

# A Self-Supervised Contrastive Learning Approach for Whole Slide Image Representation in Digital Pathology

Parsa Ashrafi Fashi <sup>1</sup>

Sobhan Hemmati <sup>1,2</sup>

Morteza Babaie <sup>1,2</sup>

Ricardo Gonzalez <sup>1,3</sup>

H.R. Tizhoosh <sup>1,2,4</sup>

PASHRAFI@UWATERLOO.CA

SOBHAN.HEMATI@UWATERLOO.CA

MORTEZA.BABAIE@UWATERLOO.CA

GONZAR5@MCMASTER.CA

TIZHOOSH.HAMID@MAYO.EDU

<sup>1</sup> *Kimia Lab, University of Waterloo, Waterloo, ON, Canada*

<sup>2</sup> *Vector Institute, MaRS Centre, Toronto, ON, Canada*

<sup>3</sup> *McMaster University, Hamilton, ON, Canada*

<sup>4</sup> *Artificial Intelligence and Informatics, Mayo Clinic, Rochester, MN, USA*

## Abstract

Image analysis in digital pathology has proven to be one of the most challenging fields in medical imaging for AI-driven classification and search tasks. Due to their gigapixel dimensions, whole slide images (WSIs) are difficult to represent for computational pathology. Self-supervised learning (SSL) has recently demonstrated excellent performance in learning effective representations on pretext objectives, which improves the generalizations of downstream tasks. Previous self-supervised representation methods rely on patch (i.e., sub-image) selection and classification such that the effect of SSL on end-to-end WSI representation is not investigated. In this paper, we propose a novel self-supervised learning scheme based on the available primary site information, in contrast to existing augmentation-based SSL methods. We also exploit a fully supervised contrastive learning setup to increase the robustness of the representations for WSI classification and search for both pretext and downstream tasks. We trained and evaluated the model on more than 6,000 WSIs from The Cancer Genome Atlas (TCGA) repository provided by the National Cancer Institute. The proposed architecture achieved excellent results on most primary sites and cancer subtypes. We also achieved the best result on validation on a lung cancer classification task.

**Keywords:** Digital Pathology, Whole Slide Image, Representation Learning, Computational Pathology, Self-Supervised Learning, Image Search, Multiple Instance Learning, Attention, Supervised Contrastive Learning.

## 1. Introduction

The emergence of digital pathology has opened new horizons in medical image analysis for diagnostic purposes. Histopathology images, also known as whole slide images (WSIs), are generally accompanied with information about the site and type of diseases and malignancies. The recent advances in digital technology enable the fast digital scanning of tissue slides to generate high-quality WSIs. As a result, the volume of WSI archives in hospitals and clinics has been drastically increasing. Consequently, the necessity of timely analysis of WSIs has become apparent to address urgent needs in daily workflow of modern pathology. Hence, the digital scanning of slides, alongside the other benefits of pathology,

has made computerized techniques a favorite approach for image analysis and diagnosis. The field of digital pathology has been drastically changing due to the recent success of artificial neural networks in the field of AI. Deep learning can facilitate various pathology tasks such as segmentation, classification of regions and nuclei, and searching among WSIs to find similar morphology. However, the representation of digitized pathology slides has proven to be rather challenging due to the large data size of WSIs (generally larger than  $50,000 \times 50,000$  pixels). Besides, the morphological characteristics that discriminate different diagnoses may be microscopically small which causes a fundamental challenge for WSI representation. Creating a single vector representation directly from a WSI is subject to research with current convolutional neural networks (CNN). A common approach is to break a WSI into many small patches, feed each patch to a CNN, and aggregate the output features to develop a single WSI representation for search and classification. Nonetheless, developing patch-based feature extraction may not be efficient due to its multi-stage architecture. Also, in aggregation stage, information about patch importance and spatial patch knowledge is often ignored. In this paper, we propose an end-to-end architecture that has two main contributions. Firstly, an end-to-end self-supervised, attention-based multiple instance learning (SS-CAMIL) method, that exploits the primary site information of each WSI, which is always available during the tissue preparation and subsequent digitization. Furthermore, we show that employing a supervised contrastive learning approach can improve the quality of model embeddings both in WSI classification and search tasks.

## 2. Related Work

**Patch-Level WSI Representation** – Early representation approaches primarily investigated patch-level classification. In (Hou et al., 2016) the authors extracted and classified patch-level features with a CNN in an iterative fashion. Then, they implemented a multi-label SVM to create a single WSI vector representation. The authors in (Coudray et al., 2018) extracted multi-magnification features from 20x and 5x magnification and aggregated the features with an average of the probabilities of the corresponding patches. Kalra et al. first cluster the entire tissue with color clustering, then select patches based on the cluster (Kalra et al., 2020b). They employed patch-level embeddings for WSI search.

**Multiple Instance Learning** – Multiple instance learning (MIL) is a specific learning scheme where a label is assigned to bag of instances (Dietterich et al., 1997). Considering the bag of patches representation for each WSI, the MIL framework takes into account multiple instances of a slide, to represent WSIs. Recently, authors in (Zaheer et al., 2017) proposed deep MIL where they demonstrated different pooling layers following a specific form can obtain permutation invariant representations. Following this paper, many MIL based WSI representation schemes have been proposed. The authors of (Ilse et al., 2018) proposed attention-based multiple instance learning to perform weighted pooling over each instance feature. Another example of attention-based pooling in MIL is (Kalra et al., 2020a) where the authors introduced memory networks (MEM) for learning permutation invariant representations. In (Adnan et al., 2020), the authors used graph CNNs to consider each instance as a node in a graph and then learned an adjacency matrix to build a graph representation of WSIs. Just recently, (Hemati et al., 2021) have exploited deep sets (Zaheer

et al., 2017) for MIL training in histopathology. They employed a conditional prediction layer where predictions of primary site labels guide the primary diagnosis predictions.

**Self-Supervised Learning** Self-supervised learning (SSL) refers to a category of deep learning methods in which a model is trained on a set of well-defined pretext tasks before being applied to a primary (or downstream) task. Pretext tasks are trained on purposefully generated “*pseudo-labels*” from the data, in order to acquire visual representations with the intention of utilizing the acquired model weights for the main task. (Gidaris et al., 2018) is among the first works in vision-based SSL, where authors define *rotation* classification as a self-supervised task and show that various computer vision tasks such as classification, detection, or segmentation generalize better with self-supervision. There have also been some patch-level self-supervision in histopathology literature. In a recent publication, (Koohbanani et al., 2021) proposed pathology-specific tasks such as magnification classification, JigMag (predicting the magnification order in a shuffled vector of different magnification), and hematoxylin channel prediction.

**Contrastive learning** Contrastive learning (CL) is another active field of research where the goal is to pull similar instances together and push the non-related samples away. Training a model with a contrastive loss can help produce a more distinct feature vector for an input. The first usage of a contrastive loss appeared in (Chopra et al., 2005). The authors proposed a similarity loss function that maps training data into a target space such that the  $L_1$  norm of the target space imitates the semantic distance of the input space. They considered pairwise input and chose to either push away or pull the samples based on similarity. In (Hoffer and Ailon, 2015), instead of two samples for comparison, authors used one instance as an anchor, one negative and one positive sample for metric learning. (Khosla et al., 2020) recently suggested a fully supervised contrastive loss that draws all clusters of points belonging to the same class together while pushing clusters of samples from other classes apart. In most recent papers, CL is implemented in a self-supervised fashion. Authors of (Chen et al., 2020) propose SimCLR (a simple framework for contrastive learning of visual representations) that uses different augmentations as positive samples. As a pathology example, in (Ciga et al., 2021) authors employed SimCLR and achieved excellent results for multiple histopathology downstream tasks, including classification, regression, and segmentation compared to baseline training methods.

### 3. Proposed Methodology

In this paper, we propose a novel end-to-end WSI level self-supervised approach based on the primary site information as the pretext task. The primary site information corresponds to the organ type of each digital slide and is always available for each WSI. Many papers have used the primary site as a “soft lable” (Riasatian et al., 2021). We show that using the primary site information for the pretext task helps the model generalize better on the diagnosis classification. We have also utilized a supervised contrastive learning loss to create a more robust representation at the WSI level. The following section provides the step-by-step explanation of the proposed method. The complete methodology is depicted in Figure 1.

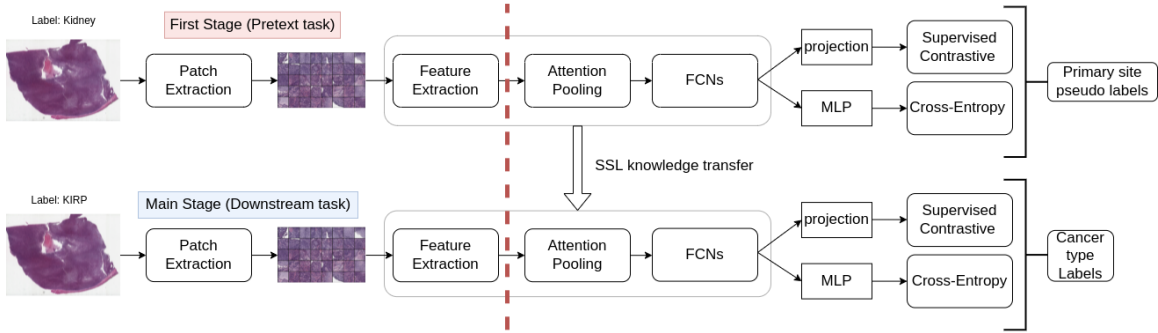


Figure 1: SS-CAMIL concept. The blocks that the transferred knowledge of pretext task (e.g., for label “kidney” as the primary site) is used for the downstream task (e.g., for label “KIRP”, *kidney renal papillary cell carcinoma*, as the primary diagnosis) are outlined with a grey line. For the LUAD/LUSC classification task, only the blocks on the right side of the dashed red line are used for we will be using pre-trained features.

**Patch Selection** For selection of the histopathology patches, we employed the patch selection method in (Kalra et al., 2020b). The authors utilized a two-step k-mean clustering. The tissue is grouped in the first step using the color histogram. The patch location is then subjected to a second k-means clustering to select spatially varied patches from each color segment. Each patch represents a different WSI location and colour. As a result, more regions of a WSI are likely to be considered during training.

**Feature Extraction** – We first modify the patch order in this phase to be fed into the feature extractor block. Suppose we have  $b$  batches, each WSI has  $n$  distinct patches, and each patch has a width  $w$  and height  $h$ . The reshape layer changes each input from the shape  $(b, n, w, h)$  to  $(b \times n, w, h)$ . The patches are now inputted to an EfficientNet B0 (Tan and Le, 2019) model for feature extraction. The features from the final convolutional block are then fed to a global max-pooling layer and a fully connected layer to extract vectors of size 1024 for each patch. Another reshape layer is then utilized to convert the output shape to  $(b, n, 1024)$ .

**Attention-based Pooling** – As displayed in Fig. 1, the feature vectors serve as the input to an attention block. Two fully connected layers plus an extension layer make up the attention block. The two dense layers produce a mask of size  $(b, n)$ , which is then duplicated to get a size of  $(b, n, 1024)$ . This is then multiplied with the attention block input and averaged to generate a 1024 vector representation of each WSI. Instead of a simple average pooling layer, the mask learns the weight of each patch (importance factor) and lets the model pick which patch is more representative of the WSI. The authors of (Ilse et al., 2018) showed that the representation of attention-based pooling is permutation-invariant, meaning that the output does not change when the input patches are reordered, hence establishing a large degree of freedom for patch selection.

### Self-supervision and contrastive learning based on primary site information –

The main contribution of this paper is to introduce the exploitation of **primary site information as pseudo-labels** in self-supervised learning setup (the first training stage). To our knowledge, previous SSL methods in pathology used data augmentation based self-supervision as pretext tasks. Primary site information of a WSI is an information that is always available and can be used as a pseudo-label. We have shown in our experiments that transferring the primary site information improves the performance of our model. To evaluate the impact of self-supervision, we conducted the experiments in two phases. First, the results of basic attention-based MIL without self-supervision (CAMIL) are reported. Then, we compare the result of primary site self-supervision on CAMIL (SS-CAMIL). Also, compared to all previous patch-based SSL methods, our self-supervision approach is performed in an **end-to-end fashion on WSI-level**. The second contribution is the utilization of supervised contrastive learning (Khosla et al., 2020) for both pretext and downstream tasks. After extracting WSI feature vectors, the features are passed to a projection head and a contrastive loss based on the primary site labels. However, using CL for a MIL setup has a bottleneck. One of the necessities of CL is large batch sizes (commonly more than 256), which may be infeasible due to the extensive set size of each WSI. To overcome this challenge, we added a cross-entropy loss to the contrastive loss function. This is because, CL cannot find enough positive samples within small batch sizes. Adding cross-entropy loss helps the positive instances to be close to a specific point in the embedding space. After the training with the above setting, the model is trained on the downstream task with diagnostic labels (i.e., primary diagnosis). After the training session, the features extracted from the last fully connected layer before the projection head are utilized for WSI search and classification.

## 4. Experiments and Results

**Dataset and Setup** We exploited 6,746 WSIs from The Cancer Genome Atlas Program (TCGA) and used 85, 5, and 10 percent (imposed by the ratios published in benchmarking literature) of the dataset for training, validation, and testing, respectively. The dataset consisted of WSIs of 24 primary sites with 30 distinct primary diagnoses. In the training stage, we set the batch size to 16 and the WSI set size to 40. We extracted patches of sizes  $1000 \times 1000$  and resized them to  $224 \times 224$  mainly due to memory limits (downsampling patches is quite common in literature (Tizhoosh and Pantanowitz, 2018; Marini et al., 2021)). For data augmentation, we applied horizontal and vertical flip, 90 degree rotation, shifting, and scaling to the data from the Albumentations library (Buslaev et al., 2020). We used an exponential decay learning rate scheduler with a base of 0.96 and a coefficient of 0.0001 ( $0.0001 \times 0.96^{\text{epoch number}}$ ). We trained each of the presented results with 150 epochs trained on three Tesla V 100 GPUs in parallel mode. We set the temperature to 0.1 for contrastive learning in both pretext and downstream tasks. For testing, we established horizontal (site identification) and vertical (subtype identification) WSI-search tasks. The precision with which we can locate a tumour type across the full test database is referred to as horizontal search. Vertical search, on the other hand, measures how well we can identify the proper cancer subtype of a tumour type from a set of slides from a single primary site, which may have a variety of initial diagnoses. For both search tasks, we employ  $k$ -

NN algorithm with  $k = 3$  to find the 3 instances closest to each test sample. We use the **leave-one-out technique** and provide the average scores due to the limited size of the test set.

In another experiment, we employed our model on a classification task of Lung Adenocarcinoma (LUAD) and Lung Squamous Cell Carcinoma (LUSC). The dataset had 2,574 lung tissues, distinct from the previous dataset. LUAD/LUSC classification is a challenging classification task that requires visual inspection of the tissue by expert pathologist (Coudray et al., 2018). We use 1,800 slides for training, and 774 slides for test (Kalra et al., 2020a). We freeze the convolutional feature extraction block to demonstrate the learned features from the previous setup. The batch size and the set size are the same as in the above setup.

**WSI search results** – Table 1 and Table 2 show the horizontal and vertical search results, respectively. We compare our performance with (Kalra et al., 2020b) and (Hemati et al., 2021). In both tables, CAMIL is the baseline attention-based MIL with CL and without self-supervision, and SS-CAMIL is the same as CAMIL setup but uses the weights of self-supervision of primary sites.

Table 1: Horizontal Search Results. F1-scores of Majority-3 (in %) are reported.

Tumor type	$n_{\text{slides}}$	Yottixel	CNN-DS	CAMIL	SS-CAMIL
Brain	46	73	91	100	100
Breast	77	45	77	91	91
Endocrine	71	61	66	86	89
Gastro.	69	50	75	84	86
Gynaec.	18	16	33	56	62
Head/neck	23	17	69	74	92
Liver	44	43	56	77	84
Melanocytic	18	16	50	61	78
Mesenchymal	12	8	100	92	92
Prostate/testis	44	47	81	91	89
Pulmonary	68	58	91	81	87
Urinary tract	112	67	76	92	95
Haematopoietic	42	0	24	50	50

For horizontal search, we can observe that the SS-CAMIL model has the best results among the four setups in 10 out of 13 cases. In one of remaining three cases (Prostate/Testis), CAMIL is the dominant model. In the rest of tumor types, SS-CAMIL has shown competitive results. One of the interesting observations is that although CNN-DS utilizes primary site information as prior information for the classification of tumor subtypes, the result of CAMIL is better in most cases. This observation demonstrates the effect of attention-based pooling compared to simple average pooling. Another observation is the improvement in performance with self-supervision on the primary sites. It can be observed that in most tumor types, SS-CAMIL has performed better than CAMIL. This observation indicates that the primary site, within the self-supervision framework, can help the model generalize better when deciding on the subtypes.

Table 2: Vertical Search Results. F1-scores of Majority-3 (in %) are reported.

Tumor Type	Subtype	$n_{\text{slides}}$	Yottixel	CNN-DS	CAMIL	SS-CAMIL
Gastrointestinal tract	COAD	22	62	69	72	73
	STAD	27	61	64	79	92
	ESCA	10	12	44	55	89
	READ	10	30	55	26	0
Pulmonary	LUAD	30	62	61	71	76
	LUSC	35	69	60	76	75
	MESO	3	0	50	50	33
Liver, pancreaticobiliary	LIHC	32	82	95	95	95
	PAAD	8	94	94	94	94
	CHOL	4	26	0	0	0
Endocrine	THCA	50	92	98	99	100
	PCPG	15	61	81	86	90
	ACC	6	25	28	50	77
Urinary tract	KIRP	25	75	84	84	88
	KIRC	47	91	87	92	92
	BLCA	31	89	95	94	98
	KICH	9	70	53	88	80
Brain	LGG	23	78	89	91	89
	GBM	23	82	89	91	90
Prostate/testis	PRAD	31	98	97	94	100
	TGCT	13	96	93	96	100
Gynaecological	OV	9	80	82	76	80
	CESC	6	92	66	44	44
	UCS	3	75	80	100	50

Table 3: LUAD/LUSC classification.

Method	Accuracy
MEM (Kalra et al., 2020a)	84%
CNN-DS (Hemati et al., 2021)	86%
CAMIL	<b>88%</b>

For the case of vertical search, in 16 subtypes of a total of 24 distinct subtypes, SS-CAMIL achieved the best F1-score. (the subtypes BLCA, THYM, HNSC, SARC, SKCM, and UVM are not included in the table since they are the only subtype in their tumor types). For five subtypes, CAMIL has performed better. Small sample sizes seem to be a recurrent pattern when our model does not perform well, meaning that the model did not have the chance to learn distinct features from these subtypes. Again, here we can see that in 15 subtypes, self-supervision has helped the model perform better than CAMIL. To show the effect of CL, we show the 2D t-SNE plot of (Hemati et al., 2021) and SS-CAMIL in Figure 2. We observe that SS-CAMIL clusters are tighter and more separable than CNN-DS.

**LUAD/LUSC Classification** The results of LUAD/LUSC classification are shown in Table 3. In this section, we use the features extracted from a DenseNet model (Huang et al., 2017), as per (Hemati et al., 2021). We can observe that our suggested strategy outperformed earlier approaches for LUAD/LUSC classification by 2% (delivering 88%), that underlines the performance of attention-pooling and contrastive learning. We have




Figure 2: t-SNE of CNN-DS (Hemati et al., 2021) (left) and SS-CAMIL (right).

Table 4: Attention pooling scores of 9 different WSIs.

Weighting	Lung			Kidney			Brain			Avg
	1	2	3	1	2	3	1	2	3	
Uniform	0.97	0.89	0.80	0.89	0.70	0.89	0.94	0.87	0.88	0.87
<b>SS-CAMIL</b>	<b>0.98</b>	<b>0.90</b>	<b>0.83</b>	<b>0.91</b>	<b>0.79</b>	<b>0.91</b>	<b>0.96</b>	<b>0.86</b>	<b>0.90</b>	<b>0.89</b>

also, employed the SS-CAMIL blocks in this task and it improved the performance to 89%, but since we are not sure whether it has seen the data in the search task (both datasets are from TCGA repository), we did not report those numbers in the table.

**Attention Pooling Effectiveness** – We have also investigated the effectiveness of the attention-pooling layer. We chose nine random WSIs from Lung, Kidney, and Brain organs. A pathology expert scored the effectiveness of the patches from each WSI with labels 1, 2 and 3, meaning “not useful”, “somewhat useful”, and “very useful”, respectively. We multiplied the normalized scores and the output of attention block for each WSI and compared the results with uniform importance (with all patches having the same weight). We then divide the scores by the optimal importance (weights of patches are proportional to effectiveness label) scores to get normalized numbers. The results are shown in Table 4. This suggests that our model has learned the relative importance of patches in the attention block.

## 5. Conclusions

In this paper, we proposed a self-supervised multiple instance learning model based on primary site information. We showed that our WSI-level representation model generalizes better on tumor subtypes comparing to two previous approaches. We demonstrated our performance on two tasks; WSI search and classification and showed that our model has a dominant performance on both tasks.



## References

- Mohammed Adnan, Shivam Kalra, and Hamid R Tizhoosh. Representation learning of histopathology images using graph neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 988–989, 2020.
- Alexander Buslaev, Vladimir I Iglovikov, Eugene Khvedchenya, Alex Parinov, Mikhail Druzhinin, and Alexandr A Kalinin. Alumentations: fast and flexible image augmentations. *Information*, 11(2):125, 2020.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, volume 1, pages 539–546. IEEE, 2005.
- Ozan Ciga, Tony Xu, and Anne Louise Martel. Self supervised contrastive learning for digital histopathology. *Machine Learning with Applications*, page 100198, 2021.
- Nicolas Coudray, Paolo Santiago Ocampo, Theodore Sakellaropoulos, Navneet Narula, Matija Snuderl, David Fenyö, Andre L Moreira, Narges Razavian, and Aristotelis Tsirigos. Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning. *Nature medicine*, 24(10):1559–1567, 2018.
- Thomas G Dietterich, Richard H Lathrop, and Tomás Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 89(1-2):31–71, 1997.
- Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations*, 2018.
- Sobhan Hemati, Shivam Kalra, Cameron Meaney, Morteza Babaie, Ali Ghodsi, and Hamid Tizhoosh. Cnn and deep sets for end-to-end whole slide image representation learning. In *Medical Imaging with Deep Learning*, 2021.
- Elad Hoffer and Nir Ailon. Deep metric learning using triplet network. In *International workshop on similarity-based pattern recognition*, pages 84–92. Springer, 2015.
- Le Hou, Dimitris Samaras, Tahsin M Kurc, Yi Gao, James E Davis, and Joel H Saltz. Patch-based convolutional neural network for whole slide tissue image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2424–2433, 2016.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.

- Maximilian Ilse, Jakub Tomczak, and Max Welling. Attention-based deep multiple instance learning. In *International conference on machine learning*, pages 2127–2136. PMLR, 2018.
- Shivam Kalra, Mohammed Adnan, Graham Taylor, and Hamid R Tizhoosh. Learning permutation invariant representations using memory networks. In *European Conference on Computer Vision*, pages 677–693. Springer, 2020a.
- Shivam Kalra, Hamid R Tizhoosh, Charles Choi, Sultaan Shah, Phedias Diamandis, Clinton JV Campbell, and Liron Pantanowitz. Yottixel—an image search engine for large archives of histopathology whole slide images. *Medical Image Analysis*, 65:101757, 2020b.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33, 2020.
- Navid Alemi Koohbanani, Balagopal Unnikrishnan, Syed Ali Khurram, Pavitra Krishnaswamy, and Nasir Rajpoot. Self-path: Self-supervision for classification of pathology images with limited annotations. *IEEE Transactions on Medical Imaging*, 2021.
- Niccolò Marini, Sebastian Otálora, Henning Müller, and Manfredo Atzori. Semi-supervised training of deep convolutional neural networks with heterogeneous data and few local annotations: An experiment on prostate histopathology image classification. *Medical image analysis*, 73:102165, 2021.
- Abtin Riasatian, Morteza Babaie, Danial Maleki, Shivam Kalra, Mojtaba Valipour, Sobhan Hemati, Mani Zaveri, Amir Safarpour, Sobhan Shafiei, Mehdi Afshari, et al. Fine-tuning and training of densenet for histopathology image representation using tcga diagnostic slides. *Medical Image Analysis*, 70:102032, 2021.
- Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114. PMLR, 2019.
- Hamid Reza Tizhoosh and Liron Pantanowitz. Artificial intelligence and digital pathology: challenges and opportunities. *Journal of pathology informatics*, 9, 2018.
- Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. *Advances in Neural Information Processing Systems*, 30, 2017.

## Appendix A. Cancer subtype Abbreviations

Table 5 explains the cancer subtype abbreviations exploited in the paper.

Table 5: Cancer subtype abbreviations.

Abbreviation	Primary Diagnosis
ACC	Adrenocortical Carcinoma
BLCA	Bladder Urothelial Carcinoma
CESC	Cervical Squamous Cell Carcinoma and Endocervical Adenoc.
CHOL	Cholangiocarcinoma
COAD	Colon Adenocarcinoma
ESCA	Esophageal Carcinoma
GBM	Glioblastoma Multiforme
KICH	Kidney Chromophobe
KIRC	Kidney Renal Clear Cell Carcinoma
KIRP	Kidney Renal Papillary Cell Carcinoma
LGG	Brain Lower Grade Glioma
LIHC	Liver Hepatocellular Carcinoma
LUAD	Lung Adenocarcinoma
LUSC	Lung Squamous Cell Carcinoma
MESO	Mesothelioma
OV	Ovarian Serous Cystadenocarcinoma
PAAD	Pancreatic Adenocarcinoma
PCPG	Pheochromocytoma and Paraganglioma
PRAD	Prostate Adenocarcinoma
READ	Rectum Adenocarcinoma
STAD	Stomach Adenocarcinoma
TGCT	Testicular Germ Cell Tumors
THCA	Thyroid Carcinoma
UCS	Uterine Carcinosarcoma