function	Input	Output	Existing	CPU	GPU
Preprocessing	Vahadane Stain Norm	I	I	✓	✓	✗
	Macenko Stain Norm	I	I	✓	✓	✗
	Tissue Mask Detection	I	M	✓	✓	✗
	Nuclei Detection	I	M	✓	✓	✓
	Nuclei Concepts	I, M	M	✓	✓	✗
	Tissue Component Detection	I	M	✗	✓	✗
	Deep Feature Extraction	I, M	X	✗	✓	✓
	Feature Cube Extraction	I	X	✗	✓	✓
	k-NN Graph Building	X, M	G	✗	✓	✗
	RAG Graph Building	X, M	G	✗	✓	✗
ML	Cell-Graph Model	G	P	✗	✓	✓
	Tissue-Graph Model	G	P	✗	✓	✓
	HACT Model	G, G, X	P	✗	✓	✓
Explainers	GNNEXPLAINER	G	S	✗	✓	✓
	GRAPHGRAD-CAM	G	S	✗	✓	✓
	GRAPHGRAD-CAM++	G	S	✗	✓	✓
	GRAPHLRP	G	S	✗	✓	✓

Tissue Detection: A WSI usually includes significant non-tissue region. Identifying the tissue regions can confine the analysis and reduce computational effort. The tissue detector in HISTOCARTOGRAPHY iteratively applies Gaussian smoothing and Otsu thresholding until the mean of non-tissue pixels is below a threshold. This module is common across HISTOLAB, SYNTAX, HISTOMICSTK and QUPATH.

Nuclei detection: This module enables to segment and locate nuclei in H&E images. Though it is well-studied in computational pathology, only a few public implementations are available. For instance, QUPATH allows to detect nuclei but requires model training and fine-tuning. While providing flexibility, the module includes only elementary ML methods. HISTOCARTOGRAPHY integrates two checkpoints from the state-of-the-art HoVerNet model (Graham et al., 2019) trained on PanNuke (Gamper et al., 2020) and MoNuSac (Ruchika et al., 2020) datasets for nuclei segmentation and classification.

Tissue Component Detection: HISTOCARTOGRAPHY includes an unsupervised superpixel based approach to segment tissue regions. First, the tissue is oversegmented into homogeneous superpixels using Simple Linear Iterative Clustering (SLIC) (Achanta et al., 2012) algorithm. Then, neighboring superpixels are hierarchically merged using color similarity to denote meaningful tissue regions, *e.g.*, epithelium and stroma regions. Superpixels depicting tissue regions are used by Bejnordi et al. (2015); Pati et al. (2020, 2021).

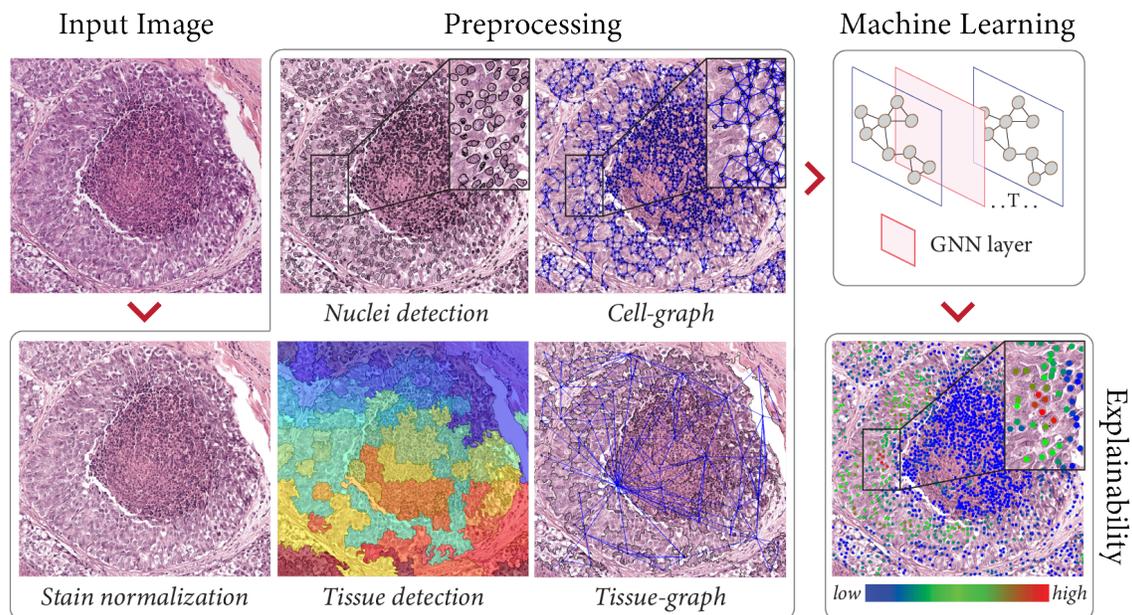


Figure 1: Overview of HISTOCARTOGRAPHY modules and functionalities.

Feature Extraction: HISTOCARTOGRAPHY includes two types of feature extractors, *i.e.*, handcrafted- and CNN-based, to encode the entity characteristics. The handcrafted feature extractor computes entity-level morphological and topological properties. Morphological features capture the shape and size, *e.g.*, entity area, eccentricity, perimeter etc., and the texture captures chromaticity using the gray-level co-occurrence matrix. Topological features capture the local entity distribution using k-NN entity density estimation. A comprehensive list is provided in the Appendix. Handcrafted features can be used for training DL algorithms (Demir et al., 2004; Zhou et al., 2019; Pati et al., 2020; Studer et al., 2021), or concept-based post-hoc explainability (Jaume et al., 2021).

The deep feature extractor allows to extract CNN features by using any pre-trained deep architecture, *e.g.*, ResNet, MobileNet, embedded in torchvision (Marcel et al., 2010). The module intakes patches centered around the entity and extracts features from the penultimate layer of the architectures. If the entity is larger than the specified patch size, then multiple patches within the entity, w/ or w/o overlapping, are processed, and the final feature is computed as the mean of the patch-level deep features, as used in Chen et al. (2020); Pati et al. (2020, 2021). Deep features can alternatively be extracted from the WSI to build a feature-cube as suggested by Shaban et al. (2020); Tellez et al. (2019b).

Graph builders: HISTOCARTOGRAPHY presents two graph builders, *i.e.*, the thresholded k-NN and the Region Adjacency Graph (RAG). The k-NN graph builder defines the graph topology by connecting each entity to its k-closest neighbors. Connections between distant entities beyond a threshold can be pruned to have spatial sparsity in the graph. We recommend this builder to connect single entities, *e.g.*, nuclei, glands. The RAG builder connects entities using spatial adjacency, *i.e.*, entities sharing a common boundary. It builds a sound topology when dealing with dense segmentation maps, *e.g.*, tissue regions. Figure 1 presents samples of cell- and tissue-graphs. Further, the module fuses the node features and the topological distribution to render a Deep Graph Library (DGL) graph for an image.

3.2 Graph Machine Learning Module

HISTOCARTOGRAPHY includes a set of DL models, based on a GNN backbone to learn from graph-structured tissue representations. It includes two state-of-the-art GNN layers, *i.e.*, Graph Isomorphism Network (GIN) (Xu et al., 2019) and Principal Neighborhood Aggregation (PNA) (Corso et al., 2020). PNA proves to outperform GIN provided more computational resources (Dwivedi et al., 2020). HISTOCARTOGRAPHY defines cell- and tissue-graph models, which are GNN-based abstractions to learn from biological entity-graphs. They offer efficient (Pati et al., 2021), scalable (Anklin et al., 2021) and explainable (Jaume et al., 2020, 2021) approaches to analyze histology images. Further, the library includes models to jointly represent and learn from cell- and tissue-graphs (Pati et al., 2021). The models in HISTOCARTOGRAPHY are organized such that they can be adapted to various GNN backbones, tasks (*e.g.*, regression, clustering, classification, segmentation), organs, and entity-types. These models provide the blueprints to accelerate the development of graph-based models in computational pathology. All the graph modules are implemented using DGL (Wang et al., 2019), a state-of-the-art library for GNNs built around PyTorch.

3.3 Explainability Module

HISTOCARTOGRAPHY includes four post-hoc feature attribution graph explainers, that can generate node-level saliency maps to highlight the node-wise contribution towards an output task. Namely, the library includes two gradient-based explainers (GRAPHGRAD-CAM (Selvaraju et al., 2017; Pope et al., 2019) and GRAPHGRAD-CAM++ (Chattopadhyay et al., 2018; Jaume et al., 2021)), a node pruning-based explainer (GNNEXPLAINER (Ying et al., 2019)), and a layer-wise relevance propagation explainer (GRAPHLRP (Schwarzenberg et al., 2019)). The saliency map can be visualized by overlaying the node importances on the input image (see Figure 2). Alternatively, entities with high importances can be extracted and studied independently to assess their relevance (Jaume et al., 2021).

4. Benchmarking HistoCartography

We benchmark HISTOCARTOGRAPHY in terms of run-time and performance for various histopathology tasks, *i.e.*, stain normalization, tissue detection, tumor classification and segmentation etc., on images of varying dimensions. The CPU and GPU compatible modules are assessed on a single-core POWER8 processor and a NVIDIA P100 GPU, respectively.

4.1 Computational Time

Analyzing the computational time for processing a histology image is imperative. We thoroughly analyze the run-time of HISTOCARTOGRAPHY functionalities on a set of Regions of Interest (RoIs) and WSIs (Table 2). The preprocessing modules are observed to be the most time-consuming. For instance, Vahadane stain normalization can take up to 3 minutes to process a $11'000 \times 11'000$ image, whereas Macenko method is $2\times$ faster for competitive result. The implementations are computationally similar to HISTOLAB and STAINTOOLS, and scale linearly w.r.to image size. The cell- and tissue-graph construction take 2.5 and 4.1 seconds respectively for a 1000×1000 image with the following parameters. Nuclei detection is performed on patches of size 256×256 with an overlap of 164 pixels. Nuclei

Table 2: Run time analysis of HISTOCARTOGRAPHY core functionalities (in seconds).

Mod.	Function / Image type	Tumor RoI			WSI		
		1000 ²	2500 ²	5000 ²	5000 ²	7500 ²	11000 ²
Preprocessing	Vahadane Stain Norm	1.77	6.46	29.03	30.67	68.27	186.10
	Macenko Stain Norm	0.80	2.86	11.19	15.98	32.37	81.72
	Tissue Mask Detection	-	-	-	1.04	2.11	8.09
	Feature Cube Extraction	0.24	1.61	5.92	5.83	11.77	30.21
	Cell-Graph Generation	2.51	13.33	50.26	-	-	-
	Tissue-Graph Generation	4.14	21.30	83.66	39.18	93.26	276.71
ML	Cell-Graph Model	0.028	0.033	0.040	-	-	-
	Tissue-Graph Model	0.011	0.015	0.026	0.039	0.056	0.069
	Hierarchical Model	0.034	0.041	0.057	-	-	-
Explainers	GNNEXPLAINER	12.00	13.09	35.33	-	-	-
	GRAPHGRAD-CAM	0.011	0.022	0.035	0.025	0.030	0.033
	GRAPHGRAD-CAM++	0.011	0.023	0.035	0.026	0.030	0.033
	GRAPHLRP	0.020	0.024	0.90	0.079	0.085	0.089

features are extracted from 72×72 patches centered around the nuclei, that are resized to 224×224 and processed by ResNet34 pretrained on ImageNet (Deng et al., 2009). Finally, thresholded k-NN topology is built with $k = 5$ and a threshold distance of 50 pixels. For the tissue-graph, SLIC is used to extract 400 superpixels per image, that are subsequently merged to provide the tissue components. Tissue features are also extracted using ResNet34 with 144×144 size patches that are resized to 224×224 . The graph buildings can be further optimized as per the task by downsampling the input image, reducing the patch overlap, or by using a lighter feature extractor. For extracting the feature cube representation, we process patches of size 144×144 resized to 224×224 w/o overlap by pretrained ResNet34.

Tumor RoIs (TRoIs) are processed using a cell- and tissue-graph model, and the hierarchical cell-to-tissue graph model (Pati et al., 2021). They consist of three PNA layers with 64 hidden units followed by an additional 2-layer MLP with 128 hidden units for classification. WSIs are processed using SEGGINI Anklin et al. (2021), a weakly supervised approach based on tissue-graphs, which contains six GIN layers with 64 hidden units followed by a 2-layer MLP with 128 hidden units. The models process in near real-time irrespective of the increment in the graph size. The graph explainers are based on GNNs with 3 GIN layers, each having a 2-layer MLP with 32 hidden units, and a 2-layer MLP head. GNNEXPLAINER is the slowest among all as it requires to optimize a mask to explain each image.

4.2 Performance Benchmark

Table 3 benchmarks the performance of HISTOCARTOGRAPHY for classification and segmentation tasks. Classification is performed on BRACS (Pati et al., 2021) and BACH (Aresta et al., 2019) datasets to characterize breast tumors using cell-graph model, tissue-graph model, and HACT-Net (Pati et al., 2021), and the performance is measured by weighted-F1 score. Segmentation is performed using SEGGINI (Anklin et al., 2021) to delineate Gleason patterns in prostate cancer images from UZH (Zhong et al., 2017) and SICAPv2 (Silva-Rodríguez et al., 2020), and the performance is measured by average Dice score. We evaluate

Table 3: Benchmarking HISTOCARTOGRAPHY for classification and segmentation (in %).

Task	Dataset	Model	Image Type	Avg. #pixels	#classes	Avg. Dice	Weighted F1
Classification	BRACS	CG-GNN	TRoI	3.9×10^6 (40 \times)	7	-	55.9 ± 1.0
	BRACS	TG-GNN	TRoI	3.9×10^6 (40 \times)	7	-	56.6 ± 1.3
	BRACS	HACT-Net	TRoI	3.9×10^6 (40 \times)	7	-	61.5 ± 0.9
	BACH	HACT-Net	TRoI	3.1×10^6 (20 \times)	4	-	90.7 ± 0.5
	SICAPv2	SEGGINI	WSI	121×10^6 (10 \times)	6	-	62.0 ± 3.6
	UZH	SEGGINI	TMA	9.6×10^6 (40 \times)	6	-	56.8 ± 1.7
Seg.	SICAPv2	SEGGINI	WSI	121×10^6 (10 \times)	4	44.3 ± 2.0	-
	UZH	SEGGINI	TMA	9.6×10^6 (40 \times)	4	66.0 ± 3.1	-

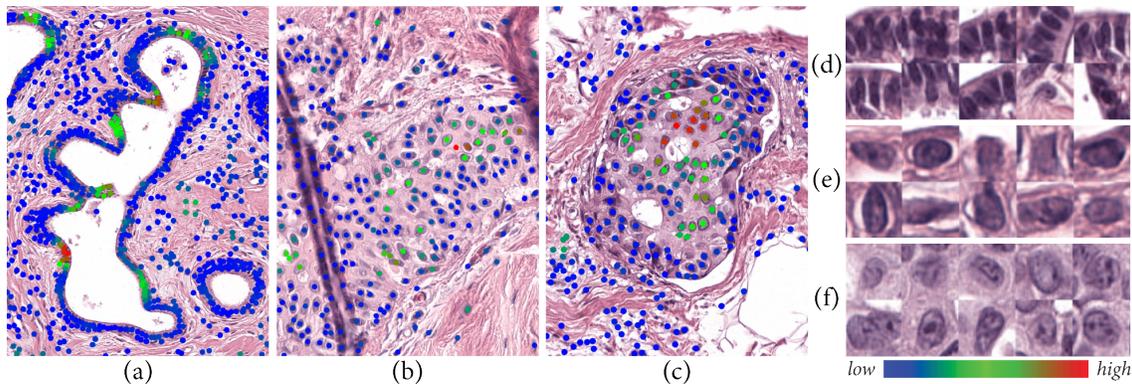


Figure 2: Qualitative explanations of sample breast RoIs: (a) Benign, (b) ADH, (c) DCIS. (d, e, f) highlight the ten most important nuclei for the respective samples.

on various image types, *i.e.*, tumor RoIs, tissue microarrays, and whole-slides, to highlight the scalability of entity-graphs in HISTOCARTOGRAPHY to arbitrary image dimensions.

Figure 2 presents the outcome of GRAPHGRADCAM function in HISTOCARTOGRAPHY to interpret a cell-graph model. This function renders per-image explanations in terms of node-level saliency maps by applying post-hoc feature attribution methods on trained cell-graph model. Further, the cell-graph model can be interpreted by characterizing the highlighted important nuclei per-image, as shown in Figure 2.

5. Conclusion

We introduced HISTOCARTOGRAPHY, the first open source library, to the best of our knowledge, to facilitate graph analytics, *i.e.*, graph representation, learning, and explainability, in computational pathology. It can potentially enable researchers to develop entity-graph based pathology workflows by leveraging the inbuilt helpers. As the library is built on python, the deep learning researchers can seamlessly customize and integrate the functionalities into their task-specific workflows. HISTOCARTOGRAPHY is constantly growing with new functionalities and improved implementations, aiming to promote the adoption of graph-based analysis in computational pathology.

References

- R. Achanta et al. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11):2274–2282, 2012.
- M. Adnan et al. Representation learning of histopathology images using graph neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 4254–4261, 2020.
- D. Ahmedt-Aristizabal et al. A survey on graph-based deep learning for computational histopathology. In *arXiv:2107.00272*, 2021.
- D. Anand et al. Histograms: graphs in histopathology. In *SPIE Medical Imaging*, volume 11320, 2020.
- V. Anklin et al. Learning whole-slide segmentation from inexact and incomplete labels using tissue graphs. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2021.
- E. Arbitrio et al. histolab, 2020. URL <https://github.com/histolab/histolab>.
- G. Aresta et al. Bach: Grand challenge on breast cancer histology images. *Medical Image Analysis*, 56:122–139, 2019.
- B. Aygunes et al. Graph convolutional networks for region of interest classification in breast histopathology. In *SPIE Medical Imaging*, volume 11320, 2020.
- P. Bankhead et al. Qupath, 2021. URL <https://qupath.github.io/>.
- J. Beezley et al. Histomicstk, 2021. URL <https://github.com/DigitalSlideArchive/HistomicsTK>.
- B.E. Bejnordi et al. A multi-scale superpixel classification approach to the detection of regions of interest in whole slide histopathology images. In *SPIE Medical Imaging*, volume 9420, 2015.
- P. Byfield et al. Staintools, 2019. URL <https://github.com/Peter554/StainTools>.
- P. Byfield et al. Syntax, 2020. URL <https://github.com/jgamper/compay-syntax/tags>.
- G. Campanella et al. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature Medicine*, 25(8):1301–1309, 2019.
- A. Chattopadhyay et al. Grad-CAM++: Generalized gradient-based visual explanations for deep convolutional networks. In *IEEE Winter Conference on Applications of Computer Vision*, pages 839–847, 2018.
- R.J. Chen et al. Pathomic fusion: An integrated framework for fusing histopathology and genomic features for cancer diagnosis and prognosis. In *IEEE Transactions on Medical Imaging*, 2020.

- G. Corso et al. Principal neighbourhood aggregation for graph nets. In *Neural Information Processing Systems (NeurIPS)*, 2020.
- C.G. Demir et al. The cell graphs of cancer. *Bioinformatics*, 20:145–151, 2004.
- J. Deng et al. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009.
- V.P. Dwivedi et al. Benchmarking graph neural networks. In *arXiv:2003.00982*, 2020.
- J. Gamper et al. Pannuke dataset extension, insights and baselines. *arXiv:2003.10778*, 2020.
- J.D. Garcíá-Arteaga et al. A lymphocyte spatial distribution graph-based method for automated classification of recurrence risk on lung cancer images. In *International Symposium on Medical Information Processing and Analysis*, volume 10956, 2017.
- B. Gilbert et al. Openslide, 2020. URL <https://github.com/openslide/openslide-python/>.
- S. Graham et al. Hover-net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images. *Medical Image Analysis*, 58, 2019.
- R. P. Harris et al. Array programming with NumPy. *Nature*, 585(7825):357–362, 2020.
- G. Jaume et al. Towards explainable graph representations in digital pathology. In *International Conference on Machine Learning (ICML), Workshop on Computational Biology*, 2020.
- G. Jaume et al. Quantifying explainers of graph neural networks in computational pathology. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- J. Levy et al. Topological feature extraction and visualization of whole slide images using graph neural networks. *Pacific Symposium on Biocomputing*, 26:285–296, 2021.
- R. Li et al. Graph cnn for survival analysis on whole slide pathological images. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pages 174–182, 2018.
- M.Y. Lu et al. Data efficient and weakly supervised computational pathology on whole slide images. *Nature Biomedical Engineering*, 5:555–570, 2021.
- M. Macenko et al. A method for normalizing histology slides for quantitative analysis. In *IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 1107–1110, 2009.
- S. Marcel et al. Torchvision the machine-vision package of torch. In *ACM International Conference on Multimedia*, page 1485–1488, 2010.
- A. Paszke et al. Pytorch: An imperative style, high-performance deep learning library. In *Neural Information Processing Systems (NeurIPS)*, pages 8024–8035, 2019.

- P. Pati et al. Hact-net: A hierarchical cell-to-tissue graph neural network for histopathological image classification. In *Medical Image Computing and Computer Assisted Intervention (MICCAI) Workshop on GRaphs in biomedicAl Image anaLysis*, 2020.
- P. Pati et al. Hierarchical graph representations in digital pathology. In *arXiv:2107.10073*, 2021.
- P.E. Pope et al. Explainability methods for graph convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10764–10773, 2019.
- G. Ruchika et al. Multi-organ nuclei segmentation and classification challenge 2020. *IEEE Transactions on Medical Imaging*, 2020.
- R. Schwarzenberg et al. Layerwise relevance visualization in convolutional text graph classifiers. *EMNLP Workshop*, pages 58–62, 2019.
- R.R. Selvaraju et al. Grad-CAM : Visual Explanations from Deep Networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 618–626, 2017.
- M. Shaban et al. Context-aware convolutional neural network for grading of colorectal cancer histology images. *IEEE Transactions on Medical Imaging*, 39(7):2395–2405, 2020.
- H. Sharma et al. Cell nuclei attributed relational graphs for efficient representation and classification of gastric cancer in digital histopathology. In *SPIE Medical Imaging*, volume 9791, 2016.
- H. Sharma et al. A comparative study of cell nuclei attributed relational graphs for knowledge description and categorization in histopathological gastric cancer whole slide images. In *IEEE Symposium on Computer-Based Medical Systems*, pages 61–66, 2017.
- J. Silva-Rodríguez et al. Going deeper through the gleason scoring scale: An automatic end-to-end system for histology prostate grading and cribriform pattern detection. *Computer Methods and Programs in Biomedicine*, 195, 2020.
- L. Studer et al. Classification of intestinal gland cell-graphs using graph neural networks. In *International Conference on Pattern Recognition (ICPR)*, 2021.
- M. Sureka et al. Visualization for histopathology images using graph convolutional neural networks. In *IEEE International Conference on Bioinformatics and Bioengineering (BIBE)*, 2020.
- D. Tellez et al. Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology. *Medical Image Analysis*, 58, 2019a.
- D. Tellez et al. Neural image compression for gigapixel histopathology image analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(2):567–578, 2019b.
- A. Vahadane et al. Structure-preserving color normalization and sparse stain separation for histological images. *IEEE Transactions on Medical Imaging*, 35(8):1962–1971, 2016.

- J. Wang et al. Weakly supervised prostate tma classification via graph convolutional networks. In *IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 239–243, 2020.
- M. Wang et al. Deep graph library: Towards efficient and scalable deep learning on graphs. In *CoRR*, volume abs/1909.01315, 2019.
- K. Xu et al. How powerful are graph neural networks? In *International Conference on Learning Representations (ICLR)*, 2019.
- R. Ying et al. GNNExplainer: Generating Explanations for Graph Neural Networks. In *Advances in Neural Information Processing Systems*, 2019.
- Y. Zhao et al. Predicting lymph node metastasis using histopathological images based on multiple instance learning with deep graph convolution. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4837–4846, 2020.
- Q. Zhong et al. A curated collection of tissue microarray images and clinical outcome data of prostate cancer patients. 4, 2017.
- Y. Zhou et al. CGC-net: Cell graph convolutional network for grading of colorectal cancer histology images. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops*, 2019.