

# Transcriptomics-guided Slide Representation Learning in Computational Pathology

Guillaume Jaume<sup>1,2</sup>\*, Lukas Oldenburg<sup>1,3</sup>\*, Anurag Vaidya<sup>1,2</sup>, Richard J. Chen<sup>1,2</sup>,  
Drew F.K. Williamson<sup>1,2</sup>†, Thomas Peeters<sup>1</sup>, Andrew H. Song<sup>1,2</sup>, Faisal Mahmood<sup>1,2</sup>

<sup>1</sup>Mass General Brigham, <sup>2</sup>Harvard University and <sup>3</sup>RWTH Aachen University

gjaume@bwh.harvard.edu, lukas.oldenburg@rwth-aachen.de, faisalmahmood@bwh.harvard.edu

## Abstract

Self-supervised learning (SSL) has been successful in building patch embeddings of small histology images (e.g.,  $224 \times 224$  pixels), but scaling these models to learn slide embeddings from the entirety of giga-pixel whole-slide images (WSIs) remains challenging. Here, we leverage complementary information from gene expression profiles to guide slide representation learning using multi-modal pre-training. Expression profiles constitute highly detailed molecular descriptions of a tissue that we hypothesize offer a strong task-agnostic training signal for learning slide embeddings. Our slide and expression (S+E) pre-training strategy, called TANGLE, employs modality-specific encoders, the outputs of which are aligned via contrastive learning. TANGLE was pre-trained on samples from three different organs: liver ( $n=6,597$  S+E pairs), breast ( $n=1,020$ ), and lung ( $n=1,012$ ) from two different species (*Homo sapiens* and *Rattus norvegicus*). Across three independent test datasets consisting of 1,265 breast WSIs, 1,946 lung WSIs, and 4,584 liver WSIs, TANGLE shows significantly better few-shot performance compared to supervised and SSL baselines. When assessed using prototype-based classification and slide retrieval, TANGLE also shows a substantial performance improvement over all baselines. Code available at <https://github.com/mahmoodlab/TANGLE>.

## 1. Introduction

Self-supervised learning (SSL) [7, 97] has recently gained significant traction in Computational Pathology (CPath) [10, 13, 60, 72, 74, 84]. SSL is particularly suited for modeling giga-pixel whole-slide images (WSIs), whose size can exceed  $150,000 \times 150,000$  pixels, and which are consequently challenging to process with Vision Transformers (ViTs) or Convolutional Neural Networks (CNNs).

\*Equal contribution

†Presently at Emory University School of Medicine

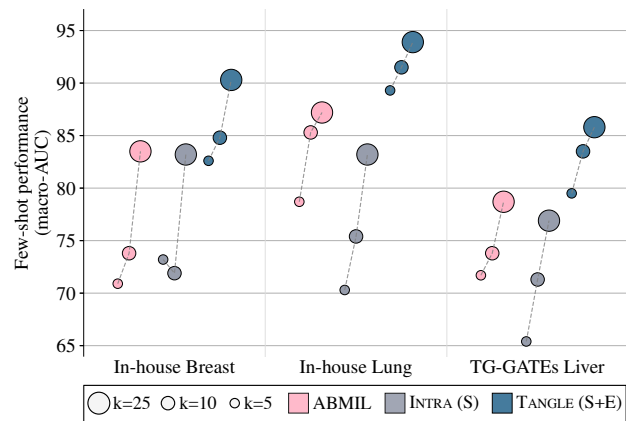


Figure 1. **Few-shot performance.** TANGLE linear probing performance compared to multiple instance learning (ABMIL) and intra-modality slide SSL (INTRA). TANGLE uses gene expression (E) to guide slide pre-training (S) using multimodal contrastive learning (S+E). Results on independent cohorts for BRCA subtyping (human breast,  $n=1,265$  WSIs), NSCLC subtyping (human lung,  $n=1,946$  WSIs), and TG-GATes lesion classification (rat liver,  $n=4,584$  WSIs).  $k$ : number of training samples per class.

Because of this size constraint, most CPath approaches adopt a divide-and-conquer strategy that consists of (1) tessellating the WSI into small patches and (2) extracting low-dimensional *patch embeddings* with a frozen pre-trained network. Until recently, the prevalent practice involved relying on networks pre-trained on ImageNet [16, 26, 55]. However, with the advent of SSL, this step is replaced by histopathology-specific visual encoders [4, 20, 80, 84] or vision-language encoders [29, 56], in most cases trained on human cancer samples. The resulting patch embeddings constituting the WSI can then be fed to weakly-supervised models for classification as done in Multiple Instance Learning [17, 32, 44, 55, 69].

SSL can also be pushed one step further to derive *slide embeddings* without requiring any human annotations [10,

42, 43, 77, 94]. The resulting slide embeddings can serve as input for various downstream tasks with minimal or no training, enabling slide classification with few-shot learning and prototyping, slide retrieval, and case stratification. In addition, as the embedding space is learned without necessitating pathologist annotations, the risk of using noisy labels inherent in inter-observer variability is greatly mitigated [24]. However, building slide embeddings with SSL remains challenging as (1) constructing slide “views” based on *patch-level augmentations* requires extracting multiple patch embeddings per patch, which is computationally expensive; (2) the visual primitives and invariances that need to be learned (such as being able to detect edges in natural images) become unclear when scaling to very large inputs; and (3) intra-slide heterogeneity can prevent deriving a consistent and strong training signal, especially when using masked image modeling.

Instead, inspired by multimodal vision-language models, we leverage gene expression data to *guide* slide representation learning into a slide-expression (S+E) pre-training model. Gene expression data, such as measured with RNA sequencing, are known to be strong indicators of disease state, with molecular signatures predictive of cancer subtype [59], patient survival [6], and drug toxicity [2], among others. Intuitively, the histology slide (S) and corresponding expression data (E) provide different *views* of the same underlying biological processes: gene expression forms a highly detailed molecular description of tissue, with as many descriptors as there are transcriptomic measurements, albeit lacking spatial information. Conversely, histology slides offer a finely detailed spatial representation of the tissue but with only two markers, namely, the hematoxylin and eosin combination represented as RGB channels. Consequently, molecular alterations, as detected through bulk transcriptomics, can be exhibited as discernible morphological patterns when examining the associated histology slides [14, 40, 41]. We hypothesize that guiding slide representation learning with expression constitutes a much stronger training signal than using slide augmentations or masking.

Here, we follow a multimodal contrastive learning paradigm where (S+E) pairs are aligned during a pre-training stage. Specifically, we address the modality heterogeneity gap by employing *modality-specific* encoders yielding a slide and expression embeddings that are aligned using a symmetric contrastive objective. Our models are based on large cohorts of publicly available (S+E) pairs, namely The Cancer Genome Atlas (TCGA) developed for studying human cancer and the Toxicogenomics Project-Genomics Assisted Toxicity Evaluation System (TG-GATES) developed for assessing drug toxicity in rat model animals. (S+E) models are trained on multiple species (*Homo sapiens* and *Rattus norvegicus*) and sites (liver, breast, and lung), that

we test on a panel of downstream tasks. To summarize, our contributions are: (1) the first self-supervised vision encoder for rat tissue trained on 15 million patches from 47,227 WSIs; (2) TANGLE, a transcriptomics-guided slide representation learning framework trained on thousands of (S+E) pairs using multimodal contrastive learning; (3) a series of few-shot classification, prototype-based classification, and slide retrieval experiments for lesion classification in rat liver and cancer subtyping in human breast and lung cancer that show the predictive capabilities of TANGLE; and (4) a post-hoc interpretability analysis that enables deriving insights about the aligned latent space.

## 2. Related work

### 2.1. Self-supervised visual representation learning

The combination of Vision Transformers (ViTs) [19, 79] and SSL [7, 97] has proved to be a powerful tool for building task-agnostic image representations. SSL can be categorized into (1) contrastive approaches [7, 66], whose underlying principle is to attract different representations of the same image (e.g., two distinct augmentations) while simultaneously pushing away representations of dissimilar images; (2) reconstruction approaches [27, 89], which aim to recover specific portions of an occluded image from the remaining parts of the same image; and (3) approaches combining both objectives [63, 97]. Representation learning in vision has also evolved towards multimodal vision-language models [1, 35, 47, 48, 50, 66, 73, 81, 93]. The same principles remain, where, for instance, the embedding of an image caption can be pulled close to the image (in a contrastive fashion), or partially masked with the objective to reconstruct the caption from the image. Vision-language models are also becoming prevalent in medical applications, by leveraging medical reports and textbooks [52, 86]. Our work aligns with this idea, where we align expression profiles with the slide representation.

### 2.2. Self-supervised learning in CPath

**Encoding histology patches:** Most works applying SSL to CPath focus on building embeddings from image patches (typically  $256 \times 256$ -pixel regions) [4, 12, 13, 20, 39, 42, 54, 80, 82, 84]. State-of-the-art methods are using a combination of contrastive- and reconstruction-based objectives trained with a student-teacher learning paradigm [20, 80]. Patch-level SSL is trained on increasingly large datasets and models (e.g., ViT-Huge trained on 1.5M slides in [80]). These can be based on public archives such as TCGA or CPTAC [4, 13, 20, 82, 84], on internal cohorts [80], or a mix of public and private datasets [38]. Recently, vision-language encoders designed for pathology have also been proposed [21, 29, 56], and rely on large-scale paired data scraped from social media, textbooks, or publications. All

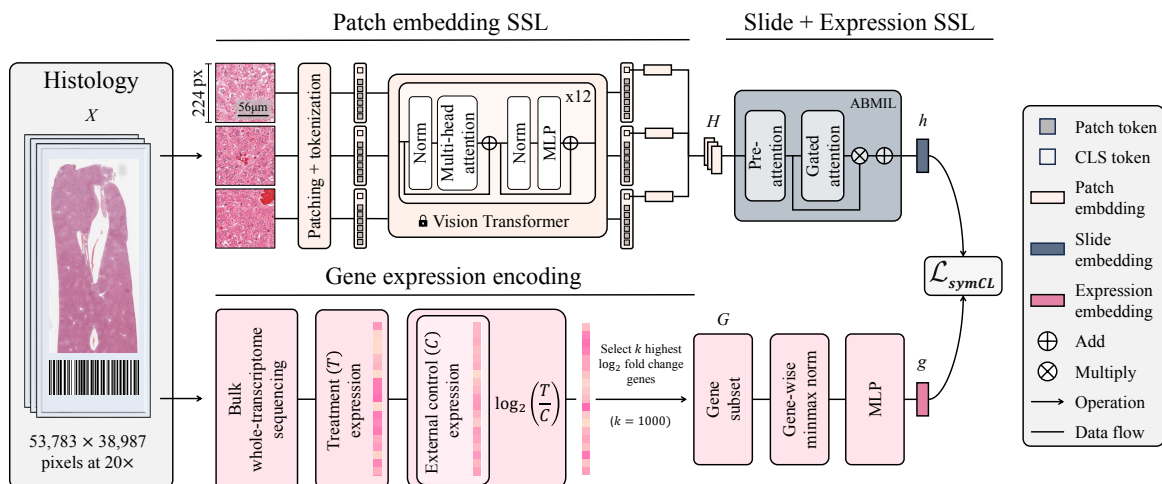


Figure 2. **Overview of TANGLE for (S+E) pre-training.** An input histology slide is tessellated into patches and encoded using a pre-trained vision encoder. The resulting patch embeddings are passed to an ABMIL module to derive a slide embedding. The corresponding gene expression data are encoded using an MLP. A symmetric contrastive objective  $\mathcal{L}_{symCL}$  learns to align embeddings from both modalities. During inference, a query slide is encoded into a slide embedding by the trained pooling module to be used for downstream tasks.

these models are solely based on human tissue, most of which are cancer samples. Here, we complement these by introducing the first vision encoder for rodent tissue microscopy, which plays a pivotal role in drug safety and biomarker discovery.

**Encoding histology slides:** Methods to build slide embeddings using SSL remain relatively scarce. Chen et al. [10] proposed a three-stage pre-training pipeline to hierarchically aggregate increasingly large tiles, from patches to patch embeddings to region embeddings to slide embeddings. Follow-up works improved pre-training using more complex training signals based on intra- and inter-slide similarity losses [43, 94], masked autoencoding [36] or patch prototyping [75].

### 2.3. Supervised learning in CPath

**Multiple Instance Learning:** MIL [17] is the current *de-facto* approach for WSI classification. In particular, Attention-based MIL (and its many extensions) has been used extensively in CPath [15, 32, 34, 46, 51, 55, 64, 65, 70, 76, 78, 83, 87, 91, 92, 95]. Context-aware extensions have also been proposed, such as based on graph neural networks [8, 44, 62] and Transformers [61, 69]. During (S+E) pre-training, we also employ MIL to pool pre-extracted patch embeddings into a slide embedding that we further use for SSL contrastive learning.

**Multimodal learning:** While the representation learning capabilities of (S+E) pre-training remain poorly understood, the multimodal integration of histology with gene expression data has been extensively studied in cancer-specific and pan-cancer works, especially for prognostication [3, 9, 11, 33, 49, 58, 68, 85, 88]. Several mechanisms

have been proposed such as late [11] or early fusion using multimodal token aggregation [33, 90, 96]. Although not directly connected to our approach, they motivate exploring the connection between gene expression profiles and tissue morphology. Notably, recent studies more closely aligned with (S+E) pre-training and demonstrated improved multimodal downstream performance through multimodal pre-training utilizing histology and expression data [18, 37, 96].

**Computational Toxicologic Pathology (CompToxPath):** The majority of work in CPath is centered around studying human cancer. CompToxPath is emerging as a new sub-field that aims to augment drug safety assessment using AI, especially at the pre-clinical stage [57]. CompToxPath has been used for organ identification [25], detecting abnormalities [5, 28, 30, 71], such as necrosis and hypertrophy detection. However, none of these works include SSL or large-scale evaluations. This work bridges this gap by applying (S+E) pre-training to large-scale toxicology datasets.

## 3. Method

Here, we present our framework, TANGLE, for TrANscriptomics-Guided sLIDE representation learning (see Figure 2). TANGLE is composed of (1) a vision encoder that encodes patches into *patch embeddings*, followed by a pooling module for learning a *slide embedding* (Section 3.1), (2) a gene expression encoder that combines transcriptomic measurements into an *expression embedding* (Section 3.2), and (3) a multimodal alignment module that learns to align both spaces (Section 3.3). TANGLE is tested on a variety of downstream tasks (Section 4).

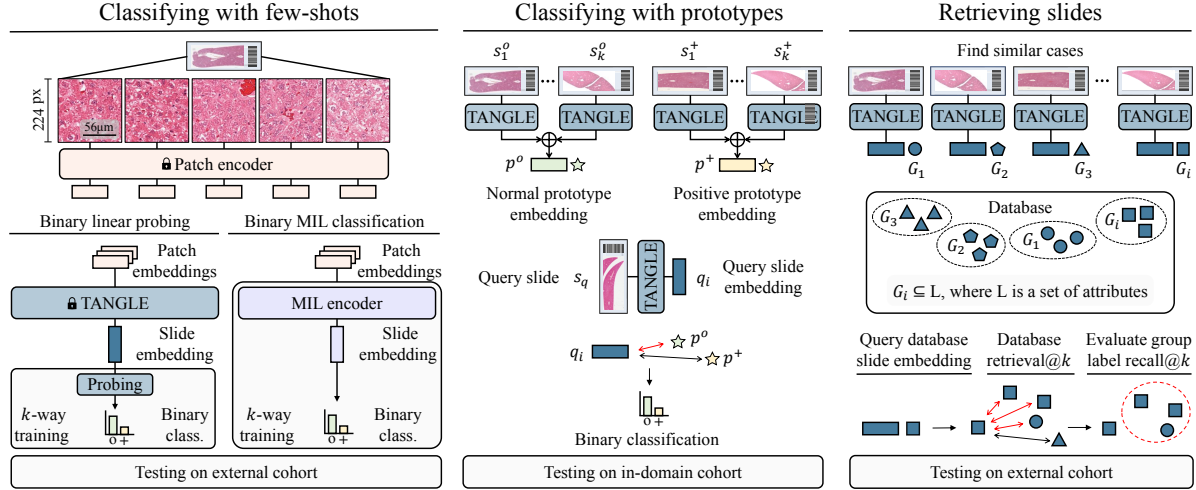


Figure 3. **Downstream tasks.** We test TANGLE and baselines on (1) few-shot, (2) prototype-based classification, and (3) slide retrieval.

### 3.1. Slide encoder

Given a histology slide  $\mathbf{X}_i \in \mathbb{R}^{d_x \times d_y \times 3}$ , we follow the MIL paradigm [17, 32, 44, 45, 55, 69], which consists of tessellating the slide into small patches, using a pre-trained vision encoder to extract patch embeddings, and pooling the resulting patch embeddings into a slide embedding.

**Pre-trained patch encoding:** For encoding rat tissue, we trained from scratch a ViT-Base (86 million parameters) with iBOT [97] on 15 million H&E patches extracted from 47,227 WSIs for 80 epochs, which we denote as iBOT-Tox. This is, to date, the largest SSL model for non-human histology tissue (see **Supplemental**). For encoding human tissue, we use CTransPath [82, 84], a state-of-the-art publicly available vision encoder. CTransPath was trained on 15 million patches from over 32,000 WSIs using a tiny Swin Transformer [53]. We denote the resulting patch embeddings of the  $i$ -th slide  $\mathbf{X}_i$  as  $\mathbf{H}_i \in \mathbb{R}^{N_{\mathcal{H}} \times d_{\mathcal{H}}}$ , where  $N_{\mathcal{H}}$  is the number of patch embeddings and  $d_{\mathcal{H}}$  their dimension.

**MIL slide encoding:** We learn a function  $f(\mathbf{H}_i) : \mathbb{R}^{N_{\mathcal{H}} \times d_{\mathcal{H}}} \rightarrow \mathbb{R}^d$  that maps the set of patch embeddings  $\mathbf{H}_i \in \mathbb{R}^{N_{\mathcal{H}} \times d_{\mathcal{H}}}$  into a slide embedding  $\mathbf{h}_i \in \mathbb{R}^d$ . Here, we use the popular attention-based MIL model (ABMIL) [32], which consists of learning patch-level attention weights used for pooling embeddings into a slide embedding.

### 3.2. Gene expression encoder

Given a set of raw transcriptomic measurements quantified across  $N_{\mathcal{G}}$  genes, we compute the log2 fold change relative to a control group, which represents gene expression deviations from a normal sample and, therefore, encode the magnitude of gene overexpression or underexpression (see **Supplemental**). The log2 fold change normalized transcriptomics associated with  $\mathbf{X}_i$ , denoted as  $\mathbf{t}_i \in \mathbb{R}^{N_{\mathcal{G}}}$ , can be seen as tabular data, which can efficiently be en-

coded with a multilayer perceptron (MLP) and named as  $\phi(\cdot)$ . Specifically, we train a 3-layer MLP to learn a mapping  $\phi(\mathbf{t}_i) : \mathbb{R}^{N_{\mathcal{G}}} \rightarrow \mathbb{R}^d$  to project a set of selected gene expressions  $\mathbf{t}_i \in \mathbb{R}^{N_{\mathcal{G}}}$  to an expression embedding  $\mathbf{g}_i \in \mathbb{R}^d$ .

### 3.3. Multimodal alignment

**Pre-training contrastive alignment:** We align the embedding space of the slide and expression encoders using a symmetric cross-modal contrastive learning objective. This is a widely employed representation learning formulation [66], especially in visual-language pre-training [56]. Formally, we define a batch as a set of  $M$  (S+E) pairs  $(\mathbf{h}_i, \mathbf{g}_i)_{i=1}^M$ , where  $\mathbf{h}_i$  and  $\mathbf{g}_i$  are the  $i$ -th slide embedding and gene expression profiles, respectively. For a given pair  $(\mathbf{h}_i, \mathbf{g}_i)$ , the objective is given by:

$$\mathcal{L}_{SymCL} = -\frac{1}{2M} \sum_{i=1}^M \log \frac{\exp(\tau \mathbf{h}_i^T \mathbf{g}_i)}{\sum_{j=1}^M \exp(\tau \mathbf{h}_i^T \mathbf{g}_j)} - \frac{1}{2M} \sum_{j=1}^M \log \frac{\exp(\tau \mathbf{g}_j^T \mathbf{h}_j)}{\sum_{i=1}^M \exp(\tau \mathbf{g}_j^T \mathbf{h}_i)} \quad (1)$$

where the first term represents the slide-to-expression contrastive loss, and the second term represents the expression-to-slide contrastive loss. Each term maximizes the dot-product similarity between embeddings from the same pair normalized (with Softmax) by negative pairs, which can be interpreted as other ‘‘classes’’.

**Complementary objective:** As an alternative to the contrastive loss, we introduce (1) an *expression reconstruction* objective  $\mathcal{L}_{REC}$  framed as an expression regression task, and (2) a vision-only intra-modality objective  $\mathcal{L}_{INTRA}$  that aims to align different random subsets of the slide (local-local alignment) and random subsets with the average patch em-



bedding (local–global alignment). We express these as,

$$\mathcal{L}_{\text{REC}} = \frac{1}{M} \sum_{i=1}^M \|\mathbf{g}_i - c(f(\mathbf{H}_i))\|_2 \quad (2)$$

$$\begin{aligned} \mathcal{L}_{\text{INTRA}} = & -\frac{1}{2M} \sum_{i=1}^M \log(\text{Softmax}(\mathbf{h}_{i,1}^T \bar{\mathbf{h}}_i, \tau)) \\ & -\frac{1}{2M} \sum_{i=1}^M \log(\text{Softmax}(\mathbf{h}_{i,2}^T \bar{\mathbf{h}}_i, \tau)) \end{aligned} \quad (3)$$

where  $c(\cdot)$  is an MLP regressor,  $\bar{\mathbf{h}}_i$  is the average patch embedding  $\bar{\mathbf{h}}_i = \frac{1}{N_{\mathcal{H}}^{(i)}} \sum_j \mathbf{H}_{ij}^{N_{\mathcal{H}}^{(i)}}$ , and where  $\mathbf{h}_{i,1}$  and  $\mathbf{h}_{i,2}$  are slide embedding views derived from different random patch embedding subsets (e.g., 2048 patches). These variants are referred to as TANGLE-REC and INTRA, respectively.

**Inference:** During inference, the query slide is passed through the vision encoder to extract patch embeddings and then to the MIL module to derive the slide embedding that encodes the morphological manifestations of the corresponding molecular signatures. We use the resulting slide embeddings for few-shot classification using linear probing and prototyping, and slide retrieval (see Figure 3).

## 4. Experiments and results

### 4.1. Dataset

**TG-GATES:** We collected all slides from the TG-GATES portal [31], which comprises 23,136 liver and 28,747 kidney slides ( $\approx 25\text{TB}$  of raw data). All slides are liver and kidney sections from Sprague–Dawley (SD) rats acquired in pre-clinical drug safety studies on 157 compounds. Each slide represents the morphological changes (lesions) observed after the administration of a particular drug dosage at a specified time point of sacrifice, denoted as a *sample group*. We manually curated the liver annotations into six classes (multi-label classification). We used a subset of 29 studies ( $n=4,584$  WSIs, liver only) as an independent test cohort. Other studies (both liver and kidney slides) are used for iBOT-Tox pre-training, (S+E) pre-training, and few-shot training. We additionally collected the corresponding gene expression profiles (microarrays) of 6,597 liver slides and selected the top 1,000 genes with the largest log2 fold change (see **Supplemental**).

**TCGA:** We collected 1,041 primary cases from the TCGA Breast Invasive Carcinoma (TCGA-BRCA) cohort, which comprises 831 Invasive Ductal Carcinoma (IDC) and 210 Invasive Lobular Carcinoma (ILC). We additionally collected 1,031 primary cases from the TCGA Non-Small Cell Lung Cancer (TCGA-NSCLC) cohort, among which

528 cases of Lung Adenocarcinoma (TCGA-LUAD) and 505 cases of Lung Squamous Cell Carcinoma (TCGA-LUSC). For each case, we downloaded the corresponding gene expression data (RNA sequencing) from the Xena database [23] that we curated using the method in [33], resulting in 4,999 gene expression per case.

**In-house:** We also collected a BRCA ( $n=1,265$  slides, 982 IDC and 283 ILC) and NSCLC ( $n=1,946$  slides,  $n=1,621$  LUAD and  $n=325$  LUSC) cohort from our in-house archives. These two cohorts are used as independent test sets for which gene expression data are not required. Slides from all datasets were processed at  $20\times$  magnification ( $0.5\mu\text{m}/\text{px}$ ).

### 4.2. Linear probing few-shot classification

We evaluate (S+E) pre-training in a few-shot classification scenario for lesion detection in liver (Table 1), and breast and lung cancer subtyping (Table 2). Following standard practice in SSL [7, 97], we employ linear probing for benchmarking TANGLE, TANGLE-REC, and INTRA. In addition, we benchmark HIPT [10] and baselines based on the average patch embeddings using different backbones (denoted as ResNet50+Avg., CTransPath+Avg. and iBOT-Tox+Avg.). Finally, we include two supervised MIL baselines, ABMIL [32] and TransMIL [69] (see Figure 3, left). Baselines are trained five times (Table 1) and ten times (Table 2), using  $k$  randomly sampled examples per class.

**TANGLE vs. supervised MIL:** TANGLE significantly outperforms all MIL baselines in the three datasets with an absolute gain of +5.9% in liver, +11.0% in breast, and +6.2% in lung compared to ABMIL for  $k=10$ . ABMIL leads to consistently better performance than TransMIL, which we hypothesize is due to (1) the use of a simpler architecture beneficial in low-data regimes and (2) tasks where the cellular morphology is more informative than the global tissue structure.

**TANGLE vs. averaging vs. MIL:** Despite the simplicity of these baselines, averaging provides performance that is on par with MIL in breast subtyping and liver lesion detection. We also observe that employing domain-specific vision encoders leads to substantial improvements, with CTransPath+Avg. outperforming ResNet50+Avg., which our iBOT-Tox+Avg. model in liver lesion detection significantly outperforms in TG-GATES (+9.6% and +11.0% compared to CTransPath+Avg. and ResNet50+Avg. for  $k=10$ ).

**TANGLE vs. INTRA vs. HIPT:** INTRA and HIPT provide similar performance in breast and lung, but are both significantly outperformed by TANGLE (+12.9% and +16.1% for  $k=10$  in breast and lung compared to INTRA). Both HIPT and INTRA are only marginally better or similar to the average patch embedding, which highlights the complexity of slide-level SSL.

\*50 or maximal available labeled samples per class

Table 1. **Few-shot lesion classification in rat liver.** Comparison of lesion classification (multi-label classification) using MIL vs. TANGLE and variations with linear probing, and evaluated using Macro-AUC (as %). All models are tested on an independent test cohort comprising 4,584 slides, without any data leakage from unimodal and multimodal pre-training. Standard deviation reported over five runs.

Model/Data	$k=1(\uparrow)$	$k=5(\uparrow)$	$k=10(\uparrow)$	$k=25(\uparrow)$	$k=50(\uparrow)^*$	
MIL	ResNet50+TransMIL [69]	53.3 ± 3.1	48.2 ± 2.9	53.2 ± 2.3	52.5 ± 3.7	52.9 ± 4.2
	CTransPath+TransMIL [69]	50.1 ± 4.1	51.1 ± 0.8	55.4 ± 3.9	58.1 ± 3.8	65.9 ± 4.2
	iBOT-Tox+TransMIL [69]	55.6 ± 6.1	66.5 ± 6.4	66.3 ± 6.2	68.6 ± 9.8	70.4 ± 10.6
	ResNet50+ABMIL [32]	56.0 ± 4.5	59.1 ± 7.1	64.1 ± 5.9	74.2 ± 8.6	80.3 ± 5.8
	CTransPath+ABMIL [32]	59.5 ± 4.4	71.7 ± 8.0	73.8 ± 9.5	78.7 ± 9.4	81.0 ± 7.3
	iBOT-Tox+ABMIL [32]	61.7 ± 5.3	73.2 ± 6.8	78.8 ± 9.3	81.6 ± 6.9	83.8 ± 8.1
Linear probing	ResNet50+Avg. [26]	55.0 ± 3.3	57.7 ± 11.8	60.5 ± 9.6	68.6 ± 8.0	72.7 ± 7.8
	CTransPath+Avg. [84]	56.9 ± 4.4	56.5 ± 10.5	61.9 ± 8.3	70.5 ± 8.1	73.9 ± 6.1
	iBOT-Tox+Avg. (ours)	53.9 ± 5.3	63.5 ± 6.9	71.5 ± 6.1	79.7 ± 5.0	81.9 ± 6.2
	iBOT-Tox+Intra (ours)	56.3 ± 7.3	62.6 ± 10.3	72.7 ± 7.4	80.2 ± 8.4	83.3 ± 8.0
	TANGLE-Rec (ours)	<b>73.8</b> ± 13.5	75.5 ± 14.3	78.3 ± 12.2	81.8 ± 10.8	82.7 ± 8.8
	TANGLE (ours)	72.1 ± 11.6	<b>80.1</b> ± 11.3	<b>84.7</b> ± 9.0	<b>86.3</b> ± 7.9	<b>86.9</b> ± 7.6

Table 2. **Few-shot cancer subtype classification in human breast and lung.** All models are tested on an independent test cohort comprising 1,265 breast slides and 1,946 lung slides and evaluated using Macro-AUC. Standard deviation reported over ten runs.

Model/Data	Breast				Lung				
	$k=1(\uparrow)$	$k=5(\uparrow)$	$k=10(\uparrow)$	$k=25(\uparrow)$	$k=1(\uparrow)$	$k=5(\uparrow)$	$k=10(\uparrow)$	$k=25(\uparrow)$	
MIL	ResNet50+TransMIL [69]	49.4 (± 13.0)	50.5 (± 7.6)	53.7 (± 8.8)	51.8 (± 4.9)	55.9 (± 5.4)	55.0 (± 5.6)	54.2 (± 6.1)	52.8 (± 5.4)
	CTransPath+TransMIL [69]	55.5 (± 9.5)	63.0 (± 9.1)	63.9 (± 7.8)	71.2 (± 12.7)	54.1 (± 8.6)	64.8 (± 8.9)	68.4 (± 10.4)	80.5 (± 10.8)
	ResNet50+ABMIL [32]	53.9 (± 14.4)	58.0 (± 9.9)	67.6 (± 9.6)	71.0 (± 3.7)	58.2 (± 7.4)	65.9 (± 6.1)	65.6 (± 4.6)	64.8 (± 1.4)
	CTransPath+ABMIL [32]	57.4 (± 14.0)	70.9 (± 10.5)	73.8 (± 7.1)	83.5 (± 8.6)	62.8 (± 9.0)	78.7 (± 11.7)	85.3 (± 4.5)	87.2 (± 3.4)
Linear probing	ResNet50+Avg. [26]	65.7 (± 17.3)	67.4 (± 13.1)	68.0 (± 13.9)	76.6 (± 8.0)	57.4 (± 6.5)	60.1 (± 4.7)	60.7 (± 4.2)	59.5 (± 2.1)
	CTransPath+Avg. [84]	<b>68.6</b> (± 16.9)	71.3 (± 11.1)	71.3 (± 14.4)	80.0 (± 7.5)	58.2 (± 6.6)	66.0 (± 6.6)	71.0 (± 2.6)	75.2 (± 3.3)
	HIPT <sub>CLS-4K</sub> [10]	62.2 (± 10.3)	63.7 (± 11.6)	71.0 (± 11.1)	78.1 (± 6.2)	59.8 (± 6.5)	70.5 (± 6.6)	74.1 (± 3.4)	79.1 (± 4.1)
	CTransPath+Intra (ours)	57.2 (± 14.7)	73.2 (± 5.5)	71.9 (± 9.1)	83.2 (± 6.8)	59.6 (± 7.0)	70.3 (± 9.8)	75.4 (± 6.7)	83.2 (± 4.4)
	TANGLE-Rec (ours)	56.3 (± 19.6)	73.6 (± 6.8)	68.3 (± 10.1)	83.4 (± 6.6)	<b>81.6</b> (± 10.3)	84.1 (± 4.9)	85.5 (± 1.8)	86.6 (± 2.3)
	TANGLE (ours)	67.3 (± 19.1)	<b>82.6</b> (± 8.0)	<b>84.8</b> (± 5.0)	<b>90.3</b> (± 3.7)	70.9 (± 6.0)	<b>89.3</b> (± 4.1)	<b>91.5</b> (± 2.1)	<b>93.9</b> (± 1.3)

**TANGLE vs. TANGLE-REC:** TANGLE-REC shows surprisingly high performance for  $k=1$ , but is outperformed for larger values of  $k$ . We hypothesize that TANGLE-REC renders simplified embeddings (*i.e.*, low-rank, see next Section), which makes one-shot learning easier but cannot express complex morphological subtleties.

**Loss ablation:** In TG-GATEs relative to TANGLE, adding a TANGLE-REC objective gives +0.05% AUC, adding INTRA on top gives -0.8% AUC, and -2.0% AUC when solely complementing TANGLE with INTRA. We hypothesize that

staining differences between train and test cause the INTRA objective to overfit, leading to worse performance. Replacing the cross-modal contrastive loss with an L1 objective gives -6.7% AUC and -7.0% AUC with an L2 (some designs conceptually similar to [18, 37, 96], see **Supplemental**).

### 4.3. Prototyping few-shot classification

We also assess the capacity of TANGLE to construct slide-level prototypes capable of predicting specific morphological characteristics. Specifically, we define a positive slide

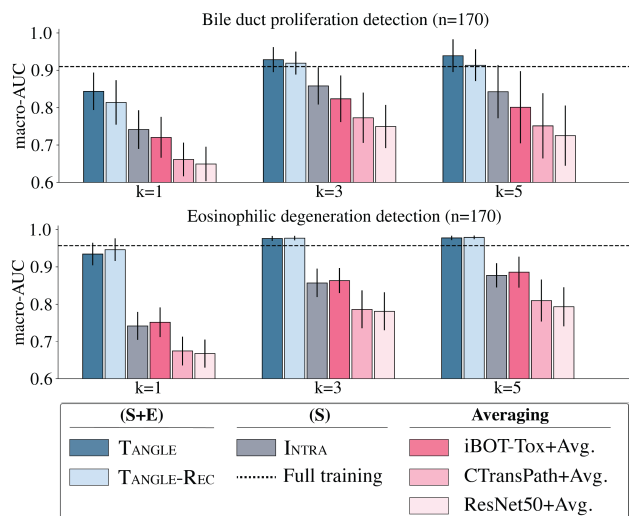


Figure 4. **Prototype-based classification.** Comparison of TANGLE and baselines for identifying study-level morphologies evaluated using macro-AUC. Prototypes are defined as the average of  $k$  slides selected from the study. Full training is an ABMIL trained on TG-GATEs train set ( $n=18,552$ ). Standard deviation reported over 100 bootstrapping iterations.

prototype  $p^+$  as the average of  $k$  ( $k=1,3,5$ ) slide embeddings with a morphology of interest. Similarly, a normal prototype  $p^0$  is defined using  $k$  normal slides, where the morphology under consideration is absent. Subsequently, we gauge the similarity between a query slide  $q_i$  and the two prototypes using the L2-distance – the distances interpreted as confidence prediction used for classification, *i.e.*,  $\|q_i - p^+\|$  and  $\|q_i - p^0\|$ , (see Figure 3, center). We apply this method to detect two types of lesions within the TG-GATEs test set, namely (1) eosinophilic degeneration in thioacetamide ( $n=170$ ), and (2) bile duct proliferation in methylene dianiline ( $n=170$ ). This setup mirrors a realistic application of AI, where the identification of a drug-induced morphology on  $k$  slides enables detecting if this morphology is present in slides from the same study, thereby enabling synergies between doctors and AI systems.

As shown in Figure 4, TANGLE and TANGLE-REC outperform all baselines in both studies. Compared to an ABMIL model trained on 100% of TG-GATEs ( $n=18,552$ ), TANGLE with  $k \geq 3$  leads to better performance. This highlights that (1) TG-GATEs includes study-specific morphologies that can be challenging to model, and (2) prototyping can help address this gap with minimal effort.

#### 4.4. Slide retrieval

We further evaluate TANGLE on slide retrieval using TG-GATEs test set. Each slide is associated with four others that share the same sample group. We extracted a subset of 594 slides with known drug-induced lesions. Our task is

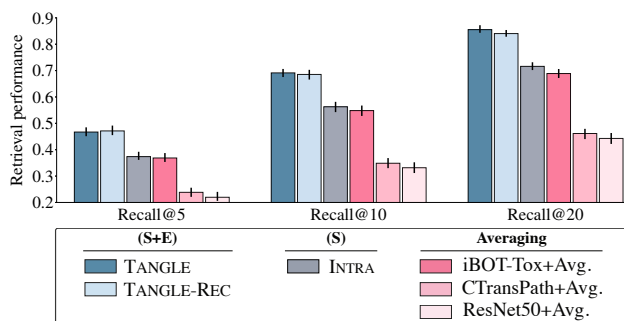


Figure 5. **Slide retrieval.** Comparison of TANGLE and baselines for retrieving slides with drug-induced lesions from the same *sample group* in TG-GATEs test. Recall@ $k$  quantifies the count of retrieved instances within the top- $k$  most similar slides normalized by the number of instances to retrieve (four per *sample group*). Standard deviation reported over 100 bootstrapping iterations.

to retrieve all slides that share the same sample group characteristics as the query, thereby demonstrating the capability of TANGLE to capture compound-, dose- and sacrifice-specific features. Specifically, we compute the Recall@ $k$  ( $k=5, 10, 20$ ), which measures the proportion of relevant slides that appear among the  $k$  most similar slides, with four being the total number of slides to retrieve in this context. The slide similarity is quantified using the cosine distance metric applied to the unnormalized slide embeddings (see Figure 3, right).

As presented in Figure 5, TANGLE reaches the best retrieval performance with on average 2.88/4 slides correctly retrieved among the top- $k=10$  instances and 3.44/4 among the top- $k=20$  instances. This result highlights that TANGLE can capture subtle morphological differences, such as those induced by administering different doses or sacrifice times.

Overall, results from Figure 4 and 5 ascertain the conclusion from the few-shot evaluation in that (1) (S+E) pre-training can capture task-agnostic features that can be used for downstream tasks, (2) intra-modality pre-training can outperform averaging, but their training signal remains weak, and (3) in-domain patch feature extractors greatly improve downstream performance. Additional experiments ablating TANGLE and INTRA losses, and showing the impact of hyper-parameters (batch size, temperature, number of sampled patches) are presented in the **Supplemental**.

#### 4.5. Interpretability

To better understand TANGLE properties, we analyzed the *rank* of the space spanned by the test slide embeddings (computed using the entropy of the  $d$  largest singular values of the embedding matrix, see **Supplemental**). Indeed, *rank* has been shown to be a predictor of downstream performance – and constitute a necessary (but not sufficient) condition for discriminative latent spaces [22]. We observe

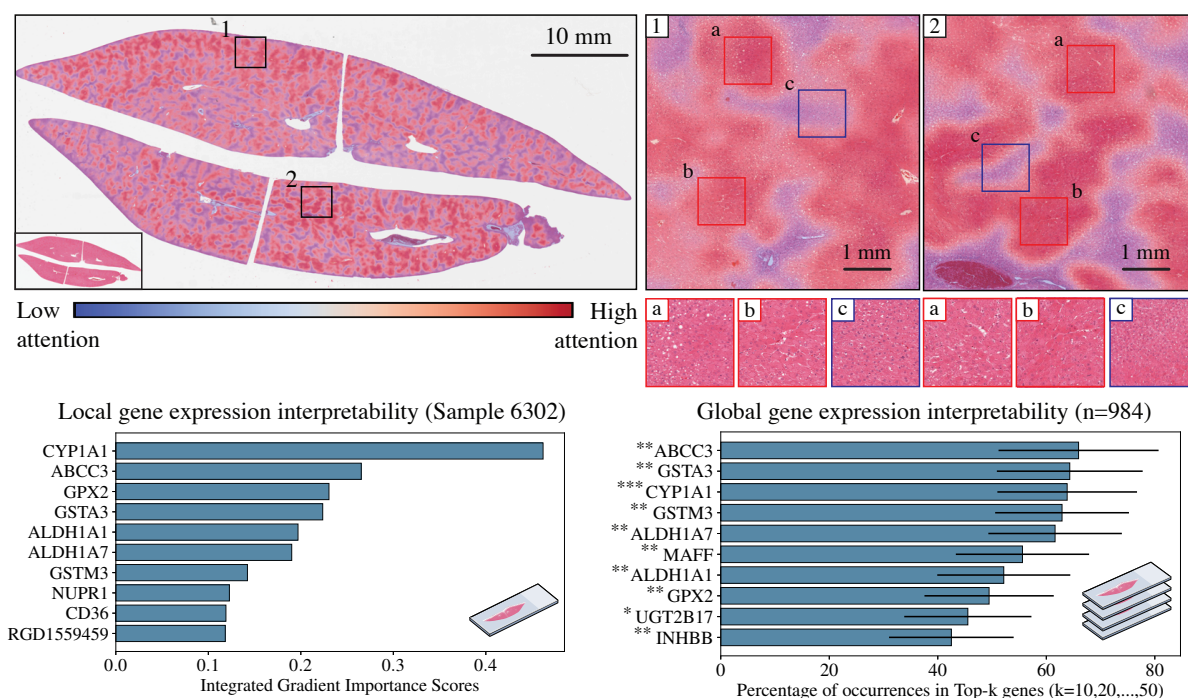


Figure 6. **Interpretability of TANGLE.** **Top:** Visualization of the attention weights of TANGLE in a TG-GATEs liver slide. High-attention regions highlight lesions (hepatocellular hypertrophy and fatty change). **Left:** Integrated Gradient (IG) scores of the gene expression encoder. High-importance genes map to well-known markers of liver toxicity, such as CYP1A1. **Right:** Percentage of occurrence of the top- $k$  genes in test. Many genes consistently appear as influential ( $>40\%$  of top- $k$  genes). \* denotes the number of publications referencing this gene as connected to drug-induced liver injury according to the CTD database (\*:  $>500$ , \*\*:  $>1,000$ , \*\*\*:  $>2,000$ ).

a strong positive correlation between rank and few-shot performance in all datasets among methods of the same family, (S+E), (S), and Averaging, as exemplified with  $k=10$  (see **Supplemental**). This confirms the importance of building domain-specific feature encoders for increased expressivity. This also suggests that reconstruction-based methods suffer from some degree of dimensionality collapse, which we hypothesize stems from over-fitting (and might disappear with larger cohorts). Finally, INTRA models have high ranks despite performing significantly worse than (S+E), which might be explained by the latent space expressing clinically irrelevant factors, such as staining variations.

Furthermore, we investigated whether salient histologic and expression features align with previously established biological findings. First, we visualized the attention weights learned during TANGLE pre-training (Figure 6, top). Important regions with high attention (visible in red) correlate with lesions (fatty change and hepatocellular hypertrophy, see **Supplemental** for heatmaps of lung and breast cancer samples). Next, we applied Integrated Gradients (IG) to derive gene-level importance scores (Figure 6, left) on TG-GATEs test samples with reported lesions. From there, we identified genes that consistently appear in the top- $k$  most influential genes, such as ABCC3

and CYP1A1 (Figure 6, right). We then quantitatively assessed their relevance by querying the Comparative Toxicogenomics Database (CTD) that aggregates all the literature on toxicology. 9/10 of the most important genes have more than 1,000 references connecting them to drug-induced liver injury, highlighting their relevance for slide representation learning.

## 5. Conclusion

In this paper, we introduced Slide+Expression (S+E) pre-training for slide representation learning. Our approach, TANGLE, was trained and tested on several species (*Homo sapiens* and *Rattus norvegicus*) and tissue sites (breast, lung, and liver). Overall, TANGLE outperforms all baselines significantly on several downstream tasks, including few-shot classification, prototype-based classification, and slide retrieval. These results highlight the potential of (S+E) pre-training and pave the way for additional developments [67]. Future work includes exploring multi-modal SSL objectives that extend beyond or synergize with, contrastive approaches, such as reconstruction of multi-modal masks. Concurrently, evaluating (S+E) pre-training on more tasks, such as predicting hormone receptor status from H&E slides, are promising research directions.



## References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. [2](#)
- [2] Benjamin Alexander-Dann, Lavinia Lorena Pruteanu, Erin Oerton, Nitin Sharma, Ioana Berindan-Neagoe, Dezső Módos, and Andreas Bender. Developments in toxicogenomics: understanding and predicting compound-induced toxicity from gene expression data. *Mol. Omics*, 14:218–236, 2018. [2](#)
- [3] Jordan Ash, Gregory Darnell, Daniel Munro, and Barbara Engelhardt. Joint analysis of expression levels and histological images identifies genes associated with tissue morphology. *Nature Communications*, 12, 2021. [3](#)
- [4] Shekoofeh Azizi, Laura Culp, Jan Freyberg, Basil Mustafa, Sebastien Baur, Simon Kornblith, Ting Chen, Nenad Tomasev, Jovana Mitrović, Patricia Strachan, et al. Robust and data-efficient generalization of self-supervised machine learning for diagnostic imaging. *Nature Biomedical Engineering*, pages 1–24, 2023. [1](#), [2](#)
- [5] Eun Bok Baek, Ji-Hee Hwang, Heejin Park, Byoung-Seok Lee, Hwa-Young Son, Yong-Bum Kim, Sang-Yeop Jun, Jun Her, Jaeku Lee, and Jae-Woo Cho. Artificial Intelligence-Assisted image analysis of Acetaminophen-Induced acute hepatic injury in Sprague-Dawley rats. *Diagnostics (Basel)*, 12(6), 2022. [3](#)
- [6] David Beer, Sharon Kardia, Chiang-Ching Huang, Thomas Giordano, Albert Levin, David Misek, Lin Lin, Guoan Chen, Tarek Gharib, Dafydd Thomas, Michelle Lizyness, Rork Kuick, Satoru Hayasaka, Jeremy Taylor, Mark Iannettoni, Mark Orringer, and Sam Hanash. Gene-expression profiles predict survival of patients with lung adenocarcinoma. *Nature medicine*, 8:816–24, 2002. [2](#)
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9630–9640, 2021. [1](#), [2](#), [5](#)
- [8] Tsai Hor Chan, Fernando Cendra, Lan Ma, Guosheng Yin, and Lequan Yu. Histopathology whole slide image analysis with heterogeneous graph representation learning. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [3](#)
- [9] Richard J Chen, Ming Y Lu, Jingwen Wang, Drew FK Williamson, Scott J Rodig, Neal I Lindeman, and Faisal Mahmood. Pathomic fusion: an integrated framework for fusing histopathology and genomic features for cancer diagnosis and prognosis. *IEEE Transactions on Medical Imaging*, 41(4):757–770, 2020. [3](#)
- [10] Richard J Chen, Chengkuan Chen, Yicong Li, Tiffany Y Chen, Andrew D Trister, Rahul G Krishnan, and Faisal Mahmood. Scaling vision transformers to gigapixel images via hierarchical self-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. [1](#), [3](#), [5](#), [6](#)
- [11] Richard J. Chen, Ming Y. Lu, Drew F.K. Williamson, Tiffany Y. Chen, Jana Lipkova, Zahra Noor, Muhammad Shaban, Maha Shady, Mane Williams, Bumjin Joo, and Faisal Mahmood. Pan-cancer integrative histology-genomic analysis via multimodal deep learning. *Cancer Cell*, 40(8): 865–878, 2022. [3](#)
- [12] Richard J. Chen, Tong Ding, Ming Y. Lu, Drew F. K. Williamson, Guillaume Jaume, Andrew H. Song, Bowen Chen, Andrew Zhang, Daniel Shao, Muhammad Shaban, Mane Williams, Lukas Oldenburg, Luca L. Weishaupt, Judy J. Wang, Anurag Vaidya, Long Phi Le, Georg Gerber, Sharifa Sahai, Walt Williams, and Faisal Mahmood. Towards a general-purpose foundation model for computational pathology. *Nature Medicine*, 2024. [2](#)
- [13] Ozan Ciga, Tony Xu, and Anne Louise Martel. Self-supervised contrastive learning for digital histopathology. *Machine Learning with Applications*, 7, 2022. [1](#), [2](#)
- [14] Nicolas Coudray, Paolo Santiago Ocampo, Theodore Sakellaropoulos, Navneet Narula, Matija Snuderl, David Fenyő, 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. [2](#)
- [15] Yufei CUI, Ziquan Liu, Xiangyu Liu, Xue Liu, Cong Wang, Tei-Wei Kuo, Chun Jason Xue, and Antoni B. Chan. Bayes-MIL: A new probabilistic perspective on attention-based multiple instance learning for whole slide images. In *The Eleventh International Conference on Learning Representations*, 2023. [3](#)
- [16] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. [1](#)
- [17] 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. [1](#), [3](#), [4](#)
- [18] Kexin Ding, Mu Zhou, Dimitris Metaxas, and Shaoting Zhang. Pathology-and-genomics multimodal transformer for survival outcome prediction. In *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pages 622–631, 2023. [3](#), [6](#)
- [19] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. [2](#)
- [20] Alexandre Filiot, Ridouane Ghermi, Antoine Olivier, Paul Jacob, Lucas Fidon, Alice Kain, Charlie Saillard, and Jean-Baptiste Schiratti. Scaling self-supervised learning for histopathology with masked image modeling. *medRxiv*, 2023. [1](#), [2](#)
- [21] Jevgenij Gamper and Nasir Rajpoot. Multiple instance captioning: Learning representations from histopathology text-

- books and articles. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16549–16559, 2021. [2](#)
- [22] Quentin Garrido, Randall Balestriero, Laurent Najman, and Yann Lecun. Rankme: Assessing the downstream performance of pretrained self-supervised representations by their rank. In *International conference on machine learning*, 2022. [7](#)
- [23] Mary J Goldman, Brian Craft, Mim Hastie, Kristupas Repečka, Fran McDade, Akhil Kamath, Ayan Banerjee, Yunhai Luo, Dave Rogers, Angela N Brooks, Jingchun Zhu, and David Haussler. Visualizing and interpreting cancer genomics data via the xena platform. *Nature biotechnology*, 38(6):675–678, 2020. [5](#)
- [24] Douglas Gomes, Simone Porto, Débora Balabram, and Helenice Gobbi. Inter-observer variability between general pathologists and a specialist in breast pathology in the diagnosis of lobular neoplasia, columnar cell lesions, atypical ductal hyperplasia and ductal carcinoma in situ of the breast. *Diagnostic pathology*, 9:121, 2014. [2](#)
- [25] Citlalli Gámez Serna, Fernando Romero-Palomo, Filippo Arcadu, Jürgen Funk, Vanessa Schumacher, and Andrew Janowczyk. Mmo-net (multi-magnification organ network): A use case for organ identification using multiple magnifications in preclinical pathology studies. *Journal of Pathology Informatics*, 13:100126, 2022. [3](#)
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [1](#), [6](#)
- [27] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. [2](#)
- [28] Holger Hoeffling, Tobias Sing, Imtiaz Hossain, Julie Boisclair, Arno Doelemeyer, Thierry Flandre, Alessandro Piaia, Vincent Romanet, Gianluca Santarossa, Chandrashegar Saravanan, Esther Sutter, Oliver Turner, Kuno Wuersch, and Pierre Moulin. Histonet: A deep learning-based model of normal histology. *Toxicologic Pathology*, 49(4):784–797, 2021. PMID: 33653171. [3](#)
- [29] Zhi Huang, Federico Bianchi, Mert Yuksekgonul, Thomas Montine, and James Zou. A visual–language foundation model for pathology image analysis using medical twitter. *Nature Medicine*, 29:1–10, 2023. [1](#), [2](#)
- [30] Ji-Hee Hwang, Minyoung Lim, Gyeongjin Han, Heejin Park, Yong-Bum Kim, Jinseok Park, Sang-Yeop Jun, Jaeku Lee, and Jae-Woo Cho. A comparative study on the implementation of deep learning algorithms for detection of hepatic necrosis in toxicity studies. *Toxicological Research*, 39(3):399–408, 2023. [3](#)
- [31] Yoshinobu Igarashi, Noriyuki Nakatsu, Tomoya Yamashita, Atsushi Ono, Yasuo Ohno, Tetsuro Urushidani, and Hiroshi Yamada. Open TG-GATes: a large-scale toxicogenomics database. *Nucleic Acids Research*, 43(D1):D921–D927, 2014. [5](#)
- [32] Maximilian Ilse, Jakub Tomczak, and Max Welling. Attention-based deep multiple instance learning. In *International conference on machine learning*, pages 2127–2136. PMLR, 2018. [1](#), [3](#), [4](#), [5](#), [6](#)
- [33] Guillaume Jaume, Anurag Vaidya, Richard Chen, Drew Williamson, Paul Liang, and Faisal Mahmood. Modeling dense multimodal interactions between biological pathways and histology for survival prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. [3](#), [5](#)
- [34] Syed Ashar Javed, Dinkar Juyal, Harshith Padigela, Amaro Taylor-Weiner, Limin Yu, and Aaditya Prakash. Additive MIL: Intrinsically interpretable multiple instance learning for pathology. In *Advances in Neural Information Processing Systems*, 2022. [3](#)
- [35] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021. [2](#)
- [36] Shuai Jiang, Liesbeth Hondelink, Arief A Suriawinata, and Saeed Hassanpour. Masked pre-training of transformers for histology image analysis. *arXiv preprint arXiv:2304.07434*, 2023. [3](#)
- [37] Ting Jin, Xingran Xie, Renjie Wan, Qingli Li, and Yan Wang. Gene-induced multimodal pre-training for image-omic classification. In *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2023. [3](#), [6](#)
- [38] Mingu Kang, Heon Song, Seonwook Park, Donggeun Yoo, and Sérgio Pereira. Benchmarking self-supervised learning on diverse pathology datasets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [2](#)
- [39] Mingu Kang, Heon Song, Seonwook Park, Donggeun Yoo, and Sérgio Pereira. Benchmarking self-supervised learning on diverse pathology datasets. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3344–3354, 2023. [2](#)
- [40] Jakob Nikolas Kather, Alexander T Pearson, Niels Halama, Dirk Jäger, Jeremias Krause, Sven H Loosen, Alexander Marx, Peter Boor, Frank Tacke, Ulf Peter Neumann, et al. Deep learning can predict microsatellite instability directly from histology in gastrointestinal cancer. *Nature Medicine*, 25(7):1054–1056, 2019. [2](#)
- [41] Jakob Nikolas Kather, Lara R Heij, Heike I Grabsch, Chiara Loeffler, Amelie Echle, Hannah Sophie Muti, Jeremias Krause, Jan M Niehues, Kai AJ Sommer, Peter Bankhead, et al. Pan-cancer image-based detection of clinically actionable genetic alterations. *Nature cancer*, 1(8):789–799, 2020. [2](#)
- [42] 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. [2](#)

- [43] Tristan Lazard, Marvin Lrousseau, Etienne Decencière, and Thomas Walter. Giga-ssl: Self-supervised learning for gigapixel images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4304–4313, 2023. 2, 3
- [44] Y. Lee, J.H. Park, S. Oh, et al. Derivation of prognostic contextual histopathological features from whole-slide images of tumours via graph deep learning. *Nat. Biomed. Eng.*, 2022. 1, 3, 4
- [45] Bin Li, Yin Li, and Kevin W Eliceiri. Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14318–14328, 2021. 4
- [46] Honglin Li, Chenglu Zhu, Yunlong Zhang, Yuxuan Sun, Zhongyi Shui, Wenwei Kuang, Sunyi Zheng, and Lin Yang. Task-specific fine-tuning via variational information bottleneck for weakly-supervised pathology whole slide image classification. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [47] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021. 2
- [48] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 2
- [49] Ruiqing Li, Xingqi Wu, Ao Li, and Minghui Wang. HFB-Surv: hierarchical multimodal fusion with factorized bilinear models for cancer survival prediction. *Bioinformatics*, 38(9):2587–2594, 2022. 3
- [50] Yanghao Li, Haoqi Fan, Ronghang Hu, Christoph Feichtenhofer, and Kaiming He. Scaling language-image pre-training via masking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23390–23400, 2023. 2
- [51] Tiancheng Lin, Hongteng Xu, Canqian Yang, and Yi Xu. Interventional multi-instance learning with deconfounded instance-level prediction. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [52] Weixiong Lin, Ziheng Zhao, Xiaoman Zhang, Chaoyi Wu, Ya Zhang, Yanfeng Wang, and Weidi Xie. Pmc-clip: Contrastive language-image pre-training using biomedical documents. In *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2023. 2
- [53] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 4
- [54] Ming Lu, Bowen Chen, Drew Williamson, Richard Chen, Ivy Liang, Tong Ding, Guillaume Jaume, Igor Odintsov, Andrew Zhang, Long Le, Georg Gerber, Anil Parwani, and Faisal Mahmood. Towards a visual-language foundation model for computational pathology. *Nature Medicine*, 2024. 2
- [55] Ming Y Lu, Drew FK Williamson, Tiffany Y Chen, Richard J Chen, Matteo Barbieri, and Faisal Mahmood. Data-efficient and weakly supervised computational pathology on whole-slide images. *Nature biomedical engineering*, 5(6):555–570, 2021. 1, 3, 4
- [56] Ming Y Lu, Bowen Chen, Andrew Zhang, Drew FK Williamson, Richard J Chen, Tong Ding, Long Phi Le, Yung-Sung Chuang, and Faisal Mahmood. Visual language pre-trained multiple instance zero-shot transfer for histopathology images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19764–19775, 2023. 1, 2, 4
- [57] Shima Mehrvar, Lauren E. Himmel, Pradeep Babburi, Andrew L. Goldberg, Magali Guffroy, Kyathanahalli Janardhan, Amanda L. Krempley, and Bhupinder Bawa. Deep learning approaches and applications in toxicologic histopathology: Current status and future perspectives. *Journal of Pathology Informatics*, 12(1):42, 2021. 3
- [58] Pooya Mobadersany, Safoora Yousefi, Mohamed Amgad, David A Gutman, Jill S Barnholtz-Sloan, José E Velázquez Vega, Daniel J Brat, and Lee AD Cooper. Predicting cancer outcomes from histology and genomics using convolutional networks. *Proceedings of the National Academy of Sciences*, 115(13):E2970–E2979, 2018. 3
- [59] Milad Mostavi, Yu-Chiao Chiu, Yufei Huang, and Yidong Chen. Convolutional neural network models for cancer type prediction based on gene expression. *BMC Med Genomics*, 2019. 2
- [60] Patience Mukashyaka, Todd Sheridan, Ali pour, and Jeffrey Chuang. Sampler: unsupervised representations for rapid analysis of whole slide tissue images. *eBioMedicine*, 99:104908, 2024. 1
- [61] A. Myronenko, Z. Xu, D. Yang, H.R. Roth, and D. Xu. Accounting for dependencies in deep learning based multiple instance learning for whole slide imaging. In *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pages 329–338, 2021. 3
- [62] Ramin Nakhli, Allen Zhang, Ali Mirabadi, Katherine Rich, Maryam Asadi, Blake Gilks, Hossein Farahani, and Ali Bashashati. Co-pilot: Dynamic top-down point cloud with conditional neighborhood aggregation for multi-gigapixel histopathology image representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 21063–21073, 2023. 3
- [63] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 2
- [64] Linhao Qu, xiaoyuan Luo, Manning Wang, and Zhijian Song. Bi-directional weakly supervised knowledge distillation for whole slide image classification. In *Advances in Neural Information Processing Systems*, 2022. 3
- [65] Linhao Qu, Zhiwei Yang, Minghong Duan, Yingfan Ma, Shuo Wang, Manning Wang, and Zhijian Song. Boosting

- whole slide image classification from the perspectives of distribution, correlation and magnification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 21463–21473, 2023. [3](#)
- [66] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [2](#), [4](#)
- [67] Md Mamunur Rahaman, Ewan K. A. Millar, and Erik Meijering. Breast cancer histopathology image-based gene expression prediction using spatial transcriptomics data and deep learning. *Sci. Rep.*, 13(13604):1–11, 2023. [8](#)
- [68] Benoit Schmauch, Alberto Romagnoni, Elodie Pronier, Charlie Saillard, Pascale Maillé, Julien Calderaro, Aurélie Kamoun, Meriem Sefta, Sylvain Toldo, Mikhail Zaslavskiy, Thomas Clozel, Matahi Moarii, Pierre Courtiol, and Gilles Wainrib. A deep learning model to predict rna-seq expression of tumours from whole slide images. *Nature Communications*, 11, 2020. [3](#)
- [69] Zhuchen Shao, Hao Bian, Yang Chen, Yifeng Wang, Jian Zhang, Xiangyang Ji, et al. Transmil: Transformer based correlated multiple instance learning for whole slide image classification. *Advances in Neural Information Processing Systems*, 34:2136–2147, 2021. [1](#), [3](#), [4](#), [5](#), [6](#)
- [70] Zhuchen Shao, Yifeng Wang, Yang Chen, Hao Bian, Shaohui Liu, Haoqian Wang, and Yongbing Zhang. Lnp1-mil: Learning from noisy pseudo labels for promoting multiple instance learning in whole slide image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 21495–21505, 2023. [3](#)
- [71] Taishi Shimazaki, Ameya Deshpande, Anindya Hajra, Tijo Thomas, Kyotaka Muta, Naohito Yamada, Yuzo Yasui, and Toshiyuki Shoda. Deep learning-based image-analysis algorithm for classification and quantification of multiple histopathological lesions in rat liver. *Journal of Toxicologic Pathology*, 35(2):135–147, 2022. [3](#)
- [72] Artem Shmatko, Narmin Ghaffari Laleh, Moritz Gerstung, and Jakob Nikolas Kather. Artificial intelligence in histopathology: enhancing cancer research and clinical oncology. *Nature Cancer*, 3(9):1026–1038, 2022. [1](#)
- [73] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15638–15650, 2022. [2](#)
- [74] Andrew H. Song, Guillaume Jaume, Drew F. K. Williamson, Ming Y. Lu, Anurag Vaidya, Tiffany R. Miller, and Faisal Mahmood. Artificial intelligence for digital and computational pathology. *Nature Reviews Bioengineering*, 2023. [1](#)
- [75] Andrew H Song, Richard J Chen, Tong Ding, Drew FK Williamson, Guillaume Jaume, and Faisal Mahmood. Morphological prototyping for unsupervised slide representation learning in computational pathology. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. [3](#)
- [76] Wenhao Tang, Sheng Huang, Xiaoxian Zhang, Fengtao Zhou, Yi Zhang, and Bo Liu. Multiple instance learning framework with masked hard instance mining for whole slide image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. [3](#)
- [77] Thomas Tavolara, Metin Gurcan, and M. Niazi. Contrastive multiple instance learning: An unsupervised framework for learning slide-level representations of whole slide histopathology images without labels. *Cancers*, 14:5778, 2022. [2](#)
- [78] Chao Tu, Yu Zhang, and Zhenyuan Ning. Dual-curriculum contrastive multi-instance learning for cancer prognosis analysis with whole slide images. In *Advances in Neural Information Processing Systems*, pages 29484–29497. Curran Associates, Inc., 2022. [3](#)
- [79] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. In *Neural Information Processing Systems (NeurIPS)*, 2017. [2](#)
- [80] Eugene Vorontsov, Alican Bozkurt, Adam Casson, George Shaikovski, Michal Zelechowski, Siqi Liu, Philippe Mathieu, Alexander van Eck, Donghun Lee, Julian Viret, Eric Robert, Yi Kan Wang, Jeremy D. Kunz, Matthew C. H. Lee, Jan Bernhard, Ran A. Godrich, Gerard Oakley, Ewan Millar, Matthew Hanna, Juan Retamero, William A. Moye, Razik Yousfi, Christopher Kanan, David Klimstra, Brandon Rothrock, and Thomas J. Fuchs. Virchow: A million-slide digital pathology foundation model, 2023. [1](#), [2](#)
- [81] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for vision and vision-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19175–19186, 2023. [2](#)
- [82] Xiyue Wang, Sen Yang, Jun Zhang, Minghui Wang, Jing Zhang, Junzhou Huang, Wei Yang, and Xiao Han. Transpath: Transformer-based self-supervised learning for histopathological image classification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 186–195. Springer, 2021. [2](#), [4](#)
- [83] Xiyue Wang, Jinxi Xiang, Jun Zhang, Sen Yang, Zhongyi Yang, Ming-Hui Wang, Jing Zhang, Yang Wei, Junzhou Huang, and Xiao Han. SCL-WC: Cross-slide contrastive learning for weakly-supervised whole-slide image classification. In *Advances in Neural Information Processing Systems*, 2022. [3](#)
- [84] Xiyue Wang, Sen Yang, Jun Zhang, Minghui Wang, Jing Zhang, Wei Yang, Junzhou Huang, and Xiao Han. Transformer-based unsupervised contrastive learning for histopathological image classification. *Medical image analysis*, 81:102559, 2022. [1](#), [2](#), [4](#), [6](#)
- [85] Zhiqin Wang, Ruiqing Li, Minghui Wang, and Ao Li. GPDBN: deep bilinear network integrating both genomic data and pathological images for breast cancer prognosis prediction. *Bioinformatics*, 37(18):2963–2970, 2021. [3](#)
- [86] Zifeng Wang, Zhenbang Wu, Dinesh Agarwal, and Jimeng Sun. Medclip: Contrastive learning from unpaired medical



- images and text. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3876–3887, 2022. 2
- [87] Jinxi Xiang and Jun Zhang. Exploring low-rank property in multiple instance learning for whole slide image classification. In *The Eleventh International Conference on Learning Representations*, 2023. 3
- [88] Ronald Xie, Kuan Pang, Sai W. Chung, Catia T. Perciani, Sonya A. MacParland, Bo Wang, and Gary D. Bader. Spatially resolved gene expression prediction from h&e histology images via bi-modal contrastive learning. In *NeurIPS*, 2023. 3
- [89] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: a simple framework for masked image modeling. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9643–9653, 2022. 2
- [90] Yingxue Xu and Hao Chen. Multimodal optimal transport-based co-attention transformer with global structure consistency for survival prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 3
- [91] Jiawei Yang, Hanbo Chen, Yuan Liang, Junzhou Huang, Lei He, and Jianhua Yao. Concl: Concept contrastive learning for dense prediction pre-training in pathology images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 523–539, 2022. 3
- [92] J. Yao, X. Zhu, J. Jonnagaddala, N. Hawkins, and J. Huang. Whole slide images based cancer survival prediction using attention guided deep multiple instance learning networks. *Medical Image Analysis*, 65, 2020. 3
- [93] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022. 2
- [94] Zhimiao Yu, Tiancheng Lin, and Yi Xu. Slpd: Slide-level prototypical distillation for wsis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 259–269. Springer, 2023. 2, 3
- [95] Hongrun Zhang, Yanda Meng, Yitian Zhao, Yihong Qiao, Xiaoyun Yang, and Yalin Zheng. Dtf-d-mil: Double-tier feature distillation multiple instance learning for histopathology whole slide image classification. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18780–18790, 2022. 3
- [96] Fengtao Zhou and Hao Chen. Cross-modal translation and alignment for survival analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 21485–21494, 2023. 3, 6
- [97] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. Image BERT pre-training with online tokenizer. In *International Conference on Learning Representations*, 2022. 1, 2, 4, 5