







HistoContours: a framework for visual annotation of histopathology whole slide images

K. Al-Thelaya¹ , F. Joad¹ [†], N. U. Gilal¹ , W. Mifsud² , G. Pintore³ , E. Gobetti³ , M. Agus¹ , and J. Schneider¹ 

¹Division of Information and Computing Technology, College of Science and Engineering, Hamad Bin Khalifa University, Qatar Foundation, Doha, Qatar

²SIDRA, Qatar Foundation, Qatar

³ Visual and Data-intensive Computing, CRS4, Italy

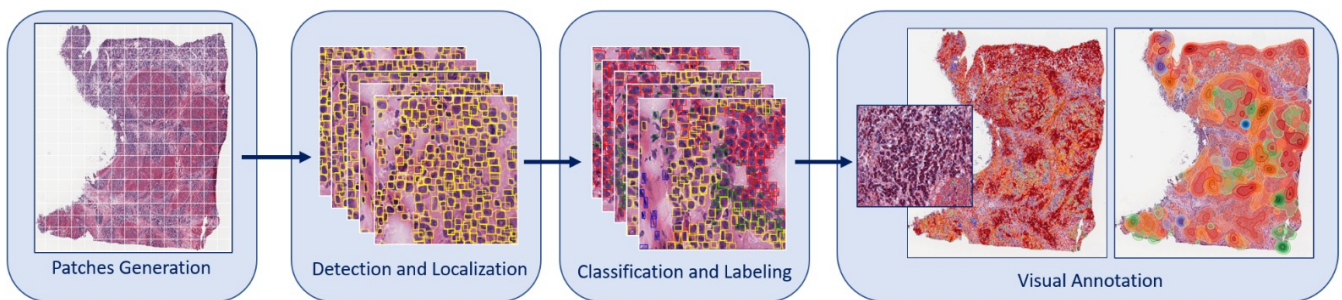


Figure 1: Our end-to-end framework allows histopathologists to work at the nucleus and tissue level at the same time, a feature rarely offered by the competition. This is enabled by slide subdivision in patches, followed by deep-learning-based localization of nuclei, and classification. The resulting class probabilities and locations of nuclei are then turned dense using kernel density estimation.

Abstract

We present an end-to-end framework for histopathological analysis of whole slide images (WSIs). Our framework uses deep learning-based localization & classification of cell nuclei followed by spatial data aggregation to propagate classes of sparsely distributed nuclei across the entire slide. We use YOLO (“You Only Look Once”) for localization instead of more costly segmentation approaches and show that using HistAuGAN boosts its performance. YOLO finds bounding boxes around nuclei at good accuracy, but the classification accuracy can be improved by other methods. To this end, we extract patches around nuclei from the WSI and consider models from the SqueezeNet, ResNet, and EfficientNet families for classification. Where we do not achieve a clear separation between highest and second-highest softmax activation of the classifier, we use YOLO’s output as a secondary vote. The result is a sparse annotation of the WSI, which we turn dense by using kernel density estimation. The result is a full vector of per pixel probabilities for each class of nucleus we consider. This allows us to visualize our results using both color-coding and isocontouring, reducing visual clutter. Our novel nuclei-to-tissue coupling allows histopathologists to work at both the nucleus and the tissue level, a feature appreciated by domain experts in a qualitative user study.

CCS Concepts

• **Applied computing** → **Imaging**; • **Computing methodologies** → **Object detection**; • **Human-centered computing** → **Information visualization**; **Heat maps**;

1. Introduction

Clinical decisions are nowadays heavily driven by medical imaging, which has become a fundamental element of the clinical analysis pipeline. Most diagnostic tasks rely on a variety of scanning technologies to detect objects, providing evidence-based conclusions. In digital pathology, the primary purpose is to remove from

the process of clinical reasoning the inadequacies of working directly on physical tissue samples. There is an abundance of studies that illustrate the effectiveness of digital distribution for collaborative diagnostic, teaching, and research purposes and show the concordance of rendering diagnoses using whole slide images (WSIs) compared to glass slides (see, e.g., [POR*19, ZGL*21] and references therein). Moreover, it has been demonstrated how digital pathology allows healthcare systems to maintain pathology operations during public health emergencies (e.g., the COVID-19 pandemic) [HRA*20]. Beyond supporting the replication of usual

[†] equal contribution

workflows using a virtual microscope, digital pathology also opens the door for disruptive supportive technologies, such as the fully automated processing of WSIs to extract meaningful information, which is considered the holy grail of pathology [DAC19]. Automation has the potential to remove repetitive tasks of high cognitive load from the busy schedule of pathologists. It ensures higher productivity, accuracy, and reproducibility. It also helps to validate the standardization of study outcomes and inclusion criteria, and to extract features supporting decisions [POR*19].

In this context, visualization of medical images is an enabling technology to gain insights from WSIs and to support various tasks ranging from characterizing the relationship between imaging and diseases to neuroscientific research [CWDW18]. Visualization techniques are commonly built upon various computational methods (e.g., signal processing, geometry processing, AI, etc.) processing image-based and auxiliary input data such as genetic markers and clinical indicators [LML*07]. Of the aforementioned methods, AI and deep learning methods are seeing rapidly increasing use due to their promising results and potential impact on digital pathology progression. In particular, deep models excel at tasks such as segmentation as well as the detection of nuclei and regions of interest within WSIs [ATAG*21]. This study addresses some of the challenges of medical imaging visualization in digital pathology. In this work, we present a visualization framework based on deep learning to help medical practitioners to explore and analyze images in histology. Our framework aims to efficiently support histologists to carry out investigations on nuclei labeling and segmentation.

Contributions We propose three kinds of annotations to support the visual analysis of WSIs using information extracted from deep learning inference:

1. micro-scale annotations of class activation maps for evaluating the behavior of inference models and correlating activations with diagnostic features;
2. micro-/meso-scale annotations using color-coded bounding boxes for representing detection and classification outcomes;
3. macro-scale annotations using weighted density estimation and isocontours to provide an overview of main tissue spots and their correlations. We also provide a preliminary qualitative evaluation of our framework by an expert pathologist.

2. Related Work

Our work targets computational histopathology and visual annotation of WSIs. In the following, we only discuss the methods that are more closely related to the proposed pipeline. For a wider coverage, we refer the reader to excellent surveys on recent developments [WCH*22], especially with respect to the application of deep learning [SCM21] and graph-based embeddings [AAAD*21].

Deep Learning and Computational Histopathology Since the release of PanNuke [GAKB*19], an open, labeled data set containing histology patches comprising 19 tissue and 5 nuclei types, various architectures have been proposed to perform different tasks on digital WSIs using information at the nucleus level. Modern deep learning technologies, such as graph convolutional networks

(GCNs), have been customized for dealing with the inherent characteristics of gigapixel slide images. Lu *et al.* [LGB*20] introduced a pipeline that constructs a graph from the nuclei level to the entire WSI-level, followed by a GCN for WSI-level prediction. Attention mechanisms have also been considered based on the assumption that the size and texture of nuclei can determine the magnification at which a patch is extracted [SCS*20]. Contrastive learning approaches have been considered that do not require a custom network [CKNH20]. Instead, they rely on a large number of minibatch instances for obtaining negative samples per training instance. This method led to the development of ALBRT [DBRM21], which trains a multi-headed architecture based on convolutional neural networks (CNNs) to predict per cell-type counts in a patch. Similarly, Ciga *et al.* [CXM22] consider contrastive learning for using residual networks for multiple downstream tasks on multi-organ digital histopathology data sets. In this work, we show that supporting a joint localization-and-classification architecture using color augmentation to alleviate intra- and inter-sample staining and tissue variations boosts localization performance significantly, but does not achieve the performance of state-of-the-art standalone classifiers. We therefore follow our localization approach by CNN-based classification models that are trained using transfer learning. Our hybrid approach significantly outperforms the state-of-the-art [GAKB*19, GKB*20].

Kernel Density Estimation in Histopathology Kernel density estimation (KDE) is a powerful regression technique for estimating a continuous probability density function $p(\mathbf{x})$ given a sequence of scattered measurements $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_k\}$ on a given domain $\mathbb{D} \subseteq \mathbb{R}^d$ [Che17]. The technique is extensively used in statistic analysis, including the analysis of histology images for a variety of tasks. Kather *et al.* [KMRA*15] use KDE for determining a continuous representation of tumor microvessel density to detect angiogenic hotspots in histological WSIs. Similarly, Yeh *et al.* [YYH*14] use a KDE for calculating the stain density across digitized pathology slides in order to perform correlation to in-vivo cellular magnetic resonance (MR) images. The same authors use kernel regression to estimate the spatial distribution of nuclear size across WSIs in order to perform automatic grading of renal cell carcinoma [YPPH14]. KDE has also been used for registration of images obtained with different modalities. In particular, Jiang *et al.* [JLP*19] use KDE as a weighting scheme for improving the accuracy in image registration for re-stained WSIs, while Theelke *et al.* [TWM*21] use a similar strategy to register specimens digitized with multiple scanning systems to compute an optimal global transformation taking care of different resolutions. In this context, KDE is used to weight local transformation estimates to mitigate the influence of registration errors. However, to the best of our knowledge, KDE has not yet been considered in the visualization process of deep learning inference applied to WSIs. In this work, we propose various weighting schemes to compute continuous density functions related to the outcomes of detection and classification models.

Visual Annotations for Histopathology Concurrently with the development of acquisition, processing, and visualization technologies for WSIs, tools to support histopathologists in annotation tasks have also been developed. Traditionally, domain scientists used manual drawing tools for creating free-form annotations

on top of virtual slides that can be represented as overlay polygons [ZGL*21, HMST21]. Recently, Jessup *et al.* [JKW*21] proposed a scalable software system for focus+context exploration and annotation of multi-channel WSIs obtained through recent high-plex imaging methods, that are able to measure the levels and sub-cellular localization of 20-60 proteins. Similarly, Somarakis *et al.* [SIK*21] developed a visual analytics system for highly multiplexed cytometric imaging data, i.e., images where every pixel contains multiple scalar values, each representing the abundance of a specific, pre-defined protein at subcellular resolution. Stoltzfus *et al.* [SFG*20] developed instead a spatial analysis toolbox called histo-cytometric multidimensional analysis pipeline (CytoMAP), incorporating data clustering, positional correlation, dimensionality reduction, and 2D/3D region reconstruction to identify localized cellular networks and reveal features of tissue organization. Manz *et al.* [MGP*22] developed instead a flexible web toolkit displaying 2D visualizations of highly multiplex and multi-scale datasets, supporting both image-based and polygonal annotations. All these systems, even if successful for supporting visual analysis, still lack in fully supporting and integrating deep learning methods [CWWF*19]. Only very recently, methods supporting the visualization of inference outputs from deep learning have been proposed to support the interpretation and analysis of slides, and they are starting to be integrated in visual analytics systems. For example, Facetto [KBJ*19] provides supervised and unsupervised learning tools for enabling the discovery of single-cell phenotypes in high-dimensional multi-channel microscopy images of human tumors and tissues, while companies like HALO AI, or VisioPharm are starting to provide custom integrated solutions that incorporate deep learning models in image viewers. Corvò *et al.* [CWDW18] created a custom visual analytics system for accelerating reporting in breast cancer by semi-automatizing the grading process according to Nottingham Histologic Grade protocol, and directly integrated in a Computer-Aided-Design viewer. Lu *et al.* [LWC*21] use heatmaps to provide interpretable visual information from a clustering-constrained-attention multiple-instance learning (CLAM) framework. The authors reason that this enables domain scientists to evaluate the relative contribution and importance of every tissue region to the predictions without using costly pixel-level annotations during training. Del Rio *et al.* [DRLA*22] include similar computationally generated annotations and visual cues to help the pathologist prioritize high-interest image areas. Class activation maps (CAMs) [ZKL*16] provide a means to post-analyse black box inference results. CAMs typically analyze the spatial feature distribution in the last layer before the first dense layer. The result is a low-resolution class probability distribution that is upsampled and overlaid on the input image, often by means of a semi-transparent heatmap color-coding. Chuang *et al.* [CCY*21] use CAMs to detect nodal micrometastasis in colorectal cancer on annotation-free WSIs. Jaume *et al.* [JPA*21] propose a toolkit for explaining the outputs of graph neural networks (GNNs) on histology images. They use two popular gradient-based CAM explainers (Graph-GradCAM [PKR*19] and Graph-GradCAM++ [JPB*21]). Our framework, likewise, exploits GradCAM [SCD*17] to provide domain experts with a tool to analyze the results of an otherwise black-box AI. In addition, we also provide tissue density contours generated automatically through KDE on top of the localization

and classification results to further enable experts to disseminate the findings of our framework.

3. Application Domain and Requirements

We provide here a brief overview of the histopathology domain and what are the requirements for providing effective annotations to support reporting. Throughout the project, we involved an expert histopathologist for gathering information about the domain, the visual analysis problems and limitations, for defining requirements, and for assessing the outcomes of the proposed annotation framework through analysis of specific use cases of interest for the specialist. We also took into account relevant surveys about the usage of visual analytics in histopathology [CWWF*19, CWDW18], in which specific needs related to the pathologist workflows are examined.

3.1. Overview of Histopathology

Histopathology refers to the visual analysis of microscopy images of living tissues, and for many diseases it is considered the diagnostic "gold standard", especially for what concerns inflammatory or neoplastic conditions. The analysis process involves the surgical extraction of sample cells or tissues. The tissue is then fixed, dehydrated, embedded, sectioned, stained and mounted before it is generally examined under a microscope by a histopathologist. In order to reveal the development and the extent of the disease, and before the microscope examination, the sections are stained with one or more pigments. The aim of staining is to reveal cellular components, while counterstains are used to provide contrast. The most commonly used staining in histopathology is obtained by combining hematoxylin and eosin (often abbreviated H&E). Hematoxylin has the effect to stain nuclei with blue color, while eosin is able to reveal cytoplasm and the extracellular connective tissue matrix through a pink color. Once properly prepared, the tissue slides are digitized through a WSI scanner, that produces high-resolution images (called virtual slides) normally at 40x resolution (0.25 μm per pixel). The original glass slides or the virtual digital scans are then sent to an histopathologist, who examines the tissue under a microscope, looking for any abnormal findings, or through a Computer Aided Diagnosis viewer. For performing these cognitive tasks, various challenges are faced: high variability of structure and appearance, depending on the tissue, staining and imaging process. Variability is relevant also for morphology of cells, clustering and aggregation, and the density inside the tissues. The histopathologist then produces a report that lists any abnormal or important findings from the biopsy. This report is finally sent and discussed with the surgeon who originally performed the biopsy on the patient, and the clinicians involved in the planning of the treatment. In most cases, the report produced by the histopathologist is decisive for the final diagnosis and for the successive planning of the treatment of the condition. Successful and accurate examination of slides requires knowledge, training, and experience and in many cases histopathologists decide to specialize on particular diseases or organs. Moreover, it happens often that they require a second opinion from other histopathologists or additional analysis when they do not feel confident about their outcomes. Considering that histopathologists normally have a workload of few thousands of

cases per year, they are increasingly getting used to automatizing processes to speed-up their analysis and reporting activities, especially for what concerns the collection of patient information, and the assessment of staining quality and tissue slide preparation. Even if the analysis process is mostly depending on the personal experience of the histopathologist, for most diseases specific guidelines and protocols are continuously compiled and refined for supporting examinations of different pathologies [Wea10, VS10]. For what concerns the examination of the slides at microscope, histopathologist use different zooms to highlight patterns at various resolutions, for analyzing cluster of cells, singular nuclei, frequency of patterns, orientation of features, shapes and textures: specifically, 5x-10x is the zoom level for detecting architectural patterns, 10x-20x is the zoom used for revealing lymphocyte invasion, 20x-40x is used for assessing cellular morphology, and finally 20x-40x is considered for the time consuming mitotic count, in case it is required by the specific protocol for grading stage and aggressivity of the pathological condition [CWWF*19].

3.2. Requirements for Annotations

As outcome of the continuous interaction with the domain expert in form of meetings, lectures and interviews, and after analysis of relevant literature (see section 2), we came out with the following requirements for providing annotations associated to the output of deep learning models to effectively support the slide examination process:

- **R1: Explicability.** The ability to explain the behavior of models is of paramount importance in many applications. In the context of histopathology, it is extremely useful to assess whether the model is performing correctly and to understand if there is a relation between statistical inference and the traditional workflow of the histopathologist. Whereas deep learning produces black box oracles, visual tools are verifiable in a-posteriori analyses.
- **R2: Multiscale Analysis.** Pathologists examine slides at various scales. A full view of the WSI is used to identify major areas, verify the tissue type, and assess the overall appearance. At the nuclei level, spatial distribution and condition of individual cells also provides important evidence. In between, the spatial relation of cell clusters matters. Thus, visual tools that support a large range of scales are of extreme interest to domain scientists. This requirement can be further refined according to the specific needs related to the examination protocols [CWDW18]:
 - **R2.1: Magnification level.** The assessment of different features need different magnification level, ranging from high level architectural patterns (5x-10x), up to small scale nuclei measurement and mitotic count (20x-40x);
 - **R2.2: Reference.** Since there are no anatomical references in the tissue samples, it is important to provide reference information for supporting the exploration and enabling the pathologist to create her own mental maps and be efficient during the exploration of the slide.
- **R3: Visual Clarity.** We consider particularly important to avoid visual clutter. Overlaid information should be strictly value added and cover or hide the slide content as little as possible, in order to do not impact negatively the cognitive process. Color

consistency is decisive for visual assessment, and particular care needs to be applied in the decision of color schemes of the overlays applied onto the visualized slides.

- **R4: Automation.** Given the high workload for histopathologists, one of the main requirements is efficiency during the examination and reporting tasks. Automatic visual tools are in particular needed for the quality assurance and analysis of outcomes of WSI process, for the mitigation of navigation efforts, for quantification, for visual comparison and for memorization of patterns.

4. Methodology

The visual inspection of WSIs is a fundamental part of digital pathology. However, tasks such as counting cells are hindered by the heterogeneity in the image as well as color and biological variability, making such routine tasks cumbersome and prone to errors and subjectivity. To address this, our end-to-end visualization framework exploits recent advances in deep transfer learning to localize and classify nuclei to construct a full map of nuclei within a WSI. The framework is schematized in Fig. 1 and it contains the following components:

- **Patches generation:** we subdivide WSIs into overlapping patches by using the features of the hierarchical image representation provided by OpenSlide [GGH*13].
- **Localization:** for each patch, we localize nuclei using a YOLO model trained on PanNuke data.
- **Classification:** we then use the resulting list of nuclei positions to generate a new smaller image patch centered around each nucleus. The latter patches are then fed into a deep classifier, again trained using PanNuke data.
- **Annotation:** we then combine nuclei positions and the classifier's softmax probabilities to obtain classification data points scattered across the WSIs domain. Finally, we use KDE to obtain a probability distribution for each pixel of the WSI.

A detailed description of each of the latter three steps is provided in the following subsections.

4.1. Detection and Localization




Figure 2: Left: Localization using plain YOLOv5. Right: HistAugAN [WKS*21] colorspace augmentation significantly boosts YOLO's performance.

Unlike many previous works, we abandon per-pixel segmentations of cell nuclei in favour of faster localization. We reason that

the additional accuracy that can be provided by segmentation is unnecessary for the task at hand, especially when spatially aggregated statistics are desired. In this work, we use YOLOv5 [JCS*22] to localize nuclei, a state-of-the-art object detection algorithm that provides high inference speed and accuracy. We trained the YOLOv5 model on PanNuke data and use it to perform nuclei localization for multiple WSIs. However, intra- and inter-sample variations in the color space severely affect its performance (also see Fig. 2, left). Color space problems are common in histopathology and are due to variations in tissue composition as well as age of the hematoxylin and eosin staining agents. The problem was recently addressed by Wagner *et al.* [WKS*21], who present a specialized image-to-image generative adversarial network, called “HistAuGAN”, that we use for data augmentation. As can be seen in Fig. 2, HistAuGAN boosts the localization performance massively, especially when performing inference on a sample that differs from the color spaces used for training.

4.2. Labeling and Classification

Due to the high inter- and intra-class variations of nuclei texture and shape, traditional methods may not be the best option for fine-grained classification and labeling [QRW*19]. Deep learning architectures such as CNNs have achieved unprecedented accuracy in different domains, including histological image classification [SRT*16, BH16, BGB*18]. ResNet [HZRS16], SqueezeNet [IHM*16], and EfficientNet [TL19] are prominent deep CNN architectures, offering models that have been pre-trained on large data sets such as ImageNet [DDS*09]. In this work, we evaluate the performance of these three architectures for the classification of nuclei.

4.3. Visual Annotations of WSIs

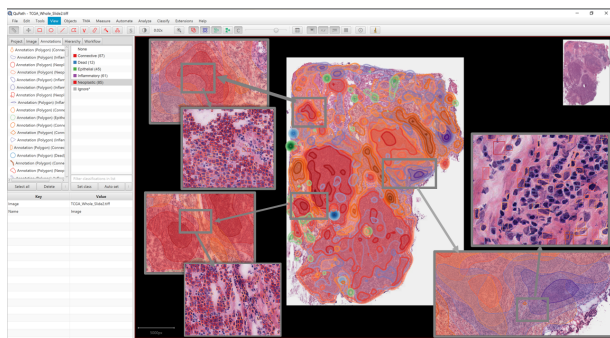


Figure 3: Annotations computed using our framework on TCGA data being explored in the open-source histopathology tool QuPath [HMST21]. Box annotations are useful for micro-scale examination of singular nuclei while density isocontours are useful for macroscale assessment of the various regions of the slides in search of specific patterns.

The input of our visual annotation scheme is the WSI along with the localization and classification results: for each nucleus, a 2D axis-aligned bounding box and a softmax probability for the targets (neoplastic, connective, epithelial, dead, inflammatory).

We designed three annotation methods to support histopathologists while addressing the requirements described in Sec. 3.2:

- **A1: Activation Maps.** Similarly to other frameworks [JPA*21], we consider the popular GradCAM method to highlight the main points of attraction for our classifier in the per-nucleus patches. This visual annotation fulfills requirement *R1*, and can be used as assessment tool of the various models applied, as well as a tool for assessing staining and WSI generation.
- **A2: Annotated Nuclear ROIs.** We overlay the outcomes of the localization and classification models in form of color-mapped frame boxes around the detected nuclei. We use a categorical ColorBrewer maps consistent with standard visual annotations considered in histology (e.g., red for neoplasia, purple for inflammatory, blue for dead, green for soft/connective, and orange for epithelial), mapping the probability of the dominant class to the color intensity. This annotation fulfills the requirement *R2.1* for microscale analysis at the nuclear level. It also supports counting, hence it fulfills requirement *R4* for automation.
- **A3: Density Isocontours.** Depending on the zoom level, A2 does not meet requirement *R2.2* and *R3*: the high density of nuclei and classification noise lead to high visual clutter to be useful for macroscale analysis. To address this, we use kernel density estimation (KDE), a data smoothing and cleaning method, to summarize the information of our scattered softmax probabilities. Given these scattered data points, KDE estimates a probability density function (pdf) across the entire WSI. Using the pdf, we apply isocontouring (see also Fig. 3) to highlight areas with varying densities of nuclei of any given class. Density isocontours can be also used as references for exploration. This meets *R2.1*, *R2.2* and *R3* at the WSI-level.

Fig. 3 shows an example of the annotations on a breast cancer slide, highlighting how box annotations are indicated for micro-scale examination of singular nuclei while density isocontour are useful for macroscale assessment of the various regions of the slides in search of specific patterns.

Kernel Density Estimation To compute density isocontours, we use a weighted KDE with radially symmetric, unimodal Gaussian kernels. A copy of the kernel is centered at every nucleus. This results in the general form

$$f(x) = \frac{1}{\sum_{i=1}^N w_i} \sum_{i=1}^N w_i K(x - x_i), \quad (1)$$

where w_i is the weight for nucleus i , K is the kernel function and N is the total number of nuclei in the slide. We then derive weights as follows. Let t denote tissue type, p_{ti} the softmax probability for tissue t at nucleus i and $\hat{t} = \operatorname{argmax}_t (p_{ti})$ the class with the highest probability \hat{p}_i at nucleus i . We first suppress non-maxima, $p_i = \hat{e}_i \cdot \max_t (\hat{p}_i)$, where \hat{e}_i denotes the i^{th} unit vector. We then use weights of the form $w_i = (\alpha_i^{-1} p_i)^\gamma$, where the α_i and γ can be used to enhance the contrast between various tissue densities. For instance, $\alpha_i = 1, \gamma = 1$ results in a standard scheme, $\alpha_i = p_{Mi}$, where M denotes the index of the second most probable class, in a contrastive scheme, and $\gamma = 0$ in an unweighted scheme. We then interpret each component of the multivariate KDE as a scalar field from which we extract per tissue isocontours starting from a prob-

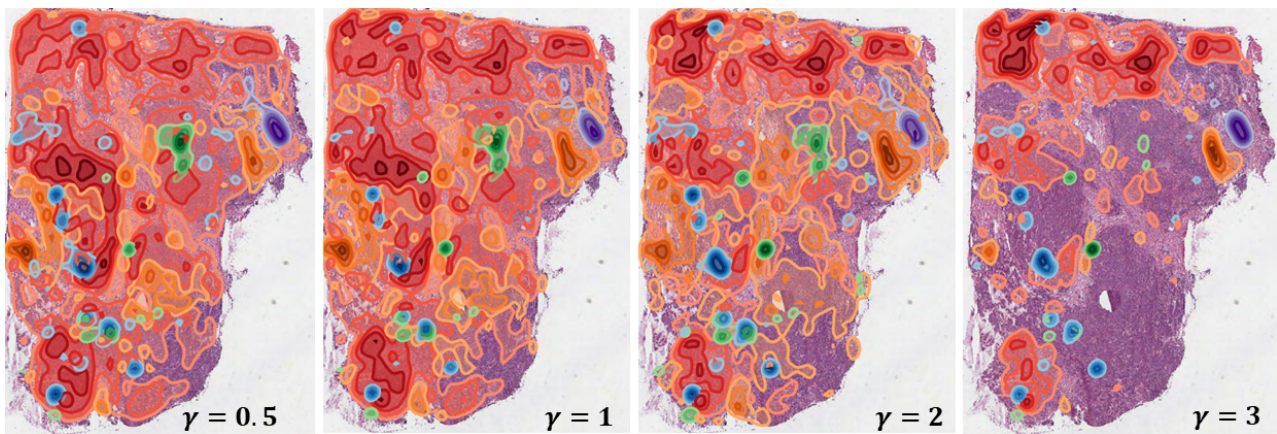


Figure 4: Examples with isocontours generated with contrast scheme on a WSI representing a breast biopsy with tumor (TCGA archive). Increasing values of exponent of γ show that higher values allow a better separation of contours representing different classes, at the cost of reduced coverage of the slide.

ability density threshold value p_t (in all results produced for this paper we considered $p_t = 0.5$)

Fig. 4 shows an example of density isocontours generated with the contrast scheme and different values of γ on one example of breast cancer extracted from The Cancer Imaging Archive Portal. Higher values of exponent γ lead to a better discrimination of contours related to different classes, at the cost of reduced coverage of WSIs. Histopathologists can choose between different parameter values according to their needs for what concerns clutter attenuation or detailed analysis of interaction and proximity between patterns related to different cell types.

5. Experiments and Results

We implemented the framework for annotation in Python. The models are created by using PyTorch, YOLOv5, and FastAI, while the annotations are created in GeoJSON format, ready to be imported in the open source software QuPath for visual exploration of virtual slides. For the slide image management, we use OpenSlide library [GGH*13], that is able to import and manage image hierarchies of various microscopes, while for visualization and exploration of images and annotations, we use QuPath [HMST21], that is currently the most popular open-source histopathology analysis framework. We performed the inference and data processing for annotating entire slides on a workstation running Linux Ubuntu 20.04 and equipped with a 16-core CPU Intel(R) Core(TM) i9-9900KS CPU with 4.00GHz frequency and GPU Nvidia Geforce RTX 2080 with 11 GB RAM. For the assessment of the framework, we considered a collection virtual slides and the corresponding metadata downloaded from the public portal Genomic Data Commons of National Cancer Institute, that is part of the National Institutes of Health at the U.S. Department of Health and Human Services. The GDC Data Portal is a robust data-driven platform that allows cancer researchers and bioinformaticians to search and download cancer data for analysis [JFGS17]. We performed a qualitative evaluation with the help of an expert histopathologist, who, in the interest of full disclosure, is also a co-author of this paper. He examined image samples from The Cancer Imaging Archive Portal taken from examples of cancer in different tissues, namely breast, colon, and

bladder. Finally, he provided his assessment on the annotations generated for a slide representing a pediatric kidney Wilms tumor with anaplasia.

5.1. Training Results

We use the PanNuke data set [GAKB*19, GKB*20], one of the most comprehensive data sets available, to train all models. PanNuke is based on several publicly available data sets [KVS*17, Vu19, LLH*18]. PanNuke uses five non-background classes which are consistent with former studies [GVR*19, HCT*19]. Unlike the binary (neoplastic, non-neoplastic) labels sometimes used in the literature, all of our classifiers predict all five classes. For Yolo, we considered the standard setup, and we trained the model for 400 iterations with the following parameters: initial and final learning rate 0.01, Adam optimizer with $\beta_1 = 0.937$ and weight decay 0.0005. To evaluate our trained YOLOv5 model, we use PanNuke's fold3 as testing data (as recommended by data providers [GKB*20]) and the intersection over union (IoU) metric to perform localization based on a variable threshold range as shown in Fig. 5 (left).

The classification results attained by YOLOv5 are shown in Fig. 5 (right). To boost this classification accuracy, we use YOLOv5 for localization, extract an image patch around each nucleus, and feed these patches to more sophisticated classifiers such as SqueezeNet, ResNet101, and EfficientNet-b7. YOLO's localization output comprises of axis-aligned bounding boxes, which we use as visual annotations of individual nuclei. Different tissue types exhibit large variations with respect to nuclei shape, size and staining adsorption. Given that this information is readily available to histopathologists and (as per our domain expert) histopathologists commonly check the tissue type at the beginning of each session, we feed this information into the classifier to further boost performance. We color-code tissue type using a bounding box around each nucleus. Thus, we achieved an absolute accuracy increase by about 2%. This results in a weighted accuracy of over 85%, which significantly outperforms prior art (e.g., Gamper *et al.* report 62% F₁ score on neoplastic nuclei whereas we achieve over 90% [GKB*20]). Table 1 summarizes our results and can be

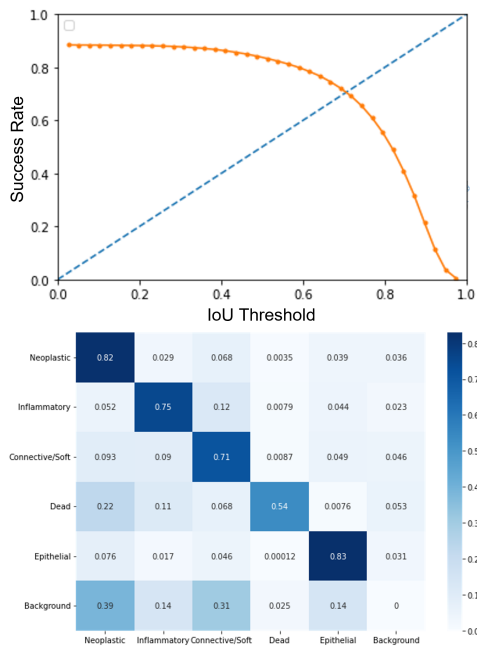


Figure 5: Top: Detection rate of nuclei based on the IoU threshold. Bottom: Confusion matrix of YOLOv5's classification.

directly compared to Gamper *et al.* [GKB*20] (Table V therein). We further examine the impact of the bounding boxes by using GradCAM [SCD*17] to highlight important regions in images with respect to the predicted class. Three examples are shown in Fig. 6. We would also like to note that, in the case of WSIs, and unlike in this figure, we center each patch around the nucleus to provide enough context around each nucleus. In contrast, PanNuke comprises images of resolution 256×256 , resulting in many nuclei being located near image boundaries.

5.2. GDC Experiments

We used the framework to perform complete annotation of a set of 73 slides extracted from the GDC portal and representing various cases of neoplasias affecting different kind of tissues, in order to compare the outcomes of the detection and classification models with metadata reported in the portal. Specifically, we selected the cases in a way that the reported percent of neoplastic nuclei range between 25% and 50%, and an average total number of nuclei of 133.9 ± 124 thousands. We obtained that our annotation has an error of $22.1 \pm 8.7\%$ compared to the reported metadata value. In some cases, the number of neoplastic nuclei is overestimated, especially for tissues that are not included in PanNuke training data. Fig. 7 top shows one example of wrong estimate over a slide representing a tissue extracted from a rectum biopsy. Our model detects 75% of neoplastic nuclei while the reported value is 25%. In this case, it is important to remind that the model is trained over PanNuke dataset, that does not contain rectum tissues. Finally, Fig. 7 bottom shows an example of correct estimate on a slide representing a stomach sample (43.6% estimate vs 45% ground truth).

For what concerns scalability, the complexity of annotation process depends from the size of the slide and the number of detected




Figure 6: Using GradCAM to visualize what our model is interested in. Left-to-right: original patch, GradCAM signal with plasma colormap overlaid onto the patch, plain GradCAM signal in gray-scale. Top-to-bottom: Neoplastic, Epithelial, Connective nuclei. Note that the bounding box used to color-code tissue receives the most attention.

nuclei. For the collection of slides that we considered, we were able to perform a full annotation in average time of 14.8 ± 9.7 minutes. The scatter plot in Fig. 8 shows the full annotation time with respect to the number of detected nuclei in the slide.

5.3. Test Case: Pediatric Kidney Wilms Tumor with Anaplasia

As additional test for the annotation framework, we fully annotated one slide with kidney anaplastic tumor with our detection and classification model trained on PanNuke data, and we submitted to our histopathologist collaborator for assessment. Between the various pediatric neoplasias, kidney Wilms tumors are considered particularly severe, since they can grow into and destroy nearby tissue, and they can also spread (metastasize) to other parts of the body [IEPJ16]. Up to 10% of Wilms tumours can be classified as anaplastic histology, that is considered a particularly unfavourable histology, since the cells in these tumours look very different from normal kidney cells and divide abnormally. Anaplasia may exhibit localized or diffuse behavior, according to the areas where it can be found. Anaplastic tumors are difficult to be recognized and even more difficult to treat because they tend to be more resistant to chemotherapy: moreover anaplastic tumors are more likely to recur than tumours with a favourable histology.

Fig. 9 shows the result of the annotation of the WSI, clearly demonstrating that the model, even if not trained to deal with rare pediatric neoplasias, is able to correctly recognize a diffuse tumor affecting the whole tissue, apart of the lower right corner containing epithelial tissue. Moreover isocontours provide the way to highlight specific patterns related to the behavior of that specific neoplasia.

Fig. 10 shows some specific patterns examined by the

Table 1: Classification results of different models trained using PanNuke data set based on color-coded tissue frames. Compare to [GKB*20]

Model	Connective/Soft			Dead			Epithelial			Inflammatory			Neoplastic			Weighted Accuracy
	precision	recall	f1-score	precision	recall	f1-score	precision	recall	f1-score	precision	recall	f1-score	precision	recall	f1-score	
ResNet34	0.76733	0.76395	0.76564	0.69433	0.47493	0.56404	0.89762	0.87777	0.88758	0.77559	0.73718	0.7559	0.88012	0.91546	0.89744	0.83486
SqueezeNet	0.79444	0.72152	0.75623	0.72169	0.35572	0.47655	0.8846	0.89549	0.89001	0.78753	0.72453	0.75472	0.85296	0.93966	0.89421	0.83251
ResNet101	0.79357	0.76572	0.7794	0.67488	0.54399	0.60241	0.90768	0.9	0.90383	0.77046	0.75751	0.76393	0.88814	0.92164	0.90458	0.84532
EfficientNet-B5	0.78422	0.7804	0.78231	0.69704	0.579	0.63256	0.91299	0.90248	0.90771	0.76595	0.76092	0.76343	0.89771	0.91149	0.90455	0.84627
EfficientNet-B7	0.80214	0.77731	0.78953	0.75448	0.55818	0.64165	0.91659	0.91038	0.91348	0.77933	0.783	0.78116	0.89599	0.92192	0.90877	0.85422




Figure 7: We tested the framework on a collection of 73 slides extracted from GDC portal and we compared the number of neoplastic nuclei detected by our model with the percent reported in GDC metadata. Top: a slide with rectum tissue with wrong neoplasia overestimate (75% vs 25%) (rectum tissues are not represented in PanNuke dataset). Bottom: a slide with stomach tissue with correct neoplastic estimate (43.6% vs 45%).

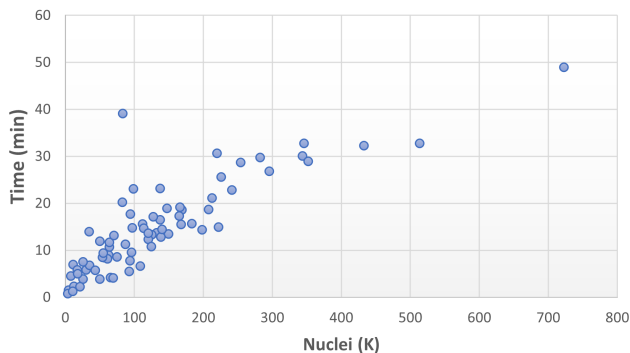


Figure 8: Annotation time for a collection of slides downloaded from GDC portal. The average computation time for full annotations is 14.8 ± 9.7 minutes and it depends on the number of nuclei detected.

histopathologist during the visual assessment, and representing particular behavior of the tumor, from replication of epithelial tissue (top left) up to formation of tubular structures (top right), and inflammatory conditions (bottom left and right).

5.4. Qualitative evaluation

From the qualitative assessment performed by the histopathologist, we got the following preliminary outcomes in form of comments and suggestions:

1. Colored bounding boxes and density contours give a good overview of various cell types across a whole slide image, and they are able to speed-up the analysis of architectural patterns, cellular morphology, and the mitotic count.
2. The estimated numbers and proportions of different nuclei can be employed to compare different slides or zones within slides, and to speed-up quantitative tasks for grading the aggressivity and the stage of the tumour.
3. By exploring the different distributions of cells scattered over slides, the various tissue structures and the interactions between different cell types are highlighted which gives references for high magnification analysis.
4. Zooming in and out could show detailed structures in a short amount of time and with less effort compared to exploring the same slide without annotations.
5. The annotations also draw attention to features that may otherwise be easily missed or would take longer to identify when examining an unannotated whole slide image, resulting in better accuracy and faster discovery (for example, clusters of inflammatory cells or tubular structures).
6. Looking at the different formations of cellular groups enables practitioners to form a preliminary picture of the distribution of cells and their different clusters by dividing images into structured, similar groups with overlapping displaying the invasive areas between tumors and other cellular formations.
7. Cancer grade and level of invasion can be easily measured and estimated. In slides where small spots of inflammation have not been easily identified, our annotations enable the pathologist to recognize them, clearly characterizing the impact of the tumor on its surrounding. They also facilitate spotting neoplastic cells in small gatherings within huge clusters of connective tissue and inflammation surrounding tumor clusters which is hard to estimate or imagine without annotations.
8. Highlighting the boundary of tumor invasion zones depicts its level of activity and aggressiveness, which can be noticed clearly using density contours that show multiple levels of density, where, at the lowest level, tumor and inflammation or connective tissue are overlapping. This can support grading and be




Figure 9: We annotated a full slide from a kidney biopsy exhibiting anaplastic Wilms tumor. The framework, even if trained with Pannuke data, is able to correctly recognize that the neoplasia is diffusely affecting the tissue. The contours are able to highlight specific patterns related to the behavior of the tumor. Left: original slide. Center: box annotations over the WSI. Right: isocontours annotations covering the slide.




Figure 10: Details from a kidney biopsy exhibiting anaplastic Wilms tumor. The framework is able to recognize different patterns in the behavior of this particular neoplastic condition, ranging from formation of tubular structures to construction of stromal structures.

used for quantification of scores according to the specific protocols.

9. Our annotations can be used to compare the different nuclei classification within each slide by describing each class separately by aligning them side by side to conceive differences in quantity and locality.
10. Moreover, it can be used to compare multiple slides to determine which one is the most useful for a second round of more careful clinical analysis to make prognosis. This process is cumbersome and very time consuming if conducted without annotations-aided visualization.
11. The GradCAM shows that in many cases both pathologist and the model look at the same image features, yet pathologist tends to prefer a wider view to perform further inspection and verification while the model use nuclei texture, appearance, color, and shape as dominating features to determine classes.

In some cases where the model fails to detect the type of nuclei based on their intrinsic features, context is found to be very important as well. Some mis-classified cases are also found difficult by pathologist due to absence of context information due to the limited image size. It was surprising that the model can recognize the type of nuclei while paying less attention to the surrounding context, especially for types where context is very crucial such as inflammatory nuclei.

6. Conclusion

In this paper, we presented a visual framework for digital histopathology. We show that localization is an apt, yet infrequently considered, substitute for more costly segmentation. We also demonstrated that color-space augmentation is necessary for localization. Finally, we found that using only the localizations of

YOLO and feeding image patches around each nucleus into a dedicated classifier significantly outperforms prior art. The feedback of domain experts in our initial qualitative user study is encouraging, but a more thorough study on our framework's usability and impact on domain experts' workflows is needed. Another future research direction is a deeper examination of different KDE weighting schemes. Furthermore, we would like to use focus+context techniques to seamlessly blend between annotation and WSI, and we would like to use a hierarchical data structure [ATAS20] for storing and summarizing the softmax activations. Finally, we plan to develop specialized protocol-based visual analytics systems for different diagnostic tasks [CWWF*19].

Acknowledgments EG and GP acknowledge the contribution of Sardinian Regional Authorities under projects SVDC and PAM. This publication was made possible by NPRP grant #0130-200207 from the Qatar National Research Fund (a member of Qatar Foundation). The findings herein reflect the work, and are solely the responsibility, of the authors.

References

- [AAAD*21] AHMEDT-ARISTIZABAL D., ARMIN M. A., DENMAN S., FOOKES C., PETERSSON L.: A survey on graph-based deep learning for computational histopathology. *Computerized Medical Imaging and Graphics* (2021). DOI: 10.1016/j.compmedimag.2021.102027. 2
- [ATAG*21] AL-THELAYA K., AGUS M., GILAL N. U., YANG Y., PINTORE G., GOBBETTI E., CALÍ C., MAGISTRETTI P. J., MIFSUD W., SCHNEIDER J.: Inshade: Invariant shape descriptors for visual 2d and 3d cellular and nuclear shape analysis and classification. *Computers & Graphics* 98 (2021), 105–125. DOI: 10.1016/j.cag.2021.04.037. 2
- [ATAS20] AL-THELAYA K., AGUS M., SCHNEIDER J.: The mixture graph-a data structure for compressing, rendering, and querying segmentation histograms. *IEEE TVCG* 27, 2 (2020), 645–655. DOI: 10.1109/TVCG.2020.3030451. 10
- [BGB*18] BASHA S. S., GHOSH S., BABU K. K., DUBEY S. R., PULABAIGARI V., ET AL.: RCCNet: An efficient convolutional neural network for histological routine colon cancer nuclei classification. In *Proc. ICARCV* (2018), pp. 1222–1227. DOI: 10.48550/arXiv.1810.02797. 5
- [BH16] BAYRAMOGLU N., HEIKKILÄ J.: Transfer learning for cell nuclei classification in histopathology images. In *Proc. ECCV* (2016), ECCV, pp. 532–539. DOI: 10.1007/978-3-319-49409-8_46. 5
- [CCY*21] CHUANG W.-Y., CHEN C.-C., YU W.-H., YEH C.-J., CHANG S.-H., ET AL.: Identification of nodal micrometastasis in colorectal cancer using deep learning on annotation-free whole-slide images. *Modern Pathology* 34, 10 (2021), 1901–1911. DOI: 10.1038/s41379-021-00838-2. 3
- [Che17] CHEN Y.-C.: A tutorial on kernel density estimation and recent advances. *Biostatistics & Epidemiology* 1, 1 (2017), 161–187. DOI: 10.1080/24709360.2017.1396742. 2
- [CKNH20] CHEN T., KORNBLITH S., NOROUZI M., HINTON G.: A simple framework for contrastive learning of visual representations. In *Proc. ICML* (2020), pp. 1597–1607. DOI: 10.48550/arXiv.2002.05709. 2
- [CWDW18] CORVÒ A., WESTENBERG M. A., DRIEL M. A. V., WIJK J. J.: Visual Analytics in Histopathology Diagnostics: a Protocol-Based Approach. In *Proc. VCBM* (2018), The Eurographics Association. DOI: 10.2312/vcbm.20181226. 2, 3, 4
- [CWWF*19] CORVÒ A., WESTENBERG M. A., WIMBERGER-FRIEDL R., FROMME S., PEETERS M. M. R., DRIEL M. A. V., WIJK J. J.: Visual Analytics in Digital Pathology: Challenges and Opportunities. In *Proc. VCBM* (2019), The Eurographics Association. DOI: 10.2312/vcbm.20191240. 3, 4, 10
- [CXM22] CIGA O., XU T., MARTEL A. L.: Self supervised contrastive learning for digital histopathology. *Machine Learning with Applications* 7 (2022). DOI: 10.1016/j.mlwa.2021.100198. 2
- [DAC19] DIMITRIOU N., ARANDJELOVIĆ O., CAIE P. D.: Deep learning for whole slide image analysis: An overview. *Frontiers in Medicine* 6 (2019). DOI: 10.3389/fmed.2019.00264. 2
- [DBRM21] DAWOOD M., BRANSON K., RAJPOOT N. M., MINHAS F.: Albri: Cellular composition prediction in routine histology images. In *Proc. ICCV* (2021), pp. 664–673. DOI: 10.48550/arXiv.2108.08306. 2
- [DDS*09] DENG J., DONG W., SOCHER R., LI L.-J., LI K., ET AL.: ImageNet: A large-scale hierarchical image database. In *Proc. CVPR* (2009), IEEE, pp. 248–255. DOI: 10.1109/CVPR.2009.5206848. 5
- [DRLA*22] DEL RIO M., LIANAS L., ASPEGREN O., BUSONERA G., VERSACI F., ZELIC R., VINCENT P. H., LEO S., PETERSSON A., AKRE O., PIREDDU L.: AI support for accelerating histopathological slide examinations of prostate cancer in clinical studies. In *Proc. ICIAP Workshops* (2022), pp. 545–556. DOI: 10.1007/978-3-031-13321-3_48. 3
- [GAKB*19] GAMPER J., ALEMI KOHBANANI N., BENET K., KHURAM A., RAJPOOT N. M.: PanNuke: An open pan-cancer histology dataset for nuclei instance segmentation and classification. In *European Congress on Digital Pathology* (2019), pp. 11–19. DOI: 10.1007/978-3-030-23937-4_2. 2, 6
- [GGH*13] GOODE A., GILBERT B., HARKES J., JUKIC D., SATYANARAYANAN M.: Openslide: A vendor-neutral software foundation for digital pathology. *Journal of pathology informatics* 4, 1 (2013), 27. DOI: 10.4103/2153-3539.119005. 4, 6
- [GKB*20] GAMPER J., KOHBANANI N. A., BENES K., GRAHAM S., JAHANIFAR M., KHURRAM S. A., AZAM A., HEWITT K., RAJPOOT N.: PanNuke dataset extension, insights and baselines. *arXiv preprint arXiv:2003.10778* (2020). arXiv:2003.10778. 2, 6, 7, 8
- [GVR*19] GRAHAM S., VU Q. D., RAZA S. E. A., AZAM A., TSANG Y. W., ET AL.: Hover-Net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images. *Medical Image Analysis* 58 (2019). DOI: 10.48550/arXiv.1812.06499. 6
- [HCT*19] HOSSEINI M. S., CHAN L., TSE G., TANG M., DENG J., ET AL.: Atlas of digital pathology: A generalized hierarchical histological tissue type-annotated database for deep learning. In *Proc. CVPR* (2019), IEEE, pp. 11747–11756. DOI: 10.1109/CVPR.2019.01202. 6
- [HMST21] HUMPHRIES M., MAXWELL P., SALTO-TELLEZ M.: QuPath: The global impact of an open source digital pathology system. *Computational and Structural Biotechnology Journal* 19 (2021), 852–859. DOI: 10.1016/j.csbj.2021.01.022. 3, 5, 6
- [HRA*20] HANNA M. G., REUTER V. E., ARDON O., KIM D., SIRINTRAPUN S. J., SCHÜFFLER P. J., BUSAM K. J., SAUTER J. L., BROGI E., TAN L. K.: Validation of a digital pathology system including remote review during the COVID-19 pandemic. *Modern Pathology* 33, 11 (2020), 2115–2127. DOI: 10.1038/s41379-020-0601-5. 1
- [HZRS16] HE K., ZHANG X., REN S., SUN J.: Deep residual learning for image recognition. In *Proc. CVPR* (2016), IEEE, pp. 770–778. DOI: 10.1109/CVPR.2016.90. 5
- [IEPJ16] IRTAN S., EHRLICH P. F., PRITCHARD-JONES K.: Wilms tumor: "state-of-the-art" update, 2016. *Seminars in pediatric surgery* 25, 5 (2016), 250–256. DOI: 10.1053/j.sempedsurg.2016.09.003. 7
- [IHM*16] IANDOLA F. N., HAN S., MOSKEWICZ M. W., ASHRAF K., DALLY W. J., ET AL.: SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and < 0.5 mb model size, 2016. DOI: 10.48550/arXiv.1602.07360. 5
- [JCS*22] JOCHER G., CHAURASIA A., STOKEN A., BOROVEC J., NANOCODE012, ET AL.: ultralytics/yolov5. *Zenodo record* 6222936, 2022. 5
- [JFGS17] JENSEN M. A., FERRETTI V., GROSSMAN R. L., STAUDT L. M.: The nci genomic data commons as an engine for precision medicine. *Blood, The Journal of the American Society of Hematology* 130, 4 (2017), 453–459. DOI: 10.1109/TVCG.2019.2934555. 6
- [JKW*21] JESSUP J., KRUEGER R., WARCHOL S., HOFFER J., MUHLICH J., RITCH C. C., GAGLIA G., COY S., CHEN Y.-A., LIN J.-R., ET AL.: Scope2screen: Focus+ context techniques for pathology tumor assessment in multivariate image data. *IEEE TVCG* 28, 1 (2021), 259–

269. DOI: [10.1109/TVCG.2021.3114786](https://doi.org/10.1109/TVCG.2021.3114786). 3
- [JLP*19] JIANG J., LARSON N. B., PRODDUTURI N., FLOTTE T. J., HART S. N.: Robust hierarchical density estimation and regression for re-stained histological whole slide image co-registration. *PLOS One* 14, 7 (2019). DOI: [10.1371/journal.pone.0220074](https://doi.org/10.1371/journal.pone.0220074). 2
- [JPA*21] JAUME G., PATI P., ANKLIN V., FONCUBIERTA A., GABRANI M.: HistoCartography: A toolkit for graph analytics in digital pathology. In *Proc. MICCAI Workshop on Computational Pathology (PMLR 156)* (2021), pp. 117–128. DOI: [10.48550/arXiv.2107.10073](https://doi.org/10.48550/arXiv.2107.10073). 3, 5
- [JPB*21] JAUME G., PATI P., BOZORGTABAR B., FONCUBIERTA A., ANNICIELLO A. M., ET AL.: Quantifying explainers of graph neural networks in computational pathology. In *Proc. CVPR* (2021), IEEE, pp. 8106–8116. DOI: [10.1109/CVPR46437.2021.00801](https://doi.org/10.1109/CVPR46437.2021.00801). 3
- [KBJ*19] KRUEGER R., BEYER J., JANG W.-D., KIM N. W., SOKOLOV A., SORGER P. K., PFISTER H.: Facetto: Combining unsupervised and supervised learning for hierarchical phenotype analysis in multi-channel image data. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2019), 227–237. DOI: [10.1109/TVCG.2019.2934547](https://doi.org/10.1109/TVCG.2019.2934547). 3
- [KMRA*15] KATHER J. N., MARX A., REYES-ALDASORO C. C., SCHAD L. R., ET AL.: Continuous representation of tumor microvessel density and detection of angiogenic hotspots in histological whole-slide images. *Oncotarget* 6, 22 (2015). DOI: [10.18632/oncotarget.4383](https://doi.org/10.18632/oncotarget.4383). 2
- [KVS*17] KUMAR N., VERMA R., SHARMA S., BHARGAVA S., VAHADANE A., ET AL.: A dataset and a technique for generalized nuclear segmentation for computational pathology. *IEEE Trans. Medical Imaging* 36, 7 (2017), 1550–1560. DOI: [10.1109/TMI.2017.2677499](https://doi.org/10.1109/TMI.2017.2677499). 6
- [LGB*20] LU W., GRAHAM S., BILAL M., RAJPOOT N., MINHAS F.: Capturing cellular topology in multi-gigapixel pathology images. In *Proc. CVPR Workshops* (2020), IEEE, pp. 260–261. DOI: [10.1109/CVPRW50498.2020.00138](https://doi.org/10.1109/CVPRW50498.2020.00138). 2
- [LLH*18] LIU J., LICHTENBERG T., HOADLEY K. A., POISSON L. M., LAZAR A. J., ET AL.: An integrated tcga pan-cancer clinical data resource to drive high-quality survival outcome analytics. *Cell* 173, 2 (2018), 400–416. DOI: [10.1158/1538-7445.AM2018-3287](https://doi.org/10.1158/1538-7445.AM2018-3287). 6
- [LML*07] LIU J., MA W., LIU F., HU Y., YANG J., XU X.: Study and application of medical image visualization technology. In *Int'l Conf. on Digital Human Modeling* (2007), ICDHM, pp. 668–677. DOI: [10.1007/978-3-540-73321-8_77](https://doi.org/10.1007/978-3-540-73321-8_77). 2
- [LWC*21] LU M. Y., WILLIAMSON D. F., CHEN T. Y., CHEN R. J., BARBIERI M., ET AL.: Data-efficient and weakly supervised computational pathology on whole-slide images. *Nature Biomedical Engineering* 5, 6 (2021), 555–570. DOI: [10.1038/s41551-020-00682-w](https://doi.org/10.1038/s41551-020-00682-w). 3
- [MGP*22] MANZ T., GOLD I., PATTERSON N. H., MCCALLUM C., KELLER M. S., HERR B. W., BÖRNER K., SPRAGGINS J. M., GEHLENBORG N.: Viv: multiscale visualization of high-resolution multiplexed bioimaging data on the web. *Nature Methods* 19, 5 (2022), 515–516. DOI: [10.1038/s41592-022-01482-7](https://doi.org/10.1038/s41592-022-01482-7). 3
- [PKR*19] POPE P. E., KOLOURI S., ROSTAMI M., MARTIN C. E., HOFFMANN H.: Explainability methods for graph convolutional neural networks. In *Proc. CVPR* (2019), IEEE, pp. 10772–10781. DOI: [10.1109/CVPR.2019.01103](https://doi.org/10.1109/CVPR.2019.01103). 3
- [POR*19] PELL R., OIEN K., ROBINSON M., PITMAN H., RAJPOOT N. M., ET AL.: The use of digital pathology and image analysis in clinical trials. *The Journal of Pathology: Clinical Research* 5, 2 (2019), 81–90. DOI: [10.1002/cjp2.127](https://doi.org/10.1002/cjp2.127). 1, 2
- [QRW*19] QU H., RIEDLINGER G., WU P., HUANG Q., YI J., ET AL.: Joint segmentation and fine-grained classification of nuclei in histopathology images. In *Proc. IEEE ISBI* (2019), IEEE, pp. 900–904. DOI: [10.1109/ISBI.2019.8759457](https://doi.org/10.1109/ISBI.2019.8759457). 5
- [SCD*17] SELVARAJU R. R., COGSWELL M., DAS A., VEDANTAM R., PARIKH D., ET AL.: Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *Proc. ICCV* (2017), IEEE, pp. 618–626. DOI: [10.1109/ICCV.2017.74](https://doi.org/10.1109/ICCV.2017.74). 3, 7
- [SCM21] SRINIDHI C. L., CIGA O., MARTEL A. L.: Deep neural network models for computational histopathology: A survey. *Medical Image Analysis* 67 (2021). DOI: [10.1016/j.media.2020.101813](https://doi.org/10.1016/j.media.2020.101813). 2
- [SCS*20] SAHASRABUDHE M., CHRISTODOULIDIS S., SALGADO R., MICHELIS S., LOI S., ANDRÉ F., PARAGIOS N., VAKALOPOULOU M.: Self-supervised nuclei segmentation in histopathological images using attention. In *Proc. MICCAI* (2020), pp. 393–402. DOI: [10.1007/978-3-030-59722-1_38](https://doi.org/10.1007/978-3-030-59722-1_38). 2
- [SFG*20] STOLTZFUS C. R., FILIPEK J., GERN B. H., OLIN B. E., LEAL J. M., WU Y., LYONS-COHEN M. R., HUANG J. Y., PAZ-STOLTZFUS C. L., PLUMLEE C. R., ET AL.: Cytomap: a spatial analysis toolbox reveals features of myeloid cell organization in lymphoid tissues. *Cell reports* 31, 3 (2020), 107523. DOI: [10.1101/769877](https://doi.org/10.1101/769877). 3
- [SIK*21] SOMARAKIS A., ISSSELSTEIJN M. E., KENKHUIS B., UNEN V. V., LUK S. J., KONING F., WEERD L. v. D., MIRANDA N. F. C. C. D., LELIEVELDT B. P. F., HÖLLT T.: Visual Analysis of Tissue Images at Cellular Level. In *EuroVis 2021 - Dirk Bartz Prize* (2021), Oeltze-Jafra S., Raidou R. G., (Eds.), The Eurographics Association. DOI: [10.2312/evm.20211074](https://doi.org/10.2312/evm.20211074). 3
- [SRT*16] SIRINUKUNWATTANA K., RAZA S. E. A., TSANG Y.-W., SNEAD D. R., CREE I. A., ET AL.: Locality sensitive deep learning for detection and classification of nuclei in routine colon cancer histology images. *IEEE Trans. Medical Imaging* 35, 5 (2016), 1196–1206. DOI: [10.1109/TMI.2016.2525803](https://doi.org/10.1109/TMI.2016.2525803). 5
- [TL19] TAN M., LE Q.: EfficientNet: Rethinking model scaling for convolutional neural networks. In *Proc. ICML* (2019), pp. 6105–6114. DOI: [10.48550/arXiv.1905.11946](https://doi.org/10.48550/arXiv.1905.11946). 5
- [TWM*21] THEELKE L., WILM F., MARZAH L. C., BERTRAM C. A., KLOPFLEISCH R., ET AL.: Iterative cross-scanner registration for whole slide images. In *Proc. ICCV Workshops* (2021), IEEE, pp. 582–590. DOI: [10.1109/ICCVW54120.2021.00071](https://doi.org/10.1109/ICCVW54120.2021.00071). 2
- [VS10] VUJANIĆ G. M., SANDSTEDT B.: The pathology of wilms' tumour (nephroblastoma): the international society of paediatric oncology approach. *Journal of clinical pathology* 63, 2 (2010), 102–109. DOI: [10.1136/jcp.2009.064600](https://doi.org/10.1136/jcp.2009.064600). 4
- [Vu19] VU Q. D. T.: Methods for segmentation and classification of digital microscopy tissue images. *Frontiers in Bioengineering and Biotechnology* (2019), 53. DOI: [10.3389/fbioe.2019.00053](https://doi.org/10.3389/fbioe.2019.00053). 6
- [WCH*22] WU Y., CHENG M., HUANG S., PEI Z., ZUO Y., ET AL.: Recent advances of deep learning for computational histopathology: Principles and applications. *Cancers* 14, 5 (2022), 1199. DOI: [10.3390/cancers14051199](https://doi.org/10.3390/cancers14051199). 2
- [Wea10] WEAVER D. L.: Pathology evaluation of sentinel lymph nodes in breast cancer: protocol recommendations and rationale. *Modern Pathology* 23, 2 (2010), S26–S32. DOI: [10.1038/modpathol.2010.36](https://doi.org/10.1038/modpathol.2010.36). 4
- [WKS*21] WAGNER S. J., KHALILI N., SHARMA R., BOXBERG M., MARR C., ET AL.: Structure-preserving multi-domain stain color augmentation using style-transfer with disentangled representations. In *Proc. MICCAI* (2021), MICCAI, pp. 257–266. DOI: [10.1007/978-3-030-87237-3_25](https://doi.org/10.1007/978-3-030-87237-3_25). 4, 5
- [YPPH14] YEH F.-C., PARWANI A. V., PANTANOWITZ L., HO C.: Automated grading of renal cell carcinoma using whole slide imaging. *Journal of pathology informatics* 5 (2014). DOI: [10.4103/2153-3539.137726](https://doi.org/10.4103/2153-3539.137726). 2
- [YYH*14] YEH F.-C., YE Q., HITCHENS T. K., WU Y. L., PARWANI A. V., ET AL.: Mapping stain distribution in pathology slides using whole slide imaging. *Journal of Pathology Informatics* 5 (2014). DOI: [10.4103/2153-3539.126140](https://doi.org/10.4103/2153-3539.126140). 2
- [ZGL*21] ZELIC R., GIUNCHI F., LIANAS L., MASCIA C., ZANETTI G., ANDRÉN O., FRIDFELDT J., CARLSSON J., DAVIDSSON S., MOLINARO L., VINCENT P., RICHIARDI L., AKRE O., FIORENTINO M., PETERSSON A.: Interchangeability of light and virtual microscopy for histopathological evaluation of prostate cancer. *Nature Scientific Reports* 11, 3257 (2021). DOI: [10.1038/s41598-021-82911-z](https://doi.org/10.1038/s41598-021-82911-z). 1, 3
- [ZKL*16] ZHOU B., KHOSLA A., LAPEDRIZA A., OLIVA A., TORRALBA A.: Learning deep features for discriminative localization. In *Proc. CVPR* (2016), pp. 2921–2929. DOI: [10.1109/CVPR.2016.319](https://doi.org/10.1109/CVPR.2016.319). 3