

TAKT: Target-Aware Knowledge Transfer for Whole Slide Image Classification

Conghao Xiong¹(✉)[0000–0003–0134–9397], Yi Lin²[0000–0002–7635–2518],
Hao Chen^{2,3}[0000–0002–8400–3780], Hao Zheng⁴(✉), Dong Wei⁴,
Yefeng Zheng⁴[0000–0003–2195–2847], Joseph J. Y. Sung⁵[0000–0003–3125–5199],
and Irwin King¹[0000–0001–8106–6447]

- ¹ Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong, China
² Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, Hong Kong, China
³ Department of Chemical and Biological Engineering, The Hong Kong University of Science and Technology, Hong Kong, China
⁴ Jarvis Research Center, Tencent YouTu Lab, Shenzhen, China
⁵ Lee Kong Chian School of Medicine, Nanyang Technological University, Singapore
chxiong21@cse.cuhk.edu.hk, howzheng@tencent.com

Abstract. Knowledge transfer from a source to a target domain is vital for whole slide image classification, given the limited dataset size due to high annotation costs. However, domain shift and task discrepancy between datasets can impede this process. To address these issues, we propose a Target-Aware Knowledge Transfer framework using a teacher-student paradigm, enabling a teacher model to learn common knowledge from both domains by actively incorporating unlabelled target images into the teacher model training. The teacher bag features are subsequently adapted to supervise the student model training on the target domain. Despite incorporating the target features during training, the teacher model tends to neglect them under inherent domain shift and task discrepancy. To alleviate this, we introduce a target-aware feature alignment module to establish a transferable latent relationship between the source and target features by solving an optimal transport problem. Experimental results show that models employing knowledge transfer outperform those trained from scratch, and our method achieves state-of-the-art performance among other knowledge transfer methods on various datasets, including TCGA-RCC, TCGA-NSCLC, and Camelyon16. Codes are released at <https://github.com/BearCleverProud/TAKT>.

Keywords: Knowledge Transfer · Computational Pathology · Whole Slide Image Classification.

1 Introduction

Whole Slide Image (WSI) refers to the digitised glass slides containing histology tissues, which is crucial for cancer diagnosis. Consequently, WSI classification

C. Xiong and Y. Lin—Equal contribution.

has become a research focus [14,17,20,26], including skin [13], lung [5], breast [19], prostate [4], and pancreas [15] cancers, while its success heavily depends on the availability of a large set of labeled samples. However, annotating WSIs is labour-intensive, and in certain cases, the patient cohort is limited, resulting in fewer WSIs in one dataset [2], limiting the data-hungry deep learning techniques.

An effective solution is to acquire knowledge from other datasets. Malignant cells share similar morphological characteristics, such as enlarged nuclei, irregular size and shape, prominent nucleoli, and intense or pale cytoplasm [3]. The shared characteristics can be leveraged to alleviate the over-fitting problem in relatively small datasets [21,29,30]. However, in practice, knowledge transfer could suffer from two challenges, namely, domain shift and task discrepancy.

Domain shift is a change in distribution, caused by differences in tumour sizes and colour tones across datasets. Models trained on one dataset are often biased towards it [6], limiting knowledge transfer. Task discrepancy is inconsistency in tasks between domains, and task-specific features may not directly transfer to another task. Previous works address domain shift from the domain adaptation perspective [1,7]. However, they focus on style differences and assume that the label space remains the same, while the label spaces are different in our case.

To tackle these issues, we propose a **Target-Aware Knowledge Transfer** (TAKT) framework to enable the teacher model to learn common knowledge across two domains in a teacher-student paradigm. Within this framework, we actively incorporate unlabelled data from the target domain during the training of the teacher model, using our proposed **Target-Aware Data Augmentation** (TADA) method. When augmenting a source feature, we retrieve the closest centroids generated by an unsupervised clustering algorithm from the target dataset and integrate them into the source features. These centroids are the most representative and relevant samples from the target domain, which enhance the diversity of the augmented dataset without compromising its original distribution. This approach has two rationales. Firstly, integrating target domain representations enhances the transferability of learned knowledge in a specified direction, *i.e.*, from the source domain to the target domain. Secondly, rich data augmentation inherently helps the model learn more generalised features.

However, it is observed that due to the domain shift and task discrepancy, the teacher model still tend to overlook the target features after TADA, and hence the teacher model remains biased towards the source domain. To address this issue, we introduce a **Target-Aware Feature Alignment** (TAFA) module, applied between features from source and target domains. This module establishes a transferable relationship between the source and target features by solving an **Optimal Transport** (OT) problem, providing a holistic view of the two distributions. This relationship enforces the teacher model to pay similar attention to the target features, enabling the model to learn more common knowledge across both domains. The contributions of this paper are summarised below:

1. We propose a TAKT framework to alleviate the impact of domain shift by actively incorporating unlabelled target data with the source domain.




Fig. 1. Illustrations of (a) our target-aware knowledge transfer framework, (b) our target-aware data augmentation method and (c) our target-aware feature alignment module and a sample optimal transport flow. We first train the teacher model and then the student model. The light areas in (c) indicate regions with higher values.

2. We propose a TAFE module to mitigate bias towards the source domain by establishing a latent relationship between source and target features.
3. We conduct extensive experiments on three datasets, including RCC, NSCLC and Camelyon16. Results demonstrate the efficacy of our proposed method.

2 Methodology

2.1 Overview of the Target-Aware Knowledge Transfer

The TAKT framework, shown in Fig. 1(a), has two parts: a TADA method and a TAFE module. Following the teacher-student paradigm, we first train a teacher model $\mathcal{M}_t(\cdot)$ on a dataset generated with TADA and then train a student model $\mathcal{M}_s(\cdot)$ on the target dataset, supervised by the adapted teacher bag features. Mathematically, the loss function for the student model $\mathcal{L}_s$ is given as,

$$\mathcal{L}_s = \mathcal{L}_{\mathcal{T}} + \alpha \mathcal{L}_{kt} = \mathcal{L}_{\mathcal{T}} + \alpha \sum \left(\text{MHA} \left(\text{sign}(h_t) \left| \frac{h_t}{T} \right|^{\frac{1}{n}} \right) - h_s \right)^2, \quad (1)$$

where $\mathcal{L}_{\mathcal{T}}$ is the loss function on the target domain $\mathcal{T}$ (in our case, cross entropy loss applied to the student logits p_s), α is a coefficient, $\text{MHA}(\cdot)$ is the **M**ulti-**H**ead **A**ttention (MHA) [25], used to adapt source-specific features from the

teacher model to target-specific ones that are more suitable for the student model to learn, $\mathbf{h}_t, \mathbf{h}_s$ are the teacher and student bag features, respectively, and $T = 0.1, n = 3$ are coefficients in **Power Temperature Scaling (PTS)** [10]. The PTS norm together with an MHA module is the $\mathcal{A}(\cdot)$ module in Fig. 1(a).

2.2 Target-Aware Data Augmentation

Overview of the TADA Method. The illustration of TADA is shown in Fig. 1(b). The augmentation is performed at the instance level between WSI bags, while for target WSI bags, clustered centroids instead of raw patch features are leveraged. During the augmentation, we find the closest centroid for each source feature. The augmentation operations are designed based on ReMix [28].

Step 1: Clustering of the Target Dataset. Each WSI bag from the target domain $\mathcal{T}$ is clustered using K -means, which discovers highly representative features within that bag. Consequently, the feature vectors are transformed into more representative centroids that can effectively represent a set of features. The source dataset remains unchanged due to the potential adverse effects of clustering on the dataset, such as a significant reduction in the number of instances, in which case, augmentation techniques would greatly change the original distribution of the source dataset, thereby negatively affecting the performance of the teacher model and its transferability to downstream tasks.

Step 2: Cross-Dataset Mix Operation. The motivation behind our TADA is to enrich the source features by imparting them with a specific “direction” to transfer, *i.e.*, towards the target domain. Trained with the target data, the teacher model grasps the general target feature distribution, enabling it to learn source features that exhibit better generalisability compared to the teachers that have not been exposed to the target data. Specifically, for each WSI bag in the source dataset, we enumerate each feature and augment the feature with a fixed probability p . For each feature $\mathbf{f}$ selected for augmentation, we first identify the closest centroid vector $\mathbf{c}$ from the stacked centroid matrix of the target domain. With the closest centroid vector $\mathbf{c}$, we employ the following operations:

1. **Append.** $\mathbf{c}$ is appended to the bag.
2. **Replace.** $\mathbf{f}$ is replaced with $\mathbf{c}$ in the bag.
3. **Interpolate.** The interpolation vector $\mathbf{f}_I = (1 - \lambda)\mathbf{f} + \lambda\mathbf{c}$ is appended to the bag, where $\lambda \in (0, 1)$ is the strength of the augmentation.
4. **Covary.** The vector $\mathbf{f}_C = \mathbf{f} + \lambda\boldsymbol{\delta}$ is appended to the bag. $\boldsymbol{\delta} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mathbf{c}})$, where $\Sigma_{\mathbf{c}}$ is the covariance matrix of $\mathbf{c}$ and $\mathcal{N}(\cdot, \cdot)$ is the normal distribution.
5. **Joint.** All the aforementioned methods are applied to the bag.

2.3 Target-Aware Feature Alignment Module

Most of the recent methods for WSI fall under the category of attention-based multiple instance learning [14, 20, 26]. These methods calculate attention scores

based on patch features using a gated attention network [14] to determine the importance of each patch in the final prediction. In our framework, we observed that attention scores for target features are often lower than those for the source domain, indicating the teacher model is still biased towards the source domain after TADA. This bias hampers the effectiveness of data augmentation and diminishes the transferability of teacher features. However, directly penalising the low attention scores is not reasonable since target features may not necessarily contain source-specific features. Therefore, we propose a target-aware feature alignment module to establish a transferable latent relationship between source and target features instead of pursuing complete consistency, and this also prevent the target features to be neglected during training.

As shown in Fig. 1, the inputs to TAFE module are bag features exclusively obtained from the source features ($\mathbf{h}_S$) and target centroids ($\mathbf{h}_T$). Within the TAFE module, how to establish the transferable latent relationship is treated as an OT problem. The OT distance is derived from the transportation plan with the lowest cost, providing a comprehensive view of the difference between two distributions in terms of shape, density, and spread, which are lacking in the traditional measures, such as Euclidean distance and cosine similarity. Mathematically, the input to the teacher model $\mathcal{M}_t(\cdot)$ is denoted as $\mathbf{F} = [\mathbf{F}_S, \mathbf{C}_T]$, where $\mathbf{F}_S, \mathbf{C}_T$ are the feature vectors and centroids from the source and target domains, respectively. The bag features from the teacher model only using $\mathbf{F}_S$ and $\mathbf{C}_T$ are denoted as $\mathbf{h}_S = \mathcal{M}_t(\mathbf{F}_S)$ and $\mathbf{h}_T = \mathcal{M}_t(\mathbf{C}_T)$, respectively. For each input pair $(\mathbf{h}_S, \mathbf{h}_T)$, we solve for the unbalanced OT problem,

$$\min_{\gamma} \langle \gamma, \mathbf{C} \rangle_F + r \cdot \Omega(\gamma) + r_1 \cdot \text{KL}(\gamma \mathbf{1}, \mathbf{a}) + r_2 \cdot \text{KL}(\gamma^T \mathbf{1}, \mathbf{b}), \quad (2)$$

where the (i, j) -th element in the cost matrix $\mathbf{C}$ is $|\mathbf{h}_{S_i} - \mathbf{h}_{T_j}|$, $\langle \cdot, \cdot \rangle_F$ is the Frobenius inner product, $\Omega(\gamma) = \text{KL}(\gamma, \mathbf{a}\mathbf{b}^T)$ is the entropic regularisation term, r, r_1, r_2 are the regularisation coefficients, $\mathbf{a}, \mathbf{b} = [1/c, \dots, 1/c]$ are two uniform distributions of size c , c is the dimension of $\mathbf{h}_S$ and $\mathbf{h}_T$, $\text{KL}(\cdot, \cdot)$ is the Kullback–Leibler divergence, and $\gamma \geq 0$ is the optimal transport flow. The optimal flow indicates how teacher features can be transformed to have the same distribution as the student feature at minimal cost, considering the amount of mass transferred. The calculated distance between these two features decreases as the mass moved decreases. The problem can be solved using the Sinkhorn-Knopp algorithm [8]. Having obtained the optimal flow γ , we calculate the distance $d = \sum \mathbf{C}\gamma$. The overall loss function of the teacher model $\mathcal{L}_t$ is given as,

$$\mathcal{L}_t = \mathcal{L}_S + \beta \mathcal{L}_{dist} = \mathcal{L}_S + \beta \sum \mathbf{C}(\mathbf{h}_S, \mathbf{h}_T) \cdot \gamma(\mathbf{h}_S, \mathbf{h}_T), \quad (3)$$

where $\mathcal{L}_S$ is the loss function on the source domain (in our case, cross entropy loss applied to the teacher logits $\mathbf{p}_t$), and $\beta > 0$ is the coefficient. Within this formulation, the latent relationship is represented as the optimal flow, and the optimisation of the regularisation term aims to enhance the transferability of features. The establishment of such a transferable relationship ensures the target centroids receive enough attention during the training of teacher model.

Table 1. Results on Camelyon16 with source domains being NSCLC or RCC. The best results are in red underlined, and the second best ones are in blue italic. The subscript in each cell is the standard deviation.

Method	NSCLC $\rightarrow$ Camelyon16						RCC $\rightarrow$ Camelyon16					
	AUC $\uparrow$		F1 $\uparrow$		Accuracy $\uparrow$		AUC $\uparrow$		F1 $\uparrow$		Accuracy $\uparrow$	
CLAM [20]	0.814	0.010	0.764	0.032	0.801	0.027	0.814	0.0010	0.764	0.032	0.801	0.027
Fine-tuning	<u>0.885</u>	0.012	<i>0.831</i>	0.027	<i>0.853</i>	0.021	<i>0.872</i>	0.044	<i>0.844</i>	0.035	<i>0.855</i>	0.024
ST [11]	0.803	0.016	0.747	0.017	0.793	0.012	0.820	0.012	0.742	0.012	0.788	0.009
NST [12]	0.824	0.080	0.758	0.063	0.796	0.043	0.814	0.023	0.742	0.007	0.788	0.005
AT [16]	0.821	0.011	0.760	0.030	0.804	0.022	0.847	0.026	0.763	0.045	0.804	0.024
PKT [22]	0.794	0.022	0.739	0.016	0.788	0.009	0.813	0.025	0.743	0.012	0.791	0.008
CC [23]	0.858	0.054	0.801	0.055	0.827	0.040	0.839	0.063	0.777	0.064	0.814	0.040
SP [24]	0.819	0.012	0.797	0.023	0.827	0.016	0.813	0.025	0.743	0.012	0.791	0.008
PTS [10]	0.835	0.037	0.780	0.031	0.811	0.020	0.795	0.020	0.754	0.006	0.793	0.005
TAKT	<u>0.952</u>	0.011	<u>0.896</u>	0.012	<u>0.904</u>	0.012	<u>0.926</u>	0.010	<u>0.854</u>	0.013	<u>0.866</u>	0.012
p-value	0.0006		0.0049		0.0049		0.0027		0.1598		0.1340	

3 Experiments and Results

3.1 Dataset Descriptions

Camelyon16. The Camelyon16 dataset [18] contains 399 WSIs of lymph nodes from women with breast cancer. The purpose of this dataset is metastasis detection. The training and test sets contain 270 and 129 WSIs, respectively. We further split the provided training set into training and validation datasets by 8:2 and compare the performances with other methods on the official test dataset.

TCGA-RCC (RCC). The RCC¹ dataset contains 940 WSIs, with 121 WSIs from 109 cases of Kidney Chromophobe Renal Cell Carcinoma (KICH), 519 WSIs from 513 cases of Kidney Renal Clear Cell Carcinoma (KIRC), and 300 WSIs from 276 cases of Kidney Renal Papillary Cell Carcinoma (KIRP). The dataset is split into training, validation, and test sets by the ratio of 6:1.5:2.5.

TCGA-NSCLC (NSCLC). The NSCLC¹ dataset includes 1,053 WSIs: 512 from 478 Lung Squamous Cell Carcinoma (LUSC) cases and 541 from 478 Lung Adenocarcinoma (LUAD) cases. The dataset split is identical to RCC.

3.2 Implementation Details

Evaluation Metrics. Area Under the Curve (AUC), F1 and accuracy are the evaluation metrics. These metrics can holistically reflect the overall performances of the models. The thresholds of F1 and accuracy scores are set to 0.5. We also conduct a significance test to further assess the significance of the difference between the means of the highest metrics and those of the second-highest metrics.

¹ <https://www.cancer.gov/tcga>

Table 2. Results on NSCLC and RCC transferring to each other.

Method	RCC $\rightarrow$ NSCLC						NSCLC $\rightarrow$ RCC					
	AUC $\uparrow$		F1 $\uparrow$		Accuracy $\uparrow$		AUC $\uparrow$		F1 $\uparrow$		Accuracy $\uparrow$	
CLAM [20]	0.936	0.015	0.859	0.007	0.859	0.007	0.970	0.003	0.859	0.010	0.886	0.007
Fine-tuning	0.944	0.002	0.876	0.008	0.876	0.008	0.979	0.002	0.874	0.010	0.899	0.009
ST [11]	0.942	0.004	0.865	0.004	0.866	0.004	0.977	0.006	0.860	0.006	0.889	0.007
NST [12]	0.943	0.007	0.874	0.007	0.875	0.007	0.980	0.003	0.871	0.006	0.900	0.003
AT [16]	0.949	0.001	0.877	0.004	0.877	0.004	0.977	0.002	0.863	0.003	0.893	0.000
PKT [22]	0.944	0.001	0.872	0.012	0.872	0.012	0.978	0.005	0.865	0.014	0.895	0.011
CC [23]	0.945	0.004	0.872	0.012	0.872	0.012	0.977	0.005	0.864	0.008	0.893	0.004
SP [24]	0.937	0.004	0.877	0.006	0.877	0.006	0.979	0.001	0.874	0.003	0.896	0.003
PTS [10]	0.942	0.003	0.864	0.009	0.864	0.009	0.970	0.002	0.865	0.004	0.885	0.004
TAKT	0.955	0.003	0.890	0.010	0.890	0.010	0.984	0.002	0.890	0.008	0.909	0.007
p-value	0.0048		0.0288		0.0288		0.0429		0.0072		0.0253	

Training Settings. Methods comparing with ours include no knowledge transfer, fine-tuning, Soft Target (ST) [11], Neuron Selectivity Transfer (NST) [12], Attention Transfer (AT) [16], Probabilistic Knowledge Transfer (PKT) [22], Correlation Congruence (CC) [23], Similarity Preserving (SP) [24] and PTS norm [10]. The base model is CLustering-constrained-Attention Multiple-instance learning (CLAM) [20]. The student model in our method is initialised with the teacher model. Both the student and teacher models are trained up to 200 and no less than 50 epochs. The training is ceased when the validation loss stops decreasing for 20 epochs. All experiments are repeated three times with different seeds. The means and standard deviations of the performances are reported.

Hyper-parameters. The learning rate, weight decay and dropout are set to 2×10^{-4} , 1×10^{-5} , and 0.25, respectively [20]. Adam optimiser is used [20]. The probability of augmenting each feature p and the strength of augmentation λ is set to 0.3 and 0.5, respectively [28]. The number of MHA heads is 8. WSIs are split into non-overlapping patches of 256×256 pixels at $20\times$ magnification [20]. ResNet-50 [9] pre-trained on ImageNet is used to extract features from them. r, r_1, r_2 are set to 0.1, 0.5, 0.5, respectively [27]. We perform sensitivity analysis on α, β in supplementary materials, and they are set to 0.1, 0.2, respectively.

3.3 Comparison Results

We compare our method with other related knowledge transfer methods in four settings: NSCLC to Camelyon16, RCC to Camelyon16, RCC to NSCLC and NSCLC to RCC. The experimental results are reported in Table 1 and Table 2. Our method achieves the best performance across every metric, especially on Camelyon16, where our method outperforms other methods by a large margin. In addition, the p-values comparing the best (ours) and second-best metrics indicate that our method significantly outperforms second-best performing

Table 3. Ablation studies on NSCLC to Camelyon16 transfer. Left: Ablation study results on our “replace” augmentation method (R), MSE, cosine similarity (cos), and OT distance with dimension reduction. Right: The augmentation methods are denoted as: Append (A), Replace (R), Interpolate (I), Covary (C), and Joint (J).

Method	NSCLC → Camelyon16			Method	NSCLC → Camelyon16		
	AUC↑	F1↑	Acc.↑		AUC↑	F1↑	Acc.↑
No Aug.	0.921 _{0.019}	0.855 _{0.046}	0.871 _{0.038}	+A,OT	0.939 _{0.018}	<i>0.885</i> _{0.031}	<i>0.897</i> _{0.025}
+R	0.932 _{0.013}	0.862 _{0.016}	0.873 _{0.016}	+R,OT	<i>0.952</i> _{0.011}	<i>0.896</i> _{0.012}	<i>0.904</i> _{0.012}
+R,MSE	0.942 _{0.007}	0.861 _{0.012}	0.873 _{0.012}	+I,OT	0.933 _{0.021}	0.854 _{0.038}	0.871 _{0.031}
+R,cos	<i>0.943</i> _{0.004}	<i>0.877</i> _{0.006}	<i>0.889</i> _{0.005}	+C,OT	0.931 _{0.012}	0.854 _{0.018}	0.868 _{0.013}
+R,OT	<i>0.952</i> _{0.011}	<i>0.896</i> _{0.012}	<i>0.904</i> _{0.012}	+J,OT	<i>0.941</i> _{0.004}	0.865 _{0.013}	0.876 _{0.013}

methods. Furthermore, due to the small tumour size and the limited number of samples, Camelyon16 is more difficult than TCGA datasets, which is reflected in the absolute value of the metrics. The results prove that our method can effectively transfer knowledge from a simpler dataset to a harder one. Another observation is that methods with additional supervision signals perform better on TCGA datasets, and fine-tuning performs better when transferring knowledge from TCGA to the Camelyon16 dataset. The potential reason for this is that TCGA datasets are more similar, as evidenced by the maximum mean discrepancy scores between these datasets. The tumour size is drastically different between Camelyon16 and TCGA datasets, leading to significant differences in attention distributions. Methods with additional supervision signals (features, attention scores and logits) all rely on the attention scores as bag features are eventually derived from them. Therefore, these methods may consistently introduce bias during student training, resulting in lower performance.

3.4 Ablation Studies

TAFa Module. We investigate the impact of “replace” augmentation with different alignment losses on the final performance, including Mean Squared Error (MSE), cosine similarity (cos) and OT distance. The experimental results are shown in the left part of Table 3. When we incorporate the “replace” augmentation during teacher model training, the AUC score increases by 1.10%. With MSE, cos and OT, the AUC score further increases by 1.00% and 1.10% and 2.00%, respectively, demonstrating the effectiveness of OT distance.

Augmentation Methods. We conduct experiments using the aforementioned augmentation methods. The experimental results are presented in the right part of Table 3. Incorporating augmentation methods results in improved AUC scores compared to the absence of augmentation, with a minimum increase of 1.00%, highlighting their effectiveness. Particularly, the “replace” augmentation method achieved the highest performance, surpassing other methods by a significant margin (with a best average AUC score of 95.2%).

4 Conclusion

In this work, we presented a TAKT framework specifically designed for WSI classification. The framework contains two main components: the TADA method and the TAFA module. The TADA method augments the source dataset by actively retrieving the closest centroids from the target domain, facilitating the teacher model to acquire common knowledge from both domains. To mitigate the bias of teacher model towards the source domain, we proposed a TAFA module to establish a latent relationship between source and target features by solving a OT problem, enforcing the teacher model to pay similar attention to features from both the source and target domains. Our experimental results demonstrated that models trained with knowledge transfer techniques outperformed those trained from scratch and our method achieved state-of-the-art performance among other adapted knowledge transfer methods.

Acknowledgments. This research was done with Jarvis Research Center, Tencent YouTu Lab and partially supported by the Research Grants Council of the Hong Kong Special Administrative Region, China (CUHK 14222922, RGC GRF 2151185).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Ahn, E., Kumar, A., Fulham, M., Feng, D., Kim, J.: Unsupervised domain adaptation to classify medical images using zero-bias convolutional auto-encoders and context-based feature augmentation. *IEEE transactions on medical imaging* **39**(7), 2385–2394 (2020) [2](#)
2. Aumpan, N., Vilaichone, R.k., Pornthisarn, B., Chonprasertsuk, S., Siramolpiwat, S., Bhanthumkomol, P., Nunanan, P., Issariyakulkarn, N., Ratana-Amornpin, S., Miftahussurur, M., et al.: Predictors for regression and progression of intestinal metaplasia (im): a large population-based study from low prevalence area of gastric cancer (im-predictor trial). *PloS one* **16**(8), e0255601 (2021) [2](#)
3. Baba, A.I., Cătoi, C.: Comparative oncology. Publishing House of the Romanian Academy (2007) [2](#)
4. Campanella, G., Hanna, M.G., Geneslaw, L., Mirafflor, A., Werneck Krauss Silva, V., Busam, K.J., Brogi, E., Reuter, V.E., Klimstra, D.S., Fuchs, T.J.: Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature Medicine* **25**(8), 1301–1309 (2019) [2](#)
5. Chen, C.L., Chen, C.C., Yu, W.H., Chen, S.H., Chang, Y.C., Hsu, T.I., Hsiao, M., Yeh, C.Y., Chen, C.Y.: An annotation-free whole-slide training approach to pathological classification of lung cancer types using deep learning. *Nature communications* **12**(1), 1193 (2021) [2](#)
6. Cui, Y., Liu, Z., Chen, Y., Lu, Y., Yu, X., Liu, X.S., Kuo, T.W., Rodrigues, M., Xue, C.J., Chan, A.: Retrieval-augmented multiple instance learning. *Advances in Neural Information Processing Systems* **36** (2024) [2](#)

7. Feng, Y., Xu, X., Wang, Y., Lei, X., Teo, S.K., Sim, J.Z.T., Ting, Y., Zhen, L., Zhou, J.T., Liu, Y., et al.: Deep supervised domain adaptation for pneumonia diagnosis from chest x-ray images. *IEEE Journal of Biomedical and Health Informatics* **26**(3), 1080–1090 (2021) [2](#)
8. Frogner, C., Zhang, C., Mobahi, H., Araya-Polo, M., Poggio, T.A.: Learning with a wasserstein loss. In: *Advances in Neural Information Processing Systems*. pp. 2053–2061 (2015) [5](#)
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 770–778 (2016) [7](#)
10. He, R., Sun, S., Yang, J., Bai, S., Qi, X.: Knowledge distillation as efficient pre-training: Faster convergence, higher data-efficiency, and better transferability. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9161–9171 (2022) [4](#), [6](#), [7](#)
11. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. *arXiv:1503.02531* (2015) [6](#), [7](#)
12. Huang, Z., Wang, N.: Like what you like: Knowledge distill via neuron selectivity transfer. *arXiv:1707.01219* (2017) [6](#), [7](#)
13. Ianni, J.D., Soans, R.E., Sankarapandian, S., Chamarthi, R.V., Ayyagari, D., Olsen, T.G., Bonham, M.J., Stavish, C.C., Motaparthi, K., Cockerell, C.J., et al.: Tailored for real-world: a whole slide image classification system validated on uncurated multi-site data emulating the prospective pathology workload. *Scientific Reports* **10**(1), 3217 (2020) [2](#)
14. Ilse, M., Tomczak, J., Welling, M.: Attention-based deep multiple instance learning. In: *International Conference on Machine Learning*. pp. 2127–2136 (2018) [2](#), [4](#), [5](#)
15. Keikhosravi, A., Li, B., Liu, Y., Conklin, M.W., Loeffler, A.G., Eliceiri, K.W.: Non-disruptive collagen characterization in clinical histopathology using cross-modality image synthesis. *Communications biology* **3**(1), 414 (2020) [2](#)
16. Komodakis, N., Zagoruyko, S.: Paying more attention to attention: improving the performance of convolutional neural networks via attention transfer. In: *5th International Conference on Learning Representations, ICLR 2017* (2017) [6](#), [7](#)
17. Lin, Y., Zhu, Z., Cheng, K.T., Chen, H.: Prompt-guided adaptive model transformation for whole slide image classification. *arXiv:2403.12537* (2024) [2](#)
18. Litjens, G., Bandi, P., Ehteshami Bejnordi, B., Geessink, O., Balkenhol, M., Bult, P., Halilovic, A., Hermesen, M., van de Loo, R., Vogels, R., et al.: 1399 h&e-stained sentinel lymph node sections of breast cancer patients: the camelyon dataset. *GigaScience* **7**(6), giy065 (2018) [6](#)
19. Litjens, G., Sánchez, C.I., Timofeeva, N., Hermesen, M., Nagtegaal, I., Kovacs, I., Hulsbergen-Van De Kaa, C., Bult, P., Van Ginneken, B., Van Der Laak, J.: Deep learning as a tool for increased accuracy and efficiency of histopathological diagnosis. *Scientific Reports* **6**(1), 26286 (2016) [2](#)
20. Lu, M.Y., Williamson, D.F., Chen, T.Y., Chen, R.J., Barbieri, M., Mahmood, F.: Data-efficient and weakly supervised computational pathology on whole-slide images. *Nature Biomedical Engineering* **5**(6), 555–570 (2021) [2](#), [4](#), [6](#), [7](#)
21. Pan, S.J., Yang, Q.: A survey on transfer learning. *IEEE Transactions on knowledge and data engineering* **22**(10), 1345–1359 (2009) [2](#)
22. Passalis, N., Tefas, A.: Learning deep representations with probabilistic knowledge transfer. In: *European Conference on Computer Vision*. pp. 268–284 (2018) [6](#), [7](#)
23. Peng, B., Jin, X., Liu, J., Li, D., Wu, Y., Liu, Y., Zhou, S., Zhang, Z.: Correlation congruence for knowledge distillation. In: *the IEEE/CVF International Conference on Computer Vision*. pp. 5007–5016 (2019) [6](#), [7](#)

24. Tung, F., Mori, G.: Similarity-preserving knowledge distillation. In: the IEEE/CVF international conference on computer vision. pp. 1365–1374 (2019) [6](#), [7](#)
25. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems. pp. 5998–6008 (2017) [3](#)
26. Xiong, C., Chen, H., Sung, J.J.Y., King, I.: Diagnose like a pathologist: Transformer-enabled hierarchical attention-guided multiple instance learning for whole slide image classification. In: International Joint Conference on Artificial Intelligence. pp. 1587–1595 (2023) [2](#), [4](#)
27. Xu, Y., Chen, H.: Multimodal optimal transport-based co-attention transformer with global structure consistency for survival prediction. In: IEEE/CVF International Conference on Computer Vision. pp. 21241–21251 (October 2023) [7](#)
28. Yang, J., Chen, H., Zhao, Y., Yang, F., Zhang, Y., He, L., Yao, J.: Remix: A general and efficient framework for multiple instance learning based whole slide image classification. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 35–45. Springer (2022) [4](#), [7](#)
29. Yu, X., Wang, J., Hong, Q.Q., Teku, R., Wang, S.H., Zhang, Y.D.: Transfer learning for medical images analyses: A survey. *Neurocomputing* **489**, 230–254 (2022) [2](#)
30. Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., He, Q.: A comprehensive survey on transfer learning. *Proceedings of the IEEE* **109**(1), 43–76 (2020) [2](#)