



A Survey on Cell Nuclei Instance Segmentation and Classification: Leveraging Context and Attention

João D. Nunes^{a,b,*}, Diana Montezuma^{c,d,e}, Domingos Oliveira^c, Tania Pereira^{a,f}, Jaime S. Cardoso^{a,b}

^aINESC TEC - Institute for Systems and Computer Engineering, Technology and Science, R. Dr. Roberto Frias, Porto, 4200-465, Portugal

^bUniversity of Porto - Faculty of Engineering, R. Dr. Roberto Frias, Porto, 4200-465, Portugal

^cIMP Diagnostics, Praça do Bom Sucesso, 4150-146 Porto, Portugal

^dCancer Biology and Epigenetics Group, Research Center of IPO Porto (CI-IPOP) / RISE@CI-IPOP (Health Research Network), Portuguese Oncology Institute of Porto (IPO Porto) / Porto Comprehensive Cancer Center (Porto.CCC), R. Dr. António Bernardino de Almeida, 4200-072, Porto, Portugal

^eDoctoral Programme in Medical Sciences, School of Medicine and Biomedical Sciences – University of Porto (ICBAS-UP), Porto, Portugal

^fFCTUC - Faculty of Science and Technology, University of Coimbra, Coimbra, 3004-516, Portugal

ARTICLE INFO

Article history:

Received 1 May 2013

Received in final form 10 May 2013

Accepted 13 May 2013

Available online 15 May 2013

Communicated by S. Sarkar

2000 MSC: 41A05, 41A10, 65D05, 65D17

Keywords: Artificial Neural Networks, Context, Attention, Nuclei Instance Segmentation and Classification, Computational Pathology

ABSTRACT

Nuclear-derived morphological features and biomarkers provide relevant insights regarding the tumour microenvironment, while also allowing diagnosis and prognosis in specific cancer types. However, manually annotating nuclei from the gigapixel Hematoxylin and Eosin (H&E)-stained Whole Slide Images (WSIs) is a laborious and costly task, meaning automated algorithms for cell nuclei instance segmentation and classification could alleviate the workload of pathologists and clinical researchers and at the same time facilitate the automatic extraction of clinically interpretable features for artificial intelligence (AI) tools. But due to high intra- and inter-class variability of nuclei morphological and chromatic features, as well as H&E-stains susceptibility to artefacts, state-of-the-art algorithms cannot correctly detect and classify instances with the necessary performance. In this work, we hypothesise context and attention inductive biases in artificial neural networks (ANNs) could increase the performance and generalization of algorithms for cell nuclei instance segmentation and classification. To understand the advantages, use-cases, and limitations of context and attention-based mechanisms in instance segmentation and classification, we start by reviewing works in computer vision and medical imaging. We then conduct a thorough survey on context and attention methods for cell nuclei instance segmentation and classification from H&E-stained microscopy imaging, while providing a comprehensive discussion of the challenges being tackled with context and attention. Besides, we illustrate some limitations of current approaches and present ideas for future research. As a case study, we extend both a general (Mask-RCNN) and a customised (HoVer-Net) instance segmentation and classification methods with context- and attention-based mechanisms and perform a comparative analysis on a multicentre dataset for colon nuclei identification and counting.

Although pathologists rely on context at multiple levels while paying attention to specific Regions of Interest (RoIs) when analysing and annotating WSIs, our findings suggest translating that domain knowledge into algorithm design is no trivial task, but to fully exploit these mechanisms in ANNs, the scientific understanding of these methods should first be addressed.

1. Introduction

Cancer is a very heterogeneous disease, and as it manifests differently in each patient, conventional general treatments such as surgery, chemotherapy, radiation, and immunotherapy may not be effective for every cancer subtype, aside from the fact that these treatments are generally considered aggressive and also target healthy tissue. Particularly, cancer subtypes are associated with distinct genetic mutations and molecular and protein alterations. Therefore, when combined with proteomics, transcriptomics, and/or genomics, the visual analysis of tumour microenvironment from digitised Haematoxylin and Eosin (H&E)-stained Whole Slide Images (WSIs) is vital to help clinicians make informed decisions. Analysis of H&E-stained WSI allows the characterization of tumour microenvironment and is of high clinical value for cancer diagnosis, treatment assessment, and prognosis. More specifically, precision and personalized medicine treatments have increasingly benefited from -omics biomarkers identified using these modalities. Furthermore, since H&E-stained WSIs are part of routine clinical practice, image-based phenotyping has the potential to replace complex and expensive diagnostic and prognostic solutions (e.g., molecular profiling) in precision and personalized medicine. In particular, the increasing migration of modern pathology towards digital workflows has allowed the introduction of novel solutions, supported by Artificial Intelligence (AI) algorithms, with unprecedented results (Cui and Zhang, 2021; Baxi *et al.*, 2022; Acs *et al.*, 2020). But computational pathology (CPath) is a broad field that involves many tasks, including diagnosis, tissue characterisation, staging, grading, and disease prognosis (Esteva *et al.*, 2021). Automatic tools that assist in these tasks usually call for Computer Vision (CV) and Machine Learning (ML) algorithms entailing regression analysis, classification, semantic segmentation, object detection, or instance segmentation and classification. Depending on the objective of the task, multiple CV and ML models can be considered, but deep neural networks usually yield better performance and allow generalization to larger and more complex datasets (Baxi *et al.*, 2022). Nonetheless, a promising strategy might be to incorporate prior clinical knowledge (Marini *et al.*, 2022; Neto *et al.*, 2022; Ahmedt-Aristizabal *et al.*, 2022) into algorithm design. CPath is thus a multidisciplinary field where pathologists and ML scientists must work in close collaboration towards the design of algorithms with clinical value and that best exploit the potentialities of WSI data with human-interpretable features.

One of these clinically important tasks is the cell nuclei instance segmentation and classification¹. This is relevant, for example, to the quantification of cells, characterization of the tumour microenvironment, and, in certain cancer types (e.g., breast cancer, lung cancer), nuclei-derived features have shown useful as diagnosis and prognosis biomarkers (Lu *et al.*, 2020a; Yan *et al.*, 2022). However, nuclei instance segmentation and

classification is not a trivial task. In the first place, there is intrinsically high intra- and inter-subject variability in tissue appearance and nuclei organization, even if we consider nuclei of the same cell type. And as different nuclei types can be morphologically similar, this task becomes even more challenging. Moreover, the digitised histology slides are highly prone to stain variability (Tellez *et al.*, 2018) and artefacts that negatively affect the generalization (Campanella *et al.*, 2019; Rolls and Farmer, 2008). Another bottleneck resides in the dimensionality of WSIs, which can be in the order of gigapixels, an undesirable attribute in hardware-limited settings. Moreover, each WSI could have thousands of nuclei instances, which in voluminous datasets could scale to millions of nuclei (or billions of pixels) to annotate. Therefore, manually annotating every nuclei instance is practically impossible, aside from high inter-observer variability and annotation inconsistencies that hinder Deep Learning (DL) models performance. To cope with these challenges, alternative and more label-efficient paradigms, like Self-supervised Learning (SSL), Weakly Supervised Learning (WSL), or Semi-supervised Learning (SmSL) naturally emerge in an attempt to design ML models robust to partially or incorrectly labelled data.

In their daily practice, to overcome most of the previously described challenges in WSI analysis and nuclei quantification, clinicians usually resort to WSI visualization tools that allow them to scroll through the tissue image while zooming in and out, and while paying attention to different RoIs and tissue concepts (e.g., cells, nuclei, etc.). Through this process, clinicians aggregate information at multiple levels to make a decision regarding nuclear boundaries and types. In essence, this is a context-based decision process that we argue translates to increased annotation quality and that we hypothesise should be exploited to devise inductive biases in neural networks. Therefore, this survey reviews works with inductive biases for the spatial and semantic correlation between tissue concepts, i.e., with context inductive biases or that explore inter-instance relations. A representation encoding the co-occurrence and long-range dependencies of different objects (i.e., instances) could be more discriminative than a representation that only considers intra-instance features. Especially in the presence of ambiguities (e.g., due to noisy or blurred data, occluded objects, small instances, etc.). These ideas have already been investigated in the CV field (Naseer *et al.*, 2021). In turn, some context variables are confounders that bias DL models to non-stable features (Dreyer *et al.*, 2023) (e.g., sky or grasslands to discriminate between birds and cows). But by using the self-attention mechanism to learn long-range dependencies (i.e., contextual information) between patch features, Vision Transformer (ViT)-based models could be less sensitive to spurious correlations when compared with Convolutional Neural Networks (ConvNets) (Zhang *et al.*, 2022a). Overall, it is not clear that contextual information will always improve object identification performance (Zhang *et al.*, 2020; Liu *et al.*, 2022a; Ouyang *et al.*, 2023). When and how to use context priors is, in general, an unsolved topic in CV and related subfields like cell nuclei instance segmentation and classification from H&E-stained brightfield microscopy images. This survey is thus an

*Corresponding author: e-mail.: joao.d.fernandes@inesctec.pt

¹Throughout this work, we use "cell nuclei instance segmentation and classification" as an umbrella term that encompasses several tasks, namely: object detection, instance counting, localization, semantic segmentation, and instance segmentation and classification.

attempt to help further researchers clarify the role of context and attention in nuclei instance segmentation and classification from H&E-stained brightfield microscopy imaging. Such knowledge could then be used to design more robust DL models.

Besides this introduction, this work is structured as follows: In section 2, we review previous related works and present our main contributions. In section 3, we provide a clinical motivation for cell nuclei instance segmentation and classification, and in section 4, we review the main metrics, datasets, and fundamental neural network architectures used in CPath for cell nuclei instance segmentation and classification. But for a comprehensive analysis of context- and attention-based mechanisms in cell nuclei instance segmentation and classification, in section 5, we start by defining these two terms, review the taxonomies proposed in the literature for each paradigm, and discuss types and levels of context and attention. Following these definitions, we then review fundamental context- and attention-based mechanisms in CV, including neural network modules, mechanisms, and even complete architectures that depend on the attention mechanism to incorporate multiple context levels, i.e., ViTs. The strategies suggested in these publications are frequently cornerstones of algorithms designed for specific domains and applications. But as these strategies are general, and often optimized for everyday scenes and objects (e.g., images of cats, cars, houses, etc.), due to the concept and covariate shift, they frequently fail in medical imaging data without fine-tuning or minor modifications. Therefore, we believe it is crucial to also understand how context and attention-based strategies are being adapted to solve medical imaging tasks. In particular, to help our discussion on cell nuclei instance segmentation and classification, in section 5.2, we review works that resemble that task, hoping that knowledge from other domains could help CPath researchers understand the potentialities of context and attention-based mechanisms. Ultimately, that knowledge could help design novel solutions tailored for CPath applications. However, the main scope of this work is a systematic review of context- and attention-based mechanisms in cell nuclei instance segmentation and classification from H&E-stained brightfield microscopy imaging (section 6). We provide a detailed analysis of state-of-the-art works, discuss how context- and attention-based mechanisms have been used to address key challenges in cell nuclei instance segmentation and classification, identify promising research directions, and present alternative research ideas. In fact, for a more complete understanding of the impact of context- and attention-based mechanisms in nuclei instance segmentation and classification, in section 7, we provide practical insights from a case study of how adapting Mask-RCNN (He et al., 2017) and HoverNet (Graham et al., 2019b) with context- and attention-based mechanisms affects the generalization of a nuclei instance segmentation and classification task. In section 8, we conclude with some final remarks regarding our findings and future directions.

2. Related Work

Medical imaging is an immensely broad field and there is an increasing interest in CV algorithms to analyse medical

data, primarily owing to the growing availability of these imaging modalities, especially in fields like dermatology, radiology, and pathology (Esteva et al., 2021). An exhaustive review of context-, and attention-based CV algorithms for medical data is thus impractical and beyond the scope of this work. Furthermore, there is already an extensive body of research discussing DL for medical image analysis (Gonçalves et al., 2022; He et al., 2022a; Zhao et al., 2022b; van der Velden et al., 2022; Rao et al., 2021; Esteva et al., 2021; Lundervold and Lundervold, 2019; Liu et al., 2019b), but we differentiate from the prior art by, in addition to reviewing nuclei instance segmentation and classification methods, summarizing works on context- and attention-based mechanisms that have been successful in medical imaging tasks that share similarities with nuclei instance segmentation and classification. More specifically, we focus on methods that incorporate several of the following features: object detection, object segmentation (instance and semantic), instance classification, imbalanced instances, small and dense object detection or segmentation, or 2D inputs. We argue that such methods can serve as motivation for solutions tailored to pathology. Nonetheless, the main focus of the present review is to provide a systematic survey on context and attention-based mechanisms in DL-based nuclei instance segmentation and classification from H&E-stained brightfield microscopy imaging. This review arises from a reflection on the methods used by pathologists during the analysis and annotation of tissue slides, as well as from domain knowledge on nuclei and tissue microenvironment in H&E-stained WSIs, which are context-rich.

Several authors aggregate and summarize the main works on cell nuclei instance segmentation and classification, complementing their review with a critical analysis and discussion of future directions. Basu et al. (2023) review, from 2017 to 2021, DL-based algorithms for cell nuclei segmentation with an emphasis on analyzing and comparing several models, namely, U-Net (Ronneberger et al., 2015), SCPP-Net (Chanchal et al., 2021), Sharp U-Net (Zunair and Ben Hamza, 2021), and LiverNet (Aatresh et al., 2021). Furthermore, the work highlights the importance of nuclei segmentation for CPath including cell counting, diagnosis, and morphological analysis Basu et al. (2023). Nasir et al. (2022), on the other hand, discuss the main public datasets, grand-challenge algorithms, and future directions in gland and nuclei instance segmentation. The authors address works on both conventional ML with manual feature extraction, and DL, including encoder-decoder neural networks and generative models. Similarly to our work, they review several attention-based DL methods, but no focus is given to context-based strategies or vision transformers. Srinidhi et al. (2021) provide a comprehensive review of deep neural networks for CPath. Indeed, the scope of their work is large, and although the paper mentions a few works on context- and attention-based mechanisms in DL for CPath, the scientific understanding of these mechanisms is still limited, and there is also the need to assess the usefulness of these methods specifically in nuclei instance segmentation and classification. Lagree et al. (2021), on the other hand, highlight several DL methods for cell nuclei segmentation in breast tissue slides, trained on a multi-organ dataset. The methodology considers U-Net (Ron-

neberger et al., 2015), Mask-RCNN (He et al., 2017), and an ensemble of various U-Net-like models. But as their results underline the challenges of highly crowded nuclei, the authors propose an ensemble network (GB-UNet) and show that transfer learning is an effective strategy for nuclei segmentation in breast tissue WSIs (Lagree et al., 2021). Furthermore, Krithiga and Geetha (2021) review different algorithms in breast CPath, namely methods for cell and nuclei instance segmentation and classification, feature extraction, image enhancement, and tissue classification, while Hayakawa et al. (2021) review nuclei instance segmentation and classification works with prominence on conventional ML and on pre-processing and post-processing techniques.

In essence, our work differentiates from related reviews as we focus on surveying context and attention-based mechanisms not only in cell nuclei instance segmentation and classification but also in related medical imaging subdomains which could provide new insights regarding the effectiveness of these methods and serve as inspiration for the development of novel solutions. We consider this to be relevant as, regardless of the imaging technique, some common challenges of the medical imaging domain make solutions tailored to other imaging modalities of potential value to H&E-stained brightfield microscopy imaging. Nonetheless, not only are digitized H&E-stained brightfield microscopy images (mostly H&E-stained WSIs), widely available, as also H&E staining is one of the brightfield microscopy imaging techniques that provide the most detailed information, allowing clear visualization of tissue structures like the cytoplasm, nucleus, organelles and extra-cellular components. Therefore, this modality offers the best cost-effectiveness when compared with other microscopy techniques, which translates to a CPath research focus on methods for H&E-stained image analysis. For these reasons, we give special attention to this imaging technique. Namely, we do a systematic review of context and attention mechanisms for DL-based cell nuclei instance segmentation and classification from H&E-stained brightfield microscopy images, including WSIs and tiles, and provide a comprehensive discussion of how the different context levels have been addressed in the literature for cell nuclei instance segmentation and classification from H&E-stained brightfield microscopy images, while illustrating the different attention mechanisms that have been explored and their intended use like enriching the estimated feature maps with contextual information or reducing sensitivity to irrelevant background information.

Following the limitations of prior art, our main contributions are as follows:

- A review of the main context- and attention-based DL algorithms for CV applications.
- A review of context- and attention-based works in medical imaging dense prediction tasks².
- A survey of 44 papers on context- and attention-based

methods in cell nuclei instance segmentation and classification from H&E-stained brightfield microscopy (including WSIs and tiles).

- A case study to empirically evaluate how extending reference cell nuclei instance segmentation and classification models with context and attention-based mechanisms affects generalization.

3. A Clinical Perspective

Nuclei evaluation has long been relevant in Anatomical Pathology, as nuclei hold critical information regarding the identity and characteristics of the cell. Different cell types display unique nuclear features that help pathologists classify cells accurately. Furthermore, morphological alterations in the nucleus occur frequently in cancer, and other diseases, and can provide important information during diagnostic interpretation (Chow et al., 2012; Gurcan et al., 2009). In fact, for more than a century, pathologists have relied on morphological abnormalities of the nucleus as a crucial feature differentiating benign from malignant cells (Fischer, 2020). These abnormalities comprise nuclear enlargement, elevated nuclear-to-cytoplasmic ratio, nuclear membrane irregularities, hyperchromasia, and altered chromatin distribution (Fischer, 2020). Still today, the Nottingham histologic score for breast cancer grading relies on three parameters: tubule formation, mitotic activity and nuclear pleomorphism (of Tumours Editorial Board, 2019). The higher the score (1-3), the more aggressively the tumour is expected to behave. Recent work, by Jaroensri et al. (2022), has investigated the use of DL for histologic grading in breast cancer. Another example where nuclei evaluation is important is renal cell carcinoma, as the ISUP/WHO score is based on nucleolar prominence/eosinophilia for grades 1-3, and nuclear anaplasia for grade 4 (Delahunt et al., 2019). In reality, nuclei evaluation is extremely relevant for most neoplastic diagnoses and plays a role in assessing disease severity, even if not included in a formal score. Moreover, nuclei evaluation holds importance beyond grading in diagnostics. For instance, diagnosing papillary thyroid carcinoma requires identifying distinctive nuclear features, concretely enlarged nuclei with irregular membranes, nuclear grooves and the characteristic optically clear nuclei (“Orphan Annie Eye” nuclei) (Ibrahim et al., 2020). There are already studies investigating ML methods for thyroid cancer classification, applying nuclear segmentation (Böhland et al., 2021). Apart from epithelial cells, nuclei evaluation is also relevant to assess/distinguish other cell types. In this context, scoring Tumour-Infiltrating Lymphocytes (TILs) has proven prognostic value for breast cancer (Salgado et al., 2015) and also in other tumour types (Hendry et al., 2017). As TILs visual scoring is a time-consuming task, limited by intra and interobserver variability, software-guided TIL analysis is currently under research, with promising results (Sun et al., 2021; Aung et al., 2022; Zhang et al., 2022c). Another domain that relies on nuclear assessment is Immunohistochemical (IHC) stain scoring, as some IHC stains target and color the cellular nuclei. Pathologists visually evaluate immunostains and estimate the percentage and intensity of staining, which encompasses diagnostic

²A set of paradigms in computer vision that demand the mapping from input images to elaborate output structures, including instance segmentation and classification, semantic segmentation, and object detection and localization.

and predictive value. As such, exploring nuclear segmentation tasks has particular relevance in this setting, with several studies addressing this topic (Priego-Torres *et al.*, 2022; Mahanta *et al.*, 2021; Lejeune *et al.*, 2013). Moreover, multiple AI/ML-tools to automatically analyze immunostains are already commercially available. Lastly, nuclear evaluation relevance exceeds oncologic diseases, as Pathologists use nuclear evaluation as part of their diagnostic toolkit in a multitude of diseases (inflammatory, infectious, degenerative, etc.), beyond just cancer. In conclusion, by examining the cell nucleus, pathologists uncover valuable insights into disease processes, enabling accurate diagnoses and appropriate follow-up recommendations for patients, underscoring the importance of ML methods focusing on nuclear segmentation and classification.

4. Preliminaries

In this section, we briefly discuss the main datasets, and relevant cell nuclei instance segmentation and classification works that, although not explicitly incorporating context and attention mechanisms, are notable in the field. We also present the most common metrics to evaluate the quality of algorithms to solve dense prediction tasks like cell nuclei instance segmentation and classification (e.g., object detection, semantic segmentation, instance segmentation and classification, etc.).

4.1. Datasets

Table 1 summarizes the main annotated datasets for cell nuclei instance segmentation and classification. Most of these datasets were initially proposed as part of ML competitions (Kumar *et al.*, 2020; Sirinukunwattana *et al.*, 2017; Graham *et al.*, 2024; Caicedo *et al.*, 2019; Verma *et al.*, 2021), whose primary goal is to gather a community of ML researchers to work on the same problem and to advance the state-of-the-art on nuclei instance segmentation and classification. In fact, data and coding competitions often translate to algorithmic performance breakthroughs, but a limitation is that little attention is paid to solving clinically relevant problems (Varoquaux and Cheplygina, 2022). Nonetheless, the competition data is often publicly shared after the end of the challenges, meaning datasets become easily accessible for future research and benchmarking. Top-performing ML models are thus often trained and evaluated in one or several of these datasets. Moreover, although most works resorting to these data focus on the task of nuclei instance segmentation and classification, there are other emerging use cases like self-supervised representation learning (Ciga *et al.*, 2022a; Song *et al.*, 2023) (KUMAR (Kumar *et al.*, 2020); CoNSeP (Graham *et al.*, 2019b)), disease diagnosis (Tourniaire *et al.*, 2023; Bai *et al.*, 2022; Zhou *et al.*, 2019b) (DigestPath (Da *et al.*, 2022); PanNuke (Gamper *et al.*, 2020a); CoNSeP (Graham *et al.*, 2019b)) and prognosis (Dodgington *et al.*, 2021; Xie *et al.*, 2022b) (KUMAR (Kumar *et al.*, 2020)), characterization of the tissue and tumour microenvironment (Lu *et al.*, 2020b; Zhao *et al.*, 2023) (PanNuke (Gamper *et al.*, 2020a)), Domain Adaptation (DA) (Li *et al.*, 2022) (DigestPath (Da *et al.*, 2022)), explainable and interpretable AI (M. *et al.*, 2020; Jaume *et al.*, 2021; Mahapatra *et al.*, 2022) (GlaS

(Sirinukunwattana *et al.*, 2017); KUMAR (Kumar *et al.*, 2020)), and biomarker discovery (Fang *et al.*, 2023) (PanNuke (Gamper *et al.*, 2020a)). Nonetheless, as these data originate from diverse geographic locations, clinical sites, scanners, and tissue types, these are also of potential value to develop and evaluate DL methods for Domain Generalization (DG), continual learning, and even federated learning. Other promising use cases for H&E-stained images and corresponding nuclei annotations are hardware-efficient DL, multimodal ML (e.g., paired text-imaging data), and identification of imaging biomarkers.

Regarding the evaluation protocol, although the scientific publications that first introduce these data and corresponding nuclei instance segmentation and classification algorithms report the metrics they use, we identify two limitations, namely, some metrics may not be appropriate to assess the segmentation quality of small instances (Cheng *et al.*, 2021a; Foucart *et al.*, 2023), and the heterogeneity across datasets of the metrics used to assess the quality of nuclei instance segmentation and classification methods. Besides, although splits are chosen carefully to not leak data from the training set to the validation and test splits, for most datasets only a portion of the tiles from the WSI are selected. Following this observation, and noticing that a common practice in nuclei instance segmentation and classification algorithms is to quantify tile-level results, the question occurs regarding the performance for the complete WSI. Besides, assuming not all tiles are effectively required to achieve a diagnosis, the question occurs regarding the optimal tile selection strategy for both training and testing. Although there is some work on sampling strategies for training DL models for colorectal dysplasia grading (Neto *et al.*, 2023), to the best of our knowledge, there is no similar strategy in other cancer types, especially considering segmentation-oriented paradigms. This is a relevant remark as reducing the subset of tiles used for training decreases a lot the computation costs, but if strong assumptions are imposed, it could also hinder DL models generalization, especially in out of distribution (OOD) data. Answering these questions could be of service to the development of nuclei instance segmentation and classification algorithms with increased performance and generalization, while also increasing data and models' reusability and trustworthiness.

Another concern regarding the data is sample selection bias, an undesirable phenomenon that occurs when certain elements or groups from a population are more likely to be included in a sample than others. This translates to datasets that are not representative of the target population, leading to flawed ML algorithms that favour spurious correlations instead of stable features. In CPath, batch effects are a common form of sample selection bias. Batch effects refer to variation between two batches of data due to differences in slide preparation, or digitizer. For instance, as different providers adopt different hardware, algorithms for image manipulation, and file formats (Janowczyk *et al.*, 2019), it is expected digitizer-related batch effects, including variations in scale and colour. Besides, colour batch effects can also occur due to differences in staining protocol, room temperature during stain preparation, the clinician that prepared the slide, variations of thickness in tissue sections, and different lots or manufacturers of staining reagents

(Niazi *et al.*, 2019). These batch effects can therefore be site-specific, thus contributing to sampling bias, aside from different patient demographics. Indeed, recent work (Asilian Bidgoli *et al.*, 2022) demonstrates that even in datasets with samples from multiple sources, DL models become biased towards site-specific signatures, learning to discriminate the cancer type from spurious associations. To help identify batch effects and sample selection bias, we suggest future dataset releases should be accompanied by detailed reporting of data acquisition protocols, along with information on preprocessing, and storage procedures. Ultimately, sample selection bias is one of the potential causes of hidden stratification. Hidden stratification occurs when a model has good average performance in the entire test set, but fails on some specific subgroups characterized by a combination of low prevalence of disease subtype, incorrect or inconsistent annotations, sparse discriminative features, and spurious correlations (Oakden-Rayner *et al.*, 2020). Hidden stratification is thus closely related to domain shift, helping clarify why models with good performance in the independent and identically distributed (IID) setting fail on OOD samples. The most effective system to mitigate hidden stratification is to provide, along with the data, metadata to help identify the relevant subsets. For instance, certain samples that originate from the TCGA database (Clark *et al.*, 2013) have some demographic variables available like gender, race, age at diagnosis, and smoking habits. These variables could help identify model biases, however, these are not provided along with the annotated nuclei subsets from TCGA (Clark *et al.*, 2013), a major drawback impeding a thorough evaluation of nuclei instance segmentation and classification algorithms. In fact, for safe ML deployments in medicine, measurement and reporting of hidden stratification effects should be part of the models' evaluation protocol (Oakden-Rayner *et al.*, 2020). However, we observe this is not addressed on datasets for cell nuclei instance segmentation and classification, as these lack metadata to help detect biases and identify relevant subsets. One form of overcoming this limitation is unsupervised learning, namely clustering techniques which could be used to both identify relevant subgroups of the data, or as part of evaluation protocols to assess human annotation of existing subclasses (Oakden-Rayner *et al.*, 2020).

A final consideration regards the annotation pipeline. While most datasets report the annotations were performed, refined, and reviewed by multiple experienced clinical annotators, DL for nuclei instance segmentation (and for medical image analysis in general) is a multidisciplinary field that demands iterative, and close collaboration between ML researchers and pathologists (Montezuma *et al.*, 2023; Graham *et al.*, 2021a) so as to reduce mislabelling or systematic errors that undermine the performance of AI algorithms. On the other hand, the lack of standard annotation protocols could translate to annotation bias as a result of the annotator's (e.g., pathologist) experience, judgment, level of fatigue, difficulty in annotating specific tissue concepts (e.g., cell nuclei) and regions (Wahab *et al.*, 2022). Moreover, while AI-driven semi-automatic tools are a valuable resource in reducing the annotation burden, these could bias the pathologist to agree with the tool on tissue regions and concepts that humans cannot easily discriminate (Ibrahim *et al.*,

2020). Moreover, following the FAIR data principles (Wilkinson, 2016) and guidelines for AI in medical imaging (Mongan *et al.*, 2020; Montezuma *et al.*, 2023), qualitative and quantitative evaluations of annotation quality should be done, as well as reporting of quality control measures. Most of the discussed datasets describe some quality control steps, but we find varied practices with different levels of quality assessment, meaning there is no standard protocol to maximize confidence in the annotations (Wahab *et al.*, 2022; Montezuma *et al.*, 2023), nor standard quantitative measures to estimate data and annotation quality in nuclei instance segmentation and classification datasets. In essence, this could translate to limitations in the knowledge drawn from algorithms trained on those data. This is something we argue should be more thoroughly addressed in future public dataset releases.

4.2. Evaluation Metrics

• Intersection Over Union

The Intersection Over Union (IoU), also known as Jaccard Index, quantifies the overlap between the detected/segmented objects and the ground truth. It is defined as

$$IoU = \frac{|P \cap G|}{|P \cup G|} \quad (1)$$

with P and G predicted and target objects.

• Aggregated Jaccard Index

The Aggregated Jaccard Index (AJI) is an extended version of the Jaccard Index designed for instance segmentation tasks. Aside from considering the overlap between each instance in the prediction and target images, it contains a term to penalize false positive detections. It, therefore, simultaneously penalizes object-level (missed detection, false positive detection) and pixel-level (under/over-segmentation) errors. Formally, we have

$$AJI = \frac{\sum_{i=1}^N |P_i \cap G_i|}{\sum_{i=1}^N |P_i \cup G_i| + \sum_{j=1}^M |FP_j|} \quad (2)$$

with P_i and G_i predicted and target objects out of N instances, and FP_j a false positive detection out of M predicted instances.

• Panoptic Quality

The Panoptic Quality (PQ) (Kirillov *et al.*, 2019) is a metric to jointly evaluate the quality of instance and semantic segmentation. It is composed of two terms, one to assess the detection quality, and the other to evaluate the segmentation quality. For a target class, t , PQ is formulated as

$$\mathcal{PQ}_t = \underbrace{\frac{|TP_t|}{|TP_t| + \frac{1}{2}|FP_t| + \frac{1}{2}|FN_t|}}_{\text{Detection Quality (DQ)}} \times \underbrace{\frac{\sum_{(x_t, y_t) \in TP} IoU(x_t, y_t)}{|TP_t|}}_{\text{Segmentation Quality (SQ)}} \quad (3)$$

with TP_t the set of predicted instances of class t that is correctly classified, FN_t the set of instances predicted as negative

Dataset	Pub. Test Set †	# WSIs	# Annot. Nuclei	Source	Scanner	Mag.	Ref.	# Cit. ‡	Year of Pub.
CoNSeP	yes	16	24,319	Univ. Hosp. Coventry and Warwickshire, UK	Omnyx VL120	40×	(Graham et al., 2019b)	687	2018
KUMAR	yes	30	21,623	TCGA	Multiple	40×	(Kumar et al., 2020)	776	2018
GlaS	no	16	55,364*	Univ. Hosp. Coventry and Warwickshire, UK	Zeiss MIRAX MIDI	20×	(Sirinukunwattana et al., 2017)	671	2016
DigestPath	no	476	168,510*	Multiple Hospitals (4), China	KFBIO FK-Pro-120	20×	(Da et al., 2022)	28	2019
TNBC	yes	11	4,022	Curie Institute, Paris, France	Philips Ultra Fast Scanner 1.6RA	40×	(Jack et al., 2018)	468	2018
CPM-17	yes	32	7,570	TCGA	Multiple	20/40×	(Vu et al., 2019)	218	2017
CCa	yes	10	29,756	Univ. Hosp. Coventry and Warwickshire, UK	Omnyx VL120	20×	(Sirinukunwattana et al., 2016)	1235	2016
Lizard	yes	NA	431,913	Multiple	Multiple	20×	(Graham et al., 2021a)	87	2019
MoNuSAC	yes	46	46,909	TCGA	Multiple	40×	(Verma et al., 2021)	86	2020
PanNuke	yes	481	205,343	Multiple	Multiple	40×	(Gamper et al., 2020a)	286	2020

Table 1: Overview of nuclei instance segmentation and classification datasets. * Annotations provided by the Lizard dataset. † H&E-stained images and cell nuclei annotations. ‡ Google Scholar citations (Last reviewed December 2023).

for class t but that should be positive, and FP_t the set of instances predicted as class t but that are not correctly classified.

PQ is computed as the average of PQ_t over all target classes. For a test dataset, \mathcal{D} , PQ is averaged over all images in \mathcal{D} . Another version of PQ exists, the multiclass \mathcal{PQ} , referred to as $m\mathcal{PQ}^+$, which is the average of \mathcal{PQ}_t computed over all instances in a dataset. This metric is more robust to inconsistencies in images where a particular class is not present, thus naturally resilient to imbalanced data.

• Mean Average Precision

The Mean Average Precision (mAP), for a n class problem, is formally described as

$$mAP = \frac{1}{n} \sum_i AP_i, \quad (4)$$

where the AP is computed as the area under the *precision-recall*, $p(r)$, curve

$$AP = \int_0^1 p(r) dr \quad (5)$$

the precision defined as the fraction of true positives among the detected instances

$$Precision = \frac{TP}{TP + FP} \quad (6)$$

and recall as the fraction of true positives that are detected:

$$Recall = \frac{TP}{TP + FN}. \quad (7)$$

The precision-recall curve represents the trade-off between Precision and Recall at different thresholds. A high Area Under the Curve (AUC) means high precision and high recall. The mAP metric is widely adopted in object detection and segmentation tasks, where a mAP at 50% Intersection over Union (IoU) means that a detected/segmented object is accepted as a true positive if at least 50% of the area covered by both objects represents their intersection.

• Dice Similarity Coefficient

The Dice Similarity Coefficient (DC), or Sørensen-Dice index, is a measure of the overlap between the prediction and ground truth segmentation maps, one of the most common metrics used to assess the quality of image segmentation algorithms. Formally, the DC is defined as

$$DC = \frac{2|P \cap G|}{|P| + |G|} \quad (8)$$

with P and G the predicted and target segmentation maps.

• Symmetric Hausdorff Distance

The symmetric Hausdorff Distance (HD) derives from the directed Hausdorff Distance (d_H), an asymmetric measure of the distance between two finite sets of points. Given two finite sets of points (e.g., predicted and ground truth segmentation maps), $\mathcal{A} = \{a_1, a_2, \dots, a_n\}$, and $\mathcal{B} = \{b_1, b_2, \dots, b_n\}$, the d_H between \mathcal{A} and \mathcal{B} is defined as the maximum distance between each point of \mathcal{A} to its closest point in \mathcal{B} . More formally, we have

$$d_H(\mathcal{A}, \mathcal{B}) = \max_{a \in \mathcal{A}} \{ \min_{b \in \mathcal{B}} \{d(a, b)\} \} \quad (9)$$

where $d(a, b)$ is a metric between these points, typically the Euclidean distance. The symmetric HD can then be expressed in terms of 9:

$$HD(\mathcal{A}, \mathcal{B}) = \max\{d_H(\mathcal{A}, \mathcal{B}), d_H(\mathcal{B}, \mathcal{A})\} \quad (10)$$

This metric is useful to assess the quality of predicted contours and, thus, particularly useful to evaluate the segmentation quality of highly overlapped instances and other fine-grained nuclei segmentation tasks.

4.2.1. Discussion

These metrics are frequently used in the literature to evaluate the performance of cell nuclei instance segmentation and classification algorithms. Nonetheless, there is no gold standard for model evaluation, which represents a bottleneck to directly comparing the results of related works. Moreover, as nuclei instances are frequently small and/or overlapped, there are a few limitations with some of these metrics. The IoU over-penalizes errors in small objects (Cheng et al., 2021a), is asymmetric with regard to penalizing over and under-estimation (Foucart et al., 2023), and does not consider object shape, i.e., nuclei morphology, which is a relevant feature in CPath. On the other hand, this also affects the PQ metric, which is based on the IoU (Foucart et al., 2023). Other limitations of the PQ include the attribution of a more aggressive penalty for good but misclassified detections than for missed detections (Foucart et al., 2023), whereas summarizing the results of a complex multi-objective task into a single metric hinders interpretability of the results. In turn, the DC does not consider the correctness of object boundary

(Reinke *et al.*, 2023), and is more sensitive to misclassifications in small instances (Reinke *et al.*, 2021). Therefore, per task and per class metrics are strongly recommended (Foucart *et al.*, 2023), while other more qualitative approaches could be considered. For instance, recent work has suggested a more robust and clinically significant evaluation protocol (Graham *et al.*, 2024) where interpretable features (morphological, colocalization, and density features) extracted from the segmented nuclei at the WSI-level are used for dysplasia grading and survival analysis.

4.3. Neural Network Architectures

In general, context and attention-based neural networks, namely ViTs (Papers With Code: COCO; Papers with Code: Cityscapes; Gu *et al.*, 2022; Jung *et al.*, 2022; Khan *et al.*, 2022) frequently correspond to top-performing methods for object detection, semantic segmentation, and even instance segmentation and classification. However, in cell nuclei instance segmentation and classification, we verify that specialized neural networks based on conventional ConvNets are the most successful approaches (Graham *et al.*, 2024; Weigert and Schmidt, 2022; Doan *et al.*, 2022), although ViTs have gained increasing interest (Haq and Huang, 2022; Luo *et al.*, 2022; Zhang *et al.*, 2023a). In this subsection, we review some of these works.

HoVer-Net (Graham *et al.*, 2019b) is a dedicated multi-branch neural network for nuclei instance segmentation and classification. It consists of a Pre-Activated ResNet-50 (He *et al.*, 2016) encoder, also called the trunk, and three decoding branches, each responsible for a different prediction. The Nuclear Classification (NC) branch predicts a nuclei class for each pixel, the Horizontal-Vertical (HoVer) branch is responsible for regressing the horizontal and vertical distance maps of each pixel to its object centre of mass, and the Nuclear Prediction (NP) branch assigns each individual pixel to either belonging to a nucleus (1) or background (0). A Sobel filter on the output of the HoVer branch is reflective of individual nuclei instances. These gradient maps have high values at the border of neighbouring nuclei and, therefore, after adequate thresholding can be used in conjunction with the NP branch output to compute markers and an energy landscape for a marker-based watershed algorithm (Graham *et al.*, 2019b). The instance segmentation predictions are then combined with the pixel class predictions, through majority voting, to achieve nuclei instance segmentation and classification.

SONNET (Doan *et al.*, 2022) is a tri-decoder architecture that predicts the Nuclear Foreground (NF) and Nuclear Type (NT), as well as performs Nuclear Ordinal Regression (NO). The ground truth for the NO task consists of a Decreasing Distance Discretization (DDD) map, computed from the Euclidean distance of each pixel to its nucleus center of mass. The DDD map discretizes the distances into K sub-intervals, with the length of the sub-intervals an exponentially decreasing function of the pixel distance to the nucleus center of mass. This approach serves a two-fold purpose. First, the pixels near the center can be used to separate overlapping instances. Second, the pixels near the boundary indicate more uncertain regions, therefore, the authors propose a loss weighting scheme such that a misclassification in pixels further from the origin is more penalized.

Besides, to obtain the final instance segmentation map, the design also includes a post-processing strategy on the outputs of the three decoders.

Liu *et al.* (2022b) introduce a weakly supervised two-stage method (SPN+IEN) that leverages estimated semantic segmentation maps to generate instance-level pseudo-masks. It consists of a Semantic Proposal Network (SPN) to estimate the background and cell nuclei masks, and an Instance Encoding Network (IEN) that outputs instance-level embeddings. Both IEN and SPN consist of U-Net-like architectures. A k-means clustering algorithm is used to generate cell nuclei pseudo-labels to train the SPN, while Voronoi partitioning of the SPN output is used to get instance pseudo-masks used in optimizing the IEN module. The final instance segmentation result derives from mean-shift clustering of the instance embeddings aggregated with the SPN output.

StarDist (Schmidt *et al.*, 2018; Weigert and Schmidt, 2022) supports itself on the principle that star-convex-polygons are well-suited to approximate cellular and nuclear shapes. Namely, for every pixel in the input image, the network regresses its distance to the furthest point in the object it belongs, along predefined equidistant radial directions. Besides, the network simultaneously assigns a probability for every pixel reflecting if it is a foreground object or background. The neural network output, after Non-Maximum Suppression (NMS), thus consists of polygon candidates, expressing cell/nuclei instances, with underlying object probabilities. The authors suggest this approach is well suited for dense nuclei instance segmentation, while having also extended the work to perform classification (Weigert and Schmidt, 2022).

CIA-Net (Zhou *et al.*, 2019c) addresses several of the nuclei instance segmentation challenges, primarily overlapping nuclei, noisy inputs, and stain variability. The proposed neural network architecture is based on an encoder-decoder network but considers two decoder branches and dense convolutional blocks in the encoder path. While one of the decoders is responsible for predicting nuclei segmentation, the other predicts nuclei contours. However, the model assumes features learned for each individual task are useful for the other task, where a learnable bidirectional Information Aggregation Module (IAM) (Zhou *et al.*, 2019c) between the two branches is designed to improve the network ability to jointly refine nuclei and contour predictions. A custom loss function is also proposed, which conditions the model to focus on informative samples to reduce sensitivity to input perturbations (*i.e.* noise).

Although typical segmentation algorithms achieve good results in nuclei semantic segmentation tasks, they become error-prone on instance segmentation, especially in the presence of highly overlapped nuclei. This condition is typical of advanced cancer stages (Alemi Koohbanani *et al.*, 2019). SpaNet is thus an instance segmentation and classification design (Alemi Koohbanani *et al.*, 2019) that leverages spatial information to compensate lack of spatial awareness in conventional algorithms (Alemi Koohbanani *et al.*, 2019). The model first predicts a binary nuclear segmentation map and nuclei centroids, using a dual-head neural network. These outputs are subsequently combined with spatial encoding, and 6 colour

channels, namely in RGB and HSV colour spaces, and used as input to SpaNet to make instance segmentation predictions. SpaNet is an encoder-decoder network, with skip connections from the encoder to the decoder, but that replaces the typical U-Net (Ronneberger et al., 2015) convolutional blocks with Multi-scale Dense Units (MSDU) (Alemi Koohbanani et al., 2019) to encode multi-resolution information into the feature maps. A post-processing routine based on clustering methods is then employed on the SpaNet output to produce the final instance segmentation maps.

As encoder-decoder neural networks are prone to loss of spatial information due to pooling operations, Qu et al. (2019) propose a full-resolution network (FullNet), composed of only dense layers with no stride or pooling and, subsequently, without any upsampling module. As a mimicry for the pooling operation, the authors suggest dilated convolutions that also allow increasing the receptive field size. Besides, the optimization procedure resorts to variance cross-entropy loss, which considers intra-instance statistics, thus working as a spatial constraint on pixels belonging to the same instance (Qu et al., 2019). The authors suggest this loss function promotes the network to be sensitive to object shapes.

Another specialized neural network is the MicroNet model (Raza et al., 2019), which is introduced as a versatile algorithm for the segmentation of several entities in microscopic images, at multiple scales, and in different pathology applications. The network receives multiple inputs, each corresponding to a downsampled version of the same tissue image, thus encoding multi-scale context to predict the underlying segmentation maps (Raza et al., 2019). More precisely, the model contains 5 multi-branch groups. The first group is an encoder module responsible for extracting multi-scale feature maps. Group 2 is a typical bridge layer between the encoder and decoder modules, whereas group 3 corresponds to the upsampling path, with lateral connections to the encoder. Group 4, on the other hand, outputs multi-scale feature maps that are subsequently aggregated and processed in group 5 to predict the segmentation masks. Overall, this model can be fine-tuned for several cell/nuclei segmentation tasks (Raza et al., 2019).

5. Context And Attention-based Methods for Dense Prediction Tasks

This section reviews context- and attention-based mechanisms in CV applications, specifically in dense prediction tasks, including medical imaging applications. Aside from summarizing the main methods, we comment on how the reviewed publications address critical points in ML research, namely considering methodological best practices, reproducibility, OOD generalization, and validity of the drawn conclusions.

5.1. Building Blocks: Computer Vision in General

The idea of incorporating attention mechanisms in deep neural networks derives from the intuition that, to recognize complex objects or scenes, biological beings with intelligence subconsciously pay more attention to discriminative parts of the

visual field while disregarding redundant or irrelevant information. Several works have reviewed attention in CV but the work of Guo et al. (2022b) has proposed a way of schematizing the different attention mechanisms in CV. They start by proposing a general definition for attention, A ,

$$A = f(g(x), x), \quad (11)$$

where $g(x)$ represents attention and $f(g(x), x)$ means processing input x based on attention. A taxonomy is then suggested to organize different kinds of attention mechanisms in CV Guo et al. (2022b):

- **Channel Attention:** In deep neural networks, we can interpret each feature map, at multiple scales, as a hierarchy of visual representations, where each channel is a unique semantic encoding of the input. As such, a channel attention mechanism defines what to pay attention to, i.e., it intends to recalibrate the importance of each semantic representation depending on what is being represented.
- **Spatial Attention:** This term is more related to the structure and spatial context than to the semantics. The spatial attention mechanism intends to refine the importance of different visual regions for representing the input by exploring where to pay attention.
- **Temporal Attention:** For time-dependent input representations (e.g., video, time-series, etc.), this mechanism is a means to establish when to pay attention, i.e., it seeks to reinforce selectivity for representations at specific time points.
- **Branch Attention:** Specific to branched neural networks, it seeks to refine the importance of each branch, thus answering which to pay attention to.
- **Spatial and Temporal/Channel Attention:** Tries to simultaneously refine spatial and temporal/channel representations.

In turn, visual context helps recognize objects more accurately and faster. Context regards all critical information of a visual scene that helps identify an object, such as location, relative size and position to other objects, co-occurrence, or scene semantics. Pathologists heavily rely on contextual information when identifying nuclei in a WSI. As cell nuclei present high intra and inter-class variability, context facilitates decision-making as nuclei rarely appear in isolation, but rather in cellular communities, with moderately well-defined positions and orientations to one another and in semantically rich scenes. Fig. 1 represents an example of colon cell nuclei in a tile extracted from a H&E-stained WSI. For instance, we observe that nuclei from epithelial cells (green) are organized in an oval geometry and that all other cells are located outwards to these and orientated in multiple directions.

But generalizing to the CV domain, a survey by Wang (2022) suggests a structured way of defining context. The authors propose context can be organized into levels and types. Levels include prior knowledge, global context, and local context. We

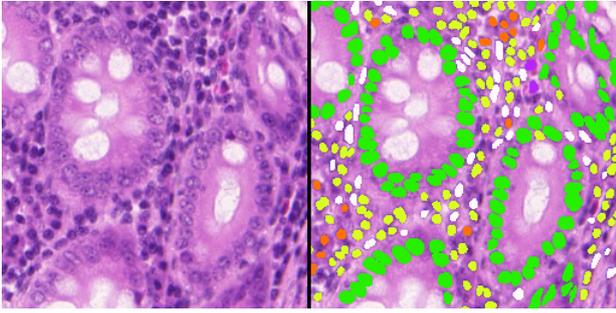


Fig. 1: Example of a WSI tile from the Lizard dataset (Colorectal Cancer (CRC)) (Graham et al., 2019a). Left: Original tile. Right: Cell nuclei overlay. We observe that digital pathology images have rich contexts, where cells organize by type in cellular communities with relatively well-defined geometries and orientations. For instance, we observe that nuclei from epithelial cells (green) are organized in an oval geometry and that all other cells are located outwards to these, and orientated in different directions.

propose to extend the subdivision of context levels also to include long-range and multi-scale contexts. Considering a visual scene, where each object and object parts can be described by a set of independent visual representations, i.e. features, we define context in CV as another set of visual representations that encodes co-occurrence, relative position, and appearance to other objects and scene and that affect the semantic meaning of the visual scene and its parts. Following this definition, we propose the following taxonomy for context levels:

- **Prior knowledge:** Refers to *a priori* knowledge obtained independently of any particular observation. In computer vision, it represents what to expect in a particular scene or domain (Wang, 2022).
- **Global Context:** Concerns visual representations of the scene as a whole, namely, global semantic and spatial features (Wang, 2022).
- **Local Context:** Considers short-range spatial and semantic relations between the visual representations of an instance or neighbouring instances (Wang, 2022).
- **Long Range Context:** Applies to the relations between the visual representations of distant objects, and object parts.
- **Multiscale Context:** Defined as the relation between the representations of the parts and the whole and between the representations of the same concept observed at different levels of magnification (e.g., relations between the visual representations of a constellation observed with 20 and 30× per inch of aperture, or the relationship between the features of a H&E-stained WSI at 10 and 20× magnification).

In turn, concerning type, context can be aggregated into spatial, temporal, and semantic context (Wang, 2022). Spatial context regards the location of an object in a scene and the relative position of other objects and visual regions. It relates to the co-occurrence of different objects, and of object parts, as well as to spatial relations, including directions (e.g., object A

is left of object B), and topological relations (e.g., object A contains object B) (Wang, 2022). Spatial semantic context, in turn, considers the likelihood of an object appearing in a particular scene and the specific semantic relationships with other objects (e.g., a pedestrian is likely to appear on top of a marked crosswalk if the traffic light for pedestrians is green). Beyond the scope of the present survey, there are also temporal and semantic temporal contexts, which consider time-dependent relationships between the features of a visual scene or its parts. Fig. 2 summarizes the characterization of context levels and types in CV.

An important caveat of the present survey is what we define as context-based mechanisms in DL. To make clear the notion of context-based mechanisms that we intend to discuss, we define it as *a set of ANN architectures, layers, modules, training recipes, or regularization strategies designed with the goal of maximizing contextual information in learned representations*. This definition stands from the need to distinguish between the implicit context learned by common ConvNets like DeepLabv3 (Chen et al., 2018) or U-Net (Ronneberger et al., 2015), and other more context-tailored strategies. In the following subsections, we review what we consider the main works in CV that propose context and attention-based mechanisms for image classification and object detection and segmentation tasks. Many other works extend these building blocks to propose specialized architectures.

5.1.1. Squeeze and Excitation Block

Motivated by the idea that the representation capacity of ConvNets could be increased by capturing correlations between features, Hu et al. (2018) introduce an explicit mechanism to model inter-channel dependencies, the so-called Squeeze and Excitation (SE) Block. The central idea of this mechanism is to perform feature refinement, supported by global information, where the most discriminative features are enhanced, and less informative features suppressed. To circumvent the intrinsic locality of convolution operations, the authors propose to squeeze the spatial dimension of the feature vectors to a single unit, through Average Pooling. The resultant feature vector, F_{avg} , is thus an aggregation of local descriptors that encodes the whole image. Then, they perform a feature recalibration operation on F_{avg} . The authors describe this step as “a gating mechanism with sigmoid activation” (Hu et al., 2018). The operation is thus an adaptive recalibration mechanism, capable of capturing channel-wise non-linearities, and non-mutually exclusive dependencies. Equation 12 offers a mathematical description of the SE Block, Where A_{SE}

$$A_{SE}(F) = \sigma(\psi(f_{avg}(F))) \quad (12)$$

means squeeze and excitation attention. $F \in \mathbb{R}^{H \times W \times C}$ is an input feature map, f_{avg} the Average Pooling operator, and ψ a two-layered Multilayer Perceptron (MLP) with Rectified Linear Unit (ReLU) hidden activation, hidden representation $h \in \mathbb{R}^{1 \times 1 \times \frac{C}{r}}$, with r the reduction ratio, and output $f \in \mathbb{R}^{1 \times 1 \times C}$.



Fig. 2: A structured way of defining context in computer vision.

5.1.2. Convolution Block Attention Module

Building on the principles of spatial and channel attention, Woo *et al.* (2018) introduced the Convolutional Block Attention Module (CBAM) that simultaneously answers “what” is meaningful and “where” to pay attention. The basic idea is to obtain channel, and spatial-wise attention scores that are subsequently used to refine the original feature map. Fig. 3 gives an overview of the proposed CBAM.

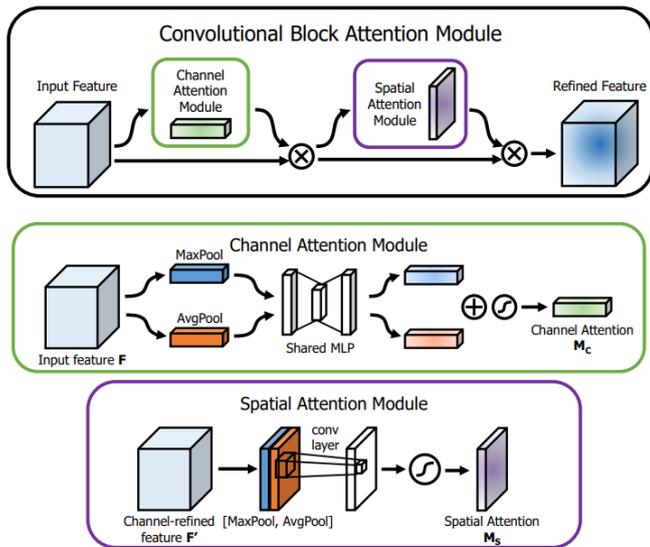


Fig. 3: Overview of CBAM. The module considers both channel, and spatial attention to refine feature maps. The spatial attention module applies 7×7 convolutions over features pooled along the channel dimension, whereas the channel attention module applies a shared MLP to average and max-pooled features along the spatial domain. Reproduced from Woo *et al.* (2018), Doi: 10.1007/978-3-030-01234-2_1

The channel attention module starts with squeezing the spatial information by parallel Max and Average Pooling. The authors argue that simultaneously using the two pooling operations leads to a more complete representation when compared with using only one of the two (Woo *et al.*, 2018). Afterwards, each vector is fed to a shared MLP with one hidden layer and the outputs are summed element-wise to obtain channel attention scores. The full channel attention, A_c , operation is formally

described by equation 13:

$$A_c(F) = \sigma(\psi(f_{avg}^c(F)) + \psi(f_{max}^c(F))) \quad (13)$$

with ψ representing a MLP, f_x the pooling functions, and σ a sigmoid non-linearity. In turn, the spatial attention module starts by squeezing the channel information by Max and Average Pooling operations along the channel dimension. The pooled vectors are then concatenated, followed by a 7×7 convolution and sigmoid non-linearity, resulting in spatial attention scores. Equation 14 is used to formally describe the spatial attention, A_s

$$A_s(F) = \sigma(\phi([f_{avg}^s(F); f_{max}^s(F)])) \quad (14)$$

with ϕ describing the 7×7 convolution, $[]$ the concatenation, and f_x^s the channel-wise pooling operators. Although the authors suggest the modules could be applied in parallel or sequentially, they experimentally demonstrate that a sequential mechanism, with the channel attention module first, provides better results (Woo *et al.*, 2018). The solution is tested with various standard classification networks and shows consistent improvements in model performance. Interestingly, they also enrich the backbones of various object detectors (e.g., Faster-RCNN (Ren *et al.*, 2015)) with CBAM and demonstrate it could lead to better precision (Woo *et al.*, 2018).

5.1.3. Attention Feature Pyramid Network

Ren *et al.* (2015) introduce Faster-RCNN, one of the most widely adopted object detection algorithms, aside from the YOLO (Redmon *et al.*, 2016) architecture. In turn, Lin *et al.* (2017) propose the Feature Pyramid Network (FPN) that, when integrated with object detectors (e.g., Faster-RCNN), enriches feature representation in a hierarchical, multiscale approach. While it explores higher-level features for detecting large objects and lower-level features for enhanced smaller object detection, it explicitly models multiscale semantics resulting in effective multiscale object detectors. However, it does not work effectively for the problem of small dense object detection, which in many ways resembles the task of nuclei detection (see Figure 1). More concretely, Min *et al.* (2022) have identified a few challenges. Lower-level features present noise that “masks”

the representation of small objects (Min et al., 2022). At the same time, although high-level features encode mostly larger objects, they also present some information for small objects, which sometimes leads to ambiguous semantics for small and large object representations (Min et al., 2022). Third, as small objects contain only a few pixels, it becomes challenging to distinguish foreground from background features in complex contexts with objects of similar shape and appearance (Min et al., 2022). They thus suggest the Attention Feature Pyramid Network (AFPN) (Min et al., 2022) for small object detection while arguing it could alleviate these issues. To enrich small object representation of lower-level features, they propose the Dynamic Texture Attention (DTA) module (Min et al., 2022). The solution aims to amplify texture information, which simultaneously emphasises small object details at lower-level layers and highlights features of larger objects at higher layers (Min et al., 2022). In turn, the proposed Foreground Aware Co-Attention (FACA) (Min et al., 2022) module aims to tackle the challenge of enhancing small object foreground features and suppressing noisy backgrounds, while the Detail Context Attention (DCA) module aims to alleviate the loss of contextual information due to RoI pooling.

The results suggest a significant improvement in object detection, especially for small object instances Min et al. (2022). The authors considered the Tsinghua-Tencent 100K (Zhu et al., 2016), Pascal VOC (Everingham et al., 2010), and MS COCO (Lin et al., 2014) datasets, demonstrating that their suggested solution surpasses most state-of-the-art algorithms while working with different model backbones, namely ResNet-50 (He et al., 2016), ResNet-101 (He et al., 2016), and ResNeXt-101 (Xie et al., 2017).

5.1.4. Relation Networks

In most cases, object instances in a visual scene are not independent and identically distributed, i.e., their context and interdependencies play an important role in recognition (e.g., a green traffic light often co-occurs with pedestrians at specific image locations). This reasoning is very appealing for cell nuclei detection tasks, as tissue samples observed from H&E-stained WSIs are highly structured and semantically rich. More precisely, it is expected for histological concepts (e.g. cells, nuclei, and tissue regions) to be related at multiple levels (Pati et al., 2022). In turn, typical object detection algorithms treat each instance individually but, inspired by the success transformers demonstrate on Natural Language Processing (NLP) tasks, Hu et al. (2018) propose an end-to-end object detection strategy that incorporates a self-attention mechanism to capture inter-object relationships. The authors suggest the proposed Object Relation Module (Hu et al., 2018) is effective, as it could capture correlations between instances without making excessive assumptions between their locations and features. By aggregating the original representation and relation features (computed from a set of appearance and geometric features), the object relation module refines the original appearance features (Hu et al., 2018). What we consider distinguishes this method is that it explores the relative geometry of objects, an interesting feature for cell nuclei detection in H&E-stained WSIs.

The Object Relation Module is then stacked on top of multiple object detectors, including a Faster-RCNN (Ren et al., 2015) with FPN (Lin et al., 2017). Fig. 4 gives an overview of the proposed Object Relation Networks.

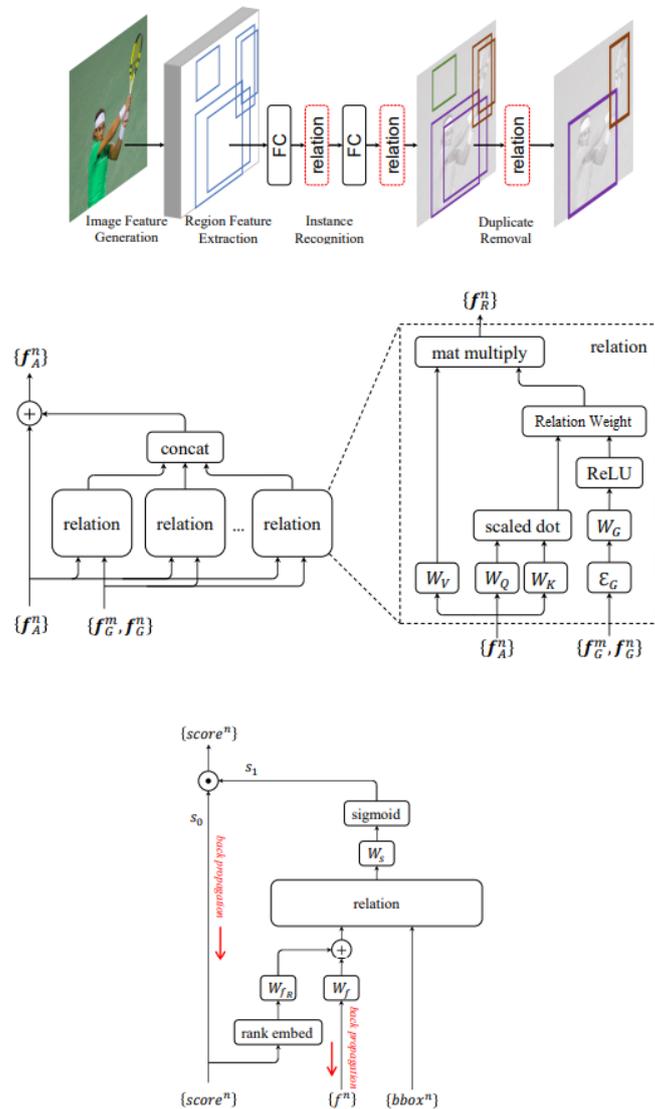


Fig. 4: Top: Overview of an Object Relation Network. Middle: Object Relation Module. Bottom: Duplicate Removal Module. Reproduced from Hu et al. (2018), Doi: 10.1109/CVPR.2018.00378

Another contribution of the method regards duplicate removal (Hu et al., 2018), which could benefit from exploring the relations between objects. To achieve the desired result, the authors define duplicate removal as a two-class classification problem. For each ground truth object, only one of the proposals is assumed to be correct. The Duplicate Removal Module (Hu et al., 2018) considers as input a set of bounding box proposal features and classification scores (s_0). The features are first fused with the classification scores and then passed through an object relation module before a linear layer for estimating the correctness of the bounding box proposal (s_1). The final classification score is established as the product between the bounding box classification score and correctness (s_0s_1). Overall, al-

though it is still unclear what is learnt in the Object Relation Module, the authors claim it allows the network to learn object inter-dependencies that are not captured with regular object detectors (Hu *et al.*, 2018).

5.1.5. Vision Transformers

Global context provides important information regarding dependencies of a particular instance with the scene in which it is inserted. Currently, the most powerful tool in CV to learn global contextual information is the global self-attention mechanism typical of ViTs. Unlike ConvNets, which use localized receptive fields, ViTs' global self-attention allows each token to attend to all other image tokens, thus enabling a more comprehensive understanding of the entire image context. While self-attention excels in capturing semantic and spatial relationships among distant image regions, its inherent flexibility also empowers the encoding of local context surrounding an instance. Consequently, ViTs stand out as an immensely flexible and adaptable family of Artificial Neural Networks (ANNs), adept at capturing local, long-range, and global dependencies, thereby acquiring representations rich in contextual information.

Unlike previous works that combine self-attention with convolutions, ViT (Dosovitskiy *et al.*, 2021) is the first successful attempt at applying transformers directly to images. To deal with the high dimensionality of images, the authors propose to represent each input as a flattened sequence of 2D patches. The patches (or tokens) are then projected (resorting to a trainable linear projection) to a vector of dimension, D , referred to as the patch embeddings and combined with position embeddings to retain positional information. These embeddings then serve as input to a transformer encoder layer consisting of alternating layers of multi-head self-attention and MLP blocks. The MLP contains two layers and GELU (Hendrycks and Gimpel, 2016) non-linearity. But although ViT is a very flexible neural network capable of attending to global information and long-range context, a limitation of this model is the weak image-specific inductive bias. This is very disadvantageous for small datasets where the ConvNets strong intrinsic inductive bias (locality, translation equivariance, etc.) becomes very useful. On the contrary, ViT demands large amounts of training data. Nonetheless, the authors demonstrate that some transformer layers attend to the entire image, whereas others are more local. Overall, the network attends to image regions that are semantically relevant to the classification task.

Although there are several systematic reviews of vision transformers (Khan *et al.*, 2022; Lin *et al.*, 2022a; Han *et al.*, 2023), we also review a few publications that build on top of ViT and that we consider being relevant for the discussion on context and attention mechanisms in cell nuclei instance segmentation and classification. We give a high-level overview of the field and present works that could be fine-tuned for medical imaging applications.

5.1.5.1 Data Efficient Image Transformer

The fact that ViT requires large volumes of annotated data is a major limitation, especially considering real-world applications, where often data acquisition and annotation are expensive

(e.g., in digital pathology) or the object instances of interest are naturally underrepresented in the dataset (e.g., defective products in manufacturing). As such, if transformers are to compete with, or even surpass ConvNets, proper strategies should be defined to attain data-efficient results.

One of the first successful attempts to efficiently train ViT is the Data Efficient Image Transformer (DeiT) proposed by Touvron *et al.* (2021a), which provides a detailed set of training mechanisms to increase data efficiency. Given the relative sensitivity of ViT to initialization, the authors demonstrate that initialization with a truncated normal distribution leads to better results Touvron *et al.* (2021a). Furthermore, they compensate for the datasets' limited size (from a transformer perspective) with extensive data augmentations, namely RandAugment (Cubuk *et al.*, 2020). At the same time, they find that AdamW (Loshchilov and Hutter, 2017) optimizer, with stochastic depth regularization during training, is also beneficial. Finally, their research suggests batch augmentation is central to improving model performance. Likewise, Steiner *et al.* (2022) extensively study cost-effectiveness in vision transformers and find that carefully tuned augmentations and regularizations allow, in general, to reduce by 10x the amount of training data while maintaining accuracy. Moreover, the authors present empirical evidence of the common belief that pretraining vision transformers on large-scale datasets followed by fine-tuning to downstream tasks is "simultaneously cheaper and gives better results" (Steiner *et al.*, 2022).

5.1.5.2 Swin Transformer

By merging patches in deeper layers, Swin Transformer introduces the hierarchical structure to ViTs (Liu *et al.*, 2021b) and closes the gap between ViTs and ConvNets. Indeed, the hierarchical nature of ConvNets makes them suitable to represent objects of varying scale and makes these networks capable of intrinsically encoding multiscale contextual information. Moreover, the hierarchical feature maps of Swin Transformer (Liu *et al.*, 2021b), in alliance with the self-attention mechanism, make it a strong backbone for state-of-the-art CV models (e.g., FPNs (Lin *et al.*, 2017) or U-Net (Ronneberger *et al.*, 2015)) to solve dense prediction tasks. Besides, the method introduces a shifted window approach to compute self-attention, which reduces the computational complexity and latency but still enables the encoding of global and long-range context. Notably, the focus on the efficiency of the Swin Transformer (Liu *et al.*, 2021b) potentiates the scaling of ViT to high-resolution images (Liu *et al.*, 2021b).

5.1.5.3 Convolution Vision Transformers

Although frequent claims are made about transformers' superiority, provided larger training datasets are given, there is evidence that ConvNets can achieve similar performance under proper design choices and training recipes (Liu *et al.*, 2022c). Furthermore, transformers most significant performance gap with ConvNets has been on classification, whereas CV entails other paradigms (e.g. object detection, instance segmentation),

which benefit from ConvNets properties (e.g., translation equivariance). Despite this, efforts have been made towards transformers for object detection and segmentation (Carion et al., 2020; Fang et al., 2021), although many still rely on convolutional feature extractors (Carion et al., 2020; Cheng et al., 2022). On the other hand, Raghu et al. (2021) compare transformers visual representations with ConvNets and demonstrate that while transformers better preserve localization, attend to global information, and possess more uniform features across the neural network layers, local information is still relevant in lower layers and larger datasets become necessary to learn this property despite the weaker inductive bias. In essence, there is no clear argument to choose one family of neural networks over the other, particularly given both paradigms possess unique strengths and can suffer from similar limitations (Pinto et al., 2022; Gonçalves et al., 2022), i.e., the tendency to learn spurious correlations, simplicity bias, and poor OOD generalization. Ultimately, the choice of neural network family is conditioned by the application and available resources. But a promising and recurrent practice is a combination of the two paradigms to exploit their strengths and mitigate individual weaknesses (Wu et al., 2021; Guo et al., 2022b; Xiao et al., 2021), which could allow the introduction of image domain-specific inductive bias into the transformers neural network, or promote self-attention and global context with ConvNets. For instance, Xiao et al. (2021) demonstrate replacing the “patchify” stem with a convolutional one improves training stability while increasing performance. Chu et al. (2021) introduce a convolutional operation in the transformer architecture to remove the positional embedding, whereas Wu et al. (2021) also consider a hierarchy of visual representations. Similarly, Guo et al. (2022a) propose a convolutional stem, depthwise separable convolutions, and an inverted residual feed-forward module with a transformer neural network, which maintains performance and increases efficiency. PatchConvNet (Touvron et al., 2021b), on the contrary, is a convolutional architecture that builds on top of the central idea of ResNets, i.e., the residual connection, but that adopts a modified convolutional stem (like Xiao et al. (2021)), and backbone with constant width and no stride/pooling (the so-called “column”). Besides, the authors propose to replace the final average pooling layer with a transformer-like cross-attention module. This translates to a model that learns a weight for each patch, depending on its similarity with a trainable class vector (analogous to ViT (Dosovitskiy et al., 2021) classification token). In essence, after aggregation and reshaping of the attended features, the attention pooling module allows visualizing attention weights (per class) according to the importance of the patch to the classification, making this model more interpretable. The results suggest PatchConvNet achieves superior performance to a ResNet-50 (He et al., 2016), and the authors demonstrate it can easily be extended to instance segmentation and classification, for example, by replacing the Mask-RCNN (He et al., 2017) backbone (without Feature Pyramid Network (Lin et al., 2017)) with a PatchConvNet.

5.1.5.4 Detection Transformer

The Detection Transformer (DETR) is introduced by Carion et al. (2020) with the intent of simplifying modern object detectors while maintaining competitiveness. The strategy directly predicts the bounding boxes without needing post-processing and other manually tuned strategies (e.g., NMS) through reasoning about the global image context and object relations. Despite consisting of a simple architecture, namely, a ConvNet feature extractor and an encoder-decoder module (see Fig. 5), the algorithm enables an end-to-end training strategy through a bipartite matching loss and parallel decoding.

Since the decoder is permutation invariant, to allow bounding box regression and classification, each N object query is embedded with a learnable positional encoding at each attention layer. The decoder thus outputs N normalized object queries that serve as input to a MLP to predict bounding box positions, defined by the centre coordinate, height, and width, relative to the input image, as well as classification scores (Carion et al., 2020). Like other modern object detectors, a big challenge while training these networks is correctly corresponding the predicted bounding boxes with the ground truth. Given that the N predictions are larger than the number of ground truth objects, the target vector is usually padded with \emptyset (no object). Adopting a fixed order matching raises a few challenges, as the loss becomes too sensitive to false negatives or positives. To avoid such cumbersome training, the authors adopt an optimum matching strategy, the Hungarian algorithm (Kuhn, 1955), that efficiently enforces permutation-invariance, and ensures a single predicted bounding box correspondence for each target. To compute the pairwise matching costs, the algorithm considers not only the error between the predicted bounding box and ground truth but also the classification scores.

5.1.5.5 Masked Attention Mask Transformer

The Masked-attention Mask Transformer (Mask2Former) (Cheng et al., 2022) extends the MaskFormer (Cheng et al., 2021b) architecture and is an attempt to replace specialized architectures for segmentation tasks by a universal architecture, while being easy and efficient to train. It comprises a backbone feature extractor, a pixel decoder, and a transformer decoder. The feature extraction module generates a low-resolution image feature map $F \in \mathbb{R}^{\frac{H}{S} \times \frac{W}{S} \times C_F}$, with C_F the dimension of the feature map, and S the stride (here, $S = 32$). In turn, the pixel decoder progressively up samples the feature maps to obtain per-pixel embeddings $\epsilon \in \mathbb{R}^{H \times W \times C_\epsilon}$. The transformer decoder module takes as input the backbone image features, F , and N learnable positional embeddings (i.e., queries), and outputs N per segment embeddings, Q . A Linear Classifier and MLP with 2 hidden layers then predict, respectively, N object class probabilities (including no object \emptyset), and mask embeddings, ϵ_{mask} , from the transformer decoder outputs. To predict the binary masks, a sigmoid non-linearity is applied after the dot product between the mask embeddings, and pixel-level embeddings.

Like DETR, MaskFormer (Cheng et al., 2021b) uses a bipartite matching loss but instead of using bounding boxes to estimate the costs between ground truth and prediction, it directly

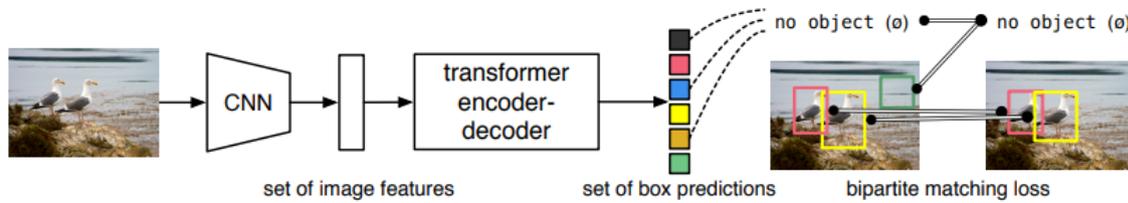


Fig. 5: Overview of the DETR. Reproduced from Carion *et al.* (2020), Doi: 10.1007/978-3-030-58452-8_13

uses class and mask predictions. Building on top of MaskFormer (Cheng *et al.*, 2021b), Mask2Former similarly leverages the concept of transformers' self- and cross-attention; however, instead of each token attending to all other tokens, it restricts the attention to localized tokens around the predicted segments. Contrary to the standard cross-attention, the proposed masked-attention considers the resized binarized output of the previous $l^{\text{th}} - 1$ layer. This allows faster convergence and better performance. Besides, to improve small object detection, Mask2Former considers multi-scale high-resolution features. More precisely, it adopts a FPN strategy where each feature scale proposed by the pixel-decoder is fed to the corresponding level at the transformer decoder, after positional and scale encoding. The transformer decoder layer thus possesses three levels, one for each feature resolution (1/8, 1/16, and 1/32). This structure is further repeated L times. Fig. 6 illustrates the Mask2Former architecture.

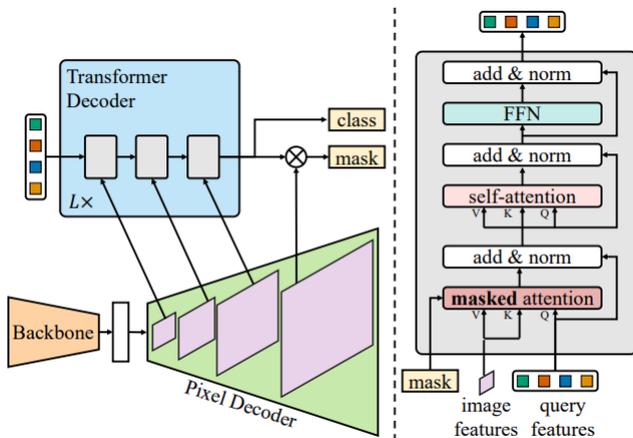


Fig. 6: Overview of Mask2Former. Reproduced from Cheng *et al.* (2022), Doi: 10.1109/CVPR52688.2022.00135

To optimize the results, the authors propose further improvements such as switching the order of self and cross attention, learnable queries that work like a Region Proposal Network (RPN), and remove dropout from the transformer decoder. Finally, to increase training efficiency, the authors suggest the use of K randomly sampled points for computing both matching and final losses, which reduces memory usage by three times. More precisely, for the matching loss, they always use the same K uniformly sampled points, whereas for the final loss, they use different prediction/ground truth pairs according to importance sampling.

5.2. Methods for Dense Prediction Tasks in Medical Imaging

Cell nuclei instance segmentation and classification is not an isolated area of knowledge. Indeed, CV researchers, in general, have many times been inspired by or benefited from scientific findings in other fields, like NLP (Dosovitskiy *et al.*, 2021), geometric DL (Bronstein *et al.*, 2021), or neuroscience (Maass, 1997). Besides, many solutions in medical imaging are built on top-of high-impact contributions designed to achieve top performance (e.g., accuracy in classification problems, or panoptic quality in instance segmentation) in big and challenging, general CV, benchmark datasets. This motivates us to the idea that context- and attention-based methods suggested in medical applications other than cell nuclei instance segmentation and classification are also worthy of careful review, as some of these contributions could inspire and guide future research in cell nuclei instance segmentation and classification from H&E-stained brightfield microscopy images.

5.2.1. Dermatology

Like cell nuclei, skin lesions possess very heterogeneous sizes, shapes, and appearances, often with irregular fuzzy boundaries, even considering lesions from the same subject. Furthermore, skin lesions' data are usually 2D RGB images with high variability (e.g., due to scanner differences), subject to occlusion (e.g., due to hairs), and sometimes with low contrast between foreground lesions and remaining healthy tissue. Other similarities with cell nuclei detection and classification include skin lesions being, to some extent, rounded, and frequently corresponding to dense, small objects in an image. Fig. 7 gives an example of a skin lesion 2D RGB image that illustrates several of these claims.



Fig. 7: Left: Example of a wide field image with multiple pigmented lesions on the back of a female subject. Right: 2-class DL-based detection of skin lesions. Non-Suspicious Pigmented Lesions of Medium Priority (NSPL-B) are signalled in yellow, whereas Suspicious Pigmented Lesions (SPL) are marked with red. It is also observable that the task of skin lesion detection and classification possesses several similarities with the task of cell nuclei detection and classification, for example, oval foreground objects, dense small objects, RGB images, and high intra-subject variability. Reproduced, with permission, from Soenksen *et al.* (2021), Doi: <https://doi.org/10.1126/scitranslmed.abb3652>

Skin lesion detection, segmentation, and classification are thus difficult tasks. Nonetheless, several works propose the use of DL with context- and attention-based mechanisms. Wang et al. (2019) resort to an encoder-decoder like neural network with attention FPN to enable a richer multiscale contextual representation useful for downstream skin lesion segmentation. Additionally, to improve the performance for small object instances, the authors suggest a scale-attention loss function, that weights the loss for objects at different scales, effectively leading to the network paying more attention to small instances. Similarly, Feng et al. (2020) introduce Context Pyramid Fusion Networks (CPFNet) for medical image segmentation tasks, which combine two pyramidal modules with a U-Net (Ronneberger et al., 2015) architecture to aggregate global and multi-scale contextual information. But as attention-based mechanisms lead to improved model performance, some works argue it could also help to achieve explainability in medical imaging tasks (van der Velden et al., 2022). For instance, in medical imaging segmentation, Gu et al. (2021a) discussed Comprehensive Attention Convolutional Neural Networks (CA-Nets) that adopt attention at multiple levels, namely, channel, spatial, and scale-wide attention, which they argue simultaneously improves accuracy and explainability. Kaul et al. (2021) also build on top of the U-Net (Ronneberger et al., 2015) architecture for skin lesion and cell nuclei segmentation. This work, however, proposes Residual Group Attention Networks (ResGANets) that couples a SE-block (Hu et al., 2018) with group convolutions for enhanced feature representations. In turn, to enhance visual representations, Xu et al. (2021) consider a neural network architecture capable of capturing global and multiscale context. The Dual Context Network (DC-Net) combines hierarchical visual representations from a ConvNet encoder with the transformers' self-attention to capture long-range contextual information. More specifically, the output features from each layer of the U-Net (Ronneberger et al., 2015) encoder (i.e., feature maps from a ResNet-34 (He et al., 2016) feature extractor) are linearly projected, along with flattened input image patches, and positional embeddings. The resulting feature sequence inputs a transformer encoder module to generate enhanced feature representations that encode long-range dependencies. The second mechanism, the Adaptive Context Fusion (ACF) Module (Kaul et al., 2021), combines multiscale contextual information from each layer of the U-Net (Ronneberger et al., 2015) decoder, supported on spatial and channel attention along with dilated convolutions, to decode the feature representations into a 2D segmentation output. But different from previous works that consider only ConvNets or hybrid ConvNet/transformer networks, He et al. (2022b) propose the Fully Transformer Network (FTN) for skin lesion segmentation, which mimics the encoder-decoder topology of widely used ConvNets (e.g., U-Net (Ronneberger et al., 2015)), while introducing a hierarchy of features into the transformer architecture and skip connections from the encoder to the decoder modules to compensate for loss of spatial resolution. Notably, this design enables local, long-range, global, and multiscale contextual information, thus being a very flexible neural network.

5.2.2. Radiology

Although radiology data often corresponds to grayscale images (reflecting physical attributes of the underlying tissues), it is reasonable to state that certain aspects of detecting, and segmenting RoIs from these data approximate the task of nuclei detection and segmentation. A few similarities support this argument. For instance, both sorts of medical imaging data possess rich contexts, often with small object instances, and high intra- and inter-subject variability. Moreover, although radiology data can be volumetric, it is often decomposed into 2D slices to ease downstream analysis. Therefore, we hypothesise contextual visual representations and attention-based algorithms previously popularized for radiology could be adapted for cell nuclei detection, and segmentation. For example, Wei et al. (2022) successfully apply an attention-based contrastive learning approach to detect pneumonia lesions. Wang and Quan (2021) propose a neural network with global and local edge attention modules for the segmentation of global and finer details in CT images. In another application, Deng et al. (2021) successfully apply a hybrid transformer to segment the left ventricle in echocardiography. However, there are other directions in medical imaging, including radiology, that are worth discussing as viable alternatives to address cell nuclei instance segmentation and classification. These include multi-task learning, annotation-efficient DL algorithms, and other active areas of research like explainability and interpretability.

Chen et al. (2019b) show that context- and attention-based multi-task learning is effective in semi-supervised brain lesion segmentation. Indeed, multi-task learning can work as an inductive bias that helps ML models generalize and could work as regularization in semi-supervised settings. Imran and Terzopoulos (2019) adopt a multi-task Generative Adversarial Network (GAN) approach for simultaneous segmentation and classification of chest radiographs. Besides, the authors consider an attention-gated block to tackle possible loss of spatial resolution in masks generated at different levels of a pyramidal encoder, which allows suppression of noisy features, leading to the neural network encoding more effectively the useful multiscale contextual information. Ullah et al. (2023) illustrate that extending a chest radiographs dataset for COVID-19 detection with additional related data in an adversarial multi-task learning paradigm enables data-efficient generalization. Although their methodology does not contemplate explicit context and attention-based mechanisms, a possible improvement could be incorporating attention mechanisms into the neural network layers, or even adapting the adversarial autoencoder to context reconstruction. Le et al. (2021), on the other hand, rely on a contextual multi-task network with atrous convolutions for tumour detection and segmentation, which works as an attention mechanism to enhance contextual information surrounding the tumour region.

Other works demonstrate that attention mechanisms simultaneously increase model performance and explainability. For instance, Karri et al. (2022) propose a multi-module semantic guided attention network (MSGANet). The self-attention and multiscale attention modules are introduced to capture global context, while the channel, edge, and location attention mod-

ules ease the aggregation of local and global representations. The model is subsequently evaluated on several standard medical imaging datasets, including multi-organ segmentation from Computerized Tomography (CT) imaging, and brain tumour segmentation from Magnetic Resonance Imaging (MRI) data. Whereas the results demonstrate higher segmentation quality, the authors overlay the attention weights each attention module provides onto the original images and demonstrate a more explainable model than other state-of-the-art alternatives.

5.2.3. Cell/Nuclei Instance Segmentation and Classification in Microscopy Imaging

Microscopy imaging allows the magnification and visualization of organelles, cells, and tissues by exploiting the optical properties of tissues and lens. Although the core of our work regards H&E-stained brightfield microscopy images, there exist other techniques in the microscopy imaging domain that provide different insights and distinct levels of detail depending on the staining or imaging modality. In this subsection, we delve into some works more tailored to microscopy imaging modalities excluding H&E-stained brightfield microscopy. We consider this to be relevant as, regardless of the imaging technique, some domain-specific attributes make solutions tailored to other imaging modalities of potential value to H&E-stained brightfield microscopy imaging. Some of the challenges in microscopy imaging orthogonal to the specificities of each technique are nuclei/cell morphology variability, class imbalance, overlapped, touching, and cluttered instances, noisy/irrelevant background information, and even artefacts, such as out-of-focus images, digitization artefacts, cracks, and air bubbles. Fig. 8 illustrates these claims. Methodologies suggested for some specific microscopy imaging modality(ies) could thus be worth considering in cell nuclei instance segmentation and classification from H&E-stained WSIs (and tiles). A successful illustration of this argument is the StarDist method (Schmidt *et al.*, 2018) that is first suggested for cell detection in fluorescence microscopy images, but that later wins the CoNIC Challenge (Weigert and Schmidt, 2022), showing its ability to generalize to H&E-stained brightfield microscopy images as well.

Wan *et al.* (2023) suggest CellT-Net, based on a Swin-transformer (Liu *et al.*, 2021b) backbone with multiscale contextual features to segment cells of different size (including small instances) in light microscopy imaging. The attention-based model enhances the semantic information of foreground regions and reduces background noise. Besides, the authors consider the adoption of deformable convolutions, which increase model robustness to cell morphology variability. Sugimoto *et al.* (2022) explore the context surrounding cells via a modified self-attention layer. Krug and Rohr (2022), on the other hand, suggest a self-supervised learning paradigm for cell segmentation in fluorescence microscopy images based on a generated attention map and supported by an auxiliary task of predicting the scale of the image (prior knowledge) using as input the image multiplied by the attention map. The method thus makes use of attention mechanisms to learn features of the most relevant regions/concepts in the input image. Wagner and Rohr (2022) propose CellCentroidFormer, a hybrid ViT model

for cell detection in microscopy imaging. The strategy suggests an EfficientNet (Tan and Le, 2019) backbone for feature extraction, and a MobileViT (Mehta and Rastegari, 2021) with ConvNet layers neck for subsequent cell centroid, width, and height regression. The suggested strategy thus combines global transformer-like context and attention-based processing, which promote the model to focus on cell shapes, with ConvNets local attention to cell centroids, thus achieving enhanced performance when compared with state-of-the-art ConvNets. Chen *et al.* (2022c) adopt a Mask-RCNN (He *et al.*, 2017) for lung cell segmentation, followed by a Swin-transformer (Liu *et al.*, 2021b) for classification. Kakumani *et al.* (2022) adopt an attention-based U-Net (Ronneberger *et al.*, 2015) with encoder weights initialized from a pretrained autoencoder and attention-gate blocks in the decoder path for semantic segmentation of cells. CellDETR (Prangemeier *et al.*, 2020) is a modified version of DETR for cell instance segmentation from microscopy images. It consists of a ConvNet backbone feature encoder, and two branches: a transformer encoder-decoder for bounding box regression and classification, and a ConvNet decoder for segmentation. The input to the segmentation decoder consists of multi-head attention-refined features from the transformer branch, aggregated with the shared ConvNet encoder features. The authors also consider skip connections from the encoder to the segmentation decoder module (similar to U-Net (Ronneberger *et al.*, 2015)). Zhou *et al.* (2019a), on the other hand, focus on the precise segmentation of overlapping instances with a method, Instance Relation Network (IRNet), supported on inter-cell contextual information (i.e. long-range context). The strategy suggests extending the Mask-RCNN (He *et al.*, 2017) detection and segmentation heads with, respectively, Duplicate Removal (Hu *et al.*, 2018) and self-attention-based Instance Relation modules (Hu *et al.*, 2018), which they show to improve the quality of the predicted instances' boundary. Yi *et al.* (2019) consider the use of attention mechanisms to address two different challenges in neural cell detection and segmentation. While the authors suggest self-attention in a detection module to refine feature maps, an attention mechanism is proposed in the skip connection of the segmentation encoder-decoder module to suppress noisy background features.

5.2.4. Discussion

The discussed publications present evidence that context- and attention-based mechanisms improve the detection and segmentation performance of neural networks, although the benefits for instance segmentation and classification are not completely clear. A notable methodological aspect of these works is that the challenge of domain generalization is reasonably addressed; namely, nearly all the adopted datasets contain data from multiple centres, spanning several regions of the globe, with samples acquired under different conditions (e.g., differences in data acquisition protocol) and scanners. This could improve DL model generalization while reducing domain shift complications, along with data augmentation and model regularization, which most of the reported publications also contemplate. Guidelines (England and Cheng, 2019; Mongan *et al.*, 2020; Lekadira *et al.*, 2021) to evaluate the models on stan-

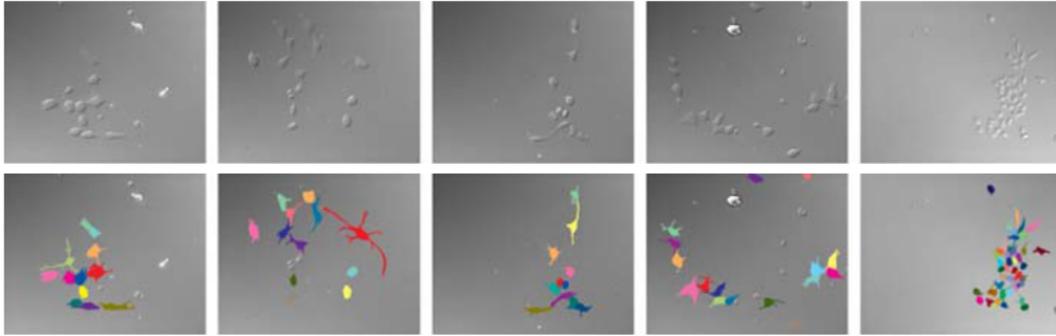


Fig. 8: Top: Sample images from a Central Nervous System (CNS) stem cells dataset (Ravin *et al.*, 2008). Bottom: Ground truth cell instance segmentation maps. This image is illustrative of some similarities between cell instance segmentation and classification in stainless brightfield microscopy imaging and H&E-stained brightfield microscopy images. We observe morphological variability, air bubble artefacts, dense and overlapped instances, as well as low contrast between foreground instances and background. Reproduced, with permission, from Yi *et al.* (2019), Doi: 10.1016/j.media.2019.05.004

dard, public access datasets have also been followed, illustrating the scientific rigour of the discussed medical image segmentation publications. Despite this, a common limitation resides in dataset size and model evaluation. Although we observe the authors are careful not to leak the same patient’s data from the training subset to the validation, and test splits, the lack of assessment of generalization to multiple target domains limits the conclusions that can be attained from the observations. Besides, and due to intrinsically balanced datasets, most publications do not address the issue of data imbalance, a common hurdle in medical imaging analysis (e.g., in rare diseases).

Another comment we make on these works is the focus on benchmark performance. Although it still is an important contribution for the ML field, Varoquaux and Cheplygina (2022) quantitatively demonstrate it could have diminishing returns. As the ultimate objective of DL research in medical imaging would be to deploy the algorithms in clinical practice, we suggest researchers should combine these efforts with other clinically meaningful tasks like biomarker research, quality control, data/label efficiency, OOD generalization, or explainability. Although there is evidence that context- and attention-based mechanisms could be useful for these topics (Chen *et al.*, 2022a; Rao *et al.*, 2021; Gu *et al.*, 2021b), we underline a tendency for publication of positive results in ML research, meaning there could be a bias in the scientific community’s understanding of the benefit of these algorithms, particularly due to limited chance to learn from negative findings to generate new hypothesis. Besides, other works (Gonçalves *et al.*, 2022; Ghaffari Laleh *et al.*, 2022) present less clear evidence of the advantages of context- and attention-based mechanisms. This question is thus a sensitive topic in ML for medical imaging research that should be thoroughly addressed.

Most previously summarised mechanisms could be adapted for cell nuclei instance segmentation and classification from H&E-stained brightfield microscopy imaging. In addition, nearly all of the discussed publications carefully address reproducibility, as they publish details on hyperparameters, data splits, or resort to standard publicly available datasets. But considering the guidelines for AI in medical imaging (England and Cheng, 2019; Mongan *et al.*, 2020; Lekadira *et al.*, 2021), we

suggest the results and conclusions should not be blindly interpreted. For example, although statistical measures are reported, they are performed only between different repetitions of the main experiments, meaning the effects of the interaction between the various modules added to the neural networks are not properly assessed, albeit the ablation studies have suggested performance improvements with each individual component. Therefore, despite accumulating evidence that context- and attention-based mechanisms could improve the quality of automated medical image segmentation, it is not possible to unequivocally conclude that these are the only factors responsible for increased capabilities, which hinders the interpretability of these methods. This is a common issue in ML research, but we understand it is difficult to assess, as previously discussed (Pineau *et al.*, 2020).

To finalize, it is worth commenting that each medical imaging modality presents unique characteristics, meaning there is no straightforward success of the same neural network and training recipe in all tasks. For instance, there are differences in contrast and noise between foreground and background objects in skin lesion images or radiology data, and H&E-stained WSIs. The latter have more complex semantic and spatial contexts, with many features of no interest for cell nuclei instance segmentation and classification but of similar morphological or chromatic appearance. This translates to noise that masks the presence of foreground instances, i.e., cell nuclei, in learned visual representations. Besides, there are also domain-specific features and artefacts that differently affect the performance of DL models trained on different types of medical imaging data.

6. Cell Nuclei Instance Segmentation and Classification in Brightfield Microscopy with H&E Staining

Due to digitized H&E-stained brightfield microscopy images (mostly H&E-stained WSIs), being part of routine clinical practice, these data are the most widely available. Moreover, H&E staining is one of the brightfield microscopy imaging techniques that provide the most detailed information, allowing clear visualization of tissue structures like the cytoplasm, nucleus, organelles and extra-cellular components. Therefore,

this modality offers the best cost-effectiveness when compared with other techniques, even holding the promise that imaging biomarkers available in these data could replace expensive laboratory tests like genetic or molecular profiling. Methods for H&E-stained image analysis have thus been the focus of CPath research. For these reasons, we give special attention to this imaging technique. Namely, we do a systematic review of context and attention mechanisms for DL-based cell nuclei instance segmentation and classification from H&E-stained brightfield microscopy images, including WSIs and tiles. We searched the Scopus database and limited the results to publications between 2019 and 2023. We then read the abstract of every entry returned and excluded publications with no mention of nuclei segmentation/detection, neither attention nor context. To further refine the results, we read the text of the returned publications and excluded papers that did not contemplate algorithms for H&E-stained WSIs. Also, we limited our analysis to papers published in Q1 journals (Scimago Journal & Country Rank) (e.g., *Medical Image Analysis*, *IEEE Transactions on Medical Imaging*, *Computers in Biology and Medicine*, etc.), and peer-reviewed conferences (ranked C or higher as per the CORE conference ranking) (*Medical Image Computing and Computer Assisted Intervention (MICCAI)*, *International Conference of the IEEE Engineering in Medicine Biology Society (EMBC)*, *International Conference on Medical Imaging with Deep Learning*, etc.). Table 2 presents the queries used to search the database, as well as the number of entries returned, inclusion and exclusion criteria, and the final number of publications reviewed.

In total, we reviewed 44 publications on context and attention mechanisms for cell nuclei instance segmentation and classification from H&E-stained microscopy images, including WSIs and tiles. Fig. 9 illustrates the number of reviewed publications, per year, from 2019 to 2022, on context and attention mechanisms for neural-network-based cell nuclei instance segmentation and classification.

We next provide a comprehensive discussion of how the different context levels have been addressed in the literature for cell nuclei instance segmentation and classification from H&E-stained brightfield microscopy images, while illustrating the different attention mechanisms that have been explored and their intended use like enriching the estimated feature maps with contextual information or reducing sensitivity to irrelevant background information.

6.1. Self-supervised Learning with Prior Knowledge

Prior knowledge is very useful in representation learning, as it enables inductive biases useful for constraining the optimization search space. Indeed, having preference for some solutions is necessary for generalization (Goyal and Bengio, 2022) and enables learning task-discriminative features with unsupervised representation learning.

SSL is a learning paradigm that transforms the unsupervised setting into a supervised one through the optimization of a loss function on a pretext task using as supervisory signal the data itself (Hastie *et al.*, 2009; Goodfellow *et al.*, 2016). As SSL enables learning from high volumes of unlabelled data (Campanella *et al.*, 2023), it is desirable in resource-constrained settings, where annotating large and diverse datasets is costly, or

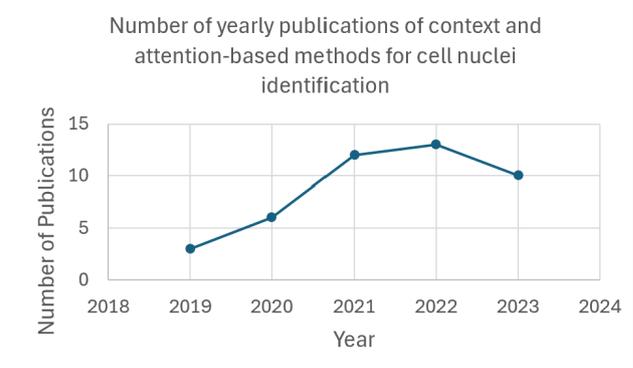


Fig. 9: Number of publications reviewed, per year, on context and attention-based mechanisms for cell nuclei instance segmentation and classification. We observe that these mechanisms are gaining momentum in CPath, namely for nuclei instance segmentation and classification from H&E-stained WSIs.

in rare target occurrence cases where the labels of some classes are naturally limited. Due to the prohibitive dimensionality of WSIs and the need for expert clinical annotators, this is especially relevant for CPath and could accelerate the clinical implementation of DL-based WSI analysis tools, including cell nuclei instance segmentation and classification. Another advantage of SSL is that it works as regularization, by constraining the search space of downstream fine-tuning optimization around the self-supervised weight initialization region in the parameter search space (Erhan *et al.*, 2010). This regularization effect improves generalization, potentially to OOD data as well. In the last years, unsupervised visual representation learning, largely supported by self-supervision with pretext tasks, contrastive learning (Hjelm *et al.*, 2018; Misra and van der Maaten, 2020; Chen *et al.*, 2020b; He *et al.*, 2020; Chen *et al.*, 2020a), knowledge distillation (Gidaris *et al.*, 2021; Grill *et al.*, 2020; Chen and He, 2021), or information maximization (Zbontar *et al.*, 2021; Bardes *et al.*, 2021; Ermolov *et al.*, 2021), has advanced such that it has been shown to compete with supervised learning in several applications (Ciga *et al.*, 2022b; Bardes *et al.*, 2021; Azizi *et al.*, 2021; Chen *et al.*, 2019a). Furthermore, self-supervision could translate to models that learn more generic and less spurious correlations, thus beneficial for domain adaptation and generalization (Hendrycks *et al.*, 2019). For example, in CPath, multi-task learning with self-supervised pretext tasks in conjunction with supervised classification allows the model to learn more discriminative and domain-invariant features under limited annotated data (Koohbanani *et al.*, 2021). On the other hand, (Srinidhi and Martel, 2021) provide evidence that while self-supervision methods result in more general visual representations, they tend to overfit the pretext tasks if the dataset is small or fail to generalize to OOD data. Nonetheless, the authors suggest a Curriculum Learning (CL) approach in downstream fine-tuning that they demonstrate to improve accuracy and generalization of patch-, and slide-level classification on benchmark histology datasets when compared with standard downstream transfer learning. In the scope of cell nuclei instance segmentation and classification, contextual information, namely prior knowledge, has been used to learn more useful

Search terms	Returned	Inclusion criteria	Exclusion criteria	Reviewed
Attention Nuclei Segmentation	302	- Reference to attention or context in the abstract/document - Trained and evaluated on H&E-stained WSIs - Q1 journal or peer-reviewed conference	- No reference to attention or context in the abstract/document - Imaging modalities other than H&E-stained WSI - Cell segmentation algorithms - Duplicates	24
Context Nuclei Segmentation	179	- Reference to attention or context in the abstract/document - Trained and evaluated on H&E-stained WSIs - Q1 journal or peer-reviewed conference	- No reference to attention or context in the abstract/document - Imaging modalities other than H&E-stained WSI - Cell segmentation algorithms - Duplicates	8
Transformer Nuclei Segmentation	58	- Vision Transformer or Hybrid Transformer/ConvNet - Q1 journal or peer-reviewed conference	- Imaging modalities other than H&E-stained WSI - Cell segmentation algorithms - Duplicates	10
Transformer Nuclei Detection	32	- Vision Transformer or Hybrid Transformer/ConvNet - Q1 journal or peer-reviewed conference	- Imaging modalities other than H&E-stained WSI - Cell segmentation algorithms - Duplicates	1
Weakly Supervised Nuclei Detection	38	- Reference to attention or context in the abstract/document - Trained and evaluated on H&E-stained WSIs - Q1 journal or peer-reviewed conference	- No reference to attention or context in the abstract/document - Imaging modalities other than H&E-stained WSI - Cell segmentation algorithms - Duplicates	1

Table 2: Details of the systematic review on context and attention-based nuclei instance segmentation and classification. The search is performed in the Scopus database.

representations with self-supervision. For instance, motivated by the annotation efficiency of SSL, Sahasrabudhe et al. (2020) introduce an attention-based unsupervised nuclei segmentation algorithm. The authors propose a multilayer convolutional neural network, based on dilated convolutions, that increase the receptive field, thus providing a larger context for the pixel classification task. This model outputs a coarse nuclei segmentation binary attention map. The attention map is then aggregated, by multiplication, with the original image and the resulting nucleus-salient H&E-stained WSI tile is used as input to a scale classification network (ResNet-34) (He et al., 2016). The networks are trained end-to-end where the total loss is a combination of the loss computed at the scale classification level, with equivariance and smoothness regularization losses at the attention mask prediction module. The authors argue that given a H&E-stained WSI tile, the level of magnification can be inferred by looking at the size and texture of the nuclei in the patch (Sahasrabudhe et al., 2020). They thus use an attention mechanism to exploit this domain-specific context, i.e. in the form of prior knowledge, to design a network capable of learning annotation-efficient and generalizable representations. The final segmentation map is obtained by post-processing the coarse segmentation image. The results demonstrate that, despite underperforming state-of-the-art supervised methods, the proposed model is competitive with other unsupervised learning methods (Sahasrabudhe et al., 2020). Another work exploring contextual information in the form of prior knowledge is that of Boserup and Selvan (2023), which consider a contrastive loss at the pixel level for annotation-efficient nuclei segmentation. In this work, it is assumed that tiles containing the same nuclei types are likely to be similar. Therefore, positive pairs are defined as semantically similar patches that are likely to contain the same object(s), while negative pairs correspond to patches that are not likely to contain the same object(s), which could be interpreted as approximating a patch to that of its local context (i.e., its neighbours) when embedded in a low-dimensional manifold. To estimate patch similarity to compute the contrastive loss, the authors consider the mean squared error or mean cross entropy and devise a confidence prediction network which outputs, for

each class, the likelihood of each pixel belonging to it. The input images are then scaled by the confidence maps (for each class), and the loss is computed at the level of these confidence-weighted vectors. To complete the nuclei segmentation task, at inference time, the predicted confidence maps are thresholded at 0.5.

Hassan et al. (2022), on the other hand, use the knowledge that cell nuclei are frequently organized into cellular “communities”. Starting from a pre-trained nuclei instance segmentation and classification method (Mask-RCNN (He et al., 2017)), the authors resort to a message-passing network to learn inter-instance relations and predict cellular communities, which they then use to refine the nuclei class predictions.

6.2. Local Context

Channel and spatial attention mechanisms are generally used to enhance foreground information while suppressing certain background features. Moreover, these are also used to refine coarse nuclei segmentation maps. Indeed, attention mechanisms are often used to suppress irrelevant contextual information allowing DL models to focus on useful features like the local context of a pixel in a nucleus. Overall, local contextual information helps refine nuclei borders and allows distinguishing between nuclei types.

Kong et al. (2020) suggest two-stage stacked U-Nets, with attention, for nuclei segmentation in H&E-stained tiles. A stacked U-Net consists of four parallel U-Net backbones in which the hidden representations are merged with an attention module. This generates attention weights for each of the four U-Net outputs that are then aggregated for the final segmentation. However, in a two-stage scheme, the authors propose that the output of the first stacked U-Net be fused with the original input image for the second round of segmentation. This allows the suppression of background noise while making nuclei more salient, which facilitates the subsequent segmentation of overlapping nuclei. By leveraging attention in this way, the strategy effectively suppresses noisy or irrelevant contextual information (i.e., background information, like extracellular matrix or irrelevant cell organelles) while enhancing the relevant local

context of the foreground pixels (i.e., nuclei). Similarly, Yoo *et al.* (2019) propose a WSL strategy based on attention to enhance local context. Their method considers an auxiliary task to learn from point annotations. The main segmentation network tends to produce coarse segmentation maps; therefore, an auxiliary network is used to guide the main network to produce more delicate edges. The adjacent network is trained with pseudo-labels generated from a Sobel filter on the predicted segmentation masks from the main network, and an attention module is used to refine the edge predictions of the auxiliary model.

CDNet (He *et al.*, 2021a) is a multitask network for nuclei instance segmentation. The architecture contemplates a feature extraction module (encoder-decoder architecture) and three branches with residual units to predict the nuclei mask, centripetal direction feature map, and nuclei centers feature map. Each branch has as input, respectively, feature maps from the feature extraction module, feature maps from the mask branch, and feature maps from the centripetal direction prediction branch. The spatial attention mechanism is then used to aggregate features from the different branches on the premise that the auxiliary tasks provide relevant contextual information to guide the instance segmentation task. Besides, at test time, a direction difference map computed from the centripetal direction feature map and nuclei centers is used to refine nuclei borders. The approach is suggested to effectively deal with ambiguities due to touching or overlapping nuclei (He *et al.*, 2021a).

The Mulvernet model (Vo and Kim, 2023) is an extension of HoVer-Net (Graham *et al.*, 2019b), where the skip connections are enriched with an attention gate and multiple filter units to enrich the representations with multiscale features. DRCA-Net, designed by Dogar *et al.* (2023), is composed of modified ResNet-50 (He *et al.*, 2016) encoder and three decoder branches, with CBAM-based (Woo *et al.*, 2018) dense attention blocks, for instance, segmentation and classification. Each branch is responsible for a single task: nuclei segmentation, regression of horizontal and vertical distance maps, and nuclei classification. A post-processing methodology, similar to HoVer-Net (Graham *et al.*, 2019b), is then adopted to combine the three outputs into a single prediction. The authors demonstrate the proposed model results in finer segmentation maps, facilitates the separation of overlapping nuclei, and reduces false positives when compared with other state-of-the-art methods on three standard open-access datasets. ConDANet (Imtiaz *et al.*, 2023) is an encoder-decoder segmentation model where it is suggested to replace average or max pooling and unpooling with, respectively, the discrete and inverse discrete wavelet transforms, an operation that better preserves image details. Moreover, the authors propose the use of a contourlet filter (Do and Vetterli, 2005) with a spatial attention mechanism to provide additional edge information, in essence, a strategy to enhance additional contextual information around the nuclei borders. Ultimately, the use of the attention mechanism allows for refining local contextual information, which translates to better nuclei instance segmentation and classification.

CPPNet (Chen *et al.*, 2023) is suggested to improve the segmentation of touching and overlapping nuclei. The method extends the StarDist (Schmidt *et al.*, 2018; Weigert and Schmidt,

2022) model by estimating the distance (along a set of predefined radial directions) from an additional set of points between the centroid and the initially predicted boundary. Furthermore, the approach contemplates a Context Enhancement Module (CEM) and a Confidence-based Weighting Module (CWM) to refine the initial centroid to boundary distance relative to these points. Besides, a shape-aware loss is suggested to penalize differences in shape features between the predicted and ground truth nucleus representations. Finally, the authors introduce a post-processing scheme to remove redundant predictions and obtain the final segmentation map. Overall, by leveraging an additional set of points, this method can capture more contextual information than StarDist, especially near the boundary.

6.3. Long-range Context

We have that CPath data represent a rich tissue microenvironment that reflects multiple clinical variables like the diagnosis, the disease subtype, tumour stage and grade, lifestyle habits of the subject, etc. The imaging features encoding these variables are spread across the tissue, even in distant regions, but their relationships and co-occurrence provide relevant anatomical and pathological information. Therefore, a representation rich in long-range context, encoding the co-occurrence and long-range dependencies of different concepts (e.g., cells, or nuclei) is more discriminative than a representation that only considers local context. Besides, such information could be especially useful in the presence of ambiguities (e.g., due to noisy or blurred data, occluded objects, small instances, etc.).

Ali *et al.* (2022) consider a U-Net-like encoder-decoder model with multi-scale attention through dense dilated convolutions at the level of the last hidden layer of the encoder, which, given the large receptive field of the hidden units in this layer, leads to learning a representation rich in long-range contextual information. Besides, they propose a multi-scale efficient channel attention module at each decoder layer, which in combination with the hierarchical structure of the decoder, increases the use of semantic and spatial context. A boundary refinement module is also proposed to refine nuclei borders.

Attention mechanisms and long-range context have also been exploited in WSL, a set of data-driven learning techniques that optimize predictive models with incomplete, inexact, or inaccurate supervision (Zhou, 2018). Indeed, WSL is a widely explored paradigm to tackle the challenges of learning from partially annotated data, label inconsistencies, noisy samples, and domain shifts. In object detection and segmentation, the goal of WSL tasks is often to propagate image (bag-level) labels to individual object instances (Zhang *et al.*, 2022b), or to generate segmentation masks from point annotations. WSL is thus a cost-effective solution for CPath as it could alleviate the burden of annotating the gigapixel WSIs at the pixel, or instance level. Notably, WSL has demonstrated useful in several CPath applications, like tissue grading (Campanella *et al.*, 2019), mitosis detection in breast tissue tiles (Li *et al.*, 2019), or prediction of molecular alterations from H&E-stained WSIs (Kather *et al.*, 2020). Thiagarajan *et al.* (2019), introduce a Knowledge Distillation (KD) and Multiple Instance Learning (MIL) framework for weakly supervised epithelial cell nuclei classification. To

perform the patch-level classification task, namely to detect epithelial nuclei, each WSI is subdivided into a set of patches and the attention-based MIL method is adopted to optimize a bag-level classifier to predict cancerous tissue. First, patch-level feature extraction takes place followed by attention-based feature aggregation to learn a global representation for the WSI. This attention-based MIL harnesses long-range inter-patch context to develop a bag classifier. Notably, given the method's flexibility in accommodating a variable number of whole-slide patches, the optimized model can be employed for predicting instance-level labels by considering each patch as a collection of nuclei instances. This is motivated by the fact that chromatic and morphological alterations in epithelial cell nuclei are correlated with cancer mutations. Hence, to fine-tune the instance-level features, a student model, employing an identical architecture to that of the teacher, undergoes optimization with supervision from the teacher model. Besides, the authors adopt virtual adversarial training as regularization, which improves robustness to bag perturbations (e.g., noisy instances), thus reducing overfitting. Overall, this is a label-efficient strategy leveraging attention and long-range contextual information in WSIs to learn a label-effective tile-level epithelial cell nuclei classification network.

LIRNet (Zhou *et al.*, 2020) is suggested as a framework for cell nuclei counting and localization. The architecture consisted of a lighter version of U-Net, with dot product attention in the bottom layer to capture inter-object relationships and long-range context. The network estimates pseudo-density maps whose ground truth is generated from point annotations, i.e., bounding box centroids. An advantage of the proposed algorithm is that it allows both Strongly Supervised Learning (SL) and WSL. While in SL the full pseudo-density map is used as ground truth, in WSL, the image is divided into patches and each patch is assigned a label representing one of the following: no nuclei, one nucleus, and at least two nuclei. Afterwards, each learning paradigm adopts a specialized loss function, supported by the integral of local densities (i.e., foreground objects), under the assumption that if the integral of the whole density map is nuclei counts, then the integral of individual foreground objects should equal one. Nuclei locations are derived by post-processing the estimated pseudo-density maps.

6.4. Global Context

Global context refers to the relationship between each instance and the scene as a whole. In nuclei instance segmentation and classification, this is achieved mostly through ViTs that allows modelling the relationship between a small tissue patch and all other patches.

Medical imaging instance segmentation and classification is an emerging application of ViTs, but the weak inductive bias and the lack of large annotated datasets slows progress and motivates the hybridization of transformers with convolutions, a more data-efficient alternative. Hybrid neural networks can potentially increase performance compared with a pure ViT or ConvNet, making them an attractive approach. Luo *et al.* (2022), develop a U-Net-based neural network with parallel self-attention blocks connected by cross-attention blocks to the

encoder module, which allows the exchange of complementary global and local information between the two encoding paths. An *et al.* (2022) define an efficient attention module, combined with dual-channel shifted MLP, and transformers self-attention to enrich foreground representations and suppress background noise. Obeid *et al.* (2022) propose the use of DETR (Carion *et al.*, 2020) for nuclei detection but introduce a realistic data augmentation policy to increase the number of negative tiles (i.e., no nuclei) and a near-duplicate merging strategy to decrease the false positive rate. Chen *et al.* (2022d) enrich the HoVer-Net (Graham *et al.*, 2019b) trunk with a transformer attention module head and extend both the residual blocks and the decoding branches with SimAM (Yang *et al.*, 2021). Valanarasu *et al.* (2021) define a two-branch encoder-decoder model with transformer-based gated axial attention. The global branch is shallower and reasons over the entire image to compute global attention, while the local branch processes image patches to capture finer details. At the final stage, the output of both branches is combined for the segmentation prediction. The gated axial transformer layers are a modified version of axial attention (Ho *et al.*, 2019) and are suggested to overcome the difficulty in learning positional encodings under limited data. Overall, by combining global and local self-attention mechanisms, the solution contemplated in this work increases data efficiency and performance in several medical imaging segmentation tasks, including cell nuclei segmentation.

Zhang *et al.* (2023a) design CircleFormer, an extension of DETR that works with circular bounding boxes ($bbox = (c, r)$, with $c = (x, y)$ the coordinates of the bounding box center and r the radius). Therefore, due to cell nuclei possessing a rounded shape, the strategy is more fitted for nuclei detection and segmentation tasks. Moreover, the work suggests circle modulated attention that takes into account the object radius, enabling the model to learn features of objects at various scales. This mechanism tailors the model to also work with smaller instances, a known shortcoming of state-of-the-art object detectors like DETR and Mask-RCNN (Dubey *et al.*, 2022; Min *et al.*, 2022).

Self-supervised learning has also been shown successful for pretraining vision transformers for nuclei segmentation. Haq and Huang (2022) propose a triplet loss computed at the embedding space of a TransUNet (Jieneng *et al.*, 2021) to promote similarity between the representations of neighbouring (8-connected) foreground patches (i.e., with cell nuclei) while contrasting the representations of these patches with background (i.e., no nuclei). At the same time, the authors adopt the task of predicting the input image (i.e., H&E-stained tile) scale (prior knowledge), with a dedicated model that receives as input the H&E-stained tile and the segmentation mask predicted by TransUNet (Jieneng *et al.*, 2021). In a subsequent stage, the model is fine-tuned in a small annotated dataset. The authors also demonstrate the widely accepted idea that pretraining benefits downstream supervised fine-tuning and that tailored self-supervised tasks are more beneficial for the nuclei segmentation task than pretraining on regular RGB images (e.g., ImageNet) (Deng *et al.*, 2009).

In turn, Pan *et al.* (2023) propose SMILE, a multitask architecture to simultaneously achieve nuclei instance and semantic

segmentation, as well as nuclei classification. The design they introduce contemplates a tri-encoder and tri-decoder U-Net but with a shared Recurrent Criss-cross Attention (RCCA) module (Huang et al., 2019) at the bottom layer of the encoders. The criss-cross attention module resorts to the affinity mechanism that is similar to that of global self-attention, but where each pixel attends only to pixels in the same row and column, instead of to all other pixels. The RCCA module consists of two stacked criss-cross attention modules. A second pass through this module produces a similar effect to that of global self-attention, thus being a computationally efficient way of capturing global context and long-range dependencies.

6.5. Multiscale Context

In CPath, common forms of multiscale context include the relationship between nuclei and cells, and the way nuclei are arranged and grouped together to form cellular communities and tissue structures. Therefore, understanding this broader context may be useful for nuclei instance segmentation and classification tasks.

U-Net-like encoder-decoder neural networks enriched with attention mechanisms are widely adopted for nuclei instance segmentation and classification. Notably, through the use of attention mechanisms or dilated convolutions in conjunction with the hierarchical structure of encoder-decoder models, these strategies enable the representation of multiscale contextual information (Li et al., 2023). Vahadane et al. (2021) propose the use of a second encoder in a U-Net-like architecture that captures input-specific attention. The multi-resolution attention maps of the attention encoder are then aggregated with the U-Net (Ronneberger et al., 2015) feature maps at an attention skip module. The dual-encoder U-Net then predicts nuclei and boundary probability maps that are post-processed to obtain the final segmentation output. He et al. (2021b), on the other hand, combine a nested U-Net, inspired by U-Net++ (Zhou et al., 2018), with a hybrid attention module (similar to CBAM (Woo et al., 2018)), and argue the resultant model can capture more multiscale context, which facilitates the segmentation of small and dense nuclei while being flexible to nuclei appearance variability. Wazir and Fraz (2022) combine a variant of the SE-block (Hu et al., 2018) with multiple losses in an encoder-decoder model to segment multiple structures, namely cells and nuclei. By adopting dilated (atrous) convolutions of multiple scales, the suggested design can cover multiple levels of context, including local, long-range, and multiscale context. Another strategy based on dilated convolutions (Li et al., 2021b) proposes to enrich a U-Net model with channel attention blocks, dilated convolutions, and a boundary smoothness loss, which promotes the separability of foreground boundaries and background representations, thus improving segmentation. FeedNet (Deshmukh et al., 2022) suggests a mechanism (FE Block) to compensate for the loss of spatial information after the pooling layers in encoder-decoder neural networks. At the same time, the authors propose depth-wise separable convolutions and design an Long Short-term Memory (LSTM) (Hochreiter and Schmidhuber, 1997) block to refine the encoder output features, which works like a chan-

nel attention mechanism. Interestingly, with fewer model parameters, the results demonstrate competitive performance on nuclei instance segmentation. ASPPU-Net (Wan et al., 2020) consists of a modified U-Net (Ronneberger et al., 2015) with atrous spatial pyramid pooling to induce multiscale representations and capture nuclei contextual information. This model aims to improve the segmentation quality with overlapped instances. Therefore, the authors combine the benefits of contextual visual representations with a concave point detection algorithm to refine the segmentation of highly cluttered nuclei. Besides, the authors exploit the properties of the H&E staining; namely, they project the RGB image into the individual hematoxylin and Eosin components on the argument that as hematoxylin is more sensitive to nucleic acids (RNA/DNA), richly present in cell nuclei, rough endoplasmatic reticulum, and ribosomes, it provides a better channel to contrast nuclei borders with surrounding tissue, thus easing nuclei segmentation. Ahmad et al. (2023) introduce DAN-NucNet, an encoder-decoder segmentation model with spatial and channel attention gates. In DSCA-Net (Shan et al., 2023), the authors propose extending a U-Net (Ronneberger et al., 2015) architecture with a channel and spatial attention module at each layer of the decoder, which, in conjunction with the hierarchical nature of the decoder, translates to a representation encoding context at multiple scales. Moreover, the method entails an edge attention module that aggregates multiscale features to refine nuclei borders.

Although FPN-based dense prediction methods have shown successful, mainly driven by their inherent ability to aggregate multi-scale contextual information, they possess some limitations like difficulty in detecting small objects. Geng et al. (2021) discuss other limitations, including the inability to transfer some relevant semantic information from higher to lower levels, and discuss the advantages of expanding FPNs (Lin et al., 2017) with attention mechanisms to capture better long-range dependencies and enhance important spatial positions. Moreover, the authors identify that feature maps at the higher level of FPNs (Lin et al., 2017) lack contextual information as they correspond to a single scale. Therefore, to enhance the expressiveness of contextual features at this level, the authors propose a context augmentation module based on dilated convolutions and the self-attention mechanism.

The architecture of the Swin transformer (Liu et al., 2021b) enables the encoding of global, long-range, and multiscale contextual information. Therefore, it is a strong backbone for H&E-stained brightfield microscopy cell instance segmentation and classification models (Zhou et al., 2023). Nevertheless, the weak inductive bias that characterizes ViTs limits the capability of these models to encode spatial context (Wang et al., 2023). To address this challenge, Wang et al. (2023) consider an architecture similar to that of Mask-RCNN (He et al., 2017), but with Swin transformer (Liu et al., 2021b) backbone and suggest dilated convolutions of multiple dilation rates at each stage of the transformer backbone, which increases the range of multiscale contextual information that is encoded. Besides, the method introduces a convolutional graph neural network-based feature fusion module that combines features from the top and bottom

layers of the Swin transformer (Liu et al., 2021b) backbone to obtain an affinity matrix that captures the extent of the dependency between two nodes. Then, this matrix is used by graph convolutions to map the features from the Swin transformer bottom layer to a representation that encodes inter-instance relations. But another method focused on improving model performance in clustered nuclei instance segmentation is proposed by Ke et al. (2023). Based on hybrid ViTs, the architecture extends the U-Net (Ronneberger et al., 2015) encoder with transformer layers and modifies the decoder path to, besides nuclei masks, output individual nuclei edges, as well as the edges of clustered nuclei. Besides, the proposed model can be optimized only from partial edge annotations and a small portion of fully annotated masks, thus being annotation-effective. A similar approach (He et al., 2023) based on a tri-decoder architecture replaces the hybrid convolutional encoder and decoders with Swin transformer (Liu et al., 2021b) layers. Besides, to increase efficiency and capture correlations between the nuclei segmentation and edge prediction tasks, the design proposes a shared multi-head self-attention mechanism in the decoders. Overall, the broader contextual understanding of these methods helps the models to better differentiate the boundaries of clustered nuclei, that frequently contain low contrast and touching or overlapping instances.

Guo et al. (2021) posit a two-stage method for cell nuclei instance segmentation, bolstered with attention, that requires only nuclei centroid annotations. From these weak annotations, the authors then estimate two types of pseudo ground truth masks: point annotations + Voronoi labels, and superpixel pseudo ground truths generated by masking superpixels from a SLIC (Achanta et al., 2010) algorithm with Voronoi cells. These two types of ground truth labels are thus complementary. The main segmentation network is an encoder-decoder architecture, with slight modifications in the encoder, like dilated convolutions to incorporate multi-scale context, but with the insertion of a CBAM-based feature aggregation module in the decoder. During stage I, the superpixel masks are used as the ground truth for updating the network weights. The authors propose an auxiliary network, the mask-guided attention network, to help refine the segmentation module predictions. The output of this network is then considered an attention weight to help denoising the nuclei segmentation predictions. A specific loss to the auxiliary task is proposed, which penalizes pixel-wise errors between the point annotations + Voronoi boundaries and the attention map. Stage II could then be considered a refinement step, where confident learning (Northcutt et al., 2021) is used to update the pseudo ground truth labels before retraining the whole framework with the revised labels (Guo et al., 2021).

6.6. Discussion

Overall, most algorithms for cell nuclei instance segmentation and classification revolve around encoder-decoder neural networks, demonstrating the approach's superiority over other alternatives. Table 3 summarizes the main works (reviewed from 2019 to 2023) on cell nuclei instance segmentation and classification. Furthermore, table 4 and Fig. 10 provide insights regarding the source code availability of the reviewed publications. We observe an increasing concern with reproducibility,

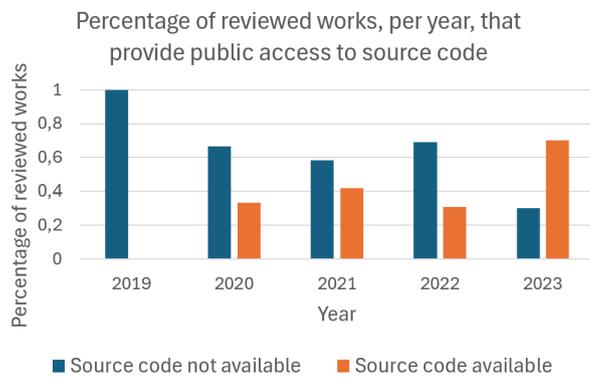


Fig. 10: Percentage of reviewed publications, per year, that provide public access to source code. We observe an increasing concern with reproducibility, although from 2019 to 2022 only less than 42.5 % of published works addressed this critical aspect of ML research.

although from 2019 to 2022 less than 37.5 % of the reviewed papers address this critical part of ML research. Figures 11, 12, and 13, on the other hand, report the performance of reviewed methods on the KUMAR (Kumar et al., 2020) and CoN-SeP (Graham et al., 2019b) datasets. As different publications choose distinct evaluation metrics, not all works are directly comparable. However, most works resort to DC, IoU (JI), AJI, or PQ, making these suitable metrics for fair comparisons. We find that most strategies show significant improvements by enriching neural networks with context and attention-based mechanisms, but we observe that while it works satisfactorily well for binary and semantic segmentation tasks, performance is usually lower in instance segmentation and classification (quantified with the PQ metric). Furthermore, although transformers are starting to prevail in CV in general, CPath, and cell nuclei instance segmentation and classification are still dominated by ConvNets (see tables 5, and 6). A possible explanation could reside in transformers' poor performance in small object instances Cheng et al. (2022); Carion et al. (2020) or in the limited availability of large-scale annotated datasets.

In summary, there is a growing trend in context- and attention-based mechanisms for cell nuclei instance segmentation and classification from H&E-stained brightfield microscopy images. In general, these methods were shown to increase overall performance, but there is no "one-size fits all" solution as different approaches tackle specific challenges. Contextual information at multiple levels in conjugation with tailored architectures and training recipes holds the promise to tackle the challenges of clustered nuclei with overlapping or touching instances. In turn, spatial and channel attention have been explored to enhance local context and relevant foreground information while suppressing noisy or irrelevant features. Other approaches consider context and attention modules to refine coarse nuclei segmentation maps, especially cell nuclei borders, whereas prior knowledge, like the image magnification level, is useful to learn from unlabelled data general but discriminative features that potentiate finetuning with scarce or incomplete annotations. Besides, there is evidence that context- and attention-based mechanisms enable more ef-

Method	Year	Base Architecture	Attention	Context Type(s)	Context Level(s)	Learning	Dataset(s)	Reference
Self-sup. Scale Pred.	2020	ConvNet	spatial	spatial	local	SSL	CoNSEP, KUMAR, TNBC	Sahasrabudhe et al. (2020)
Patch Contrast. Learn	2023	ConvNet	spatial	spatial	local	SSL	CoNSEP, KUMAR	Boserup and Selvan (2023)
MSAL-Net	2022	U-Net	channel	semantic	long-range	SL	KUMAR	Ali et al. (2022)
DEAU	2021	U-Net	channel, self-attention	semantic	long-range, multiscale	SL	KUMAR, ConSEp, CPM-17	Vahadane et al. (2021)
SUNet	2020	stacked U-Net	spatial	spatial	local	SL	KUMAR, TNBC	Kong et al. (2020)
HistoSeg	2022	U-Net	channel, spatial	semantic, spatial	local, long-range	SL	GlaS, KUMAR	Wazir and Fraz (2022)
CAB-Net	2021	U-Net	channel	semantic	long-range	SL	KUMAR	Li et al. (2021b)
HAN-Net	2021	U-Net	channel, spatial	semantic, spatial	local, long-range	SL	KUMAR	He et al. (2021b)
DRCA-Net	2023	U-Net	channel, spatial	semantic, spatial	local, long-range	SL	ConSEp, KUMAR, CPM-17, PanNuke	Dogar et al. (2023)
HyLT	2022	Hybrid Transformer	self-global	semantic, spatial	global, long-range, multiscale	SL	GlaS, KUMAR	Luo et al. (2022)
HEA-Net	2022	Hybrid Transformer	channel, spatial, self-global	semantic, spatial	global, long-range, multiscale	SL	GlaS, KUMAR	An et al. (2022)
TransNuSS	2022	Hybrid Transformer	self-global	semantic	global, long-range	SSL	TNBC, KUMAR	Haq and Huang (2022)
NucDETR	2022	Hybrid Transformer	self-global	semantic, spatial	global, long-range, multiscale	SL	ConSEp, PanNuke	Obeid et al. (2022)
THSVNet	2022	Hybrid Transformer	self-global, channel, spatial	semantic, spatial	global, long-range	SL	ConSEp, PanNuke	Chen et al. (2022d)
MedT	2021	Hybrid Transformer	self-global	semantic, spatial	global, long-range, multiscale	SL	GlaS, KUMAR	Valanarasu et al. (2021)
Mask Guided Attention	2021	Encoder-Decoder (ConvNet)	spatial, channel	semantic, spatial	local, long-range	WSL	KUMAR, TNBC	Guo et al. (2021)
Pseudo-edge Net	2019	Encoder-Decoder (ConvNet)	spatial	spatial	local	WSL	KUMAR, TNBC	Yoo et al. (2019)
Distill-to-Label	2019	ConvNet	self-attention (gated)	semantic	local, long-range	WSL	CCA, BCa	Thiagarajan et al. (2019)
LIRNet	2020	U-Net	self-attention	semantic, spatial	local, long-range	WSL	CA	Zhou et al. (2020)
DAN-NucNet	2023	U-Net	channel, spatial	semantic, spatial	local, long-range	SL	PanNuke, KUMAR, CPM-17	Ahmad et al. (2023)
Mulvernet	2023	HoVer-Net	spatial	spatial	local	SL	MoNuSAC, GlySAC, ConSEp	Vo and Kim (2023)
DSCA-Net	2023	U-Net	channel, spatial	semantic, spatial	local, multiscale, long-range	SL	TNBC	Shan et al. (2023)
REU-Net	2022	U-Net	branch, spatial	semantic, spatial	local, long-range	SL	KUMAR, ConSEp, HUSTS, CPM-17	Qin et al. (2022)
AL-Net	2022	U-Net	spatial, channel	semantic, spatial	local, long-range	SL	KUMAR, PanNuke, ConSEp	Zhao et al. (2022a)
RCSAU-Net	2022	U-Net	channel, spatial	semantic, spatial	local, long-range, global	SL	KUMAR, PanNuke	Wang et al. (2022a)
MCNS-U-Net	2022	U-Net	channel, spatial	semantic, spatial	local, long-range	SL	MoNuSAC	Sheikh et al. (2022)
HBA-Net	2021	ConvNet	channel, spatial	semantic, spatial	local, long-range	SL	KUMAR, CPM-17	Cheng et al. (2021c)
MASG-GAN	2021	Encoder-Decoder (ConvNet)	channel, spatial, branch	semantic, spatial	local, long-range	SL	KUMAR, TNBC, ConSEp, KIRK, CPM-17	Zhang et al. (2021)
Feat. Att. Net	2021	Encoder-Decoder (ConvNet)	channel, spatial	semantic, spatial	local, long-range	SL	KUMAR, ConSEp	Murtaza Dogar et al. (2021)
NucleiSegNet	2021	U-Net	channel	semantic	local, long-range	SL	KMC, KUMAR	Lal et al. (2021)
ACDNN	2020	ConvNet	channel	semantic, spatial	local, long-range	SL	-	Yi et al. (2020)
NAS-SCAM	2020	ConvNet	channel, spatial	semantic, spatial	local, long-range	SL	MoNuSAC	Liu et al. (2020)
RIC-U-Net	2019	Encoder-Decoder (ConvNet)	channel	semantic	local, long-range	SL	TCGA, CPM-17	Zeng et al. (2019)
MDC-Net	2021	Encoder-Decoder (ConvNet)	-	spatial, semantic	long-range, multiscale	SL	TNBC, TCGA	Liu et al. (2021a)
AA-Net	2021	Encoder-Decoder (ConvNet)	channel	semantic	long-range, multi-scale	SL	KUMAR	Geng et al. (2021)
ASPPU-Net	2020	U-Net	-	semantic, spatial	long-range, multi-scale	SL	TNBC, TCGA	Wan et al. (2020)
Color-space HoVer-Net	2022	ConvNet/HoVer-Net	channel, spatial	semantic, spatial	local, long-range, multi-scale	SL	Lizard	Azzuni et al. (2022)
Message Passing Net.	2022	multiple	-	semantic, spatial	local, long-range	SL	ConSEp, PanNuke, Lizard, CRCHisto.	Hassan et al. (2022)
CDNet	2021	Encoder-Decoder (ConvNet)	channel, semantic	semantic, spatial	local	SL	KUMAR, CPM-17	(He et al., 2021a)
ConDANet	2023	U-Net	channel, spatial	semantic, spatial	local	SL	KUMAR, CPM-17	(Imtiaz et al., 2023)
CircleFormer	2023	Hybrid Transformer	self, global	semantic, spatial	local, global, long-range	SL	KUMAR	Zhang et al. (2023a)
SMILE	2023	Hybrid Transformer	self, global	semantic/spatial	local, global, long-range	SL	ConSEp, MoNuSAC	(Pan et al., 2023)
ClusterSeg	2023	Hybrid Transformer	self, global	semantic, spatial	local, global, long-range, multiscale	SL	KUMAR	(Ke et al., 2023)
TransNuSeg	2023	Hybrid Transformer	self, global	semantic, spatial	local, global, long-range, multiscale	SL	KUMAR	(He et al., 2023)

Table 3: Summary of the reviewed, from 2019 to 2023, context- and attention-based cell nuclei instance segmentation and classification methods. SSL: Self-supervised Learning; SL: Supervised Learning; WSL: Weakly Supervised Learning.

fective multi-task learning approaches, therefore allowing the encoding of more stable features simultaneously useful, for instance, segmentation, instance classification, and semantic segmentation. Notably, by reducing feature map noise, and by capturing inter-instance dependencies, context (e.g., long-range, and multiscale) and attention methods have been extensively used to improve the ability of DL-based instance segmentation and classification methods in detecting small instances, a recurring challenge with both ConvNets and ViTs, but an important one as cell nuclei instances are very likely to be small when compared with the total image size (in pixels).

The results of the present review also reveal some misleading practices. For example, comparing neural networks with large differences in the number of parameters or optimizing the hyperparameters of new methodologies while not that of the baseline ConvNets. Another frequent practice that leads to misleading conclusions is comparing novel methodologies with works known to have inferior performance than the current top-performing solutions (e.g., comparing with UNet (Ronneberger et al., 2015), or UNet++ (Zhou et al., 2018), instead of with HoVer-Net (Graham et al., 2019b)). From a scientific standpoint, this can be of interest and contribute to a better understanding of the different methods. However, considering future clinical applications, algorithms with the best performance will always be preferred. Data availability is also a concern in nuclei instance segmentation and classification. While a pool of datasets is available, many derive from the same source/set of sources. This poses a challenge for assessing model robustness as while algorithms may have similar performance when

evaluated in IID data, the generalization to OOD settings may be undermined in some of the models. Other future directions worthy of careful consideration include data, and annotation efficient DL. As a frequent hurdle in medical image analysis is data and annotation availability, improving data efficiency could not only increase the discriminative power of DL models with smaller datasets but as well is a promising strategy to leverage some of the already existing large-scale, but unlabelled, datasets. This is where self-supervised, semi-supervised, and WSL take a major role, and although the literature reports some significant contributions, we argue most works regard image classification paradigms, whereas dense prediction tasks are mostly open ground. We hold the next breakthroughs in dense prediction in CPath could come from exploiting domain knowledge, and from introducing context- and attention-based mechanisms in algorithm design. These ideas can translate to more intrinsically interpretable models, increase model performance, and solve the challenge of generalizing to OOD data, a requirement if the models are to be trusted and deployed in clinical practice.

6.7. Future Research Directions

Certain methods that potentiate the encoding of contextual information seem to translate to increased generalization (e.g., ViTs (Ke et al., 2023)). However, we observe limited advantages with other approaches (Murtaza Dogar et al., 2021). This suggests that although context and attention-based mechanisms hold the potential to increase performance and generalization, it is still not clear what context priors are more useful or how

Method	Authors	Year	Public Repository	Reference
TransNuSeg	He et al.	2023	https://github.com/zhenqi-he/transnuseg	He et al. (2023)
ClusterSeg	Ke et al.	2023	https://github.com/lu-yizhou/ClusterSeg	Ke et al. (2023)
SMILE	Pan et al.	2023	https://github.com/panxipeng/nuclear_segandcls	Pan et al. (2023)
CircleFormer	Zhang et al.	2023	https://github.com/zhanghx-iiim-ahu/CircleFormer	Zhang et al. (2023a)
Patch. Contrast. Learn.	Boserup et al.	2023	https://github.com/nickeopti/bach-contrastive-segmentation	Boserup and Selvan (2023)
DRCA-Net	Dogar et al.	2023	https://github.com/Ghulam111/DRCA_Net	Dogar et al. (2023)
DAN-NucNet	Ahmad et al.	2023	https://github.com/ibtihajahmadkhan/DAN-Nuc-Net	Ahmad et al. (2023)
MSAL-Net	Ali et al.	2022	https://github.com/haideralimughal/MSAL-Net	Ali et al. (2022)
HistoSeg	Wazir and Fraz	2022	https://github.com/saadwazir/HistoSeg	Wazir and Fraz (2022)
HyLT	Luo et al.	2022	https://github.com/Roypic/LTUNet	Luo et al. (2022)
RCSAU-Net	Wang et al.	2022	https://github.com/antifen/Nuclei-Segmentation	Wang et al. (2022a)
CDNet	He et al.	2021	https://github.com/honglianghe/CDNet	He et al. (2021a)
MedT	Valanarasu et al.	2021	https://github.com/jeya-maria-jose/Medical-Transformer	Valanarasu et al. (2021)
Mask Guided Attention	Guo et al.	2021	https://github.com/RuoyuGuo/MaskGA_Net	Guo et al. (2021)
MASG-GAN	Zhang et al.	2021	https://github.com/wangpengyu0829/MASG-GAN	Zhang et al. (2021)
NucleiSegNet	Lal et al.	2021	https://github.com/shyamfec/NucleiSegNet	Lal et al. (2021)
Self-sup. Scale Pred	Sahasrabudhe et al.	2020	https://github.com/msahasrabudhe/miccai2020_self_sup_nuclei_seg	Sahasrabudhe et al. (2020)
NAS-SCAM	Liu et al.	2020	https://github.com/ZuhaoLiu/NAS-SCAM	Liu et al. (2020)

Table 4: Summary of the reviewed, from 2019 to 2023, context- and attention-based cell nuclei instance segmentation and classification methods, and that provide public access to the source code.

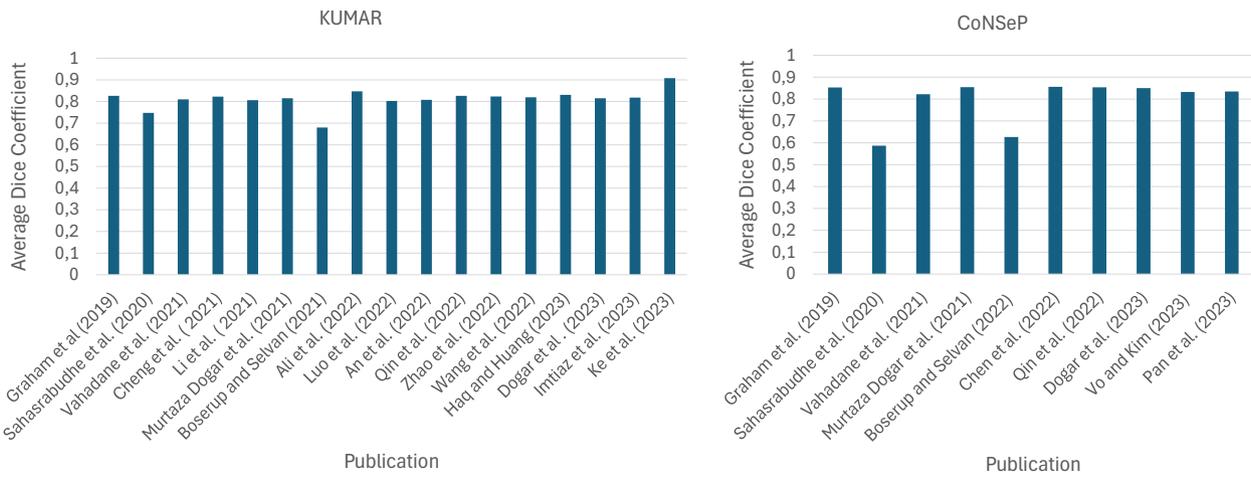


Fig. 11: Comparison of DC on the CoNSEP (Graham et al., 2019b) and KUMAR (Kumar et al., 2020) (test) datasets.

Reference	Type of Segmentation	DC (avg)	F1	J1	AJ1	PQ
Graham et al. (2019) Graham et al. (2019b)	Instance	0.8260	-	-	0.6180	0.5970
Zhou et al. (2019) Zhou et al. (2019c)	Instance	-	0.8458	-	0.6306	-
Qu et al. (2019) Qu et al. (2019)	Instance	0.8054	0.8639	-	0.6164	-
Liu et al. (2022) (Liu et al., 2022b)	Instance	0.7394	0.7863	0.6494	0.5430	-
Sahasrabudhe et al. (2020) Sahasrabudhe et al. (2020)	Semantic	0.7477	-	-	0.5354	-
Vahadane et al. (2021) Vahadane et al. (2021)	Instance	0.8100	-	-	0.5160	0.4400
Cheng et al. (2021) Cheng et al. (2021a)	Semantic	0.8227	-	-	0.6215	-
Lal et al. (2021) Lal et al. (2021)	Semantic	-	0.8136	0.6888	-	-
Li et al. (2021) Li et al. (2021b)	Semantic	0.8061	0.8602	-	0.6133	-
Murtaza Dogar et al. (2021) Murtaza Dogar et al. (2021)	Instance	0.8150	-	-	0.6150	0.5960
Boserup and Selvan (2022) Boserup and Selvan (2023)	Semantic	0.6797	-	-	-	-
Ali et al. (2022) Ali et al. (2022)	Semantic	0.8470	-	0.7060	-	-
Luo et al. (2022) Luo et al. (2022)	Instance	0.8025	-	0.6711	-	0.6021
An et al. (2022) An et al. (2022)	Instance	0.8080	-	0.6826	-	0.6040
Qin et al. (2022) Qin et al. (2022)	Semantic	0.8260	-	-	0.6360	-
Zhao et al. (2022) Zhao et al. (2022a)	Semantic	0.8230	-	-	0.6490	0.6100
Wang et al. (2022) Wang et al. (2022a)	Semantic	0.8200	0.8670	-	0.6190	-
Haq and Huang (2023) Haq and Huang (2022)	Semantic	0.8307	-	0.6872	-	-
Dogar et al. (2023) Dogar et al. (2023)	Instance	0.8150	-	-	0.6150	0.5960
Ahmad et al. (2023) Ahmad et al. (2023)	Semantic	-	0.8414	0.7403	-	-
Imtiaz et al. (2023) Imtiaz et al. (2023)	Semantic	0.8181	-	0.7089	-	-
Ke et al. (2023) (Ke et al., 2023)	Instance	0.9080	-	-	0.8230	0.7870
He et al. (2023) (He et al., 2023)	Instance	0.9081	0.8152	0.6949	-	-

Table 5: Summary of state-of-the-art results of context and attention-based neural networks (reviewed from 2019 to 2023) on the KUMAR (Kumar et al., 2020) dataset (test). HoVer-Net Graham et al. (2019b), CIA-Net Zhou et al. (2019c), and FullNet Qu et al. (2019) are included as a reference for comparison.

to optimally explore contextual information in cell nuclei instance segmentation and classification. Furthermore, although domain knowledge is relevant to impose useful inductive biases, we observe contextual information has been disregarded during data collection and annotation, which hinders downstream model optimization. For instance, cell nuclei annotations could be accompanied by pixel-level annotations of tissues or tissue

Reference	Type of Segmentation	DC (avg)	AJ1	PQ
Graham et al. (2019) Graham et al. (2019b)	Instance	0.8530	0.5710	0.5470
Doan et al. (2022) Doan et al. (2022)	Instance	0.8440	0.5860	0.5400
Sahasrabudhe et al. (2020) Sahasrabudhe et al. (2020)	Semantic	0.5870	0.1980	-
Vahadane et al. (2021) Vahadane et al. (2021)	Instance	0.8220	0.4850	0.3680
Murtaza Dogar et al. (2021) Murtaza Dogar et al. (2021)	Instance	0.8500	0.5650	0.5460
Boserup and Selvan (2022) Boserup and Selvan (2023)	Semantic	0.6266	-	-
Chen et al. (2022) Chen et al. (2022d)	Instance	0.8560	0.5580	0.5460
Qin et al. (2022) Qin et al. (2022)	Instance	0.8540	0.5610	0.5680
Dogar et al. (2023) Dogar et al. (2023)	Instance	0.8500	0.5650	0.5460
Vo and Kim (2023) Vo and Kim (2023)	Instance	0.8330	0.5150	0.4820
Pan et al. (2023) (Pan et al., 2023)	Instance	0.8350	0.5500	0.5300

Table 6: Summary of state-of-the-art results of context and attention-based neural networks (reviewed from 2019 to 2023) on the CoNSEP (Graham et al., 2019b) dataset (test). HoVer-Net Graham et al. (2019b) and SONNET Doan et al. (2022) are included as a reference for comparison.

concepts (e.g., glands), which could then be used to develop supervised methods that consider contextual features. In turn, debiasing the effects of context confounders is also relevant for domain generalization, therefore, providing along with the dataset and annotations metadata to help identify irrelevant context variables (e.g., stain manufacturer, scanner provider, patient age and gender, etc.) could aid in developing DL models more robust to spurious correlations between instances and their surrounding tissue context. In general, there is a need to understand the most relevant context priors for cell nuclei instance segmentation and classification. Nonetheless, there are a few

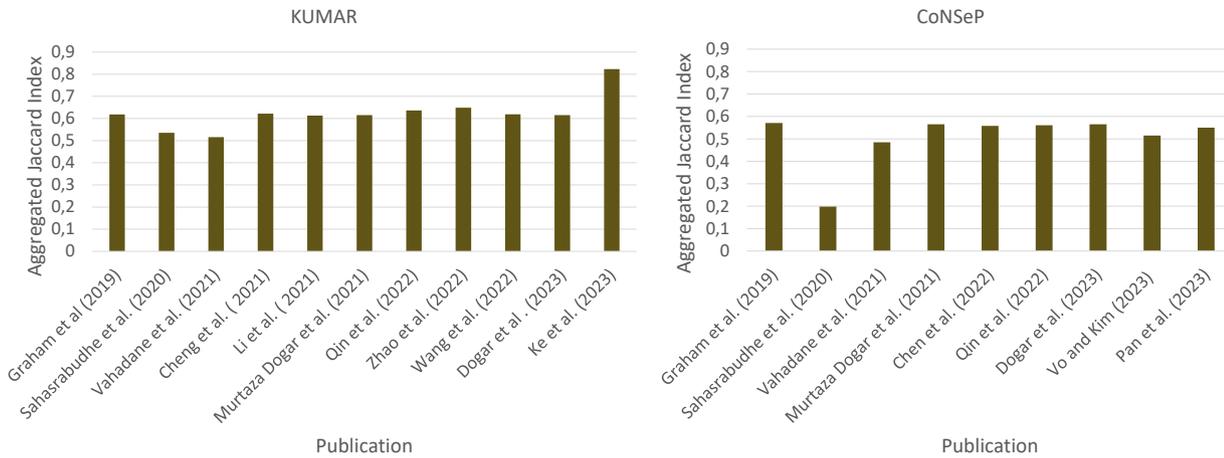


Fig. 12: Comparison of AJI on the CoNSEP (Graham et al., 2019b) and KUMAR (Kumar et al., 2020) (test) datasets.

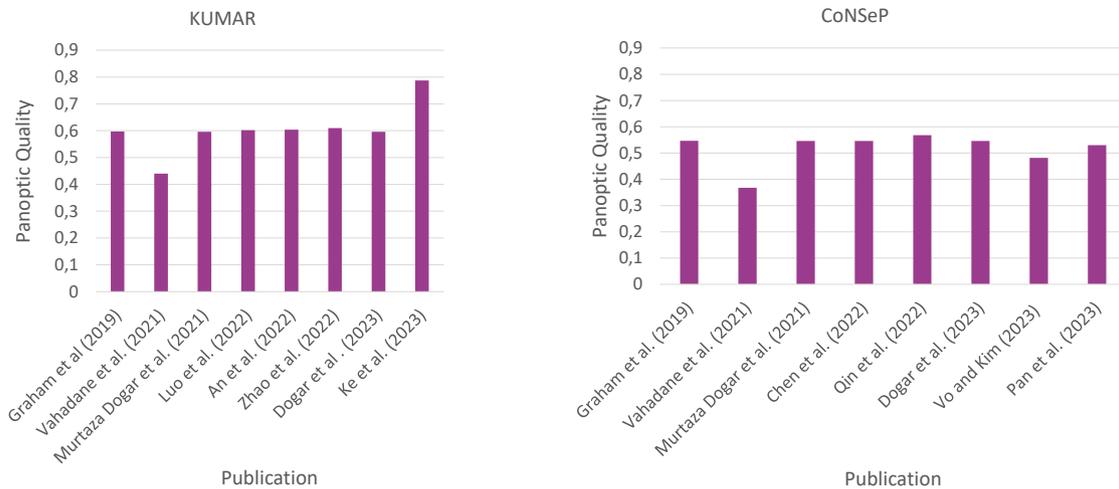


Fig. 13: Comparison of PQ on the CoNSEP (Graham et al., 2019b) and KUMAR (Kumar et al., 2020) (test) datasets.

promising research topics that could benefit from context and attention like causal representation learning, multimodal ML, and annotation-efficient learning.

6.7.1. Causal Representation Learning

Current DL methods are based on statistical learning, which relies on the assumption of identically distributed data. Indeed, although we frequently observe generalizability increases with the scale and diversity of training environments, the assumption of identical distribution can still be violated in the testing domains (e.g., due to sample selection bias, interventions, etc.). Furthermore, statistical models are incomplete descriptors of a system, as they rely on the distribution of observed random variables instead of the true generative factors (unobserved). In fact, statistical models could easily learn spurious correlations and thus catastrophically fail under covariate shift (Arjovsky et al., 2019). On the contrary, causal learning aims to uncover the causal structure and underlying causal variables from observations and outcomes, interventions, and distribution shifts (Pearl, 2009; Peters et al., 2017; Schölkopf et al., 2021). As an emerging research direction, causal representation learning

could thus lead to more robust predictors, invariant to distribution shifts (Peters et al., 2016). For instance, Huang et al. (2021) suggest a Structural Causal Model (SCM) for object detection tasks and consider that certain context factors are confounding variables for instances in a given observation. They thus suggest a causal intervention approach to deconfound the effect of context features that they demonstrate to improve generalization. In fact, suitable assumptions for a structural causal model representing the data-generating mechanism of tissue concepts and their corresponding context could lead to a representation that better approximates the ground truth (although unobserved and unknown) latent factors. Bashir et al. (2024) are among the first to suggest a regularization approach to learn context-invariant representations for tissues and cell nuclei semantic segmentation tasks. A causal perspective would allow us to argue that the learned representation is provably (at least partially) causal, as the causal features are the most likely to remain stable in varying tissue contexts. Likewise, preliminary work on causal segmentation in medical imaging as shown some degree of success. For example, a recurring inconvenience in medical imaging is that often training data originates from a single

source. This adds to the previously mentioned challenges of DL (e.g., batch effects; data imbalance, etc.) thus making single source domain generalization one of the most challenging objectives of DL-based segmentation algorithms. But robust learning with causal data augmentation could simulate distribution shifts (Ouyang *et al.*, 2023) and deconfound the effects of spurious context correlations (Ouyang *et al.*, 2023).

6.7.2. Graph Neural Networks

Not all learning tasks are defined in the Euclidean space (e.g., social network-based recommender systems, molecular structure prediction, etc.), meaning algorithms operating in grid-like data (e.g., ConvNets) are not appropriate for these tasks. At its core, WSI analysis and cell nuclei instance segmentation and classification could also be defined in a non-Euclidean space, as tissues and cells possess intricate relationships and interdependencies. With this view, ViTs are a more flexible alternative as they allow reasoning over long-range context, but ViTs are a special case of a more general and flexible family of parametric hypothesis, namely Graph Neural Networks (GNNs) (Veličković, 2023). GNNs represent data in a graph structure defined as a set of node vectors (objects) connected by edges (modelling inter-object relationships). Multiple tasks are thus allowed with GNNs such as node classification, link prediction, or node clustering. For instance, Hassan *et al.* (2022) use domain knowledge and, starting from a set of pre-detected and segmented nuclei, propose a GNN (Message Passing Network (MPN)) where each node represents a nucleus and edges encode the relation from each nucleus to its neighbours. Through reasoning over spatial and semantic context, the suggested GNN method allows the clustering of cell nuclei, which is used to refine nuclei classification on the premise that each nucleus in a cluster should be of the same cell type.

Due to the ability of GNNs to learn better object-level features (Xu *et al.*, 2022) and inter-instance relations (Yang *et al.*, 2020), these enable a more accurate instance segmentation in challenging settings like touching and overlapping instances (Xu *et al.*, 2022; Yang *et al.*, 2020). Likewise, by encoding inter-instance relations (context inductive biases), GNNs demonstrate the potential to refine instance masks in cluttered scenarios (Xie *et al.*, 2022a), a commonly found condition in H&E-stained microscopy images. But these are still largely underexplored in CPath, leaving room for further research in this direction to design algorithms that can effectively encode contextual priors for reliable cell nuclei instance segmentation and classification.

6.7.3. Multimodal Machine Learning

Until recently, most ML models are optimized to achieve a decision using data from a single modality. However, when making decisions, clinical experts commonly consider data from multiple streams of information (e.g., imaging data, clinical reports, molecular profiling data, etc.). Therefore, to close the gap between artificial and biological intelligence, ML algorithms should also consider integrating data from multiple modalities. The complementarity of multiple data types thus gives more context (i.e., semantic context/ prior-knowledge) to

the learned representations, which could increase generalization. Nonetheless, there are a few challenges, namely regarding *what* types of data to combine, *when* to use multimodality vs single modality and *how* to best fuse the multiple information streams. In a recent example, Li *et al.* (2024) design an attention-based multi-modal approach for unsupervised medical imaging segmentation tasks, where a transformer-based text-guided model is shown to improve the quality of generated pseudo-labels for unlabeled medical images. Interestingly, the authors also show the effectiveness of the approach in a nuclei segmentation task.

Nonetheless, a challenge with multimodal ML is the availability of datasets with paired samples of multiple data types. For instance, representation learning with language and imaging data (Huang *et al.*, 2023; Lu *et al.*, 2023; Lai *et al.*, 2023; Naseem *et al.*, 2023) is an emerging direction in CPath. But if we consider other modalities, like radiology-histopathology (Hamidinekoo *et al.*, 2021), genomics-histopathology (Chen *et al.*, 2022b), or radiology-genomics-histopathology (Chen *et al.*, 2021) data integration, literature and available datasets are only slowly gaining interest (Vanguri *et al.*, 2022). But to advance future efforts, experts should start by collecting, curating, and annotating large-scale multimodal datasets that could then be used to devise more robust algorithms, not only for nuclei instance segmentation and classification, but also for downstream applications that integrate nuclear features with other data modalities. Depending on the available data, (e.g., radiology-histopathology), this could even correspond to learning a representation expressing contextual information like global, hierarchical and prior context. Moreover, aside from integrated language and vision data (Wu *et al.*, 2023), we observe there is no current literature reporting the integration of multiple data modalities to improve nuclei instance segmentation and classification tasks. Therefore, although multimodal approaches are believed to improve model performance and OOD generalization, there is no consensus if such an approach would benefit nuclei instance segmentation and classification tasks. If we bring context and attention to the discussion, further questions arise concerning primarily if and how these mechanisms could be used to better combine multiple data types.

6.7.4. Self-supervised Learning

Self-supervised pretraining on large CV datasets or on large volumes of unlabeled histopathology data followed by task-specific fine-tuning has been shown to improve model performance (Ciga *et al.*, 2022a; Kang *et al.*, 2023). Indeed, self-supervised pathology-tailored foundation models are effective in solving multiple tasks, even generalizing to distributions beyond those seen during training (Filiot *et al.*, 2023; Huang *et al.*, 2023; Wang *et al.*, 2022b). But while SSL methods were shown to generalize better than SL on OOD data (Shi *et al.*, 2023), these are still prone to learning spurious correlations (Hamidieh *et al.*, 2022; Chhipa *et al.*, 2023), particularly on subgroups least represented in the dataset (Hamidieh *et al.*, 2022). However, as H&E-stained brightfield microscopy images are rich in contextual information, a viable solution to improve generalization could stem from using context priors and attention mechanisms

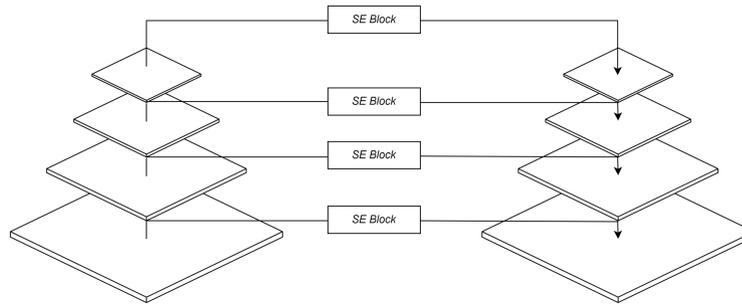


Fig. 14: Squeeze and Excitation-Feature Pyramid Network (SE-FPN). We add a SE-block before every lateral connection of the FPN module.

to impose inductive biases that enable SSL algorithms more robust to distribution shifts. For instance, the self-attention mechanism in ViTs has been shown to decrease model bias towards spurious correlations (Paul and Chen, 2022; Ghosal and Li, 2023), especially considering larger models and large-scale pretraining (Zhang et al., 2022a; Ghosal and Li, 2023), which could be attributed to learning better contextual dependencies (Paul and Chen, 2022).

In turn, some SSL paradigms promote invariance to synthetic data transformations, which increases robustness to real world transforms (i.e., distribution shifts). The effectiveness of these methods is higher the better the chosen transforms simulate expected spurious factors of variation, but as the challenge of generalizing cell nuclei instance segmentation and classification algorithms is still unsolved, more effective (domain specific) data augmentation strategies need to be devised. Here, context priors could also play a major role as certain transformations could corrupt the semantics of the data. In fact, there is work showing the effectiveness of context-driven data augmentation not only for dense prediction tasks in CV (Dvornik et al., 2021), but also for cell nuclei instance segmentation (Lin et al., 2022b), where a GAN-based approach is used to achieve context consistency after mixing the instances of two images.

6.7.5. Weakly Supervised Learning

One of the key strengths of WSL is the ability to harness large-scale datasets with minimal manual labeling efforts, easing the need for exhaustive annotation. This is particularly valuable in CPath, where collecting fully labeled data is impractical. Several strategies can be adopted for weakly-supervised cell nuclei instance segmentation and classification. For example, generative models show promise in solving several segmentation tasks (Li et al., 2021a; Lei et al., 2020; Nema et al., 2020; Majurski et al., 2019), but could be adapted to WSL. DetectGAN is suggested for data-efficient object detection in medical imaging (Liu et al., 2019a). Generative adversarial learning and

template matching allow localizing objects in an image with weak supervision (Diba et al., 2019), and another work shows the potential of GANs for weakly supervised brain tumour segmentation (Yoo et al., 2022). An idea we suggest is Variational Autoencoder (VAE) regularization, which has been successful in SL brain tumour segmentation (Myronenko, 2019). The design could be extended with context encoding (Zimmerer et al., 2018) and adapted to weak supervision (e.g., point annotations). This could increase model robustness to input perturbations. Moreover, besides promoting OOD generalization, generative models also have the advantage of producing disentangled and interpretable representations, especially if jointly optimized for reconstruction fidelity and interpretability (Adel et al., 2018). Another viable future direction in weakly supervised nuclei instance segmentation and classification is prior knowledge to guide the learning task (Zhang et al., 2022b). For instance, predicting spatial semantic context, like cellular communities' orientation along an image (or patch in MIL), could be exploited in a multi-task setting as a regularization strategy.

In turn, Segment Anything (SAM) (Kirillov et al., 2023) is introduced as a foundation model for image segmentation. But SAM performance is lower when directly applied to histopathology images (Zhang et al., 2023b; Huang et al., 2024), as the covariate and concept shifts limit the zero-shot performance in histopathology segmentation tasks. Notably, this is exacerbated in cell nuclei segmentation (Deng et al., 2023), underscoring the effectiveness of domain-specific models and domain-specific pretraining as the prevailing paradigms. Yet, there is evidence that generalization increases with the number of input prompts (Deng et al., 2023; Xu et al., 2024), hinting that SAM and equivalent foundation models could alleviate the annotation burden by predicting cell nuclei pseudo-masks from point or bounding box annotations for WSL methods. This unveils a particularly interesting application of SAM and equivalent foundation models: predicting relevant context priors in histopathology images from only a few input prompts,

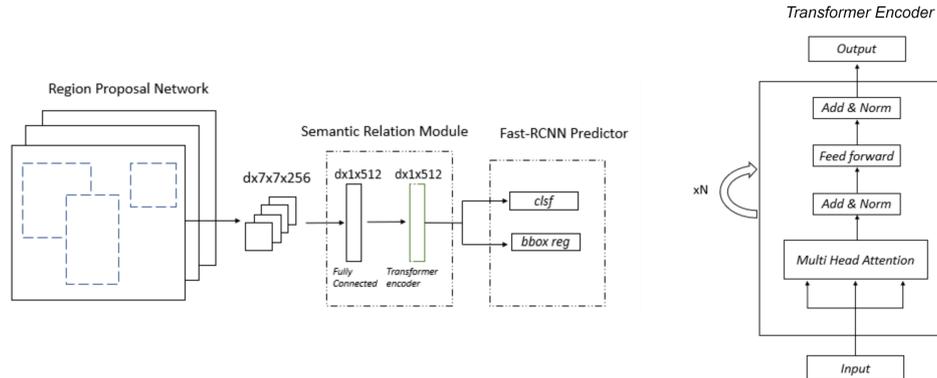


Fig. 15: Left: Semantic Relation Module. We replace the second fully connected layer of the Multi-layer Perceptron Head (MLP-Head) with a transformer encoder. Right: Transformer Encoder.

like cellular communities, tumours, etc. These priors could then be used to optimize cell nuclei instance segmentation and classification models, for example, as pseudo-labels or as additional input variables (i.e., attention maps). Another promising use case for these foundation models is data augmentation. For instance, using the U-Net (Ronneberger et al., 2015) architecture in a cell nuclei segmentation task, Zhang et al. (2023c) augment H&E-stained microscopy images with boundary and segmentation prior attention maps generated from SAM, which lead to improved performance.

Overall, despite the promises of WSL, a major limitation is the generation of noisy/incorrect pseudo-masks. This could be attributed to confounding contextual information, but recent work demonstrates the potential of causal interventions to enable more effective WSL as this approach leads to more reliable pseudo-masks (Chen et al., 2022e; Zhang et al., 2020).

7. Colon Nuclei Instance Segmentation and Classification: A Multi-centre Comparative Study

In this section, we consider a case study with a comparative analysis to assess how extending an instance segmentation and classification model, namely Mask-RCNN (He et al., 2016), with context and attention-based mechanisms impacts model performance and generalization. We aim to assess if simple mechanisms offer any advantages or if more tailored solutions would be required. We thus extend a baseline Mask-RCNN (He et al., 2016) with SE-blocks (Hu et al., 2018) as well as transformer layers.

7.1. Datasets

This work makes use of the Lizard dataset (Graham et al., 2021a) consisting of H&E-stained WSI tiles, from six different sources at 20x objective magnification. The data contains annotations for six nuclei types, namely, epithelial cells, connective tissue, plasma, lymphocytes, neutrophils, and eosinophils.

This is one of the biggest fully annotated public datasets for cell nuclei segmentation, with a total of 431,913 annotated nuclei. Despite this, the data are highly imbalanced. For example, the lymphocytes category contains 210,372 annotated nuclei, whereas the eosinophils class represents less than 1% of all instances. For training and validation, we consider data from the 5 centres on the training set of the CoNIC challenge Graham et al. (2021b). The experiments in this work thus include 5 data sources: GLaS (Sirinukunwattana et al., 2017), CRAG (Graham et al., 2019a), CoNSeP (Graham et al., 2019b), Digest-Path (Da et al., 2022), and PanNuke (Gamper et al., 2020b). To evaluate the generalizability of the context- and attention-based mechanisms developed through this work, we resort to *leave-one-domain-out* validation (Zhou et al., 2022), where each individual source is treated as a domain.

Model	Batch Size	LR	W Decay	Optimizer	N Epochs	Dropout
Mask-RCNN-ResNet50-FPN	1	0.001	$5e^{-4}$	SGD	100	-
Mask-RCNN-SE-ResNet50-FPN	1	0.001	$5e^{-4}$	SGD	100	-
Mask-RCNN-ResNet50-SE-FPN	1	0.001	$5e^{-4}$	SGD	100	-
Mask-RCNN-ResNet50-FPN-SRM	1	$5e^{-6}$	$5e^{-4}$	AdamW	100	0.5
Mask-RCNN-SRM-ResNet50-FPN	1	$5e^{-5}$	$5e^{-4}$	AdamW	100	0.4

Table 7: Summary of implementation details.

7.2. Methods

We first consider Mask-RCNN (He et al., 2017) with a ResNet-50 (He et al., 2016) backbone and FPN (Lin et al., 2017) as a baseline model (Mask-RCNN-ResNet50-FPN). We then propose two variants of Mask-RCNN (He et al., 2017) with attention-based mechanisms. On the one hand, we add a SE-block (Hu et al., 2018) before every lateral connection of the FPN module, which we call the SE-FPN (Mask-RCNN-ResNet50-SE-FPN). We hypothesised that this feature refinement step between the FPN top-down and bottom-up paths would reduce the FPN sensitivity to noise from non-foreground

Comparison of Mask-RCNN-ResNet50-FPN with Mask-RCNN-ResNet50-SE-FPN, and Mask-RCNN-ResNet50-FPN-SRM														
model	val set	map ₅₀ bbox	map ₅₀ segm	neu	epi	lym	pls	eos	con	pq	pq det	pq seg	multi	pq+
mask-rcnn-resnet50-fpn	glas	0.7348	0.7464	0.7504	0.5107	0.5439	0.5383	0.2935	0.4510	0.5309	0.6868	0.7658	0.3790	0.3430
	dpath	0.5826	0.6200	0.6200	0.5228	0.5307	0.5307	0.2142	0.4610	0.5792	0.7480	0.7685	0.3430	0.3430
	cons	0.8707	0.9642	0.9645	0.4998	0.5457	0.4333	0.4947	0.5691	0.5654	0.7414	0.7620	0.4845	0.4845
	pann	0.4819	0.4831	0.4612	0.5254	0.5186	0.3184	0.3241	0.5378	0.5689	0.7344	0.7450	0.4476	0.4476
	crag	0.8091	0.8017	0.8088	0.5237	0.5152	0.3539	0.3010	0.5004	0.5197	0.6710	0.7247	0.3792	0.3792
	overall	0.6958 ± 0.1439	0.7168 ± 0.1683	0.7168 ± 0.1683	0.2146 ± 0.1695	0.5165 ± 0.0098	0.5308 ± 0.0125	0.3486 ± 0.0449	0.3255 ± 0.0923	0.5039 ± 0.0448	0.5528 ± 0.0232	0.7163 ± 0.0313	0.7532 ± 0.0164	0.4066 ± 0.0516
mask-rcnn-resnet50-se-fpn	glas	0.7496	0.7534	0.7534	0.5199	0.5419	0.3471	0.3114	0.4459	0.5317	0.6886	0.7661	0.3884	0.3884
	dpath	0.5746	0.5408	0.5408	0.5285	0.5586	0.3249	0.2168	0.4709	0.5859	0.7556	0.7692	0.3528	0.3528
	cons	0.8912	0.9355	0.9355	0.4951	0.5453	0.4464	0.4779	0.5731	0.5629	0.7389	0.7607	0.4807	0.4807
	pann	0.6500	0.6500	0.4400	0.5300	0.5200	0.3200	0.3400	0.5400	0.5700	0.7300	0.7500	0.4500	0.4500
	crag	0.7542	0.7630	0.0661	0.5304	0.5245	0.3820	0.2778	0.5100	0.5235	0.6747	0.7271	0.3818	0.3818
	overall	0.7239 ± 0.0977	0.7285 ± 0.1198	0.7285 ± 0.1198	0.2068 ± 0.1480	0.5208 ± 0.0122	0.5381 ± 0.0129	0.3641 ± 0.0426	0.3248 ± 0.0793	0.5080 ± 0.0418	0.5548 ± 0.0215	0.7176 ± 0.0281	0.7546 ± 0.0139	0.4106 ± 0.0432
mask-rcnn-resnet50-fpn-srm	glas	0.6427	0.3052	0.0791	0.4520	0.3735	0.2445	0.0631	0.2681	0.4317	0.5694	0.7469	0.2467	0.2467
	dpath	0.5498	0.5498	0.0205	0.5507	0.4668	0.2251	0.1005	0.3751	0.5477	0.6971	0.7782	0.2898	0.2898
	cons	0.8575	0.8908	0.2656	0.4589	0.5008	0.3413	0.4366	0.5222	0.5314	0.6996	0.7580	0.4209	0.4209
	pann	0.8574	0.8676	0.3715	0.5032	0.4823	0.2685	0.2957	0.4891	0.5357	0.6935	0.7423	0.4015	0.4015
	crag	0.7078	0.7124	0.045	0.4992	0.4604	0.2979	0.1373	0.4530	0.4860	0.6342	0.7180	0.3155	0.3155
	overall	0.7230 ± 0.1102	0.6652 ± 0.1987	0.6652 ± 0.1987	0.1563 ± 0.1259	0.4928 ± 0.0324	0.4568 ± 0.0401	0.2755 ± 0.0374	0.2066 ± 0.1275	0.4215 ± 0.0831	0.5065 ± 0.0391	0.6588 ± 0.0464	0.7487 ± 0.0180	0.3349 ± 0.0606
mask-rcnn-se-resnet50-fpn	glas	0.6365	0.6617	0.1266	0.3689	0.4779	0.2797	0.1089	0.3923	0.4652	0.6047	0.763	0.2924	0.2924
	dpath	0.4506	0.4655	0.0124	0.3476	0.5284	0.1402	0.1272	0.3426	0.5019	0.6538	0.7613	0.2497	0.2497
	cons	0.8886	0.9510	0.3999	0.4556	0.5315	0.2688	0.3237	0.4972	0.5632	0.7357	0.7649	0.4128	0.4128
	pann	0.5010	0.5180	0.4633	0.4860	0.4814	0.2975	0.3035	0.5149	0.5574	0.7202	0.7441	0.4244	0.4244
	crag	0.6654	0.6794	0.0308	0.3441	0.4553	0.3106	0.0325	0.3471	0.4151	0.537	0.7154	0.2534	0.2534
	overall	0.6284 ± 0.1530	0.6551 ± 0.1691	0.6551 ± 0.1691	0.2066 ± 0.1888	0.4004 ± 0.0589	0.4949 ± 0.0300	0.2594 ± 0.0613	0.1792 ± 0.1145	0.4188 ± 0.0735	0.5006 ± 0.0561	0.6503 ± 0.0736	0.7497 ± 0.0187	0.3265 ± 0.0767
mask-rcnn-srm-resnet50-fpn	glas	0.6751	0.1529	0	0.455	0.5157	0.0851	0	0.311	0.4825	0.6304	0.7577	0.2278	0.2278
	dpath	0.4766	0.3293	0.1898	0.4739	0.4702	0.309	0.1219	0.4009	0.5484	0.7109	0.765	0.3276	0.3276
	cons	0.7075	0.5339	0	0.2673	0.4485	0.0502	0	0.2034	0.3507	0.4826	0.7217	0.1616	0.1616
	pann	0.8238	0.6291	0	0.3049	0.2915	0.1237	0	0.2049	0.3812	0.5041	0.74	0.1542	0.1542
	crag	0.6508	0.4806	0.1564	0.4633	0.5072	0.2515	0.0811	0.3975	0.4607	0.5885	0.7142	0.3095	0.3095
	overall	0.6654 ± 0.1121	0.5393 ± 0.1122	0.5393 ± 0.1122	0.0308 ± 0.0855	0.3441 ± 0.0882	0.4553 ± 0.0813	0.3106 ± 0.0995	0.0325 ± 0.0514	0.3471 ± 0.0873	0.4151 ± 0.0711	0.5370 ± 0.0836	0.7154 ± 0.0197	0.2534 ± 0.0722

Table 8: Comparison of Mask-RCNN-ResNet50-FPN with Mask-RCNN-ResNet50-SE-FPN, Mask-RCNN-SE-ResNet50-FPN, Mask-RCNN-SRM-ResNet50-FPN, and Mask-RCNN-ResNet50-FPN-SRM in a colon nuclei instance segmentation and classification task. The initials “neu”, “epi”, “lym”, “pls”, “eos”, and “con” indicate mPQ^+ for each nucleus type. Results in bold indicate the best mPQ^+ .

Paired Samples T-Test for the Mask-RCNN Experiments																					
null hypothesis	val set	neu		epi		lym		pls		eos		con		pq_det		pq_seg		pq		mpq+	
		t-stat	p-val	t-stat	p-val	t-stat	p-val	t-stat	p-val												
mask-rcnn-resnet50-se-fpn > mask-rcnn-resnet50-fpn	glas	0.7320	0.4644	5.2427	<0.0001	-0.2195	0.8263	1.5973	0.1106	2.8566	0.0044	-1.4217	0.1555	3.6597	0.0003	-1.0845	0.2785	3.4555	0.0006	3.1655	0.0016
	dpath	-0.9670	0.3337	6.2508	<0.0001	12.0822	<0.0001	2.6934	0.0071	0.8739	0.2281	5.8705	<0.0001	9.0679	<0.0001	1.8547	0.0638	10.9622	<0.0001	9.9595	<0.0001
	cons	0.8869	0.3785	-1.3743	0.1742	-0.8178	0.4166	0.6014	0.5498	-0.9826	0.3296	0.8242	0.4129	-0.5589	0.5782	-1.0243	0.3096	-0.7573	0.4517	0.1941	0.8467
	pann	-0.6228	0.5347	1.8002	0.0745	2.1394	0.0346	-1.3756	0.1717	-0.5498	0.5836	0.9819	0.3283	-0.4647	0.643	2.1738	0.0318	0.0467	0.9628	0.4549	0.6501
	crag	-3.0945	0.0020	5.5218	<0.0001	7.8009	<0.0001	5.2249	<0.0001	-2.9030	0.0037	7.4053	<0.0001	3.8205	0.0001	3.6464	0.0003	5.3357	<0.0001	8.5734	<0.0001
mask-rcnn-resnet50-se-fpn > mask-rcnn-resnet50-fpn-srm	glas	6.8748	<0.0001	11.7056	<0.0001	22.5822	<0.0001	11.2663	<0.0001	12.8686	<0.0001	37.4220	<0.0001	27.3278	<0.0001	12.235	<0.0001	27.7169	<0.0001	31.26	<0.0001
	dpath	-1.0570	0.2906	-7.0787	<0.0001	26.6229	<0.0001	17.5991	<0.0001	10.6613	<0.0001	30.3162	<0.0001	30.0842	<0.0001	-14.3029	<0.0001	21.8959	<0.0001	13.921	<0.0001
	cons	1.9370	0.0572	2.6935	0.7973	4.5987	0.0090	4.4060	<0.0001	0.2580	<0.0001	4.6524	<0.0001	4.0888	0.0001	0.9253	0.3584	4.6578	<0.0001	13.091	<0.0001
	pann	1.4672	0.1451	2.1132	0.0368	3.8403	<0.0001	4.0339	<0.0001	1.2108	0.2281	5.8705	<0.0001	4.9897	<0.0001	1.6873	0.0943	5.3082	<0.0001	16.886	<0.0001
	crag	0.8642	0.3876	5.0746	<0.0001	18.4799	<0.0001	12.1518	<0.0001	8.0579	<0.0001	22.1590	<0.0001	27.0217	<0.0001	9.7992	<0.0001	32.5868	<0.0001	33.256	<0.0001
mask-rcnn-resnet50-fpn > mask-rcnn-resnet50-fpn-srm	glas	6.8530	<0.0001	9.8997	<0.0001	23.0282	<0.0001	10.4970	<0.0001	12.1580	<0.0001	37.4541	<0.0001	26.2797	<0.0001	12.0827	<0.0001	26.6727	<0.0001	31.115	<0.0001
	dpath	-0.4312	0.6639	-9.2711	<0.0001	19.4217	<0.0001	15.8981	<0.0001	10.7513	<0.0001	28.8324	<0.0001	26.9082	<0.0001	-13.7927	<0.0001	18.5482	<0.0001	10.098	<0.0001
	cons	1.5949	0.1157	2.9231	0.0048	4.2319	<0.0001	4.3483	<0.0001	0.4117	0.6820	4.2142	<0.0001	4.4034	<0.0001	1.3173	0.1925	5.0382	<0.0001	13.621	<0.0001
	pann	2.0367	0.0441	1.8431	0.0680	2.4267	0.0168	4.8150	<0.0001	1.4769	0.1425	5.6564	<0.0001	5.5451	<0.0001	1.1294	0.2611	5.6305	<0.0001	16.707	<0.0001
	crag	3.2289	0.0013	1.6125	0.1070	13.3727	<0.0001	8.6679	<0.0001	9.4167	<0.0001	16.9443	<0.0001	25.5689	<0.0001	9.1502	<0.0001	30.5842	<0.0001	27.634	<0.0001

Table 9: Paired samples t-test for the Mask-RCNN experiments at the level of the FPN. Values in bold indicate significant differences (p-value < 0.05). The experiments for Mask-RCNN with context and attention mechanisms at the level of the ResNet-50 backbone always demonstrated significant differences (p-value < 0.0001). The initials “neu”, “epi”, “lym”, “pls”, “eos”, and “con” indicate mPQ^+ for each nucleus type.

objects, especially in lower-resolution feature maps. Second, we consider an object relation network, without geometric features, that we term the Semantic Relation Module (SRM) (Mask-RCNN-ResNet50-FPN-SRM). The SRM considers a transformer-encoder to replace the second fully connected layer of the MLP-Head. The objective was to enrich the semantics of every object candidate with cross-attention and long range-dependencies, while also exploring the potential of hybridizing ConvNets with transformer neural networks. Figs. 14 and 15 illustrate, respectively, the proposed SE-FPN and SRM. In addition, not only do we experiment with replacing the ResNet-50 backbone of Mask-RCNN with SE-blocks, the designated Mask-RCNN-SE-ResNet50-FPN, as also suggest extending the last convolutional block of ResNet-50 with a SRM, the so-called Mask-RCNN-SRM-ResNet50-FPN.

But a known limitation of Mask-RCNN (He et al., 2017) is the difficulty in detecting small instances and that are densely distributed in challenging contexts (Min et al., 2022). Therefore, we also extend the domain-specific HoVer-Net Graham et al. (2019b) with context and attention mechanisms to understand if it would compare favourably with the Mask-RCNN-based models and if incorporating attention and context would increase generalization. As before, we extend the HoVer-Net ResNet-50 backbone with the SE-blocks at every block and

with a SRM at the lowest level, that we term, respectively, SE-HoVer-Net and SRM-HoVer-Net.

Implementation Details

The Mask-RCNN experiments were implemented in *python* with *Pytorch* version 1.10.2. For the Mask-RCNN-ResNet50-FPN, Mask-RCNN-ResNet50-SE-FPN, and Mask-RCNN-SE-ResNet50-FPN we observed that a batch size of 1 resulted in better results, with a learning rate of 0.001, weight decay equal to 5×10^{-4} , and Stochastic Gradient Descent (*SGD*) optimizer (momentum=0.9). For the Mask-RCNN-Resnet-50-FPN-SRM experiments, we considered 8 attention heads in the transformer layer, with a dropout (Srivastava et al., 2014) probability of 0.5 and GeLU activation (Hendrycks and Gimpel, 2016). We use the Adaptive Moment with decoupled weight decay (*AdamW* (Loshchilov and Hutter, 2017)) optimizer with a learning rate of 5×10^{-6} and weight decay equal to 5×10^{-4} . In turn, for the Mask-RCNN-SRM-ResNet50-FPN we observe a dropout probability of 0.4 is the optimal choice, with a learning rate of 5×10^{-5} , weight decay equal to 5×10^{-4} , and *SGD* optimizer (momentum=0.9). For all Mask-RCNN variants, we consider 100 training epochs and opt for a linear learning rate warm-up in the first 1000 iterations of the first epoch (starting at 1×10^{-6} for Mask-RCNN-ResNet50-FPN, Mask-RCNN-ResNet50-SE

Comparison of HoVer-Net with SE-HoVer-Net, and SRM-HoVer-Net											
model	val set	neu	epi	lym	pls	eos	con	pq	pq det	pq seg	multi pq+
hoover-net	glas	0.0491	0.4799	0.5861	0.2924	0.1686	0.3439	0.5423	0.6952	0.7686	0.3197
	dpath	0.0175	0.3600	0.5225	0.2688	0.0762	0.3895	0.5359	0.6798	0.7797	0.2724
	cons	0.1556	0.499	0.5941	0.4075	0.2219	0.4587	0.5322	0.6840	0.7635	0.3895
	pann	0.2394	0.5095	0.5368	0.3227	0.0642	0.4438	0.5581	0.7115	0.7541	0.3528
	crag	0.0172	0.4574	0.5623	0.3646	0.1031	0.4044	0.4860	0.6169	0.7359	0.3182
	overall	0.0342 ± 0.0880	0.4539 ± 0.0536	0.5586 ± 0.0272	0.3239 ± 0.0498	0.1024 ± 0.0597	0.4038 ± 0.0408	0.5297 ± 0.0241	0.6759 ± 0.0322	0.7601 ± 0.0148	0.3259 ± 0.0390
se-hoover-net	glas	0.1041	0.4853	0.5848	0.3239	0.1963	0.3860	0.5448	0.6985	0.7681	0.3467
	dpath	0.0461	0.4095	0.5229	0.3069	0.1116	0.4321	0.5657	0.7085	0.7905	0.3049
	cons	0.2132	0.5055	0.5883	0.4017	0.3887	0.4664	0.5386	0.6928	0.7758	0.4273
	pann	0.1078	0.5057	0.5559	0.3125	0.1382	0.4420	0.5591	0.7120	0.7613	0.3437
	crag	0.0297	0.4443	0.5600	0.3944	0.1017	0.3971	0.4871	0.6195	0.7335	0.3211
	overall	0.0633 ± 0.0644	0.4669 ± 0.0377	0.5614 ± 0.0236	0.3432 ± 0.0414	0.1484 ± 0.1059	0.4227 ± 0.0295	0.5375 ± 0.0277	0.6844 ± 0.0341	0.7654 ± 0.0189	0.3442 ± 0.0422
srm-hoover-net	glas	0.0994	0.4826	0.5838	0.355	0.1981	0.3244	0.5407	0.6918	0.7698	0.3406
	dpath	0.0448	0.4143	0.5199	0.3002	0.1073	0.4288	0.5661	0.7122	0.7869	0.3026
	cons	0.1822	0.5115	0.5870	0.4482	0.3413	0.4704	0.5465	0.7037	0.7736	0.4234
	pann	0.2435	0.5249	0.5772	0.3459	0.166	0.4607	0.5671	0.7206	0.7638	0.3864
	crag	0.0203	0.4868	0.5705	0.4402	0.1235	0.4391	0.5065	0.6420	0.7388	0.3467
	overall	0.0548 ± 0.0838	0.4808 ± 0.0382	0.5666 ± 0.0246	0.3692 ± 0.0573	0.1591 ± 0.0834	0.4171 ± 0.0523	0.5445 ± 0.0221	0.6929 ± 0.0277	0.7662 ± 0.0158	0.3552 ± 0.0414

Table 10: Comparison of HoVer-Net with SE-HoVer-Net, and SRM-HoVer-Net in a colon nuclei instance segmentation and classification task. The initials “neu”, “epi”, “lym”, “pls”, “eos”, and “con” indicate mPQ^+ for each nucleus type. Results in bold indicate the best mPQ^+ .

Paired Samples T-Test for the HoVer-Net Experiments																					
null hypothesis	val set	neu		epi		lym		pls		eos		con		pq.det		pq_seg		pq		mpq+	
		t-stat	p-val	t-stat	p-val	t-stat	p-val	t-stat	p-val	t-stat	p-val	t-stat	p-val	t-stat	p-val	t-stat	p-val	t-stat	p-val	t-stat	p-val
se-hoover-net > hoover-net	glas	3.1444	0.0017	1.7008	0.0894	-0.0829	0.9339	0.4781	0.6327	1.4526	0.1468	11.9248	<0.0001	2.5512	0.0109	-1.1093	0.2677	2.5790	0.0101	6.4410	<0.0001
	dpath	2.7979	0.0052	18.9057	<0.0001	-8.4985	<0.0001	10.3395	<0.0001	4.3424	<0.0001	16.2891	<0.0001	23.8710	<0.0001	27.0760	<0.0001	30.4551	<0.0001	15.5477	<0.0001
	cons	1.2768	0.2064	0.6534	0.5159	0.9567	0.3424	-0.0828	0.9343	3.1539	0.0025	0.4014	0.6895	0.8514	0.3970	1.0310	0.3064	0.8308	0.4092	3.2680	0.0018
	pann	-0.4930	0.6230	-1.5815	0.1166	-0.9473	0.3455	-2.1102	0.0371	2.2412	0.0270	1.1240	0.2634	0.0690	0.9446	1.0641	0.2896	0.1777	0.8593	0.1069	0.9150
	crag	3.4205	0.0006	-7.0200	<0.0001	-7.4141	<0.0001	2.8204	0.0048	-1.1873	0.2352	0.3356	0.7372	1.5795	0.1144	-2.1720	0.0299	0.9240	0.3556	-4.0071	<0.0001
srm-hoover-net > se-hoover-net	glas	-3.1949	0.0015	-0.8353	0.4038	4.8248	<0.0001	1.4852	0.1379	-3.2505	0.0012	-16.7237	<0.0001	-4.6036	<0.0001	3.6949	0.0002	-3.8037	0.0002	-5.6549	<0.0001
	dpath	0.5985	0.5516	1.5707	0.1164	-5.4150	<0.0001	-5.0276	<0.0001	-0.0069	0.9945	-1.6899	0.0912	3.2503	0.0012	10.9651	<0.0001	0.5211	0.6023	-4.7306	<0.0001
	cons	-1.4960	0.1396	0.1797	0.8380	-1.4831	0.1430	0.4376	0.6632	-1.2644	0.2107	1.3096	0.1951	2.0413	0.0454	-0.6245	0.5345	1.8040	0.0760	-0.6977	0.4880
	pann	1.3364	0.1841	3.5645	<0.0001	2.8376	0.0054	0.4112	0.6817	0.7398	0.4610	0.6538	0.5146	2.0001	0.0479	1.0702	0.2869	2.3040	0.0231	2.8960	0.0045
	crag	-1.9654	0.0495	16.6869	<0.0001	4.7051	<0.0001	6.7274	<0.0001	3.5850	<0.0001	13.1578	<0.0001	13.6052	<0.0001	4.8841	<0.0001	14.8626	<0.0001	15.0817	<0.0001
srm-hoover-net > hoover-net	glas	-0.8114	0.4174	0.8109	0.4177	4.9793	<0.0001	2.3828	0.0174	-2.2134	0.0272	-6.5097	<0.0001	-2.7318	0.0065	2.1846	0.0292	-1.8074	0.0711	-0.2850	0.7750
	dpath	3.3095	<0.0001	20.8255	<0.0001	-12.8211	<0.0001	5.0426	<0.0001	5.3716	<0.0001	15.6021	<0.0001	19.8131	<0.0001	28.4857	<0.0001	32.5700	<0.0001	10.2096	<0.0001
	cons	0.2614	0.7946	0.8104	0.4207	-0.2047	0.8384	0.4504	0.6540	2.8279	0.0063	1.4373	0.1556	2.0012	0.0497	1.1079	0.2721	2.2837	0.0258	12.7128	<0.0001
	pann	1.2746	0.2051	3.5341	0.0006	1.9661	0.0518	-1.1428	0.2556	2.7019	0.0080	1.3939	0.1661	1.3300	0.1862	1.7367	0.0852	2.0335	0.0444	2.8870	0.0047
	crag	1.1424	0.2534	14.2186	<0.0001	-2.3907	0.0169	9.5351	<0.0001	2.2800	0.0227	12.9627	<0.0001	14.1342	<0.0001	2.7589	0.0058	14.5832	<0.0001	1.9801	0.0521

Table 11: Paired samples t-test for the HoVer-Net experiments. Values in bold indicate significant differences (p-value < 0.05). The initials “neu”, “epi”, “lym”, “pls”, “eos”, and “con” indicate mPQ^+ for each nucleus type.

FPN, and Mask-RCNN-SE-ResNet50-SE-FPN, and at 5×10^{-9} for the Mask-RCNN-ResNet50-FPN-SRM and 5×10^{-8} for Mask-RCNN-SRM-ResNet50-FPN). The learning rate was decreased by a factor of 0.1 if the training loss did not improve for 10 epochs. Table 7 presents a summary of the hyperparameter details for each experiment. A simple data augmentation policy of random horizontal flips with a probability of 0.5 was considered. To deal with data imbalance in the classification task, we resorted to the focal loss Lin et al. (2020), weighted by the intra-tile effective number of samples (Cui et al., 2019) (equation 15) which promotes the cross-entropy loss to give a higher penalty for harder examples and for the instances that are least represented in a tile.

$$FL(p_t) = -\frac{1-\beta}{1-\beta^n} (1-p_t)^\gamma \log(p_t) \quad (15)$$

With t the target class, n the number of intra-tile instances of class t , β and γ hyperparameters, and p_t the posterior probability, $p(t|X)$. In our experiments, we set $\beta = 0.9$ and $\gamma = 0.9$. For validation, we selected the model saved at the last epoch.

Regarding the HoVer-Net experiments, we replicate the baseline from the official code base provided for the CoNIC Challenge (Graham et al., 2024), only extending the method with the SE-blocks and SRM as previously described. We train the model for 50 epochs, and observe the best results with batch sizes of 2 and 6 for, respectively, HoVer-Net/SE-HoVer-Net and SRM-HoVer-Net, with Adaptive Moment Estimation (ADAM) as optimizer and a learning rate of 1×10^{-4} . Besides, the learn-

ing rate is decreased by a factor of 0.1 every 25 epochs. As is done with the Mask-RCNN experiments, we replace the standard cross-entropy loss with the focal loss. Furthermore, in SRM-HoVer-Net we use 8 attention heads and dropout probability of 0.2.

The source code is available at <https://github.com/joao-nunes/context-attention-nuclei-instance-segmentation.git>.

7.3. Results

Tables 8, 10, and 12 present the results of the models evaluated in this work. In turn, tables 9 and 11 present the results of a paired samples t-test. The experiments for Mask-RCNN with context and attention mechanisms at the level of the ResNet-50 backbone always demonstrated significant differences (p-value < 0.0001).

7.4. Discussion

Pathologists rely on contextual information at multiple levels when analysing WSIs. For instance, they could consider the relationship between tissues and cells (Ryu et al., 2023), the co-occurrence of the same nuclei types (Hassan et al., 2022), or the correlation of tissue subregions with high-diagnostic value (Shao et al., 2021). Therefore, extending ANNs with context- and attention-based inductive biases could translate to more discriminative features. In this work, we extend Mask-RCNN and

Results: IID validation. Data source: CRAG Graham et al. (2019a)											
maps ₅₀ bbox	maps ₅₀ segm	neu	epi	lym	pls	eos	con	pq	pq det	pq seg	multi pq+
mask-rcnn-resnet50-fpn											
0.781 ± 0.090	0.804 ± 0.077	0.256 ± 0.057	0.561 ± 0.021	0.551 ± 0.032	0.423 ± 0.068	0.367 ± 0.072	0.549 ± 0.038	0.557 ± 0.028	0.713 ± 0.034	0.734 ± 0.034	0.451 ± 0.026
mask-rcnn-resnet50-se-fpn											
0.806 ± 0.074	0.825 ± 0.084	0.230 ± 0.047	0.559 ± 0.025	0.544 ± 0.045	0.437 ± 0.058	0.352 ± 0.081	0.564 ± 0.019	0.555 ± 0.031	0.711 ± 0.038	0.733 ± 0.033	0.448 ± 0.024

Table 12: Results of the Mask-RCNN-ResNet50-FPN, and Mask-RCNN-ResNet50-SE-FPN when trained and evaluated on the same data source. The initials “neu”, “epi”, “lym”, “pls”, “eos”, and “con” indicate mPQ^+ for each nucleus type.

HoVer-Net with context and attention-based mechanisms to assess if it could improve the performance and generalization of the models. However, while most authors suggest extending ConvNets with context and attention-based mechanisms improves segmentation performance, our findings are inconclusive. While in the HoVer-Net experiments, we observe the SE-blocks and SRM improve mPQ^+ over the baseline by, respectively, 1.83 % and 2.93 %, the results in Mask-RCNN are not stable. Indeed, this is confirmed by the inconsistency of the significance tests.

This is contrary to our intuition that H&E-stained WSIs are context-rich and that incorporating context and attention-based inductive bias into algorithm design would enrich the learned representations. Furthermore, the surveyed papers demonstrate performance improvement in instance segmentation and classification with context and attention-based mechanisms (Vahadane et al., 2021; Chen et al., 2022d; Dogar et al., 2023), which is not consistent with our observations. As a matter of fact, the results suggest that depending on how context and attention modules are integrated into the neural network architecture, these could even deteriorate performance. Nonetheless, the performance degradation could be due to overfitting, as we could end with an overparameterized network relative to dataset size.

Our findings thus suggest extending state-of-the-art models with simple context and attention-based mechanisms is not sufficient, and that more tailored solutions are required. Incorporating domain knowledge in algorithm design is thus no trivial task, which requires extensive experimentation, or even explicit context annotations. However, the lack of datasets with annotations referring to nuclei context (e.g., tissues, cells, etc.) hinders our ability to explicitly model contextual information. Indeed, while there is a recent effort towards annotating the context of cells (i.e. tissues (hierarchical/global context)) for improving cell detection algorithms (Ryu et al., 2023), we find there is no similar approach for nuclei instance segmentation and classification, besides that other levels of explicit contextual information could also be explored (e.g., long-range context). On the other hand, there is the hypothesis that when the data semantics is context-rich, traditional ConvNets could already learn contextual representations (e.g., through the FPNs), meaning extending these works with explicit context mechanisms (e.g., transformer layers) would not provide additional information. This hypothesis should be addressed in future work to help clarify these yet unanswered topics while contributing to improving algorithm performance and trustworthiness. It would provide some new insights regarding the inner workings of ANNs that could guide future developments towards more precise and reliable ML in CPath.

8. Conclusion

We have reviewed context and attention-based neural networks in medical image analysis, with a special focus on cell nuclei instance segmentation and classification. Despite frequent claims that these mechanisms increase model performance, we found most works regard semantic segmentation, whereas the performance is usually lower in instance segmentation and classification. State-of-the-art models are frequently based on encoder-decoder neural networks, even in transformer neural networks, as these often correspond to hybrids with ConvNet encoder and transformer decoder, or with transformers cross-attention incorporated in the ConvNet layers. Moreover, we find that the challenge of learning with limited annotations is being addressed in CPath, where several works suggest exploring domain knowledge in the form of context and attention mechanisms to tackle the gap when compared with the fully supervised setting. Overall, context and attention-based mechanisms are promising approaches to solve many of the challenges of cell nuclei instance segmentation and classification from H&E-stained brightfield microscopy images like mitigating the effects of spurious contexts (e.g., low contrast between foreground and background instances) or refine instance segmentation masks of touching and overlapping instances. On the other hand, our comparative study of context and attention mechanisms does not support the hypothesis that simple context and attention-based models could increase the performance and generalization of nuclei instance segmentation and classification tasks, meaning more tailored solutions would be required. How to optimally exploit this domain knowledge thus remains an open question, but there is also the hypothesis that state-of-the-art ConvNets (e.g. Mask-RCNN (He et al., 2017)) naturally learn contextual features from the training data. Moreover, while there is work demonstrating ConvNets can achieve the same results as models with an intrinsic context and attention design, namely ViT,(Liu et al., 2022c) a more detailed analysis could allow us to more objectively address this topic, especially in cell nuclei instance segmentation and classification. Although pathologists rely on multiscale context while paying attention to specific RoIs when analysing and annotating WSIs, our findings suggest translating that domain knowledge into ANNs requires careful design and extensive experimentation, but to exploit these mechanisms fully, the scientific understanding of these methods should first be addressed.

Acknowledgments

The authors would like to thank Isabel Rio Torto, Pedro Costa and Professor Aurélio Campilho for their contributions

to the experimental part of this work, for their ideas and critical perspective.

This work is financed by National Funds through the Portuguese funding agency, FCT - Fundação para a Ciência e a Tecnologia, within the PhD grant 2022.12385.BD.

List of Acronyms

ACF	Adaptive Context Fusion	IEN	Instance Encoding Network
AFPN	Attention Feature Pyramid Network	IHC	Immunohistochemical
AI	Artificial Intelligence	IID	independent and identically distributed
AJI	Aggregated Jaccard Index	IoU	Intersection Over Union
ANN	Artificial Neural Network	IRNet	Instance Relation Network
AUC	Area Under the Curve	IoU	Intersection over Union
CA-Net	Comprehensive Attention Convolutional Neural Network	KD	Knowledge Distillation
CBAM	Convolutional Block Attention Module	LSTM	Long Short-term Memory
CEM	Context Enhancement Module	mAP	Mean Average Precision
CL	Curriculum Learning	MIL	Multiple Instance Learning
CPFNet	Context Pyramid Fusion Network	MLP-Head	Multi-layer Perceptron Head
CPath	computational pathology	MLP	Multilayer Perceptron
CRC	Colorectal Cancer	ML	Machine Learning
CT	Computerized Tomography	MPN	Message Passing Network
CV	Computer Vision	MPN	Message Passing Network
CWM	Confidence-based Weighting Module	MRI	Magnetic Resonance Imaging
ConvNet	Convolutional Neural Network	MSDU	Multi-scale Dense Units
DA	Domain Adaptation	NC	Nuclear Classification
DCA	Detail Context Attention	NF	Nuclear Foreground
DC-Net	Dual Context Network	NLP	Natural Language Processing
DC	Dice Similarity Coefficient	NMS	Non-Maximum Suppression
DDD	Decreasing Distance Discretization	NO	Nuclear Ordinal Regression
DETR	Detection Transformer	NP	Nuclear Prediction
DG	Domain Generalization	NT	Nuclear Type
d_H	directed Hausdorff Distance	OOD	out of distribution
DL	Deep Learning	PQ	Panoptic Quality
DTA	Dynamic Texture Attention	RCCA	Recurrent Criss-cross Attention
DeiT	Data Efficient Image Transformer	RPN	Region Proposal Network
FACA	Foreground Aware Co-Attention	ReLU	Rectified Linear Unit
FPN	Feature Pyramid Network	ResGANet	Residual Group Attention Network
FTN	Fully Transformer Network	RoI	Region of Interest
GAN	Generative Adversarial Network	SAM	Segment Anything
GNN	Graph Neural Network	SCM	Structural Causal Model
HD	Hausdorff Distance	SE	Squeeze and Excitation
H&E	Haematoxylin and Eosin	SL	Strongly Supervised Learning
HoVer	Horizontal-Vertical	SPN	Semantic Proposal Network
IAM	Information Aggregation Module	SRM	Semantic Relation Module
		SSL	Self-supervised Learning
		SmSL	Semi-supervised Learning
		TIL	Tumour-Infiltrating Lymphocyte
		VAE	Variational Autoencoder
		ViT	Vision Transformer
		WSI	Whole Slide Image
		WSL	Weakly Supervised Learning

References

- Aatresh, A.A., Alabhya, K., Lal, S., Kini, J., Saxena, P.P., 2021. Liver-net: efficient and robust deep learning model for automatic diagnosis of sub-types of liver hepatocellular carcinoma cancer from h&e stained liver histopathology images. *International Journal of Computer Assisted Radiology and Surgery* 16, 1549–1563. URL: <https://doi.org/10.1007/s11548-021-02410-4>, doi:10.1007/s11548-021-02410-4.
- Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Süsstrunk, S., 2010. Slic superpixels, 15URL: <http://infoscience.epfl.ch/record/149300>.
- Acs, B., Rantalainen, M., Hartman, J., 2020. Artificial intelligence as the next step towards precision pathology. *Journal of Internal Medicine* 288, 62–81. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/joim.13030>, doi:https://doi.org/10.1111/joim.13030, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/joim.13030
- Adel, T., Ghahramani, Z., Weller, A., 2018. Discovering interpretable representations for both deep generative and discriminative models, in: Dy, J., Krause, A. (Eds.), *Proceedings of the 35th International Conference on Machine Learning*, PMLR. pp. 50–59. URL: <https://proceedings.mlr.press/v80/adel18a.html>.
- Ahmad, I., Xia, Y., Cui, H., Islam, Z., 2023. Dan-nucnet: A dual attention based framework for nuclei segmentation in cancer histology images under wild clinical conditions. *Expert Systems with Applications* 213. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85139854050&doi=10.1016%2Fj.eswa.2022.118945&partnerID=40&md5=8522c564a33bd8b272c764728e9c9e44>, doi:10.1016/j.eswa.2022.118945.
- Ahmedt-Aristizabal, D., Armin, M.A., Denman, S., Fookes, C., Petersson, L., 2022. A survey on graph-based deep learning for computational histopathology. *Computerized Medical Imaging and Graphics* 95, 102027. URL: <https://www.sciencedirect.com/science/article/pii/S0895611121001762>, doi:https://doi.org/10.1016/j.compedimag.2021.102027.
- Alemi Koohbanani, N., Jahanifar, M., Gooya, A., Rajpoot, N., 2019. Nuclear instance segmentation using a proposal-free spatially aware deep learning framework, in: Shen, D., Liu, T., Peters, T.M., Staib, L.H., Esert, C., Zhou, S., Yap, P.T., Khan, A. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, Springer International Publishing, Cham. pp. 622–630. doi:https://doi.org/10.1007/978-3-030-32239-7_69.
- Ali, H., Haq, I., Cui, L., Feng, J., 2022. Msal-net: improve accurate segmentation of nuclei in histopathology images by multiscale attention learning network. *BMC Medical Informatics and Decision Making* 22, 90. URL: <https://doi.org/10.1186/s12911-022-01826-5>, doi:10.1186/s12911-022-01826-5.
- An, L., Wang, L., Li, Y., 2022. Hea-net: Attention and mlp hybrid encoder architecture for medical image segmentation. *Sensors* 22. URL: <https://www.mdpi.com/1424-8220/22/18/7024>, doi:10.3390/s22187024.
- Arjovsky, M., Bottou, L., Gulrajani, I., Lopez-Paz, D., 2019. Invariant risk minimization. arXiv preprint arXiv:1907.02893.
- Asilian Bidgoli, A., Rahnamayan, S., Dehkharghanian, T., Grami, A., Tizhoosh, H.R., 2022. Bias reduction in representation of histopathology images using deep feature selection. *Scientific Reports* 12, 19994. URL: <https://doi.org/10.1038/s41598-022-24317-z>, doi:10.1038/s41598-022-24317-z.
- Aung, T.N., Shafi, S., Wilmott, J.S., Nourmohammadi, S., Vathiotis, I., Gavrielatou, N., Fernandez, A., Yaghoobi, V., Sinnberg, T., Amaral, T., Ikenberg, K., Khosrotehrani, K., Osman, I., Acs, B., Bai, Y., Martinez-Morilla, S., Moutafi, M., Thompson, J.F., Scolyer, R.A., Rimm, D.L., 2022. Objective assessment of tumor infiltrating lymphocytes as a prognostic marker in melanoma using machine learning algorithms. *eBioMedicine* 82, 104143. URL: <https://www.sciencedirect.com/science/article/pii/S2352396422003243>, doi:https://doi.org/10.1016/j.ebiom.2022.104143.
- Azizi, S., Mustafa, B., Ryan, F., Beaver, Z., Freyberg, J., Deaton, J., Loh, A., Karthikesalingam, A., Kornblith, S., Chen, T., Natarajan, V., Norouzi, M., 2021. Big self-supervised models advance medical image classification, in: *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 3458–3468. doi:10.1109/ICCV48922.2021.00346.
- Azzuni, H., Ridzuan, M., Xu, M., Yaqub, M., 2022. Color space-based hover-net for nuclei instance segmentation and classification, in: *2022 IEEE International Symposium on Biomedical Imaging Challenges (ISBIC)*, pp. 1–4. doi:10.1109/ISBIC56247.2022.9854725.
- Bai, Y., Mi, Y., Su, Y., Zhang, B., Zhang, Z., Wu, J., Huang, H., Xiong, Y., Gong, X., Wang, W., 2022. A scalable graph-based framework for multi-organ histology image classification. *IEEE Journal of Biomedical and Health Informatics* 26, 5506–5517. doi:10.1109/JBHI.2022.3199110.
- Bardes, A., Ponce, J., LeCun, Y., 2021. Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *CoRR abs/2105.04906*. URL: <https://arxiv.org/abs/2105.04906>, arXiv:2105.04906.
- Bashir, R.M.S., Qaiser, T., Raza, S.E.A., Rajpoot, N.M., 2024. Consistency regularisation in varying contexts and feature perturbations for semi-supervised semantic segmentation of histology images. *Medical Image Analysis* 91, 102997. URL: <https://www.sciencedirect.com/science/article/pii/S1361841523002578>, doi:https://doi.org/10.1016/j.media.2023.102997.
- Basu, A., Senapati, P., Deb, M., Rai, R., Dhal, K.G., 2023. A survey on recent trends in deep learning for nucleus segmentation from histopathology images. *Evolving Systems* URL: <https://doi.org/10.1007/s12530-023-09491-3>, doi:10.1007/s12530-023-09491-3.
- Baxi, V., Edwards, R., Montalto, M., Saha, S., 2022. Digital pathology and artificial intelligence in translational medicine and clinical practice. *Modern Pathology* 35, 23–32. URL: <https://doi.org/10.1038/s41379-021-00919-2>, doi:10.1038/s41379-021-00919-2.
- Boserup, N., Selvan, R., 2023. Efficient self-supervision using patch-based contrastive learning for histopathology image segmentation. *Proceedings of the Northern Lights Deep Learning Workshop 4*. URL: <https://doi.org/10.7557%2F18.6798>, doi:10.7557/18.6798.
- Bronstein, M.M., Bruna, J., Cohen, T., Velicković, P., 2021. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. arXiv preprint arXiv:2104.13478.
- Böhland, M., Tharun, L., Scherr, T., Mikut, R., Hagenmeyer, V., Thompson, L.D.R., Perner, S., Reischl, M., 2021. Machine learning methods for automated classification of tumors with papillary thyroid carcinoma-like nuclei: A quantitative analysis. *PLOS ONE* 16, 1–21. URL: <https://doi.org/10.1371/journal.pone.0257635>, doi:10.1371/journal.pone.0257635.
- Caicedo, J.C., Goodman, A., Karhohs, K.W., Cimini, B.A., Ackerman, J., Haghghi, M., Heng, C., Becker, T., Doan, M., McQuin, C., Rohban, M., Singh, S., Carpenter, A.E., 2019. Nucleus segmentation across imaging experiments: the 2018 data science bowl. *Nature Methods* 16, 1247–1253. URL: <https://doi.org/10.1038/s41592-019-0612-7>, doi:10.1038/s41592-019-0612-7.
- Campanella, G., Hanna, M.G., Geneslaw, L., Mirafior, A., Werneck Krauss Silva, V., Busam, K.J., Brogi, E., Reuter, V.E., Klimstra, D.S., Fuchs, T.J., 2019. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature Medicine* 25, 1301–1309. URL: <https://doi.org/10.1038/s41591-019-0508-1>, doi:10.1038/s41591-019-0508-1.
- Campanella, G., Kwan, R., Fluder, E., Zeng, J., Stock, A., Veremis, B., Polydorides, A.D., Hedvat, C., Schoenfeld, A., Vanderbilt, C., et al., 2023. Computational pathology at health system scale—self-supervised foundation models from three billion images. arXiv preprint arXiv:2310.07033.
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S., 2020. End-to-end object detection with transformers, in: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (Eds.), *Computer Vision – ECCV 2020*, Springer International Publishing, Cham. pp. 213–229. doi:https://doi.org/10.1007/978-3-030-58452-8_13.
- Chanchal, A.K., Kumar, A., Lal, S., Kini, J., 2021. Efficient and robust deep learning architecture for segmentation of kidney and breast histopathology images. *Computers & Electrical Engineering* 92, 107177. URL: <https://www.sciencedirect.com/science/article/pii/S0045790621001786>, doi:https://doi.org/10.1016/j.compeleceng.2021.107177.
- Chen, H., Gomez, C., Huang, C.M., Unberath, M., 2022a. Explainable medical imaging ai needs human-centered design: guidelines and evidence from a systematic review. *npj Digital Medicine* 5, 156. URL: <https://doi.org/10.1038/s41746-022-00699-2>, doi:10.1038/s41746-022-00699-2.
- Chen, L., Bentley, P., Mori, K., Misawa, K., Fujiwara, M., Rueckert, D., 2019a. Self-supervised learning for medical image analysis using image context restoration. *Medical Image Analysis* 58, 101539. URL: <https://www.sciencedirect.com/science/article/pii/S1361841518304699>, doi:https://doi.org/10.1016/j.media.2019.101539.
- Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H., 2018. Encoder-

- decoder with atrous separable convolution for semantic image segmentation, in: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (Eds.), *Computer Vision – ECCV 2018*, Springer International Publishing, Cham. pp. 833–851. doi:https://doi.org/10.1007/978-3-030-01234-2_49.
- Chen, R.J., Lu, M.Y., Weng, W.H., Chen, T.Y., Williamson, D.F., Manz, T., Shady, M., Mahmood, F., 2021. Multimodal co-attention transformer for survival prediction in gigapixel whole slide images, in: 2021 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 3995–4005. doi:10.1109/ICCV48922.2021.00398.
- Chen, R.J., Lu, M.Y., Williamson, D.F., Chen, T.Y., Lipkova, J., Noor, Z., Shaban, M., Shady, M., Williams, M., Joo, B., Mahmood, F., 2022b. Pan-cancer integrative histology-genomic analysis via multimodal deep learning. *Cancer Cell* 40, 865–878.e6. URL: <https://www.sciencedirect.com/science/article/pii/S1535610822003178>, doi:<https://doi.org/10.1016/j.ccell.2022.07.004>.
- Chen, S., Bortsova, G., García-Uceda Juárez, A., van Tulder, G., de Bruijne, M., 2019b. Multi-task attention-based semi-supervised learning for medical image segmentation, in: Shen, D., Liu, T., Peters, T.M., Staib, L.H., Essert, C., Zhou, S., Yap, P.T., Khan, A. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, Springer International Publishing, Cham. pp. 457–465. doi:https://doi.org/10.1007/978-3-030-32248-9_51.
- Chen, S., Ding, C., Liu, M., Cheng, J., Tao, D., 2023. Cpp-net: Context-aware polygon proposal network for nucleus segmentation. *IEEE Transactions on Image Processing* 32, 980–994. doi:10.1109/TIP.2023.3237013.
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G., 2020a. A simple framework for contrastive learning of visual representations, in: *Proceedings of the 37th International Conference on Machine Learning*, JMLR.org.
- Chen, X., Fan, H., Girshick, R., He, K., 2020b. Improved baselines with momentum contrastive learning. arXiv preprint arXiv:2003.04297.
- Chen, X., He, K., 2021. Exploring simple siamese representation learning, in: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 15745–15753. doi:10.1109/CVPR46437.2021.01549.
- Chen, Y., Feng, J., Liu, J., Pang, B., Cao, D., Li, C., 2022c. Detection and classification of lung cancer cells using swin transformer. *Journal of Cancer Therapy* 13, 464–475. doi:10.4236/jct.2022.137041.
- Chen, Y., Jia, Y., Zhang, X., Bai, J., Li, X., Ma, M., Sun, Z., Pei, Z., 2022d. Tshvnet: Simultaneous nuclear instance segmentation and classification in histopathological images based on multiattention mechanisms. *BioMed Research International* 2022. doi:<https://doi.org/10.1155/2022/7921922>.
- Chen, Z., Tian, Z., Zhu, J., Li, C., Du, S., 2022e. C-cam: Causal cam for weakly supervised semantic segmentation on medical image, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11676–11685.
- Cheng, B., Girshick, R., Dollár, P., Berg, A.C., Kirillov, A., 2021a. Boundary iou: Improving object-centric image segmentation evaluation, in: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 15329–15337. doi:10.1109/CVPR46437.2021.01508.
- Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R., 2022. Masked-attention mask transformer for universal image segmentation, in: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1280–1289. doi:10.1109/CVPR52688.2022.00135.
- Cheng, B., Schwing, A.G., Kirillov, A., 2021b. Per-pixel classification is not all you need for semantic segmentation. CoRR abs/2107.06278. URL: <https://arxiv.org/abs/2107.06278>, arXiv:2107.06278.
- Cheng, Y., Qiao, P., He, H., Song, G., Chen, J., 2021c. Hard-boundary attention network for nuclei instance segmentation, Association for Computing Machinery. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85123051772&doi=10.1145%2f3469877.3490602&partnerID=40&md5=4b9af9406067d7d61cbba26fd6cb2a5c>, doi:10.1145/3469877.3490602.
- Chhipa, P.C., Holmgren, J.R., De, K., Saini, R., Liwicki, M., 2023. Can self-supervised representation learning methods withstand distribution shifts and corruptions?, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pp. 4467–4476.
- Chow, K.H., Factor, R.E., Ullman, K.S., 2012. The nuclear envelope environment and its cancer connections. *Nature Reviews Cancer* 12, 196–209.
- Chu, X., Tian, Z., Zhang, B., Wang, X., Shen, C., 2021. Conditional positional encodings for vision transformers, in: *International Conference on Learning Representations*.
- Ciga, O., Xu, T., Martel, A.L., 2022a. Self supervised contrastive learning for digital histopathology. *Machine Learning with Applications* 7, 100198. URL: <https://www.sciencedirect.com/science/article/pii/S2666827021000992>, doi:<https://doi.org/10.1016/j.mlwa.2021.100198>.
- Ciga, O., Xu, T., Martel, A.L., 2022b. Self supervised contrastive learning for digital histopathology. *Machine Learning with Applications* 7, 100198. URL: <https://www.sciencedirect.com/science/article/pii/S2666827021000992>, doi:<https://doi.org/10.1016/j.mlwa.2021.100198>.
- Clark, K., Vendt, B., Smith, K., Freymann, J., Kirby, J., Koppel, P., Moore, S., Phillips, S., Maffitt, D., Pringle, M., Tarbox, L., Prior, F., 2013. The cancer imaging archive (TCIA): maintaining and operating a public information repository. *J. Digit. Imaging* 26, 1045–1057. doi:10.1007/s10278-013-9622-7.
- Cubuk, E.D., Zoph, B., Shlens, J., Le, Q., 2020. Randaugment: Practical automated data augmentation with a reduced search space, in: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (Eds.), *Advances in Neural Information Processing Systems*, Curran Associates, Inc., pp. 18613–18624. URL: <https://proceedings.neurips.cc/paper/2020/file/d85b63ef0ccb114d0a3bb7b7d808028f-Paper.pdf>.
- Cui, M., Zhang, D.Y., 2021. Artificial intelligence and computational pathology. *Laboratory Investigation* 101, 412–422. URL: <https://doi.org/10.1038/s41374-020-00514-0>, doi:10.1038/s41374-020-00514-0.
- Cui, Y., Jia, M., Lin, T.Y., Song, Y., Belongie, S., 2019. Class-balanced loss based on effective number of samples, in: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 9260–9269. doi:10.1109/CVPR.2019.00949.
- Da, Q., Huang, X., Li, Z., Zuo, Y., Zhang, C., Liu, J., Chen, W., Li, J., Xu, D., Hu, Z., Yi, H., Guo, Y., Wang, Z., Chen, L., Zhang, L., He, X., Zhang, X., Mei, K., Zhu, C., Lu, W., Shen, L., Shi, J., Li, J., S. S., Krishnamurthi, G., Yang, J., Lin, T., Song, Q., Liu, X., Graham, S., Bashir, R.M.S., Yang, C., Qin, S., Tian, X., Yin, B., Zhao, J., Metaxas, D.N., Li, H., Wang, C., Zhang, S., 2022. Digestpath: A benchmark dataset with challenge review for the pathological detection and segmentation of digestive-system. *Medical Image Analysis* 80, 102485. URL: <https://www.sciencedirect.com/science/article/pii/S1361841522001323>, doi:<https://doi.org/10.1016/j.media.2022.102485>.
- Delahunt, B., Eble, J.N., Egevad, L., Samaratunga, H., 2019. Grading of renal cell carcinoma. *Histopathology* 74, 4–17.
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database, in: 2009 IEEE Conference on Computer Vision and Pattern Recognition, pp. 248–255. doi:10.1109/CVPR.2009.5206848.
- Deng, K., Meng, Y., Gao, D., Bridge, J., Shen, Y., Lip, G., Zhao, Y., Zheng, Y., 2021. Transbridge: A lightweight transformer for left ventricle segmentation in echocardiography, in: Noble, J.A., Aylward, S., Grimwood, A., Min, Z., Lee, S.L., Hu, Y. (Eds.), *Simplifying Medical Ultrasound*, Springer International Publishing, Cham. pp. 63–72. doi:https://doi.org/10.1007/978-3-030-87583-1_7.
- Deng, R., Cui, C., Liu, Q., Yao, T., Remedios, L.W., Bao, S., Landman, B.A., Wheless, L.E., Coburn, L.A., Wilson, K.T., et al., 2023. Segment anything model (sam) for digital pathology: Assess zero-shot segmentation on whole slide imaging. arXiv preprint arXiv:2304.04155.
- Deshmukh, G., Susladkar, O., Makwana, D., R. S.C.T., S. N.K., Mittal, S., 2022. Feednet: a feature enhanced encoder-decoder lstm network for nuclei instance segmentation for histopathological diagnosis. *Physics in Medicine & Biology* 67, 195011. URL: <https://dx.doi.org/10.1088/1361-6560/ac8594>, doi:10.1088/1361-6560/ac8594.
- Diba, A., Sharma, V., Stiefelhagen, R., Van Gool, L., 2019. Weakly supervised object discovery by generative adversarial & ranking networks, in: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 601–610. doi:10.1109/CVPRW.2019.00086.
- Do, M., Vetterli, M., 2005. The contourlet transform: an efficient directional multiresolution image representation. *IEEE Transactions on Image Processing* 14, 2091–2106. doi:10.1109/TIP.2005.859376.
- Doan, T.N.N., Song, B., Vuong, T.T.L., Kim, K., Kwak, J.T., 2022. Sonnet: A self-guided ordinal regression neural network for segmentation and classification of nuclei in large-scale multi-tissue histology images. *IEEE Journal of Biomedical and Health Informatics* 26, 3218–3228. doi:10.1109/JBHI.2022.3149936.
- Dodington, D.W., Lagree, A., Tabbarah, S., Mohebbpour, M., Sadeghi-Naini, A., Tran, W.T., Lu, F.I., 2021. Analysis of tumor nuclear features using arti-

- cial intelligence to predict response to neoadjuvant chemotherapy in high-risk breast cancer patients. *Breast Cancer Research and Treatment* 186, 379–389. URL: <https://doi.org/10.1007/s10549-020-06093-4>, doi:10.1007/s10549-020-06093-4.
- Dogar, G.M., Shahzad, M., Fraz, M.M., 2023. Attention augmented distance regression and classification network for nuclei instance segmentation and type classification in histology images. *Biomedical Signal Processing and Control* 79, 104199. URL: <https://www.sciencedirect.com/science/article/pii/S174680942200653X>, doi:<https://doi.org/10.1016/j.bspc.2022.104199>.
- Dosovitskiy, A., Beyler, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houshy, N., 2021. An image is worth 16x16 words: Transformers for image recognition at scale, in: *International Conference on Learning Representations*. URL: <https://openreview.net/forum?id=YicbFdNTTy>.
- Dreyer, M., Achibat, R., Wiegand, T., Samek, W., Lapuschkin, S., 2023. Revealing hidden context bias in segmentation and object detection through concept-specific explanations, in: *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 3829–3839. doi:10.1109/CVPRW59228.2023.00397.
- Dubey, S., Olimov, F., Rafique, M.A., Jeon, M., 2022. Improving small objects detection using transformer. *Journal of Visual Communication and Image Representation* 89, 103620. URL: <https://www.sciencedirect.com/science/article/pii/S1047320322001432>, doi:<https://doi.org/10.1016/j.jvcir.2022.103620>.
- Dvornik, N., Mairal, J., Schmid, C., 2021. On the importance of visual context for data augmentation in scene understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43, 2014–2028. doi:10.1109/TPAMI.2019.2961896.
- England, J.R., Cheng, P.M., 2019. Artificial intelligence for medical image analysis: A guide for authors and reviewers. *American Journal of Roentgenology* 212, 513–519. URL: <https://doi.org/10.2214/AJR.18.20490>, doi:10.2214/AJR.18.20490, arXiv:<https://doi.org/10.2214/AJR.18.20490>.
- Erhan, D., Bengio, Y., Courville, A., Manzagol, P.A., Vincent, P., Bengio, S., 2010. Why does unsupervised pre-training help deep learning? *Journal of Machine Learning Research* 11, 625–660. URL: <http://jmlr.org/papers/v11/erhan10a.html>.
- Ermolov, A., Siarohin, A., Sangineto, E., Sebe, N., 2021. Whitening for self-supervised representation learning, in: *International Conference on Machine Learning*, PMLR, pp. 3015–3024.
- Esteva, A., Chou, K., Yeung, S., Naik, N., Madani, A., Mottaghi, A., Liu, Y., Topol, E., Dean, J., Socher, R., 2021. Deep learning-enabled medical computer vision. *npj Digital Medicine* 4, 5. URL: <https://doi.org/10.1038/s41746-020-00376-2>, doi:10.1038/s41746-020-00376-2.
- Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A., 2010. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision* 88, 303–338. URL: <https://doi.org/10.1007/s11263-009-0275-4>, doi:10.1007/s11263-009-0275-4.
- Fang, Y., Liao, B., Wang, X., Fang, J., Qi, J., Wu, R., Niu, J., Liu, W., 2021. You only look at one sequence: Rethinking transformer in vision through object detection, in: *Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (Eds.), Advances in Neural Information Processing Systems*, Curran Associates, Inc., pp. 26183–26197. URL: https://proceedings.neurips.cc/paper_files/paper/2021/file/dc912a253d1e9ba40e2c597ed2376640-Paper.pdf.
- Fang, Z., Liu, Y., Wang, Y., Zhang, X., Chen, Y., Cai, C., Lin, Y., Han, Y., Wang, Z., Zeng, S., et al., 2023. Deep learning predicts biomarker status and discovers related histomorphology characteristics for low-grade glioma. arXiv preprint arXiv:2310.07464.
- Feng, S., Zhao, H., Shi, F., Cheng, X., Wang, M., Ma, Y., Xiang, D., Zhu, W., Chen, X., 2020. Cpfnet: Context pyramid fusion network for medical image segmentation. *IEEE Transactions on Medical Imaging* 39, 3008–3018. doi:10.1109/TMI.2020.2983721.
- Filiot, A., Ghermi, R., Olivier, A., Jacob, P., Fidon, L., Mac Kain, A., Saillard, C., Schiratti, J.B., 2023. Scaling self-supervised learning for histopathology with masked image modeling. medRxiv, 2023–07.
- Fischer, E.G., 2020. Nuclear morphology and the biology of cancer cells. *Acta cytologica* 64, 511–519.
- Foucart, A., Debeir, O., Decaestecker, C., 2023. Panoptic quality should be avoided as a metric for assessing cell nuclei segmentation and classification in digital pathology. *Scientific Reports* 13, 8614. URL: <https://doi.org/10.1038/s41598-023-35605-7>, doi:10.1038/s41598-023-35605-7.
- Gamper, J., Koohbanani, N.A., Benes, K., Graham, S., Jahanifar, M., Khurram, S.A., Azam, A., Hewitt, K., Rajpoot, N., 2020a. Pannuke dataset extension, insights and baselines. arXiv preprint arXiv:2003.10778.
- Gamper, J., Koohbanani, N.A., Rajpoot, N., 2020b. Multi-task learning in histo-pathology for widely generalizable model. arXiv preprint arXiv:2005.08645.
- Geng, R., Huang, Z., Chen, J., 2021. Adaptive feature aggregation network for nuclei segmentation, Association for Computing Machinery, Inc. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85105877498&doi=10.1145/2f3444685.3446271&partnerID=40&md5=44c270ade684ce9bd3080c2734371d4c>, doi:10.1145/3444685.3446271.
- Ghaffari Laleh, N., Muti, H.S., Loeffler, C.M.L., Echle, A., Saldanha, O.L., Mahmood, F., Lu, M.Y., Trautwein, C., Langer, R., Dislich, B., Buelow, R.D., Grabsch, H.I., Brenner, H., Chang-Claude, J., Alwers, E., Brinker, T.J., Khader, F., Truhn, D., Gaisa, N.T., Boor, P., Hoffmeister, M., Schulz, V., Kather, J.N., 2022. Benchmarking weakly-supervised deep learning pipelines for whole slide classification in computational pathology. *Medical Image Analysis* 79, 102474. URL: <https://www.sciencedirect.com/science/article/pii/S1361841522001219>, doi:<https://doi.org/10.1016/j.media.2022.102474>.
- Ghosal, S.S., Li, Y., 2023. Are vision transformers robust to spurious correlations? *International Journal of Computer Vision*, 1–21.
- Gidariss, S., Bursuc, A., Puy, G., Komodakis, N., Cord, M., Pérez, P., 2021. Obow: Online bag-of-visual-words generation for self-supervised learning, in: *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6826–6836. doi:10.1109/CVPR46437.2021.00676.
- Gonçalves, T., Rio-Torto, I., Teixeira, L.F., Cardoso, J.S., 2022. A survey on attention mechanisms for medical applications: are we moving toward better algorithms? *IEEE Access* 10, 98909–98935. doi:10.1109/ACCESS.2022.3206449.
- Goodfellow, I., Bengio, Y., Courville, A., 2016. *Deep learning*. MIT press.
- Goyal, A., Bengio, Y., 2022. Inductive biases for deep learning of higher-level cognition. *Proceedings of the Royal Society A* 478, 20210068.
- Graham, S., Chen, H., Gamper, J., Dou, Q., Heng, P.A., Snead, D., Tsang, Y.W., Rajpoot, N., 2019a. Mild-net: Minimal information loss dilated network for gland instance segmentation in colon histology images. *Medical Image Analysis* 52, 199–211. URL: <https://www.sciencedirect.com/science/article/pii/S1361841518306030>, doi:<https://doi.org/10.1016/j.media.2018.12.001>.
- Graham, S., Jahanifar, M., Azam, A., Nimir, M., Tsang, Y.W., Dodd, K., Hero, E., Sahota, H., Tank, A., Benes, K., Wahab, N., Minhas, F., Raza, S.E.A., El Daly, H., Gopalakrishnan, K., Snead, D., Rajpoot, N., 2021a. Lizard: A large-scale dataset for colonic nuclear instance segmentation and classification, in: *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pp. 684–693. doi:10.1109/ICCVW54120.2021.00082.
- Graham, S., Jahanifar, M., Vu, Q.D., Hadjigeorgiou, G., Leech, T., Snead, D., Raza, S.E.A., Minhas, F., Rajpoot, N., 2021b. Conic: Colon nuclei identification and counting challenge 2022. arXiv preprint arXiv:2111.14485.
- Graham, S., Vu, Q.D., Jahanifar, M., Weigert, M., Schmidt, U., Zhang, W., Zhang, J., Yang, S., Xiang, J., Wang, X., Rumberger, J.L., Baumann, E., Hirsch, P., Liu, L., Hong, C., Aviles-Rivero, A.I., Jain, A., Ahn, H., Hong, Y., Azzuni, H., Xu, M., Yaqub, M., Blache, M.C., Piégu, B., Vernay, B., Scherr, T., Böhlend, M., Löffler, K., Li, J., Ying, W., Wang, C., Snead, D., Raza, S.E.A., Minhas, F., Rajpoot, N.M., 2024. Conic challenge: Pushing the frontiers of nuclear detection, segmentation, classification and counting. *Medical Image Analysis* 92, 103047. URL: <https://www.sciencedirect.com/science/article/pii/S1361841523003079>, doi:<https://doi.org/10.1016/j.media.2023.103047>.
- Graham, S., Vu, Q.D., Raza, S.E.A., Azam, A., Tsang, Y.W., Kwak, J.T., Rajpoot, N., 2019b. Hover-net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images. *Medical Image Analysis* 58, 101563. URL: <https://www.sciencedirect.com/science/article/pii/S1361841519301045>, doi:<https://doi.org/10.1016/j.media.2019.101563>.
- Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., Piot, B., kavukcuoglu, k., Munos, R., Valko, M., 2020. Bootstrap your own latent - a

- new approach to self-supervised learning, in: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (Eds.), *Advances in Neural Information Processing Systems*, Curran Associates, Inc., pp. 21271–21284. URL: https://proceedings.neurips.cc/paper_files/paper/2020/file/f3ada80d5c4ee70142b17b8192b2958e-Paper.pdf.
- Gu, R., Wang, G., Song, T., Huang, R., Aertsen, M., Deprest, J., Ourselin, S., Vercauteren, T., Zhang, S., 2021a. Ca-net: Comprehensive attention convolutional neural networks for explainable medical image segmentation. *IEEE Transactions on Medical Imaging* 40, 699–711. doi:10.1109/TMI.2020.3035253.
- Gu, R., Zhang, J., Huang, R., Lei, W., Wang, G., Zhang, S., 2021b. Domain composition and attention for unseen-domain generalizable medical image segmentation, in: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, Springer International Publishing, Cham, pp. 241–250. doi:https://doi.org/10.1007/978-3-030-87199-4_23.
- Gu, W., Bai, S., Kong, L., 2022. A review on 2d instance segmentation based on deep neural networks. *Image and Vision Computing* 120, 104401. URL: <https://www.sciencedirect.com/science/article/pii/S0262885622000300>, doi:<https://doi.org/10.1016/j.imavis.2022.104401>.
- Guo, J., Han, K., Wu, H., Tang, Y., Chen, X., Wang, Y., Xu, C., 2022a. Cmt: Convolutional neural networks meet vision transformers, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12175–12185.
- Guo, M.H., Xu, T.X., Liu, J.J., Liu, Z.N., Jiang, P.T., Mu, T.J., Zhang, S.H., Martin, R.R., Cheng, M.M., Hu, S.M., 2022b. Attention mechanisms in computer vision: A survey. *Computational Visual Media* 8, 331–368. URL: <https://doi.org/10.1007/s41095-022-0271-y>, doi:10.1007/s41095-022-0271-y.
- Guo, R., Pagnucco, M., Song, Y., 2021. Learning with noise: Mask-guided attention model for weakly supervised nuclei segmentation, in: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, Springer International Publishing, Cham, pp. 461–470. doi:https://doi.org/10.1007/978-3-030-87196-3_43.
- Gurcan, M.N., Boucheron, L.E., Can, A., Madabhushi, A., Rajpoot, N.M., Yener, B., 2009. Histopathological image analysis: A review. *IEEE Reviews in Biomedical Engineering* 2, 147–171. doi:10.1109/RBME.2009.2034865.
- Hamidieh, K., Zhang, H., Ghassemi, M., 2022. Evaluating and improving robustness of self-supervised representations to spurious correlations, in: *ICML 2022: Workshop on Spurious Correlations, Invariance and Stability*.
- Hamidinekoo, A., Pieciak, T., Afzali, M., Akanyeti, O., Yuan, Y., 2021. Glioma classification using multimodal radiology and histology data, in: Crimi, A., Bakas, S. (Eds.), *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, Springer International Publishing, Cham, pp. 508–518.
- Han, K., Wang, Y., Chen, H., Chen, X., Guo, J., Liu, Z., Tang, Y., Xiao, A., Xu, C., Xu, Y., Yang, Z., Zhang, Y., Tao, D., 2023. A survey on vision transformer. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 87–110. doi:10.1109/TPAMI.2022.3152247.
- Haq, M.M., Huang, J., 2022. Self-supervised pre-training for nuclei segmentation, in: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, Springer Nature Switzerland, Cham, pp. 303–313. doi:https://doi.org/10.1007/978-3-031-16434-7_30.
- Hassan, T., Javed, S., Mahmood, A., Qaiser, T., Werghi, N., Rajpoot, N., 2022. Nucleus classification in histology images using message passing network. *Medical Image Analysis* 79, 102480. URL: <https://www.sciencedirect.com/science/article/pii/S136184152200127X>, doi:<https://doi.org/10.1016/j.media.2022.102480>.
- Hastie, T., Tibshirani, R., Friedman, J., Hastie, T., Tibshirani, R., Friedman, J., 2009. Overview of supervised learning. *The elements of statistical learning: Data mining, inference, and prediction*, 9–41.
- Hayakawa, T., Prasath, V.B.S., Kawanaka, H., Aronow, B.J., Tsuruoka, S., 2021. Computational nuclei segmentation methods in digital pathology: A survey. *Archives of Computational Methods in Engineering* 28, 1–13. URL: <https://doi.org/10.1007/s11831-019-09366-4>, doi:10.1007/s11831-019-09366-4.
- He, H., Huang, Z., Ding, Y., Song, G., Wang, L., Ren, Q., Wei, P., Gao, Z., Chen, J., 2021a. Cdnnet: Centripetal direction network for nuclear instance segmentation, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4026–4035.
- He, H., Zhang, C., Chen, J., Geng, R., Chen, L., Liang, Y., Lu, Y., Wu, J., Xu, Y., 2021b. A hybrid-attention nested unet for nuclear segmentation in histopathological images. *Frontiers in Molecular Biosciences* 8. URL: <https://www.frontiersin.org/articles/10.3389/fmolb.2021.614174>, doi:10.3389/fmolb.2021.614174.
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R., 2020. Momentum contrast for unsupervised visual representation learning, in: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9726–9735. doi:10.1109/CVPR42600.2020.00975.
- He, K., Gan, C., Li, Z., Reikik, I., Yin, Z., Ji, W., Gao, Y., Wang, Q., Zhang, J., Shen, D., 2022a. Transformers in medical image analysis: A review. *Intelligent Medicine* URL: <https://www.sciencedirect.com/science/article/pii/S2667102622000717>, doi:<https://doi.org/10.1016/j.imed.2022.07.002>.
- He, K., Gkioxari, G., Dollár, P., Girshick, R., 2017. Mask r-cnn, in: *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 2980–2988. doi:10.1109/ICCV.2017.322.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778. doi:10.1109/CVPR.2016.90.
- He, X., Tan, E.L., Bi, H., Zhang, X., Zhao, S., Lei, B., 2022b. Fully transformer network for skin lesion analysis. *Medical Image Analysis* 77, 102357. URL: <https://www.sciencedirect.com/science/article/pii/S136184152200010X>, doi:<https://doi.org/10.1016/j.media.2022.102357>.
- He, Z., Unberath, M., Ke, J., Shen, Y., 2023. Transnuseg: A lightweight multi-task transformer for nuclei segmentation, in: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*, Springer Nature Switzerland, Cham, pp. 206–215.
- Hendry, S., Salgado, R., Gevaert, T., Russell, P.A., John, T., Thapa, B., Christie, M., Van De Vijver, K., Estrada, M.V., Gonzalez-Ericsson, P.I., et al., 2017. Assessing tumor infiltrating lymphocytes in solid tumors: a practical review for pathologists and proposal for a standardized method from the international immuno-oncology biomarkers working group: Part 2: Tils in melanoma, gastrointestinal tract carcinomas, non-small cell lung carcinoma and mesothelioma, endometrial and ovarian carcinomas, squamous cell carcinoma of the head and neck, genitourinary carcinomas, and primary brain tumors. *Advances in anatomic pathology* 24, 235–2510. doi:10.1097/PAP.000000000000161.
- Hendrycks, D., Gimpel, K., 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Hendrycks, D., Mazeika, M., Kadavath, S., Song, D., 2019. Using self-supervised learning can improve model robustness and uncertainty, in: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (Eds.), *Advances in Neural Information Processing Systems*, Curran Associates, Inc. URL: https://proceedings.neurips.cc/paper_files/paper/2019/file/a2b15837edac15df90721968986f7f8e-Paper.pdf.
- Hjelm, R.D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., Bengio, Y., 2018. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*.
- Ho, J., Kalchbrenner, N., Weissenborn, D., Salimans, T., 2019. Axial attention in multidimensional transformers. *arXiv preprint arXiv:1912.12180*.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural computation* 9, 1735–1780.
- Hu, H., Gu, J., Zhang, Z., Dai, J., Wei, Y., 2018. Relation networks for object detection, in: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3588–3597. doi:10.1109/CVPR.2018.00378.
- Huang, W., Jiang, M., Li, M., Meng, B., Ren, J., Zhao, S., Bai, R., Yang, Y., 2021. Causal intervention for object detection, in: *2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (ICTAI)*, pp. 770–774. doi:10.1109/ICTAI52525.2021.00122.
- Huang, Y., Yang, X., Liu, L., Zhou, H., Chang, A., Zhou, X., Chen, R., Yu, J., Chen, J., Chen, C., Liu, S., Chi, H., Hu, X., Yue, K., Li, L., Grau, V., Fan, D.P., Dong, F., Ni, D., 2024. Segment anything model for medical images? *Medical Image Analysis* 92, 103061. URL: <https://www.sciencedirect.com/science/article/pii/S1361841523003213>, doi:<https://doi.org/10.1016/>

- j.media.2023.103061.
- Huang, Z., Bianchi, F., Yuksekgonul, M., Montine, T.J., Zou, J., 2023. A visual-language foundation model for pathology image analysis using medical twitter. *Nature Medicine* 29, 2307–2316. URL: <https://doi.org/10.1038/s41591-023-02504-3>, doi:10.1038/s41591-023-02504-3.
- Huang, Z., Wang, X., Huang, L., Huang, C., Wei, Y., Liu, W., 2019. Cnet: Criss-cross attention for semantic segmentation, in: 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 603–612. doi:10.1109/ICCV.2019.00069.
- Ibrahim, A., Gamble, P., Jaroensri, R., Abdelsamea, M.M., Mermel, C.H., Chen, P.H.C., Rakha, E.A., 2020. Artificial intelligence in digital breast pathology: Techniques and applications. *The Breast* 49, 267–273. URL: <https://www.sciencedirect.com/science/article/pii/S0960977619312147>, doi:<https://doi.org/10.1016/j.breast.2019.12.007>.
- Imran, A.A.Z., Terzopoulos, D., 2019. Semi-supervised multi-task learning with chest x-ray images, in: Suk, H.I., Liu, M., Yan, P., Lian, C. (Eds.), *Machine Learning in Medical Imaging*, Springer International Publishing, Cham. pp. 151–159. doi:https://doi.org/10.1007/978-3-030-32692-0_18.
- Intiaz, T., Fattah, S.A., Saquib, M., 2023. Condanet: Contourlet driven attention network for automatic nuclei segmentation in histopathology images. *IEEE Access* 11, 129321–129330. doi:10.1109/ACCESS.2023.3321799.
- Jack, N.P., Thomas, W., Marick, L., Fabien, R., 2018. Segmentation of Nuclei in Histopathology Images by deep regression of the distance map. URL: <https://doi.org/10.5281/zenodo.1175282>, doi:10.5281/zenodo.1175282.
- Janowczyk, A., Zuo, R., Gilmore, H., Feldman, M., Madabhushi, A., 2019. Histoqc: An open-source quality control tool for digital pathology slides. *JCO Clinical Cancer Informatics* , 1–7 URL: <https://doi.org/10.1200/CCI.18.00157>, doi:10.1200/CCI.18.00157, arXiv:<https://doi.org/10.1200/CCI.18.00157>.
- Jaroensri, R., Wulczyn, E., Hegde, N., Brown, T., Flament-Auvigne, I., Tan, F., Cai, Y., Nagpal, K., Rakha, E.A., Dabbs, D.J., Olson, N., Wren, J.H., Thompson, E.E., Seetao, E., Robinson, C., Miao, M., Beckers, F., Corrado, G.S., Peng, L.H., Mermel, C.H., Liu, Y., Steiner, D.F., Chen, P.H.C., 2022. Deep learning models for histologic grading of breast cancer and association with disease prognosis. *npj Breast Cancer* 8, 113. URL: <https://doi.org/10.1038/s41523-022-00478-y>, doi:10.1038/s41523-022-00478-y.
- Jaume, G., Pati, P., Bozorgtabar, B., Foncubierta, A., Anniciello, A.M., Feroce, F., Rau, T., Thiran, J.P., Gabrani, M., Goksel, O., 2021. Quantifying explainers of graph neural networks in computational pathology, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8106–8116.
- Jieneng, C., Yongyi, L., Qihang, Y., Xiangde, L., Ehsan, A., Yan, W., Le, L., L., Y.A., Yuyin, Z., 2021. Transunet: Transformers make strong encoders for medical image segmentation. arXiv preprint arXiv:2102.04306 .
- Jung, S., Heo, H., Park, S., Jung, S.U., Lee, K., 2022. Benchmarking deep learning models for instance segmentation. *Applied Sciences* 12. URL: <https://www.mdpi.com/2076-3417/12/17/8856>, doi:10.3390/app12178856.
- Kakumani, A.K., Sree, L.P., Krishna, C.S., Uppalapati, G., Pavithra, G.S.S., Harshini, S., 2022. Semantic segmentation of cells in microscopy images via pretrained autoencoder and attention u-net, in: 2022 International Conference on Machine Learning, Computer Systems and Security (MLCSS), pp. 94–99. doi:10.1109/MLCSS57186.2022.00025.
- Kang, M., Song, H., Park, S., Yoo, D., Pereira, S., 2023. Benchmarking self-supervised learning on diverse pathology datasets, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3344–3354.
- Karri, M., Annavarapu, C.S.R., Acharya, U.R., 2022. Explainable multi-module semantic guided attention based network for medical image segmentation. *Computers in Biology and Medicine* 151, 106231. URL: <https://www.sciencedirect.com/science/article/pii/S0010482522009398>, doi:<https://doi.org/10.1016/j.combiomed.2022.106231>.
- Kather, J.N., Heij, L.R., Grabsch, H.I., Loeffler, C., Echle, A., Muti, H.S., Krause, J., Niehues, J.M., Sommer, K.A.J., Bankhead, P., Kooreman, L.F.S., Schulte, J.J., Cipriani, N.A., Buelow, R.D., Boor, P., Ortiz-Brüchle, N., Hanby, A.M., Speirs, V., Kochanny, S., Patnaik, A., Srisuwananukorn, A., Brenner, H., Hoffmeister, M., van den Brandt, P.A., Jäger, D., Trautwein, C., Pearson, A.T., Luedde, T., 2020. Pan-cancer image-based detection of clinically actionable genetic alterations. *Nature Cancer* 1, 789–799. URL: <https://doi.org/10.1038/s43018-020-0087-6>, doi:10.1038/s43018-020-0087-6.
- Kaul, C., Pears, N., Dai, H., Murray-Smith, R., Manandhar, S., 2021. Focusnet++: Attentive aggregated transformations for efficient and accurate medical image segmentation, in: 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI), pp. 1042–1046. doi:10.1109/ISBI48211.2021.9433918.
- Ke, J., Lu, Y., Shen, Y., Zhu, J., Zhou, Y., Huang, J., Yao, J., Liang, X., Guo, Y., Wei, Z., Liu, S., Huang, Q., Jiang, F., Shen, D., 2023. Clusterseg: A crowd cluster pinpointed nucleus segmentation framework with cross-modality datasets. *Medical Image Analysis* 85, 102758. URL: <https://www.sciencedirect.com/science/article/pii/S1361841523000191>, doi:<https://doi.org/10.1016/j.media.2023.102758>.
- Khan, S., Naseer, M., Hayat, M., Zamir, S.W., Khan, F.S., Shah, M., 2022. Transformers in vision: A survey. *ACM Comput. Surv.* 54. URL: <https://doi.org/10.1145/3505244>, doi:10.1145/3505244.
- Kirillov, A., He, K., Girshick, R., Rother, C., Dollár, P., 2019. Panoptic segmentation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9404–9413.
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al., 2023. Segment anything. arXiv preprint arXiv:2304.02643 .
- Kong, Y., Genchev, G.Z., Wang, X., Zhao, H., Lu, H., 2020. Nuclear segmentation in histopathological images using two-stage stacked u-nets with attention mechanism. *Frontiers in Bioengineering and Biotechnology* 8. URL: <https://www.frontiersin.org/articles/10.3389/fbioe.2020.573866>, doi:10.3389/fbioe.2020.573866.
- Koohbanani, N.A., Unnikrishnan, B., Khurram, S.A., Krishnaswamy, P., Rajpoot, N., 2021. Self-path: Self-supervision for classification of pathology images with limited annotations. *IEEE Transactions on Medical Imaging* 40, 2845–2856. doi:10.1109/TMI.2021.3056023.
- Krithiga, R., Geetha, P., 2021. Breast cancer detection, segmentation and classification on histopathology images analysis: A systematic review. *Archives of Computational Methods in Engineering* 28, 2607–2619. URL: <https://doi.org/10.1007/s11831-020-09470-w>, doi:10.1007/s11831-020-09470-w.
- Krug, C., Rohr, K., 2022. Unsupervised cell segmentation in fluorescence microscopy images via self-supervised learning, in: El Yacoubi, M., Granger, E., Yuen, P.C., Pal, U., Vincent, N. (Eds.), *Pattern Recognition and Artificial Intelligence*, Springer International Publishing, Cham. pp. 236–247. doi:https://doi.org/10.1007/978-3-031-09037-0_20.
- Kuhn, H.W., 1955. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly* 2, 83–97. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/nav.3800020109>, doi:<https://doi.org/10.1002/nav.3800020109>, arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/nav.3800020109>.
- Kumar, N., Verma, R., Anand, D., Zhou, Y., Onder, O.F., Tsougenis, E., Chen, H., et al., 2020. A multi-organ nucleus segmentation challenge. *IEEE Transactions on Medical Imaging* 39, 1380–1391. doi:10.1109/TMI.2019.2947628.
- Lagrec, A., Mohebpour, M., Meti, N., Saednia, K., Lu, F.I., Slodkowska, E., Gandhi, S., Rakovitch, E., Shenfield, A., Sadeghi-Naini, A., Tran, W.T., 2021. A review and comparison of breast tumor cell nuclei segmentation performances using deep convolutional neural networks. *Scientific Reports* 11, 8025. URL: <https://doi.org/10.1038/s41598-021-87496-1>, doi:10.1038/s41598-021-87496-1.
- Lai, Z., Li, Z., Oliveira, L.C., Chauhan, J., Dugger, B.N., Chuah, C.N., 2023. Clipath: Fine-tune clip with visual feature fusion for pathology image analysis towards minimizing data collection efforts, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pp. 2374–2380.
- Lal, S., Das, D., Alabhya, K., Kanfode, A., Kumar, A., Kini, J., 2021. Nucleisegnet: Robust deep learning architecture for the nuclei segmentation of liver cancer histopathology images. *Computers in Biology and Medicine* 128. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85095915871&doi=10.1016/j.combiomed.2020.104075&partnerID=40&md5=9eac78a7ce3c120e83c1894ea3d2ed5e>, doi:10.1016/j.combiomed.2020.104075.
- Le, N., Yamazaki, K., Quach, K.G., Truong, D., Savvides, M., 2021. A multi-

- task contextual atrous residual network for brain tumor detection & segmentation, in: 2020 25th International Conference on Pattern Recognition (ICPR), pp. 5943–5950. doi:10.1109/ICPR48806.2021.9412414.
- Lei, B., Xia, Z., Jiang, F., Jiang, X., Ge, Z., Xu, Y., Qin, J., Chen, S., Wang, T., Wang, S., 2020. Skin lesion segmentation via generative adversarial networks with dual discriminators. *Medical Image Analysis* 64, 101716. URL: <https://www.sciencedirect.com/science/article/pii/S1361841520300803>, doi:<https://doi.org/10.1016/j.media.2020.101716>.
- Lejeune, M., Gestí, V., Tomás, B., Korzyńska, A., Roso, A., Callau, C., Bosch, R., Baucells, J., Jaén, J., López, C., 2013. A multistep image analysis method to increase automated identification efficiency in immunohistochemical nuclear markers with a high background level. *Diagnostic Pathology* 8, S13. URL: <https://doi.org/10.1186/1746-1596-8-S1-S13>, doi:10.1186/1746-1596-8-S1-S13.
- Lekadir, K., Osuala, R., Gallin, C.S., Lazrak, N., Kushibar, K., Tsakou, G., Auss'ou, S., Alberich, L.C., Marias, K., Tskinakis, M., Colantonio, S., Papanikolaou, N., Salahuddin, Z., Woodruff, H., Lambin, P., Mart'i-Bonmat'i, L., 2021. Future-ai: Guiding principles and consensus recommendations for trustworthy artificial intelligence in medical imaging. *ArXiv abs/2109.09658*.
- Li, C., Liu, D., Li, H., Zhang, Z., Lu, G., Chang, X., Cai, W., 2022. Domain adaptive nuclei instance segmentation and classification via category-aware feature alignment and pseudo-labelling, in: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, Springer Nature Switzerland, Cham, pp. 715–724.
- Li, C., Wang, X., Liu, W., Latecki, L.J., Wang, B., Huang, J., 2019. Weakly supervised mitosis detection in breast histopathology images using concentric loss. *Medical Image Analysis* 53, 165–178. URL: <https://www.sciencedirect.com/science/article/pii/S1361841519300118>, doi:<https://doi.org/10.1016/j.media.2019.01.013>.
- Li, D., Yang, J., Kreis, K., Torralba, A., Fidler, S., 2021a. Semantic segmentation with generative models: Semi-supervised learning and strong out-of-domain generalization, in: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 8296–8307. doi:10.1109/CVPR46437.2021.00820.
- Li, M., Fan, H., Yang, D., 2021b. Cab-net: Channel attention block network for pathological image cell nucleus segmentation, in: Peng, Y., Hu, S.M., Gabbouj, M., Zhou, K., Elad, M., Xu, K. (Eds.), *Image and Graphics*, Springer International Publishing, Cham, pp. 654–663. doi:https://doi.org/10.1007/978-3-030-87358-5_53.
- Li, X., Pi, J., Lou, M., Qi, Y., Li, S., Meng, J., Ma, Y., 2023. Multi-level feature fusion network for nuclei segmentation in digital histopathological images. *The Visual Computer* 39, 1307–1322. URL: <https://doi.org/10.1007/s00371-022-02407-3>, doi:10.1007/s00371-022-02407-3.
- Li, Z., Li, Y., Li, Q., Wang, P., Guo, D., Lu, L., Jin, D., Zhang, Y., Hong, Q., 2024. Lvit: Language meets vision transformer in medical image segmentation. *IEEE Transactions on Medical Imaging* 43, 96–107. doi:10.1109/TMI.2023.3291719.
- Lin, T., Wang, Y., Liu, X., Qiu, X., 2022a. A survey of transformers. *AI Open* 3, 111–132. URL: <https://www.sciencedirect.com/science/article/pii/S2666651022000146>, doi:<https://doi.org/10.1016/j.aiopen.2022.10.001>.
- Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S., 2017. Feature pyramid networks for object detection, in: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 936–944. doi:10.1109/CVPR.2017.106.
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P., 2020. Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42, 318–327. doi:10.1109/TPAMI.2018.2858826.
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L., 2014. Microsoft coco: Common objects in context, in: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (Eds.), *Computer Vision – ECCV 2014*, Springer International Publishing, Cham, pp. 740–755. doi:https://doi.org/10.1007/978-3-319-10602-1_48.
- Lin, Y., Wang, Z., Cheng, K.T., Chen, H., 2022b. Insmix: Towards realistic generative data augmentation for nuclei instance segmentation, in: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, Springer Nature Switzerland, Cham, pp. 140–149. doi:https://doi.org/10.1007/978-3-031-16434-7_14.
- Liu, L., Muellly, M., Deng, J., Pfister, T., Li, L.J., 2019a. Generative modeling for small-data object detection, in: 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 6072–6080. doi:10.1109/ICCV.2019.00617.
- Liu, R., Liu, H., Li, G., Hou, H., Yu, T., Yang, T., 2022a. Contextual debiasing for visual recognition with causal mechanisms, in: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 12745–12755. doi:10.1109/CVPR52688.2022.01242.
- Liu, S., Wang, Y., Yang, X., Lei, B., Liu, L., Li, S.X., Ni, D., Wang, T., 2019b. Deep learning in medical ultrasound analysis: A review. *Engineering* 5, 261–275. URL: <https://www.sciencedirect.com/science/article/pii/S2095809918301887>, doi:<https://doi.org/10.1016/j.eng.2018.11.020>.
- Liu, W., He, Q., He, X., 2022b. Weakly supervised nuclei segmentation via instance learning, in: 2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI), pp. 1–5. doi:10.1109/ISBI52829.2022.9761644.
- Liu, X., Guo, Z., Cao, J., Tang, J., 2021a. Mdc-net: A new convolutional neural network for nucleus segmentation in histopathology images with distance maps and contour information. *Computers in Biology and Medicine* 135, 104543. URL: <https://www.sciencedirect.com/science/article/pii/S0010482521003371>, doi:<https://doi.org/10.1016/j.compbiomed.2021.104543>.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B., 2021b. Swin transformer: Hierarchical vision transformer using shifted windows, in: 2021 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 9992–10002. doi:10.1109/ICCV48922.2021.00986.
- Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S., 2022c. A convnet for the 2020s, in: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 11966–11976. doi:10.1109/CVPR52688.2022.01167.
- Liu, Z., Wang, H., Zhang, S., Wang, G., Qi, J., 2020. Nas-scam: Neural architecture search-based spatial and channel joint attention module for nuclei semantic segmentation and classification. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 12261 LNCS, 263–272. URL: https://www.scopus.com/inward/record.uri?eid=2-s2.0-85093101118&doi=10.1007/978-3-030-59710-8_26&partnerID=40&md5=aa8eb9b13b7f18ddcc8265c1ef14421b, doi:10.1007/978-3-030-59710-8_26.
- Loshchilov, I., Hutter, F., 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Lu, C., Bera, K., Wang, X., Prasanna, P., Xu, J., Janowczyk, A., Beig, N., Yang, M., Fu, P., Lewis, J., Choi, H., Schmid, R.A., Berezowska, S., Schalper, K., Rimm, D., Velcheti, V., Madabhushi, A., 2020a. A prognostic model for overall survival of patients with early-stage non-small cell lung cancer: a multicentre, retrospective study. *The Lancet Digital Health* 2, e594–e606. URL: <https://www.sciencedirect.com/science/article/pii/S2589750020302259>, doi:[https://doi.org/10.1016/S2589-7500\(20\)30225-9](https://doi.org/10.1016/S2589-7500(20)30225-9).
- Lu, M.Y., Chen, B., Williamson, D.F., Chen, R.J., Ikamura, K., Gerber, G., Liang, I., Le, L.P., Ding, T., Parwani, A.V., et al., 2023. A foundational multimodal vision language ai assistant for human pathology. *arXiv preprint arXiv:2312.07814*.
- Lu, W., Graham, S., Bilal, M., Rajpoot, N., Minhas, F., 2020b. Capturing cellular topology in multi-gigapixel pathology images, in: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 1049–1058. doi:10.1109/CVPRW50498.2020.00138.
- Lundervold, A.S., Lundervold, A., 2019. An overview of deep learning in medical imaging focusing on mri. *Zeitschrift für Medizinische Physik* 29, 102–127. URL: <https://www.sciencedirect.com/science/article/pii/S0939388918301181>, doi:<https://doi.org/10.1016/j.zemedi.2018.11.002>.
- Luo, H., Changdong, Y., Selvan, R., 2022. Hybrid ladder transformers with efficient parallel-cross attention for medical image segmentation, in: Konukoglu, E., Menze, B., Venkataraman, A., Baumgartner, C., Dou, Q., Albarqouni, S. (Eds.), *Proceedings of The 5th International Conference on Medical Imaging with Deep Learning*, PMLR, pp. 808–819. URL: <https://proceedings.mlr.press/v172/luo22a.html>.
- M., G., V., A., S., M.M., H., M., 2020. Concept attribution: Explaining cnn decisions to physicians. *Computers in Biology and Medicine* 123, 103865. URL: <https://www.sciencedirect.com/science/>

- article/pii/S0010482520302225, doi:<https://doi.org/10.1016/j.compbio.2020.103865>.
- Maass, W., 1997. Networks of spiking neurons: the third generation of neural network models. *Neural networks* 10, 1659–1671.
- Mahanta, L.B., Hussain, E., Das, N., Kakoti, L., Chowdhury, M., 2021. Ihc-net: A fully convolutional neural network for automated nuclear segmentation and ensemble classification for allred scoring in breast pathology. *Applied Soft Computing* 103, 107136. URL: <https://www.sciencedirect.com/science/article/pii/S1568494621000594>, doi:<https://doi.org/10.1016/j.asoc.2021.107136>.
- Mahapatra, D., Poellinger, A., Reyes, M., 2022. Interpretability-guided inductive bias for deep learning based medical image. *Medical Image Analysis* 81, 102551. URL: <https://www.sciencedirect.com/science/article/pii/S1361841522001980>, doi:<https://doi.org/10.1016/j.media.2022.102551>.
- Majurski, M., Manescu, P., Padi, S., Schaub, N., Hotaling, N., Simon, C., Bajcsy, P., 2019. Cell image segmentation using generative adversarial networks, transfer learning, and augmentations, in: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 1114–1122. doi:10.1109/CVPRW.2019.00145.
- Marini, N., Marchesin, S., Otálora, S., Wodzinski, M., Caputo, A., van Rijthoven, M., Aswolinskiy, W., Bokhorst, J.M., Podareanu, D., Petters, E., Boytcheva, S., Buttafuoco, G., Vatrano, S., Frassetto, F., van der Laak, J., Agosti, M., Ciompi, F., Silvello, G., Muller, H., Atzori, M., 2022. Unleashing the potential of digital pathology data by training computer-aided diagnosis models without human annotations. *npj Digital Medicine* 5, 102. URL: <https://doi.org/10.1038/s41746-022-00635-4>, doi:10.1038/s41746-022-00635-4.
- Mehta, S., Rastegari, M., 2021. Mobilevit: light-weight, general-purpose, and mobile-friendly vision transformer. arXiv preprint arXiv:2110.02178.
- Min, K., Lee, G.H., Lee, S.W., 2022. Attentional feature pyramid network for small object detection. *Neural Networks* 155, 439–450. URL: <https://www.sciencedirect.com/science/article/pii/S089360802200329X>, doi:<https://doi.org/10.1016/j.neunet.2022.08.029>.
- Misra, I., van der Maaten, L., 2020. Self-supervised learning of pretext-invariant representations, in: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 6706–6716. doi:10.1109/CVPR42600.2020.00674.
- Mongan, J., Moy, L., Kahn, C.E., 2020. Checklist for artificial intelligence in medical imaging (claim): A guide for authors and reviewers. *Radiology: Artificial Intelligence* 2, e200029. URL: <https://doi.org/10.1148/ryai.2020200029>, doi:10.1148/ryai.2020200029, arXiv:<https://doi.org/10.1148/ryai.2020200029>.
- Montezuma, D., Oliveira, S.P., Neto, P.C., Oliveira, D., Monteiro, A., Cardoso, J.S., Macedo-Pinto, I., 2023. Annotating for artificial intelligence applications in digital pathology: A practical guide for pathologists and researchers. *Modern Pathology*, 100086URL: <https://www.sciencedirect.com/science/article/pii/S0893395222055260>, doi:<https://doi.org/10.1016/j.modpat.2022.100086>.
- Murtaza Dogar, G., Fraz, M., Javed, S., 2021. Feature attention network for simultaneous nuclei instance segmentation and classification in histology images, Institute of Electrical and Electronics Engineers Inc. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85107670840&doi=10.1109%2FICoDT252288.2021.9441474&partnerID=40&md5=465c65c8c9c304f4861036ae250d82ab>, doi:10.1109/ICoDT252288.2021.9441474.
- Myronenko, A., 2019. 3d mri brain tumor segmentation using autoencoder regularization, in: Crimi, A., Bakas, S., Kuijff, H., Keyvan, F., Reyes, M., van Walsum, T. (Eds.), *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, Springer International Publishing, Cham, pp. 311–320. doi:https://doi.org/10.1007/978-3-030-11726-9_28.
- Naseem, U., Khushi, M., Kim, J., 2023. Vision-language transformer for interpretable pathology visual question answering. *IEEE Journal of Biomedical and Health Informatics* 27, 1681–1690. doi:10.1109/JBHI.2022.3163751.
- Naseer, M.M., Ranasinghe, K., Khan, S.H., Hayat, M., Shahbaz Khan, F., Yang, M.H., 2021. Intriguing properties of vision transformers, in: Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (Eds.), *Advances in Neural Information Processing Systems*, Curran Associates, Inc., pp. 23296–23308. URL: https://proceedings.neurips.cc/paper_files/paper/2021/file/c404a5adbf90e09631678b13b05d9d7a-Paper.pdf.
- Nasir, E.S., Parvaiz, A., Fraz, M.M., 2022. Nuclei and glands instance segmentation in histology images: a narrative review. *Artificial Intelligence Review* URL: <https://doi.org/10.1007/s10462-022-10372-5>, doi:10.1007/s10462-022-10372-5.
- Nema, S., Dudhane, A., Murala, S., Naidu, S., 2020. Rescuenet: An unpaired gan for brain tumor segmentation. *Biomedical Signal Processing and Control* 55, 101641. URL: <https://www.sciencedirect.com/science/article/pii/S1746809419302228>, doi:<https://doi.org/10.1016/j.bspc.2019.101641>.
- Neto, P.C., Montezuma, D., Oliveira, S.P., Oliveira, D., Fraga, J., Monteiro, A., Monteiro, J., Ribeiro, L., Gonçalves, S., Reinhard, S., et al., 2023. A cad system for colorectal cancer from wsi: A clinically validated interpretable ml-based prototype. arXiv preprint arXiv:2301.02608.
- Neto, P.C., Oliveira, S.P., Montezuma, D., Fraga, J., Monteiro, A., Ribeiro, L., Gonçalves, S., Pinto, I.M., Cardoso, J.S., 2022. imil4path: A semi-supervised interpretable approach for colorectal whole-slide images. *Cancers* 14. URL: <https://www.mdpi.com/2072-6694/14/10/2489>, doi:10.3390/cancers14102489.
- Niazi, M.K.K., Parwani, A.V., Gurcan, M.N., 2019. Digital pathology and artificial intelligence. *The Lancet Oncology* 20, e253–e261. URL: <https://www.sciencedirect.com/science/article/pii/S1470204519301548>, doi:[https://doi.org/10.1016/S1470-2045\(19\)30154-8](https://doi.org/10.1016/S1470-2045(19)30154-8).
- Northcutt, C., Jiang, L., Chuang, I., 2021. Confident learning: Estimating uncertainty in dataset labels. *J. Artif. Int. Res.* 70, 1373–1411. URL: <https://doi.org/10.1613/jair.1.12125>, doi:10.1613/jair.1.12125.
- Oakden-Rayner, L., Dunnmon, J., Carneiro, G., Re, C., 2020. Hidden stratification causes clinically meaningful failures in machine learning for medical imaging, in: *Proceedings of the ACM Conference on Health, Inference, and Learning*, Association for Computing Machinery, New York, NY, USA, p. 151–159. URL: <https://doi.org/10.1145/3368555.3384468>, doi:10.1145/3368555.3384468.
- Obeid, A., Mahbub, T., Javed, S., Dias, J., Werghe, N., 2022. Nucdetr: End-to-end transformer for nucleus detection in histopathology images, in: Qin, W., Zaki, N., Zhang, F., Wu, J., Yang, F. (Eds.), *Computational Mathematics Modeling in Cancer Analysis*, Springer Nature Switzerland, Cham, pp. 47–57. doi:https://doi.org/10.1007/978-3-031-17266-3_5.
- Ouyang, C., Chen, C., Li, S., Li, Z., Qin, C., Bai, W., Rueckert, D., 2023. Causality-inspired single-source domain generalization for medical image segmentation. *IEEE Transactions on Medical Imaging* 42, 1095–1106. doi:10.1109/TMI.2022.3224067.
- Pan, X., Cheng, J., Hou, F., Lan, R., Lu, C., Li, L., Feng, Z., Wang, H., Liang, C., Liu, Z., Chen, X., Han, C., Liu, Z., 2023. Smile: Cost-sensitive multi-task learning for nuclear segmentation and classification with imbalanced annotations. *Medical Image Analysis* 88, 102867. URL: <https://www.sciencedirect.com/science/article/pii/S1361841523001275>, doi:<https://doi.org/10.1016/j.media.2023.102867>.
- Papers with Code: Cityscapes, . Cityscapes test benchmark (instance segmentation). URL: <https://paperswithcode.com/sota/instance-segmentation-on-cityscapes>, accessed: 01-08-2023.
- Papers With Code: COCO, . Coco test-dev benchmark (instance segmentation). URL: <https://paperswithcode.com/sota/instance-segmentation-on-coco>, accessed: 01-08-2023.
- Pati, P., Jaume, G., Foncubierta-Rodríguez, A., Feroce, F., Annicciello, A.M., Scognamiglio, G., Brancati, N., Fiche, M., Dubruc, E., Riccio, D., Di Bonito, M., De Pietro, G., Botti, G., Thiran, J.P., Frucci, M., Goksel, O., Gabrani, M., 2022. Hierarchical graph representations in digital pathology. *Medical Image Analysis* 75, 102264. URL: <https://www.sciencedirect.com/science/article/pii/S1361841521003091>, doi:<https://doi.org/10.1016/j.media.2021.102264>.
- Paul, S., Chen, P.Y., 2022. Vision transformers are robust learners, in: *Proceedings of the AAAI conference on Artificial Intelligence*, pp. 2071–2081.
- Pearl, J., 2009. Causal inference in statistics: An overview.
- Peters, J., Bühlmann, P., Meinshausen, N., 2016. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 78, 947–1012.
- Peters, J., Janzing, D., Schölkopf, B., 2017. *Elements of causal inference: foundations and learning algorithms*. The MIT Press.
- Pineau, J., Vincent-Lamarre, P., Sinha, K., Larivière, V., Beygelzimer, A.,

- d'Alché Buc, F., Fox, E., Larochelle, H., 2020. Improving reproducibility in machine learning research (a report from the neurips 2019 reproducibility program). *J. Mach. Learn. Res.* 22.
- Pinto, F., Torr, P.H.S., K. Dokania, P., 2022. An impartial take to the cnn vs transformer robustness contest, in: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (Eds.), *Computer Vision – ECCV 2022*, Springer Nature Switzerland, Cham, pp. 466–480. doi:https://doi.org/10.1007/978-3-031-19778-9_27.
- Prangemeier, T., Reich, C., Koeppl, H., 2020. Attention-based transformers for instance segmentation of cells in microstructures, in: *2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 700–707. doi:[10.1109/BIBM49941.2020.9313305](https://doi.org/10.1109/BIBM49941.2020.9313305).
- Priego-Torres, B.M., Lobato-Delgado, B., Atienza-Cuevas, L., Sanchez-Morillo, D., 2022. Deep learning-based instance segmentation for the precise automated quantification of digital breast cancer immunohistochemistry images. *Expert Systems with Applications* 193, 116471. URL: <https://www.sciencedirect.com/science/article/pii/S0957417421017516>, doi:<https://doi.org/10.1016/j.eswa.2021.116471>.
- Qin, J., He, Y., Zhou, Y., Zhao, J., Ding, B., 2022. Reu-net: Region-enhanced nuclei segmentation network. *Computers in Biology and Medicine* 146. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85129478540&doi=10.1016%2Fj.combiomed.2022.105546&partnerID=40&md5=646f80b2044c877527691fc93c1bedd>, doi:[10.1016/j.combiomed.2022.105546](https://doi.org/10.1016/j.combiomed.2022.105546).
- Qu, H., Yan, Z., Riedlinger, G.M., De, S., Metaxas, D.N., 2019. Improving nuclei/gland instance segmentation in histopathology images by full resolution neural network and spatial constrained loss, in: Shen, D., Liu, T., Peters, T.M., Staib, L.H., Essert, C., Zhou, S., Yap, P.T., Khan, A. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, Springer International Publishing, Cham, pp. 378–386. doi:https://doi.org/10.1007/978-3-030-32239-7_42.
- Raghu, M., Unterthiner, T., Kornblith, S., Zhang, C., Dosovitskiy, A., 2021. Do vision transformers see like convolutional neural networks?, in: Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (Eds.), *Advances in Neural Information Processing Systems*, Curran Associates, Inc., pp. 12116–12128. URL: https://proceedings.neurips.cc/paper_files/paper/2021/file/652cf38361a209088302ba2b8b7f51e0-Paper.pdf.
- Rao, A., Park, J., Woo, S., Lee, J.Y., Aalami, O., 2021. Studying the effects of self-attention for medical image analysis, in: *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pp. 3409–3418. doi:[10.1109/ICCVW54120.2021.00381](https://doi.org/10.1109/ICCVW54120.2021.00381).
- Ravin, R., Hoepfner, D.J., Munno, D.M., Carmel, L., Sullivan, J., Levitt, D.L., Miller, J.L., Athaide, C., Panchision, D.M., McKay, R.D., 2008. Potency and fate specification in cns stem cell populations in vitro. *Cell Stem Cell* 3, 670–680. URL: <https://www.sciencedirect.com/science/article/pii/S1934590908004827>, doi:<https://doi.org/10.1016/j.stem.2008.09.012>.
- Raza, S.E.A., Cheung, L., Shaban, M., Graham, S., Epstein, D., Pelengaris, S., Khan, M., Rajpoot, N.M., 2019. Micro-net: A unified model for segmentation of various objects in microscopy images. *Medical Image Analysis* 52, 160–173. URL: <https://www.sciencedirect.com/science/article/pii/S1361841518300628>, doi:<https://doi.org/10.1016/j.media.2018.12.003>.
- Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You only look once: Unified, real-time object detection, in: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 779–788. doi:[10.1109/CVPR.2016.91](https://doi.org/10.1109/CVPR.2016.91).
- Reinke, A., Maier-Hein, L., Müller, H., 2021. Common limitations of performance metrics in biomedical image analysis. *Proceedings of the Medical Imaging with Deep Learning (MIDL 2021)*.
- Reinke, A., Tizabi, M.D., Baumgartner, M., Eisenmann, M., Heckmann-Nötzel, D., Kavur, A.E., Rädtsch, T., Sudre, C.H., Acion, L., Antonelli, M., Arbel, T., Bakas, S., Benis, A., Blaschko, M.B., Buettner, F., Cardoso, M.J., Chelygina, V., Chen, J., Christodoulou, E., Cimini, B.A., Collins, G.S., Farahani, K., Ferrer, L., Galdran, A., VAN Ginneken, B., Glocker, B., Godau, P., Haase, R., Hashimoto, D.A., Hoffman, M.M., Huisman, M., Isensee, F., Jannin, P., Kahn, C.E., Kainmueller, D., Kainz, B., Karargyris, A., Karthikesalingam, A., Kennigott, H., Kleesiek, J., Kofler, F., Kooi, T., Kopp-Schneider, A., Kozubek, M., Kreshuk, A., Kurc, T., Landman, B.A., Litjens, G., Madani, A., Maier-Hein, K., Martel, A.L., Mattson, P., Meijering, E., Menze, B., Moons, K.G.M., Müller, H., Nichyporuk, B., Nickel, F., Petersen, J., Rafelski, S.M., Rajpoot, N., Reyes, M., Riegler, M.A., Rieke, N., Saez-Rodriguez, J., Sánchez, C.I., Shetty, S., Summers, R.M., Taha, A.A., Tlupin, A., Tsaftaris, S.A., VAN Calster, B., Varoquaux, G., Yaniv, Z.R., Jäger, P.F., Maier-Hein, L., 2023. Understanding metric-related pitfalls in image analysis validation.
- Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster r-cnn: Towards real-time object detection with region proposal networks, in: Cortes, C., Lawrence, N., Lee, D., Sugiyama, M., Garnett, R. (Eds.), *Advances in Neural Information Processing Systems*, Curran Associates, Inc. URL: <https://proceedings.neurips.cc/paper/2015/file/14bfa6bb14875e45bba028a21ed38046-Paper.pdf>.
- Rolls, G., Farmer, N., 2008. Artifacts in Histological and Cytological Preparations. *Scientia leica Microsystems education series*, 1–112 URL: http://www.leica-microsystems.com/fileadmin/academy/2011/Artifacts_f_Handbook.pdf.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation, in: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, Springer, pp. 234–241. doi:https://doi.org/10.1007/978-3-319-24574-4_28.
- Ryu, J., Puche, A.V., Shin, J., Park, S., Brattoli, B., Lee, J., Jung, W., Cho, S.I., Paeng, K., Ock, C.Y., et al., 2023. Ocelot: Overlapped cell on tissue dataset for histopathology, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 23902–23912.
- Sahasrabudhe, M., Christodoulidis, S., Salgado, R., Loi, S., André, F., Paragios, N., Vakalopoulou, M., 2020. Self-supervised nuclei segmentation in histopathological images using attention, in: Martel, A.L., Abolmaesumi, P., Stoyanov, D., Mateus, D., Zuluaga, M.A., Zhou, S.K., Racoceanu, D., Joskowicz, L. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, Springer International Publishing, Cham, pp. 393–402. doi:https://doi.org/10.1007/978-3-030-59722-1_38.
- Salgado, R., Denkert, C., Demaria, S., Sirtaine, N., Klauschen, F., Pruneri, G., Wienert, S., Van den Eynden, G., Baehner, F., Penault-Llorca, F., Perez, E., Thompson, E., Symmans, W., Richardson, A., Brock, J., Criscitiello, C., Bailey, H., Ignatiadis, M., Floris, G., Sparano, J., Kos, Z., Nielsen, T., Rimm, D., Allison, K., Reis-Filho, J., Loibl, S., Sotiriou, C., Viale, G., Badve, S., Adams, S., Willard-Gallo, K., Loi, S., 2015. The evaluation of tumor-infiltrating lymphocytes (tils) in breast cancer: recommendations by an international tils working group 2014. *Annals of Oncology* 26, 259–271. URL: <https://www.sciencedirect.com/science/article/pii/S0923753419313699>, doi:<https://doi.org/10.1093/annonc/mdu450>.
- Schmidt, U., Weigert, M., Broaddus, C., Myers, G., 2018. Cell detection with star-convex polygons, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, pp. 265–273. doi:https://doi.org/10.1007/978-3-030-00934-2_30.
- Schölkopf, B., Locatello, F., Bauer, S., Ke, N.R., Kalchbrenner, N., Goyal, A., Bengio, Y., 2021. Toward causal representation learning. *Proceedings of the IEEE* 109, 612–634. doi:[10.1109/JPR0C.2021.3058954](https://doi.org/10.1109/JPR0C.2021.3058954).
- Shan, T., Yan, J., Cui, X., Xie, L., 2023. Dsca-net: A depthwise separable convolutional neural network with attention mechanism for medical image segmentation. *Mathematical Biosciences and Engineering* 20, 365–382. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85140914178&doi=10.3934%2fmbe.2023017&partnerID=40&md5=60a766589358704279749a6898d4d1b9>, doi:[10.3934/mbe.2023017](https://doi.org/10.3934/mbe.2023017).
- Shao, Z., Bian, H., Chen, Y., Wang, Y., Zhang, J., Ji, X., et al., 2021. Transmil: Transformer based correlated multiple instance learning for whole slide image classification. *Advances in neural information processing systems* 34, 2136–2147.
- Sheikh, T.S., Shim, J., Cho, M., 2022. A multi-class nuclei segmentation model using attention modules, in: *2022 7th International Conference on Intelligent Informatics and Biomedical Science (ICIIBMS)*, pp. 213–217. doi:[10.1109/ICIIBMS55689.2022.9971499](https://doi.org/10.1109/ICIIBMS55689.2022.9971499).
- Shi, Y., Daunhawer, I., Vogt, J.E., Torr, P.H., Sanyal, A., 2023. How robust is unsupervised representation learning to distribution shift?, in: *The Eleventh International Conference on Learning Representations*, OpenReview.
- Sirinukunwattana, K., Pluim, J.P., Chen, H., Qi, X., Heng, P.A., Guo, Y.B., Wang, L.Y., Matuszewski, B.J., Bruni, E., Sanchez, U., Böhm, A., Ronneberger, O., Cheikh, B.B., Racoceanu, D., Kainz, P., Pfeiffer, M., Urschler,

- M., Snead, D.R., Rajpoot, N.M., 2017. Gland segmentation in colon histology images: The glas challenge contest. *Medical Image Analysis* 35, 489–502. URL: <https://www.sciencedirect.com/science/article/pii/S1361841516301542>, doi:<https://doi.org/10.1016/j.media.2016.08.008>.
- Sirinukunwattana, K., Raza, S.E.A., Tsang, Y.W., Snead, D.R.J., Cree, I.A., Rajpoot, N.M., 2016. Locality sensitive deep learning for detection and classification of nuclei in routine colon cancer histology images. *IEEE Transactions on Medical Imaging* 35, 1196–1206. doi:10.1109/TMI.2016.2525803.
- Soenksen, L.R., Kassis, T., Conover, S.T., Marti-Fuster, B., Birkenfeld, J.S., Tucker-Schwartz, J., Naseem, A., Stavert, R.R., Kim, C.C., Senna, M.M., Avilés-Izquierdo, J., Collins, J.J., Barzilay, R., Gray, M.L., 2021. Using deep learning for dermatologist-level detection of suspicious pigmented skin lesions from wide-field images. *Science Translational Medicine* 13, eabb3652. doi:10.1126/scitranslmed.abb3652.
- Song, Z., Du, P., Yan, J., Li, K., Shou, J., Lai, M., Fan, Y., Xu, Y., 2023. Nucleus-aware self-supervised pretraining using unpaired image-to-image translation for histopathology images. *IEEE Transactions on Medical Imaging*, 1–1. doi:10.1109/TMI.2023.3309971.
- Srinidhi, C.L., Ciga, O., Martel, A.L., 2021. Deep neural network models for computational histopathology: A survey. *Medical Image Analysis* 67, 101813. URL: <https://www.sciencedirect.com/science/article/pii/S1361841520301778>, doi:<https://doi.org/10.1016/j.media.2020.101813>.
- Srinidhi, C.L., Martel, A.L., 2021. Improving self-supervised learning with hardness-aware dynamic curriculum learning: An application to digital pathology, in: 2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW), pp. 562–571. doi:10.1109/ICCVW54120.2021.00069.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1929–1958. URL: <http://jmlr.org/papers/v15/srivastava14a.html>.
- Steiner, A.P., Kolesnikov, A., Zhai, X., Wightman, R., Uszkoreit, J., Beyer, L., 2022. How to train your vit? data, augmentation, and regularization in vision transformers. *Transactions on Machine Learning Research* URL: <https://openreview.net/forum?id=4nPsWr1KCP>.
- Sugimoto, T., Ito, H., Teramoto, Y., Yoshizawa, A., Bise, R., 2022. Multi-class cell detection using modified self-attention, in: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 1854–1862. doi:10.1109/CVPRW56347.2022.00202.
- Sun, P., He, J., Chao, X., Chen, K., Xu, Y., Huang, Q., Yun, J., Li, M., Luo, R., Kuang, J., et al., 2021. A computational tumor-infiltrating lymphocyte assessment method comparable with visual reporting guidelines for triple-negative breast cancer. *EBioMedicine* 70.
- Tan, M., Le, Q., 2019. Efficientnet: Rethinking model scaling for convolutional neural networks, in: International conference on machine learning, PMLR. pp. 6105–6114.
- Tellez, D., Balkenhol, M., Karssemeijer, N., Litjens, G., van der Laak, J., Ciompi, F., 2018. H and E stain augmentation improves generalization of convolutional networks for histopathological mitosis detection, in: Tomaszewski, J.E., Gurcan, M.N. (Eds.), *Medical Imaging 2018: Digital Pathology*, International Society for Optics and Photonics. SPIE. p. 105810Z. URL: <https://doi.org/10.1117/12.2293048>, doi:10.1117/12.2293048.
- Thiagarajan, J.J., Kashyap, S., Karargyris, A., 2019. Distill-to-label: Weakly supervised instance labeling using knowledge distillation, in: 2019 18th IEEE International Conference on Machine Learning And Applications (ICMLA), pp. 902–907. doi:10.1109/ICMLA.2019.00156.
- Tourniaire, P., Ilie, M., Hofman, P., Ayache, N., Delingette, H., 2023. Ms-clam: Mixed supervision for the classification and localization of tumors in whole slide images. *Medical Image Analysis* 85, 102763. URL: <https://www.sciencedirect.com/science/article/pii/S1361841523000245>, doi:<https://doi.org/10.1016/j.media.2023.102763>.
- Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jegou, H., 2021a. Training data-efficient image transformers & distillation through attention, in: International Conference on Machine Learning, pp. 10347–10357.
- Touvron, H., Cord, M., El-Nouby, A., Bojanowski, P., Joulin, A., Synnaeve, G., Jégou, H., 2021b. Augmenting convolutional networks with attention-based aggregation. *arXiv preprint arXiv:2112.13692*.
- of Tumours Editorial Board, W.C., 2019. WHO Classification of Tumours, 5th Edition, Volume 2: Breast Tumours. IARC Publications.
- Ullah, Z., Usman, M., Gwak, J., 2023. Mtss-aae: Multi-task semi-supervised adversarial autoencoding for covid-19 detection based on chest x-ray images. *Expert Systems with Applications* 216, 119475. URL: <https://www.sciencedirect.com/science/article/pii/S0957417422024940>, doi:<https://doi.org/10.1016/j.eswa.2022.119475>.
- Vahadane, A., B, A., Majumdar, S., 2021. Dual encoder attention u-net for nuclei segmentation, in: 2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), pp. 3205–3208. doi:10.1109/EMBC46164.2021.9630037.
- Valanarasu, J.M.J., Oza, P., Hachililoglu, I., Patel, V.M., 2021. Medical transformer: Gated axial-attention for medical image segmentation, in: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, Springer International Publishing, Cham. pp. 36–46. doi:https://doi.org/10.1007/978-3-030-87193-2_4.
- van der Velden, B.H., Kuijff, H.J., Gilhuijs, K.G., Viergever, M.A., 2022. Explainable artificial intelligence (xai) in deep learning-based medical image analysis. *Medical Image Analysis* 79, 102470. URL: <https://www.sciencedirect.com/science/article/pii/S1361841522001177>, doi:<https://doi.org/10.1016/j.media.2022.102470>.
- Vanguri, R.S., Luo, J., Aukerman, A.T., Egger, J.V., Fong, C.J., Horvat, N., Pagano, A., Araujo-Filho, J.d.A.B., Geneslaw, L., Rizvi, H., Sosa, R., Boehm, K.M., Yang, S.R., Bodd, F.M., Ventura, K., Hollmann, T.J., Ginsberg, M.S., Gao, J., Vanguri, R., Hellmann, M.D., Sauter, J.L., Shah, S.P., Consortium, M.S.K.M.I.N.D., 2022. Multimodal integration of radiology, pathology and genomics for prediction of response to pd-(l)1 blockade in patients with non-small cell lung cancer. *Nature Cancer* 3, 1151–1164. URL: <https://doi.org/10.1038/s43018-022-00416-8>, doi:10.1038/s43018-022-00416-8.
- Varoquaux, G., Cheplygina, V., 2022. Machine learning for medical imaging: methodological failures and recommendations for the future. *npj Digital Medicine* 5, 48. URL: <https://doi.org/10.1038/s41746-022-00592-y>, doi:10.1038/s41746-022-00592-y.
- Veličković, P., 2023. Everything is connected: Graph neural networks. *Current Opinion in Structural Biology* 79, 102538.
- Verma, R., Kumar, N., Patil, A., Kurian, N.C., Rane, S., Graham, S., Vu, Q.D., Zwager, M., et al., 2021. Monusac2020: A multi-organ nuclei segmentation and classification challenge. *IEEE Transactions on Medical Imaging* 40, 3413–3423. doi:10.1109/TMI.2021.3085712.
- Vo, V.T., Kim, S.H., 2023. Mulvernet: Nucleus segmentation and classification of pathology images using the hover-net and multiple filter units. *Electronics (Switzerland)* 12. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85146793207&doi=10.3390%2felectronics12020355&partnerID=40&md5=6258cac2926aa5af9a6711c8b2b81923>, doi:10.3390/electronics12020355.
- Vu, Q.D., Graham, S., Kurc, T., To, M.N.N., Shaban, M., Qaiser, T., Koohbanani, N.A., Khurram, S.A., Kalpathy-Cramer, J., Zhao, T., Gupta, R., Kwak, J.T., Rajpoot, N., Saltz, J., Farahani, K., 2019. Methods for segmentation and classification of digital microscopy tissue images. *Frontiers in Bioengineering and Biotechnology* 7. URL: <https://www.frontiersin.org/articles/10.3389/fbioe.2019.00053>, doi:10.3389/fbioe.2019.00053.
- Wagner, R., Rohr, K., 2022. Cellcentroidformer: Combining self-attention and convolution for cell detection, in: Yang, G., Aviles-Rivero, A., Roberts, M., Schönlieb, C.B. (Eds.), *Medical Image Understanding and Analysis*, Springer International Publishing, Cham. pp. 212–222. doi:https://doi.org/10.1007/978-3-031-12053-4_16.
- Wahab, N., Miligy, I.M., Dodd, K., Sahota, H., Toss, M., Lu, W., Jahanifar, M., Bilal, M., Graham, S., Park, Y., Hadjigeorgiou, G., Bhalerao, A., Lashen, A.G., Ibrahim, A.Y., Katayama, A., Ebili, H.O., Parkin, M., Sorell, T., Raza, S.E.A., Hero, E., Eldaly, H., Tsang, Y.W., Gopalakrishnan, K., Snead, D., Rakha, E., Rajpoot, N., Minhas, F., 2022. Semantic annotation for computational pathology: multidisciplinary experience and best practice recommendations. *The Journal of Pathology: Clinical Research* 8, 116–128. URL: <https://pathsocjournals.onlinelibrary.wiley.com/doi/abs/10.1002/cjp2.256>, doi:<https://doi.org/10.1002/cjp2.256>, arXiv:<https://pathsocjournals.onlinelibrary.wiley.com/doi/pdf/10.1002/cjp2.256>.
- Wan, T., Zhao, L., Feng, H., Li, D., Tong, C., Qin, Z., 2020. Robust nuclei segmentation in histopathology using asppu-net and

- boundary refinement. *Neurocomputing* 408, 144–156. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85082497177&doi=10.1016%2fj.neucom.2019.08.103&partnerID=40&md5=e4205e8105f57102df65d0c900a6f138>, doi:10.1016/j.neucom.2019.08.103.
- Wan, Z., Li, M., Wang, Z., Tan, H., Li, W., Yu, L., Samuel, D.J., 2023. Cellt-net: A composite transformer method for 2-d cell instance segmentation. *IEEE Journal of Biomedical and Health Informatics* , 1–15doi:10.1109/JBHI.2023.3265006.
- Wang, H., Wang, G., Sheng, Z., Zhang, S., 2019. Automated segmentation of skin lesion based on pyramid attention network, in: Suk, H.I., Liu, M., Yan, P., Lian, C. (Eds.), *Machine Learning in Medical Imaging*, Springer International Publishing, Cham. pp. 435–443. doi:https://doi.org/10.1007/978-3-030-32692-0_50.
- Wang, H., Xu, G., Pan, X., Liu, Z., Lan, R., Luo, X., 2022a. Multi-task generative adversarial learning for nuclei segmentation with dual attention and recurrent convolution. *Biomedical Signal Processing and Control* 75. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85124284705&doi=10.1016%2fj.bspc.2022.103558&partnerID=40&md5=79452a04cb5fc5cad2c21cb87949b1da>, doi:10.1016/j.bspc.2022.103558.
- Wang, J., Zhang, Z., Wu, M., Ye, Y., Wang, S., Cao, Y., Yang, H., 2023. Nuclei instance segmentation using a transformer-based graph convolutional network and contextual information augmentation. *Computers in Biology and Medicine* 167, 107622. URL: <https://www.sciencedirect.com/science/article/pii/S0010482523010879>, doi:<https://doi.org/10.1016/j.combiomed.2023.107622>.
- Wang, N., Quan, H., 2021. Glunet: Global-local fusion u-net for 2d medical image segmentation, in: Farkaš, I., Masulli, P., Otte, S., Wermter, S. (Eds.), *Artificial Neural Networks and Machine Learning – ICANN 2021*, Springer International Publishing, Cham. pp. 74–85. doi:https://doi.org/10.1007/978-3-030-86380-7_7.
- Wang, X., 2022. Context understanding in computer vision: A survey. *SSRN Electronic Journal* doi:10.2139/ssrn.4081162.
- Wang, X., Yang, S., Zhang, J., Wang, M., Zhang, J., Yang, W., Huang, J., Han, X., 2022b. Transformer-based unsupervised contrastive learning for histopathological image classification. *Medical Image Analysis* 81, 102559. URL: <https://www.sciencedirect.com/science/article/pii/S1361841522002043>, doi:<https://doi.org/10.1016/j.media.2022.102559>.
- Wazir, S., Fraz, M.M., 2022. Histoseg: Quick attention with multi-loss function for multi-structure segmentation in digital histology images, in: 2022 12th International Conference on Pattern Recognition Systems (ICPRS), pp. 1–7. doi:10.1109/ICPRS54038.2022.9854067.
- Wei, X., Bai, H., Zhang, X.S., Li, Y., 2022. Deep pneumonia: Attention-based contrastive learning for class-imbalanced pneumonia lesion recognition in chest x-rays. *ArXiv abs/2207.11393*.
- Weigert, M., Schmidt, U., 2022. Nuclei instance segmentation and classification in histopathology images with stardist, in: 2022 IEEE International Symposium on Biomedical Imaging Challenges (ISBIC), IEEE. pp. 1–4.
- Wilkinson, M.D.e.a., 2016. The fair guiding principles for scientific data management and stewardship. *Scientific Data* 3, 160018. URL: <https://doi.org/10.1038/sdata.2016.18>, doi:10.1038/sdata.2016.18.
- Woo, S., Park, J., Lee, J.Y., Kweon, I.S., 2018. Cbam: Convolutional block attention module, in: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (Eds.), *Computer Vision – ECCV 2018*, Springer International Publishing, Cham. pp. 3–19. doi:https://doi.org/10.1007/978-3-030-01234-2_1.
- Wu, H., Xiao, B., Codella, N., Liu, M., Dai, X., Yuan, L., Zhang, L., 2021. Cvt: Introducing convolutions to vision transformers, in: 2021 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 22–31. doi:10.1109/ICCV48922.2021.00009.
- Wu, Y., Zhou, Y., Saiyin, J., Wei, B., Lai, M., Shou, J., Fan, Y., Xu, Y., 2023. Zero-shot nuclei detection via visual-language pre-trained models, in: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*, Springer Nature Switzerland, Cham. pp. 693–703.
- Xiao, T., Singh, M., Mintun, E., Darrell, T., Dollár, P., Girshick, R., 2021. Early convolutions help transformers see better, in: Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (Eds.), *Advances in Neural Information Processing Systems*, Curran Associates, Inc.. pp. 30392–30400. URL: https://proceedings.neurips.cc/paper_files/paper/2021/file/ff1418e8cc993fe8abcfe3ce2003e5c5-Paper.pdf.
- Xie, C., Mousavian, A., Xiang, Y., Fox, D., 2022a. Rice: Refining instance masks in cluttered environments with graph neural networks, in: *Conference on Robot Learning*, PMLR. pp. 1655–1665.
- Xie, J., Pu, X., He, J., Qiu, Y., Lu, C., Gao, W., Wang, X., Lu, H., Shi, J., Xu, Y., Madabhushi, A., Fan, X., Chen, J., Xu, J., 2022b. Survival prediction on intrahepatic cholangiocarcinoma with histomorphological analysis on the whole slide images. *Computers in Biology and Medicine* 146, 105520. URL: <https://www.sciencedirect.com/science/article/pii/S0010482522003122>, doi:<https://doi.org/10.1016/j.combiomed.2022.105520>.
- Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K., 2017. Aggregated residual transformations for deep neural networks, in: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 5987–5995. doi:10.1109/CVPR.2017.634.
- Xu, K., Goetz, L., Rajpoot, N., 2024. On generalisability of segment anything model for nuclear instance segmentation in histology images. *arXiv preprint arXiv:2401.14248*.
- Xu, R., Li, Y., Wang, C., Xu, S., Meng, W., Zhang, X., 2022. Instance segmentation of biological images using graph convolutional network. *Engineering Applications of Artificial Intelligence* 110, 104739. URL: <https://www.sciencedirect.com/science/article/pii/S0952197622000446>, doi:<https://doi.org/10.1016/j.engappai.2022.104739>.
- Xu, R., Wang, C., Xu, S., Meng, W., Zhang, X., 2021. De-net: Dual context network for 2d medical image segmentation, in: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, Springer International Publishing, Cham. pp. 503–513. doi:https://doi.org/10.1007/978-3-030-87193-2_48.
- Yan, R., Ren, F., Li, J., Rao, X., Lv, Z., Zheng, C., Zhang, F., 2022. Nuclei-guided network for breast cancer grading in he-stained pathological images. *Sensors* 22. URL: <https://www.mdpi.com/1424-8220/22/11/4061>, doi:10.3390/s22114061.
- Yang, L., Zhang, R.Y., Li, L., Xie, X., 2021. Simam: A simple, parameter-free attention module for convolutional neural networks, in: *International conference on machine learning*, PMLR. pp. 11863–11874.
- Yang, Y., Li, H., Li, X., Zhao, Q., Wu, J., Lin, Z., 2020. Sognet: Scene overlap graph network for panoptic segmentation. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 12637–12644. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/6955>, doi:10.1609/aaai.v34i07.6955.
- Yi, C., Chen, X., Quan, L., Lu, C., 2020. Attentional dilated convolution neural network for nuclei segmentation in histopathology images, *Institute of Electrical and Electronics Engineers Inc.* pp. 6737–6740. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85100947100&doi=10.1109%2fCAC51589.2020.9326630&partnerID=40&md5=6c25b5590a2604953244f41396514a8e>, doi:10.1109/CAC51589.2020.9326630.
- Yi, J., Wu, P., Jiang, M., Huang, Q., Hoepfner, D.J., Metaxas, D.N., 2019. Attentive neural cell instance segmentation. *Medical Image Analysis* 55, 228–240. URL: <https://www.sciencedirect.com/science/article/pii/S1361841518308442>, doi:<https://doi.org/10.1016/j.media.2019.05.004>.
- Yoo, I., Yoo, D., Paeng, K., 2019. Pseudodengen: Nuclei segmentation only with point annotations, in: Shen, D., Liu, T., Peters, T.M., Staib, L.H., Essert, C., Zhou, S., Yap, P.T., Khan, A. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, Springer International Publishing, Cham. pp. 731–739. doi:https://doi.org/10.1007/978-3-030-32239-7_81.
- Yoo, J.J., Namdar, K., Wagner, M.W., Nobre, L., Tabori, U., Hawkins, C., Ertl-Wagner, B.B., Khalvati, F., 2022. A novel gan-based paradigm for weakly supervised brain tumor segmentation of mr images. *arXiv preprint arXiv:2211.05269*.
- Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S., 2021. Barlow twins: Self-supervised learning via redundancy reduction, in: *International Conference on Machine Learning*, PMLR. pp. 12310–12320.
- Zeng, Z., Xie, W., Zhang, Y., Lu, Y., 2019. Ric-unet: An improved neural network based on unet for nuclei segmentation in histology images. *IEEE Access* 7, 21420–21428. doi:10.1109/ACCESS.2019.2896920.
- Zhang, C., Zhang, M., Zhang, S., Jin, D., Zhou, Q., Cai, Z., Zhao, H.,

- Liu, X., Liu, Z., 2022a. Delving deep into the generalization of vision transformers under distribution shifts, in: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 7267–7276. doi:10.1109/CVPR52688.2022.00713.
- Zhang, D., Han, J., Cheng, G., Yang, M.H., 2022b. Weakly supervised object localization and detection: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 5866–5885. doi:10.1109/TPAMI.2021.3074313.
- Zhang, D., Zhang, H., Tang, J., Hua, X.S., Sun, Q., 2020. Causal intervention for weakly-supervised semantic segmentation, in: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (Eds.), *Advances in Neural Information Processing Systems*, Curran Associates, Inc., pp. 655–666. URL: https://proceedings.neurips.cc/paper_files/paper/2020/file/07211688a0869d995947a8fb11b215d6-Paper.pdf.
- Zhang, H., Liang, P., Sun, Z., Song, B., Cheng, E., 2023a. Circleformer: Circular nuclei detection in whole slide images with circle queries and attention, in: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*, Springer Nature Switzerland, Cham, pp. 493–502.
- Zhang, H., Liu, J., Yu, Z., Wang, P., 2021. Masg-gan: A multi-view attention superpixel-guided generative adversarial network for efficient and simultaneous histopathology image segmentation and classification. *Neurocomputing* 463, 275–291. URL: <https://www.sciopus.com/inward/record.uri?eid=2-s2.0-85113589295&doi=10.1016/j.neucom.2021.08.039&partnerID=40&md5=204acbc88155f422da17451e87962550>, doi:10.1016/j.neucom.2021.08.039.
- Zhang, J., Ma, K., Kapse, S., Saltz, J., Vakalopoulou, M., Prasanna, P., Samaras, D., 2023b. Sam-path: A segment anything model for semantic segmentation in digital pathology, in: Celebi, M.E., Salekin, M.S., Kim, H., Albarqouni, S., Barata, C., Halpern, A., Tschandl, P., Combalia, M., Liu, Y., Zamzmi, G., Levy, J., Rangwala, H., Reinke, A., Wynn, D., Landman, B., Jeong, W.K., Shen, Y., Deng, Z., Bakas, S., Li, X., Qin, C., Rieke, N., Roth, H., Xu, D. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023 Workshops*, Springer Nature Switzerland, Cham, pp. 161–170.
- Zhang, X., Zhu, X., Tang, K., Zhao, Y., Lu, Z., Feng, Q., 2022c. Ddtnet: A dense dual-task network for tumor-infiltrating lymphocyte detection and segmentation in histopathological images of breast cancer. *Medical image analysis* 78, 102415.
- Zhang, Y., Zhou, T., Wang, S., Liang, P., Zhang, Y., Chen, D.Z., 2023c. Input augmentation with sam: Boosting medical image segmentation with segmentation foundation model, in: Celebi, M.E., Salekin, M.S., Kim, H., Albarqouni, S., Barata, C., Halpern, A., Tschandl, P., Combalia, M., Liu, Y., Zamzmi, G., Levy, J., Rangwala, H., Reinke, A., Wynn, D., Landman, B., Jeong, W.K., Shen, Y., Deng, Z., Bakas, S., Li, X., Qin, C., Rieke, N., Roth, H., Xu, D. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023 Workshops*, Springer Nature Switzerland, Cham, pp. 129–139.
- Zhao, J., He, Y.J., Zhao, S.Q., Huang, J.J., Zuo, W.M., 2022a. Al-net: Attention learning network based on multi-task learning for cervical nucleus segmentation. *IEEE Journal of Biomedical and Health Informatics* 26, 2693–2702. URL: <https://www.sciopus.com/inward/record.uri?eid=2-s2.0-85122078106&doi=10.1109/JBHI.2021.3136568&partnerID=40&md5=57fd0211fc1dfbb00ea6d47b2f7c2559>, doi:10.1109/JBHI.2021.3136568.
- Zhao, J., Hou, X., Pan, M., Zhang, H., 2022b. Attention-based generative adversarial network in medical imaging: A narrative review. *Computers in Biology and Medicine* 149, 105948. URL: <https://www.sciencedirect.com/science/article/pii/S0010482522006837>, doi:https://doi.org/10.1016/j.combiomed.2022.105948.
- Zhao, S., Chen, D.P., Fu, T., Yang, J.C., Ma, D., Zhu, X.Z., Wang, X.X., Jiao, Y.P., Jin, X., Xiao, Y., Xiao, W.X., Zhang, H.Y., Lv, H., Madabhushi, A., Yang, W.T., Jiang, Y.Z., Xu, J., Shao, Z.M., 2023. Single-cell morphological and topological atlas reveals the ecosystem diversity of human breast cancer. *Nature Communications* 14, 6796. URL: <https://doi.org/10.1038/s41467-023-42504-y>, doi:10.1038/s41467-023-42504-y.
- Zhou, K., Liu, Z., Qiao, Y., Xiang, T., Loy, C.C., 2022. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–20doi:10.1109/TPAMI.2022.3195549.
- Zhou, X., Cheng, Z., Gu, M., Chang, F., 2020. Lirnet: Local integral regression network for both strongly and weakly supervised nuclei detection, in: 2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), pp. 945–951. doi:10.1109/BIBM49941.2020.9313265.
- Zhou, Y., Chen, H., Xu, J., Dou, Q., Heng, P.A., 2019a. Inet: Instance relation network for overlapping cervical cell segmentation, in: Shen, D., Liu, T., Peters, T.M., Staib, L.H., Essert, C., Zhou, S., Yap, P.T., Khan, A. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, Springer International Publishing, Cham, pp. 640–648. doi:https://doi.org/10.1007/978-3-030-32239-7_71.
- Zhou, Y., Graham, S., Alemi Koohbanani, N., Shaban, M., Heng, P.A., Rajpoot, N., 2019b. Cgc-net: Cell graph convolutional network for grading of colorectal cancer histology images, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*.
- Zhou, Y., Li, W., Yang, G., 2023. Scts: Instance segmentation of single cells using a transformer-based semantic-aware model and space-filling augmentation, in: 2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), pp. 5933–5942. doi:10.1109/WACV56688.2023.00589.
- Zhou, Y., Onder, O.F., Dou, Q., Tsougenis, E., Chen, H., Heng, P.A., 2019c. Cia-net: Robust nuclei instance segmentation with contour-aware information aggregation, in: Chung, A.C.S., Gee, J.C., Yushkevich, P.A., Bao, S. (Eds.), *Information Processing in Medical Imaging*, Springer International Publishing, Cham, pp. 682–693. doi:https://doi.org/10.1007/978-3-030-20351-1_53.
- Zhou, Z., Rahman Siddiquee, M.M., Tajbakhsh, N., Liang, J., 2018. Unet++: A nested u-net architecture for medical image segmentation, in: Stoyanov, D., Taylor, Z., Carneiro, G., Syeda-Mahmood, T., Martel, A., Maier-Hein, L., Tavares, J.M.R., Bradley, A., Papa, J.P., Belagiannis, V., Nascimento, J.C., Lu, Z., Conjeti, S., Moradi, M., Greenspan, H., Madabhushi, A. (Eds.), *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, Springer International Publishing, Cham, pp. 3–11. doi:https://doi.org/10.1007/978-3-030-00889-5_1.
- Zhou, Z.H., 2018. A brief introduction to weakly supervised learning. *National science review* 5, 44–53.
- Zhu, Z., Liang, D., Zhang, S., Huang, X., Li, B., Hu, S., 2016. Traffic-sign detection and classification in the wild, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2110–2118. doi:10.1109/CVPR.2016.232.
- Zimmerer, D., Kohl, S.A., Petersen, J., Isensee, F., Maier-Hein, K.H., 2018. Context-encoding variational autoencoder for unsupervised anomaly detection. arXiv preprint arXiv:1812.05941.
- Zunair, H., Ben Hamza, A., 2021. Sharp u-net: Depthwise convolutional network for biomedical image segmentation. *Computers in Biology and Medicine* 136, 104699. URL: <https://www.sciencedirect.com/science/article/pii/S0010482521004935>, doi:https://doi.org/10.1016/j.combiomed.2021.104699.