



One model is all you need: Multi-task learning enables simultaneous histology image segmentation and classification

Simon Graham^{a,b,*}, Quoc Dang Vu^a, Mostafa Jahanifar^a, Shan E Ahmed Raza^a, Fayyaz Minhas^a, David Snead^{b,c}, Nasir Rajpoot^{a,b,c}

^a Tissue Image Analytics Centre, Department of Computer Science, University of Warwick, United Kingdom

^b Histofy Ltd, United Kingdom

^c Department of Pathology, University Hospitals Coventry & Warwickshire, United Kingdom

ARTICLE INFO

Keywords:

Computational pathology
Multi-task learning
Deep learning

ABSTRACT

The recent surge in performance for image analysis of digitised pathology slides can largely be attributed to the advances in deep learning. Deep models can be used to initially localise various structures in the tissue and hence facilitate the extraction of interpretable features for biomarker discovery. However, these models are typically trained for a single task and therefore scale poorly as we wish to adapt the model for an increasing number of different tasks. Also, supervised deep learning models are very *data hungry* and therefore rely on large amounts of training data to perform well. In this paper, we present a multi-task learning approach for segmentation and classification of nuclei, glands, lumina and different tissue regions that leverages data from multiple independent data sources. While ensuring that our tasks are aligned by the same tissue type and resolution, we enable meaningful simultaneous prediction with a single network. As a result of feature sharing, we also show that the learned representation can be used to improve the performance of additional tasks via transfer learning, including nuclear classification and signet ring cell detection. As part of this work, we train our developed *Cerberus* model on a huge amount of data, consisting of over 600 thousand objects for segmentation and 440 thousand patches for classification. We use our approach to process 599 colorectal whole-slide images from TCGA, where we localise 377 million, 900 thousand and 2.1 million nuclei, glands and lumina respectively. We make this resource available to remove a major barrier in the development of explainable models for computational pathology.

1. Introduction

In recent years, there has been a progressive shift towards the digitisation of histopathology slides, enabling the development of computer vision (CV) techniques for automated tissue analysis. In particular, with the rapidly increasing amount of pathology image data and computing power, deep learning has revolutionised the field of computational pathology (CPath). A large proportion of deep CV models used in CPath are convolutional neural networks (CNNs) and have proven to be successful when applied to a plethora of tasks, including cancer detection (Bejnordi et al., 2017; Veeling et al., 2018), cancer grading (Bulten et al., 2020; Shaban et al., 2020) and survival analysis (Bychkov et al., 2018; Shaban et al., 2019). However, *black box* CNNs may have a poor level of interpretability when used directly in slide classification, as they are typically non-linear and are governed by millions of parameters.

To help overcome the above challenge, deep learning can be leveraged as an initial step to localise regions of interest in the tissue

before extracting human-interpretable and biologically meaningful features for downstream slide analysis. This fosters the development of more explainable pipelines for CPath, which is particularly important for boosting the confidence of pathologists in their diagnoses. Explainable models increase model transparency by providing an output that is easy to understand and therefore may be utilised as a tool to help facilitate the decision of the pathologist. Furthermore, using human-interpretable features (Diao et al., 2021) in explainable pipelines can help reveal new biomarkers for complex tasks, such as survival analysis and image-based detection of clinically actionable genetic alterations (Kather et al., 2020).

Previous works on the localisation of various regions and constructs in the tissue typically involve the use of a devoted CNN for a particular task. For example, Diao et al. (2021) trained a network to detect and classify nuclei and an additional network for the segmentation of different tissue types before feature extraction. This strategy is often needed

* Corresponding author at: Tissue Image Analytics Centre, Department of Computer Science, University of Warwick, United Kingdom.
E-mail address: simon.graham@warwick.ac.uk (S. Graham).

because conventional training pipelines require target objects to be labelled across the entire dataset. Optimising networks separately also ensures that training procedures do not need to be tuned over multiple tasks. However, using one model per task may not be scalable as we wish to localise an increasing number of tissue constructs. Instead, utilising a single model that can perform multiple tasks *simultaneously* is preferable because it reduces the computational overhead, can benefit from shared features between tasks and also saves the time needed for training separate models.

Multi-Task Learning (MTL) (Caruana, 1997; Zhang and Yang, 2017) has previously been explored in CPath, but with a key focus on obtaining a generalisable encoder (Gamper et al., 2020b) for transfer learning (Mormont et al., 2020) or image compression (Tellez et al., 2020). Here, each task may include image regions from different tissue types, with varying resolution and potentially with different staining. Therefore, such models cannot be used to perform simultaneous prediction for a single input image, due to inconsistency between assumed inputs for each task. Instead, it may be beneficial to utilise an MTL approach with *aligned* tasks, thereby enabling simultaneous prediction with a single model. Not only does this enable a meaningful simultaneous prediction, but using aligned tasks also helps ease optimisation by reducing the potential for conflicting gradients to exist between tasks. Therefore, this holds potential for achieving a strong performance across all considered tasks, which was previously not possible with other MTL approaches in CPath (Gamper et al., 2020b).

To perform well, deep learning approaches require an abundance of labelled data, which is a current bottleneck in the development of CPath models because the collection of annotations requires significant pathologist input (Wahab et al., 2021; Tizhoosh and Pantanowitz, 2018). Without large and diverse datasets, any developed model may struggle to generalise to unseen examples, limiting potential usage in a clinical setting. MTL may help to overcome this challenge by utilising a shared encoder, enabling features to be learned over multiple tasks (Crawshaw, 2020). Yet, it is desirable for each task to have access to a large amount of annotated data from multiple sources to ensure strong performance.

In this paper, we propose a multi-task learning approach for simultaneous segmentation and classification of nuclei, glands, lumina and different tissue regions. With the use of a novel sampling strategy, the proposed *Cerberus* model leverages data from multiple independent sources during training, enabling a competitive performance compared to single-task alternatives while only requiring a single network. Therefore, our approach is the first MTL method in CPath that can sustain a strong performance at the output of the network across all considered tasks. We demonstrate that Cerberus learns a strong feature representation that can help improve the performance of additional tasks via transfer learning, including nuclear classification and signet ring cell detection. To ensure that Cerberus has sufficient labelled examples from multiple different data sources, we create a large and diverse colorectal dataset containing over 600 thousand object boundaries for segmentation and 440 thousand patches for classification. In addition, we process all colorectal slides from The Cancer Genome Atlas (TCGA) and make the results accessible. We hope that doing this will remove a major barrier for the development of explainable approaches for CPath.

The main contributions of this work are listed as follows:

- We present a multi-task network, named Cerberus, that is capable of performing simultaneous prediction with superior performance as compared to single-task learning methods.
- We show that Cerberus learns a strong feature representation that can be leveraged during transfer learning to boost the performance of additional tasks.
- We use our method to process all colorectal slides from TCGA, resulting in 377 million, 900 thousand and 2.1 million nuclei, glands and lumina respectively and make the results available¹.

¹ The results on TCGA and the pretrained Cerberus backbone can be downloaded here: <https://github.com/TissueImageAnalytics/cerberus>

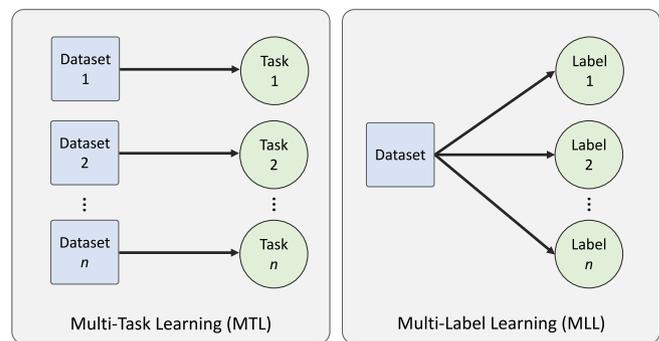


Fig. 1. Multi-Task Learning (MTL) vs Multi-Label Learning (MLL): MTL is a training paradigm where a model is trained over different predictive tasks. Here, labels are task-dependent. MLL requires labels to be provided for every input sample.

2. Related work

2.1. Multi-task learning vs multi-label learning

MTL for computer vision enables multiple independent datasets to be leveraged during training time, which offers advantages like improved data efficiency, reduced overfitting through shared representations and fast learning by leveraging auxiliary information (Crawshaw, 2020). There are various ways in which MTL can be implemented, including weight sharing (Caruana, 1997; Mormont et al., 2020) and architecture adaptations (Strezoski et al., 2019; Misra et al., 2016). For deep learning, MTL is typically done via hard or soft parameter sharing of hidden layers (Ruder, 2017). Hard parameter sharing is the most common strategy used for MTL, where hidden layers are shared between tasks before using several task-specific output layers. In soft parameter sharing, each task has its own model, where the associated parameters are constrained (Duong et al., 2015; Yang and Hospedales, 2016) so that they are similar for each task. Therefore, using hard parameter sharing may be preferable because using separate models may not be feasible when the number of tasks is large.

It is important to distinguish MTL from Multi-Label Learning (MLL) (Zhang and Zhou, 2013), where each data point is associated with multiple labels. If each label is treated as a separate task, MLL can be viewed as a special case of MTL where different tasks always share the same data (Zhang and Yang, 2017). For clarity, we show the differences between these two learning strategies in Fig. 1.

2.2. Multi-label learning for CPath

As mentioned above, it is possible to predict multiple tasks at the output of the network using MLL. For example, Shephard et al. (2021) performed nuclear segmentation and classification along with segmentation of intra-epithelial layers in oral tissue. Similarly Fraz et al. (2020) performed simultaneous segmentation of blood vessels and nerves. However, because the two above-mentioned approaches use MLL, they expect every input image to have associated labels for all tasks. This limits the data that can be used because it is common for available datasets in CPath to only consider a single object (Sirinukunwattana et al., 2017; Gamper et al., 2020a) or label (Kather et al., 2019). Also, tasks require different levels of annotation effort (eg. nuclei take longer to annotate than glands) and therefore exhaustively labelling data across tasks may not be feasible. As a result of the challenges above and to enable the effective utilisation of multiple independent datasets, in this paper we focus on MTL, rather than MLL.

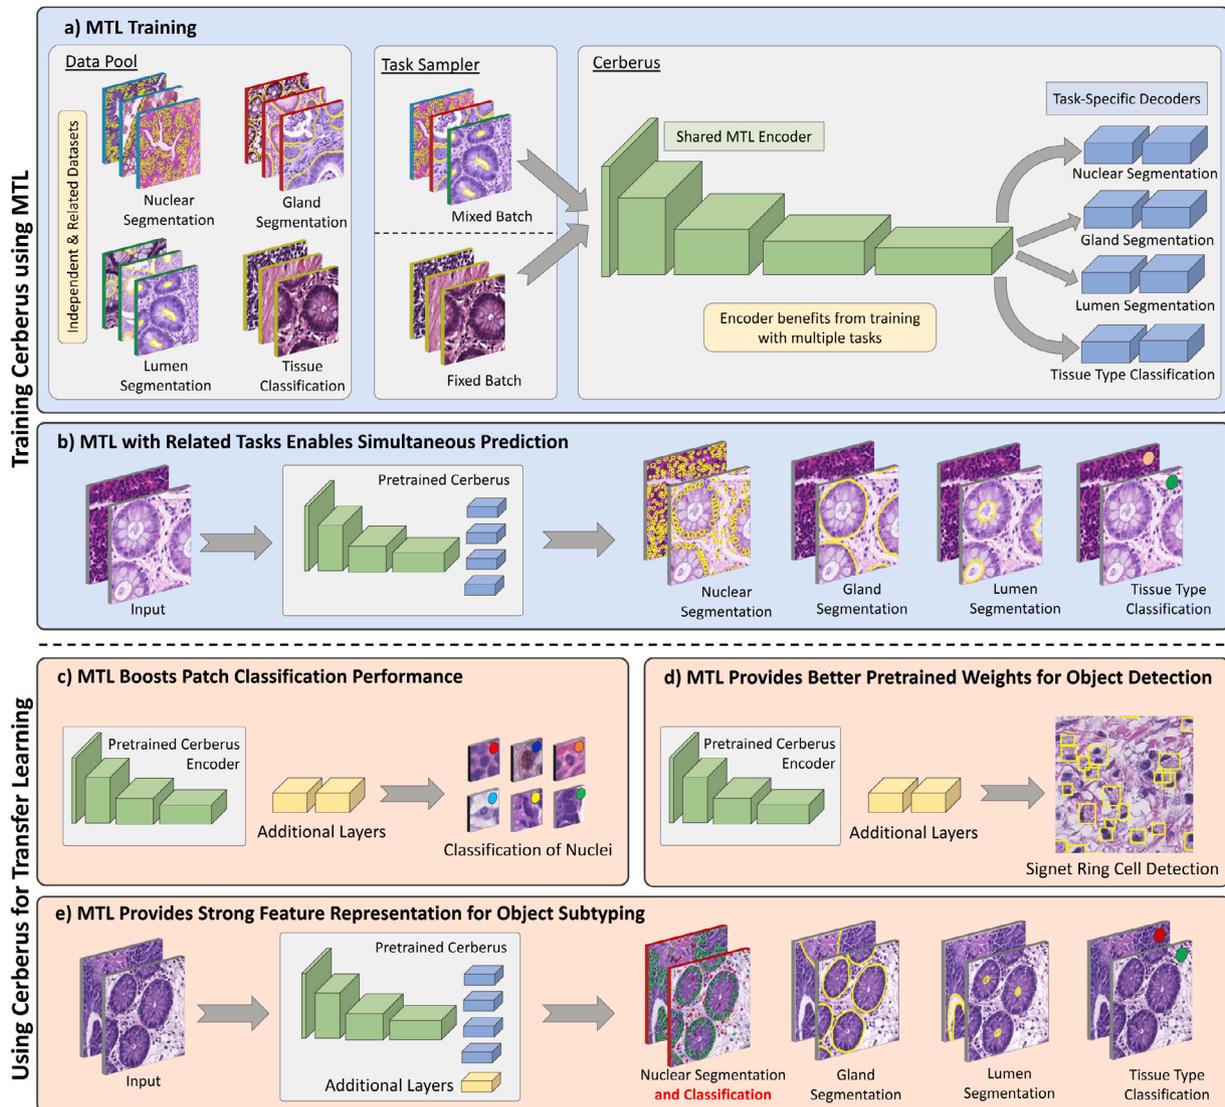


Fig. 2. MTL training strategy (a), where patches are extracted from each task-level dataset by the task sampler and are fed into Cerberus. The sampler either chooses a fixed or mixed batch, where input samples are drawn from the same task or a mixture of tasks, respectively. Training with Cerberus (b) enables meaningful predictions to be made simultaneously. Using the trained Cerberus encoder can also help improve the performance of other tasks, such as patch classification (c), object detection (d) and object subtyping (e). For (c) and (d), only the pretrained Cerberus encoder is leveraged and additional classification or detection layers are added. For (e), the entire Cerberus architecture is used and extended by adding additional task-dependent decoders.

2.3. Multi-task learning for CPath

Several recent CPath models utilise MTL to obtain a representative set of features for transfer learning. For example, [Mormont et al. \(2020\)](#) considered 22 classification tasks within a simple MTL framework, where the learned weights are used to better initialise single task networks for transfer learning. Similarly, [Tellez et al. \(2020\)](#) used MTL to optimise 4 tasks simultaneously and used the learned weights to encode a low-dimensional feature representation that generalised well to new tasks. [Gamper and Rajpoot \(2021\)](#) used MTL in combination with multiple instance captioning to learn representations from histopathology textbooks. Here, it was again shown that using MTL provided a strong feature encoder for transfer learning. Despite the demonstrated benefit of using MTL for transfer learning, subsequent training may still require a single model per task, which comes at an extra computational cost and demands additional training time.

Instead, it may also be desirable to perform well for all considered MTL tasks at the *output* of the network. In this way, MTL provides a mechanism to learn from multiple independent datasets and enables

simultaneous predictions to be made. However, this is only recommended if all considered tasks are aligned. For example, in work by [Mormont et al. \(2020\)](#) one task uses images stained via immunohistochemistry (IHC) for identification of tumour and stroma regions in colorectal tissue, while another task uses Haematoxylin & Eosin (H&E) stained bone marrow tissue for cell classification. Therefore, unexpected results may be obtained at the output of the cell classification branch if presented with IHC stained data. This is because different input data assumptions exist for each task. In fact, none of the above-mentioned studies use MTL for the primary purpose of simultaneous prediction. With the use of different input assumptions, there may be task entanglement and consequently gradient conflicts ([Liu et al., 2021](#); [Sener and Koltun, 2018](#)) during network optimisation. This is a major reason why recent MTL methods in CPath cannot match single-task performance at the network output ([Gamper et al., 2020b](#)).

2.4. Supervised learning datasets for CPath

MTL can help improve data efficiency because tasks with limited examples can benefit from the data provided by other tasks. Despite

this, it is desirable to ensure that a sufficient number of examples from a variety of different sources are available per task, to increase the likelihood of strong generalisation.

Over recent years, there has been a growing effort in generating labelled datasets for CPath (Graham et al., 2019c; Kumar et al., 2019; Verma et al., 2020; Amgad et al., 2021; Naylor et al., 2018; Gamper et al., 2020a; Graham et al., 2021a). Due to differences in the tissue type, resolution and stain, there may exist significant variation between these datasets. For example, Amgad et al. (2019) generated a dataset of annotated regions in H&E stained breast tissue at 20× magnification, whereas Ciompi et al. (2019) created a dataset of lymphocyte counts in tissue with IHC staining at 40× magnification. This variation makes it difficult to aggregate datasets that enable simultaneous prediction, where it is necessary for task-level datasets to have the same input assumptions. As well as this, manual annotation makes it difficult to generate large-scale datasets for segmentation tasks in computational pathology. Iterative labelling strategies (Tavolaro et al., 2020; Graham et al., 2021a; Jaber et al., 2021), which use a human-in-the-loop mechanism, have recently been used in the literature and can help significantly decrease the time required for annotation, while still ensuring a high level of accuracy.

If we assume that all data irrespective of the task should be H&E stained, from colorectal tissue and at 20× objective magnification, then there exists a large amount of available data. Graham et al. (2021a) generated the largest dataset for nuclear instance segmentation and classification in CPath, consisting of 495,179 labelled nuclei from multiple sources. As part of the GlAS challenge, Sirinukunwattana et al. (2017) curated a dataset of 1,530 annotated gland boundaries originating from University Hospitals Coventry and Warwickshire (UHCW). This gland segmentation data was extended by Graham et al. (2019a) into the CRAG dataset, with an additional 3,054 annotated glands from UHCW. As well as segmentation data in colorectal tissue, there exists available patch-based classification data that enables the development of further supervised learning methods. For example, Kather et al. (2019) and Javed et al. (2020) both created large datasets for tissue type classification consisting of 100 thousand and 280 thousand patches, respectively. As a result of the number of existing available datasets, we use H&E stained images from colorectal tissue in all experiments within this paper.

3. Methods

In this section, we describe our proposed MTL-based Cerberus model that enables the simultaneous prediction of aligned tasks in CPath. We report how the learned feature representation can be used to help improve the performance of additional tasks with transfer learning. We then describe the data generation strategy to enable the collection of large datasets for segmentation tasks in CPath. This strategy is used to significantly increase the amount of data used in experiments performed in Section 4.3.

3.1. The Cerberus model architecture

Cerberus is a fully convolutional neural network with a shared encoder Φ and T independent decoders Ψ_t , which make a prediction for each task t . Utilising a shared encoder ensures that a general representation is learned, where each task can benefit from features learned from other tasks. This can be especially useful when certain tasks do not have access to a large amount of data.

We use a ResNet34 encoder (He et al., 2016) for ease of use of weights after training with MTL in downstream applications. For segmentation tasks, we use a U-Net (Ronneberger et al., 2015) style decoder that incrementally upsamples the features by a factor of 2. After each upsampling operation, we incorporate features from the encoder with skip connections, followed by 2 convolutions (3×3 kernel) with batch normalisation (Ioffe and Szegedy, 2015). This is

repeated until the features have the same spatial dimensions as the input. For patch classification tasks, we use global average pooling to reduce the features at the output of the encoder to a k dimensional vector, which is then followed by 2 fully connected (FC) layers. In particular, we set k to 256 and use Dropout (Srivastava et al., 2014) between each FC layer, with a dropout rate of 0.3. An overview of the architecture can be seen in Fig. 2a.

3.2. Cerberus training

Our overall MTL training strategy for Cerberus incorporates data from multiple independent sources and makes a simultaneous prediction at the output of the network. This training pipeline consists of the following steps: (i) sampling patches from each task, (ii) feeding patches through a convolutional neural network with a shared encoder and (iii) task-level loss aggregation.

3.2.1. Task sampler

Our overall dataset D consists of T task-dependent datasets D_t , which must be sampled appropriately during batch aggregation. This is because the network parameters are updated according to the data that is present in a particular batch. Careful consideration may be needed when there exist significant differences in the number of samples within each dataset D_t or if certain tasks are more challenging than others. More information on optimisation will be provided in Section 3.2.2.

When aggregating a batch, we need to ensure that all patches are the same size, so that they can be processed by the network. However, images for different tasks may have different spatial dimensions. It is possible to perform resizing or cropping to ensure that patches across all tasks have the same dimensions, but this may have an unwanted effect on the image quality or spatial context. Therefore, we group tasks into *super tasks*, which can contain multiple sub-tasks with the same input dimensions. If considering multiple *super tasks*, then one must be first randomly selected before batch aggregation. In our case, segmentation tasks (input dimensions 448×448) and tissue type classification (input dimensions 144×144) are treated as different *super tasks*, which are selected with probabilities of 0.7 and 0.3 respectively.

In this work, we investigate the effect of using two different batch aggregation methods: *mixed batch* and *fixed batch* sampling. A mixed batch contains patches from multiple tasks, whereas a fixed batch contains data from a single task. For each patch within a mixed batch, we first select a task-level dataset D_t from the entire multi-task dataset D with a probability of p_t . Once a task is selected, a patch is randomly chosen from the corresponding dataset. This two-step process is repeated n times, where n is the number of samples in the batch. Rather than selecting a random task for each sample in a batch, for a *fixed batch* we sample a task dataset D_t once per batch and then select n random samples from that given dataset. In all experiments, we set $p_t = \frac{1}{T}$.

In Fig. 2a, we show a simple example of the output of our task sampler, with a batch size of 3. Here, each coloured border denotes a separate task.

3.2.2. Dynamic training and loss aggregation

Previous work has shown that optimisation of MTL models can be difficult, especially when tasks are in conflict with each other (Lin et al., 2019). In CPath, this may involve using images from different tissue types or with different pixel resolution. This challenge can be exacerbated when each task has a different associated loss function, which makes balancing each task-level loss and consequently overall network optimisation quite challenging. To help ease the training of our model, we ensure that all considered tasks have the same input data assumptions and use a cross-entropy-based loss function to optimise the model. Cross-entropy loss is widely used for both segmentation and classification tasks and therefore is an appropriate choice if we want to use a similar loss function across different tasks.

When training the network, we use dynamic weight freezing that is dependent on how batches are aggregated by the task sampler. After passing a batch through the network, the weights in the encoder Φ are always updated irrespective of how the batch is sampled, whereas the weights of decoder Ψ_t are only updated when at least one example from task t is present.

Hereby we define the loss function when we jointly train Cerberus:

$$\mathcal{L} = \sum_{t \in [1, T]} \sum_{\rho \in D_t} \mathcal{L}_t(\{\Phi, \Psi_t\}, x_\rho, y_\rho) \quad (1)$$

In this equation, $\rho \in D_t$ denotes a sample patch ρ which belongs to the dataset D_t of the task t . Meanwhile, x_ρ and y_ρ are predictions generated by the decoder Ψ_t and the ground truth for the sample ρ respectively. In practice, the loss of each task t is computed by masking samples from other tasks with zeros. In other words, we multiply the resulting loss of all samples with a mask to set their subsequent gradient calculation to zero.

3.3. Cerberus for transfer learning

As a result of training Cerberus on a large amount of data, we can leverage the encoder to improve the performance of additional tasks with transfer learning. For example, the Cerberus feature representation can be used to enhance the performance of classification tasks or the model parameters can be used as pretrained weights for better model initialisation. The learned representation may also be used to extend the trained model and enable the subtyping of initially localised objects. To highlight the advantage of utilising Cerberus for transfer learning, we experiment with 3 different applications. First, we see whether the feature representation learned by Cerberus can improve both patch classification and object subtyping and then we see whether Cerberus can provide better pretrained weights for object detection.

An overview of each of the tasks mentioned above is provided in Fig. 2c-e. These are treated as single-task problems and therefore a simple batch aggregation strategy is used, rather than utilising our MTL task sampler, as described in Section 3.2.1.

3.3.1. Patch classification

To test the strength of the features learned by Cerberus, we use them for the task of subtyping small patches centred at individual objects, such as nuclei. In particular, we utilise the trained Cerberus encoder Φ and add several output layers to predict the corresponding category. For this purpose, we aggregate features from several positions in the encoder before applying global average pooling and two fully connected layers to obtain the prediction. We account for features at different scales within the encoder in order to demonstrate the feature strength at multiple points within the network. We also freeze the encoder weights and train the output layers in isolation to capture the contribution of the feature representation.

3.3.2. Object subtyping

To give further merit to the Cerberus feature representation, we show how it can be leveraged to sub-categorise each of the initially considered binary segmentation tasks. For this, we add an additional decoder Ψ_{T+1} to Cerberus, dedicated to the task of subtyping a particular output, such as the nuclei. Then, we train Ψ_{T+1} in isolation, while freezing both Φ and the decoders $\Psi_t \forall t \in [1, T]$. We treat this task as a pixel-based classification problem and therefore utilise the same decoder architecture used by the other segmentation tasks. After training the additional subtyping branch, Cerberus can predict all tasks simultaneously.

Training the subtyping branch in isolation means that we do not have to cater for balancing the different components of the loss function between each task. Therefore, we use a combination of cross-entropy and Dice loss, which can also help counter class imbalance in the data (Graham et al., 2019c). Also, we only calculate the loss within

foreground pixels by performing a masking operation on the loss function. We do this because the foreground pixels have already been localised by the associated binary segmentation branch and therefore differentiating between foreground and background is not needed. Concretely, the loss at the output of the subtyping branch Ψ_{T+1} for input X is defined as:

$$\mathcal{L}_{T+1}(X) = \mathcal{L}_\alpha(X) + \mathcal{L}_\beta(X), \quad (2)$$

where \mathcal{L}_α and \mathcal{L}_β refer to the masked cross-entropy and Dice losses respectively. Specifically, $\mathcal{L}_\alpha(X)$ and $\mathcal{L}_\beta(X)$ are defined as follows:

$$\mathcal{L}_\alpha(X) = - \sum_{k \in K} \frac{1}{|v_k|} \sum_{i \in v_k} y_{i,k}(X) \cdot \log \hat{y}_{i,k}(X) \quad (3)$$

$$\mathcal{L}_\beta(X) = \sum_{k \in K} 1 - \frac{2 \cdot \sum_{i \in v_k} (y_{i,k}(X) + \hat{y}_{i,k}(X)) + \epsilon}{\sum_{i \in v_k} y_{i,k}(X) + \sum_{i \in v_k} \hat{y}_{i,k}(X) + \epsilon}. \quad (4)$$

Here, K is the number of classes, y denotes the prediction and \hat{y} is the ground truth. Therefore, $y_{i,k}(X)$ corresponds to the prediction for input image X for class k and pixel i . v_k denotes the set of foreground pixels belonging to class k and ϵ is a smoothing parameter used in the Dice loss to avoid division by zero.

3.3.3. Object detection

Instead of adding additional decoders so that Cerberus performs an increasing number of tasks, we may want to utilise the weights in the trained Cerberus encoder for improved weight initialisation in a separate network. It is a common strategy to initialise networks with weights obtained from sampling from the normal distribution (Glorot and Bengio, 2010; He et al., 2015) or from initially training on large datasets, such as ImageNet. However, we demonstrate that utilising weights obtained from training Cerberus can provide better initialisation for certain tasks in CPath, leading to improved performance. To showcase this, we train RetinaNet (Lin et al., 2017), with our pretrained Cerberus encoder as the backbone, for the task of object detection. RetinaNet is a region proposal network (RPN) (Ren et al., 2015) that performs object detection by predicting the bounding boxes within an input image. RPNs are designed to efficiently predict region proposals with a wide range of scales and aspect ratios, which are then further processed to predict the location of object bounding boxes. For this experiment, our goal is to understand whether Cerberus can provide better initial weights, rather than isolating the contribution of the feature representation. Therefore, all weights in the network are optimised together.

3.4. Collecting large datasets for segmentation

Manually annotating datasets for deep segmentation models is a very time-consuming procedure. Many deep learning models can provide an accurate segmentation, especially in trivial areas with a distinct morphology. Therefore, it may be a waste of effort to still rely on manual annotation in areas where deep learning models can perform well. Therefore, to enable the collection of a large amount of accurately labelled data we employ a pathologist-in-the-loop annotation collection pipeline, which is a simplified version of the strategy implemented by Graham et al. (2021a). Our method consists of three main steps: (i) initial segmentation model training with manually annotated data, (ii) pathologist-in-the-loop refinement and retraining, (iii) final manual refinement.

First, we train an initial deep segmentation model on manually annotated data. This may be on existing available data if we are annotating objects with known publicly available datasets such as glands (Sirinukunwattana et al., 2017; Graham et al., 2019a). Otherwise, if annotating an object with no available dataset in a given tissue type, then a manually annotated dataset needs to be created.

After training the initial segmentation model, we process all data that we wish to annotate with the algorithm. With the input of a pathologist, we then extract image regions of size $1,000 \times 1,000$

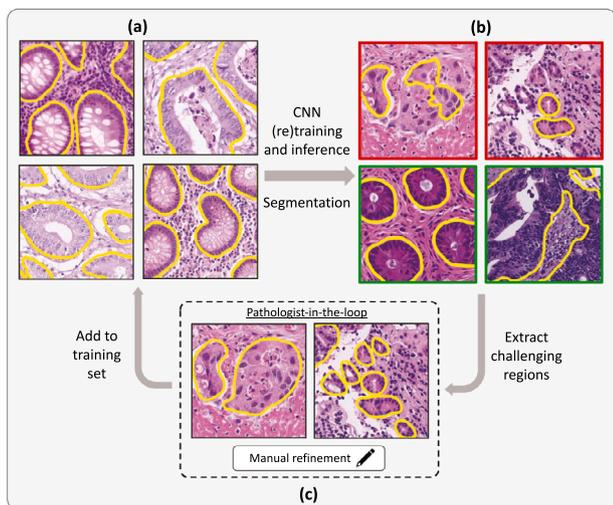


Fig. 3. Annotating datasets at scale for segmentation: (a) first train an initial segmentation model with available data, (b) process unlabelled data and extract difficult regions where the model does not perform well and (c) difficult regions are then refined and added to the training set before retraining the model. This sequence of steps is repeated until the data is accurately annotated. Difficult regions are shown with red borders in (b), otherwise they are shown with green borders.

where the algorithm does not perform well and refine the results. These images, along with the corresponding refined annotations, are then added to the initial training set before re-training the model. This *extract-refine-retrain* step is repeated until there is no noticeable improvement with the addition of further data. As a final step, we perform manual refinement of all data, which is then verified by a qualified pathologist. Segmented glands and lumina were presented for pathologist review when the data scientists were not sure about the visual quality of segmentation or noticed regions (glands or lumina) that could be potential false negatives. As a result, the number of actual glands/lumina reviewed and/or annotated by the pathologist was much less than the total number of glands/lumina. We provide an overview of our iterative labelling strategy in Fig. 3 and a summary of the number of collected annotations in Table 1.

4. Experiments and results

4.1. The datasets

In our MTL experiments, we utilise four task-level datasets. Specifically, we utilise a gland, lumen, nuclei and tissue type classification dataset, where each one contains image regions taken from H&E stained colon tissue samples at 20× objective magnification (around 0.5 μm/pixel).

Gland dataset: GlaS (Sirinukunwattana et al., 2017) and CRAG (Graham et al., 2019a) are two of the most widely used available datasets for gland instance segmentation in colon histology images, containing a total of 1,602 and 3,209 annotated glandular boundaries, respectively. Despite GlaS and CRAG proving instrumental in the development of the recent automated gland segmentation models (Chen et al., 2017; Graham et al., 2019a), both datasets are from a single centre (UHCW) and therefore models trained on these may not generalise well to external data. Therefore, in addition, we annotate glands from two further datasets using our iterative pathologist-in-the-loop annotation method described in Section 3.4. Specifically, we label 46,346 additional glands within images extracted from the DigestPath and TCGA datasets, respectively.

Lumen dataset: There has been recent work on the collection of large annotated datasets for gland segmentation within colon tissue,

Table 1

Overview of the number of annotated objects used for segmentation tasks in this paper.

	Gland	Lumen	Nuclei
Existing data	4,811	–	495,179
Generated data	46,346	56,358	–
Total	51,157	56,358	495,179

but there are not many works on lumen segmentation, which can be beneficial for diagnosing certain conditions such as cribriform architecture and serrated polyps. To combat this, we annotate lumina within images taken from the GlaS, CRAG, DigestPath and TCGA datasets using our iterative labelling approach and accordingly use the data in associated experiments. For this, we manually label 6,863 lumen regions in GlaS and CRAG images as our initial annotated dataset and then use our iterative scheme to generate 49,495 lumen boundaries in the DigestPath and TCGA datasets.

Nuclei dataset: We use the Lizard dataset (Graham et al., 2021a,b), which is the largest known instance segmentation and classification dataset in CPath containing nearly half a million labelled nuclei. Each nucleus is labelled according to the cell that it belongs to. Specifically, the labels used in the dataset are epithelial, lymphocyte, plasma, neutrophil, eosinophil and connective tissue. Here, the connective tissue is a broader category that includes endothelial cells, fibroblasts and muscle cells. When creating the dataset, image regions are extracted from the following 6 sources: CRAG (Graham et al., 2019a), GlaS (Sirinukunwattana et al., 2017), CoNSeP (Graham et al., 2019c), PanNuke (Gamper et al., 2020a), DigestPath (Da et al., 2022) and TCGA. Therefore, Lizard is a diverse dataset and models trained on it may likely generalise to unseen examples.

Tissue type dataset: We utilise 3 datasets for tissue type classification in colorectal images, that we refer to as CRC-TP2, Kather5K (Kather et al., 2016) and Kather100K (Kather et al., 2019), which contain 100,000, 5,000 and 333,362 patches respectively. Here, CRC-TP2 contains image patches from UHCW, whereas Kather5K and Kather100K contain patches from the pathology archive at the University Medical Center Mannheim. CRC-TP2 is an extension of the dataset by Javed et al. (2020), where we add an extra 53,362 patches from an internal colon biopsy dataset (Wahab et al., 2021). Overall, the tissue types present in the dataset are background, adipose (fat), abnormal epithelium, normal epithelium, inflammatory, stroma, muscle, debris and mucous. Example image patches for each class can be seen in Fig. 4.

Data split: For segmentation tasks, we use TCGA data as an external test set and split the remaining data into 3 folds for cross-validation. For tissue type classification, we do not use an external test set but similarly split the data into 3 folds. When creating each split, we separate the data on a patient level to ensure that test data remains completely unseen. Then, for a given fold we assign each split as either a training, validation or testing subset. We also make sure that no patient overlap occurs between the subsets. For example, if a patch from fold 1 of the nuclei dataset belongs to a given patient, then that patient should not exist in any of the other datasets in folds 2 or 3.

4.2. Evaluation metrics

Segmentation: We utilise the Dice score (Dice, 1945) and Panoptic Quality (PQ) (Kirillov et al., 2019) to measure the binary segmentation performance. Here, the Dice score measures how well a model can separate foreground pixels from the background and is defined as:

$$\text{Dice} = \frac{2 \times (|Y \cap \hat{Y}|)}{|Y| + |\hat{Y}|}, \quad (5)$$

where Y and \hat{Y} are the prediction and ground truth maps respectively. However, the Dice score does not indicate whether neighbouring objects are correctly separated. Therefore, we also use PQ to provide a measure for the quality of instance segmentation. PQ matches

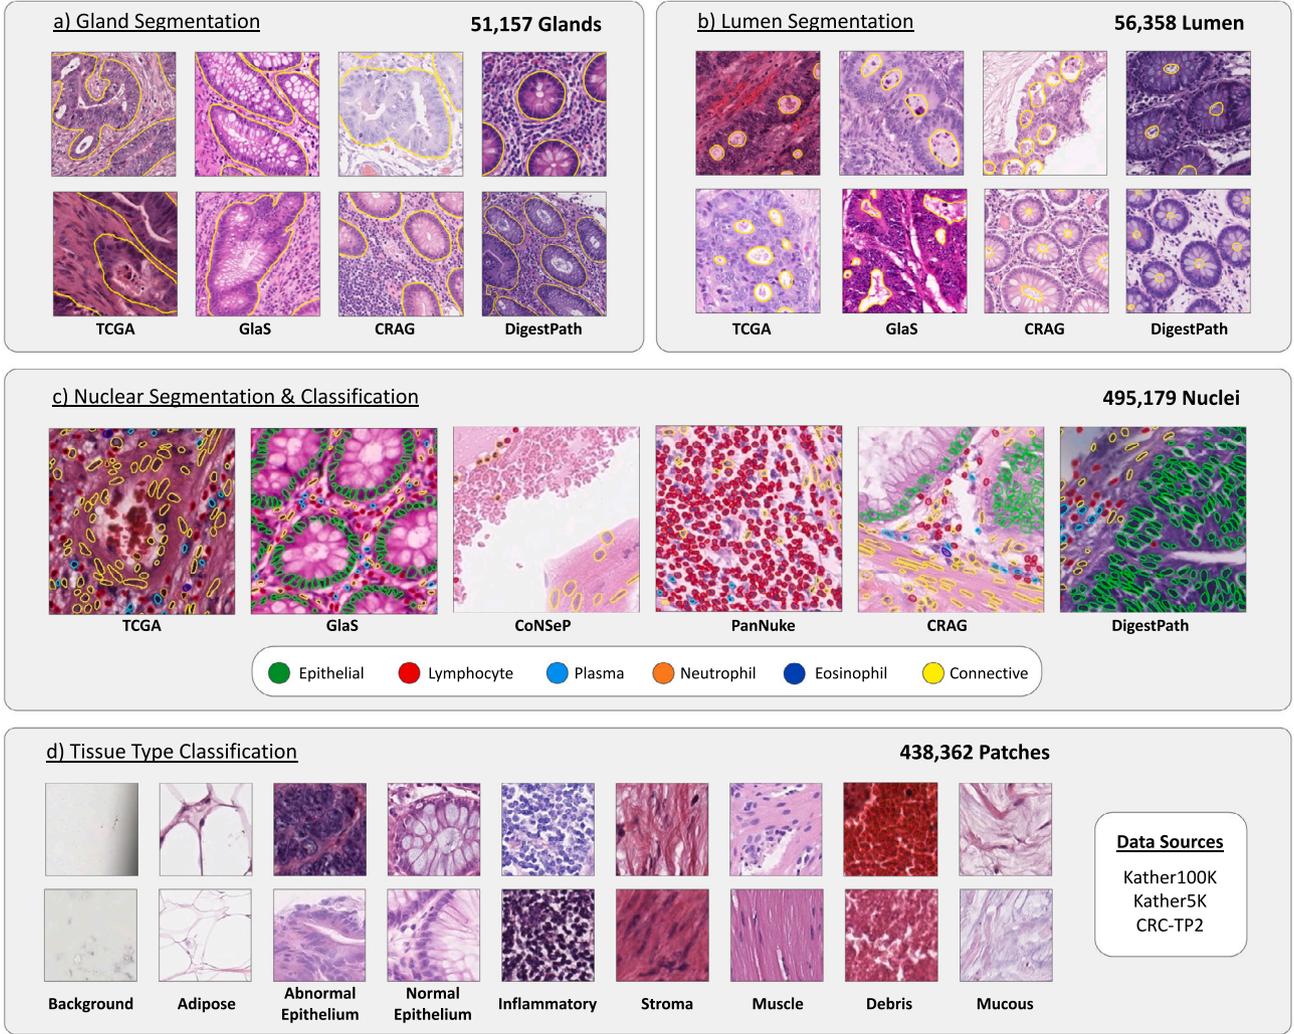


Fig. 4. Overview of the datasets used for (a) gland segmentation, (b) lumen segmentation, (c) nuclear segmentation & classification and (d) tissue type classification. We also show the origin of the image regions used for each task.

predicted objects with ground truth objects based on whether their intersection-over-union (IoU) exceeds 0.5. This matching criterion is used to separate objects into true positives (TP), false positives (FP) and false negatives (FN). Then, PQ is defined as:

$$PQ = \underbrace{\frac{|TP|}{|TP| + \frac{1}{2}|FP| + \frac{1}{2}|FN|}}_{\text{Detection Quality(DQ)}} \times \underbrace{\frac{\sum_{(y,\hat{y}) \in TP} IoU(y,\hat{y})}{|TP|}}_{\text{Segmentation Quality(SQ)}}, \quad (6)$$

where \hat{y} denotes a ground truth instance and y denotes a predicted instance. We calculate PQ individually for each image and then average the results to get the overall statistic.

When assessing multi-class instance segmentation, we calculate PQ independently for each class and then report the average score. In previous work (Gamper et al., 2020a; Graham et al., 2021a), mPQ was defined by calculating PQ for each image and for each class before the overall results were averaged. However, the result of a particular image was not considered in the class average if that class was not present in the ground truth. To help avoid skipping images, we also report mPQ^+ (Graham et al., 2021b), which aggregates the statistics over all images before taking the class average.

Classification: To assess the classification performance, we utilise mean average precision (mAP) and mean F_1 (mF_1). For both metrics, either average precision (AP) or F_1 is calculated and the average over the classes is reported to give mAP and mF_1 . AP and F_1 are defined as:

$$Pr_n = \frac{TP_n}{TP_n + FP_n} \quad (7)$$

$$Re_n = \frac{TP_n}{TP_n + FN_n} \quad (8)$$

$$AP = \sum_n (Re_n - Re_{n-1}) Pr_n \quad (9)$$

$$F_1^n = \frac{Pr_n * Re_n}{Pr_n + Re_n} \quad (9)$$

Upon using a threshold n to convert a probability to a binary label, the predictions and ground truth are split into true positive (TP), false positive (FP) and false negative (FN) sets. The number of items in each of these sets is then utilised to calculate the precision (Pr) and recall (Re) values. Depending on the use case, n may have a different meaning. For AP , n denotes the n th threshold level that is associated with a specific threshold value. For $n-1$, this extends to be the previous threshold level that is smaller than n . However, for F_1 , we set the prediction label to be the one with the highest probability, rather than using a specific threshold value.

4.3. Results

4.3.1. Cerberus vs state-of-the-art approaches

First, we compare the segmentation performance of Cerberus with recent state-of-the-art approaches in CPath and provide the results in

Table 2

Simultaneous nuclear segmentation and classification results of our proposed approach compared to recent state-of-the-art methods. Cerberus predicts the eroded instance target, whereas Cerberus⁺ also predicts the object boundary.

Model	Cross validation			External test		
	Binary PQ	mPQ	mPQ ⁺	Binary PQ	mPQ	mPQ ⁺
U-Net	0.564 ± 0.020	0.275 ± 0.014	0.324 ± 0.018	0.493 ± 0.051	0.235 ± 0.018	0.258 ± 0.034
HoVer-Net	0.583 ± 0.014	0.295 ± 0.018	0.409 ± 0.027	0.514 ± 0.026	0.285 ± 0.018	0.335 ± 0.017
Cerberus	0.588 ± 0.008	0.323 ± 0.020	0.396 ± 0.022	0.560 ± 0.028	0.308 ± 0.015	0.353 ± 0.009
Cerberus ⁺	0.612 ± 0.010	0.358 ± 0.011	0.425 ± 0.019	0.568 ± 0.009	0.332 ± 0.011	0.388 ± 0.003

Table 3

Gland and lumen segmentation results of our proposed approach compared to recent top-performing methods. Cerberus predicts the eroded instance target, whereas Cerberus⁺ also predicts the object boundary.

Model	Gland		Lumen	
	Cross val	External test	Cross val	External test
U-Net	0.622 ± 0.030	0.459 ± 0.024	0.501 ± 0.008	0.240 ± 0.109
MILD-Net	0.647 ± 0.046	0.526 ± 0.027	0.522 ± 0.034	0.353 ± 0.088
Rota-Net	0.662 ± 0.048	0.573 ± 0.034	0.569 ± 0.042	0.436 ± 0.054
Cerberus	0.677 ± 0.028	0.640 ± 0.012	0.589 ± 0.006	0.525 ± 0.027
Cerberus ⁺	0.674 ± 0.021	0.650 ± 0.004	0.590 ± 0.018	0.530 ± 0.005

Tables 2 and 3. Here, we only consider MTL with mixed batch sampling, due to its superior performance as shown in Table 4. For this, we assess the binary segmentation performance of nuclei, glands and lumina using PQ and then quantify the multi-class segmentation performance of nuclei using both mPQ and mPQ⁺.

For nuclei, we compare the performance of our proposed approach with U-Net (Ronneberger et al., 2015) and HoVer-Net (Graham et al., 2019c). HoVer-Net is a recent top-performing method for combined nuclear instance segmentation and classification (Verma et al., 2020; Gamper et al., 2020a) and therefore serves as a strong benchmark. To enable full comparison amongst all models for nuclear segmentation and classification, we add extra layers to the output of U-Net for nuclear subtyping. For both gland and lumen segmentation, we compare the performance of Cerberus with U-Net, MILD-Net (Graham et al., 2019a) and Rota-Net (Graham et al., 2019b). MILD-Net and Rota-Net are competitive methods for both of these tasks and therefore serve as strong baselines. Here, Rota-Net segments the glands and lumina simultaneously, like in the original publication. All Cerberus and U-Net models use a simple 2-class eroded instance target, whereas MILD-Net and Rota-Net also incorporate the object boundary as an additional class. HoVer-Net uses a more complex instance segmentation target, consisting of horizontal and vertical distance maps. Therefore, if Cerberus can outperform competing methods using a simple instance segmentation target, then this holds great promise for achieving even better results when adding more complex targets. As an extra experiment, we see whether we can further boost the performance of Cerberus by also incorporating the object boundary. We denote this as Cerberus⁺ in Tables 2 and 3. For all U-Net models, we use batch normalisation, which was not used in the original publication (Ronneberger et al., 2015).

In Table 2, we see that Cerberus outperforms both U-Net and HoVer-Net on nearly all measures for both binary and multi-class nuclear instance segmentation. This performance is further improved when using the object boundary as an additional target of the network and leads to significantly better performance for all measures compared to both U-Net and HoVer-Net. It may be worth noting that we report slightly different results than the HoVer-Net results reported by Graham et al. (2021a) because we ensure that the exact same input patches are used for a fair comparison. It is interesting to see that the boost in performance is more pronounced for the external test set, suggesting that MTL can help with generalisation. For gland and lumen segmentation, we observe from Table 3 that Cerberus exceeds the performance of U-Net, MILD-Net and Rota-Net on all measures. Again, this is especially prominent on the external test set. Despite the increase in performance not being large, we see a slight increase in

performance when considering the object boundary, especially on the completely unseen data source. Therefore, this justifies predicting the object boundary for instance segmentation in all segmentation tasks considered by our network.

4.3.2. Multi-task learning vs single-task learning for segmentation

We perform extensive experiments to fully understand the contribution of multi-task learning for segmentation. For this experiment, we only consider the binary output for segmentation tasks and consider the eroded map as the instance target for simplicity. We report results for further subtyping of the segmentation output in Section 4.3.3. In particular, we compare single-task learning (STL) methods, where a devoted network is trained individually per task, with multi-task networks, where tasks are optimised simultaneously. We also experiment with two different sampling strategies for MTL: (i) sampling a single task per batch and (ii) sampling a mixture of tasks per batch, referred to as *fixed* and *mixed* batches, respectively in 3.2.1. We investigate the above strategies using encoder-decoder based model architectures with hard parameter sharing. Specifically, we experiment with both a basic U-Net style encoder (Ronneberger et al., 2015) with batch normalisation and a ResNet34 encoder (He et al., 2016). We assess the performance of multiple models to remove any bias from the model architecture and instead solely focus on the contribution of the training scheme. After gaining insight into the effect of MTL for segmentation tasks, we investigate whether ImageNet pretrained weights and utilising patch classification as an auxiliary task can further increase the segmentation performance. As described in Section 3.2.1, a batch must contain a single *super task*, which needs to be initially selected if considering both segmentation and patch classification.

Our major aim is that we want the performance of networks trained using MTL to be **at least as good** as those trained using STL. This would reduce the computational requirements during training and inference because just a single network could be used to perform tasks simultaneously. We can see from Table 4 that in fact we achieve superior performance for most tasks considered during our experiments. In particular, we observe that both U-Net and ResNet34 alternatives benefit from MTL, especially for gland and lumen segmentation. This observation is more pronounced on the external test set, suggesting that the features learned during MTL can help with generalisation. The greatest improvement from training with MTL is observed for lumen segmentation, especially on the external test set. This can potentially be attributed to learning complementary information regarding the location of glandular boundaries. Despite not observing an obvious increase in performance for nuclear segmentation, results are comparable, suggesting that we do not require a separate network to perform

Table 4

Comparison of results obtained for binary segmentation using MTL and STL. Sampler indicates how patches are aggregated into batches. A *Fixed* batch contains patches from a single task and a *Mixed* batch contains patches from multiple tasks (Section 3.2.1). IN refers to ImageNet pretrained models and PC denotes patch classification. When using MTL with PC, a batch contains either only patches for segmentation or PC (Section 3.2.1).

(a) Cross validation results.								
	Task	Sampler	Nuclei		Gland		Lumen	
			Dice	PQ	Dice	PQ	Dice	PQ
U-Net	STL	Fixed	0.767 ± 0.013	0.566 ± 0.021	0.891 ± 0.013	0.622 ± 0.030	0.659 ± 0.014	0.501 ± 0.008
	MTL	Fixed	0.759 ± 0.004	0.560 ± 0.006	0.894 ± 0.006	0.636 ± 0.025	0.688 ± 0.017	0.536 ± 0.022
	MTL	Mixed	0.762 ± 0.007	0.558 ± 0.011	0.897 ± 0.011	0.628 ± 0.015	0.697 ± 0.003	0.543 ± 0.011
ResNet-34	STL	Fixed	0.772 ± 0.002	0.581 ± 0.006	0.902 ± 0.007	0.669 ± 0.014	0.666 ± 0.060	0.508 ± 0.091
	MTL	Fixed	0.769 ± 0.007	0.572 ± 0.009	0.903 ± 0.001	0.677 ± 0.022	0.726 ± 0.015	0.585 ± 0.013
	MTL	Mixed	0.770 ± 0.003	0.574 ± 0.005	0.905 ± 0.005	0.674 ± 0.016	0.721 ± 0.018	0.590 ± 0.014
	IN-MTL	Mixed	0.779 ± 0.004	0.588 ± 0.009	0.907 ± 0.002	0.683 ± 0.013	0.733 ± 0.019	0.592 ± 0.019
	IN-MTL+PC	Mixed	0.778 ± 0.006	0.589 ± 0.008	0.909 ± 0.008	0.677 ± 0.028	0.735 ± 0.007	0.589 ± 0.006

(b) External test results.								
	Task	Sampler	Nuclei		Gland		Lumen	
			Dice	PQ	Dice	PQ	Dice	PQ
U-Net	STL	Fixed	0.713 ± 0.067	0.499 ± 0.058	0.741 ± 0.012	0.459 ± 0.024	0.281 ± 0.089	0.240 ± 0.109
	MTL	Fixed	0.708 ± 0.061	0.486 ± 0.058	0.770 ± 0.100	0.471 ± 0.077	0.420 ± 0.051	0.376 ± 0.063
	MTL	Mixed	0.738 ± 0.017	0.505 ± 0.017	0.824 ± 0.061	0.492 ± 0.039	0.446 ± 0.063	0.377 ± 0.059
ResNet-34	STL	Fixed	0.758 ± 0.021	0.554 ± 0.030	0.820 ± 0.104	0.566 ± 0.082	0.388 ± 0.016	0.313 ± 0.041
	MTL	Fixed	0.739 ± 0.032	0.527 ± 0.042	0.864 ± 0.031	0.602 ± 0.020	0.594 ± 0.046	0.470 ± 0.010
	MTL	Mixed	0.742 ± 0.027	0.531 ± 0.038	0.881 ± 0.013	0.617 ± 0.004	0.622 ± 0.053	0.492 ± 0.013
	IN-MTL	Mixed	0.773 ± 0.031	0.560 ± 0.037	0.875 ± 0.072	0.619 ± 0.064	0.637 ± 0.057	0.516 ± 0.056
	IN-MTL+PC	Mixed	0.774 ± 0.024	0.560 ± 0.028	0.908 ± 0.010	0.640 ± 0.012	0.666 ± 0.014	0.525 ± 0.027

Table 5

Comparison of performance with models trained using multi-task and single-task learning for tissue type classification.

Task	mAP	mF ₁
IN-STL	0.949 ± 0.002	0.885 ± 0.001
IN-MTL+PC	0.948 ± 0.006	0.883 ± 0.007

this task. We suspect that this may be because we are already using a very large dataset for nuclear segmentation and so the addition of complementary tasks may not be particularly advantageous.

Although we achieve similar performance during cross-validation, it can be seen that utilising a mixed batch leads to significantly better performance on the external test set. Therefore, for the remainder of the experiments in this paper, we only consider MTL with mixed batch sampling. We also hypothesise that mixed batch sampling will scale better when using a large number of tasks because for a given batch, there is a greater probability of at least one sample being present for a particular task. However, this assumption only holds if the batch size is greater than the number of tasks. Further investigation into the relationship between batch types (mixed/fixed) and batch size may be required in future work to fully understand what kind of batches work best for MTL.

For all tasks, we observe that using a ResNet encoder results in a significantly better performance compared to using a conventional U-Net approach. We also observe that initialising the ResNet encoder with ImageNet weights, rather than randomly, gives a further boost in the performance. From Table 4 we can see that with the addition of patch classification, results are comparable for cross-validation, but have improved on the external test set. This suggests that utilising the additional data from the patch classification task leads to the shared encoder learning more representative features, and hence leads to better performance on external data. Of course, as an additional benefit, the network can predict the tissue type at the output of the network, which can potentially lead to more powerful downstream pipelines. It must be noted that results using ImageNet-pretrained MTL models were obtained to assess potential performance gains over MTL. Therefore, they should not be compared directly with STL models. We report

the patch classification performance in Table 5, where we observe that the multi-task and single-task models achieve similar performance. Therefore, performance is not compromised when considering multiple tasks alongside patch classification. For this experiment, we compare our approach with STL pretrained with ImageNet to provide a fair comparison.

4.3.3. Cerberus feature representation for patch-based nuclear classification

After training Cerberus, our next experiment involves understanding whether the learned features can be used to improve the performance of patch-level nuclear classification. We first assess whether the representation captured from training Cerberus is superior to that obtained from training on ImageNet in a supervised setting. Using ImageNet-trained networks is a common strategy to obtain a deep representation in the field of computer vision and therefore serves as a good benchmark for this experiment. Then, we compare single-task and self-supervised learning alternatives with histology-specific data. In particular, contrastive learning is a self-supervised approach that has shown great promise, without requiring labelled data. Therefore, we compare Cerberus features with those obtained using a recent contrastive learning method called SimCLR (Chen et al., 2020) on our combined gland, lumen and nuclei datasets. However, this may not fully unlock the potential of self-supervised methods that are only bounded by the amount of available image data. Therefore, we also add results using SimCLR trained with over 2 million H&E colon input patches at 20× objective magnification and denote it as SimCLR⁺. For this, in addition to the data that was used for MTL, we sampled many image patches from the PAIP challenge² and TCGA datasets. As an additional benchmark, we also utilise a model provided by Ciga et al. (2022) that was optimised by training SimCLR on 57 histology datasets.

In Table 6, we show results for nuclear classification using different feature representations. Each representation is obtained by passing images through a pretrained ResNet encoder, where weights were initially obtained using one of the strategies described above. Features are passed to several trainable output layers to predict the nuclear category.

² The PAIP dataset is available to download here: <https://paip2020.grand-challenge.org/>

Table 6

Comparison of different deep features for nuclear classification. All models consider a patch of size 32×32 extracted at the centre of each nucleus. We only consider MTL with *mixed batch* sampling due to its superior performance in Table 4. For all models, we freeze the ResNet weights and only train the classification layers. Models above the dashed line use a ResNet34 backbone, whereas models below the dashed line use ResNet18 to enable cross-comparison with Ciga et al. (2022), which uses SimCLR. The difference between this and SimCLR/SimCLR+ reported on rows 2 and 3 is the data used for training. Here, SimCLR uses the same data used during MTL and SimCLR+ uses over 2 million colon H&E image patches from many datasets.

Features	Mode	Cross validation		External test	
		mAP	mF ₁	mAP	mF ₁
IN	Patch	0.535 ± 0.018	0.491 ± 0.026	0.385 ± 0.013	0.355 ± 0.020
SimCLR	Patch	0.508 ± 0.005	0.464 ± 0.014	0.432 ± 0.029	0.403 ± 0.026
SimCLR+	Patch	0.592 ± 0.012	0.557 ± 0.012	0.488 ± 0.023	0.442 ± 0.009
STL Nuclei	Patch	0.635 ± 0.033	0.588 ± 0.029	0.572 ± 0.022	0.538 ± 0.020
STL Gland	Patch	0.542 ± 0.027	0.499 ± 0.041	0.495 ± 0.042	0.456 ± 0.056
STL Lumen	Patch	0.520 ± 0.022	0.478 ± 0.038	0.452 ± 0.059	0.418 ± 0.073
MTL	Patch	0.587 ± 0.032	0.543 ± 0.039	0.558 ± 0.005	0.506 ± 0.010
IN-MTL	Patch	0.607 ± 0.033	0.562 ± 0.038	0.540 ± 0.035	0.492 ± 0.041
IN-MTL+PC	Patch	0.651 ± 0.067	0.601 ± 0.049	0.573 ± 0.013	0.540 ± 0.019
IN-MTL+PC*	Patch	0.628 ± 0.030	0.577 ± 0.038	0.586 ± 0.024	0.518 ± 0.034
(Ciga et al., 2022)	Patch	0.615 ± 0.004	0.566 ± 0.019	0.540 ± 0.035	0.492 ± 0.041

Table 7

Results for nuclear classification with different loss strategies by making a prediction per pixel. Pixel_A performs nuclear classification by adding an extra classification layer at the output of the nuclear segmentation branch, whereas Pixel_B uses a devoted upsampling branch.

Features	Mode	Loss	Cross validation		External test	
			mAP	mF ₁	mAP	mF ₁
IN-MTL+PC	Pixel _A	CE	0.493 ± 0.021	0.477 ± 0.013	0.447 ± 0.025	0.432 ± 0.012
IN-MTL+PC	Pixel _A	CE & Dice	0.537 ± 0.017	0.545 ± 0.015	0.494 ± 0.030	0.497 ± 0.012
IN-MTL+PC	Pixel _A	Masked CE & Dice	0.576 ± 0.012	0.566 ± 0.008	0.529 ± 0.032	0.509 ± 0.020
IN-MTL+PC	Pixel _B	Masked CE & Dice	0.678 ± 0.035	0.710 ± 0.030	0.634 ± 0.013	0.669 ± 0.003

We aggregate features from several positions of the encoder before applying global average pooling and two fully connected layers to obtain the final prediction. For most entries in Table 6 above the dashed line, we utilise a ResNet34 model for feature extraction. However, Ciga et al. (2022) does not have this model available and instead provides a ResNet18 alternative. Therefore, to enable a like-for-like comparison, we also retrain Cerberus with a ResNet18 backbone and compare the corresponding results.

We extract a dataset of patches of size 32×32 pixels centred at each nucleus within our nuclear segmentation dataset and use the same splits that were initially used for feature extraction to ensure that all test data remains completely unseen. We observe from Table 6 that all features obtained using single or multi-task learning achieve a better performance than ImageNet-derived features for nuclear subtyping. They are also superior to those features obtained from training with SimCLR on our gland, lumen and nuclei datasets. Using MTL trained on segmentation tasks leads to stronger features as compared to single-task approaches for gland and lumen segmentation. However, using features from MTL does not surpass the performance of single task nuclear segmentation features. This may be unsurprising because the features learned for nuclear segmentation are most probably directly related to this particular task of nuclear subtyping. However, when adding patch classification as an auxiliary task during MTL, we achieve the best overall performance. MTL with auxiliary patch classification (using a ResNet18 encoder) also obtains superior features for nuclear classification as compared to the model supplied by Ciga et al. (2022).

4.3.4. Cerberus feature representation for object subtyping

Cerberus is trained for binary segmentation and multi-class patch classification. However, it may be desirable to further subtype the segmentation output. For example, the output of the gland segmentation branch could be further split into benign and malignant glands, whereas the nuclear segmentation output could be split into different types of nuclei. As described in Section 3.3.2, this can be done by adding a dedicated decoder to Cerberus, which performs pixel-wise classification within objects for a particular task. Majority voting can

then be used to assign a single category to each object. When utilising this strategy, the original Cerberus weights are frozen and only the additional decoder is trained. Therefore, the original model parameters remain unchanged and so performance is not compromised. Because the weights are frozen, it is essential for the encoder to have initially learned a strong feature representation to perform well on this task. In Section 3.3.2, we also describe our proposed loss function that focuses only on foreground pixels and aims to combat class imbalance.

For this experiment, we explore whether the strategy described above can be used to improve the performance achieved in Table 6. The added benefit of extending Cerberus in this way is that it enables subtyping to be done jointly with other tasks and therefore a single network can be used. On the other hand, if our patch classification pipeline for nuclear classification was used, then an additional network would be needed. To give justification for our chosen subtyping network architecture and loss function, we provide an ablation study in Table 7.

Rather than having an entirely separate branch for subtyping, we initially split from the binary nuclear segmentation branch after the final upsampling operation to leverage already existing features and save on computational cost. Using this strategy, we experiment with 3 different loss configurations. We first use cross-entropy (CE) and then combine CE with Dice loss to see whether this can help counter the class imbalance present in the dataset. We calculate the loss for *all* pixels, which is not needed if we already perform object localisation in the segmentation branch. Therefore, we also experiment with a masked loss that computes the loss only in foreground pixels. After analysing the best loss configuration, we then see whether utilising the devoted upsampling branch can further boost the performance.

In Table 7, we observe that adding Dice loss to CE loss leads to an improvement in performance for all measures. There is also a further boost in the performance when using masked loss, which helps the network to learn discriminative features within the nuclei, rather than learning to separate foreground from background. However, it is evident from Table 6 that the performance is still inferior to the patch-classification counterpart that uses the same features. When adding the

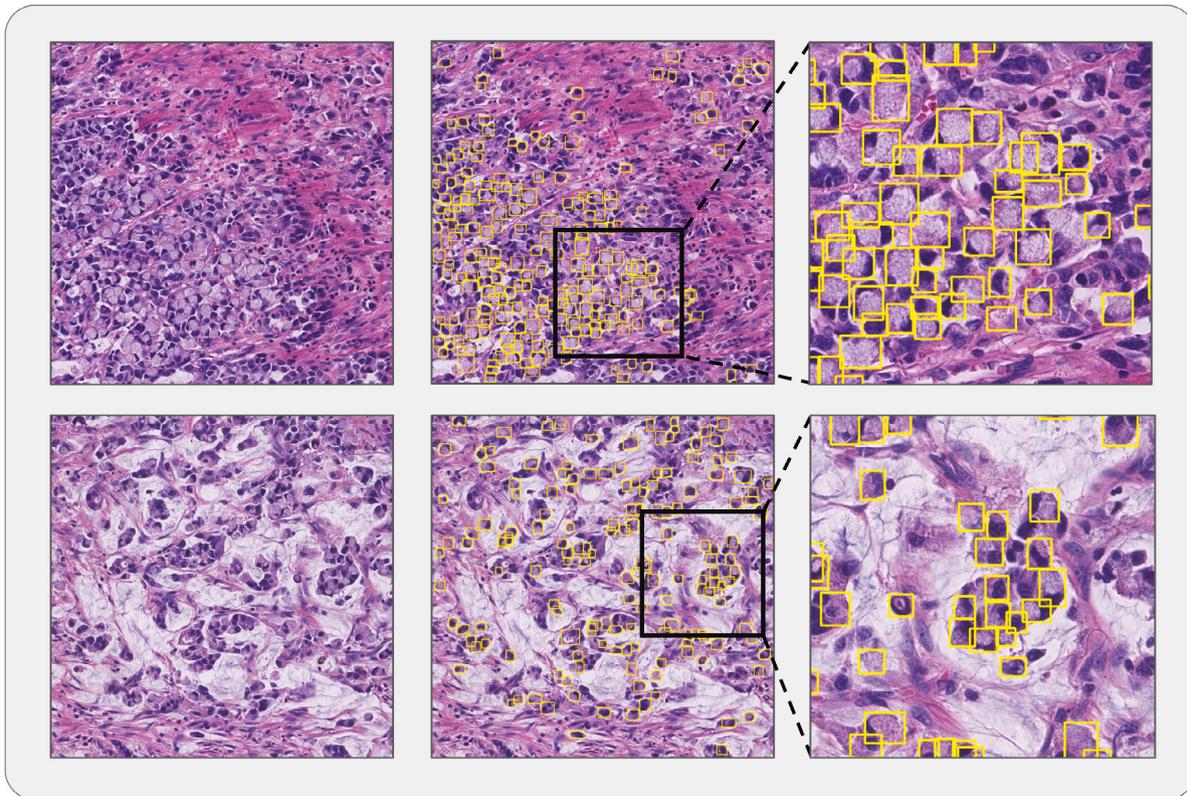


Fig. 5. Signet ring cell detection visual results using the pretrained Cerberus encoder.

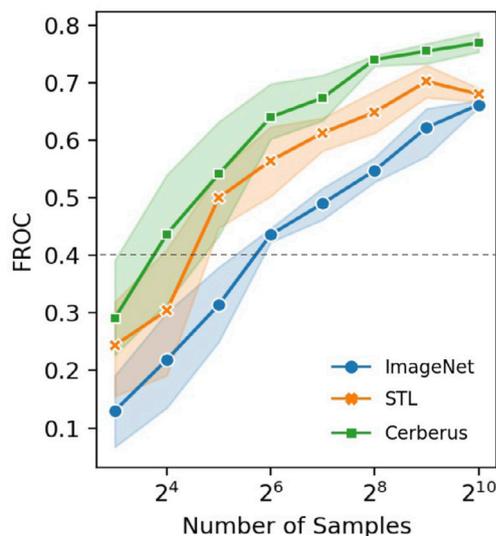


Fig. 6. Signet ring cell detection performance with different numbers of input samples. We compare Cerberus features with STL (trained on nuclei) and ImageNet-based features. *FROC*: free-response receiver operating characteristic.

devoted upsampling branch, we observe a significant increase in performance, that even surpasses all patch-based alternatives. Therefore, it is clear that the use of a full branch for subtyping is required to ensure a strong performance.

4.3.5. Cerberus pretrained weights for object detection

To understand whether our trained multi-task model leads to better weight initialisation for new tasks, we assess whether utilising the Cerberus pretrained weights can help boost the performance of

signet ring cell detection. For this, we utilise data provided as part of the DigestPath challenge, but perform additional labelling due to incomplete annotations. Here, the images with signet ring cell annotations are independent to DigestPath images used during MTL. We also want to understand whether leveraging Cerberus for enhanced weight initialisation can lead to improved data efficiency on novel tasks. For this, we train multiple signet ring cell detection models with increasing numbers of random input samples and assess the impact on performance. Specifically, we train RetinaNet (Lin et al., 2017) with a ResNet backbone to predict the bounding box around each signet ring cell. Here, the backbone is initialised with parameters either obtained from training on ImageNet, from using STL on nuclei histology data (STL-Nuclei from Table 4) or from training Cerberus. As opposed to our nuclei subtyping experiment, we train *all* model weights after initialisation. Example visual results can be seen in Fig. 5.

Results from training with the different weight initialisation strategies are displayed in Fig. 6, where the performance is measured in terms of area under the free-response receiver characteristic (*FROC*) curve. We choose this metric as it was used as part of the DigestPath challenge to evaluate different submissions. We repeat each experiment 4 times and therefore the plot shows the mean result along with the calculated confidence intervals. It is important to note that the *x*-axis in the plot is in log-scale. We can see that leveraging the Cerberus encoder leads to an improvement in results over ImageNet and STL pretrained weights, irrespective of the number of input samples. When we are presented with a small number of samples, then using our Cerberus encoder as a pretraining mechanism can lead to greater data efficiency. We make this clear by drawing a horizontal dashed line on the plot denoting a score of 0.4. It can be seen that much fewer samples are needed to achieve this score when using Cerberus pretrained weights. This holds potential for other tasks when not provided with a large amount of data.

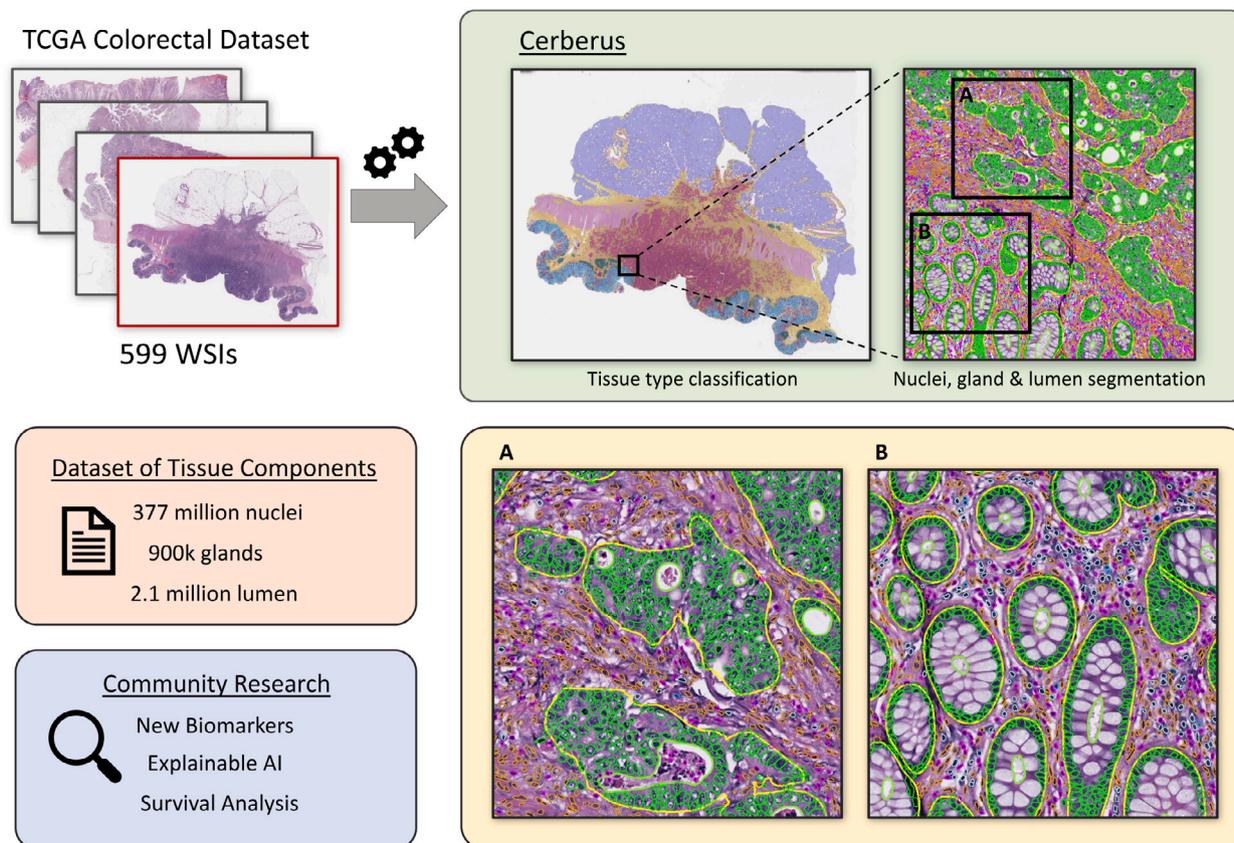


Fig. 7. Overview of the pipeline for processing colorectal TCGA WSIs with Cerberus. In total, we process 599 WSIs and obtain 377 million different nuclei, 900 thousand glands and 2.1 million lumina. We release the processed dataset to enable research in the CPath community. We provide zoomed-in image regions of the corresponding boxes denoted by A and B.

4.3.6. Processing the TCGA colorectal database

A major bottleneck in the development of explainable models for CPath is the localisation of various regions in the tissue. This enables the subsequent extraction of interpretable features which can then be used in downstream predictive pipelines for tasks such as biomarker discovery and survival analysis. Therefore, to accelerate research in CPath, we process all formalin-fixed paraffin-embedded colorectal WSIs from TCGA with Cerberus and make the results publicly available. We provide results for gland, lumen and nuclear segmentation, as well as patch-level tissue type classification. We also utilise our subtyping strategy (Section 3.3.2) and predict the category for each nucleus, which can enable effective modelling of the tumour micro-environment.

In total, we localise **377 million** different types of nuclei, **900 thousand** glands and **2.1 million** lumina,³ making it by far the largest available resource of its kind in CPath. In Fig. 7, we display some example results obtained from processing sample TCGA WSIs with the proposed approach.

4.4. Implementation and training details

We implemented our framework with PyTorch (Paszke et al., 2019) version 1.9 and utilised 3 NVIDIA V100 GPUs, each with 32 GB RAM. During training, we applied the following random data augmentations: flip, rotation, Gaussian blur, median blur and colour perturbation. For segmentation tasks, networks were fed input patches of size 448×448 pixels, whereas for patch classification tasks the input patch size was 144×144 . We trained all networks with a batch size of 9 samples

³ This dataset is available to download at: <https://github.com/TissueImageAnalytics/cerberus>.

per GPU (a total of 27 samples with 3 GPUs). All models were trained for 90,000 steps using Adam optimisation (Kingma and Ba, 2014) with an initial learning rate of 10^{-3} , which was reduced to 10^{-4} after 70,000 steps. All networks were trained with an RGB input, normalised between 0 and 1.

5. Discussion and conclusions

Localisation of structures of interest in WSIs enables the extraction of interpretable features for pipelines in CPath. For this, typically a deep neural network is used per region/structure that we wish to localise. However, this may not scale well as we consider an increasing number of tissue constructs. Therefore, in this paper, we propose a multi-task learning model, named Cerberus, that performs simultaneous prediction of multiple tasks at the output of the network, without compromising on performance. In fact, we show that Cerberus achieves superior performance compared to single-task alternatives for segmentation and the learned feature representation can be leveraged to improve the performance of additional tasks via transfer learning.

Most MTL models are developed to learn a generalisable representation that is achieved by training on various independent datasets. Thereby, the collective knowledge learned can lead to more general features compared to when networks are trained on small datasets on a specific task. Despite their capability in learning a strong feature representation, datasets used in standard MTL pipelines are not necessarily aligned. For example, in CPath the datasets may have different staining, pixel resolution or belong to different tissue types. This means that, despite the network being able to make a simultaneous prediction for single image input, the outputs may not make sense. This is because each task may have its own input data assumptions. On the other hand, Cerberus ensures that each independent task-level dataset has the same

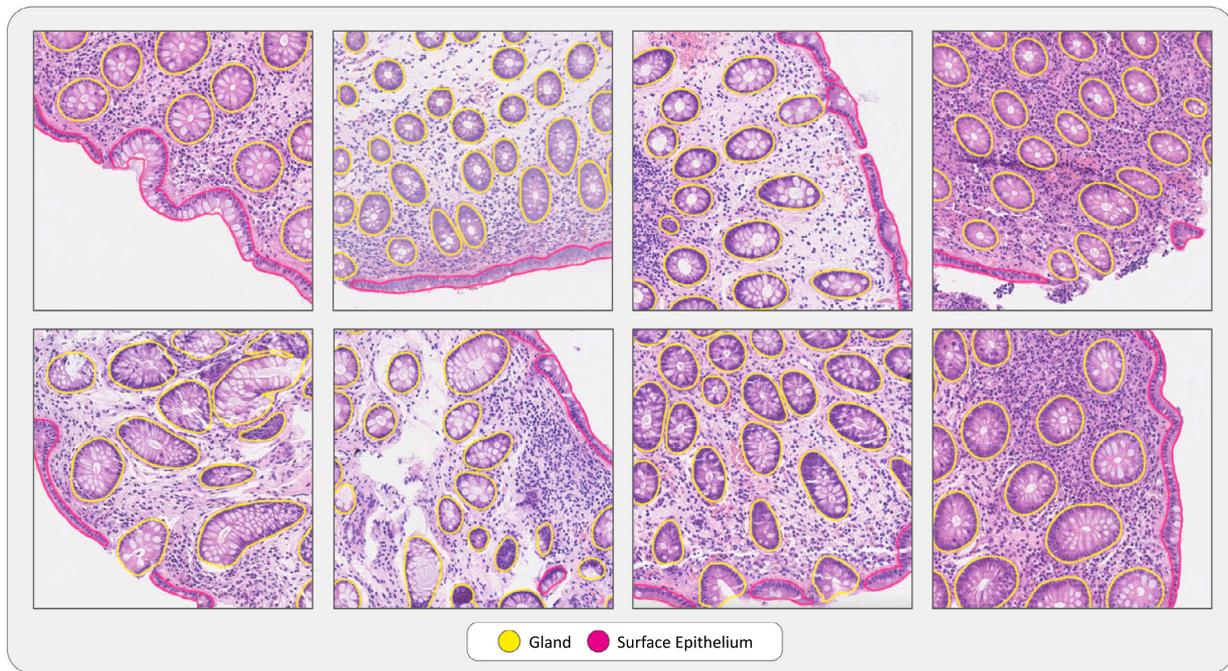


Fig. A.1. Visual results demonstrating classification of glands as either surface epithelium or not.

input assumptions and therefore a meaningful simultaneous prediction can be made for all output branches of the network. Not only this but aligning the input helps ease network optimisation. In this paper, we use H&E stained colon tissue at 20 \times objective magnification (around 0.5 $\mu\text{m}/\text{pixel}$), but the method is not necessarily limited to this data configuration. Despite the fact that MTL is usually primarily used to gain a generalisable feature representation, our main focus is enabling a strong performance for simultaneous prediction by leveraging the collective knowledge from multiple available datasets. To the best of our knowledge, this is the first approach in CPath that utilises MTL with the main focus of predicting tasks simultaneously with a single network.

We have shown that Cerberus learns a strong feature representation that can be used for transfer learning. We demonstrated that the learned representation can be used to increase the performance of patch classification tasks, such as nuclear subtyping and can be used to accurately subtype objects initially segmented by the network. In addition, we show that Cerberus can be used for improved weight initialisation for tasks such as signet ring cell detection and can even lead to better data efficiency. Even though we showcased the merit of utilising Cerberus for transfer learning, there are still plenty of other applications that have not been explored. Therefore, even though Cerberus helped improve the performance of the 3 transferred tasks considered in this paper, we cannot conclusively say that it *always* leads to an increase in performance when applied to novel tasks. In particular, it would be interesting to see whether the feature representation learned by Cerberus can be used with other tissue types.

For the purpose of this paper, we focussed on nuclear, gland and lumen segmentation as well as tissue type classification. Of course, it would be interesting to integrate additional tasks into our multi-task network, which can lead to a greater number of interpretable features that we can extract from the model output and may potentially increase the strength of the learned feature representation. For example, it would be interesting to consider instance segmentation of nerves and blood vessels as tasks within our framework, which would enable us to extract features relating to perineural invasion, lymphovascular invasion and tumour angiogenesis. It would also be interesting to instead use semantic segmentation of different tissue types, rather than relying

on patch classification. This would enable more precise localisation of different regions, which can lead to more accurate quantification of features. As discussed previously, typical MTL pipelines in CPath do not achieve a good performance across tasks at the output of the network. However, we believe that if we ensure that the input data is aligned by the same tissue type, stain type and magnification, we can sustain a good level of performance when adding further tasks, beyond those experimented with in this paper.

During the design of our experiments, we decided which tasks to use for MTL and which to use for assessment of the learned feature representation. In theory, there is no reason why a different number of upstream and additional tasks could not be chosen for training and transfer learning of the Cerberus model. For example, we could have used signet ring cell detection during MTL and nuclear segmentation to help assess the strength of the learned features. Also, we could have added signet ring cell detection as an *extra* branch during MTL. However, our chosen experimental setup allowed us to train on a large amount of data and enabled appropriate assessment of both multi-task output performance and the learned feature representation.

It would be natural to assume that due to the consideration of multiple tasks, the size of Cerberus would be very large. This is because each task has its own independent decoder. However, due to the use of relatively lightweight decoders (each containing 1.3 million parameters), the model only contains 27.1 million parameters when considering the tasks of gland segmentation, lumen segmentation, nuclear segmentation, nuclear classification and tissue type classification. This is approximately 2 times fewer than the number of parameters in MILD-Net and 3 times fewer than the number in HoVer-Net. Therefore, despite us using high-end GPUs in our experiments, this is not a requirement to fit the model into memory. When using an NVIDIA GeForce 1080 Ti GPU with a batch size of 25, the inference time for a TCGA WSI (dimensions of 28,451 \times 56,782 at a resolution of 0.5 $\mu\text{m}/\text{pixel}$) was 11.7 minutes.⁴ This will be much quicker than using separate networks for each task.

⁴ A link to the WSI that we processed can be found here: <https://portal.gdc.cancer.gov/files/3b88b579-a27a-4ac4-8533-48dded7d11c7>

Table A.1

Breakdown of nuclear classification performance (average precision) for each category. Ciga et al. (2022) uses SimCLR for model training. The difference between this and SimCLR/SimCLR⁺ reported on rows 2 and 3 is the data used during optimisation. Here, SimCLR uses the same data used during MTL and SimCLR⁺ uses over 2 million colon H&E image patches from many datasets.

(a) Cross validation results.							
Features	Mode	Epithelial	Lymphocyte	Plasma	Neutrophil	Eosinophil	Connective
IN	Patch	0.954 ± 0.002	0.742 ± 0.025	0.341 ± 0.033	0.157 ± 0.083	0.286 ± 0.073	0.728 ± 0.031
SimCLR	Patch	0.953 ± 0.004	0.761 ± 0.035	0.386 ± 0.018	0.146 ± 0.070	0.066 ± 0.033	0.736 ± 0.029
SimCLR ⁺	Patch	0.971 ± 0.002	0.778 ± 0.029	0.455 ± 0.005	0.256 ± 0.101	0.302 ± 0.090	0.791 ± 0.020
STL-Nuclei	Patch	0.963 ± 0.001	0.816 ± 0.020	0.434 ± 0.018	0.242 ± 0.129	0.542 ± 0.111	0.811 ± 0.025
STL-Gland	Patch	0.958 ± 0.005	0.743 ± 0.027	0.337 ± 0.020	0.098 ± 0.036	0.393 ± 0.130	0.724 ± 0.049
STL-Lumen	Patch	0.948 ± 0.005	0.726 ± 0.044	0.330 ± 0.024	0.103 ± 0.046	0.308 ± 0.117	0.703 ± 0.036
MTL	Patch	0.968 ± 0.001	0.793 ± 0.019	0.410 ± 0.019	0.156 ± 0.057	0.399 ± 0.137	0.793 ± 0.031
IN-MTL	Patch	0.969 ± 0.003	0.789 ± 0.023	0.427 ± 0.027	0.208 ± 0.099	0.482 ± 0.133	0.768 ± 0.050
IN-MTL+PC	Patch	0.974 ± 0.014	0.807 ± 0.049	0.463 ± 0.079	0.228 ± 0.117	0.608 ± 0.097	0.828 ± 0.058
IN-MTL+PC ⁺	Patch	0.970 ± 0.002	0.812 ± 0.029	0.439 ± 0.021	0.233 ± 0.089	0.513 ± 0.129	0.798 ± 0.034
(Ciga et al., 2022)	Patch	0.962 ± 0.004	0.763 ± 0.017	0.403 ± 0.011	0.256 ± 0.078	0.553 ± 0.055	0.755 ± 0.035
Cerberus	Pixel	0.984 ± 0.002	0.775 ± 0.019	0.430 ± 0.017	0.358 ± 0.134	0.635 ± 0.065	0.886 ± 0.014
(b) External test results.							
Features	Mode	Epithelial	Lymphocyte	Plasma	Neutrophil	Eosinophil	Connective
IN	Patch	0.682 ± 0.017	0.517 ± 0.026	0.275 ± 0.028	0.055 ± 0.019	0.334 ± 0.047	0.483 ± 0.053
SimCLR	Patch	0.885 ± 0.022	0.674 ± 0.026	0.401 ± 0.049	0.029 ± 0.011	0.044 ± 0.009	0.560 ± 0.096
SimCLR ⁺	Patch	0.917 ± 0.022	0.644 ± 0.036	0.336 ± 0.057	0.049 ± 0.023	0.351 ± 0.066	0.630 ± 0.030
STL-Nuclei	Patch	0.932 ± 0.009	0.743 ± 0.043	0.395 ± 0.045	0.060 ± 0.038	0.571 ± 0.032	0.729 ± 0.048
STL-Gland	Patch	0.912 ± 0.018	0.606 ± 0.132	0.367 ± 0.035	0.016 ± 0.006	0.450 ± 0.051	0.620 ± 0.099
STL-Lumen	Patch	0.900 ± 0.016	0.681 ± 0.065	0.326 ± 0.056	0.017 ± 0.005	0.211 ± 0.139	0.575 ± 0.095
MTL	Patch	0.939 ± 0.011	0.745 ± 0.010	0.429 ± 0.018	0.034 ± 0.010	0.496 ± 0.044	0.708 ± 0.043
IN-MTL	Patch	0.933 ± 0.016	0.657 ± 0.125	0.432 ± 0.037	0.040 ± 0.017	0.480 ± 0.037	0.698 ± 0.038
IN-MTL+PC	Patch	0.937 ± 0.018	0.714 ± 0.047	0.405 ± 0.018	0.062 ± 0.010	0.493 ± 0.030	0.726 ± 0.051
IN-MTL+PC ⁺	Patch	0.939 ± 0.014	0.750 ± 0.028	0.510 ± 0.027	0.048 ± 0.015	0.559 ± 0.029	0.710 ± 0.079
(Ciga et al., 2022)	Patch	0.931 ± 0.008	0.713 ± 0.020	0.329 ± 0.042	0.030 ± 0.006	0.558 ± 0.061	0.682 ± 0.033
Cerberus	Pixel	0.974 ± 0.004	0.688 ± 0.036	0.447 ± 0.029	0.116 ± 0.033	0.720 ± 0.006	0.858 ± 0.014

In this paper, we demonstrated how we can accurately subtype objects initially localised by our network. Specifically, we subcategorised segmented nuclei by adding an extra decoder devoted to the task of pixel classification. This can help with better quantification of important phenomena in the tumour microenvironment, which can potentially help reveal new biomarkers for downstream diagnostic and prognostic tasks. In future work, we will investigate other segmentation tasks considered within this paper. For example, our proposed subtyping approach could be used to determine whether glands are at the surface epithelium or not. We show example visual results for this in Fig. A.1.

As part of this work, we processed 599 colorectal WSIs from the TCGA database with our proposed model. As a result, we localised a large number of nuclei and other tissue constructs, which enables the extraction of interpretable features for downstream CPath pipelines. We hope that releasing this dataset will accelerate research for the development of explainable models using interpretable features because the time-consuming localisation step is no longer required. After further inspection of the results on TCGA, we found that the model found lumen segmentation in cancerous glands particularly challenging. This may be because of the limited available context when working at 20× magnification. We chose to work at this single resolution because it was reasonable for all tasks considered by the model. However, the multi-scale analysis may yield improved performance on the task of lumen segmentation. The model sometimes also struggled to distinguish between muscle and stroma tissue for patch classification. This may be due to the very small patch size that was used, with limited context. Instead, using a semantic segmentation branch that uses a larger patch as input, as described earlier in this section, would help provide additional context and therefore improve results.

It has been described numerous times that localisation of histological primitives enables the extraction of interpretable tissue-based biomarkers, such as those shown to be predictive of clinical outcome in the literature (Awan et al., 2017; Lu et al., 2018). However, this

is only the case if the segmentation performance is good enough. For example, if we generate morphological features from poorly segmented objects, then this will lead to low-quality features, which has negative implications on generated explanations. Therefore, it is imperative to perform a large-scale assessment of the quality of the localisation across multiple slides and centres. Only then can we be confident that the features that we extract are clinically meaningful.

In future work, we would like to utilise the segmentation and classification outputs of Cerberus to extract interpretable features for explainable machine learning pipelines in CPath. For example, some features that we are able to extract from the output of Cerberus include gland, lumen and nuclear morphology; the number of inflammatory cells infiltrating into glands; intra-gland epithelial nuclei organisation and quantification of cells in the stroma. In this work, we have processed the formalin-fixed paraffin-embedded colorectal TCGA database, which has associated clinical data, including whether there exist any genetic alterations. In particular, an application of interest, especially in colon tissue, would be to understand which clinically meaningful features are associated with the microsatellite instability status. Another interesting application would be to perform survival analysis on the colorectal TCGA dataset using features, like the ones we described above.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Generated results will be made publicly available.

Acknowledgements

Simon Graham, Mostafa Jahanifar, David Snead, Shan Raza, Fayyaz Minhas and Nasir Rajpoot are part of the PathLAKE digital pathology consortium, which is funded by the Data to Early Diagnosis and Precision Medicine strand of the governments Industrial Strategy Challenge Fund, managed and delivered by UK Research and Innovation (UKRI). Nasir Rajpoot was also supported by the UK Medical Research Council (grant award MR/P015476/1) and the Alan Turing Institute.

Appendix

Breakdown of nuclear subtyping results

To provide a more detailed analysis, we give the *AP* per nuclear category for subtyping tasks in Table A.1. Here, we show the results for patch-based and pixel-based subtyping, as described in Sections 4.3.3 and 4.3.4 respectively. For patch classification, we provide all methods that were given in the main paper, whereas for pixel-level classification we only provide the results obtained by Cerberus.

When using our multi-task features with auxiliary patch classification of the tissue type, we regularly outperform all other methods for patch-based subtyping. This performance is further improved when using pixel-level subtyping. We observe that, in particular, using our pixel-based subtyping method had the greatest impact on the neutrophil and eosinophil performance. For these classes, there are much fewer examples and therefore our proposed loss function may have been able to better deal with the class imbalance in the dataset. Nevertheless, the performance of epithelial cell subtyping also improved with pixel-level subtyping. Despite this, the neutrophil performance is still relatively low on the external test set. In future work, we would like to explore why this is low and understand how we can further improve the performance.

Subtyping gland segmentation output

The subtyping method that we described in Section 3.3.2 is not limited to nuclear classification. We use the same method to classify glands as either surface epithelium or glands located within the tissue. In fact, the surface epithelium is not typically denoted as being glandular and therefore subtyping in this way is important. Also, analysis of glandular morphology should typically be performed at glands that have been cut in the transverse plane and not at the surface epithelium. Therefore, this subcategorisation of the gland segmentation output also enables the extraction of better features for subsequent analysis.

In Fig. A.1 we display some example visual results, where we observe that our method can well differentiate surface epithelium from glands within the tissue.

References

- Amgad, M., Atteya, L.A., Hussein, H., Mohammed, K.H., Hafiz, E., Elsebaie, M.A., Alhusseiny, A.M., AlMoslemany, M.A., Elmatboly, A.M., Pappalardo, P.A., et al., 2021. Nucls: A scalable crowdsourcing, deep learning approach and dataset for nucleus classification, localization and segmentation. *arXiv preprint arXiv:2102.09099*.
- Amgad, M., Elfandy, H., Hussein, H., Atteya, L.A., Elsebaie, M.A., Abo Elnasr, L.S., Sakr, R.A., Salem, H.S., Ismail, A.F., Saad, A.M., et al., 2019. Structured crowdsourcing enables convolutional segmentation of histology images. *Bioinformatics* 35 (18), 3461–3467.
- Awan, R., Sirinukunwattana, K., Epstein, D., Jefferyes, S., Qidwai, U., Aftab, Z., Mujeeb, I., Snead, D., Rajpoot, N., 2017. Glandular morphometrics for objective grading of colorectal adenocarcinoma histology images. *Sci. Rep.* 7 (1), 1–12.
- Bejnordi, B.E., Veta, M., Van Diest, P.J., Van Ginneken, B., Karssemeijer, N., Litjens, G., Van Der Laak, J.A., Hermens, M., Manson, Q.F., Balkenhol, M., et al., 2017. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *JAMA* 318 (22), 2199–2210.

- Bulten, W., Pinckaers, H., van Boven, H., Vink, R., de Bel, T., van Ginneken, B., van der Laak, J., Hulsbergen-van de Kaa, C., Litjens, G., 2020. Automated deep-learning system for gleason grading of prostate cancer using biopsies: a diagnostic study. *Lancet Oncol.* 21 (2), 233–241.
- Bychkov, D., Linder, N., Turkki, R., Nordling, S., Kovanen, P.E., Verrill, C., Wallinder, M., Lundin, M., Haglund, C., Lundin, J., 2018. Deep learning based tissue analysis predicts outcome in colorectal cancer. *Sci. Rep.* 8 (1), 1–11.
- Caruana, R., 1997. Multitask learning. *Mach. Learn.* 28 (1), 41–75.
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G., 2020. A simple framework for contrastive learning of visual representations. In: *International Conference on Machine Learning*. PMLR, pp. 1597–1607.
- Chen, H., Qi, X., Yu, L., Dou, Q., Qin, J., Heng, P.-A., 2017. DCAN: Deep contour-aware networks for object instance segmentation from histology images. *Med. Image Anal.* 36, 135–146.
- Ciga, O., Xu, T., Martel, A.L., 2022. Self supervised contrastive learning for digital histopathology. *Mach. Learn. Appl.* 7, 100198.
- Ciampi, F., Jiao, Y., van der Laak, J., 2019. Lymphocyte assessment hackathon (LYSTO). <http://dx.doi.org/10.5281/zenodo.3513571>.
- Crawshaw, M., 2020. Multi-task learning with deep neural networks: A survey. *arXiv preprint arXiv:2009.09796*.
- Da, Q., Huang, X., Li, Z., Zuo, Y., Zhang, C., Liu, J., Chen, W., Li, J., Xu, D., Hu, Z., et al., 2022. DigestPath: a benchmark dataset with challenge review for the pathological detection and segmentation of digestive-system. *Med. Image Anal.* 102485.
- Diao, J.A., Wang, J.K., Chui, W.F., Mountain, V., Gullapally, S.C., Srinivasan, R., Mitchell, R.N., Glass, B., Hoffman, S., Rao, S.K., et al., 2021. Human-interpretable image features derived from densely mapped cancer pathology slides predict diverse molecular phenotypes. *Nature Commun.* 12 (1), 1–15.
- Dice, L.R., 1945. Measures of the amount of ecologic association between species. *Ecology* 26 (3), 297–302.
- Duong, L., Cohn, T., Bird, S., Cook, P., 2015. Low resource dependency parsing: Cross-lingual parameter sharing in a neural network parser. In: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. pp. 845–850.
- Fraz, M., Khurram, S.A., Graham, S., Shaban, M., Hassan, M., Loya, A., Rajpoot, N.M., 2020. Fagnet: Feature attention-based network for simultaneous segmentation of microvessels and nerves in routine histology images of oral cancer. *Neural Comput. Appl.* 32 (14), 9915–9928.
- 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*.
- Gamper, J., Rajpoot, N., 2021. Multiple instance captioning: Learning representations from histopathology textbooks and articles. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16549–16559.
- Glorot, X., Bengio, Y., 2010. Understanding the difficulty of training deep feedforward neural networks. In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics. JMLR Workshop and Conference Proceedings*, pp. 249–256.
- 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. *Med. Image Anal.* 52, 199–211.
- Graham, S., Epstein, D., Rajpoot, N., 2019b. Rota-net: Rotation equivariant network for simultaneous gland and lumen segmentation in colon histology images. In: *European Congress on Digital Pathology*. Springer, pp. 109–116.
- Graham, S., Jahanifar, M., Azam, A., Nimir, M., Tsang, Y.-W., Dodd, K., Hero, E., Sahota, H., Tank, A., Benes, K., et al., 2021a. Lizard: A large-scale dataset for colonic nuclear instance segmentation and classification. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 684–693.
- 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., Raza, S.E.A., Azam, A., Tsang, Y.W., Kwak, J.T., Rajpoot, N., 2019c. Hover-net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images. *Med. Image Anal.* 58, 101563.
- He, K., Zhang, X., Ren, S., Sun, J., 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 1026–1034.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 770–778.
- Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: *International Conference on Machine Learning*. PMLR, pp. 448–456.
- Jaber, M.I., Song, B., Beziaeva, L., Szeto, C.W., Spilman, P., Yang, P., Soon-Shiong, P., 2021. A deep learning-based iterative digital pathology annotation tool. *BioRxiv*.

- Javed, S., Mahmood, A., Fraz, M.M., Koohbanani, N.A., Benes, K., Tsang, Y.-W., Hewitt, K., Epstein, D., Snead, D., Rajpoot, N., 2020. Cellular community detection for tissue phenotyping in colorectal cancer histology images. *Med. Image Anal.* 63, 101696.
- Kather, J.N., Heij, L.R., Grabsch, H.I., Loeffler, C., Echle, A., Muti, H.S., Krause, J., Niehues, J.M., Sommer, K.A., Bankhead, P., et al., 2020. Pan-cancer image-based detection of clinically actionable genetic alterations. *Nat. Cancer* 1 (8), 789–799.
- Kather, J.N., Krisam, J., Charoentong, P., Luedde, T., Herpel, E., Weis, C.-A., Gaiser, T., Marx, A., Valous, N.A., Ferber, D., et al., 2019. Predicting survival from colorectal cancer histology slides using deep learning: A retrospective multicenter study. *PLoS Med.* 16 (1), e1002730.
- Kather, J.N., Weis, C.-A., Bianconi, F., Melchers, S.M., Schad, L.R., Gaiser, T., Marx, A., Zöllner, F.G., 2016. Multi-class texture analysis in colorectal cancer histology. *Sci. Rep.* 6 (1), 1–11.
- Kingma, D.P., Ba, J., 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- 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.
- Kumar, N., Verma, R., Anand, D., Zhou, Y., Onder, O.F., Tsougenis, E., Chen, H., Heng, P.-A., Li, J., Hu, Z., et al., 2019. A multi-organ nucleus segmentation challenge. *IEEE Trans. Med. Imaging* 39 (5), 1380–1391.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., Dollár, P., 2017. Focal loss for dense object detection. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 2980–2988.
- Lin, X., Zhen, H.-L., Li, Z., Zhang, Q.-F., Kwong, S., 2019. Pareto multi-task learning. *Adv. Neural Inf. Process. Syst.* 32, 12060–12070.
- Liu, B., Liu, X., Jin, X., Stone, P., Liu, Q., 2021. Conflict-averse gradient descent for multi-task learning. *Adv. Neural Inf. Process. Syst.* 34, 18878–18890.
- Lu, C., Romo-Bucheli, D., Wang, X., Janowczyk, A., Ganesan, S., Gilmore, H., Rimm, D., Madabhushi, A., 2018. Nuclear shape and orientation features from H&E images predict survival in early-stage estrogen receptor-positive breast cancers. *Lab. Invest.* 98 (11), 1438–1448.
- Misra, I., Srivastava, A., Gupta, A., Hebert, M., 2016. Cross-stitch networks for multi-task learning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3994–4003.
- Mormont, R., Geurts, P., Marée, R., 2020. Multi-task pre-training of deep neural networks for digital pathology. *IEEE J. Biomed. Health Inf.* 25 (2), 412–421.
- Naylor, P., Laé, M., Reyat, F., Walter, T., 2018. Segmentation of nuclei in histopathology images by deep regression of the distance map. *IEEE Trans. Med. Imaging* 38 (2), 448–459.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al., 2019. Pytorch: An imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* 32, 8026–8037.
- Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Adv. Neural Inf. Process. Syst.* 28.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 234–241.
- Ruder, S., 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*.
- Sener, O., Koltun, V., 2018. Multi-task learning as multi-objective optimization. *Adv. Neural Inf. Process. Syst.* 31.
- Shaban, M., Awan, R., Fraz, M.M., Azam, A., Tsang, Y.-W., Snead, D., Rajpoot, N.M., 2020. Context-aware convolutional neural network for grading of colorectal cancer histology images. *IEEE Trans. Med. Imaging* 39 (7), 2395–2405.
- Shaban, M., Khurram, S.A., Fraz, M.M., Alsubaie, N., Masood, I., Mushtaq, S., Hassan, M., Loya, A., Rajpoot, N.M., 2019. A novel digital score for abundance of tumour infiltrating lymphocytes predicts disease free survival in oral squamous cell carcinoma. *Sci. Rep.* 9 (1), 1–13.
- Shephard, A.J., Graham, S., Bashir, S., Jahanifar, M., Mahmood, H., Khurram, A., Rajpoot, N.M., 2021. Simultaneous nuclear instance and layer segmentation in oral epithelial dysplasia. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 552–561.
- 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., et al., 2017. Gland segmentation in colon histology images: The glas challenge contest. *Med. Image Anal.* 35, 489–502.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15 (1), 1929–1958.
- Strezoski, G., Noord, N.v., Worring, M., 2019. Many task learning with task routing. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 1375–1384.
- Tavolara, T.E., Niazi, M.K.K., Beamer, G., Gurcan, M.N., 2020. Segmentation of mycobacterium tuberculosis bacilli clusters from acid-fast stained lung biopsies: a deep learning approach. In: *Medical Imaging 2020: Digital Pathology*. Vol. 11320, SPIE, pp. 92–98.
- Tellez, D., Höppner, D., Verhoef, C., Grünhagen, D., Nierop, P., Drozdal, M., Laak, J., Ciompi, F., 2020. Extending unsupervised neural image compression with supervised multitask learning. In: *Medical Imaging with Deep Learning*. PMLR, pp. 770–783.
- Tizhoosh, H.R., Pantanowitz, L., 2018. Artificial intelligence and digital pathology: challenges and opportunities. *J. Pathol. Inform.* 9.
- Veeling, B.S., Linmans, J., Winkens, J., Cohen, T., Welling, M., 2018. Rotation equivariant cnns for digital pathology. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 210–218.
- Verma, R., Kumar, N., Patil, A., Kurian, N.C., Rane, S., Sethi, A., 2020. Multi-organ nuclei segmentation and classification challenge 2020. *IEEE Trans. Med. Imaging* 39, 1380–1391.
- Wahab, N., Miligy, I.M., Dodd, K., Sahota, H., Toss, M., Lu, W., Jahanifar, M., Bilal, M., Graham, S., Park, Y., et al., 2021. Semantic annotation for computational pathology: Multidisciplinary experience and best practice recommendations. *arXiv preprint arXiv:2106.13689*.
- Yang, Y., Hospedales, T.M., 2016. Trace norm regularised deep multi-task learning. *arXiv preprint arXiv:1606.04038*.
- Zhang, Y., Yang, Q., 2017. A survey on multi-task learning. *arXiv preprint arXiv:1707.08114*.
- Zhang, M.-L., Zhou, Z.-H., 2013. A review on multi-label learning algorithms. *IEEE Trans. Knowl. Data Eng.* 26 (8), 1819–1837.