



Tissue concepts: Supervised foundation models in computational pathology

Till Nicke ^a,* Jan Raphael Schäfer ^a, Henning Höfener ^a, Friedrich Feuerhake ^{c,e}, Dorit Merhof ^{a,b}, Fabian Kießling ^{a,d}, Johannes Lotz ^a

^a Fraunhofer Institute for Digital Medicine MEVIS, Bremen/Lübeck/Aachen, Germany

^b Institute of Image Analysis and Computer Vision, University of Regensburg, Regensburg, Germany

^c Institute for Pathology, Hannover Medical School, Hannover, Germany

^d Institute for Experimental Molecular Imaging, RWTH Aachen University, Aachen, Germany

^e Institute of Neuropathology, Medical Center – University of Freiburg, Freiburg, Germany



ARTICLE INFO

Keywords:

Foundation models
Computational pathology
Multi-task learning
Representation learning

ABSTRACT

Due to the increasing workload of pathologists, the need for automation to support diagnostic tasks and quantitative biomarker evaluation is becoming more and more apparent. Foundation models have the potential to improve generalizability within and across centers and serve as starting points for data efficient development of specialized yet robust AI models. However, the training of foundation models themselves is usually very expensive in terms of data, computation, and time.

This paper proposes a supervised training method that drastically reduces these expenses. The proposed method is based on multi-task learning to train a joint encoder, by combining 16 different classification, segmentation, and detection tasks on a total of 912,000 patches. Since the encoder is capable of capturing the properties of the samples, we term it the *Tissue Concepts* encoder.

To evaluate the performance and generalizability of the *Tissue Concepts* encoder across centers, classification of whole slide images from four of the most prevalent solid cancers – breast, colon, lung, and prostate – was used. The experiments show that the *Tissue Concepts* model achieve comparable performance to models trained with self-supervision, while requiring only 6% of the amount of training patches. Furthermore, the *Tissue Concepts* encoder outperforms an ImageNet pre-trained encoder on both in-domain and out-of-domain data.

The pre-trained models and will be made available under <https://github.com/FraunhoferMEVIS/MedicalMultitaskModeling>.

1. Introduction

The need for diagnostic systems to help pathologists manage the anticipated workload increases as cancer cancers worldwide are on the rise [1]. As Sung et al. [2] estimate, breast, colorectal, prostate, and lung cancers are among the six most common cancers types. Projections suggest that cases of these cancers will continue to increase, posing significant challenges due to time-consuming diagnosis, increased demand for tumor subtyping, and personalized treatment [3–5]. Deep learning (DL) has made significant progress in medical imaging, particularly in the field of computational pathology (CPATH). Some studies have demonstrated that DL models even surpass human performance in certain tasks, making them effective tools to help pathologists cope with the increasing workload [6,7]. However, the unavailability of the required large data sets and the needed investment of time and effort limits the effectiveness and impact of DL models in pathology.

Recent advances in self-supervised learning have enabled the training of deep neural networks on large amounts of unlabeled medical data, resulting in the creation of foundation models in computer vision [8]. These models are pre-trained on a wide range of images, primarily using self-supervision through contrastive learning or masked image modeling. They have been shown to perform well in downstream tasks, including patch classification, and weakly labeled whole slide image (WSI) classification [9,10]. Projects such as the Tissue Cancer Genome Atlas Program (TCGA) provide an openly available data source of thousands of WSIs for training these networks on real-world data. This vast amount of data is necessary for self-supervised trained networks to reach their full potential [8]. However, the amount of resources required to create, train, and deploy such models has raised concerns among researchers about environmental and other impacts [11–13]. In addition, extended training periods of several weeks

* Correspondence to: Maria-Goeppert Str. 3, 23562, Lübeck, Germany.

E-mail address: till.nicke@mevis.fraunhofer.de (T. Nicke).

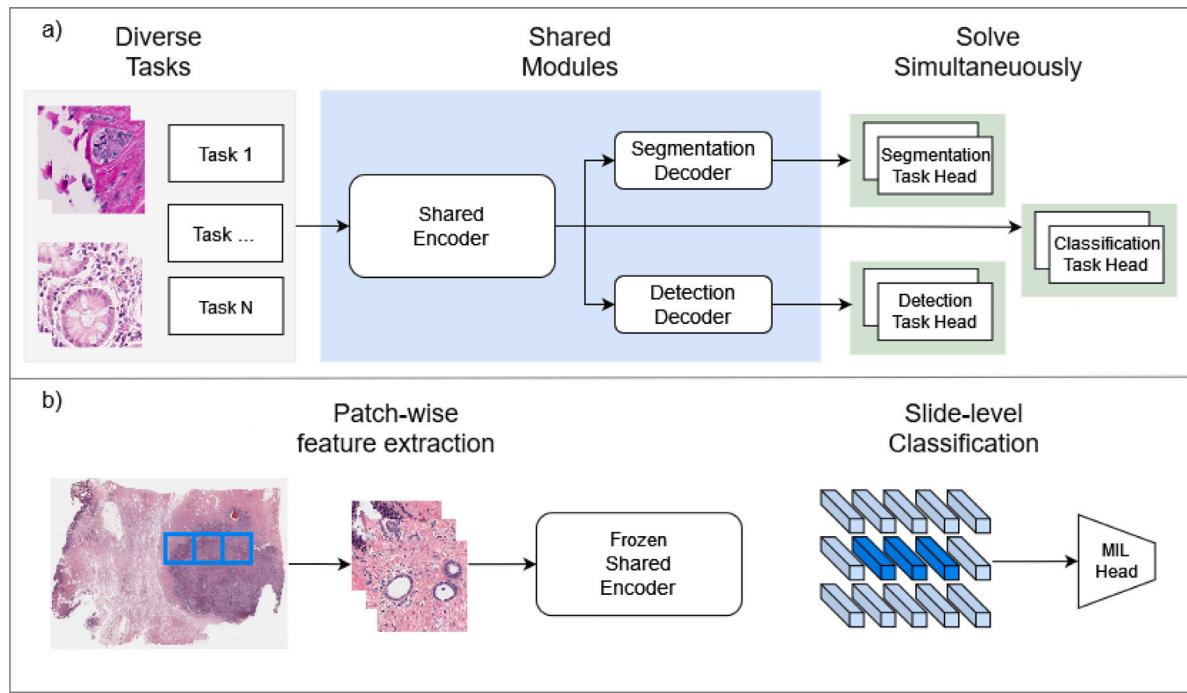


Fig. 1. Overview of the study. (a) Different pre-training of Tissue Concepts using multi-task learning on 16 different tasks. (b) the shared encoder is evaluated using multiple-instance learning on WSI classification. From each WSI, patches of size 224×224 are extracted in an iterative windowing fashion and the latent representation is positioned at the same spatial location as the patches. A simple CNN is trained on the latent WSIs to learn the label at the slice level.

impede development cycles and prolong research time.

Supervised learning, on the other hand, has been shown to outperform models trained on self-supervision in some tasks [14,15]. Although there are many annotated datasets available through challenges or other benchmarks, these datasets vary in size and contain annotations with varying degrees of detail. This variability between the datasets makes it challenging to condense the knowledge they contain into a single model. One approach to integrating all of these label types is to use multi-task learning (MTL) [16]. In [17], we recently proposed a learning framework that combines the information contained in different labeling strategies, including detection, segmentation, and classification, and use it to train a single shared backbone model on a large corpus of images. In the study, images from different medical imaging domains, such as CT, X-ray, and microscopic images but also non-medical images were included.

This paper demonstrates that training a foundation model on supervised signals in CPath using MTL requires less data, time, and energy compared to models trained with self-supervision. At the same time, the measured performance is similar to that obtained from models trained on about 17 times more data without supervision. Following the MTL training scheme presented in Fig. 1a, this paper presents *Tissue Concepts* (TC), a robust encoder that is trained on a mixture of diverse annotations from small and medium-sized datasets in CPath to learn different concepts related to tissue. Considering the future prediction of cancer cases and clinical workflow, we evaluated the performance of the encoder on the four major cancer types, breast, colon, lung, and prostate, for classification of entire slide images, as shown in Fig. 1b. In addition, since models trained on one site are known to perform worse when evaluated on different sites, we tested the performance of *Tissue Concepts* using a cross-center evaluation scheme [18].

The main contributions of the paper can be summarized as follows.

- We show that diverse pre-training using MTL learns robust representations and drastically reduces the required amount of data compared to self-supervised approaches.
- Our evaluation of the *Tissue Concepts* encoder on four of the most prevalent cancer types across multiple centers highlights the

generalizability of our approach.

2. Related work

First approaches using MTL in CPath were presented by Mormont et al. [19], and Graham et al. [20]. Mormont and colleagues converted different datasets into 22 classification tasks to train a shared network and contrasted the learned encoders against ImageNet weights. The latent representations of the encoder were used to train an SVM. They found that the representations perform equally or better than the baseline ImageNet weights. Graham et al. then used MTL on segmentation and classification tasks. This research focused on specific tasks that were present in the pre-training. However, the evaluation of the general-purpose encoder and the corresponding latent representations based on whole slide image classification combined with cross-center evaluation is still an unexplored area. In addition, general purpose encoders in the form of foundation models have not been considered by Graham et al. [20].

In [17], we presented a first approach using MTL to train supervised foundation models. Using expert knowledge in the form of multi-task learning we trained a shared model, called UMedPT, which can be applied to various medical images. To achieve this, different imaging domains, such as CT, X-ray, and microscopic images, were used to train a shared backbone on classification, segmentation, and detection tasks. Currently, the impact of tasks outside the histopathology domain remains unclear due to the diverse pre-training of the encoder. This impact on performance and robustness requires further investigation.

The following sections focus in more detail on two topics discussed in this paper. Although foundation models are still largely unexplored in terms of their application and performance, some approaches are mentioned below.

2.1. Foundation models

A foundation model is broadly defined as being trained on a wide variety of data and being easily adaptable to many different downstream tasks [21,22].

Wang et al. [9] used data from the TCGA in combination with data from the pathology AI platform, PAIP, to train a modified swin transformer, called CTransPath (CTP), using self-supervision on 15 million patches. They presented an adapted contrastive loss, based on MoCo v3 [23], which uses a memory bank to retrieve the top S semantically relevant entries. These entries were used as additional positive examples for the loss calculation. The authors evaluated their model using patch classification, image retrieval, and weakly labeled WSI classification. Due to the large number of training images and the slow convergence of self-supervised training, they reported a training time of 250 h on 48 GPUs (12.000 GPU-hours).

Campanella et al. [10], presented a comparable approach, training a tiny vision transformer (ViT) using standard DINO, and a ViT base model using standard masked autoencoding (MAE), both trained on about 3 billion patches. The authors report a training time of over 3000 GPU-hours for the models that were evaluated on a variety of tasks ranging from disease detection to outcome prediction. The evaluation also included images scanned at a different hospital from the training slides.

Chen et al. [24] present a general purpose foundation model that leverages over 100 million patches from 100.000 WSI slides across 20 major cancer types. They train a large ViT on patches, that were collected in an internal dataset, using DINOv2 [25]. They evaluate the model on 34 Tasks and find that it surpasses the previous baselines on most of them. The model was trained on 24 80 GB GPUs.

Overall, all of the presented models rely on large image databases and require long training times, which contributes to increased CO₂ emissions. The presented TC encoder and MTL training aim to reduce the need for large amounts of data while maintaining the desired performance. In addition, cross-center evaluation is needed to accurately assess models' performances.

2.2. Weakly labeled WSI classification

Learning from WSIs that are only labeled on a case basis, or that have only one endpoint, is challenging because training on the entire image at once typically exceeds the GPU memory. In addition, since a WSI provides only one sample, many WSIs are needed to effectively train a deep learning model. Classification of such gigapixel images is therefore typically performed using multiple instance learning (MIL) [26,27]. Using MIL involves two parts: first, extracting features from patches of the WSI using a pre-trained encoder to convert them into their latent representations, and second, aggregating features from a WSI using a trainable MIL head to predict the given label [28]. Therefore, robust encoders are needed to obtain patch representations that facilitate the second step of MIL [29]. In this paper, MIL is used as an evaluation procedure to test the representativeness of the encoder's features. The following presents the most commonly used approaches that focus on solving the second stage of MIL using either attention or convolution-based methods.

Lu et al. [30] introduced CLAM, a clustering-constrained attention MIL algorithm. The authors trained an attention-based head on features extracted from patches of a WSI to classify the corresponding labels. The attention was then used to identify subregions of high diagnostic value, which in turn were used to classify the entire slide. In addition, instance-level clustering was applied over the representative regions to constrain and refine the feature space.

Shao et al. [31] proposed TransMIL, an attention-based correlation method for solving weakly labeled classification tasks. The method uses differently sized convolutional layers to apply additional pyramid position encoding information between the attention modules. This allows the attention layers to aggregate morphological features, while the Pyramid Position Encoding Generator (PPEG) encodes spatial information.

Tellez et al. [32] proposed neural image compression to train on entire WSIs. The authors trained an autoencoder on patches and used

the resulting encoder for feature extraction. The patches extracted in the image domain were encoded and their latent representations were placed in the same spatial location. This effectively compressed the entire WSI into a smaller latent image with more channels, while preserving the spatial relationship between the individual patches. A small CNN was then trained on the compressed WSIs to predict the label of the WSI. In a second version of this approach the same authors used multi-task learning on four classification tasks to train the feature extractor [33]. The effect of segmentation and detection tasks, as well as more diverse pre-training, remained a point of further investigation and are part of the research presented in this paper.

All presented methods propose different aggregation methods to learn the desired label predictions based on the extracted features and thus work with the features extracted by the TC encoder. As an evaluation method, we adapted the convolution-based aggregation method presented by Tellez et al. [32] and also applied an attention-based approach based on Ilse et al. [34], which are further described in Section 4.3.

3. Methods

Multi-task learning was used to train the Tissue Concepts encoder on 14 small and medium-sized datasets. The pre-training datasets and the procedure used during this phase are described in the following sections. In addition, the evaluation datasets, MIL head, and corresponding training are explained in detail.

3.1. Pre-training datasets

To pre-train the shared encoder, a total of 14 data sources were collected and distributed over 16 tasks. All data sources with corresponding number of patches and WSIs, as well as their tasks are presented in Table 1. The NCT-CRC-HE 100k, Panda, TUH, Breakhis, and Arvaniti datasets have been designated as patch-level classification tasks by their respective curators. Data from the PANDA dataset was used for classification and segmentation tasks, where the patches in PANDA were 20 patches per WSI, sampled of a total of 4000 WSIs. To increase data diversity, the Breakhis dataset was used at both 40x and 100x magnification. Conic and MiDoG were included as detection tasks, while Conic also served as a segmentation task. Other datasets used for segmentation were the SemiCOL training dataset, Arvaniti, Peso, Schoemig-Markiewka, PANDA, TIGER, CRAG, BCSS, and HubMap. Each of the mentioned datasets is described in more detail in Appendix A. In total, about 912,000 patches from around 7000 WSIs were used during the training. The training data consisted of about 100,000 patches from colorectal, 680,000 patches from prostate, and 10,000 patches from breast tissue. In addition, a small number of spleen, liver, and skin tissue slides were included, which are marked as "various" in Table 1.

All patches were extracted or scaled to 224 × 224 pixels, resulting in a resolution of approximately 1 to 0.5 micron per pixel (MPP). In addition, standard augmentation, such as random rotation, distortion, blurring, brightness, contrast, and hue changes were applied during pre-training.

3.2. Encoder training

The shared backbone was trained using multi-task learning with all of the above data sources [17]. In the MTL pipeline, each task was solved by a shallow task-specific head (θ_t) that received input from larger, shared modules (θ_{shared}). Fig. 1a gives a brief overview of the different blocks in the architecture. The shared blocks are shown in the blue area, while the individual tasks are shown in green. During training, all tasks were treated equally and were processed iteratively within the training loop. For each task the task-specific loss was computed and accumulated on the total loss. Formally, the total loss $\mathcal{L}_{total}(X, Y)$ for a set of tasks $X = (X_1, \dots, X_T)$ of task specific images $X_t = (x_{t,1}, \dots, x_{t,N})$

Table 1

Overview of the datasets used in the training of the “Tissue Concepts” encoder (purpose “training”), the ablation study (“ablation”), and in the validation experiments (“validation”).

Dataset	# WSIs	# Patches	Organ	Task Type	Purpose
NCT-CRC-HE 100K [35]	86	100 000	Colorectal	Classification	training
Crag [36,37]	14	173	Colorectal	Segmentation	training
SemiCOL [38]	20	1759	Colorectal	Segmentation	training
Conic [39,40]	291	4981	Colon	Detection	training
Conic [39,40]	291	4981	Colon	Segmentation	training
Arvaniti [41]	641	2000	Prostate	Classification	training
Peso [42]	62	392	Prostate	Segmentation	training
Schoemig-Markiewka [43]	1177	600 000	Prostate	Classification	training
PANDA [44]	2000	40 000	Prostate	Classification	training
PANDA [44]	2000	40 000	Prostate	Segmentation	training
TUH [45]	136	3720	Breast	Classification	training
TIGER [46]	151	1876	Breast	Segmentation	training
BCSS [47]	18	139	Breast	Segmentation	training
BreakHis [48]	82	4076	Breast	Classification	training
MiDoG [49]	49	405	Various	Detection	training
HubMap [50]	315	1260	Various	Segmentation	training
Total	7042	912 157	Various	Various	training
Bach [51]	400	400	Breast	Classification	ablation
BRACS [52]	547	5 566 906	Breast	MIL	validation
Panda [44]	10 000	98 928 234	Prostate	MIL	validation
SemiCOL [38]	499	1 436 065	Colorectal	MIL	validation
TCGA-NSCLC [53,54]	1006	12 641 437	Lung	MIL	validation
Total	12 052	118 572.642	Various	MIL	validation

and corresponding labels $Y_t = (y_{t,1}, \dots, y_{t,N})$ is computed as

$$\mathcal{L}_{total}(X, Y) = \sum_{t=1}^T \mathcal{L}_t(\theta_t(\theta_{shared}(X_t)), Y_t) \quad (1)$$

where the loss of the task t is calculated based on the output of the task specific head θ_t that received its’ input from a shared network structure θ_{shared} . X_t and Y_t describe the task-specific batch input images and the corresponding labels respectively. While the task-specific parameters θ_t in Eq. (1) changed, depending on which task was being processed, the shared parameters θ_{shared} were task-independent. Since all tasks contributed equally to the total loss, multiple cycles through all tasks can be performed by performing gradient accumulation before an optimization step is performed.

The tasks themselves were divided into classification, segmentation, and detection. All tasks shared the same encoder, which yielded a feature pyramid with four feature maps at different perceptual levels. The feature pyramid was further processed depending on the task.

Classification tasks received a globally average pooled version of the lowest feature pyramid level from the encoder. Each task-specific classification head consisted of 20% dropout followed by a linear layer, mapping from the dimension of the latent space to the number of classes. By minimizing the cross-entropy loss, each head learned to predict the task-specific classes.

Segmentation tasks consisted of a shared U-Net style decoder that received features from all levels of the shared encoder. Each individual segmentation task was then processed by a 1×1 convolutional layer operating on the output of the last decoder level. A combination of dice and focal loss was minimized by each segmentation head.

The **detection tasks** also consisted of a shared U-Net style decoder that received features from all levels of the shared encoder to extract spatially relevant feature maps. Unlike the segmentation tasks head, the detection task head consisted of an FCOS detection head, which minimized the corresponding losses.

Since transformer-based approaches are currently dominant in the literature we chose a tiny swin transformer as shared encoder [55,56] (further denoted as TC-Swin). The encoder shares the same latent dimension with CTransPath, 768, and has 27.5M trainable parameters. Additionally, we compared the attention-based backbone to a tiny ConvNeXt architecture that shares a similar number of parameters and latent dimension but follows a convolution-based approach [57]

(further denoted as TC-Conv). Both networks were initialized with the corresponding ImageNet-1k weights which were imported from the torchvision library [58]. During pre-training, the networks were optimized for 5 million steps until convergence, where one step was defined as computing the loss for one batch of one task. AdamW optimizer [59] was used with a constant learning rate of 10^{-4} , a weight decay of 0.01, and a gradient accumulation of 128 steps. In total, the models were trained for 160 h on one NVidia RTX A5000 GPU and the final models were used for evaluation.

4. Evaluation

This section first establishes a comparison between UMedPT and TC using a sample efficiency experiment as an ablation study. Then, the evaluation datasets and corresponding tasks are presented, and the MIL head training procedure is described.

4.1. Ablation study

To determine whether a multi-domain encoder, UMedPT, or a histology specific encoder, TC-Swin, performs better on unseen histology data, we measured the sample efficiency on an unseen, in-domain dataset. In this experiment, the frozen encoder generated latent representations for all patches in the downstream dataset. These latent representations were then used as input to a random forest model. Sample efficiency was measured by establishing of a fixed training and validation split. Different subsets of images per class were systematically sampled from the training split to serve as training examples for the random forest. Each model was consecutively evaluated on the designated test split. This training-validation process was repeated 10 times with different seed values for robustness in the analyses and the F1-score was reported. We compared TC-Swin to the base Swin transformer encoder presented in [17]. Additionally, ImageNet weights from torchvision were used as baseline for a tiny swin transformer.

The **BACH** challenge dataset [51] is a breast tumor classification dataset consisting of 400 images of 2048×1536 pixels at 0.25 MPP, which are equally divided into 4 classes. Images were center cropped to 1024×1024 pixels and downsampled to 224×224 pixels. The dataset was split 80/20 into a training and test subset. From the training split of the BACH dataset, sets of 1, 3, 5, 10, 25, and \max 46 images per class

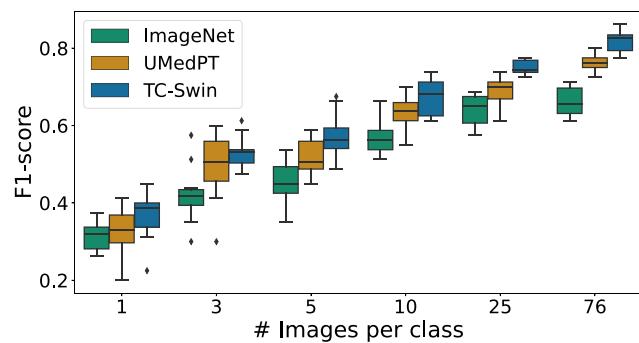


Fig. 2. Sample efficiency for encoders of different specificity on the BACH patch classification dataset. Each boxplot represents 10 repetitions. The TC-Swin encoder is compared to the tiny swin transformer ImageNet weights and a small swin transformer of UMedPT [17]. The F1-scores are plotted against increasing numbers of images per class to examine sample efficiency.

were sampled using 10 different seeds. For each training, the random forest was evaluated on the left out test set.

Fig. 2 shows that the domain-specific encoder outperforms the multi-domain encoder and the pre-trained ImageNet encoder in terms of sample efficiency. The TC-Swin encoder showed higher F1-scores compared to the multi-domain encoder and to ImageNet. With three images per class, TC-Swin's mean F1-score was at 0.531, while UMedPT and ImageNet were at 0.496 and 0.423, respectively. With 10 images per class, the TC-Swin encoder showed a mean F1-score of 0.673, while UmedPT and ImageNet showed values of 0.628 and 0.566, respectively. Using the maximum number of available training images per class, the differences in reported performance are similar. While TC-Swin achieved 0.816, UmedPT and ImageNet achieved mean F1-scores of 0.762 and 0.662, respectively.

The TC-Swin encoder outperforms the multi-domain medical encoder UMedPT even though the later was trained with more images overall. In particular, with 5 or more images per class, the features obtained from the TC encoder provide better performance than the ImageNet and UMedPT features. When considering 1 and 3 images per class, UMedPT shows similar performance to the TC encoder, but also shows a higher variance over the 10 seeds. Therefore, we excluded the UMedPT encoder from further evaluation and focused on the TC encoder instead.

4.2. Evaluation datasets

We compared the TC encoder to the current state of the art for each organ, referred to as the external baseline, as well as to the results obtained using the features of the publicly available CTransPath encoder [9]. In addition, we use the Swin transformer model with ImageNet weights without further training as an additional baseline.

The learned representations of Tissue Concepts were evaluated using weakly labeled WSI classification tasks. Breast, prostate, and caolorectal tissue represent tissues within the training domain, lung cancer tissue was deliberately chosen to represent tissues outside the training domain. The architecture and learning schedule as described in Section 4.3, was used during the training of each MIL head. Where possible, multi-site datasets were split so that evaluation sites were not included in the training. The feature extractors were evaluated on the following datasets for each organ:

Breast Cancer: The BReAst Cancer Subtyping (BRACS) dataset consists of 547 WSIs from 189 patients which are already divided into training (395), validation (87) and test (65) splits [52]. According to Pati et al. [60], the WSIs were obtained at the Department of Pathology at the National Cancer Institute - IRCCS-Fondazione Pascale, Naples, Italy, and were scanned with an Aperio AT2 scanner

at a resolution of $0.25 \mu\text{m}$ per pixel (mpp). Each WSI is assigned to one of 7 different classes: normal, benign, usual ductal hyperplasia (UDH), atypical ductal hyperplasia (ADH), flat epithelial atypia (FEA), ductal carcinoma in situ (DCIS), and invasive. In addition, Pati et al. [60] grouped the different types of breast carcinoma into four coarser classes: normal, non-cancerous, precancerous, and cancerous. Furthermore, binary classification between invasive and non-invasive breast cancer was also performed by Pati et al. [60]. The authors also presented a trained model, HACT-Net, and applied it to this dataset. They reported the weighted F1-score for all of the classification problems mentioned above which we used as an external baseline and also reported the weighted F1-score in the BRACS experiments. Additionally, following Maier-Hein et al. [61], we reported accuracy as multi-class metric and area under the receiver operating curve (AUC) as multi-threshold metric. For the evaluation of the TC-encoder, three experiments were created from the dataset, similarly to the approach proposed by Pati et al. [60]. The first experiment consisted of differentiating between normal and cancerous tissue slides. The second experiment focused on more fine-grained classification between normal, non-cancerous, precancerous, and cancerous tissues. Finally, the third experiment for breast cancer aimed at predicting one of the seven sub-types provided.

Prostate Cancer: The Prostate cANcer graDe Assessment (PANDA) challenge provides 10 616 openly available WSIs from two different centers, Karolinska and Radboud [44]. Each of the WSIs was assigned an ISUP score resulting from the primary and secondary Gleason pattern. The score ranges from 0 to 5, with 0 representing normal tissue and higher numbers representing more severe prostate cancer. This information was used to create binary classification task or evaluation, where the task consists of differentiating between normal (ISUP 0) and cancerous tissue (ISUP > 0). To test cross-center performance and overall performance, three different experiments were created from this dataset. First, the entire set of images obtained from the Radboud UMC was used as both the training and validation set, using a split ratio of 90% for training and 10% for validation. All images provided by the Karolinska Institute were used as the hold-out test set. In a subsequent experiment, the roles of training and test centers were swapped. In addition, a 5-fold cross-validation experiment involving both centers was conducted. We note that a small percentage of the dataset in the form of pre-selected patches were used during pre-training of the encoder, however, the task-specific head was discarded. To evaluate the effect of possible data leakage [62] an additional encoder was trained without the PANDA dataset in pre-training and evaluated only on this task. Mahdi Behzadi et al. [63] used this dataset in a k-fold cross-validation to evaluate the performance of their model and reported the accuracy and area under the curve for binary classification. Their reported results serve as an external baseline. Mahdi Behzadi et al. [63] reported the accuracy and AUC of their algorithm on the dataset. To be comparable, we also report these metrics in our experiments. Additionally, we report the F1-score as per-class metric, as suggested by Maier-Hein et al. [61].

Colorectal Cancer: During the SemiCOL challenge, a dataset of 499 WSIs from 4 different centers was provided as training dataset. The challenge task was to predict whether a given slide contained cancerous tissue or normal tissue. During the challenge, we used a modified model presented in [17] as an initial starting point and further fine-tuned on several colon and colorectal tasks. Using the learned encoder to extract features from the WSI, a MIL head was successively trained on these features. This resulted in an external test AUC of 0.99. In the experiments presented in this paper, the challenge training data was split by slide provider, effectively creating four smaller subsets of the original challenge dataset. For the evaluation presented in this paper, four different experiments were conducted. Images from one center were used as training examples (90/10), while images from the other three centers served as the hold out test set. By treating each center as training provider once, and using the other three centers as test

centers, four different configurations were created. We note that a small percentage of patch data of the dataset was used during pre-training of the encoder, however, the task-specific head was discarded. To the best of our knowledge, there is no other publication reporting results on this specific dataset, so we are unable to report an external baseline. Since the dataset is balanced with regard to the class distribution for each center in all splits, we reported the Accuracy, F1-Score and AUC as metrics, following Maier-Hein et al. [61].

Lung Cancer: The TCGA-NSCLC dataset contains two subtypes of lung cancer which are lung squamous cell carcinoma (TCGA-LUSC [53]), and lung adenocarcinoma (TCGA-LUAD [54]) in 1006 (512 LUSC, 494 LUAD) slides from over 40 different sites. Two training and testing scenarios were created from the TCGA-NSCLC dataset. One experiment focused on cross-center evaluation. In this scenario, the three largest contributors (Johns Hopkins, International Genomics Consortium, and Asterand) were selected to serve as the training and validation subset providers. The remaining 39 sites were selected to serve as hold out test set. This split resulted in 340 training slides (90/10 split) and 666 test slides. In addition, a 5-fold cross-validation experiment was performed on the entire dataset to increase training set diversity, again using a 90/10 training and validation split. The objective of each experiment was to distinguish between slides containing lung adenocarcinoma and those containing squamous cell carcinoma (LUAD vs. LUSC). The obtained results were compared to the results reported by Shao et al. [31] and Wang et al. [9]. To be comparable, the accuracy and AUC, as reported by the external baselines, were selected as metrics. Additionally, the F1-score was reported.

4.3. Training of the MIL head

Following medical diagnoses, WSIs are often only labeled at the patient or slide level. To address this, we used MIL to train slide-level classification heads for the respective evaluation tasks. The WSIs were compressed into latent images, by the respective frozen encoder, as suggested by Tellez et al. [32]. During this process, all patches of a WSI were extracted with a size of 224×224 at about 0.5 MPP.

Using the exact same architecture for aggregation as presented by Tellez et al. [32], we found that training the proposed MIL head left room for improvement in terms of training stability and convergence speed. The selection and adaptation were motivated by the prior performance on the SemiCOL challenge [38]. There, we fine-tuned UMedPT and trained the small MIL head on WSI classification. On an external, multi-center test set, this pipeline achieved an AUC of 0.998.

To aggregate the latent WSIs, a small CNN with global max pooling, was used. The initial layer utilized a 1×1 convolution, reducing the latent WSIs' 768 channel to 16. Following this, 3×3 convolutions with a stride of 2, and padding of 1 doubling in depth each time were applied. Instance normalization, leaky ReLU activation, and 10% 2D dropout were applied in the convolutional layers. The lowest feature pyramid of 64 channels was globally max pooled, followed by one hidden linear layer and a classification layer. This resulted in 19 250 total trainable parameters. The convolutional layers were initialized by drawing from a Kaiming normal distribution with a fixed seed. Additionally, we tested an attention-based approach by replacing the max pooling layer with an attention pooling layer, as suggested by Ilse et al. [34]. Overall, we found similar results with both approaches across all organs. The results using an ABMIL aggregation head can be found in [Appendix B.1](#).

The same architecture was used over all evaluation datasets. To avoid overfitting, a label smoothing of 0.1 was used while minimizing the cross-entropy loss. The MIL head was trained using the AdamW optimizer with a constant learning rate of 10^{-4} and a weight decay of 0.01 for 100 epochs, while monitoring the validation loss. The best performing model was used for testing. During training, additional augmentations were applied to each latent WSI. Since each latent WSI was treated as a 768-channel image, random flipping, mirroring, and resizing with nearest neighbor interpolation were applied during training. This resulted in high-dimensional latent WSI representations with sizes ranging from 32×32 to 224×224 .

Table 2

Breast: Tissue-specific model performance when applied to the BRACS challenge data. The model consist of organ-specific head and frozen encoders: ImageNet, CTP, TC-Swin, TC-Conv. The table shows mean weighted F1 score, AUC and standard deviation over 5 different runs for different problem formulations. The best score is shown in bold.

Problem	Model	wF1	AUC
2-class	Pati et al. [60]	0.96	–
	ImageNet	0.960 \pm 0.02	0.994 ± 0.003
	CTP	0.954 ± 0.02	0.997 \pm 0.001
	TC-Swin	0.951 ± 0.02	0.992 ± 0.001
4-class	TC-Conv	0.921 ± 0.02	0.989 ± 0.006
	Pati et al. [60]	0.84	–
	ImageNet	0.490 ± 0.033	0.777 ± 0.01
	CTP	0.613 \pm 0.03	0.829 ± 0.01
7-class	TC-Swin	0.545 ± 0.04	0.83 ± 0.01
	TC-Conv	0.555 ± 0.03	0.837 \pm 0.02
	Pati et al. [60]	0.69	–
	ImageNet	0.349 ± 0.01	0.711 ± 0.002
7-class	CTP	0.3989 ± 0.02	0.755 ± 0.01
	TC-Swin	0.405 \pm 0.02	0.769 \pm 0.02
	TC-Conv	0.378 ± 0.02	0.764 ± 0.02

5. Results

An overview of the obtained results is given in [Fig. 3](#), where all tasks are presented with corresponding test AUC scores over 5 runs. The results show that the TC encoder, although trained with less data, performs as well as, or better than the CTransPath encoder, given the selected tasks. The Swin transformer with ImageNet weights exhibit the lowest overall performance and the largest variance across all tasks. All reported results were obtained using the default parameters of the corresponding functions from the scikit-learn library [64]. In the following, the results on the different organs are examined in more detail.

5.1. Breast cancer

[Table 2](#) shows the results for the three sub-problems of the BRACS classification challenge, including the weighted F1-score and the AUC score for each. For the binary classification problem, ImageNet weights achieve the highest weighted F1-score on average with 0.960, which is comparable to the external baseline of 0.96. CTP and TC-Swin show similar performance with average wF1-scores of 0.954 and 0.951, respectively. The TC-Conv encoder performs almost three percentage points worse than the TC-Swin encoder. Overall, all tested encoders result in similar AUC values.

Considering the more fine-grained classification of the given WSIs, none of the encoders matches the reported performance of the external baseline HACT-Net with a weighted F1-score of 0.84. The two TC encoders outperform the ImageNet weights on the four- and seven-class problem formulations. On the four-class problem, CTP achieves a mean weighted F1-score of 0.613, approximately 6 percentage points higher than the two TC encoders. However, when considering the AUC, all of the foundation models perform equally with a mean AUC of about 0.83. For the more detailed seven-class problem, the TC-Swin encoder achieves the highest mean weighted F1 and AUC scores out of all the tested models with 0.40 and 0.76, respectively.

Overall, the performance of all four encoders tested decreases significantly as the problem formulation becomes more fine-grained. Although the external baseline performance also declines, the decline is not as steep as that observed with the foundation models and ImageNet weights.

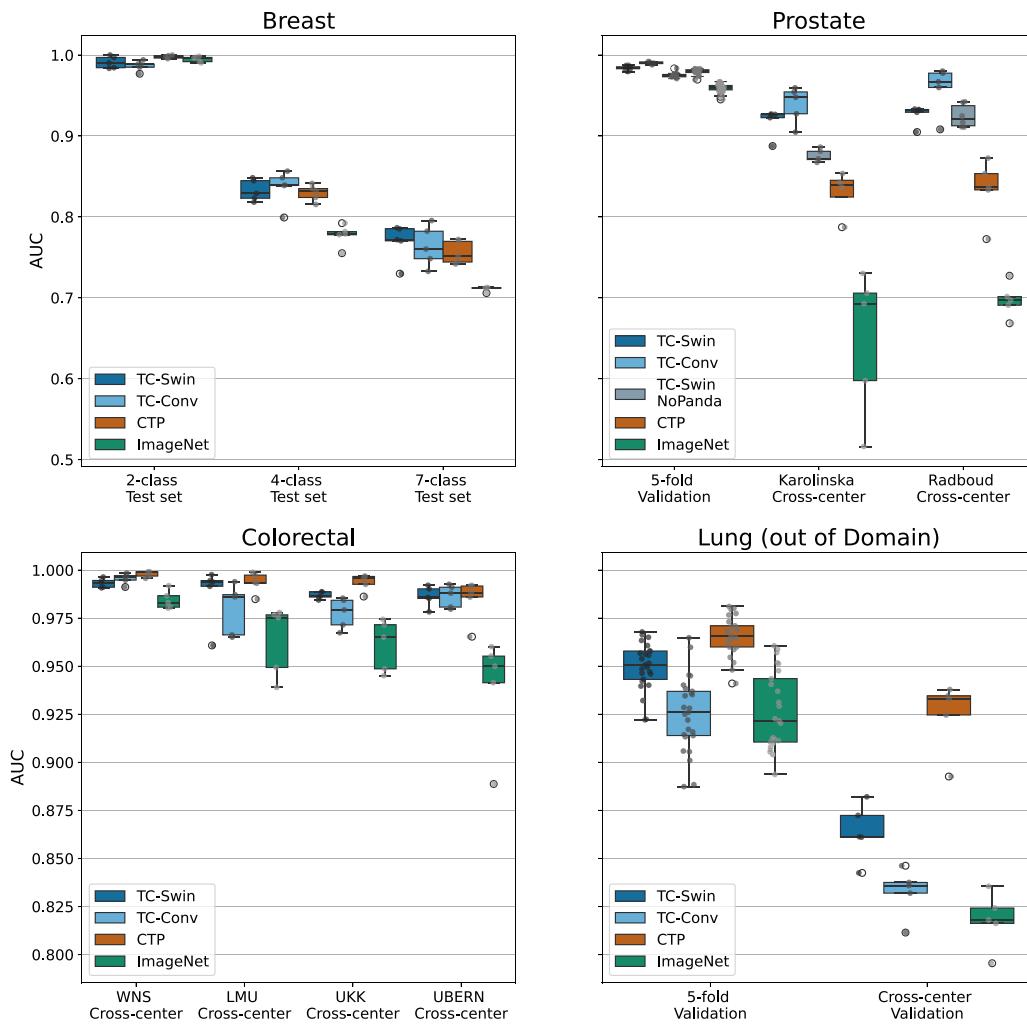


Fig. 3. Organ-specific model performance when applied to downstream data. Models consist of organ-specific head and frozen encoders: ImageNet, CTP, TC-Swin, TC-Conv. While breast, colorectal, and prostate tissue was included in the pre-training, no lung tissue has been used during encoder pre-training of TC. “Breast” shows the results on three different classification tasks based on the BRACS dataset. “Colon” shows the cross-center performance between sites of the SemiCol challenge when trained on one center and evaluated on the others. “Lung” shows a 5-fold cross-validation and a cross-center evaluation. “Prostate” shows the performance on two cross-center evaluations as well as a 5-fold cross-validation from the Panda-challenge dataset. Boxes cover the inter-quartile-range, the median is marked by a horizontal line.

Table 3

Prostate: Tissue-specific model performance when applied to the PANDA challenge data. The model consist of organ-specific head and frozen encoders: ImageNet, CTP, TC-Swin, TC-Conv, TC-Swin-NoPanda. The table shows mean weighted Accuracy, F1 score, AUC, and standard deviations over 5 different runs for 5-fold cross-validation and cross-center transfer performance. The best score is shown in bold.

Clinic	Model	Acc	F1	AUC
External	Mahdi Behzadi et al. [63]	0.95	0.93	–
5 fold	ImgNet	0.900 ± 0.01	0.877 ± 0.01	0.958 ± 0.01
	CTP	0.936 ± 0.01	0.921 ± 0.01	0.979 ± 0.01
	TC-Swin	0.942 ± 0.01	0.928 ± 0.01	0.983 ± 0.01
	TC-Conv	0.957 ± 0.05	0.947 ± 0.06	0.990 ± 0.01
	TC-Swin-NoPand	0.932 ± 0.01	0.918 ± 0.01	0.975 ± 0.003
Karolinska → Radboud	ImgNet	0.658 ± 0.15	0.529 ± 0.08	0.648 ± 0.08
	CTP	0.819 ± 0.01	0.524 ± 0.04	0.825 ± 0.03
	TC-Swin	0.896 ± 0.01	0.829 ± 0.01	0.914 ± 0.02
	TC-Conv	0.900 ± 0.05	0.845 ± 0.06	0.958 ± 0.01
	TC-Swin-NoPanda	0.865 ± 0.07	0.758 ± 0.08	0.924 ± 0.01
Radboud → Karolinska	ImgNet	0.646 ± 0.01	0.577 ± 0.04	0.697 ± 0.02
	CTP	0.678 ± 0.03	0.512 ± 0.11	0.834 ± 0.03
	TC-Swin	0.846 ± 0.02	0.842 ± 0.02	0.926 ± 0.01
	TC-Conv	0.764 ± 0.05	0.646 ± 0.06	0.958 ± 0.01
	TC-Swin-NoPanda	0.821 ± 0.07	0.807 ± 0.06	0.875 ± 0.007

Table 4

Colorectal: Tissue-specific model performance when applied to the SemiCOL challenge data. The model consist of organ-specific MIL head and frozen encoders: ImageNet, CTP, TC-Swin, TC-Conv. The table shows mean accuracy, F1 score, AUC, and the respective standard deviations over 5 different runs. Each model was trained on images from one site and evaluated on all images from the remaining sites. The best score is shown in bold.

Clinic	Model	Acc	F1	Auc
LMU	ImgNet	0.877 ± 0.05	0.875 ± 0.06	0.961 ± 0.01
	CTP	0.967 ± 0.01	0.964 ± 0.01	0.991 ± 0.004
	TC-Swin	0.948 ± 0.03	0.948 ± 0.03	0.987 ± 0.01
	TC-Conv	0.921 ± 0.02	0.919 ± 0.02	0.979 ± 0.01
UBERN	ImgNet	0.734 ± 0.09	0.709 ± 0.11	0.939 ± 0.02
	CTP	0.860 ± 0.08	0.854 ± 0.09	0.985 ± 0.01
	TC-Swin	0.943 ± 0.02	0.944 ± 0.02	0.987 ± 0.01
	TC-Conv	0.943 ± 0.01	0.943 ± 0.01	0.986 ± 0.01
UKK	ImgNet	0.890 ± 0.02	0.889 ± 0.02	0.960 ± 0.01
	CTP	0.962 ± 0.01	0.962 ± 0.01	0.993 ± 0.01
	TC-Swin	0.955 ± 0.01	0.953 ± 0.01	0.989 ± 0.01
	TC-Conv	0.925 ± 0.02	0.925 ± 0.01	0.978 ± 0.01
WNS	ImgNet	0.936 ± 0.02	0.926 ± 0.02	0.984 ± 0.004
	CTP	0.989 ± 0.01	0.989 ± 0.01	0.998 ± 0.001
	TC-Swin	0.967 ± 0.01	0.966 ± 0.01	0.993 ± 0.01
	TC-Conv	0.973 ± 0.01	0.973 ± 0.01	0.995 ± 0.01
Average	ImgNet	0.859 ± 0.09	0.852 ± 0.1	0.962 ± 0.02
	CTP	0.944 ± 0.06	0.943 ± 0.01	0.992 ± 0.01
	TC-Swin	0.953 ± 0.02	0.953 ± 0.02	0.989 ± 0.01
	TC-Conv	0.941 ± 0.02	0.940 ± 0.02	0.984 ± 0.01

5.2. Prostate cancer

[Table 3](#) shows the results obtained of five runs on the PANDA dataset with cross-center and mixed evaluation. The ‘clinic’ column indicates different settings, where “Karolinska → Radboud” refers to the MIL head being trained on data from the Karolinska Institute and evaluated on data from the Radboud Clinic. The additional encoder TC-Swin-NoPanda was trained without PANDA data in the pretraining.

When comparing the results obtained in the 5-fold cross-validation with the external baseline, both TC encoders show performances that were close to the external baseline, with TC-Conv slightly surpassing it with a mean accuracy of 0.957. The mean accuracies of CTP and ImageNet are 2 and 5 percentage points below the selected baseline, respectively. All foundation models outperform the pure ImageNet encoder by at least three percentage points. In terms of accuracy and AUC scores, two TC encoders slightly outperform the CTP weights by one percentage point with mean acc of 0.936 and mean auc of 0.979. When training the TC-Swin encoder without the PANDA dataset, a slight drop in performance can be observed, resulting in similar metrics compared to the CTP encoder.

In the cross-center evaluation, the performance of all models decreases compared to an in-center evaluation. CTP achieves a mean accuracy of 0.819 when transferring decision boundaries from the Karolinska site to Radboud subset. The opposite scenario results in a decrease of 14 percentage points (0.678 mean acc). Furthermore, an even steeper drop in mean accuracy and AUC from 5-fold cross-validation to cross-center validation is observed when using the ImageNet weights (from 0.900 to 0.658). All TC encoders are capable of extracting robust features. The TC-Swin encoder (from 0.942 to 0.846) shows more stable performance than the TC-Conv encoder (from 0.957 to 0.764) when compared to the metrics reported in the cross-center evaluation. The TC-Swin encoder trained without panda shows slightly lower performance on average to the TC-Swin encoder trained with PANDA. Independently of including the PANDA dataset in the pre-training, all TC-based encoders show a stronger cross-center performance over CTP.

5.3. Colorectal cancer

[Table 4](#) presents the results per trained center on the SemiCOL challenge dataset. The ‘clinic’ column indicates the training center.

Table 5

Lung: Tissue-specific model performance when applied to the TCGA-NSCLC dataset. The model consist of organ-specific MIL head and frozen encoders: ImageNet, CTP, TC-Swin, TC-Conv. The table shows mean weighted accuracy, F1 score, AUC, and the respective standard deviations over 5 different runs. Each model was trained in a 5-fold cross-validation and in a cross-center setting using the 3 main contributing centers for training and the remaining as hold out test set.

Evaluation	Model	Acc	F1	AUC
External	Shao et al. [31]	0.884	–	0.961
	Wang et al. [9]	0.912	–	0.973
5-fold	ImgNet	0.838 ± 0.02	0.834 ± 0.02	0.926 ± 0.02
	CTP	0.902 ± 0.02	0.899 ± 0.02	0.965 ± 0.01
	TC-Swin	0.878 ± 0.02	0.875 ± 0.02	0.949 ± 0.01
	TC-Conv	0.852 ± 0.01	0.847 ± 0.01	0.925 ± 0.01
Cross-center	ImgNet	0.733 ± 0.01	0.721 ± 0.01	0.817 ± 0.01
	CTP	0.854 ± 0.01	0.850 ± 0.01	0.924 ± 0.01
	TC-Swin	0.785 ± 0.01	0.778 ± 0.01	0.863 ± 0.01
	TC-Conv	0.763 ± 0.01	0.759 ± 0.01	0.832 ± 0.01

The averaged results are shown at the bottom of the table for orientation. On average the TC-Swin encoder showed higher mean accuracy (0.953) compared to the TC-Conv (0.941) and CTP (0.944) encoders, as well as higher F1-scores. In terms of AUC, all three encoders show approximately the same performance on average.

When looking at the performance on individual clinics, the transfer from UBERN to other clinics seems to be the most challenging. ImageNet (0.734) and CTP (0.860) both showed a performance about 10 percentage points lower than the performance on other centers. However, both TC encoders were able to extract features that allowed the transfer of the learned classification to other centers (TC-Swin: 0.943, TC-Conv: 0.943). The performance when trained on data from the UBERN clinic is similar to that when trained on data from other clinics.

Looking at performance on other clinics, there is a slight discrepancy in mean accuracy and AUC score between CTP and TC. CTP consistently achieves slightly higher scores of approximately one percentage point over the TC-Swin encoder.

5.4. Lung cancer

[Table 5](#) shows the results of the 5-fold cross-validation and the cross-center evaluation on the TCGA-NSCLC dataset. Despite the absence of lung cancer images in the pre-training, both TC encoders outperform the ImageNet weights in the cross-center and 5-fold cross-validation tasks. The CTP encoder achieves the highest mean accuracy of 0.902, which is about 2 percentage points better than the TC-Swin of 0.878. The CTP encoder was pretrained using the TCGA-NSCLC dataset and its performance achieved when evaluating the CTP in our pipeline is similar to the external baseline reported by Wang et al. [9] (0.902 vs. 0.912). The TC-Swin encoder shows comparable performance earlier results reported by Shao et al. [31], with an accuracy of 0.878 vs. 0.884 and an AUC of 0.949 vs. 0.961.

As expected, both models TC-models show weaker performance when evaluated across centers, compared to mixed-center 5-fold cross-validation. Both TC encoders outperform the ImageNet-based encoder. The CTP encoder was pre-trained to extract features from this dataset and outperforms all other models in the cross-center evaluation.

As an additional experiment, the UNI [24] encoder was compared to the other models using the ABMIL aggregation head and showed similar results to the other encoders. The detailed results can be found in [Appendix B.1](#).

6. Discussion

This paper demonstrates that MTL is an effective method for training foundation models on supervised signals in CPath. We tested the

MTL-trained Tissue Concepts encoders on breast, prostate, colorectal, and lung cancers and found that the performance of the encoders is comparable to a self-supervised model. This comparable performance was achieved with only 6% of the data and resources: training TC required only 912,000 patches, compared to 15 million patches used in [9]’s self-supervised approach. Less data and shorter training enable faster development and research cycles.

Further development and scaling of the foundation models in CPath will inevitably contribute to extended CO₂ emissions. The TC encoders were trained for 160 h on a single Nvidia RTX A5000 in Europe. This corresponds to an estimated 18.91 kg of CO₂ emissions [65]. Most self-supervised approaches exceed these emissions by orders of magnitude due to larger training datasets, longer convergence times, and the use of multiple GPUs. Training such a model for the same amount of time on 48 NVIDIA V100s, a common number of GPUs for training large foundation models, results in 2004 kg of CO₂ during training. Overall, depending on the GPU, one training of an MTL-based model produces only 0.9% to 2.25% of the CO₂ emissions of an SSL-based training while providing similar performance.

Even though, performance generally is comparable, there are two differences between TC and CTP in organ specific performance. One is the difference on the prostate dataset. While a small amount of patches of this dataset were present during the pre-training of original TC encoder, the vast amount of patches used during testing were unseen. Overall, a lot of prostate patches were used during the pre-training of TC, as mentioned in Section 3.1. This large proportion of prostate tissue during pre-training may have led to the better performance on this organ and it likely prevented a larger reduction of the encoder’s performance when omitting the small amount of PANDA data during pre-training. This raises the question if organ-specific fine-tuning of TC encoders can lead to better performance on specific organs when considering cross-center evaluation and testing on unseen datasets. This question is further justified by the differences in organ specific performance between foundation models and external baselines.

Another clear difference in performance between CTransPath and Tissue Concepts was observed for lung cancer in the 5-fold cross-validation and cross-center validation. This difference most likely arises due to the out-of-domain nature of lung cancer tissue in TC pre-training. However, training on pathology images instead of non-medical images improved performance when comparing Tissue Concepts weights to ImageNet weights. Scaling the Tissue Concepts pre-training with more data should be considered to overcome the current limitations. Importantly, balancing different tasks, tissue types, and magnifications during pre-training is important to effectively scale and apply multi-task learning in the CPath domain.

In the present study, only frozen encoders were considered. While fine-tuning these foundation models can lead to better adapting to specific tasks, more research is needed on the balance between data specificity and variation, and the amount of data is required. In [17] we found that a fine-tuned foundation model encoder outperformed a fine-tuned ImageNet-base model on all tasks in terms of data efficiency and F1-score. Similar results can be expected for TC encoders and CPath specific problems, but this remains an open question to be investigated further. Additionally, this study focuses on solving MIL problems which occur frequently in diagnostic or prognostic tasks. While MIL tasks are clinically motivated, the individual performance of the segmentation, and detection branches in the multi-task trained model need to be considered in future evaluations.

Overall, the results show that the more domain-specific pre-training of Tissue Concepts encoders is advantageous for solving domain-specific tasks compared to the more general, multi-domain pre-trained encoders. A question that remains to be answered is the performance difference between foundation models and models trained on one specific organ. This will be part of future research.

7. Conclusion

In this paper, we propose to train a foundation model for CPath on supervised signals using multi-task learning to reduce the need for large corpora of data, computation time, and resources during training. The proposed method shows comparable results to a model trained using self-supervision while relying on a fraction of the training patches.

In addition, we found that, although trained on large amounts of data, existing models still exhibit a loss of performance when applied across centers. Better cross-center generalization is crucial to facilitate broader clinical application and needs to be further addressed in future research.

In order to effectively scale multi-task learning of foundation models, questions regarding the balance of tasks and tissue types need to be answered. Organ-specific fine-tuning might create robust and high-performing encoders for specific problems but needs to be further explored.

CRediT authorship contribution statement

Till Nicke: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Jan Raphael Schäfer:** Software, Methodology, Investigation, Data curation, Conceptualization. **Henning Höfener:** Writing – review & editing, Supervision, Project administration, Investigation, Funding acquisition. **Friedrich Feuerhake:** Writing – review & editing, Supervision, Project administration, Funding acquisition. **Dorit Merhof:** Writing – review & editing, Supervision, Project administration, Funding acquisition. **Fabian Kießling:** Writing – review & editing, Supervision, Project administration, Conceptualization. **Johannes Lotz:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: The applicant ‘Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung eingetragener Verein’ has a patent pending related to the training algorithm and neural architecture components presented in this Article (patent application no. EP23209015.9; names of inventors, R.S., T.N., H.H., J.L., F.K.). If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research was in part funded by the German ministry of education and research (BMBF) through the projects SynDICAD (01IS21067) and PROSurvival (01KD2213). The authors are responsible for the content of this publication. The results show in this paper are in part based upon data generated by the TCGA Research Network: <https://www.cancer.gov/tcga>.

Appendix A. Training datasets

NCT-CRC-HE 100k

The NCT-CRC-HE 100k dataset consists of 100,000 non-overlapping patches from 86 H&E stained human colorectal cancer and normal WSIs. The patches are of size 224 × 224 and have a resolution of 0.5 μm per pixel (MPP). Nine different tissue classes were present in this classification dataset. The dataset is available under <https://zenodo.org/records/1214456>.

Conic

This dataset is based in the Lizard Dataset [40]. It contains 4981 H&E stained, non-overlapping patches of colon tissue at 256×256 pixels. The segmentation masks were generated using the Hover-Net. The dataset is available under <https://conic-challenge.grand-challenge.org/>.

CRAG

The CRAG dataset consists of 213 H&E stained patches from colon tissue of 38 WSIs. The patches are of size 1512×1512 at a resolution of 0.5 MPP. During preprocessing the patches and masks were further cropped into non overlapping 224×224 tiles with 1MPP resolution. The original 173/40 patch train/val split was maintained. The dataset is available under <https://warwick.ac.uk/fac/sci/dcs/research/tia/data/mildnet> with corresponding login.

SemiCol

The training data from the SemiCOL challenge is divided into two parts. One contains manual segmentation masks, which were used in the pre-training of the encoder. The other part contains weakly labeled data, which was described in 3.1. The 20 sparsely annotated slides were obtained from the university hospital in Cologne (Hamamatsu S360) and from the university hospital LMU in Munich (Leica GT450). Both sides scanned at 0.5 MPP and provided 10 slides each. Patches with annotations were extracted from the provided images and scaled to 1 MPP and 224×224 pixels. More details about the data can be found here <https://www.semicol.org/data/>.

Arvaniti

The Arvaniti TMA dataset consists of prostate tissue microarrays, which were scanned at 0.23 MPP resolution at the University Hospital Zurich using a Hamamatsu NanoZoomer-XR Digital slide scanner. Patches of 1024×1024 were cut out from the foreground and scaled to 224×224 . The major class of the mask was selected to serve as patch label. The dataset is available under <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/OCYCM>.

Peso

The PESO dataset consists of 102 WSIs, which were scanned at 20x around 0.48 MPP with a 3DHistech Pannoramic Flash II 250 scanner. From the H&E stained prostate slides were, patches were extracted and scaled to 1 MPP of size 224×224 . The corresponding masks were used as targets. The dataset is available under <https://zenodo.org/records/1485967>.

Schoemig-Markiewka

This prostate dataset contains 6 sub-datasets, each containing 120.000 patches. All of the patches were scanned at approximately 0.25 MPP with at least 4 different scanners. Each sub-dataset contains 50.000 patches with tumor tissue, 50.000 patches with non-neoplastic glandular prostate tissue and 20.000 patches with non-glandular tissue. In our training the patches were scaled accordingly. 5 of the sub-datasets were used for pre-training, while 1 was used for validation. The dataset is available under: <http://zenodo.org>, Deposits: 4 789 576 (Dataset 1–4) and 4 904 569 (Datasets 5–6).

Panda

This prostate dataset was already described in 3.1. From the slides, patches of 224×224 were extracted. The labels for the patches were generated from the corresponding masks. The dataset is available under <https://www.kaggle.com/c/prostate-cancer-grade-assessment/data>.

TUH

The Temple University digital pathology corpus consists of over 3505 annotated images of breast tissue. The labels range from artifact and background annotations to specific breast cancer annotations like invasive ductal carcinoma. All WSIs of the 296 patients are scanned at 0.5 MPP with corresponding annotations. From the 3505 slides 136 exhibited quality annotations that were used in the training. Fixed 20 patches per WSI were sampled. Details about the download can be found here https://isip.piconepress.com/projects/nsf_dpath/html/downloads.shtml.

Tiger

This breast cancer dataset contains H&E stained patches from tumor infiltrating lymphocytes. The patches were extracted from 151 TCGA-BRCA slides at 0.25 MPP. Manual annotations for 7 different classes were used as masks. Patches and masks were scaled to 224×224 pixels. The dataset is available under <https://tiger.grand-challenge.org/Data/>.

BCSS

The BCSS dataset was contained as part of the TIGER challenge. Annotations provided through the challenge were used as segmentations masks. These annotations were derived from the original BCSS dataset. All images were scaled same as patches from the Tiger dataset mentioned above. 124 slides from BCSS and NuCLS were annotated, where some of the annotations were grouped. The dataset is available under <https://tiger.grand-challenge.org/Data/>.

BreakHis

The breakhis dataset contains H&E stained patches from 82 patients at different magnifications. The patches are of size 700×460 and were scaled accordingly to 1 MPP resolution. The classes for this dataset were derived from the original dataset. A two-class benign/malignant split is possible, however, the more fine-grained 8 class dataset was used. This dataset is available under <https://web.inf.ufpr.br/vri/databases/breast-cancer-histopathological-database-breakhis/>.

MiDoG

The MiDOG 2022 challenge dataset contains 405 tumor cases across six different tumor types. During training the 44 lung cancer cases were excluded. The cases were patchified and scaled to 224×224 pixels with approximately 1 MPP. The dataset is available under <https://zenodo.org/records/6547151>.

HubMap

The HubMap dataset contains 351 cases from different organs. Patches were extracted from the foreground containing annotation masks and scaled to 1 MPP of 224×224 pixels. Only large intestine was used during training. The dataset can be found under <https://www.kaggle.com/datasets/dingyan/hubmap-data>.

Appendix B. MIL head architecture

To aggregate the latent WSIs, as described in Section 2.2, a small CNN with global max pooling, based on Tellez et al. [32] was used. The initial layer utilized a 1×1 convolution, reducing the latent WSIs' 768 channel to 16. Following this, 3×3 convolutions with a stride of 2, and padding of 1 doubling in depth each time were applied. Instance normalization, leaky ReLU activation, and 10% 2D dropout were applied in the convolutional layers. The lowest feature pyramid of 64 channels was globally max pooled, followed by one hidden linear layer and a classification layer. This resulted in 19 250 total trainable parameters. The convolutional layers were initialized by drawing from a Kaiming normal distribution with a fixed seed.

B.1. ABMIL results and comparison to the UNI encoder

A different aggregation approach was proposed by Ilse et al. [34] under the term “Attention-based Deep Multiple Instance Learning”. For comparison, the following tables present the results obtained by using an attention based multiple instance learning aggregation head instead of a maximum pooling aggregation head. The results were collected over three runs and averaged. Additionally, the experiments were repeated using the UNI [24] foundation model.

Overall, we observed similar results between the different encoders and aggregation techniques across all organs (see Tables B.6–B.9).

Data availability

The trained models, scripts, and notebooks to reproduce the results will be made publicly available upon publication. The pre-training framework to build the pipeline will be available as a python package. The links to all datasets are mentioned in the appendix with a corresponding description.

Table B.6

Breast: Tissue-specific model performance when applied to the BRACS challenge data. The model consist of organ-specific head and frozen encoders: ImageNet, CTP, TC-Swin, TC-Conv, UNI. The table shows mean F1 score, AUC, Balanced Accuracy, and standard deviation over 3 different runs for different problem formulations. All results were obtained using an ABMIL aggregation head.

Problem	Model	F1	AUC	b ACC
2-class	Pati et al. [60]	0.96	–	–
	ImageNet	0.888 ± 0.04	0.968 ± 0.002	0.865 ± 0.05
	CTP	0.884 ± 0.05	0.990 ± 0.001	0.853 ± 0.06
	TC-Swin	0.874 ± 0.02	0.983 ± 0.02	0.833 ± 0.05
	TC-Conv	0.893 ± 0.09	0.990 ± 0.001	0.870 ± 0.001
	UNI	0.930 ± 0.03	0.992 ± 0.006	0.901 ± 0.04
4-class	Pati et al. [60]	0.84	–	–
	ImageNet	0.283 ± 0.01	0.677 ± 0.05	0.357 ± 0.02
	CTP	0.329 ± 0.01	0.759 ± 0.01	0.414 ± 0.01
	TC-Swin	0.313 ± 0.04	0.735 ± 0.05	0.381 ± 0.04
	TC-Conv	0.276 ± 0.01	0.699 ± 0.01	0.347 ± 0.02
	UNI	0.287 ± 0.02	0.711 ± 0.06	0.357 ± 0.02
7-class	Pati et al. [60]	0.69	–	–
	ImageNet	0.189 ± 0.01	0.667 ± 0.001	0.275 ± 0.01
	CTP	0.186 ± 0.007	0.698 ± 0.001	0.278 ± 0.01
	TC-Swin	0.180 ± 0.01	0.685 ± 0.02	0.275 ± 0.009
	TC-Conv	0.115 ± 0.06	0.580 ± 0.06	0.194 ± 0.08
	UNI	0.183 ± 0.03	0.674 ± 0.02	0.269 ± 0.01

Table B.7

Prostate: Tissue-specific model performance when applied to the PANDA challenge data. The model consist of organ-specific head and frozen encoders: ImageNet, CTP, TC-Swin, TC-Conv. The table shows mean weighted Accuracy, F1 score, AUC, balanced Accuracy, and standard deviations over 3 different runs for 5-fold cross-validation and cross-center transfer performance. All results were obtained using an ABMIL aggregation head.

Clinic	Model	Acc	F1	AUC	Balanced Acc
External	Mahdi Behzadi et al. [63]	0.95	0.93	–	–
5 fold	ImgNet	0.823 ± 0.07	0.793 ± 0.06	0.900 ± 0.03	0.819 ± 0.06
	CTP	0.865 ± 0.03	0.834 ± 0.05	0.924 ± 0.04	0.847 ± 0.06
	TC-Swin	0.923 ± 0.01	0.905 ± 0.02	0.970 ± 0.01	0.913 ± 0.03
	TC-Conv	0.923 ± 0.02	0.904 ± 0.02	0.9690.02	0.914 ± 0.03
	UNI	0.944 ± 0.01	0.931 ± 0.01	0.982 ± 0.01	0.939 ± 0.02
Karolinska → Radboud	ImgNet	0.595 ± 0.07	0.576 ± 0.06	0.662 ± 0.07	0.618 ± 0.07
	CTP	0.819 ± 0.02	0.605 ± 0.02	0.825 ± 0.01	0.627 ± 0.02
	TC-Swin	0.902 ± 0.01	0.857 ± 0.01	0.948 ± 0.01	0.904 ± 0.01
	TC-Conv	0.726 ± 0.01	0.684 ± 0.01	0.919 ± 0.01	0.816 ± 0.01
	UNI	0.748 ± 0.01	0.702 ± 0.01	0.893 ± 0.02	0.825 ± 0.01
Radboud → Karolinska	ImgNet	0.572 ± 0.07	0.572 ± 0.07	0.668 ± 0.04	0.626 ± 0.07
	CTP	0.676 ± 0.03	0.583 ± 0.07	0.729 ± 0.01	0.595 ± 0.03
	TC-Swin	0.896 ± 0.01	0.891 ± 0.01	0.969 ± 0.04	0.915 ± 0.01
	TC-Conv	0.654 ± 0.05	0.676 ± 0.06	0.898 ± 0.01	0.902 ± 0.01
	UNI	0.662 ± 0.01	0.661 ± 0.01	0.871 ± 0.02	0.738 ± 0.02

Table B.8

Colorectal: Tissue-specific model performance when applied to the SemiCOL challenge data. The model consist of organ-specific MIL head and frozen encoders: ImageNet, CTP, TC-Swin, TC-Conv, UNI. The table shows mean accuracy, F1 score, AUC, Balanced Accuracy, and the respective standard deviations over 3 different runs. Each model was trained on images from one site and evaluated on all images from the remaining sites. All results were obtained using an ABMIL aggregation head.

Clinic	Model	Acc	F1	Auc	Balanced Acc
LMU	ImgNet	0.678 ± 0.12	0.646 ± 0.16	0.741 ± 0.16	0.678 ± 0.12
	CTP	0.925 ± 0.01	0.925 ± 0.01	0.985 ± 0.005	0.925 ± 0.01
	TC-Swin	0.873 ± 0.06	0.871 ± 0.07	0.934 ± 0.07	0.873 ± 0.06
	TC-Conv	0.913 ± 0.06	0.912 ± 0.06	0.981 ± 0.02	0.913 ± 0.06
	UNI	0.807 ± 0.03	0.802 ± 0.04	0.939 ± 0.02	0.809 ± 0.03
UBERN	ImgNet	0.793 ± 0.18	0.793 ± 0.18	0.838 ± 0.20	0.793 ± 0.18
	CTP	0.820 ± 0.03	0.819 ± 0.04	0.914 ± 0.02	0.820 ± 0.03
	TC-Swin	0.915 ± 0.01	0.915 ± 0.01	0.973 ± 0.01	0.915 ± 0.01
	TC-Conv	0.947 ± 0.01	0.947 ± 0.01	0.984 ± 0.01	0.947 ± 0.01
	UNI	0.905 ± 0.04	0.904 ± 0.04	0.956 ± 0.03	0.905 ± 0.04
UKK	ImgNet	0.865 ± 0.064	0.864 ± 0.064	0.927 ± 0.067	0.866 ± 0.064
	CTP	0.901 ± 0.104	0.897 ± 0.109	0.979 ± 0.024	0.901 ± 0.103
	TC-Swin	0.809 ± 0.099	0.799 ± 0.112	0.965 ± 0.027	0.809 ± 0.098

(continued on next page)

Table B.8 (continued).

Clinic	Model	Acc	F1	Auc	Balanced Acc
WNS	TC-Conv	0.892 ± 0.050	0.892 ± 0.050	0.953 ± 0.050	0.892 ± 0.050
	UNI	0.835 ± 0.082	0.833 ± 0.082	0.935 ± 0.68	0.834 ± 0.081
Average	ImgNet	0.801 ± 0.22	0.801 ± 0.22	0.857 ± 0.22	0.801 ± 0.22
	CTP	0.962 ± 0.01	0.962 ± 0.01	0.992 ± 0.02	0.962 ± 0.01
	TC-Swin	0.949 ± 0.02	0.949 ± 0.02	0.993 ± 0.02	0.949 ± 0.02
	TC-Conv	0.960 ± 0.01	0.960 ± 0.01	0.994 ± 0.01	0.960 ± 0.01
	UNI	0.964 ± 0.02	0.964 ± 0.02	0.993 ± 0.01	0.964 ± 0.02
Lung	ImgNet	0.783 ± 0.16	0.774 ± 0.16	0.839 ± 0.15	0.783 ± 0.14
	CTP	0.902 ± 0.07	0.901 ± 0.07	0.968 ± 0.04	0.902 ± 0.07
	TC-Swin	0.887 ± 0.09	0.884 ± 0.09	0.966 ± 0.07	0.887 ± 0.09
	TC-Conv	0.900 ± 0.08	0.899 ± 0.08	0.955 ± 0.04	0.900 ± 0.08
	UNI	0.88 ± 0.08	0.880 ± 0.08	0.958 ± 0.04	0.882 ± 0.08

Table B.9

Lung: Tissue-specific model performance when applied to the TCGA-NSCLC dataset. The model consist of organ-specific MIL head and frozen encoders: ImageNet, CTP, TC-Swin, TC-Conv, UNI. The table shows mean accuracy, F1 score, AUC, Balanced Accuracy, and the respective standard deviations over 3 different runs. Each model was trained in a 5-fold cross-validation and in a cross-center setting using the 3 main contributing centers for training and the remaining as hold out test set. All results were obtained using an ABMIL aggregation head.

Evaluation	Model	Acc	F1	AUC	Balanced Acc
External	Shao et al. [31]	0.884	–	0.961	–
	Wang et al. [9]	0.912	–	0.973	–
5-fold	ImgNet	0.803 ± 0.0001	0.807 ± 0.0001	0.0890 ± 0.0001	0.801 ± 0.0001
	CTP	0.889 ± 0.001	0.888 ± 0.001	0.954 ± 0.001	0.887 ± 0.001
	TC-Swin	0.843 ± 0.001	0.830 ± 0.001	0.947 ± 0.001	0.840 ± 0.001
	TC-Conv	0.865 ± 0.001	0.864 ± 0.001	0.933 ± 0.001	0.865 ± 0.001
	UNI	0.924 ± 0.001	0.923 ± 0.001	0.970 ± 0.001	0.920 ± 0.001
Cross-center	ImgNet	0.723 ± 0.0001	0.720 ± 0.0001	0.777 ± 0.0001	0.720 ± 0.0001
	CTP	0.804 ± 0.0001	0.802 ± 0.0001	0.891 ± 0.0001	0.801 ± 0.0001
	TC-Swin	0.786 ± 0.0001	0.772 ± 0.0001	0.859 ± 0.0001	0.771 ± 0.0001
	TC-Conv	0.720 ± 0.0001	0.718 ± 0.0001	0.784 ± 0.0001	0.719 ± 0.0001
	UNI	0.857 ± 0.0001	0.855 ± 0.0001	0.917 ± 0.0001	0.852 ± 0.0001

References

- M. Moscalu, R. Moscalu, C.G. Dascălu, V. Țarcă, E. Cojocaru, I.M. Costin, E. Țarcă, I.L. Șerban, Histopathological images analysis and predictive modeling implemented in digital pathology—Current affairs and perspectives, *Diagnostics* 13 (14) (2023) 2379, <http://dx.doi.org/10.3390/diagnostics13142379>, URL <https://www.mdpi.com/2075-4418/13/14/2379>, Number: 14 Publisher: Multidisciplinary Digital Publishing Institute.
- H. Sung, J. Ferlay, R.L. Siegel, M. Laversanne, I. Soerjomataram, A. Jemal, F. Bray, Global cancer statistics 2020: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries, *CA: Cancer J. Clin.* 71 (3) (2021) 209–249, <http://dx.doi.org/10.3322/caac.21660>, URL <https://onlinelibrary.wiley.com/doi/abs/10.3322/caac.21660> _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.3322/caac.21660>.
- L. Rahib, M.R. Wehner, L.M. Matrisian, K.T. Nead, Estimated projection of US cancer incidence and death to 2040, *JAMA Netw. Open* 4 (4) (2021) e214708, <http://dx.doi.org/10.1001/jamanetworkopen.2021.4708>, URL <https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2778204>.
- I. Soerjomataram, F. Bray, Planning for tomorrow: global cancer incidence and the role of prevention 2020–2070, *Nat. Rev. Clin. Oncol.* 18 (10) (2021) 663–672, <http://dx.doi.org/10.1038/s41571-021-00514-z>, URL <https://www.nature.com/articles/s41571-021-00514-z>, Publisher: Nature Publishing Group.
- R.L. Siegel, K.D. Miller, N.S. Wagle, A. Jemal, Cancer statistics, 2023, *CA: Cancer J. Clin.* 73 (1) (2023) 17–48, <http://dx.doi.org/10.3322/caac.21763>, URL <https://acsjournals.onlinelibrary.wiley.com/doi/10.3322/caac.21763>.
- Z. Zhang, P. Chen, M. McGough, F. Xing, C. Wang, M. Bui, Y. Xie, M. Sapkota, L. Cui, J. Dhillon, N. Ahmad, F.K. Khalil, S.I. Dickinson, X. Shi, F. Liu, H. Su, J. Cai, L. Yang, Pathologist-level interpretable whole-slide cancer diagnosis with deep learning, *Nat. Mach. Intell.* 1 (5) (2019) 236–245, <http://dx.doi.org/10.1038/s42256-019-0052-1>, URL <https://www.nature.com/articles/s42256-019-0052-1>, Publisher: Nature Publishing Group.
- Y. Liu, A. Jain, C. Eng, D.H. Way, K. Lee, P. Bui, K. Kanada, G. de Oliveira Marinho, J. Gallegos, S. Gabriele, V. Gupta, N. Singh, V. Natarajan, R. Hofmann-Wellenhof, G.S. Corrado, L.H. Peng, D.R. Webster, D. Ai, S.J. Huang, Y. Liu, R.C. Dunn, D. Coz, A deep learning system for differential diagnosis of skin diseases, *Nature Med.* 26 (6) (2020) 900–908, <http://dx.doi.org/10.1038/s41591-020-0842-3>, URL <https://www.nature.com/articles/s41591-020-0842-3>, Publisher: Nature Publishing Group.
- S. Zhang, D. Metaxas, On the challenges and perspectives of foundation models for medical image analysis, 2023, <http://dx.doi.org/10.48550/arXiv.2306.05705>, URL <http://arxiv.org/abs/2306.05705> arXiv:2306.05705 [cs, eess].
- X. Wang, S. Yang, J. Zhang, M. Wang, J. Zhang, W. Yang, J. Huang, X. Han, Transformer-based unsupervised contrastive learning for histopathological image classification, *Med. Image Anal.* 81 (2022) 102559, <http://dx.doi.org/10.1016/j.media.2022.102559>, URL <https://linkinghub.elsevier.com/retrieve/pii/S1361841522002043>.
- G. Campanella, R. Kwan, E. Fluder, J. Zeng, A. Stock, B. Veremis, A.D. Polydorides, C. Hedvat, A. Schoenfeld, C. Vanderbilt, P. Kovatch, C. Cordon-Cardo, T.J. Fuchs, Computational pathology at health system scale – self-supervised foundation models from three billion images, 2023, <http://dx.doi.org/10.48550/arXiv.2310.07033>, URL <http://arxiv.org/abs/2310.07033> arXiv:2310.07033 [cs, eess].
- P. Dhar, The carbon impact of artificial intelligence, *Nat. Mach. Intell.* 2 (8) (2020) 423–425, <http://dx.doi.org/10.1038/s42256-020-0219-9>, URL <https://www.nature.com/articles/s42256-020-0219-9>, Publisher: Nature Publishing Group.
- L.H. Kaack, P.L. Donti, E. Strubell, G. Kamiya, F. Creutzig, D. Rolnick, Aligning artificial intelligence with climate change mitigation, *Nature Clim. Change* 12 (6) (2022) 518–527, <http://dx.doi.org/10.1038/s41558-022-01377-7>, URL <https://www.nature.com/articles/s41558-022-01377-7>, Publisher: Nature Publishing Group.
- T.L.D. Health, Curbing the carbon footprint of health care, *Lancet Digit. Health* 5 (12) (2023) e848, [http://dx.doi.org/10.1016/S2589-7500\(23\)00229-7](http://dx.doi.org/10.1016/S2589-7500(23)00229-7), URL <https://www.sciencedirect.com/science/article/pii/S2589750023002297>.
- D. Pototsky, A. Sultan, L. Schmidt-Thieme, Does self-supervised pretraining really match ImageNet weights? in: 2022 IEEE 14th Image, Video, and Multidimensional Signal Processing Workshop (IVMSP), 2022, pp. 1–5, <http://dx.doi.org/10.1109/IVMSP54334.2022.9816238>, URL <https://ieeexplore.ieee.org/abstract/document/9816238>.
- Y. Huang, X. Yang, L. Liu, H. Zhou, A. Chang, X. Zhou, R. Chen, J. Yu, J. Chen, C. Chen, S. Liu, H. Chi, X. Hu, K. Yue, L. Li, V. Grau, D.-P. Fan, F. Dong, D. Ni, Segment anything model for medical images? *Med. Image Anal.* 92 (2024) 103061, <http://dx.doi.org/10.1016/j.media.2023.103061>, URL <https://www.sciencedirect.com/science/article/pii/S1361841523003213>.
- R. Caruana, Multitask learning, *Mach. Learn.* 28 (1) (1997) 41–75, <http://dx.doi.org/10.1023/A:1007379606734>.
- R. Schäfer, T. Nicke, H. Höfener, A. Lange, D. Merhof, F. Feuerhake, V. Schulz, J. Lotz, F. Kiessling, Overcoming data scarcity in biomedical imaging with a foundational multi-task model, 2023, <http://dx.doi.org/10.48550/arXiv.2311.09847>, URL <http://arxiv.org/abs/2311.09847> [cs, eess].
- F.M. Howard, J. Dolezal, S. Kochanny, J. Schulte, H. Chen, L. Heij, D. Huo, R. Nanda, O.I. Olopade, J.N. Kather, N. Cipriani, R.L. Grossman, A.T. Pearson,

The impact of site-specific digital histology signatures on deep learning model accuracy and bias, *Nature Commun.* 12 (1) (2021) 4423, <http://dx.doi.org/10.1038/s41467-021-24698-1>, URL <https://www.nature.com/articles/s41467-021-24698-1>, Number: 1 Publisher: Nature Publishing Group.

[19] R. Mormont, P. Geurts, R. Marée, Multi-task pre-training of deep neural networks for digital pathology, *IEEE J. Biomed. Health Inf.* 25 (2) (2021) 412–421, <http://dx.doi.org/10.1109/JBHI.2020.2992878>, URL <http://arxiv.org/abs/2005.02561> [cs, eess].

[20] S. Graham, Q.D. Vu, M. Jahanifar, S.E.A. Raza, F. Minhas, D. Snead, N. Rajpoot, One model is all you need: Multi-task learning enables simultaneous histology image segmentation and classification, *Med. Image Anal.* 83 (2023) 102685, <http://dx.doi.org/10.1016/j.media.2022.102685>, URL <https://www.sciencedirect.com/science/article/pii/S1361841522003139>.

[21] R. Bommansani, D.A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M.S. Bernstein, J. Bohg, A. Bosselut, E. Brunsell, E. Brynjolfsson, S. Buch, D. Card, R. Castellon, N. Chatterji, A. Chen, K. Creel, J.Q. Davis, D. Demszky, C. Donahue, M. Doumbouya, E. Durmus, S. Ermon, J. Etchemendy, K. Ethayarajh, L. Fei-Fei, C. Finn, T. Gale, L. Gillespie, K. Goel, N. Goodman, S. Grossman, N. Guha, T. Hashimoto, P. Henderson, J. Hewitt, D.E. Ho, J. Hong, K. Hsu, J. Huang, T. Icard, S. Jain, D. Jurafsky, P. Kalluri, S. Karamcheti, G. Keeling, F. Khani, O. Khattab, P.W. Koh, M. Krass, R. Krishnan, R. Kuditipudi, A. Kumar, F. Ladakh, M. Lee, T. Lee, J. Leskovec, I. Levent, X.L. Li, X. Li, T. Ma, A. Malik, C.D. Manning, S. Mirchandani, E. Mitchell, Z. Munyikwa, S. Nair, A. Narayan, D. Narayanan, B. Newman, A. Nie, J.C. Niebles, H. Nilforoshan, J. Nyarko, G. Ogun, L. Orr, I. Papadimitriou, J.S. Park, C. Piech, E. Portelance, C. Potts, A. Raghunathan, R. Reich, H. Ren, F. Rong, Y. Roohani, C. Ruiz, J. Ryan, C. Ré, D. Sadigh, S. Sagawa, K. Santhanam, A. Shih, K. Srinivasan, A. Tamkin, R. Taori, A.W. Thomas, F. Tramèr, R.E. Wang, W. Wang, B. Wu, J. Wu, Y. Wu, S.M. Xie, M. Yasunaga, J. You, M. Zaharia, M. Zhang, T. Zhang, X. Zhang, Y. Zhang, L. Zheng, K. Zhou, P. Liang, On the opportunities and risks of foundation models, 2022, URL <http://arxiv.org/abs/2108.07258> arXiv:2108.07258 [cs].

[22] N. Fei, Z. Lu, Y. Gao, G. Yang, Y. Huo, J. Wen, H. Lu, R. Song, X. Gao, T. Xiang, H. Sun, J.-R. Wen, Towards artificial general intelligence via a multimodal foundation model, *Nature Commun.* 13 (1) (2022) 3094, <http://dx.doi.org/10.1038/s41467-022-30761-2>, URL <https://www.nature.com/articles/s41467-022-30761-2>, Publisher: Nature Publishing Group.

[23] X. Chen, S. Xie, K. He, An empirical study of training self-supervised vision transformers, 2021, <http://dx.doi.org/10.48550/arXiv.2104.02057>, URL <http://arxiv.org/abs/2104.02057> arXiv:2104.02057 [cs].

[24] R.J. Chen, T. Ding, M.Y. Lu, D.F. Williamson, G. Jaume, A.H. Song, B. Chen, A. Zhang, D. Shao, M. Shaban, et al., Towards a general-purpose foundation model for computational pathology, *Nature Med.* 30 (3) (2024) 850–862.

[25] M. Oquab, T. Darct, T. Moutakanni, H.V. Vo, M. Szafrańiec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby, R. Howes, P.-Y. Huang, H. Xu, V. Sharma, S.-W. Li, W. Galuba, M. Rabbat, M. Assran, N. Ballas, G. Synnaeve, I. Misra, H. Jegou, J. Mairal, P. Labatut, A. Joulin, P. Bojanowski, DINOv2: Learning robust visual features without supervision, 2023.

[26] G. Campanella, M.G. Hanna, L. Geneslaw, A. Miraflor, V. Werneck Krauss Silva, K.J. Busam, E. Brogi, V.E. Reuter, D.S. Klimstra, T.J. Fuchs, Clinical-grade computational pathology using weakly supervised deep learning on whole slide images, *Nature Med.* 25 (8) (2019) 1301–1309, <http://dx.doi.org/10.1038/s41591-019-0508-1>, URL <https://www.nature.com/articles/s41591-019-0508-1>, Publisher: Nature Publishing Group.

[27] N. Ghaffari Laleh, H.S. Muti, C.M.L. Loeffler, A. Echle, O.L. Saldanha, F. Mahmood, M.Y. Lu, C. Trautwein, R. Langer, B. Dislich, R.D. Buelow, H.I. Grabsch, H. Brenner, J. Chang-Claude, E. Alwers, T.J. Brinker, F. Khader, D. Truhn, N.T. Gaisa, P. Boor, M. Hoffmeister, V. Schulz, J.N. Kather, Benchmarking weakly-supervised deep learning pipelines for whole slide classification in computational pathology, *Med. Image Anal.* 79 (2022) 102474, <http://dx.doi.org/10.1016/j.media.2022.102474>, URL <https://www.sciencedirect.com/science/article/pii/S1361841522001219>.

[28] P. Liu, L. Ji, F. Ye, B. Fu, AdvMIL: Adversarial multiple instance learning for the survival analysis on whole-slide images, *Med. Image Anal.* 91 (2024) 103020, <http://dx.doi.org/10.1016/j.media.2023.103020>, URL <https://www.sciencedirect.com/science/article/pii/S1361841523002803>.

[29] L. Godson, N. Alemi, J. Nsengimana, G.P. Cook, E.L. Clarke, D. Treanor, D.T. Bishop, J. Newton-Bishop, A. Gooya, D. Magee, Immune subtyping of melanoma whole slide images using multiple instance learning, *Med. Image Anal.* 93 (2024) 103097, <http://dx.doi.org/10.1016/j.media.2024.103097>, URL <https://www.sciencedirect.com/science/article/pii/S1361841524000227>.

[30] M.Y. Lu, D.F.K. Williamson, T.Y. Chen, R.J. Chen, M. Barbieri, F. Mahmood, Data-efficient and weakly supervised computational pathology on whole-slide images, *Nat. Biomed. Eng.* 5 (6) (2021) 555–570, <http://dx.doi.org/10.1038/s41551-020-00682-w>, URL <https://www.nature.com/articles/s41551-020-00682-w>, Number: 6 Publisher: Nature Publishing Group.

[31] Z. Shao, H. Bian, Y. Chen, Y. Wang, J. Zhang, X. Ji, y. zhang, TransMIL: Transformer based correlated multiple instance learning for whole slide image classification, in: *Advances in Neural Information Processing Systems*, Vol. 34, Curran Associates, Inc., 2021, pp. 2136–2147, URL https://proceedings.neurips.cc/paper_files/paper/2021/hash/10c272d06794d3e5785d5e7c5356e9ff-Abstract.html.

[32] D. Tellez, G. Litjens, J. Van Der Laak, F. Ciompi, Neural image compression for gigapixel histopathology image analysis, *IEEE Trans. Pattern Anal. Mach. Intell.* 43 (2) (2021) 567–578, <http://dx.doi.org/10.1109/TPAMI.2019.2936841>, URL <https://ieeexplore.ieee.org/document/8809829>.

[33] D. Tellez, D. Höppener, C. Verhoef, D. Grünhagen, P. Nierop, M. Drozdzał, J. Laak, F. Ciompi, Extending unsupervised neural image compression with supervised multitask learning, in: *Proceedings of the Third Conference on Medical Imaging with Deep Learning*, PMLR, 2020, pp. 770–783, URL <https://proceedings.mlr.press/v121/tellez20a.html>, ISSN: 2640-3498.

[34] M. Ilse, J. Tomczak, M. Welling, Attention-based deep multiple instance learning, in: *International Conference on Machine Learning*, PMLR, 2018, pp. 2127–2136.

[35] J.N. Kather, J. Krisam, P. Charoentong, T. Luedde, E. Herpel, C.-A. Weis, T. Gaiser, A. Marx, N.A. Valous, D. Ferber, L. Jansen, C.C. Reyes-Aldasoro, I. Zörnig, D. Jäger, H. Brenner, J. Chang-Claude, M. Hoffmeister, N. Halama, Predicting survival from colorectal cancer histology slides using deep learning: A retrospective multicenter study, *PLoS Med.* 16 (1) (2019) e1002730, <http://dx.doi.org/10.1371/journal.pmed.1002730>, URL <https://journals.plos.org/plosmedicine/article?id=10.1371/journal.pmed.1002730>, Publisher: Public Library of Science.

[36] S. Graham, H. Chen, J. Gamper, Q. Dou, P.-A. Heng, D. Snead, Y.W. Tsang, N. Rajpoot, MILD-Net: Minimal information loss dilated network for gland instance segmentation in colon histology images, *Med. Image Anal.* 52 (2019) 199–211, <http://dx.doi.org/10.1016/j.media.2018.12.001>, URL <http://arxiv.org/abs/1806.01963> [cs].

[37] R. Awan, K. Sirinunukwattana, D. Epstein, S. Jefferyes, U. Qidwai, Z. Aftab, I. Mujeeb, D. Snead, N. Rajpoot, Glandular morphometrics for objective grading of colorectal adenocarcinoma histology images, *Sci. Rep.* 7 (1) (2017) 16852, <http://dx.doi.org/10.1038/s41598-017-16516-w>, URL <https://www.nature.com/articles/s41598-017-16516-w>, Publisher: Nature Publishing Group.

[38] Y. Tolkach, A. Mukhopadhyay, A. Sanner, N. Zerbe, Home | SemiCOL challenge — semicol.org, 2023, <https://www.semicol.org/>. (Accessed 08 March 2024).

[39] S. Graham, Q.D. Vu, M. Jahanifar, M. Weigert, U. Schmidt, W. Zhang, J. Zhang, S. Yang, J. Xiang, X. Wang, J.L. Rumberger, E. Baumann, P. Hirsch, L. Liu, C. Hong, A.I. Aviles-Rivero, A. Jain, H. Ahn, Y. Hong, H. Azzuni, M. Xu, M. Yaquib, M.-C. Blache, B. Piégu, B. Vernay, T. Scherr, M. Böhlund, K. Löffler, J. Li, W. Ying, C. Wang, D. Snead, S.E.A. Raza, F. Minhas, N.M. Rajpoot, CoNIC challenge: Pushing the frontiers of nuclear detection, segmentation, classification and counting, *Med. Image Anal.* 92 (2024) 103047, <http://dx.doi.org/10.1016/j.media.2023.103047>, URL <https://www.sciencedirect.com/science/article/pii/S1361841523003079>.

[40] S. Graham, M. Jahanifar, A. Azam, M. Nimir, Y.-W. Tsang, K. Dodd, E. Hero, H. Sahota, A. Tank, K. Benes, N. Wahab, F. Minhas, S.E.A. Raza, H. El Daly, K. Gopalakrishnan, D. Snead, N.M. Rajpoot, Lizard: A large-scale dataset for colonic nuclear instance segmentation and classification, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, 2021, pp. 684–693, URL https://openaccess.thecvf.com/content/ICCV2021W/CDPath/html/Graham_Lizard_A_Large-Scale_Dataset_for_Colonic_Nuclear_Instance_Segmentation_and_ICCV_2021_paper.html.

[41] E. Arvaniti, K.S. Fricker, M. Moret, N. Rupp, T. Hermanns, C. Fankhauser, N. Wey, P.J. Wild, J.H. Rüschoff, M. Claassen, Automated Gleason grading of prostate cancer tissue microarrays via deep learning, *Sci. Rep.* 8 (1) (2018) 12054, <http://dx.doi.org/10.1038/s41598-018-30535-1>, URL <https://www.nature.com/articles/s41598-018-30535-1>, Publisher: Nature Publishing Group.

[42] W. Bulten, P. Bánki, J. Hoven, R.v.d. Loo, J. Lotz, N. Weiss, J.v.d. Laak, B.v. Ginneken, C. Hulsbergen-van de Kaa, G. Litjens, Epithelium segmentation using deep learning in H&E-stained prostate specimens with immunohistochemistry as reference standard, *Sci. Rep.* 9 (1) (2019) 864, <http://dx.doi.org/10.1038/s41598-018-37257-4>, URL <https://www.nature.com/articles/s41598-018-37257-4>, Publisher: Nature Publishing Group.

[43] B. Schömg-Markiewka, A. Pryalukhin, W. Hulla, A. Bychkov, J. Fukuo, A. Madabhushi, V. Achter, L. Nieroda, R. Büttner, A. Quaas, Y. Tolkach, Quality control stress test for deep learning-based diagnostic model in digital pathology, *Mod. Pathol.* 34 (12) (2021) 2098–2108, <http://dx.doi.org/10.1038/s41379-021-00859-x>, URL <https://linkinghub.elsevier.com/retrieve/pii/S0893395222003702>.

[44] W. Bulten, K. Kartasalo, P.-H.C. Chen, P. Ström, H. Pinckaers, K. Nagpal, Y. Cai, D.F. Steiner, H. van Boven, R. Vink, C. Hulsbergen-van de Kaa, J. van der Laak, M.B. Amin, A.J. Evans, T. van der Kwast, R. Allan, P.A. Humphrey, H. Grönberg, H. Samarantunga, B. Delahunt, T. Tsuzuki, T. Häkkinen, L. Egevad, M. Demkin, S. Dane, F. Tan, M. Valkonen, G.S. Corrado, L. Peng, C.H. Mermel, P. Ruusuvuori, G. Litjens, M. Eklund, Artificial intelligence for diagnosis and Gleason grading of prostate cancer: the PANDA challenge, *Nature Med.* 28 (1) (2022) 154–163, <http://dx.doi.org/10.1038/s41591-021-01620-2>, URL <https://www.nature.com/articles/s41591-021-01620-2>, Number: 1 Publisher: Nature Publishing Group.

[45] Z. Wevodau, B. Doshna, N. Jhala, I. Akhtar, I. Obeid, J. Picone, The temple university digital pathology corpus: The breast tissue subset, in: *2021 IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*, IEEE, Philadelphia, PA, USA, 2021, pp. 1–4, <http://dx.doi.org/10.1109/SPMB52430.2021.9672275>, URL <https://ieeexplore.ieee.org/document/9672275>.

[46] A. Shephard, M. Jahanifar, R. Wang, M. Dawood, S. Graham, K. Sidlauskas, S.A. Khurram, N. Rajpoot, S.E.A. Raza, TiAger: Tumor-infiltrating lymphocyte scoring in breast cancer for the TiGER challenge, 2022, <http://dx.doi.org/10.48550/arXiv.2206.11943>, URL <http://arxiv.org/abs/2206.11943> [cs, eess].

[47] M. Amgad, H. Elfandy, H. Hussein, L.A. Atteya, M.A.T. Elsebaie, L.S. Abo Elnasr, R.A. Sakr, H.S.E. Salem, A.F. Ismail, A.M. Saad, J. Ahmed, M.A.T. Elsebaie, M. Rahman, I.A. Ruhban, N.M. Elgazar, Y. Alagha, M.H. Osman, A.M. Alhusseiny, M.M. Khalaf, A.-A.F. Younes, A. Abdulkarim, D.M. Younes, A.M. Gadallah, A.M. Elkashash, S.Y. Fala, B.M. Zaki, J. Beezley, D.R. Chittajallu, D. Manthey, D.A. Gutman, L.A.D. Cooper, Structured crowdsourcing enables convolutional segmentation of histology images, *Bioinformatics* 35 (18) (2019) 3461–3467, <http://dx.doi.org/10.1093/bioinformatics/btz083>.

[48] F.A. Spanhol, L.S. Oliveira, C. Petitjean, L. Heutte, A dataset for breast cancer histopathological image classification, *IEEE Trans. Biomed. Eng.* 63 (7) (2016) 1455–1462, <http://dx.doi.org/10.1109/TBME.2015.2496264>, URL <http://ieeexplore.ieee.org/document/7312934>.

[49] M. Aubreville, N. Stathonikos, C.A. Bertram, R. Klopffleisch, N. ter Hoeve, F. Giompi, F. Wilm, C. Marzah, T.A. Donovan, A. Maier, J. Breen, N. Ravikumar, Y. Chung, J. Park, R. Nateghi, F. Pourakpour, R.H.J. Fick, S. Ben Hadj, M. Jahanifar, A. Shephard, J. Dexl, T. Wittenberg, S. Kondo, M.W. Lafarge, V.H. Koelzer, J. Liang, Y. Wang, X. Long, J. Liu, S. Razavi, A. Khademi, S. Yang, X. Wang, R. Erber, A. Klang, K. Lipnik, P. Bolfa, M.J. Dark, G. Wasinger, M. Veta, K. Breninger, Mitosis domain generalization in histopathology images — The MIDOG challenge, *Med. Image Anal.* 84 (2023) 102699, <http://dx.doi.org/10.1016/j.media.2022.102699>, URL <https://www.sciencedirect.com/science/article/pii/S1361841522003279>.

[50] Y. Jain, L.L. Godwin, S. Joshi, S. Mandarapu, T. Le, C. Lindskog, E. Lundberg, K. Börner, Segmenting functional tissue units across human organs using community-driven development of generalizable machine learning algorithms, *bioRxiv* (2023) <http://dx.doi.org/10.1101/2023.01.05.522764>, URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9881902/>, 2023.01.05.522764.

[51] G. Aresta, T. Araújo, S. Kwok, S.S. Chennamsetty, M. Safwan, V. Alex, B. Marami, M. Prastawa, M. Chan, M. Donovan, G. Fernandez, J. Zeineh, M. Kohl, C. Walz, F. Ludwig, S. Braunewell, M. Baust, Q.D. Vu, M.N.N. To, E. Kim, J.T. Kwak, S. Galal, V. Sanchez-Freire, N. Brancati, M. Frucci, D. Riccio, Y. Wang, L. Sun, K. Ma, J. Fang, I. Kone, L. Boulmane, A. Campilho, C. Eloy, A. Polónia, P. Aguiar, BACH: Grand challenge on breast cancer histology images, 56, (ISSN: 13618415) 2019, pp. 122–139, <http://dx.doi.org/10.1016/j.media.2019.05.010>, URL <http://arxiv.org/abs/1808.04277> arXiv:1808.04277 [cs],

[52] N. Brancati, A.M. Anniciello, P. Pati, D. Riccio, G. Scognamiglio, G. Jaume, G. De Pietro, M. Di Bonito, A. Foncubierta, G. Botti, M. Gabrani, F. Feroce, M. Frucci, BRACS: A dataset for breast carcinoma subtyping in H&E histology images, *Database* 2022 (2022) baac093, <http://dx.doi.org/10.1093/database/baac093>.

[53] S. Kirk, Y. Lee, P. Kumar, J. Filippini, B. Albertina, M. Watson, K. Rieger-Christ, J. Lemmerman, The cancer genome atlas lung squamous cell carcinoma collection (TCGA-LUSC) (version 4) [Data set], 2016, <http://dx.doi.org/10.7937/K9/TCIA.2016.TYKKFMQ>, <https://www.cancerimagingarchive.net/collection/tcga-lusc/>.

[54] B. Albertina, M. Watson, C. Holback, R. Jarosz, S. Kirk, Y. Lee, K. Rieger-Christ, J. Lemmerman, The cancer genome atlas lung adenocarcinoma collection (TCGA-LUAD) (Version 4) [Data set], 2016, <http://dx.doi.org/10.7937/K9/TCIA.2016.JGNIHEP5>, <https://www.cancerimagingarchive.net/collection/tcga-luad/>.

[55] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, B. Guo, Swin Transformer: Hierarchical vision transformer using shifted windows, 2021, URL <http://arxiv.org/abs/2103.14030> arXiv:2103.14030 [cs].

[56] R. Azad, A. Kazerouni, M. Heidari, E.K. Aghdam, A. Molaei, Y. Jia, A. Jose, R. Roy, D. Merhof, Advances in medical image analysis with vision Transformers: A comprehensive review, *Med. Image Anal.* 91 (2024) 103000, <http://dx.doi.org/10.1016/j.media.2023.103000>, URL <https://www.sciencedirect.com/science/article/pii/S1361841523002608>.

[57] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, S. Xie, A ConvNet for the 2020s, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR*, 2022, pp. 11976–11986, URL https://openaccess.thecvf.com/content/CVPR2022/html/Liu_A_ConvNet_for_the_2020s_CVPR_2022_paper.html.

[58] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A.C. Berg, L. Fei-Fei, ImageNet large scale visual recognition challenge, *Int. J. Comput. Vis.* 115 (3) (2015) 211–252, <http://dx.doi.org/10.1007/s11263-015-0816-y>.

[59] I. Loshchilov, F. Hutter, Decoupled weight decay regularization, 2019, <http://dx.doi.org/10.48550/arXiv.1711.05101>, URL <http://arxiv.org/abs/1711.05101> arXiv:1711.05101 [cs, math].

[60] P. Pati, G. Jaume, A. Foncubierta-Rodríguez, F. Feroce, A.M. Anniciello, G. Scognamiglio, N. Brancati, M. Fiche, E. Dubruc, D. Riccio, M. Di Bonito, G. De Pietro, G. Botti, J.-P. Thiran, M. Frucci, O. Goksel, M. Gabrani, Hierarchical graph representations in digital pathology, *Med. Image Anal.* 75 (2022) 102264, <http://dx.doi.org/10.1016/j.media.2021.102264>, URL <https://www.sciencedirect.com/science/article/pii/S1361841521003091>.

[61] L. Maier-Hein, A. Reinke, P. Godau, M.D. Tizabi, F. Buettner, E. Christodoulou, B. Glocker, F. Isensee, J. Kleesiek, M. Kozubek, M. Reyes, M.A. Riegler, M. Wiesenfarth, A.E. Kavur, C.H. Sudre, M. Baumgartner, M. Eisenmann, D. Heckmann-Nötzel, T. Rädsch, L. Acion, M. Antonelli, T. Arbel, S. Bakas, A. Benis, M.B. Blaschko, M.J. Cardoso, V. Cheplygina, B.A. Cimini, G.S. Collins, K. Farahani, L. Ferrer, A. Galdran, B. van Ginneken, R. Haase, D.A. Hashimoto, M.M. Hoffman, M. Huisman, P. Jannin, C.E. Kahn, D. Kainmueller, B. Kainz, A. Karargyris, A. Karthikesalingam, F. Kofler, A. Kopp-Schneider, A. Kreshuk, T. Kurc, B.A. Landman, G. Litjens, A. Madani, K. Maier-Hein, A.L. Martel, P. Mattsson, E. Meijering, B. Menze, K.G.M. Moons, H. Müller, B. Nichyporuk, F. Nickel, J. Petersen, N. Rajpoot, N. Rieke, J. Saez-Rodriguez, C.I. Sánchez, S. Shetty, M. van Smeden, R.M. Summers, A.A. Taha, A. Tiulpin, S.A. Tsaftaris, B. Van Calster, G. Varoquaux, P.F. Jäger, Metrics reloaded: recommendations for image analysis validation, *Nature Methods* 21 (2) (2024) 195–212, <http://dx.doi.org/10.1038/s41592-023-02151-z>, URL <https://www.nature.com/articles/s41592-023-02151-z>, Publisher: Nature Publishing Group.

[62] S. Kapoor, A. Narayanan, Leakage and the reproducibility crisis in machine-learning-based science, *Patterns* 4 (9) (2023) <http://dx.doi.org/10.1016/j.patter.2023.100804>, URL [https://www.cell.com/patterns/abstract/S2666-3899\(23\)00159-9](https://www.cell.com/patterns/abstract/S2666-3899(23)00159-9), Publisher: Elsevier.

[63] M. Mahdi Behzadi, M. Madani, H. Wang, J. Bai, A. Bhardwaj, A. Tarakanova, H. Yamase, G. Hie Nam, S. Nabavi, Weakly-supervised deep learning model for prostate cancer diagnosis and gleason grading of histopathology images, 2022, <http://dx.doi.org/10.48550/arXiv.2212.12844>, URL <https://ui.adsabs.harvard.edu/abs/2022arXiv221212844M>, Publication Title: arXiv e-prints ADS Bibcode: 2022arXiv221212844M.

[64] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, Scikit-learn: Machine learning in Python, *Mach. Learn. Python*.

[65] L. Lannelongue, J. Grealey, M. Inouye, Green algorithms: Quantifying the carbon footprint of computation, *Adv. Sci.* 8 (12) (2021) 2100707, <http://dx.doi.org/10.1002/advs.202100707>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/advs.202100707> _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/advs.202100707>.