
VIRCHOW 2: SCALING SELF-SUPERVISED MIXED MAGNIFICATION MODELS IN PATHOLOGY

A PREPRINT

Eric Zimmermann¹, Eugene Vorontsov², Julian Viret², Adam Casson², Michal Zelechowski², George Shaikovski²,
Neil Tenenholtz¹, James Hall¹, Thomas Fuchs², Nicolò Fusi¹, Siqi Liu^{‡2}, and Kristen Severson^{‡1}

¹Microsoft Research, Cambridge, MA United States

²Paige, NYC, NY United States

ABSTRACT

Foundation models are rapidly being developed for computational pathology applications. However, it remains an open question which factors are most important for downstream performance with data scale and diversity, model size, and training algorithm all playing a role. In this work, we present the result of scaling both data and model size, surpassing previous studies in both dimensions, and introduce two new models: Virchow 2, a 632M parameter vision transformer, and Virchow 2G, a 1.85B parameter vision transformer, each trained with 3.1M histopathology whole slide images. To support this scale, we propose domain-inspired adaptations to the DINOv2 training algorithm, which is quickly becoming the default method in self-supervised learning for computational pathology. We achieve state of the art performance on twelve tile-level tasks, as compared to the top performing competing models. Our results suggest that data diversity and domain-specific training can outperform models that only scale in the number of parameters, but, on average, performance benefits from domain-tailoring, data scale, and model scale.

1 Introduction

Stained tissue microscopy slides, also known as whole slide images (WSIs), are routinely collected as part of the standard of care in cancer. The digitization of these slides, at increasingly large scales, has enabled a revolution in computational pathology (CPath) towards the use of foundation models. The interest in and success of foundation models are primarily driven by the ability to learn general representations from vast amounts data that include diverse stains, tissue types, and disease without task-specific labels via self-supervised learning approaches. The use of these models represents a paradigm shift in CPath, where previously, expert systems were trained on a single tissue type with expensive, curated labels for a single task. The generalizability and robustness of a foundation model is therefore desirable for the pathology domain as there are many tasks such as diagnosis, disease subtyping, biomarker quantification, estimation of treatment response, and survival prediction. Motivated by these factors, there have been recent efforts to collect large pathology image datasets and subsequently several foundation models have been proposed [53, 17, 3, 12, 50, 21, 24, 57].

Scaling up the dataset and model size has been observed to significantly boost the performance and transferability of foundation models in other domains [34, 60, 41]. In the natural image domain, foundation models use millions of images (e.g. ImageNet [20], JFT-300M [47] and LVD-142M [42]) to train models with hundreds of millions to billions of parameters (e.g. ViT [22]). Early works [12, 50] have suggested these trends exist in CPath as well, both in terms of dataset and model size, however the topic remains underexplored. Furthermore, the extent to which scaling trends observed in the natural image domain hold without adaptation to the unique aspects of CPath data is unknown.

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

[‡]Corresponding author. kseverson AT microsoft DOT com

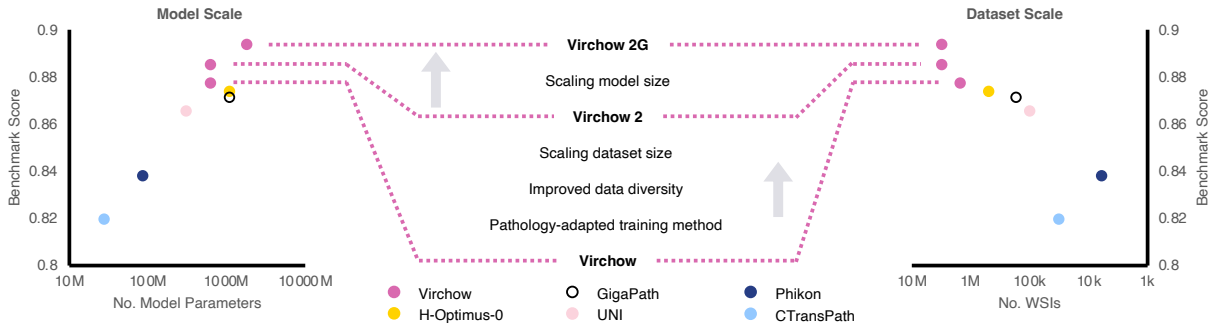


Figure 1: The average weighted F1 score across 8 tile-level benchmarks using public data is shown with respect to two scaling axes: model size (left), in terms of the number of parameters, and dataset size (right), in terms of the number of whole slide images (WSIs). Virchow 2 improves over Virchow with pathology-specific modifications to the training method while scaling up the dataset from 1.5M to 3.1M WSIs, with increased diversity. Virchow 2G then scales the model size from 632M to 1.8B parameters (ViT-H to ViT-G). Benchmark performance across foundation models in pathology appears to scale with model and dataset size.

In this work, we explore scaling along both the data and model size axes, while adjusting the training method with domain-inspired modifications. We present Virchow 2 and Virchow 2G, mixed magnification foundation models for computational pathology. Virchow 2, a 632M parameter vision transformer (ViT-H) expands on our work developing Virchow [50], scaling the dataset size from 1.5M WSIs from a single institution with routine hemotoxylin and eosin (H&E) staining, to 3.1M WSIs from globally diverse institutions with diverse staining such as immunohistochemistry (IHC). Virchow 2G builds on Virchow 2 by increasing the model size to 1.8B parameters (ViT-G).

We evaluate performance on various in-domain and out-of domain benchmarks at multiple magnifications and demonstrate the benefits of our domain-inspired adaptations at scale.

2 Background & Related Work

2.1 Self-Supervised Learning

Self-supervision is a learning paradigm where generalizable features are learned from data using a pretext task, an objective that is constructed from implicit information contained in an input. Joint-embedding self-supervised learning (JE-SSL) methods are a subclass of SSL approaches that pose the learning objective in terms of alignment and diversity [7, 52]. Alignment is accomplished by encouraging features, otherwise known as embeddings, of pairs or sets of samples generated from the same source image via the application of an augmentation policy to be close to one another [13, 59, 5]. Diversity provides the necessary support to learn representations that avoid collapse or trivial solutions across the entire set of observations [52, 14, 11, 25, 59]. These methods, in particular self-distillation variants, have emerged as powerful techniques in no-tuning transfer, both in natural images and in CPath [25, 11, 42, 28, 50, 33].

2.2 CPath Foundation Models

In the past two years, at least ten CPath foundation models have been proposed in the literature (see Tab. A1 for an overview). The CTransPath model [53] was the first and trained a 28M parameter Swin transformer model using MoCoV3 [16] using 15M tiles from 32K WSIs. Since that time, two works [17, 3] have proposed models using SimCLR [13], but the vast majority [24, 12, 50, 57, 40, 44, 21] have used DINOv2 [42] or the related iBOT [61]. The prevalence of DINOv2 appears to be largely motivated by its success with natural images as few studies have released comparative analyses for the choice of algorithm. Kang et al. [33] compared four algorithms, DINO [11], MoCo v2 [15], SwAV [10] and Barlow Twins [59] and found no clear best method. Conversely, Chen et al. [12] found that the choice of algorithm has a large effect on performance, observing that MoCoV3 significantly underperformed DINOv2 in CPath, despite the two approaches having comparable performance on natural image tasks.

Although CPath foundation models have largely aligned on their use of DINOv2, they are quite diverse in terms of the training dataset. Earlier works [53, 17, 3, 24, 33] relied primarily on open access data, specifically The Cancer Genome Atlas (TCGA) [55], a repository of approximately 30K WSIs. Models were trained using this data with varying numbers of tile samples ranging from 4-50M. More recently [12, 21, 50, 44, 40, 31], proprietary datasets

ranging from 100K-1M WSIs have emerged, allowing for training dataset sizes to reach billions of tiles. It is worth noting that characterizing data scale in CPath is non-obvious as various works have quantified it in terms of WSIs, tiles (crops of WSIs used for training), and patients. There are further data dimensions which are also of interest such as source institution, disease status, tissue of origin, stain, and scanner.

3 Learning in the pathology domain

Unlike natural images, which are a 2D projection of a 3D scene, CPath images are acquired from essentially 2D slices of tissue that are stained and scanned at various resolutions or magnifications (e.g. $5\times$, $10\times$, $20\times$, $40\times$) through a digitization process. The resulting image features vastly differ from those found in the natural image domain, as tissue patterns are repetitive, non-unique, pose-invariant, and contain meaningful but minimal color variation due to routine staining procedures, such as H&E. CPath foundation models are trained with self-supervised learning on tissue tiles, i.e. image crops, as WSIs are gigapixel images that are costly to train with directly. Tile size is selected based on compute cost, and the quality of tile content can be highly variable. Despite these differences from the natural image domain, most CPath foundation models have directly used training algorithms developed for natural images and instead focus on the dataset as the primary differentiator. Given the crucial impact of augmentation, algorithm, and model size on foundation model performance, we present modifications to the DINOv2 recipe at scale. In the following sections, we address domain differences by adopting more domain-relevant augmentations and regularization techniques through the lens of feature alignment and diversification.

3.1 Domain-specific augmentations

In the context of joint-embedding self-supervised learning, feature invariance and equivariance are learned through selective alignment of image pairs, also known as views, that are generated by perturbing an input with a set of random augmentations. The quality of features is dependent on the complexity of the pretext task as well as its relationship to the downstream objective.

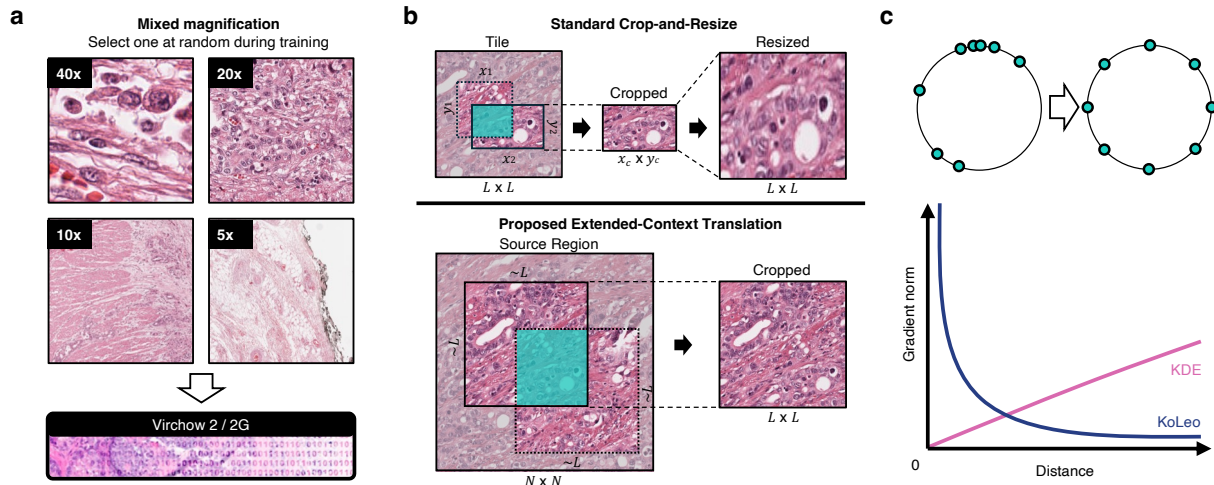


Figure 2: **a.** Virchow 2 and 2G are trained with mixed magnification samples; slides are tiled at 40x, 20x, 10x, or 5x. **b.** Illustration of the difference between standard crop-and-resize augmentation (top) and the proposed extended-context translation (ECT) augmentation (bottom). ECT is expected to be more suitable for pathology images as it avoids distorting the cell morphology. We tuned the ECT approach to retain the average overlap between crops (shaded, teal). **c.** A diversity regularization objective seeks to push apart tile embeddings on a hypersphere (top). For Virchow 2 and 2G, we replace the unstable KoLeo regularizer used by Virchow with a kernel density estimator (KDE) that does not produce exploding gradients for similar tiles (bottom), which is particularly important for pathology whole slide images as they contain many highly correlated tiles.

Augmentations are typically categorized as either photometric or geometric, and most augmentation strategies have been designed for object-centric natural images. In the majority of instances, these augmentations are directly applied in pathology applications, but there are a small number of exceptions. Examples of pathology-specific photometric augmentations include stain jitter and transfer [46, 26], where colors are augmented within and across tiles, aiming to account for differences in staining protocols and facilitate the learning of color invariances. In benchmarking

studies, Tellez et al. [48], Ciga et al. [17], and Gullapally et al. [26] investigated the impact of color augmentation and demonstrated improvements in performance using domain-specific approaches. Since then, several models have employed domain-orientated approaches [53, 33, 21]. However, the benefit of these augmentations when applied to mixed stain datasets, e.g. datasets that contain IHC, is unclear. Other photometric augmentation methods such as image solarization have been omitted in some CPath works [23, 33, 21], as this augmentation is hypothesized to generate color profiles that are not useful for learning relevant invariances.

The random crop operation is arguably the most important geometric augmentation [39]. This augmentation introduces feature co-occurrences within and across samples. These co-occurrences may be explicitly generated based on the augmentation parameters which in turn govern the amount of expected overlap [27, 29] or implicitly generated based on the data distribution. The crop augmentation is paired with a resize operation due to a fixed input size and a desire for scale invariance (the resulting combination is referred to as crop-and-resize). As cell morphology plays an important role in understanding tissue structure and disease, depending on the aspect ratio of the crop, a resize may introduce unwanted distortions, affecting tissue and cell shapes. Indeed, Ciga et al. [17] investigated the impact of random cropping on model performance and found less random cropping generally improved performance although the largest observed improvement in any setting was 5%. Chen et al. [12] deviates from DINOv2’s default parameters and explores less aggressive crop-and-resize scale parameters without changing aspect ratio. As noted above, one unique aspect of pathology data is that tiles are sourced from larger WSIs. As a result of this workflow, image size is a design variable and boundaries can be extended without the need for padding or resizing. Because of this, we propose a translation-like augmentation that uses a larger source field of view to create crops with minimal resizing thereby minimizing morphological distortions while maintaining the same output size and expected overlap between image views. This operation differs from typical translations as it does not introduce imaging artifacts along boundaries, i.e. padded values. It also benefits from extending the relative field-of-view of the WSI without computational overhead. To demonstrate how this is possible, consider replacing the $L \times L$ -sized training tile with an $N \times N$ -sized source region and a target $L \times L$ tile to be sub-sampled within the larger window, where $N > L$ (see Fig. 2). The particular value for N can be selected based on the desired intersection over union of views. This alternative augmentation, which we refer to as extended-context translation (ECT), can serve as a drop-in replacement for traditional crop-and-resize approaches and can be combined with other augmentations such as photometric augmentations.

3.2 Accounting for Tissue Redundancy

Self-supervised alignment tasks are inherently unstable and require additional regularization to encourage diversity and avoid dimension collapse. DINOv2 encourages diversity through contrastive and non-contrastive objectives. Non-contrastive diversity is achieved using asymmetric centered and sharpened distillation between the student and teacher models, while contrastive diversity aims to maximize differential entropy using KoLeo. Different diversity regularization objectives often share similar optima such as uniformity of the hypersphere; however, the dynamics of each method throughout training are considerably different and introduce practical challenges that must be addressed.

Feature diversity is dependent on the training data distribution and training methodology. In CPath, the likelihood of contrasting two similar tissue tiles may be high, unlike when contrasting images from large uncurated natural image datasets. Therefore, any contrastive diversity estimators should be constructed under the assumption that features are not independent, may be very close together, and cannot be arbitrarily separated. DINOv2 approximates differential entropy using KoLeo defined as

$$H_{\text{KoLeo}}(f) = \frac{1}{n} \sum_{i=1}^n \min_{j \neq i} \log d(\mathbf{z}_i, \mathbf{z}_j), \quad (1)$$

where i indexes the samples in a batch of size n , d is a distance measure and \mathbf{z} is an embedding normalized to the hypersphere [43]. The inclusion of an entropy maximization objective using KoLeo aims to spread out learned embeddings via a convex regularizer. While this method has demonstrated value in natural image pretraining settings for linear and nearest neighbour classification objectives where sample diversity is high [42], it is clear that when two samples are very similar, and distance approaches zero, the loss approaches infinity. In practice, KoLeo is stabilized with the addition of an $\epsilon > 0$ and is implemented such that samples are not compared across devices. These factors introduce numeric and hardware dependent hyperparameter that are unexplored.

Any entropy estimator can be used as an effective replacement in order to mitigate the issues associated with KoLeo in a setting where features may be clustered, if the estimator is bounded for nearby points. We select the kernel density estimator (KDE), defined as

$$H_{\text{KDE}}(f) = -\frac{1}{n} \sum_{i=1}^n \log \sum_{j=1}^n k(\mathbf{z}_i, \mathbf{z}_j), \quad (2)$$

where k is an appropriate kernel function and all other terms are as defined above. Self-comparisons ensure the the gradients are bounded. Similar objectives have been used in various other methods [58, 52]. The kernel density estimator has different regularization effects for local behaviors, i.e. for samples with large and small distances, but maintains the population effect of promoting diversity. Examples of kernels include Gaussian, von Mises-Fisher, inverse multi-quadratic, or Laplacian kernel each of which allows for different repulsion characteristics that have bounded gradients which can also vanish for a given pair. In our work, density estimation for entropy regularization is computed using an unnormalized von Mises-Fisher (vMF) kernel,

$$k_{\text{vMF}}(\mathbf{x}, \mathbf{y}) = \exp(\kappa \mathbf{x}^\top \mathbf{y}), \quad (3)$$

where κ is a scaling constant also known as the concentration. The vMF kernel is selected because of its favorable computational qualities and its demonstrated success in encouraging diverse embeddings [52, 8].

4 Stabilizing self-supervised vision transformers at scale

Vision transformers are notoriously challenging to train. Architecture based instability in self-supervised learning protocols have been empirically correlated with high gradient norms occurring in the patch embedding layer at the stem of the network for large batch sizes [16]. As model size is increased, large gradient norms become increasingly problematic across layers of the ViT and partial failures may not be recoverable without the manual intervention of a complete rollback prior to the event. This is costly, time consuming, and may not guarantee success. To address the architectural challenges associated with training large models, dual patch normalization (DPN) [37] and query-key normalization (QKN) [19] have been introduced in the literature and are adopted in our work. Both normalization techniques aim to limit problematic gradients, aside from gradient clipping techniques, by standardizing activations throughout the network. The inclusion of DPN circumvents the need to freeze the randomly initialized weights of the linear projection in the patch embedding layer.

Optimizing ViTs requires adaptive methods like AdamW and its variations [?]. The stability of AdamW depends on the efficacy of the warm-up phase, on the learning rate, and on the selection of momentum parameters. Gradient spikes have been correlated with the ratio of first and second moments for an update step which provide insights on when a collapse event is likely to occur [38, 56]. For example, instability has been observed in the weights of the patch embedding layer [16, 56]. It is possible to reduce these issues by lowering the second moment hyperparameter but this may incur reduced downstream task performance. StableAdamW is explored as a means of stabilizing weight updates when the ratio between the first and second moments are too large (such as when nearly converged weights suddenly need a correction during a long training run) by mitigating large update steps with root-mean-square clipping [56].

Aside from architectural and optimization instability, self-supervised methods commonly suffer from partial or complete dimension collapse as a result of the imbalance between alignment and diversity objectives. Like DINO, DINOv2 primarily avoids collapse over class and patch tokens using an asymmetric centered and sharpened cross-entropy loss between the student and teacher networks. Stability over the distillation objective is highly dependent on the sharpness of the distributions which is controlled by temperature parameters. DINO studied the effects of varying teacher temperature for a fixed student temperature. Results demonstrated that higher teacher temperatures yield better transfer performance but require a warm-up period to avoid collapse, whereas a lower choice of temperature produced marginal performance drops with increased stability [11]. Moreover, registers [18] may be added mitigate the possibility of encountering high norm tokens as model size increases. High norm tokens are explored in object-centric settings where background information is redundant and of lesser importance. In contrast, pathology image tiles filtered to the tissue foreground have no such bias, thus the addition of registers is precautionary.

5 Ablations investigating learning in the pathology domain

We present an ablation study to evaluate modeling choices for the pathology domain. In order to understand the performance impacts of various augmentations and regularization methods prior to model and data scaling, a sweep is performed at a smaller scale on limited data to assess the viability of each technique. The goal of the ablation is to gain better insights into how different elements of a complex training procedure are coupled, despite the possibility that any performance gains or losses attributed to these changes may not be a factor at scale for long training horizons. The axes of exploration are in terms of geometric and photometric augmentations through the use of crop-and-resize and ECT, as well as the presence of solarization, while regularization is explored through the use of KoLeo and KDE.

5.1 Training data

The ablation training dataset is comprised of 1.5M H&E stained WSI from 120K MSKCC patients scanned at $5\times$, $10\times$, $20\times$, $40\times$ resolutions across 17 tissue types (the original Virchow training set [50]). WSI are subdivided

into 224×224 or 392×392 non-overlapping tiles and filtered to include a minimum tissue coverage of 45% as determined by a hue, saturation, and value (HSV) filter with an acceptance criterion in ranges [90, 180], [8, 255], [103, 255], respectively.

5.2 Pretraining methodology

A ViT-B/16 is trained using variations of DINOv2 [42] on 224×224 global views and 96×96 local views sampled from image tiles. The method uses a batch size of 1024, a stochastic depth drop rate of 0.4, and the AdamW optimizer with a learning rate equal to 2×10^{-4} on a cosine schedule for approximately 112K iterations, or equivalently 115M tiles sampled from the dataset detailed in section 5.1, using 16 Nvidia V100 GPUs. The student DINO and iBOT temperature is set to 0.1 while the teacher temperature is warmed up from 0.04 to 0.07 over a period of 12K iterations.

The augmentation policy is composed of horizontal and vertical flips, grayscale, and color jitter, while the crop-and-resize is interchanged with ECT. The crop-and-resize operation is applied to 224×224 tiles and follows the default parameters used in DINOv2, which include scale ranges of (0.32, 1.0) and (0.05, 0.32) for global and local views with an aspect ratio range of (0.75, 1.33). ECT is applied to regions of size 392×392 pixels and uses a base scale range of (0.9, 1.1) and aspect ratio range of (0.95, 1.05). The scale range is then adjusted using the relative ratio of the source and target size for both global and local views.

Entropy regularization is performed using KoLeo, a differential method, or KDE, a resubstitution method. Density estimation is performed using the vMF kernel and the concentration is set to 5 for all experiments based on empirical results in the literature.

5.3 Tile benchmarks

A mix of seven in-distribution (ID) and out-of-distribution (OOD) tile level classification tasks across various magnifications are used to evaluate the quality and separability of the learned features of the teacher network via a linear probe protocol.

PanMSK \times is an in-distribution cancer detection task performed at $5\times$, $10\times$, and $20\times$ magnifications. 1,196,171 224×224 pixel tile samples, sourced from 3,999 WSI from MSKCC representing 17 tissue types, are labeled either cancer or benign [50].

PCam is a public (PatchCamelyon) dataset of 327,680 lymph node images labeled as cancer or benign [6, 49]. Images are upsampled from 96×96 at $10\times$ magnification to 224×224 for analysis.

MHIST is a public dataset of 3152 colorectal polyp images ($5\times$ magnification 224×224 pixels) labeled benign or precursor [54], used as a “minimalist histopathology image analysis dataset” (MHIST). Since a validation data split is not provided, we randomly split 10% of the training set into a validation set.

CRC is a public dataset with colorectal cancer (CRC), composed of 100,000 colorectal images ($20\times$ magnification 224×224 pixels) classified into 9 morphologies [35].

MIDOG is a binary mitotic cell detection task that we created from the public MItoSis DOmain Generalization (MIDOG++) dataset, containing 21,806 mitotic and non-mitotic events labeled on $503.7K \times 5K$ pixel WSI regions from several tumor, species, and scanner types at $40\times$ magnification [2, 50]. This data was converted into a binary classification task by expanding each 50×50 pixel annotation to 224×224 regions and then randomly shifting in the horizontal and vertical regions such that the event is not centered in the tile. All negative instances that overlapped with positive instances were removed from the dataset.

5.4 Linear evaluation protocol

Performance is assessed by training a linear classifier on z-score normalized frozen embeddings generated from non-augmented samples. The classifier is trained with a batch size of 4096 using stochastic gradient descent with a cosine learning rate schedule from 1×10^{-2} to 0 for a total of 12.5K iterations, to maximize convergence. The classifier checkpoint achieving the lowest validation loss was selected for evaluation on a holdout test set. All benchmark datasets are partitioned to include a validation set if not provided by the public resource. If no validation set exists, a random split of 10% is created from the training set.

5.5 Ablation results

Tab. 1 presents the results of the ablation analysis. In a majority of tasks (5/7), using all three domain-inspired changes, results in the best performance as well as the best average performance on in-distribution tasks. Using only ECT and

KDE regularization has the best performance of 3/7 tasks and the highest average performance on out-of-distribution tasks, although the gap with all three is only 0.002. Interestingly, we observe that using only ECT leads to a small drop in performance while using only KDE regularization leads to a small improvement in performance on out-of-distribution tasks. Combining the two leads to a pronounced improvement, as noted above. All changes improved in-distribution tasks.

Geo. Aug.	Phot. Aug.	Reg.	PanMSK 20×	PanMSK 10×	PanMSK 5×	PCam	CRC	MHIST	MIDOG	ID Avg.	OOD Avg.
			0.860	0.888	0.890	0.833	0.940	<i>0.793</i>	0.628	0.879	0.798
✓			0.866	0.895	0.899	<i>0.848</i>	0.933	0.766	0.635	0.886	0.796
		✓	0.880	0.901	0.900	0.841	0.949	0.777	<i>0.664</i>	0.894	0.808
✓		✓	<i>0.894</i>	<i>0.920</i>	<i>0.922</i>	0.867	<i>0.953</i>	0.804	0.665	<i>0.912</i>	0.822
✓	✓	✓	0.899	0.924	0.933	0.867	0.958	0.792	0.663	0.919	<i>0.820</i>

Table 1: Weighted F1 results for seven evaluation tasks and five model settings with a ViT-B/16. For each change, a checkmark indicates a deviation from the standard DINOv2 training protocol. In-distribution evaluation tasks include PanMSK XX, while out-of-distribution tasks include PCam, CRC, MHIST, and MIDOG. Boldface indicates the best performance and italics indicates the second best performance. In a majority of tasks and on average, ECT, KDE, and no solarization regularization performs best.

6 Results

We improve on the training recipe of Virchow with pathology specific tuning and scale up the training first along the data axis (Virchow 2) and then along the model axis (Virchow 2G). We evaluate these models against other foundation models in pathology on tile-level benchmark tasks.

6.1 Training data

Virchow 2 is trained on a dataset of 3.1M biopsy and resection WSIs processed from 225,401 patients across 5×, 10×, 20×, 40× magnifications, with both hematoxylin and eosin (H&E) and immunohistochemical (IHC) staining. This dataset extends the 1.5M WSI dataset that Virchow was trained on and is no longer restricted to 17 tissue types (indeed, it contains nearly 200 recorded tissue types). In addition to samples collected by the Memorial Sloan Kettering Cancer Center (MSKCC) 15% of the WSIs and 57% of the patients are sourced from diverse institutions worldwide that submitted challenging cases to MSKCC for consultation.

Non-overlapping region tiles of size 392 × 392 pixels with at least 65% tissue by area are used for model training. Note that Virchow 2 and Virchow 2G extract 224 × 224 crops from these tiles. Tissue segmentation is performed with a trained fully-convolutional network, with post processing via Otsu thresholding with thresholds of (0.4, 0.5).

Due to the long-tail nature of the data, balancing is explored in terms of tissue type, diagnosis, stain, and magnification. Metadata is not complete for all samples and the presence of cancer, precursor, and neoplastic tissue does not imply the lack of benign tiles within these large samples. Given these sources of noise, we aim to balance with the best available data. Cancer is sampled at 40% and precursor is sampled at 15%, which closely matches the observed frequency of precursor samples in the dataset (17.9%, where 12.3% contain both cancer and precursor and 5.6% contain only precursor). Benign and benign non-neoplasms are sampled at approximately half their observed frequencies for a total of 8%. Benign neoplasms are sampled at 2%. Samples with unknown diagnosis are sampled at 35% (see Fig. 2c for the observed distribution). Tissue type balancing is performed to flatten out the distribution while maintaining the proportion of IHC in an H&E dominant dataset. This approach primarily decreases the prevalence of breast and increases the prevalence of tissues with small sample size and unknown tissue (see Fig. 3d). Magnification balancing is performed to increase uniformity in batches during training under the assumption that each resolution may contain unique feature sets that should be contrasted against one another. 40% of WSI do not have 40× magnifications, therefore these slides are sampled 1.5× more often than those which do not. Sampling of the magnifications is also balanced to yield a final approximate distribution of 20%, 40%, 20%, and 20% for 40×, 20×, 10×, 5×, respectively.

6.2 Pretraining methodology

Two models are trained for a full schedule using insights from the ablations detailed in section 5.2. Both models include ECT instead of crop-and-resize augmentations as well as KDE rather than KoLeo for diversity regularization.




Figure 3: Overview of the training dataset. **a**. A breakdown of the scale of the training dataset with 3.1M whole slide images (WSIs), showing the distribution of slides with hemotoxylin and eosin (H&E) and immunohistochemical (IHC) stains. **b**. The data are globally diverse, with 15% of the slides and 57% of the patients sourced from many institutions across all continents. These slides were submitted to MSKCC for review. **c**. The distribution of most severe diagnoses of the WSIs, in decreasing order of severity: cancer, precursor, benign, unknown. **d**. The top identified tissue types shown as a fraction of the WSIs. There are over 150 additional tissue types in the long tail, summarized here as ‘Other Tissues’ (shown gray, along with ‘Unknown Tissue’, on the right). Online data balancing when training Virchow 2G affected these tissue distributions as shown by the magenta dotted line.

Global tiles of size 224×224 pixels and local tiles of size 98×98 are extracted from the 392×392 extended context source regions, following the DINOv2 multi-view recipe. Each model is trained on 512 Nvidia V100 GPUs. Further configuration details are found in appendix A2.

Virchow 2, a ViT-H/14 with 4 registers, explores the effects of increased data scale and diversity across a mix of magnifications paired with the algorithmic changes proposed. Virchow 2 is trained with solarization on a cosine schedule and is optimized with AdamW for a base learning rate equal to 2×10^{-4} with a batch size of 4096 over a total of 2B unbalanced tiles.

Following the modifications made to the ViT-g architecture to maximize compute efficiency [42], we introduce Virchow 2G, a ViT-G/14 that adopts similar architectural changes. Virchow 2G increases the embedding dimension from 1,664 to 1,792 to be a multiple of 256 and increases the number of attention heads from 16 to 28 in order to have 64 dimensions per head. Additionally, Virchow 2G adopts additional normalization via DPN and QKN and has 8 registers. Virchow 2G is trained without solarization on a cosine schedule and is optimized with StableAdamW for a base learning rate equal to 1×10^{-4} and a batch size of 3072 over a total of 2B balanced tiles. For stability, the second moment optimizer hyperparameter was reduced from 0.999 to 0.95. Additionally, both iBOT and DINO head teacher temperatures are reduced from 0.07 to 0.04. Further configuration details are found in appendix A3.

6.3 Embedding generation

For a 224×224 input tile image, a Virchow embedding is defined as the concatenation of the class token and the mean across all 256 of the patch tokens. This produces an embedding size of 2,560 ($1,280 \times 2$) for Virchow, or 3,584

(1,792 × 2) for Virchow 2 / 2G. The class token is used for H-optimus-0 [44], Prov-GigaPath [57], UNI [12], and Phikon [24]. For CTransPath [53], the mean of all tokens is used as there is no class token.

6.4 Tile benchmarks

The benchmarking protocol is the same as described in Sec. 5.3. In addition to the aforementioned tasks, we add an additional four out-of-distribution tasks.

CRC No-Norm is a variation of the CRC task where the testing set does not use Macenko stain normalization. This simulates a distribution shift between the training or validation sets and the testing set.

WILDS is a public dataset (Camelyon17-WILDS) of 455,954 images, each at a resolution of 96 × 96 pixels at 10× magnification upsampled to 224 × 224 derived from the larger Camelyon17 dataset [36]. The images are labeled to indicate the presence or absence of tumor. The validation and testing data are from institutions not represented in the training data.

DLBCL is a binary 5-year survival prediction task from the public DLBCL-Morph dataset containing 42 tissue microarrays from 209 Diffuse Large B-Cell Lymphoma (DLBCL) cases/patients, with H&E staining and five IHC stains (CD10, BCL6, MUM1, BCL2, and MYC.) at 40× magnification [51]. Due to damaged or missing cores, some stains are missing for some patients. Training, validation, and testing sets are created on the patient level.

TCGA-TILS is a public dataset of 304,097 images, each at a resolution of 100 × 100 pixels at 20× magnification. The images are labeled to indicate the presence of tumor-infiltrating lymphocytes (TILs) if at least two TILs are present [32, 1, 45].

HEST-Benchmark is a public dataset of 74 spatial transcriptomics profiles from 48 patients grouped into 10 tasks based on organ [30]. The aim of the task is to predict the expression levels of the 50 most variable genes from 112 × 112 μm H&E stained image patches, each centered on a spatial transcriptomics spot. Here we present the results of a random forest classifier following the protocol described in Jaume et al. [30], as opposed to the linear protocol described in Sec. 5.3.

6.5 Scaling results

Virchow 2 and 2G, along with Virchow and other baseline foundation models, were evaluated on in-distribution (ID) tile-level benchmarks (Tab. 2) and out-of-distribution (OOD) tile-level benchmarks (Tab. 3). The pathology-specific modifications to the training recipe of Virchow 2, along with increased training data scale and diversity, yielded a substantial boost in ID performance, raising the average weighted F1 score from 0.944 (Virchow) to 0.966 (Virchow 2). Gains in ID performance from further scaling the model size with Virchow 2G were more modest, further raising the score to 0.971. OOD average weighted F1 score increased from 0.877 with Virchow to 0.885 with Virchow 2 and 0.894 with Virchow 2G, demonstrating the benefits of model and data scale, as well as algorithmic improvements. Indeed, Fig. 1 demonstrates a log-linear trend in average OOD performance with model and data scale, across all models.

Although H-optimus-0 and Prov-GigaPath both use a ViT-g with 1.1B parameters, Virchow with ViT-H and 632M parameters matches or exceeds their performance on average, being surpassed only on PanMSK 20×, MHIST, PCam, and WILDS by H-optimus-0 and on PanMSK 5×, PCam, and WILDS by Prov-GigaPath. Virchow 2, also with 632M parameters but with a larger training set and improved training recipe, outperforms H-optimus-0 and Prov-Gigapath on all ID (PanMSK) tasks and MHIST, and bridges the gap on WILDS. Scaling up the model size to ViT-G (1.8B parameters) with Virchow 2G, results in top performance on all benchmark tasks. We note that while Virchow 2G achieves a weighted F1 score of 0.948 on TILS, as compared to the top score of 0.949 (Virchow, H-optimus-0), these scores are similar and performance on this task appears to have saturated. A major source of improvement lies in the addition of mixed magnifications. Virchow 2 and Virchow 2G both significantly improve on ID PanMSK tasks at 10× and 5× magnification. High resolution 40× OOD tasks such as MIDOG and DLBCL further illustrate the benefits of training on magnifications other than 20×.

While we use the concatenation of the class [CLS] token with the mean of the patch tokens for all Virchow models, experiments are also performed using only the class token, with models denoted as “Virchow [CLS]”, “Virchow 2 [CLS]”, and “Virchow 2G [CLS]”. While Virchow 2 [CLS] performs similarly to Virchow 2, Virchow and Virchow 2G benefit from the concatenated embedding. Using Virchow 2 [CLS] may be favourable to using Virchow 2 in practice due to the reduced storage cost for smaller embeddings; however, further task-specific validation is necessary to confirm that performance is retained in this configuration.

Tab. 4 presents the Pearson correlation coefficient results for the HEST-benchmark. Here we observe less consistency in the performance of the models, possibly because the training protocol is fixed and the hyperparameters are not tuned. Virchow 2G has the best performance on 1/10 tasks, second best on 5/10 tasks, and the best average performance. H-optimus-0 has the best performance on 3/10 tasks, second best on 1/10 tasks, and the second best average performance.

Model	PanMSK 20×	PanMSK 10×	PanMSK 5×	Average
CTransPath	0.897	0.892	0.894	0.894
Phikon	0.932	0.926	0.901	0.920
UNI	0.943	0.950	0.936	0.943
Prov-GigaPath	0.940	0.947	0.944	0.944
H-optimus-0	0.952	0.947	0.932	0.944
Virchow [CLS]	0.950	0.949	0.924	0.941
Virchow 2 [CLS]	<i>0.964</i>	0.967	<i>0.967</i>	0.966
Virchow 2G [CLS]	0.966	0.972	0.975	0.971
Virchow	0.950	0.948	0.933	0.944
Virchow 2	<i>0.964</i>	0.966	<i>0.967</i>	0.966
Virchow 2G	0.966	<i>0.971</i>	0.975	<i>0.970</i>

Table 2: Weighted F1 score results for the in-distribution tile-level benchmark tasks. Best results are bolded and second best are in italics. For the Virchow models, embeddings are defined as the concatenation of class and mean patch tokens. [CLS] indicates the use of class tokens only.

Model	CRC	CRC No-Norm	MHIST	PCam	WILDS	TILS	DLBCL	MIDOG	Average
CTransPath	0.962	0.844	0.817	0.871	0.947	0.931	0.542	0.643	0.820
Phikon	0.959	0.888	0.794	0.904	0.972	0.944	0.544	0.700	0.838
UNI	0.963	0.943	0.843	0.934	0.983	0.945	0.565	0.749	0.866
Prov-GigaPath	0.967	0.968	0.822	0.941	0.976	0.941	0.573	0.782	0.871
H-optimus-0	0.963	0.943	0.848	0.937	<i>0.987</i>	0.949	0.582	0.781	0.874
Virchow [CLS]	0.970	0.932	0.836	0.934	0.966	<i>0.948</i>	0.591	0.760	0.867
Virchow 2 [CLS]	0.976	0.971	<i>0.860</i>	0.935	0.985	0.947	0.615	0.800	0.886
Virchow 2G [CLS]	<i>0.974</i>	<i>0.970</i>	0.857	<i>0.944</i>	0.983	0.947	<i>0.628</i>	<i>0.804</i>	<i>0.888</i>
Virchow	0.973	0.968	0.836	0.933	0.971	0.949	0.602	0.787	0.877
Virchow 2	<i>0.974</i>	0.969	0.859	0.935	<i>0.987</i>	<i>0.948</i>	0.606	<i>0.804</i>	0.885
Virchow 2G	0.973	<i>0.970</i>	0.864	0.947	0.988	<i>0.948</i>	0.629	0.836	0.894

Table 3: Weighted F1 score results for the out-of-distribution tile-level benchmark tasks. Best results are bolded and second best results are in italics. For the Virchow models, embeddings are defined as the concatenation of class and mean patch tokens. [CLS] indicates the use of class tokens only.

7 Discussion

Our results suggest that foundation model performance in computational pathology continues to benefit from model and data scale. Nevertheless, substantial gains can still be achieved by methodological improvements even at small scales. We therefore built Virchow 2 and Virchow 2G, which scale the dataset size and then the model size, respectively, on top of modifications of the DINOv2 training setup, tailored to the pathology domain.

Our ablation results demonstrate that such modifications can significantly improve benchmark performance. Changing the KoLeo regularizer to a kernel density estimator (KDE) also boosted performance. This is likely because the conditions that produce instability with KoLeo are particularly common in pathology images: that is, tiles within a minibatch are much more likely to be similar to each other compared to ImageNet [20] images within a minibatch. Furthermore, the proposed extended-context translation (ECT) is more fitting for pathology data than crop-and-resize because it avoids distorting the cell morphology which is crucial for interpreting pathology images. Although decreasing random resizing during training may reduce the model’s scale invariance, we can retain it for pathology by training

	IDC	PRAD	PAAD	SKCM	COAD	READ	ccRCC	HCC	LUAD	LIDC	Avg.
CTransPath	0.476	0.363	0.416	0.557	0.118	0.093	0.177	0.063	0.560	0.255	0.308
Phikon	0.468	0.383	0.397	0.581	0.136	0.138	0.175	0.043	0.566	0.264	0.315
UNI	0.498	0.371	0.437	0.637	0.142	0.160	0.178	0.052	0.565	0.261	0.330
Prov-GigaPath	0.514	0.386	0.437	0.577	0.144	0.156	0.188	0.057	0.568	<i>0.274</i>	0.330
H-optimus-0	0.536	0.378	0.446	0.649	0.159	0.160	0.210	0.068	<i>0.591</i>	0.271	<i>0.347</i>
Virchow [CLS]	0.529	0.364	0.431	0.612	<i>0.150</i>	0.120	0.205	<i>0.065</i>	0.584	0.274	0.334
Virchow 2 [CLS]	0.563	0.379	0.402	0.631	0.143	<i>0.169</i>	0.224	0.060	<i>0.591</i>	0.275	0.344
Virchow 2G [CLS]	0.547	0.375	0.420	<i>0.638</i>	0.140	0.146	0.213	0.056	0.590	0.271	0.339
Virchow	0.545	0.372	0.465	0.624	0.159	0.133	0.211	<i>0.065</i>	0.601	0.269	0.344
Virchow 2	0.539	0.382	0.425	0.617	0.127	0.168	0.226	0.056	0.586	0.273	0.340
Virchow 2G	<i>0.559</i>	<i>0.385</i>	<i>0.458</i>	0.632	0.139	0.175	<i>0.222</i>	0.062	0.588	<i>0.274</i>	0.350

Table 4: Pearson correlation results for the HEST-Benchmark tasks. Best results are bolded and second best results are in italics. For the Virchow models, embeddings are defined as the concatenation of class and mean patch tokens. [CLS] indicates the use of class tokens only.

across multiple magnifications. Similarly, removing solarization boosted performance (especially on PCam), as also previously observed [33, 21]. Indeed, it has been shown that the effects of augmentation are often data dependent and vary across classes [4], however, with a sufficiently large and diverse training dataset, the importance of photometric data augmentations diminishes [39]. Nevertheless, suitable image transforms, like cropping, are still necessary between the two views of an input in joint embedding models like DINOv2.

Further adjustments to the training recipe were required to allow scaling Virchow both along the data size and the model size axes. Without KDE, DINOv2 training collapsed before enough of the 3.1M slide dataset could be processed. Additional stabilizations were required to then scale up the model from ViT-H (632M parameters) to ViT-G (1.8B parameters). We note that largest models previously trained with DINOv2 had 1.1B parameters (ViT-g) [42, 44, 57]. The types of challenges encountered are classified into numerical issues and collapse events. Numerical issues are determined when a training run fails due to NaNs that could not be caught and mitigated and may be due to the use of FP16 rather than BF16. These are addressed by DPN, QKN, and gradient and update clipping in StableAdamW. Collapse events are categorized by sudden spikes in a loss component when features collapse to a lower dimensional subspace that is correlated with the degradation in downstream performance. Collapse primarily occurs in the iBOT head and is mitigated by reducing the learning rate, reducing the β_2 momentum parameter in StableAdamW from 0.999 to 0.95, and reducing the teacher temperature from 0.07 to 0.04 based on the guidance in DINO [11]. No collapse was observed in the DINO head. It is possible that diversity regularization provided sufficient support to avoid catastrophic events where asymmetric centering and sharpening was insufficient. These modifications reduce the convergence rate but allow training longer.

Based on the scaling plots (Fig. 1) and performance of Virchow 2 and Virchow 2G, benchmark performance continues to improve with model and data scale; however, additional benchmarks on which performance has not yet saturated may be necessary to gauge further gains (like MIDOG, DLBCL, or PanMSK, all of which contain magnifications other than 20 \times). No substantial improvements with model scale were observed on WILDS, CRC, CRC-No-Norm, or TILS beyond ViT-H (632M parameters). Indeed, F1 scores for these tasks, as well as PCam and PanMSK are between 0.95 and 0.99 (where 1.00 is the maximum), suggesting that further performance gains may be insignificant or unmeasurable due to label noise. MHIST results are difficult to interpret due to their high variance. Although PanMSK [50] covers 17 tissue types (all the tissue types that Virchow was trained on), it does not evaluate the long tail of tissue types learned by Virchow 2 and Virchow 2G and it could be made more challenging by enriching it for rare cancers. Indeed, in order to quantify the benefits and impact of model and data scale, it is imperative to prepare diverse benchmarks across different tasks, as some benchmarks may saturate more quickly than others. Without a reliable suite of downstream tasks, confident claims cannot be made about the quality of a model.

Ultimately, there is still room to improve benchmark performance along the axes of model scale, data scale, and algorithmic improvements. We demonstrate the value of pathology-specific improvements, as well as training on mixed magnifications, allowing improved performance on diverse benchmark tasks towards a unified image-based foundation model.

8 Acknowledgements

We gratefully thank Philip Rosenfield from Microsoft and Djamila Dierov from Paige for their contributions in making this collaboration possible, Philippe Mathieu for distributed inference support, Mark Fleishman for data support, and Wayne Hendricks and Alexander van Eck at Paige and Jim Jernigan, Lifeng Li, Ben Huntley, Wyman Chong, Oleg Losinets, and the rest of the team at MSR for infrastructure support.

References

- [1] Shahira Abousamra, Rajarsi Gupta, Le Hou, Rebecca Batiste, Tianhao Zhao, Anand Shankar, Arvind Rao, Chao Chen, Dimitris Samaras, Tahsin Kurc, et al. Deep learning-based mapping of tumor infiltrating lymphocytes in whole slide images of 23 types of cancer. *Frontiers in Oncology*, 11:806603, 2022.
- [2] Marc Aubreville, Frauke Wilm, Nikolas Stathonikos, Katharina Breininger, Taryn A Donovan, Samir Jabari, Mitko Veta, Jonathan Ganz, Jonas Ammeling, Paul J van Diest, et al. A comprehensive multi-domain dataset for mitotic figure detection. *Scientific Data*, 10(1):484, 2023.
- [3] 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*, 7(6):756–779, 2023.
- [4] Randall Balestriero, Leon Bottou, and Yann LeCun. The effects of regularization and data augmentation are class dependent, 2022.
- [5] Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.
- [6] Babak Ehteshami Bejnordi, Mitko Veta, Paul Johannes Van Diest, Bram Van Ginneken, Nico Karssemeijer, Geert Litjens, Jeroen AWM Van Der Laak, Meyke Hermsen, Quirine F Manson, Maschenka Balkenhol, et al. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *JAMA*, 318(22):2199–2210, 2017.
- [7] Florian Bordes, Randall Balestriero, and Pascal Vincent. Towards democratizing joint-embedding self-supervised learning. *arXiv preprint arXiv:2303.01986*, 2023.
- [8] Sergiy V Borodachov, Douglas P Hardin, and Edward B Saff. *Discrete energy on rectifiable sets*, volume 4. Springer, 2019.
- [9] Gabriele Campanella, Ricky Kwan, Eugene Fluder, Jennifer Zeng, Aryeh Stock, Brandon Veremis, Alexandros D Polydorides, Cyrus Hedvat, Adam Schoenfeld, Chad Vanderbilt, et al. Computational pathology at health system scale—self-supervised foundation models from three billion images. *arXiv preprint arXiv:2310.07033*, 2023.
- [10] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Advances in Neural Information Processing Systems*, 2020.
- [11] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.
- [12] Richard J Chen, Tong Ding, Ming Y Lu, Drew FK Williamson, Guillaume Jaume, Andrew H Song, Bowen Chen, Andrew Zhang, Daniel Shao, Muhammad Shaban, et al. Towards a general-purpose foundation model for computational pathology. *Nature Medicine*, pages 1–13, 2024.
- [13] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, 2020.
- [14] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [15] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020.
- [16] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.
- [17] Ozan Ciga, Tony Xu, and Anne Louise Martel. Self supervised contrastive learning for digital histopathology. *Machine Learning with Applications*, 7:100198, 2022.

- [18] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. *arXiv preprint arXiv:2309.16588*, 2023.
- [19] Mostafa Dehghani, Josip Djolonga, Basil Mustafa, Piotr Padlewski, Jonathan Heek, Justin Gilmer, Andreas Peter Steiner, Mathilde Caron, Robert Geirhos, Ibrahim Alabdulmohsin, et al. Scaling vision transformers to 22 billion parameters. In *International Conference on Machine Learning*, 2023.
- [20] 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*, 2009.
- [21] Jonas Dippel, Barbara Feulner, Tobias Winterhoff, Simon Schallenberg, Gabriel Dernbach, Andreas Kunft, Stephan Tietz, Philipp Jurmeister, David Horst, Lukas Ruff, et al. Rudolfv: A foundation model by pathologists for pathologists. *arXiv preprint arXiv:2401.04079*, 2024.
- [22] 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. *arXiv preprint arXiv:2010.11929*, 2020.
- [23] Khrystyna Faryna, Jeroen van der Laak, and Geert Litjens. Tailoring automated data augmentation to h&e-stained histopathology. In *Proceedings of the Fourth Conference on Medical Imaging with Deep Learning*.
- [24] Alexandre Filiot, Ridouane Ghermi, Antoine Olivier, Paul Jacob, Lucas Fidon, Alice Mac Kain, Charlie Saillard, and Jean-Baptiste Schiratti. Scaling self-supervised learning for histopathology with masked image modeling. *medRxiv preprint 23292757*, 2023.
- [25] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent-a new approach to self-supervised learning. In *Advances in Neural Information Processing Systems*, 2020.
- [26] Sai Chowdary Gullapally, Yibo Zhang, Nitin Kumar Mittal, Deeksha Kartik, Sandhya Srinivasan, Kevin Rose, Daniel Shenker, Dinkar Juyal, Harshith Padigela, Raymond Biju, et al. Synthetic domain-targeted augmentation (S-DOTA) improves model generalization in digital pathology. *arXiv preprint arXiv:2305.02401*, 2023.
- [27] Lu Han, Han-Jia Ye, and De-Chuan Zhan. Augmentation component analysis: Modeling similarity via the augmentation overlaps. *arXiv preprint arXiv:2206.00471*, 2022.
- [28] 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*, 2022.
- [29] Weiran Huang, Mingyang Yi, Xuyang Zhao, and Zihao Jiang. Towards the generalization of contrastive self-supervised learning. *arXiv preprint arXiv:2111.00743*, 2021.
- [30] Guillaume Jaume, Paul Doucet, Andrew H Song, Ming Y Lu, Cristina Almagro-Pérez, Sophia J Wagner, Anurag J Vaidya, Richard J Chen, Drew FK Williamson, Ahrong Kim, et al. Hest-1k: A dataset for spatial transcriptomics and histology image analysis. *arXiv preprint arXiv:2406.16192*, 2024.
- [31] Dinkar Juyal, Harshith Padigela, Chintan Shah, Daniel Shenker, Natalia Harguindeguy, Yi Liu, Blake Martin, Yibo Zhang, Michael Nercessian, Miles Markey, et al. Pluto: Pathology-universal transformer. *arXiv preprint arXiv:2405.07905*, 2024.
- [32] JR Kaczmarzyk, S Abousamra, T Kurc, R Gupta, and J Saltz. Dataset for tumor infiltrating lymphocyte classification (304,097 image patches from tcga), 2022. URL <https://doi.org/10.5281/2022>.
- [33] 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*, 2023.
- [34] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- [35] Jakob Nikolas Kather, Niels Halama, and Alexander Marx. 100,000 histological images of human colorectal cancer and healthy tissue. *Zenodo*, April 2018. doi: 10.5281/zenodo.1214456. URL <https://doi.org/10.5281/zenodo.1214456>.
- [36] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, 2021.
- [37] Manoj Kumar, Mostafa Dehghani, and Neil Houlsby. Dual patchnorm. *arXiv preprint arXiv:2302.01327*, 2023.

- [38] Igor Molybog, Peter Albert, Moya Chen, Zachary DeVito, David Esiobu, Naman Goyal, Punit Singh Koura, Sharan Narang, Andrew Poulton, Ruan Silva, et al. A theory on adam instability in large-scale machine learning. *arXiv preprint arXiv:2304.09871*, 2023.
- [39] Théo Moutakanni, Maxime Oquab, Marc Szafraniec, Maria Vakalopoulou, and Piotr Bojanowski. You don't need data-augmentation in self-supervised learning. *arXiv preprint arXiv:2406.09294*, 2024.
- [40] Dmitry Nechaev, Alexey Pchelnikov, and Ekaterina Ivanova. Hibou: A family of foundational vision transformers for pathology. *arXiv preprint arXiv:2406.05074*, 2024.
- [41] OpenAI. GPT-4 Technical Report. Preprint at <https://doi.org/10.48550/arXiv.2303.08774>, 2023.
- [42] 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.
- [43] Alexandre Sablayrolles, Matthijs Douze, Cordelia Schmid, and Hervé Jégou. Spreading vectors for similarity search. *arXiv preprint arXiv:1806.03198*, 2018.
- [44] Charlie Saillard, Rodolphe Jenatton, Felipe Llinares-López, Zeldia Mariet, David Cahané, Eric Durand, and Jean-Philippe Vert. H-optimus-0, 2024. URL <https://github.com/bioptimus/releases/tree/main/models/h-optimus/v0>.
- [45] Joel Saltz, Rajarsi Gupta, Le Hou, Tahsin Kurc, Pankaj Singh, Vu Nguyen, Dimitris Samaras, Kenneth R Shroyer, Tianhao Zhao, Rebecca Batiste, et al. Spatial organization and molecular correlation of tumor-infiltrating lymphocytes using deep learning on pathology images. *Cell Reports*, 23(1):181–193, 2018.
- [46] Yiqing Shen, Yulin Luo, Dinggang Shen, and Jing Ke. Randstainna: Learning stain-agnostic features from histology slides by bridging stain augmentation and normalization. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2022.
- [47] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017.
- [48] David Tellez, Geert Litjens, Péter Bándi, Wouter Bulten, John-Melle Bokhorst, Francesco Ciompi, and Jeroen Van Der Laak. Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology. *Medical image analysis*, 58:101544, 2019.
- [49] Bastiaan S Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. Rotation equivariant cnns for digital pathology. In *Medical Image Computing and Computer Assisted Intervention*, 2018.
- [50] Eugene Vorontsov, Alican Bozkurt, Adam Casson, George Shaikovski, Michal Zelechowski, Kristen Severson, Eric Zimmermann, James Hall, Neil Tenenholtz, Nicolo Fusi, et al. A foundation model for clinical-grade computational pathology and rare cancers detection. *Nature Medicine*, pages 1–12, 2024.
- [51] Damir Vrabac, Akshay Smit, Rebecca Rojansky, Yasodha Natkunam, Ranjana H Advani, Andrew Y Ng, Sebastian Fernandez-Pol, and Pranav Rajpurkar. Dlbcl-morph: morphological features computed using deep learning for an annotated digital dlbc image set. *Scientific Data*, 8(1):135, 2021.
- [52] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, 2020.
- [53] 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.
- [54] Jerry Wei, Arief Suriawinata, Bing Ren, Xiaoying Liu, Mikhail Lisovsky, Louis Vaickus, Charles Brown, Michael Baker, Naofumi Tomita, Lorenzo Torresani, et al. A petri dish for histopathology image analysis. In *Artificial Intelligence in Medicine: 19th International Conference on Artificial Intelligence in Medicine*, 2021.
- [55] John N Weinstein, Eric A Collisson, Gordon B Mills, Kenna R Shaw, Brad A Ozenberger, Kyle Ellrott, Ilya Shmulevich, Chris Sander, and Joshua M Stuart. The cancer genome atlas pan-cancer analysis project. *Nature Genetics*, 45(10):1113–1120, 2013.
- [56] Mitchell Wortsman, Tim Dettmers, Luke Zettlemoyer, Ari Morcos, Ali Farhadi, and Ludwig Schmidt. Stable and low-precision training for large-scale vision-language models. In *Advances in Neural Information Processing Systems*, 2023.
- [57] Hanwen Xu, Naoto Usuyama, Jaspreet Bagga, Sheng Zhang, Rajesh Rao, Tristan Naumann, Cliff Wong, Zelalem Gero, Javier González, Yu Gu, Yanbo Xu, Mu Wei, Wenhui Wang, Shuming Ma, Furu Wei, Jianwei Yang, Chunyuan Li, Jianfeng Gao, Jaylen Rosemon, Tucker Bower, Soohee Lee, Roshanthi Weerasinghe, Bill J. Wright, Ari

- Robicsek, Brian Piening, Carlo Bifulco, Sheng Wang, and Hoifung Poon. A whole-slide foundation model for digital pathology from real-world data. *Nature*, 630(8015):181–188, 2024.
- [58] Chun-Hsiao Yeh, Cheng-Yao Hong, Yen-Chi Hsu, Tyng-Luh Liu, Yubei Chen, and Yann LeCun. Decoupled contrastive learning. In *European Conference on Computer Vision*, 2022.
- [59] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, 2021.
- [60] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [61] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021.

A Appendix

A.1 Foundation models in CPath

Model	Data source	Data size		Model architecture	Model size	Objective function
		WSI	Tiles			
Virchow 2G	MSKCC	3.4M	1.8B	ViT-G	1.8B	DINOv2
Virchow 2	MSKCC	3.4M	1.7B	ViT-H	632M	DINOv2
Virchow [50]	MSKCC	1.5M	2B	ViT-H	632M	DINOv2
Hibou [40]	Proprietary	1.1M	1.2B	ViT-L	307M	DINOv2
H-optimus-0 [44]	Proprietary	500K		ViT-g	1.1B	DINOv2
Campanella et al. [9]	Mount Sinai	400K	3B	ViT-S	22M	DINO, MAE
Prov-GigaPath [57]	Providence	170K	1.3B	ViT-g	1.1B	DINOv2
PLUTO [31]	Proprietary	158K	195M	ViT-S	22M	DINOv2 + MAE + Fourier
RudolfV [21]	TCGA + Proprietary	103K	750M	ViT-L	307M	DINOv2
UNI [12]	Mass-100K	100K	100M	ViT-L	307M	DINOv2
Lunit [33]	TCGA + TULIP	37K	33M	ViT-S	22M	Various
CTransPath [53]	TCGA + PAIP	32K	15M	Swin Transformer	28M	MoCoV3
Remedis [3]	TCGA	29K	50M	ResNet-152	232M	SimCLR
Ciga et al. [17]	TCGA + CPTAC ++	25K	4.2M	ResNet	11-45M	SimCLR
Phikon [24]	TCGA	6K	43M	ViT-B	86M	iBOT

Table A1: Summary of proposed foundation models in computational pathology highlighting the size of the training data, size of the model architecture, and training objective. Works are ordered based on the number of training WSI.

A.2 Hyperparameters

Virchow 2 Hyperparameter	Value
Vision Transformer	
Patch Size	14
Embedding Dimension	1280
Layers	32
Heads	16
MLP Ratio	4
MLP Activation	SwiGLU
MLP Bias	True
QKV Bias	True
Registers	4
Qk Normalization	False
Dual PatchNorm	False
Projection Heads	
Shared Heads	False
Layers	3
Bottleneck Dimension	384
Hidden Dimension	2048
Loss Functions	
Sinkhorn Centering	True
Student Temperature	0.1
Teacher Temperature	(0.07, 0.04)
Prototypes	131072
Regularizer	KDE
Regularizer Parameter	5
Regularizer Weight	0.05
Augmentations	
Tile Context Size	392
Global View Size	224
Local View Size	98
Method	ECT
Aspect Ratio Range	(0.95, 1.05)
Scale Range	(0.9, 1.1)
Solarization	True
Vertical Flips	True
Optimizer	
Optimizer	AdamW
Optimizer Momentum	(0.9, 0.999)
Optimizer Epsilon	1×10^{-8}
Optimizer Weight Decay	(0.04, 0.2)
Learning rate	2×10^{-4}
Learning Rate Schedule	Cosine
Learning Rate Scaling	Square Root-1024
Teacher Momentum	(0.994, 1.0)
Batch size	4096
Gradient clipping norm	3.0
Drop Rate	0.4
Precision	FP16

Table A2: Virchow 2 adapted hyperparameters from DINOv2 trained on 512 NVIDIA V100 32GB GPUs. Default augmentations are not included and are left unchanged.

Virchow 2G Hyperparameter	Value
Vision Transformer	
Patch Size	14
Embedding Dimension	1792
Layers	48
Heads	28
MLP Ratio	4
MLP Activation	SwiGLU
MLP Bias	True
QKV Bias	True
Registers	8
Qk Normalization	True
Dual PatchNorm	True
Projection Heads	
Shared Heads	False
Layers	3
Bottleneck Dimension	384
Hidden Dimension	2048
Loss Functions	
Sinkhorn Centering	True
Student Temperature	0.1
Teacher Temperature	0.04
Prototypes	131072
Regularizer	KDE
Regularizer Parameter	5
Regularizer Weight	0.05
Augmentations	
Tile Context Size	392
Global View Size	224
Local View Size	98
Method	ECT
Aspect Ratio Range	(0.95, 1.05)
Scale Range	(0.9, 1.1)
Solarization	False
Vertical Flips	True
Optimizer	
Optimizer	StableAdamW
Optimizer Momentum	(0.9, 0.95)
Optimizer Epsilon	1×10^{-8}
Optimizer Weight Decay	(0.04, 0.2)
Learning rate	1×10^{-4}
Learning Rate Schedule	Cosine
Learning Rate Scaling	Square Root-1024
Teacher Momentum	(0.994, 1.0)
Batch size	3072
Gradient clipping norm	3.0
Drop Rate	0.4
Precision	FP16

Table A3: Virchow 2G adapted hyperparameters from DINOv2 trained on 512 NVIDIA V100 32GB GPUs. Default augmentations are not include and are left unchanged.