# PRISM: A Multi-Modal Generative Foundation Model for Slide-Level Histopathology

George Shaikovski[†1], Adam Casson[†1], Kristen Severson[2], Eric Zimmermann[2], Yi Kan Wang[1], Jeremy D. Kunz[1], Juan A. Retamero[1], Gerard Oakley[1], David Klimstra[1], Christopher Kanan[1,3], Matthew Hanna[1,4], Michal Zelechowski[1], Julian Viret[1], Neil Tenenholtz[2], James Hall[2], Nicolò Fusi[2], Razik Yousfi[1], Peter Hamilton[1], William A. Moye[1], Eugene Vorontsov[1], Siqi Liu[‡1], and Thomas J. Fuchs[1]

[1]Paige, NYC, NY United States
[2]Microsoft Research, Cambridge, MA United States
[3]University of Rochester, Rochester, NY United States
[4]Memorial Sloan Kettering Cancer Center, NYC, NY United States

## Abstract

Foundation models in computational pathology promise to unlock the development of new clinical decision support systems and models for precision medicine. However, there is a mismatch between most clinical analysis, which is defined at the level of one or more whole slide images, and foundation models to date, which process the thousands of image tiles contained in a whole slide image separately. The requirement to train a network to aggregate information across a large number of tiles in multiple whole slide images limits these models' impact. In this work, we present a slide-level foundation model for H&E-stained histopathology, *PRISM*, that builds on Virchow tile embeddings and leverages clinical report text for pre-training. Using the tile embeddings, PRISM produces slide-level embeddings with the ability to generate clinical reports, resulting in several modes of use. Using text prompts, PRISM achieves zero-shot cancer detection and sub-typing performance approaching and surpassing that of a supervised aggregator model. Using the slide embeddings with linear classifiers, PRISM surpasses supervised aggregator models. Furthermore, we demonstrate that fine-tuning of the PRISM slide encoder yields label-efficient training for biomarker prediction, a task that typically suffers from low availability of training data; an aggregator initialized with PRISM and trained on as little as 10% of the training data can outperform a supervised baseline that uses all of the data.

## 1 Introduction

Recent progress in computational pathology, fueled by deep learning and extensive pathology image datasets, has facilitated the creation of clinical decision support systems with the primary goal of detecting cancer in whole slide images (WSIs) [6, 34, 39, 31]. Already, computational pathology tools are available for use in clinical practice [33]. Motivated by this success, the field has grown to encompass a broad range of clinically-relevant tasks including cancer sub-typing, microenvironment characterisation, biomarker detection [13, 44], treatment response, and overall survival prediction. This research has the potential to revolutionize the field of pathology and, ultimately, clinical oncology practices.

To build on this progress, recent efforts have focused on creating foundation models capable of operating on tile images (small sub-regions of the WSI), trained with self-supervised learning on datasets ranging from tens of thousands to

---

[‡]Corresponding author. siqi.liu AT paige DOT ai

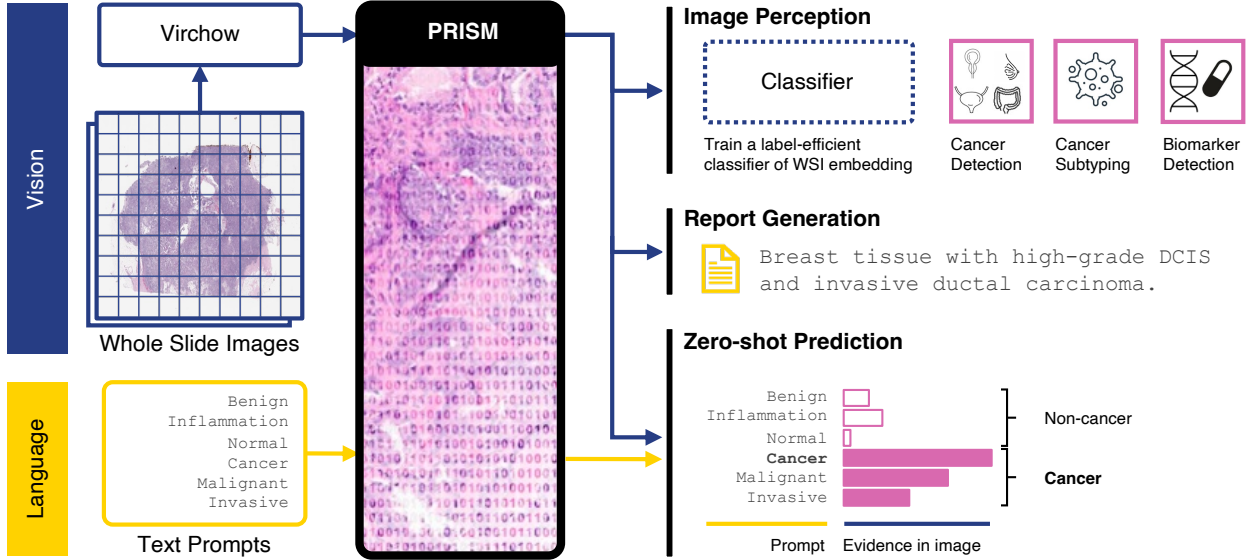[†]These authors contributed equally to this work.

Figure 1: An overview of the capabilities enabled by the slide-level foundation model (PRISM), built on Virchow [43] tile embeddings. Whereas Virchow produces an embedding for each foreground tile of a set of whole slide images, PRISM aggregates these embeddings into a single slide embedding that can be used for image perception by training a linear classifier for downstream tasks including cancer detection, cancer sub-typing, and biomarker detection. Optionally, the model can be fine-tuned for the classification task. Language-enabled capabilities of PRISM include the training-free "zero-shot" prediction via text prompting, and generation of interpretable free text clinical reports.

millions of WSIs [46, 17, 10, 43]. These models are constrained to tile images because WSIs are gigapixels in size, or approximately five orders of magnitude larger than a typical natural image. At such a scale, hardware constraints limit the applicability of many deep learning architectures. While foundation models can act as building blocks for a wide range of downstream tasks, most ground truth labels obtainable from clinical databases without manual annotations—such as diagnosis, biomarker status, treatment response, or survival—are linked to a whole slide image or a collection of WSIs created from the same tissue sample. Therefore, these foundation models must be paired with multiple instance learning (MIL) aggregator models [6, 20, 24, 26, 37] to map the tile-level embeddings to slide or specimen-level ground truth labels. Although tile-level foundation models have been pre-trained, aggregator models are typically trained from scratch on each slide-level task. Given a large number of high-dimensional embeddings per WSI, these aggregator models are prone to overfitting, particularly when trained on small datasets. We believe that large-scale pre-training on whole slide images with natural language supervision can tackle the challenge of high dimensionality and limited label availability, and create a generally-applicable slide-level foundation model for pathology.

In this work, we present a multi-modal slide-level foundation model named *PRISM*, for Pathology Report and Image Summarization Model. To pre-train the model, we use clinical report data as WSI-level supervisory signal. We show that the proposed pre-training method improves performance on downstream tasks, compared to fully-supervised models. Importantly, we demonstrate that pre-training can also benefit specific downstream tasks not covered by the report text. As clinical reports can be obtained without further expensive manual annotation, the proposed pre-training approach can be easily scaled to millions of WSIs. PRISM is pre-trained using 587 thousand WSIs and 195 thousand associated clinical text reports. The model is capable of generating text-based diagnosis reports for WSIs. While the primary objective of this work is not to produce clinically usable reports, we demonstrate that this pre-training method produces the model that can accurately identify various WSI features and match the slides to a correct prompt (zero-shot prediction) or generate a text caption describing the features (report generation), without further training. These include tissue types, the presence and sub-types of cancer, non-neoplastic patterns such as Crohn's disease, inflammation, polyps, and other diagnostic information. Besides report generation and zero-shot prediction, a *slide embedding* can be obtained from the slide encoder of PRISM for WSI-level linear classification with specific ground truth labels; alternatively, the model weights can also be further fine-tuned on smaller datasets (image perception). See Fig 1 for the overview of PRISM's capabilities.

We evaluate PRISM on cancer detection, tissue sub-typing, and biomarker prediction tasks using zero-shot classification, linear probing, and fine-tuning. Zero-shot cancer classification and sub-typing performance approaches or surpasses that of the supervised cancer detection baselines [43]. Fine-tuned on various biomarkers, PRISM shows su-

perior performance in the low data regime compared to the supervised baselines, and with limited training data, often outperforms supervised training on the full dataset. We also perform a qualitative evaluation of randomly selected generated reports to contextualize PRISM performance. Our contributions can be summarised as follows:

1. A generative slide-level pathology foundation model pre-trained with clinical report supervision.
2. A memory-efficient training methodology to enable WSI-level pre-training.
3. Evaluation of the model's diagnostic capabilities and label efficiency on clinically relevant tasks.

## 2   Related Work

**Tile-level foundation models.** Vision-based foundation models have recently emerged in computational pathology for WSI tiles. These self-supervised models are trained with contrastive learning methods. Initial works employed public datasets, including The Cancer Genome Atlas (TCGA), with tiles extracted from on the order of 10,000 WSIs to train models with up to 307 million parameters, based on both convolutional and transformer architectures [46, 12, 2, 17, 23]. More recent works trained a vision transformer (ViT) [16] (22 million to 632 million parameters) on larger proprietary datasets, including 100,000 to 1.5 million WSIs [7, 15, 10, 43]. The Uni work [10] demonstrates that model performance scales with dataset size. Furthermore, Virchow [43] demonstrates that performance can be further improved by scaling both the model and dataset sizes, approaching clinically relevant cancer prediction performance for specimen-level aggregators trained on Virchow tile embeddings. Nevertheless, the need to train an aggregator network makes it difficult to substantially improve performance on low data tasks like biomarker prediction. This motivates pre-training of the aggregator network in a more complete foundation model.

Some recent works incorporate language for pre-training tile-level models [18, 19, 27, 28]. Gamper et al. [18] encode all tissue tiles with a convolutional neural network (CNN) and pass to a transformer for autoregressive caption decoding. Later methods build on the CLIP [36] approach which introduces a contrastive objective between the representations of a tile and its corresponding description. PLIP [19] extends this to pathology images collected from educational materials. CONCH [27] adopts the CoCa framework [47] that adds a generative objective to CLIP, allowing the model to learn autoregressive caption generation from tiles. Both PLIP and CONCH can be used at the slide-level by simple top-k max-pooling (or more complex regional graph-based pooling) of tile-level predictions. Our work also builds on the CoCa framework but takes it further by training an aggregator at the specimen level (including one or more WSIs) to predict summarized clinical reports. The Perceiver architecture [21] is used to make it possible to aggregate over hundreds of thousands of tiles.

**Slide-level foundation models.** Most work on pre-trained pathology models has focused on tile-level representations. There are two notable exceptions: the hierarchical image pyramid transformer (HIPT) [9] and LongViT [45]. HIPT proposes a hierarchical two-stage model where a ViT learns tile representations in the first stage and another ViT aggregates the class token representations of regional tiles. Whereas HIPT only summarizes local regions, LongViT summarizes a full WSI using LongNet [14] attention to allow for very long context lengths. This relies on dilated attention in each head looking at a unique subset of a sequence of tiles. These works differ from the proposed method in that they rely on self-supervised objectives that lack clinical-report-based supervision. We posit that this supervision may help to disentangle more clinically relevant features. Clinical report generation in the proposed method also enables zero-shot prediction.

**Pathology report generation.** Early pathology report generation work for WSI used CNN tile encoding and recurrant neural network (RNN) caption decoding [48, 40, 32]. Later work improved the spatial context inferred from the image by using HIPT instead of a CNN. These methods relied on randomly subsampling the tiles in WSI and using memory-efficient RNNs which are slower to train and harder to parallelize at scale than the transformers commonly used in state of the art language models.

## 3   Slide-Level Foundation Model

This work combines natural-language supervision and cross-attention-based resampling to overcome the computational challenges of training with whole slide images while effectively directing the learning algorithm to find important, generalizable features in the slides. PRISM, the slide-level foundation model, contains two components: a slide encoder which leverages a Perceiver network [21], and a language decoder which leverages the BioGPT language model [29]. Using paired samples of clinical reports, which are rewritten using GPT-4 [1], and WSIs, which are pre-processed to create a sequence of tile-level embeddings using our Virchow foundation model [43], PRISM is trained following the CoCa methodology [47] which uses two objectives: (1) alignment of the the encoded report embedding from the pre-trained BioGPT [29] language decoder with the slide latent embedding, and (2) prediction
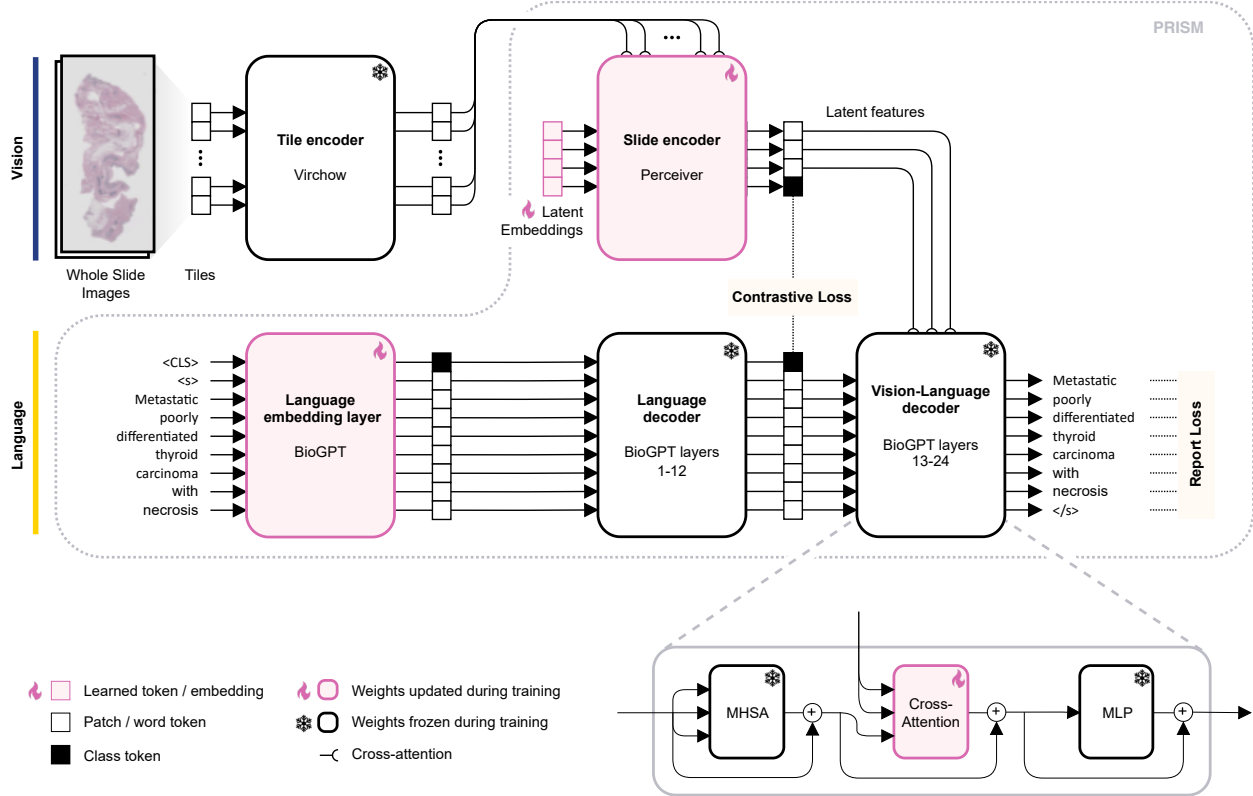
Figure 2: The training methodology for the slide-level foundation model (PRISM). All trained weights are initialized to random values except for the BioGPT word embeddings. Whole slide images and clinical report latent embeddings are aligned with a contrastive loss. Report generation is trained with a generative loss using teacher forcing. Layers 13-24 of the BioGPT decoder are modified to cross-attend to vision embeddings.

of the generated caption by BioGPT to match the ground-truth rewritten report. The full pre-training framework is shown in Fig. 2 and described in detail below. After pre-training, the combined Virchow-PRISM model can be used to produce a single embedding of one or more WSIs, and generate reports to describe histological features in the slides.

**Pre-processing of clinical reports.** A training sample for PRISM contains multiple whole slide images grouped into a specimen, and the associated natural-language clinical report. The reports typically contain a mixture of pre-defined structured fields (synoptic reports) and free-text fields, although the formats can vary significantly, particularly for different sites of tissue samples. The reports are first processed to extract the clinical diagnosis section using heuristic text matching rules. If available, immunohistochemistry (IHC) or molecular test results are appended to the diagnosis section.

The concatenated reports are rewritten into concise text summaries using GPT-4 [1] to reduce their length, and increase the density of relevant diagnostic information per report. The rewritten summaries may include, but are not limited to, sites of tissue sample, presence of cancer, its sub-type and grade if available, non-cancerous tissue, IHC status, and potential molecular testing results. Each specimen report is rewritten five times by GPT-4 (Section A.7) to be sampled randomly during training to prevent the network from memorising specimen-to-report associations.

**Tile embeddings.** During training and inference, each WSI is divided into a uniform grid of 224×224 pixel image tiles at a 20× magnification level, corresponding to a resolution of 0.5 µm/pixel. Tiles with at least 25% tissue by area are used during training and the rest are discarded. To estimate the amount of tissue in a tile, each WSI is downsampled 16× with bilinear interpolation and every pixel of the downsampled image is evaluated as to whether its hue, saturation, and value are within [90, 180], [8, 255], and [103, 255], respectively. These tiles are subsequently processed with Virchow [43], a tile-level foundation model, to produce tile embeddings. The Virchow model is a 632M parameter vision transformer (ViT-H/14) [16] trained using DINOv2 [30] on over 2B image tiles extracted from 1.5 million hematoxylin and eosin (H&E)-stained WSIs. This large scale pre-training leads to embeddings which can capture the large diversity of tissue and morphological patterns [43]. The resulting embedding is a 2560-dimensional

4

vector constructed by concatenating the output class token (1280 dimensions) with the mean across all output patch tokens (1280 dimensions).

**WSI embeddings.** Encoding whole slide images into a single global representation is highly compressive, given that each specimen can consist of hundreds of thousands of tiles. We use the memory-efficient Perceiver network [21] to aggregate the long sequence of specimen tile embeddings into a much smaller fixed-length set of latent features using trainable latent embeddings ("latents") and an asymmetric cross-attention mechanism.

The Perceiver slide aggregator consists of 8 blocks (see Section A.9 for implementation details). Each block has a cross-attention module followed by a 6-layer latent self-attention transformer [42]. A cross-attention module takes tile embeddings and latents as inputs, and outputs latents. A latent transformer processes the latents returned by the preceding cross-attention module. The weights of the cross-attention modules from the 2nd through the 8th Perceiver blocks are shared. Similarly, the weights of the latent transformers from the 1st through the 8th Perceiver blocks are also shared. The context key-value pairs computed from the input embeddings are cached between the weight-sharing cross-attention modules to decrease memory consumption.

In this study, the Perceiver network uses 513 learned latents with 1280 dimensions. After being processed by the network, 512 output latents are used as latent context features in the language decoder to generate reports. The remaining output latent is the *slide embedding* used for the contrastive loss described below. We found that using an extra slide latent, instead of mean-pooling introduced by Jaegle et al. [21], produced more accurate reports. The slide embedding serves as a global slide-level or specimen-level representation in downstream applications.

**Text embeddings and report generation.** We use the pre-trained BioGPT [29] as the language network for PRISM. BioGPT is a decoder-only language model tailored to biology applications and is based on the 345M parameter variant of the GPT-2 architecture [35]. The total BioGPT vocabulary is 42,384 tokens. We add a class token (<CLS> in Figure 2) to the vocabulary to produce a global clinical report representation. We also add cross-attention modules for context-aware report generation using latent features from the Perceiver model. We use the BioGPT tokenizer and the embedding layer to encode the training text prompt as a sequence of word embeddings with positional information.

Following Yu et al. [47], we split the decoder into uni-modal and multi-modal parts. The uni-modal part includes the layers 1 through 12. We append the class (<CLS>) token to the input sequence. Self-attention over language tokens uses causal masking to prevent look-ahead when generating the next tokens, while the <CLS> token attends to all language tokens and to itself. After layer 12 we separate the <CLS> token embedding from the sequence. This embedding represents the ground truth clinical report and is used in computing the contrastive loss via alignment with the corresponding slide embedding elaborated below. Further in the text we refer to it as report embedding or prompt embedding depending on the model application context. The remaining language tokens proceed to the multi-modal part of the decoder, which are the layers 13 through 24. Each layer is augmented with a cross-attention module, inserted between the self-attention module and the MLP network, which take 512 latent features from the Perceiver as key-value context.

We initialize the network with pre-trained BioGPT weights and freeze all of them except the initial word embedding layer. The initial class token embedding and the weights of cross-attention modules in the multi-modal part are initialized to random values from the Normal distribution.

**Training objective.** Following CoCa [47], the model is trained using contrastive and generative objectives. To compute the contrastive loss, the slide embedding from the Perceiver and the paired report embedding from BioGPT are projected using a linear layer with 5120 output dimensions and normalized with the Euclidean $\ell_2$ norm. The objective maximizes the cosine similarity between the paired projections, while minimizing the similarity of the unmatched pairs. Specifically, it scales logits with a learned temperature parameter as in CoCa, and computes the symmetric cross-entropy loss:

$$\mathcal{L}_{con} = -\frac{1}{N}\left(\sum_{i=1}^{N}\log\frac{\exp(\boldsymbol{v}_i^\top \boldsymbol{t}_i/\tau)}{\sum_{j=1}^{N}\exp(\boldsymbol{v}_i^\top \boldsymbol{t}_j/\tau)} + \sum_{i=1}^{N}\log\frac{\exp(\boldsymbol{t}_i^\top \boldsymbol{v}_i/\tau)}{\sum_{j=1}^{N}\exp(\boldsymbol{t}_i^\top \boldsymbol{v}_j/\tau)}\right), \tag{1}$$

where $(\boldsymbol{v}_i, \boldsymbol{t}_i)$ are the vision and language projections respectively, $\tau$ is a temperature parameter learned during training, and $N$ is the batch size.

The report (captioning) objective for report generation minimizes the negative log-likelihood of individual tokens using the factorized joint distribution with teacher-forcing:
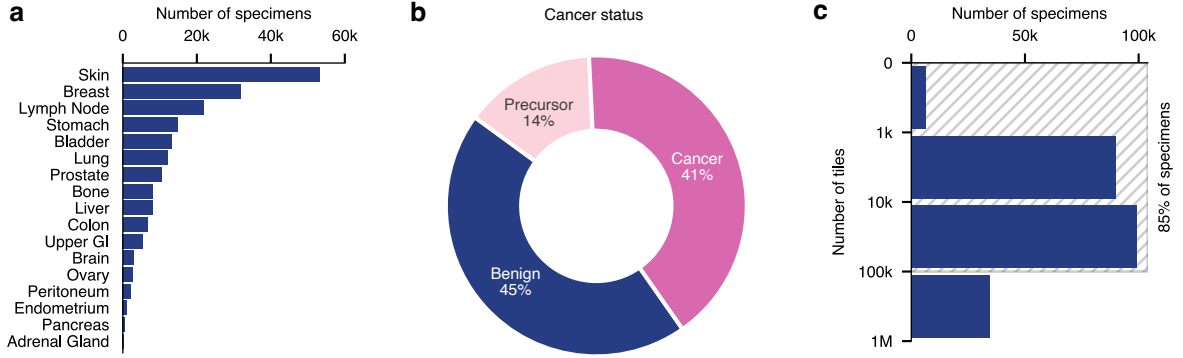
Figure 3: Statistics on the specimen-level pre-training dataset for PRISM. Note that a specimen may contain one or more whole slide images (WSIs). **a.** Distribution of specimens by the site of tissue origin. **b.** Proportion of data with the most severe diagnosis being cancer, precursor to caner, or benign. Note for example that a specimen with cancer may also have a precursor to cancer. **c.** The histogram of tile counts per specimen. 85% of the specimens (195,344 specimens) have fewer than 100 thousand tiles (plots **a** and **b** describe this subset).

$$\mathcal{L}_{rep} = -\sum_{t=1}^{T} \log p(y_t|y_{<t}, \boldsymbol{X}),$$ (2)

where $y_t|y_{<t}$ is the autoregressive token prediction for token $t$ of $T$ and $\boldsymbol{X}$ is the latent features from the slide encoder.

Besides learning the report generation capability, Yu et al. [47] show that the captioning objective improves instance-level representations compared to using the contrastive objective alone.

The final training objective is a weighted sum of the two the contrastive and generative losses:

$$\mathcal{L}_{tot} = \lambda_{con}\mathcal{L}_{con} + \lambda_{rep}\mathcal{L}_{rep}.$$ (3)

## 4 Experimental Methods

### 4.1 Training protocols

We collected a dataset of 587,196 whole slide images corresponding to 195,344 specimens, where each specimen is a collection of one or more WSIs with a corresponding clinical report.

Applying our tiling scheme to our dataset results in more than 500,000 tiles for the largest specimen. We constrain our dataset to specimens with up to 100,000 tile embeddings to increase the batch size and speed up training. This constraint does not limit the applicability of the model to clinically-relevant tasks since during application, inference is usually performed on individual slides rather than on the whole specimen and the slides in our dataset have less than 100,000 foreground tiles. See Fig. 3 for the slide and specimen distributions over site of tissue origin and tumor types in the training dataset respectively.

We train the model for 10 epochs on 16 NVIDIA V100 32GB GPU's using fp16 precision with a global batch size of 64 and gradient accumulation over 4 iterations. The model is updated using AdamW with the base learning rate of $2 \cdot 10^{-4}$ and a cosine decay schedule over 75,000 iterations with 2,000 warm-up iterations. Optimizer weight decay is held fixed at $1 \cdot 10^{-6}$.

### 4.2 Evaluation protocols

**Image perception** is performed using either linear probing of fine-tuning. Linear probing maps the slide embedding from the slide encoder of PRISM (the Perceiver network) to the slide or specimen-level class label. The Perceiver weights are not updated. We implement linear probing using logistic regression with 5-fold cross validation and a search over a range of $\ell_2$-regularization coefficients from 1.0 to $10^6$ in powers of 10.

To fine-tune the slide encoder we attach a linear classifier as a task head on top of the slide embedding. The classifier's weights are initialized to random values while the slide encoder's weights are pre-trained. We use the same data splits as in linear probing. The Virchow tile-level foundation model is kept frozen.

**Report generation**. PRISM can generate a clinical report for a slide or a specimen using autoregressive decoding. If tumor is present in the slides, the report describes its type and sub-type, whether it is malignant and/or metastatic, the cancer grade, non-cancerous but relevant features in the surrounding tissue, and the site of tissue origin. In addition, it may predict the results of IHC staining or molecular tests. Otherwise, the model describes the observed benign tissue.

To generate the reports, the slide encoder (Perceiver) takes a sequence of specimen tile embeddings while the language model is prompted with the special token that indicates the start of a sentence. The model outputs a probability distribution over its vocabulary and selects the token with the highest output probability (top-1). The selected token becomes the ground truth for the next token prediction. The model iteratively decodes the report, halting once the end of sentence token is produced.

**Zero-shot prediction**. The contrastive training objective allows for zero-shot evaluation, showcasing how effectively the model can learn clinically relevant tasks directly from diagnostic reports without explicit task-specific training. It establishes the baseline performance of pre-trained PRISM on novel tasks without further training.

Zero-shot evaluation can be applied to binary classification tasks by creating sets of negative and positive prompts. Specimen slides are processed by the Virchow tile encoder and the slide encoder of PRISM to produce a slide embedding ($\ell_2$-normalized to $\mathbf{v}$), while the language decoder encodes every prompt as a prompt embedding ($\ell_2$-normalized to $\mathbf{t}$), with $\boldsymbol{T}$ being the set of all prompt embeddings. We compute the cosine similarity between each prompt embedding and the slide embedding and apply a softmax activation to the resulting logits with temperature $\tau$ learned during training. Each entry in the vector corresponds to a probability of that prompt matching the slide. In the case of multiple negative and/or positive prompts we perform prompt ensembling by marginalizing probability scores over each set of prompts. Specifically, a subset of prompt embeddings that belong to class $c$ is defined as $\boldsymbol{T}_c$, and the probability of the class $c$ being the true class given the slide embedding $v$ is

$$p(C = c \mid \boldsymbol{v}) = \sum_{\boldsymbol{t}_c \in \boldsymbol{T}_c} \frac{\exp(\boldsymbol{v}^\top \boldsymbol{t}_c / \tau)}{\sum_{\boldsymbol{t} \in \boldsymbol{T}} \exp(\boldsymbol{v}^\top \boldsymbol{t} / \tau)}. \tag{4}$$

**Supervised baselines trained from scratch.** We include a fully-supervised baseline performance for every task to contextualize PRISM's capabilities. Supervised baselines mirror the fine-tuning protocol except the slide encoder weights are initialized to random values.

### 4.3  Evaluation tasks and data

**Cancer detection.** We evaluated cancer detection performance as a binary classification of the presence of cancer in a specimen of multiple WSIs. We stratified the results in two settings: *All Cancers* and *Rare Cancers* (see [43]). *All Cancers* contains 22,932 WSIs from 6,142 specimens corresponding to 16 cancer types. The samples are sourced from Memorial Sloan Kettering Cancer Center (MSKCC) (49%) and other institutions (51%). *Rare Cancers* are a subset of *All Cancers* for specimens which meet the National Cancer Institute definition of rare, an incidence of less 15 per 100K per year in the United States. *Rare Cancers* comprises 8753 slides from 2595 specimens corresponding to 7 of the 16 cancer types.

**Cancer sub-typing.** We evaluated the cancer sub-typing capabilities of the pre-trained network on three binary tasks: invasive ductal carcinoma (IDC) versus invasive lobular carcinoma (ILC) with TCGA breast cancer (BRCA) data, lung adenocarcinoma (LUAD) versus lung squamous cell carcinoma (LUSC) with TCGA non-small cell lung cancer (NSCLC) data, and ductal carcinoma in situ (DCIS) vs IDC on internal MSKCC data.

**Biomarker prediction.** We evaluated the biomarker prediction performance on 9 different biomarkers originally identified by MSK-IMPACT [11], a targeted test for genetic mutations, described in detail in Appendix A.5 and Table A.5.1.

Cancer detection and sub-typing with PRISM are evaluated with zero-shot classification, linear probing, and fine-tuning and compared to a baseline slide encoder (Perceiver) trained from scratch. Biomarker prediction is evaluated only with and without slide encoder pre-training (fine-tuning vs supervised).

Most evaluation datasets contain out-of-distribution (OOD) data: both Virchow and PRISM were trained only on data internal to MSKCC, whereas 49% of the cancer detection dataset in Table 1 was composed of cases submitted

| Evaluation | Pre-trained | Cancer sub-typing | | | Cancer detection | |
|---|---|---|---|---|---|---|
| | | BRCA[†] | NSCLC[†] | DCIS/IDC[‡] | Rare Cancers[‡] | All Cancers[‡] |
| Baseline | | 0.949 | 0.980 | 0.876 | 0.925 | 0.947 |
| Zero-shot | ✓ | 0.952 | 0.975 | 0.908 | 0.868 | 0.906 |
| Linear probe | ✓ | **0.958** | **0.983** | **0.939** | 0.931 | 0.949 |
| Fine-tuned | ✓ | 0.957 | 0.978 | 0.924 | **0.938** | **0.952** |

Table 1: Area under (the receiver operating characteristic) curve (AUC) performance achieved by PRISM pre-trained slide encoder and compared to a slide encoder baseline trained from scratch on cancer sub-typing and detection tasks (from TCGA[†] or MSKCC[‡]). Note that all evaluation methods used pre-trained Virchow embeddings for WSI tiles.

to MSKCC from external sites; data from TCGA was not seen during model pre-training. We do not apply stain normalization during evaluation.

## 5  Results

### 5.1  PRISM demonstrates diagnostic capability

PRISM pre-trained with pathology reports supervision learns diagnostic capabilities demonstrated on the cancer detection and sub-typing tasks (see Table 1). In all cases, performance using the pre-trained model is better than training a model from scratch. Below we describe the results for the zero-shot, linear probing, and fine-tuned settings in detail.

Zero-shot classification performance is comparable to the supervised baseline on the TCGA cancer sub-typing tasks: ILC vs IDC in BRCA (0.952 AUC compared to 0.949, Table 1) and LUSC from LUAD in NSCLC (0.975 AUC compared to 0.980, Table 1). Zero-shot classification outperforms the supervised baseline on DCIS vs IDC sub-typing in breast on our internal MSKCC dataset. These results were achieved without tuning or ensembling text prompts—each class in all sub-typing tasks was encoded using a single prompt selected to resemble the wording in the training data for these sub-types (Table A.1.1).

On MSKCC cancer detection, zero-shot classification performed well but worse than the supervised baseline (0.868 AUC vs 0.925 on rare cancers and 0.906 AUC vs 0.947 on all cancers, Table 1). There are a few key differences that likely explain the differences in zero-shot performance across the tasks. First, the cancer detection dataset is much larger than the sub-typing datasets (40,402 specimens in the training set), making it likely that supervised training would perform well. Second, the pan-cancer detection dataset includes 16 cancer types spanning diverse tissue types, 20% of which were not included in the training set of either Virchow or PRISM. This distribution shift has implications for zero-shot classification as it is possible that some classes have never been observed. Unlike in the cancer sub-typing setting, for cancer detection we ensemble each class over a set of seven prompts (see Tab. A.1.1). For the set of



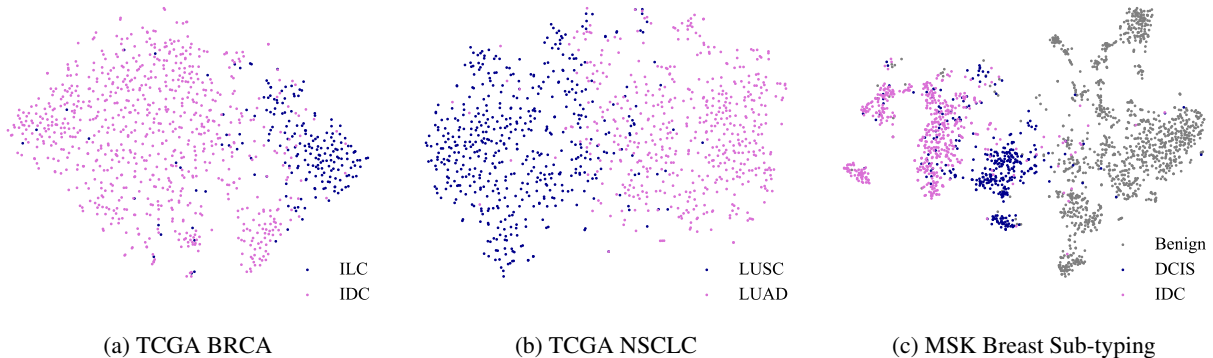(a) TCGA BRCA          (b) TCGA NSCLC          (c) MSK Breast Sub-typing

Figure 4: t-SNE plots of slide embeddings for cancer sub-typing datasets. IDC and ILC are types of breast cancer. LUSC and LUAD are types of non-small cell lung cancer. DCIS is an early stage non-invasive breast cancer and a precursor to IDC; some IDC slides can contain DCIS regions, however IDC takes precedence in the diagnosis as a higher stage cancer. All three plots suggest distinct clusters of slide embeddings in higher dimensions along the cancer sub-types labels.
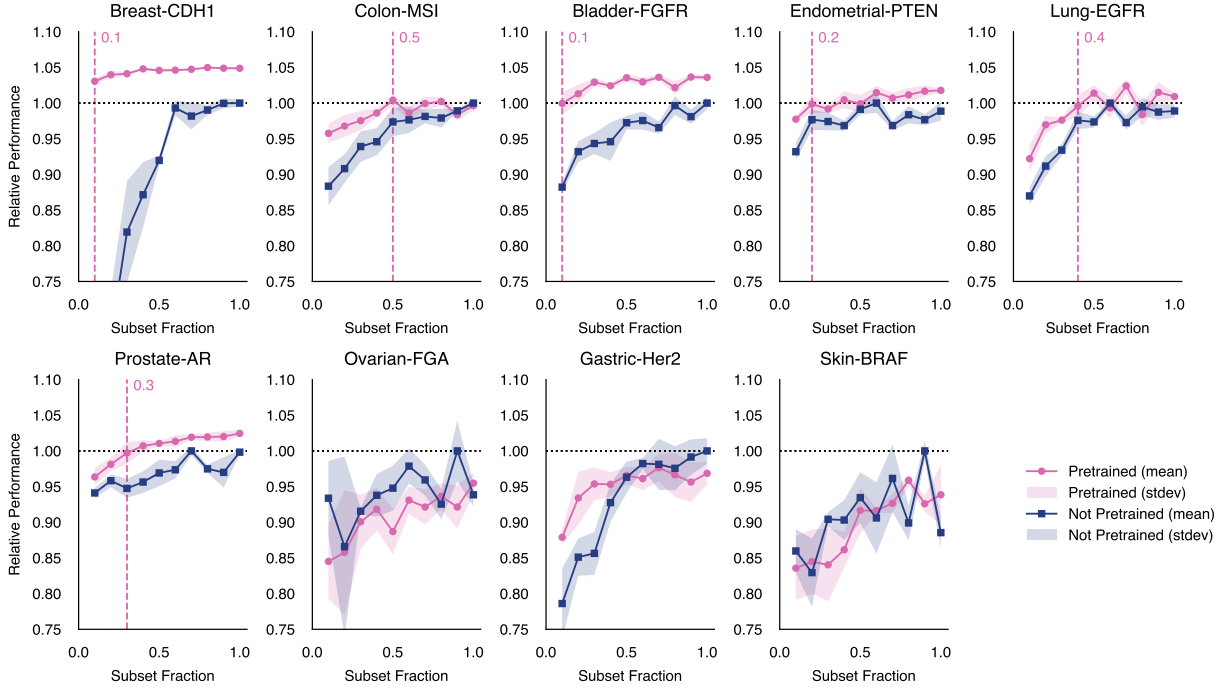
Figure 5: Fine-tuning pre-trained slide encoder (PRISM Perceiver) improves data efficiency compared to training the same model from scratch. The mean (standard deviation) performance across 3 experimental runs is plotted as a solid line (shaded region), relative to the highest area under (the receiver operating characteristic) curve (AUC) achieved without pre-training. The subset fraction denotes the fraction of the available training data used to fine-tune or train the model. A different random subset is selected with each experimental run. The vertical dashed line (magenta) denotes the minimal subset fraction required when fine-tuning PRISM to reach at least 99.5% of the AUC that can be achieved without pre-training. Note that in many cases, pre-training yields better performance than can be achieved when training from scratch.

positive prompts, we select seven most common cancer types in the pre-training dataset. This is because the reports in the training dataset do not mention "cancer", instead using more precise terms describing its type, such as lymphoma, sarcoma, adenocarcinoma, and others. Likewise, the reports don't explicitly state the absence of cancer. Instead, they mention the type of tissue present in the WSIs, which can include everything from normal tissue described as normal or unremarkable, to benign tumors, inflammation and special cases of inflammation such as Crohn's disease, suspicious or atypical hyperplasia, cancer precursors, and carcinoma in situ. The non-cancer prompts for this task include the seven types of most commonly observed non-cancer histology in the pre-training dataset. Both sets of prompts do not entirely cover the histology present in the evaluation cancer detection dataset, possibly contributing to reduced performance relative to other evaluation methods. Prompt tuning, using other zero-shot methods like caption scoring, as well as enabling zero-shot cancer detection without the need for ensembling, may improve zero-shot performance and is left as future work.

Linear probing performance surpasses the supervised baseline and zero-shot prediction in all settings. In the cancer sub-typing-tasks, linear probing also outperforms fine-tuning. We use t-SNE to illustrate the separability of the pre-trained slide embeddings of the sub-typing tasks (see Fig. 4). Linear probing performance on the MSKCC cancer detection is competitive with, but slightly below, the fine-tuned model.

Fine-tuning pre-trained Perceiver leads to better performance than training from scratch and zero-shot and approaches linear probing on the cancer sub-typing tasks. On the MSKCC cancer detection task fine-tuning is the best performing method. We suspect the performance gap between linear probe and fine-tuning for sub-typing could be due to overfitting on the relatively small datasets; fine-tuning has superior performance on the much larger MSKCC detection dataset. Whole slide aggregation models can be sensitive to the hyperparameters used for training [5] and that hyperparameter tuning could result in better fine-tuning performance.

9

## 5.2    PRISM allows label-efficient training

Biomarker prediction from routine H&E stained WSI can reduce the delays for patients by removing the need for additional molecular and IHC testing. We selected nine biomarkers that play an important role in the diagnosis and treatment of various cancers (see Section A.5 for details). The biomarker datasets encode a binary status of genetic alterations ("biomarker") as measured by DNA extraction and MSK-Integrated Mutation Profiling of Actionable Targets (MSK-IMPACT) sequencing.

Among nine biomarkers, we compare the performance between fine-tuning the pre-trained slide encoder of PRISM and training the same slide encoder from scratch. As shown in Fig. 5, pre-training yielded higher detection performance (AUC) as compared to not pre-training in 6 of the 9 prediction tasks, by as much as 5% in the case of Breast-CDH1. Furthermore, fine-tuning the pre-trained model was much more label efficient in all 6 tasks. Biomarker prediction was tested with random subsets of each training set in 10% increments from 10% to 100% of the available data. For breast-CDH1 and bladder-FGFR, only 10% of training data was necessary to meet or exceed the maximum performance attained without pre-training; for endometrial-PTEN, prostate-AR, lung-EGFR, and colon-MSI, 20%, 30%, 40%, and 50% of the training data was necessary, respectively. Finally, pre-training appears to lower the variance across experiment re-runs for some biomarkers, yielding more consistent results.

## 5.3    PRISM predictions are interpretable

To contextualize the performance of PRISM and evaluate the explainability of generated reports, we evaluate the 16 tiles most attended by the latent slide embedding. High attention scores (in the last cross-attention layer of the Perceiver slide encoder) indicate that when pooling information across the slide, these tiles have more impact on the final slide embedding than other tiles with lower scores. Four specimens were randomly selected for a thorough review. Fig. 6 presents the analysis of the tiles and the corresponding generated reports, performed by a certified pathologist. The pathologist had access to the entire WSI at the highest resolution to confirm the diagnosis.

For the given WSIs, we found a correspondence between the histological features identified by the pathologist in the selected tiles and the text of the generated report. For example, when the generated report mentions IDC, many top tiles show invasive cancer cells (see Fig 6, Example 1). Although the pathologist cannot confidently conclude IDC in the selected tiles without more context, they can confirm it by looking at the larger region shown on the right. In another example (see Fig 6, Example 3), while the generated report correctly labels the sample as cancerous and many tiles show cancerous tissue, the pathologist identified non-cancerous tissue in some tiles, such as benign colorectal glands. We hypothesize the presence of these glands among the most attended tiles helps the model to identify the site of tissue origin. In the generated report, the model predicts rectum as the origin site, which identifies the right organ (the large intestine) but not the correct location (ascending colon), which is histologically indistinguishable from the rectum.

Overall, the model correctly predicts cancer presence across all four examples and its sub-type in three examples, making a mistake in the example where the sub-type cannot be reliably inferred from the slide alone. It is correct when determining the cancer grade in three examples, making a mistake in one example between a well differentiated and moderately differentiated IDC. It sometimes struggles with the site of tissue origin when it cannot be inferred from the slide alone, consequently making a mistake in the sub-type assessment.

## 6    Discussion

Whole slide image analysis is a computationally expensive task that, due to the gigapixel-scale of WSIs, requires compression of hundreds of thousands of image embeddings containing a vast diversity of features into a single slide embedding. Moreover, typical downstream classification tasks, which include but are not limited to cancer detection, cancer sub-typing, and biomarker prediction, rely on noisy categorical labels which may not provide adequate or robust signals to learn and interpret. Furthermore, tasks such as the prediction of biomarkers or treatment response suffer from an additional challenge of label scarcity. Such factors motivate the design of PRISM, where natural-language clinical reports can be leveraged to learn more nuanced relationships between features and predictions that may also generalize across a wide range of tasks via fine-tuning slide encoder or training a linear classifier on the slide embeddings. Additionally, the ability to generate reports allows for training-free zero-shot classification and enhances the interpretability of the model.

To date, a growing number of works have attempted to learn vision-language pairings for pathology but stayed on the level of tiles. Scaling to the WSI level is inhibited by (1) the lack of high-level text descriptions in pre-training datasets and (2) computational bottlenecks of the selected architectures. To address (1), we use WSI-level clinical reports summarized by GPT-4 [1]. To address (2), we use a memory-efficient Perceiver architecture for tile aggre-
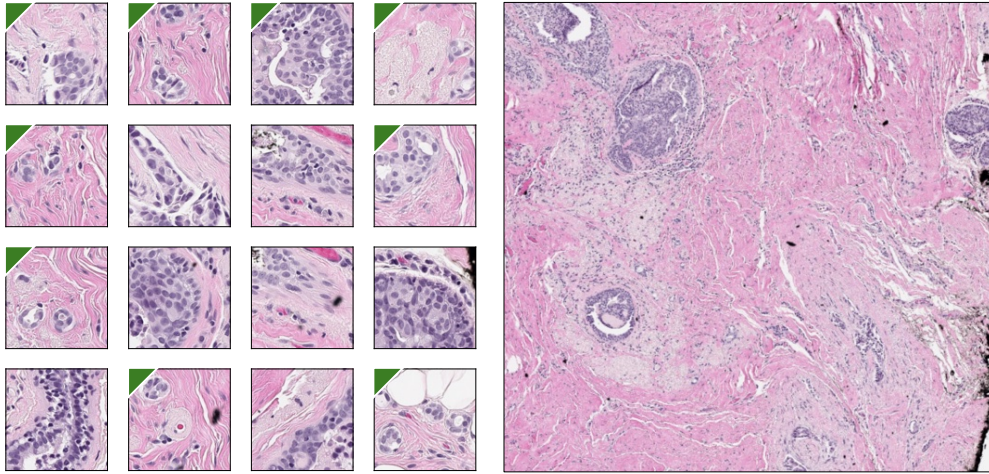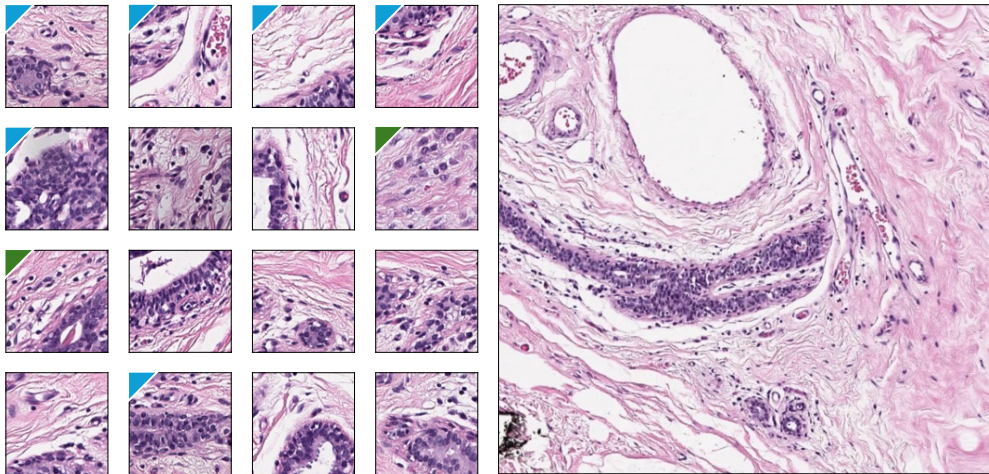
*Generated caption:*
**Invasive ductal carcinoma, well-differentiated, with positive estrogen and progesterone receptors.**

*Expert analysis:*
The original report confirms **IDC**. The green mark highlights tiles selected by a pathologist as suspicious for IDC. The report mentions that the tumor is **moderately differentiated** instead of well differentiated, with histologic grade 3, nuclear grade 2, and overall tumor grade 2. The report confirms **ER/PR positivity**; however, their expression cannot be confirmed histologically in the given tissue samples.



*Generated caption:*
**Diagnosis: Invasive lobular carcinoma in left breast, ER and PR positive.**

*Expert analysis:*
The original report confirms **ILC**. The report also mentions **LCIS** (lobular carcinoma in situ), a precursor to ILC. The green mark highlights tiles selected by a pathologist as suspicious for ILC, the blue mark – for LCIS. We hypothesise that the model merges all salient cancer features in a slide into a single slide embedding based on their hierarchy and importance to predict the correct report. In this example, the blue-marked tiles point to **lobular** subtype because of their specific cellular architecture, while the green-marked tiles indicate **invasive** carcinoma without having any distinct lobular features. Prediction of **left breast** is a guess since there's no possibility to know by looking at the slide. The report confirms **ER positivity** but indicates **PR negativity**.
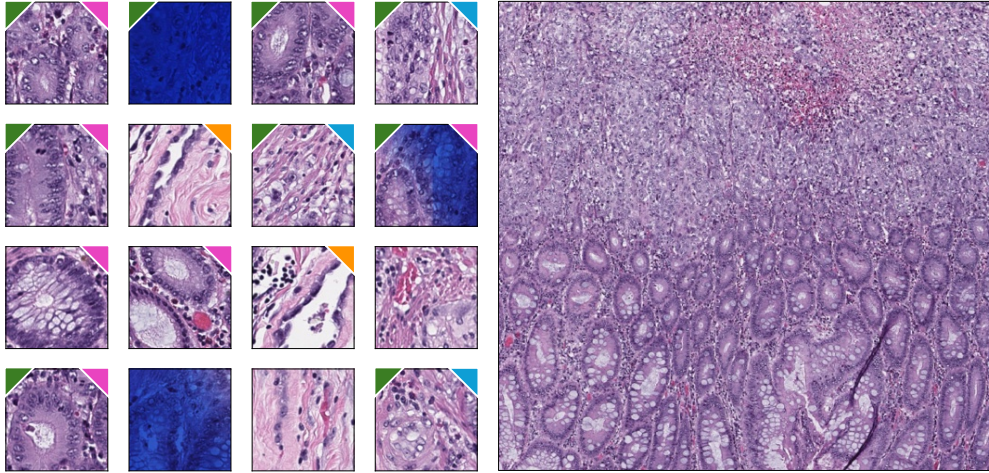
*Generated caption:*
**Adenocarcinoma of the rectum, poorly differentiated, with perineural and lymphovascular invasion.**

*Expert analysis:*
The original report confirms **carcinoma** in the colorectal tract; however, it states the precise location of the sample as ascending **colon rather than rectum**. Rectum and colon are almost indistinguishable histologically. Magenta-marked tiles show the presence of **colorectal glands**. Some glands are non-cancerous (first and second tiles on the third row). We hypothesise that the presence of tiles with non-cancerous glands indicates that the model learns features that help to identify tissue origin and other non-cancerous attributes present in the reports. Blue-marked tiles show tumor cells that lost glandular structure, pointing to **poorly differentiated** cancer. The original report confirms **lymphovascular** invasion. Orange-marked tiles show vessels, with the third-row tile showing what might be an artefactual tumor fragment inside the vessel, making it suspicious for lymphovascular invasion, although it can also be debris. The report indicates the absence of **perineural** invasion. Tiles with the dark blue tint come from a pen marked region.
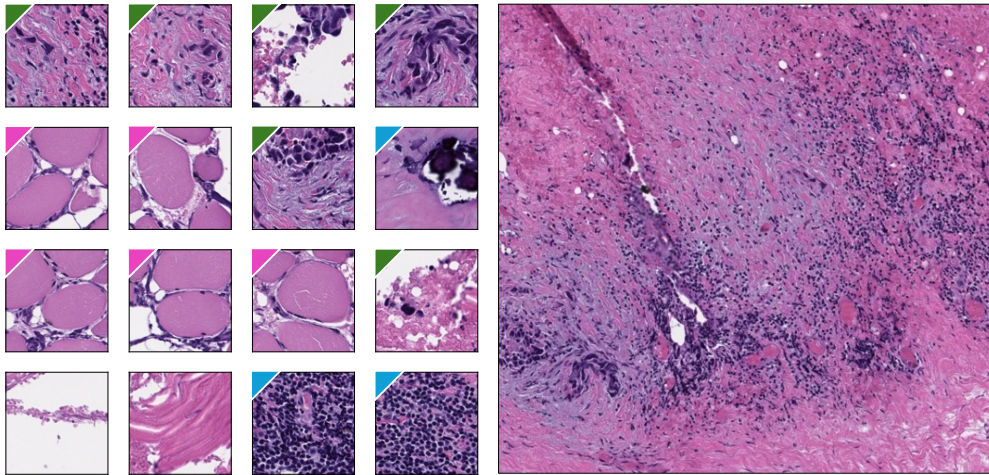


*Generated caption:*
**Diagnosis: Metastatic high-grade serous carcinoma in examined tissue.**

*Expert analysis:*
The slide has only few malignant cells which were successfully found by the model. The green-marked tiles show **high-grade cancer** cells. The sample shows fibroconnective tissue which can be found anywhere in the body. The magenta-marked tiles show benign **muscle tissue**. The type of high-grade cancer shown here does not arise in either tissues thus suggesting **metastasis**, which the model predicted, although the report doesn't confirm it. The blue-marked tiles show **inflammation and calcifications** – both are sometimes present in the reports, further supporting our claim that the model learned to identify variety of tissues mentioned in the reports, including non-cancerous, and to combine them into a single embedding to match to a correct report. The model predicts **serous** carcinoma suggesting the tissue comes from ovary; the original report instead states that this is a pleura biopsy (pleura is a membrane surrounding lungs), and that the cancer is most likely a non-small cell carcinoma which originates in lungs.

Figure 6: Analysis of generated reports and the most attended tiles. The top 16 tiles are shown, weighted by their attention weights in the last cross-attention layer of the Perceiver to the slide embedding latent. The slide embedding is the representation of the WSI which learns to summarize all relevant histological features in the slide. Below, a report generated for the corresponding WSI and its analysis in relation to the presented tiles. The analysis of the two in conjunction aims to demonstrate that the model learns to select all cancer and non-cancer features in the slide to produce the slide embedding, and at the same time to generated the report that is aligned with the slide embedding. On the right, a region of the WSI that contains some of the tiles is shown for context. The image of the region was not used by the model; it is provided only to contextualize model's prediction on the lower resolution level.

gation. PRISM is trained on diverse pathology reports and demonstrates the effectiveness of making predictions on a WSI level. Aligning to rich pathology reports boosts in-domain transfer performance as the slide embeddings are encouraged to be adequately diverse, instead of tightly packed about a categorical target, which is illustrated by linear probing and fine-tuning benchmark results. It also benefits transfer to the tasks not covered by the reports (Fig. 5), enabling the types of generalizability that foundation models aim to achieve.

With sufficient data scale, zero-shot and semi-supervised transfer can enable significant improvements in domains where data may be limited, however, it must be noted that a major limitation of this scaling is the availability and quality of clinical reports that cover large variations in the population as well as rare cases not seen during training. Recent research suggests that linear improvements for zero-shot performance requires an exponential increase in multi-modal data for the concepts of interest [41]. Furthermore, diagnostic reports, while more informative than simple categorical labels, may fail to capture all the necessary details desired in downstream applications and may provide information that is impossible to infer from imaging features derived from WSIs, e.g. origin site of the tissue sample, introducing noise in the training data. Furthermore, while PRISM appears to perform well on biomarkers that are not represented in the training distribution, evaluation on more tasks is needed to better understand its capabilities. Finally, it has been shown that multi-modal models can over-rely on linguistic priors present in the training data and mostly ignore the information provided by the vision module [25]. For example, the model may generate a report indicating the presence of a biomarker that is statistically correlated with the input tissue or disease state but not actually present.

There are several areas of this work that warrant future investigation. First, increasing the scale and diversity of the data could result in a better performing and more capable system. Specifically, including richer information about specimens from the diagnostic reports and beyond may allow the model to predict more fine-grained tissue characteristics and understand relationships to molecular and genomic data. This, coupled with further iteration on the GPT-4 [1] rewriting procedure, could improve downstream capabilities including zero-shot generalization. Another axis of investigation is the scaling of the constituent models. Using a larger and more capable language model should allow for more complex report generation and a better text embedding, while a larger slide embedding model may allow for improved visual features. Moreover, contrastive methods benefit from large and diverse batch sizes [36] that have not been explored in the current iteration.

Clinical practice has already benefited from computational pathology. Moving towards slide-level models, which more closely match clinical needs, promises to accelerate development even further.

# 7   Acknowledgements

# References

[1]  Josh Achiam et al. "Gpt-4 technical report". In: *arXiv preprint arXiv:2303.08774* (2023).

[2]  Shekoofeh Azizi et al. "Robust and data-efficient generalization of self-supervised machine learning for diagnostic imaging". In: *Nature Biomedical Engineering* 7.6 (2023), pp. 756–779.

[3]  Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. "Layer normalization". In: *arXiv preprint arXiv:1607.06450* (2016).

[4]  Geert Berx and Frans Van Roy. "The E-cadherin/catenin complex: an important gatekeeper in breast cancer tumorigenesis and malignant progression". In: *Breast Cancer Research* 3 (2001), pp. 289–293. DOI: 10.1186/bcr309.

[5]  Gustav Bredell et al. "Aggregation Model Hyperparameters Matter in Digital Pathology". In: *arXiv preprint arXiv:2311.17804* (2023).

[6]  Gabriele Campanella et al. "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images". In: *Nature medicine* 25.8 (2019), pp. 1301–1309.

[7]  Gabriele Campanella et al. "Computational Pathology at Health System Scale–Self-Supervised Foundation Models from Three Billion Images". In: *arXiv preprint arXiv:2310.07033* (2023).

[8]  Debyani Chakravarty et al. "OncoKB: a precision oncology knowledge base". In: *JCO precision oncology* 1 (2017), pp. 1–16.

[9]  Richard J Chen et al. "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, pp. 16144–16155.

[10]  Richard J Chen et al. "Towards a general-purpose foundation model for computational pathology". In: *Nature Medicine* (2024), pp. 1–13.

[11]  Diana T Cheng et al. "Memorial Sloan Kettering-Integrated Mutation Profiling of Actionable Cancer Targets (MSK-IMPACT): A Hybridization Capture-Based Next-Generation Sequencing Clinical Assay for Solid Tumor Molecular Oncology". In: *Journal of Molecular Diagnostics* 17.3 (2015), pp. 251–264. DOI: 10.1016/j.jmoldx.2014.12.006.

[12]  Ozan Ciga, Tony Xu, and Anne Louise Martel. "Self supervised contrastive learning for digital histopathology". In: *Machine Learning with Applications* 7 (2022), p. 100198.

[13]  Nicolas Coudray et al. "Classification and mutation prediction from non–small cell lung cancer histopathology images using deep learning". In: *Nature medicine* 24.10 (2018), pp. 1559–1567.

[14]  Jiayu Ding et al. "Longnet: Scaling transformers to 1,000,000,000 tokens". In: *arXiv preprint arXiv:2307.02486* (2023).

[15]  Jonas Dippel et al. "RudolfV: A Foundation Model by Pathologists for Pathologists". In: *arXiv preprint arXiv:2401.04079* (2024).

[16]  Alexey Dosovitskiy et al. "An image is worth 16x16 words: Transformers for image recognition at scale". In: *arXiv preprint arXiv:2010.11929* (2020).

[17]  Alexandre Filiot et al. "Scaling Self-Supervised Learning for Histopathology with Masked Image Modeling". In: *medRxiv* (2023), pp. 2023–07.

[18]  Jevgenij Gamper and Nasir Rajpoot. "Multiple instance captioning: Learning representations from histopathology textbooks and articles". In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2021, pp. 16549–16559.

[19]  Zhi Huang et al. "A visual–language foundation model for pathology image analysis using medical twitter". In: *Nature medicine* 29.9 (2023), pp. 2307–2316.

[20]  Maximilian Ilse, Jakub Tomczak, and Max Welling. "Attention-based deep multiple instance learning". In: *International conference on machine learning*. PMLR. 2018, pp. 2127–2136.

[21]  Andrew Jaegle et al. "Perceiver: General perception with iterative attention". In: *International conference on machine learning*. PMLR. 2021, pp. 4651–4664.

[22]  Gregory P. Kalemkerian et al. "Molecular Testing Guideline for the Selection of Patients With Lung Cancer for Treatment With Targeted Tyrosine Kinase Inhibitors: American Society of Clinical Oncology Endorsement of the College of American Pathologists/International Association for the Study of Lung Cancer/Association for Molecular Pathology Clinical Practice Guideline Update". In: *Journal of Clinical Oncology* 36(9) (2018). DOI: 10.1200/JCO.2017.76.7293.

[23]  Mingu Kang et al. "Benchmarking self-supervised learning on diverse pathology datasets". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2023, pp. 3344–3354.

[24]  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*. 2021, pp. 14318–14328.

[25]  Fuxiao Liu et al. "Aligning Large Multi-Modal Model with Robust Instruction Tuning". In: *arXiv preprint arXiv:2306.14565* (2023).

[26]  Ming Y Lu et al. "Data-efficient and weakly supervised computational pathology on whole-slide images". In: *Nature biomedical engineering* 5.6 (2021), pp. 555–570.

[27]  Ming Y Lu et al. "Towards a visual-language foundation model for computational pathology". In: *arXiv preprint arXiv:2307.12914* (2023).

[28]  Ming Y Lu et al. "Visual language pretrained multiple instance zero-shot transfer for histopathology images". In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2023, pp. 19764–19775.

[29]  Renqian Luo et al. "BioGPT: generative pre-trained transformer for biomedical text generation and mining". In: *Briefings in bioinformatics* 23.6 (2022), bbac409.

[30]  Maxime Oquab et al. "DINOv2: Learning robust visual features without supervision". In: *arXiv preprint arXiv:2304.07193* (2023).

[31]  Sudhir Perincheri et al. "An independent assessment of an artificial intelligence system for prostate cancer detection shows strong diagnostic accuracy". In: *Modern Pathology* 34.8 (2021), pp. 1588–1595.

[32]  Wenkang Qin et al. "What a whole slide image can tell? subtype-guided masked transformer for pathological image captioning". In: *arXiv preprint arXiv:2310.20607* (2023).

[33]  Patricia Raciti et al. "Clinical validation of artificial intelligence–augmented pathology diagnosis demonstrates significant gains in diagnostic accuracy in prostate cancer detection". In: *Archives of Pathology & Laboratory Medicine* 147.10 (2023), pp. 1178–1185.

[34]  Patricia Raciti et al. "Novel artificial intelligence system increases the detection of prostate cancer in whole slide images of core needle biopsies". In: *Modern Pathology* 33.10 (2020), pp. 2058–2066.

[35]  Alec Radford et al. "Language models are unsupervised multitask learners". In: *OpenAI blog* 1.8 (2019), p. 9.

[36]  Alec Radford et al. "Learning transferable visual models from natural language supervision". In: *International conference on machine learning*. PMLR. 2021, pp. 8748–8763.

[37]  Zhuchen Shao et al. "Transmil: Transformer based correlated multiple instance learning for whole slide image classification". In: *Advances in neural information processing systems* 34 (2021), pp. 2136–2147.

[38]  Noam Shazeer. "Glu variants improve transformer". In: *arXiv preprint arXiv:2002.05202* (2020).

[39]  Leonard M da Silva et al. "Independent real-world application of a clinical-grade automated prostate cancer detection system". In: *The Journal of pathology* 254.2 (2021), pp. 147–158.

[40]  Masayuki Tsuneki and Fahdi Kanavati. "Inference of captions from histopathological patches". In: *International Conference on Medical Imaging with Deep Learning*. PMLR. 2022, pp. 1235–1250.

[41]  Vishaal Udandarao et al. "No" zero-shot" without exponential data: Pretraining concept frequency determines multimodal model performance". In: *arXiv preprint arXiv:2404.04125* (2024).

[42]  Ashish Vaswani et al. "Attention is all you need". In: *Advances in neural information processing systems* 30 (2017).

[43]  Eugene Vorontsov et al. *Virchow: A Million-Slide Digital Pathology Foundation Model*. 2024. arXiv: 2309.07778 [eess.IV].

[44]  S. J. Wagner et al. "Transformer-based biomarker prediction from colorectal cancer histology: A large-scale multicentric study". In: *Cancer Cell* 41.9 (Sept. 2023), pp. 1650–1661.

[45]  Wenhui Wang et al. "When an Image is Worth 1,024 x 1,024 Words: A Case Study in Computational Pathology". In: *arXiv preprint arXiv:2312.03558* (2023).

[46]  Xiyue Wang et al. "Transformer-based unsupervised contrastive learning for histopathological image classification". In: *Medical image analysis* 81 (2022), p. 102559.

[47]  Jiahui Yu et al. "Coca: Contrastive captioners are image-text foundation models". In: *arXiv preprint arXiv:2205.01917* (2022).

[48]  Renyu Zhang et al. "Evaluating and interpreting caption prediction for histopathology images". In: *Machine Learning for Healthcare Conference*. PMLR. 2020, pp. 418–435.

# A  Appendix

## A.1  Zero-shot prompts

| Task | Classes | Class 1 Prompts | Class 2 Prompts |
|---|---|---|---|
| TCGA BRCA | ILC / IDC | lobular carcinoma, invasive | ductal carcinoma, invasive |
| TCGA NSCLC | LUSC / LUAD | lung squamous cell carcinoma | lung adenocarcinoma |
| MSK DCIS/IDC | DCIS / IDC | ductal carcinoma in situ | invasive ductal carcinoma |
| Pan-cancer detection | negative / positive | benign<br>inflammation<br>normal<br>infectious<br>cyst<br>polyp<br>unremarkable | cancer<br>carcinoma<br>adenocarcinoma<br>malignant<br>metastatic<br>invasive<br>sarcoma |

Table A.1.1: Prompts used in zero-shot evaluations. All tasks are binary classification. The pan-cancer detection dataset uses multiple prompts per class; they are ensembled by taking a sum of probabilities for each prompt in the class to produce the final class score.

## A.2   Hyperparameters

| Hyperparameter | Value |
| --- | --- |
| Perceiver | |
| Layers | 8 |
| Number of latents | $512 + 1$ |
| Latent dimensions | 1280 |
| Tile embedding dimensions | 2560 |
| Cross-attention | |
| KQV dimensions | 1280 |
| Heads | 1 |
| MLP activation | GEGLU |
| MLP inner dimensions | 1280 |
| Layers sharing weights | 2-8 |
| Latent Transformer | |
| Layers | 6 |
| KQV dimensions | 1280 |
| Heads | 8 |
| MLP activation | GEGLU |
| MLP inner dimensions | 1280 |
| Layers sharing weights | 1-8 |
| BioGPT | |
| Token embedding dimension | 768 |
| Uni-modal layers | $1 - 12$ |
| Multi-modal layers | $13 - 24$ |
| Contrastive Projector | Linear |
| Contrastive embedding dimension | 5120 |
| Contrastive loss weight | 1.0 |
| Report loss weight | 2.0 |
| AdamW $\beta$ | $(0.9, 0.999)$ |
| Batch size | 64 |
| Warm up steps | 2000 |
| Total steps | 24000 |
| Learning rate schedule | Cosine |
| Learning rate (start) | $2 \cdot 10^{-4}$ |
| Weight decay | $1 \cdot 10^{-6}$ |
| Gradient accumulation iterations | 4 |
| Gradient clipping norm | 3.0 |
| Automatic mixed precision | FP16 |

Table A.2.1: PRISM hyperparameters trained on 16 NVIDIA V100 32GB GPU's.

## A.3  Pancancer AUCs

| Models | Perceiver (from scratch) | PRISM (zero-shot) | PRISM (linear probe) | PRISM (fine-tuned) |
|---|---|---|---|---|
| Common origins | | | | |
| Breast | 0.980 | 0.901 | **0.983** | 0.974 |
| Prostate | 0.964 | **0.975** | 0.929 | 0.944 |
| Lung | 0.970 | 0.958 | 0.968 | **0.975** |
| Colon | **0.998** | 0.991 | 0.996 | 0.996 |
| Skin | 0.917 | 0.860 | 0.911 | **0.939** |
| Bladder | 0.933 | 0.943 | **0.958** | 0.950 |
| Uterus | 0.957 | 0.928 | 0.934 | **0.963** |
| Pancreas | **0.985** | 0.963 | 0.981 | 0.976 |
| Head and Neck | 0.987 | 0.966 | 0.985 | **0.996** |
| Rare origins | | | | |
| Liver | **0.985** | 0.934 | 0.984 | 0.977 |
| Stomach | 0.995 | 0.995 | **0.997** | 0.995 |
| Brain | 0.947 | 0.809 | **0.961** | 0.955 |
| Ovary | 0.962 | 0.960 | 0.965 | **0.969** |
| Cervix | 0.853 | 0.870 | 0.852 | **0.888** |
| Testis | 0.952 | 0.904 | 0.901 | **0.954** |
| Bone | 0.800 | 0.683 | 0.818 | **0.856** |

Table A.3.1: AUC per site of tumor origin in MSKCC cancer detection task.

### A.4    Sub-typing tasks

**TCGA BRCA** was curated by selecting only diagnostic slides and then selecting slides whose diagnosis was either IDC or ILC. Furthermore, we discarded slides that were not readable due to missing metadata or file corruption. This resulted in a dataset with 1002 slides with 798 (80%) IDC and 204 (20%) ILC.

**TCGA NSCLC** was curated by selecting diagnostic slides from the TCGA LUAD and LUSC projects. As in TCGA BRCA, we discarded slides that were not readable due to missing metadata or file corruption. This resulted in a dataset with 1043 slides with 531 (51%) LUAD and 512 (49%) LUSC.

**MSKCC DCIS/IDC** was curated by selecting from a subset of internal breast specimens that were diagnosed with either DCIS or IDC. This resulted in a dataset with 896 specimens with 346 (39%) DCIS and 550 (61%) IDC. The dataset includes rare invasive sub-types.

### A.5    Biomarker prediction tasks

The training, validation, and testing distribution is shown in Supplementary Tab. A.5.1 for each biomarker dataset. The biomarkers are described below.

| Biomarker | Subset | Cases | Slides | Positive |
|---|---|---|---|---|
| Breast-CDH1 | train | 648 | 673 | 0.13 |
|  | tune | 215 | 220 | 0.13 |
|  | test | 214 | 228 | 0.13 |
| Colon-MSI | train | 4609 | 11027 | 0.10 |
|  | tune | 481 | 1417 | 0.14 |
|  | test | 482 | 1446 | 0.14 |
| Bladder-FGFR | train | 520 | 542 | 0.24 |
|  | tune | 259 | 275 | 0.29 |
|  | test | 259 | 270 | 0.25 |
| Endometrial-PTEN | train | 983 | 1038 | 0.48 |
|  | tune | 164 | 170 | 0.43 |
|  | test | 164 | 178 | 0.41 |
| Lung-EGFR | train | 2186 | 2858 | 0.28 |
|  | tune | 356 | 457 | 0.29 |
|  | test | 358 | 457 | 0.28 |
| Prostate-AR | train | 1051 | 1461 | 0.18 |
|  | tune | 348 | 480 | 0.20 |
|  | test | 347 | 480 | 0.16 |
| Ovarian-FGA | train | 679 | 791 | 0.91 |
|  | tune | 115 | 134 | 0.90 |
|  | test | 111 | 126 | 0.88 |
| Gastric-Her2 | train | 968 | 968 | 0.19 |
|  | tune | 170 | 170 | 0.23 |
|  | test | 161 | 161 | 0.17 |
| Skin-BRAF | train | 782 | 868 | 0.25 |
|  | tune | 131 | 137 | 0.21 |
|  | test | 131 | 138 | 0.13 |

Table A.5.1: Statistics of the case-level biomarker target datasets, including the number of cases ("Cases"), the number of slides ("Slides"), and the proportion of positive labels ("Positive").

**Breast-CDH1** was assessed via the presence of inactivating mutations associated with loss-of-heterozygosity (LOH) or a second somatic loss-of-function mutation as determined by MSK-IMPACT sequencing test results. Bi-allelic loss of Cadherin 1 (CDH1) is strongly correlated with lobular breast cancer, and a distinct histologic phenotype and biologic behavior [4]. Samples with other types of variants, i.e. mono-allelic mutations, were excluded.

**Colon-MSI** was assessed using both IHC and MSK-IMPACT sequencing for deficient mismatch repair (dMMR) and high-frequency MSI (MSI-H) detection, prioritizing IHC results when both test outcomes are available. MSI-H is present in approximately 15% of colorectal cancers (CRCs), often linked to germline mutations that elevate hereditary cancer risk. Consequently, routine microsatellite instability (MSI) or IHC-based dMMR screening is recommended for all primary colorectal carcinoma samples.

**Bladder-FGFR** was assessed via the presence of FGFR3 p.S249C, p.R248C, p.Y373C, p.G370C mutations, FGFR3-TACC3 fusions and FGFR2 p.N549H, pN549K, p.N549S, p.N549T mutations using MSK-IMPACT data. The fibroblast growth factor receptor (FGFR) alterations screening in bladder carcinoma allows the identification of patients targetable by FGFR inhibitors.

**Endometrial-PTEN** was assessed via the oncogenic status of phosphatase and tensin homolog (PTEN) mutation using MSK-IMPACT data and OncoKB annotation [8]. The variants associated with any oncogenic effect (including predicted/likely oncogenic) were defined as positive label for PTEN mutations, and variants with unknown oncogenic status were excluded. PTEN is the most frequently mutated tumor suppressor gene in endometrial cancer and the presence of PTEN mutation has been shown to be significantly associated with poorer prognosis in survival and disease recurrence.

**Lung-EGFR** was assessed via the oncogenic status of epidermal growth factor receptor (EGFR) mutation using MSK-IMPACT data and OncoKB annotation [8]. EGFR mutations with any oncogenic effect (including predicted/likely oncogenic) were defined as positive label, and EGFR mutation with unknown oncogenic status were excluded. The EGFR oncogenic mutation screening in NSCLC is essential to determine eligibility for targeted therapies in late stage NSCLC [22].

**Prostate-AR** was assessed based on whether the fold change of copy number in androgen receptor (AR) was greater than 2 as measured by MSK-IMPACT data. AR amplification/over-expression is found in 30%-50% of castration resistant prostate cancers (CRPCs), and is associated with the resistance to androgen deprivation therapy (ADT).

**Ovarian-FGA** was assessed based on whether the fraction of genome altered was $\geq$ 30% as measured by MSK-IMPACT data. This cutoff has been established to enrich for TP53 mutations, a factor in the characterization of High-grade serous ovarian cancer (HGSOC) and previously reported to correlate to increased FGA.

**Gastric-HER2** amplification was assessed based on human epidermal growth factor receptor 2 (HER2) IHC 2+/FISH+ or IHC 3+. Approximate 20% of gastric cancer patients are found to correlate with HER2 overexpression / high-level amplification and would likely benefit from treatment with an anti-HER2 antibody therapy.

**Skin-BRAF** was assessed based on oncogenic mutation status and the presence of V600E variant using MSK-IMPACT data and OncoKB annotation [8]. B-Raf Proto-Oncogene (BRAF) is one of the most frequently mutated genes in melanoma, and V600E mutation is the most common variant, which leads to constitutive activation of the BRAF/MEK/ERK (MAPK) signalling pathway. Targeted therapy with BRAF inhibitors showed better survival outcome in patients with BRAF V600-mutated melanoma.

## A.6  Biomarker results

| Biomarker | Not pre-trained | | Pre-trained | |
|---|---|---|---|---|
| | AUC | f | AUC | f |
| Breast-CDH1 | 0.938 (0.009) | 1.0 | **0.984** (0.004) | 0.8 |
| Colon-MSI | 0.966 (0.012) | 1.0 | **0.969** (0.015) | 0.5 |
| Bladder-FGFR | 0.847 (0.015) | 1.0 | **0.878** (0.005) | 0.9 |
| Endometrial-PTEN | 0.833 (0.023) | 0.6 | **0.847** (0.005) | 1.0 |
| Lung-EGFR | 0.770 (0.006) | 0.6 | **0.789** (0.015) | 0.7 |
| Prostate-AR | 0.800 (0.010) | 0.7 | **0.820** (0.007) | 1.0 |
| Ovarian-FGA | **0.825** (0.070) | 0.9 | 0.788 (0.025) | 1.0 |
| Gastric-Her2 | **0.812** (0.029) | 1.0 | 0.793 (0.032) | 0.7 |
| Skin-BRAF | **0.736** (0.020) | 0.9 | 0.706 (0.012) | 0.8 |

Table A.6.1: The top area under (the receiver operating characteristic) curve (AUC) achieved for various biomarker tasks when training the tile aggregator with and without pre-training with PRISM, shown as mean (standard deviation) across 3 experimental runs. The fraction f of the training set used to produce this result is also shown.

## A.7 GPT-4 prompt for clinical report rewriting

| Parameter | Value |
|---|---|
| Model | gpt-4-1106-preview |
| Temperature | 0.7 |
| Top P | 0.95 |

Table A.7.1: API parameters for GPT-4 report rewriting

System prompt:

You are a pathology lab assistant. You are given an unstructured pathology report describing a tissue sample. Follow these instructions carefully and to the letter:
- Extract a detailed summary of the diagnosis and the examined tissue from the report in a sentence under 20 words.
- Mention tissue type if it is mentioned in the report.
- If the report includes results of immunohistochemical studies or molecular tests, include a short summary of the most important positive results. Keep the overall sentence length below 20 words.
- Do not mention any negative test results or absence of something.
- Do not mention any numbers.
- Do not mention cm or mm measurements.
- Minimize prose, be succinct, use as few words as possible.
- Following the instructions above, rewrite the report 5 times, for a total of 5 sentences, each with a summary written in a slightly different way. Make sure all summaries are consistent with each other.
- Output the resulting 5 sentences as a list in JSON format. Don't output anything else.

Examples of rewritten reports:

Metastatic urothelial carcinoma in one lymph node with extranodal extension.
One lymph node shows metastatic urothelial carcinoma with extranodal spread.
Diagnosis: Urothelial carcinoma metastasis in lymph node, extranodal extension present.
Lymph node involved by metastatic urothelial carcinoma; extranodal extension observed.
Urothelial carcinoma with metastasis to a lymph node and extranodal extension identified.

Diagnosed cutaneous neuroendocrine Merkel cell carcinoma with lymphovascular invasion.
Merkel cell carcinoma of the skin with infiltrative growth and lymphovascular invasion found.
Skin Merkel cell carcinoma confirmed, showing infiltrative pattern and lymphovascular invasion.
Cutaneous Merkel cell carcinoma with lymphovascular invasion, infiltrative pattern diagnosed.
Pathology reveals skin Merkel cell carcinoma, lymphovascular invasion present.

Acral skin biopsy; reactive changes, no melanoma, PRAME positive, SOX10 positive.
Examined acral skin; found reactive changes, PRAME and SOX10 immunopositivity, no melanoma.
Acral skin showing reactive changes; melanoma absent; positive for PRAME and SOX10.
Reactive changes in acral skin tissue; PRAME and SOX10 positive; melanoma not detected.
Acral skin with reactive changes; tests positive for PRAME and SOX10; no evidence of melanoma.

Benign breast parenchyma diagnosed.
Diagnosis: Benign condition of breast tissue.
Breast tissue exam reveals benign parenchyma.
Benign parenchymal findings in breast tissue.
Examined breast tissue is benign.

## A.8  Examples of generated reports

| | |
|---|---|
| Relevant report sections | Left ovary and fallopian tube: - Left ovary with follicular cysts and focal dystrophic calcification. - Left fallopian tube with tubal cysts. - No tumor seen. |
| Generated report | Benign ovarian and fallopian tube tissue with calcifications. |
| Expert comment | Correct diagnosis and site |
| Relevant report sections | Core biopsy right distal femur: - Enchondroma |
| Generated report | Diagnosis: Enchondroma in bone tissue. |
| Expert comment | Correct diagnosis and site |
| Relevant report sections | Liver, core biopsy: - Hepatocellular carcinoma, moderately differentiated, see note Note: Submitted reticulin stain supports the diagnosis. Submitted immunohistochemical stain for glypican 3 is focally positive in tumor cells. |
| Generated report | Liver biopsy: Well-differentiated hepatocellular carcinoma identified. |
| Expert comment | Correct diagnosis and site |
| Relevant report sections | Prostate, right medial base, mid, apex: a. Prostate biopsy location: Right Base medial Benign Processes: Benign prostatic tissue, Acute inflammation, Chronic inflammation b. Prostate biopsy location: Right Mid medial Benign Processes: Benign prostatic tissue, Acute inflammation, Chronic inflammation c. Prostate biopsy location: Right Apex medial Benign Processes: Benign fibromuscular stroma; no prostatic glands are identified |
| Generated report | Diagnosis: Benign prostatic tissue with acute and chronic inflammation. |
| Expert comment | Correct diagnosis and site |
| Relevant report sections | Esophagus and proximal stomach (fs): Tumor Type: Invasive adenocarcinoma Treatment effect (Status Post neoadjuvant therapy): Treatment response involves 80% of the tumor Histologic Grade: Moderately differentiated Tumor Location: Lower thoracic portion esophagus ($<$ 5cm above EGJ) extends 1% into the stomach Siewert Type: Type I ($>$ 75% above GEJ) Tumor Size: Length is 1.5 cm, Width is 1.4 cm, Maximal thickness is 0.5 cm Depth of Tumor Invasion: Tumor invades the muscularis propria, Foci of acellular mucin likely representing regressed tumor are present in the subserosal tissue. Vascular Invasion: Not identified Perineural Invasion: Not identified Tumor Multicentricity: Not identified Adjacent Mucosa: Unremarkable squamous mucosa, Gastric mucosa shows focal intestinal metaplasia, negative for dysplasia Proximal Margin: Free of carcinoma or dysplasia, with 4.5 cm clearance Distal Margin: Free of carcinoma or dysplasia Deep Margin: Free of invasive carcinoma with a clearence of 0.5 cm Lymph Nodes: Number of nodes examined: 17, Number of metastatic nodes: 0 Tumor Stage (AJCC 7th Edition): pT2 (Invades the muscularis propria) Lymph Node Stage (AJCC 7th Edition): pN0 (No regional node metastasis) |
| Generated report | Diagnosis: Metastatic adenocarcinoma in a lymph node. |
| Expert comment | Diagnosis is correct. Discordance: no lymph nodes are present in the specimen. |

| | |
|---|---|
| Relevant report sections | Left partial glossectomy, floor of mouth resection: Procedure: Resection, Other: Partial glossectomy Histologic Type: Squamous cell carcinoma, keratinizing Histologic Grade: G1: Well differentiated Tumor Location: Floor of mouth Tumor Laterality: Left Tumor Focality: Single focus Tumor Size: Greatest dimension is 2.9 cm (per report) Tumor Thickness: Tumor thickness: 8 mm, Depth of invasion: 4.8 mm Microscopic Tumor Extension: Invades skeletal muscle In situ carcinoma: Not Identified Specimen Margins: Margins involved by invasive tumor (specify margin): Deep margin Tumor Bed Margins: Not applicable Lymph-Vascular Invasion: Not identified Perineural Invasion: Not identified Lymph Nodes: No lymph nodes submitted or found Extranodal Extension (ENE): Not applicable Additional Pathologic Findings: None identified |
| Generated report | Diagnosis: Well-differentiated squamous cell carcinoma with perineural invasion in the dermis. |
| Expert comment | Diagnosis and site are correct. Perineural invasion is in caption but not identified in the ground truth. |
| Relevant report sections | Rectum, 10 cm, biopsy: - Invasive adenocarcinoma, moderately differentiated, see note Note: Immunohistochemical stains for MMR proteins will be reported in an addendum. |
| Generated report | Moderately differentiated invasive adenocarcinoma in rectum; MLH1, MSH2, MSH6, PMS2 proteins retained. |
| Expert comment | Diagnosis and site are correct. MMR biomarkers are not reported in original report provided but it would be of interest to know if all were retained (common). |
| Relevant report sections | Lung, superior segment LLL; segmentectomy: Tumor Location: Lung, left lower lobe, Superior segment Procedure: Segmentectomy Tumor Type: Adenocarcinoma Predominant pattern: Micropapillary 60% Other patterns: Acinar 30%, Solid 10%, The acinar component includes 10% cribriform pattern Tumor Focality: Unifocal Tumor Size: Total tumor size 1.8 cm Histologic Grade: Poorly differentiated (G3) Lymphovascular Invasion: Present (extensive) Perineural Invasion: Not identified Spread Through Air Spaces: Present Pleural Invasion: The tumor invades into the visceral pleura beyond the elastic layer (PL1) Bronchial Margin: For final margin status, please see part 3 Vascular Margin: For final margin status, please see part 3 Staple Line: For final margin status, please see part 3 Distance from margin: Cannot be assessed, See part 3 Peribronchial or intrapulmonary Lymph Nodes: Number with metastasis: 1 |
| Generated report | Adenocarcinoma with mucinous features in lung; CK7, CK20, CDX2 positive. |
| Expert comment | Partially correct: site and tumor are correct, however no mention of any mucinous component. |
| Relevant report sections | Skin and soft tissue, right leg; wide excision: - Metastatic melanoma. |
| Generated report | Diagnosis: Metastatic melanoma in lymph node with extranodal extension. |
| Expert comment | Morphological diagnosis is correct however the report does not mention any lymph nodes. |
| Relevant report sections | Left hemiglossectomy: Procedure: Glossectomy (hemiglossectomy) Histologic Type: Squamous cell carcinoma, keratinizing Histologic Grade: G2: Moderately differentiated Tumor Location: Tongue Tumor Laterality: Left Tumor Focality: Single focus Tumor Size: Greatest dimension is 5.0 cm per outside report Tumor Thickness: Depth of invasion 18 mm, per outside report Microscopic Tumor Extension: Invades skeletal muscle In situ carcinoma: Not Identified in the submitted material Specimen Margins: Cannot be assessed, Location of closest margin, per orientation: cannot be assessed in the submitted material Tumor Bed Margins: Not applicable Lymph-Vascular Invasion: Not identified in the submitted material Perineural Invasion: Not identified in the submitted material Lymph Nodes: No lymph nodes submitted or found in the submitted material |
| Generated report | Lymph node with metastatic keratinizing squamous carcinoma; extranodal extension present. |
| Expert comment | Morphological diagnosis is correct however the report mentions no lymph nodes submitted. |

| | |
|---|---|
| Relevant report sections | Right mandible, biopsy: Small fragment of benign skeletal muscle, fibrous tissue, bone, and dentin. No epithelium-lined cyst seen |
| Generated report | Bone and muscle tissue show benign pathology. |
| Expert comment | Correct, all benign findings. |
| Relevant report sections | Breast, left lateral calcifications; core biopsy: In Situ Carcinoma: Markedly atypical ductal hyperplasia focally reaching the level of ductal carcinoma in situ (DCIS) DCIS, Architectural Pattern: Cribriform Nuclear Grade: Low Non-neoplastic Breast Tissue: A few foci suspicious for atypical lobular hyperplasia (ALH) Calcifications: Rare, in DCIS and in stroma |
| Generated report | Diagnosis: Ductal carcinoma in situ (DCIS) in breast tissue with cribriform pattern. |
| Expert comment | Correct diagnosis of DCIS. |
| Relevant report sections | Skin, left chest: Procedure: Excisional biopsy Tumor Type: Melanoma Histologic Type: Superficial spreading type Breslow thickness: 0.95 mm Ulceration: Not identified Mitotic index: 1 /mm2 Clark level: IV Growth phase: Vertical Infiltrating lymphocytes: Non-brisk Regression: identified, involving less than 75% of lesion Lymphovascular Invasion: Not identified Perineural Invasion: Not identified Microscopic satellite: Not identified Associated nevus: Not identified Surgical Margins: Not involved, The distance of melanoma in situ from the nearest side margin is:1mm. Tumor Staging: pT1b (The melanoma is less than or equal to 1 mm thick. It is ulcerated and/or the mitotic rate is equal to or greater than 1/mm2) |
| Generated report | Diagnosis: Superficial spreading melanoma, Clark level IV, with ulceration. |
| Expert comment | Correct diagnosis, minor discordance in one part where the GT says no ulceration but it describes ulceration. Otherwise diagnosis and grade are correct. |
| Relevant report sections | Right Fallopian Tube and Ovary, Excision/Resection: Serous cystadenoma (4.5 cm) Fallopian tube within normal limits |
| Generated report | Diagnosis: Serous cystadenoma of the ovary, fallopian tube unremarkable. |
| Expert comment | Correct diagnosis and site. |
| Relevant report sections | Uterus, cervix, bilateral tubes and ovaries; total hysterectomy and bilateral salpingo-oophorectomy: - Atrophic endometrium with focal glandular crowding - Benign endometrial polyp - Adenomyosis - Benign ovaries with simple cysts - Benign fallopian tubes with paratubal cysts |
| Generated report | Diagnosis: Endometrioid adenocarcinoma, FIGO grade 1, with myometrial invasion. |
| Expert comment | Incorrect diagnosis, major discordance. |
| Relevant report sections | Ileal mucosa with no significant pathologic abnormalities. |
| Generated report | Small lymphocytic lymphoma / chronic lymphocytic leukemia identified in lymphoid tissue. |
| Expert comment | Incorrect diagnosis, major discordance. |
| Relevant report sections | Distal pancreas for tumor bank Specimen(s) Received: Pancreas, tail Procedure: Distal pancreatectomy Histologic Type: Residual adenocarcinoma exhibiting changes consistent with treatment effect (status post neoadjuvant therapy) Histologic Grade: Poorly differentiated Tumor Location: Tail of the pancreas Overall Tumor Size: Greatest dimension: 5.0 cm, Per outside report Microscopic Tumor Extension: Tumor invades peripancreatic soft tissues Lymph-Vascular invasion: Present Perineural Invasion: Present Margins: Uninvolved by invasive carcinoma/pre-invasive neoplasm per report Anterior Peritoneal Surface: Uninvolved by invasive carcinoma Posterior Peritoneal Surface: Free of invasive carcinoma/pre-invasive neoplasm Histologic Evidence of Treatment Effect: Present, Involves approximately 60Adjacent Pancreas: Exhibits atrophy, Exhibits fibrosis Lymph Nodes: Number of lymph nodes examined in this specimen: 16, Number of lymph nodes with metastasis in this specimen: 2 TNM descriptors: y (posttreatment) Primary Tumor (pT) (AJCC 8th Edition): pT3: Tumor >4 cm in greatest dimension Regional Lymph Nodes (pN) (AJCC 8th Edition): pN1: Metastasis in one to three regional lymph nodes |
| Generated report | Metastatic high-grade serous carcinoma found in lymph node tissue. |
| Expert comment | Metastatic carcinoma is correct however saying "serous carcinoma" is specific to mullerian origin (e.g., ovary) and thus incorrect. |

24

| | |
|---|---|
| Relevant report sections | Gastric cancer. History of met gastric signet ring cell cancer. Pylorus, biopsies: Invasive adenocarcinoma, poorly differentiated. |
| Generated report | Diagnosis: Poorly differentiated adenocarcinoma with signet ring features in stomach. |
| Expert comment | Correct diagnosis and site. |
| Relevant report sections | Benign cystic tissue consistent with clinical impression of hydrocele. |
| Generated report | Examined tissue: Benign ovarian cyst. |
| Expert comment | Benign diagnosis is correct, incorrect site of origin. |
| Relevant report sections | Breast, left; lumpectomy: Invasive Carcinoma: Mucinous Size: Microscopic measurement 1.5 cm Histologic Grade: III/III: minimal or no tubule formation ($< 10\%$ of tumor) [score 3] Nuclear Grade: II/III (moderate variation in size and shape) [score 2] Mitotic Count: $< 8$ mitoses per 10 high power fields [score 1] Overall Tumor Grade: Combined score 6: grade II/III (moderately differentiated) Ductal Carcinoma in Situ (DCIS): Identified, Extends into lobules DCIS Architectural Pattern: Solid, Cribriform DCIS Nuclear Grade: Intermediate Calcification: In DCIS, In benign epithelium Lymphovascular Invasion: Not identified Surgical Margins (Invasive Carcinoma): For final margin status see separately submitted margins Surgical Margins (DCIS): For final margin status see separately submitted margins Non-neoplastic Breast Tissue: Biopsy site changes |
| Generated report | Breast tissue shows mucinous carcinoma, Grade II-III, with mucin production. |
| Expert comment | Correct diagnosis and site. |

Table A.8.4: We compare relevant report sections of the original clinical reports with the reports generated by PRISM. The reports are on specimen level, many spanning multiple WSIs. The specimens come from the evaluation cancer detection dataset (Section 4.3.

### A.9  Perceiver implementation details

The cross-attention module takes tile embeddings as key-value tokens and resamples them with latent tokens as queries using single-head attention [42]. Layer normalization [3] is applied to latents but not to tile embeddings in order to save memory. The latent self-attention transformer is a 6-layer transformer encoder that closely follows Vaswani et al. [42], with 8 heads per attention module. The MLP networks in the cross-attention modules and latent transformers consist of a GEGLU layer [38] followed by a linear layer, with 1280 inner dimensions. We don't apply position encoding to tile embeddings.

```python
def perceiver(
    latents: Tensor,  # learned latent embeddings (513, 1280)
    embeddings: Tensor,  # Virchow tile embeddings (N, 2560)
    xattn_layer_0: nn.Module,  # cross-attention module with MLP
    xattn_layer_1: nn.Module,  # cross-attention module with MLP
    latent_transformer: nn.Module,  # 6-layer transformer
    perceiver_depth: int = 8,
) -> tuple[Tensor, Tensor]:  # return slide embedding and latent features
    for i in range(perceiver_depth):
        # Cross-attention
        if i == 0:
            latents, kv_cache = xattn_layer_0(
                latents, context=embeddings, kv_cache=None
            )
        elif i == 1:
            latents, kv_cache = xattn_layer_1(
                latents, context=embeddings, kv_cache=None
            )
        else:
            # Layers 2 through 8 share the weights with layer 1
            # and take KV cache instead of tile embeddings
            latents, kv_cache = xattn_layer_1(
                latents, context=None, kv_cache=kv_cache
            )

        # Latent transformer; layers 1 through 8 share the weights
        latents = latent_transformer(latents)

    return (
        latents[0, :],  # slide embedding (1280,)
        latents[1:, :],  # latent features to language decoder (512, 1280)
    )
```

Listing 1: Perceiver slide encoder pseudocode

```python
def xattn_layer_i(
    latents: Tensor,  # learned latent embeddings (513, 1280)
    context: Tensor | None,  # Tile embeddings (N, 2560), or None if kv_cache
    kv_cache: Tensor | None,  # Keys and values (2, N, 1280), or None if context
    attention: nn.Module,  # QKV-attention with W_q, W_k, W_v, W_o linear layers
    layer_norm_1: nn.Module,  # Layer normalization layer before attention
    mlp: nn.Module,  # 2-layer MLP with GEGLU non-linearity and inner dim 1280
    layer_norm_2: nn.Module,  # Layer normalization layer before MLP
) -> tuple[Tensor, Tensor]:  # return new latents and new kv_cache
    output, kv_cache = attention(q=layer_norm_1(latents), kv=context or kv_cache)
    latents = latents + output
    latents = latents + mlp(layer_norm_2(latents))
    return latents, kv_cache
```

Listing 2: Cross-attention module pseudocode

### A.10  Acronyms

| | |
|---|---|
| **ADT** | androgen deprivation therapy |
| **AR** | androgen receptor |
| **AUC** | area under (the receiver operating characteristic) curve |
| **BRAF** | B-Raf Proto-Oncogene |
| **BRCA** | breast cancer |
| **CDH1** | cadherin 1 |
| **CNN** | convolutional neural network |
| **CRC** | colorectal cancer |
| **CRPC** | castration resistant prostate cancer |
| **dMMR** | deficient mismatch repair |
| **EGFR** | epidermal growth factor receptor |
| **FGA** | fraction of genome altered |
| **FGFR** | fibroblast growth factor receptor |
| **HER2** | human epidermal growth factor receptor 2 |
| **H&E** | hematoxylin and eosin |
| **HGSOC** | high-grade serous ovarian cancer |
| **HIPT** | the hierarchical image pyramid transformer |
| **IHC** | immunohistochemistry |
| **LOH** | loss-of-heterozygosity |
| **MIL** | multiple instance learning |
| **MSI-H** | high-frequency MSI |
| **MSI** | microsatellite instability |
| **MSK-IMPACT** | MSK-Integrated Mutation Profiling of Actionable Targets |
| **MSKCC** | Memorial Sloan Kettering Cancer Center |
| **NSCLC** | non-small cell lung cancer |
| **OOD** | out-of-distribution |
| **PTEN** | phosphatase and tensin homolog |
| **RNN** | recurrant neural network |
| **TCGA** | The Cancer Genome Atlas |
| **ViT** | vision transformer |
| **WSI** | whole slide image |
| **LUSC** | lung squamous cell carcinoma |
| **LUAD** | lung adenocarcinoma |
| **IDC** | invasive ductal carcinoma |
| **ILC** | invasive lobular carcinoma |
| **DCIS** | ductal carcinoma in situ |