

# Distilling foundation models for robust and efficient models in digital pathology

Alexandre Filiot<sup>1,†,\*</sup>, Nicolas Dop<sup>\*</sup>, Oussama Tchita<sup>1</sup>, Auriane Riou<sup>1</sup>, Rémy Dubois<sup>1</sup>, Thomas Peeters<sup>2</sup>, Daria Valter<sup>2</sup>, Marin Scalbert<sup>2</sup>, Charlie Saillard<sup>2</sup>, Geneviève Robin<sup>1</sup>, Antoine Olivier<sup>1,†</sup>.

<sup>1</sup> Owkin, Inc, <sup>2</sup> Biopimus, Inc.

† Corresponding authors.

\* Equal contribution.

{alexandre.filiot, nicolas.dop, oussama.tchita, auriane.riou, remy.dubois, genevieve.robin, antoine.olivier}@owkin.com, {thomas.peeters, dasha.walter, marin.scalbert, charlie.saillard}@biopimus.com

**Abstract.** In recent years, the advent of foundation models (FM) for digital pathology has relied heavily on scaling the pre-training datasets and the model size, yielding large and powerful models. While it resulted in improving the performance on diverse downstream tasks, it also introduced increased computational cost and inference time. In this work, we explore the distillation of a large foundation model into a smaller one, reducing the number of parameters by several orders of magnitude. Leveraging distillation techniques, our distilled model, H0-mini, achieves comparable performance to large FMs at a significantly reduced inference cost on HEST and EVA public benchmarks. Additionally, we conduct robustness analyses on the PLISM-WSI dataset and a multi-scanner, multi-staining private breast cancer cohort. We demonstrate that our distilled model reaches excellent robustness to variations in staining and scanning conditions, significantly outperforming other state-of-the-art models. This opens new perspectives to design lightweight and robust models for digital pathology, without compromising on performance. We publicly release H0-mini<sup>1</sup> along with `plismbench`<sup>2</sup>, the first robustness benchmark of pathology foundation models based on the PLISM dataset.

**Keywords:** Digital pathology · Self-supervised learning · Distillation · Foundation models.

## 1 Introduction

In recent years, representation learning has revolutionized computational pathology (CPath) by providing efficient models that can serve for a variety of tasks, including biomarker prediction, gene expression prediction, whole slide image (WSI) and tissue classification or survival analysis [23]. Modern CPath frameworks share the common idea to leverage foundation models [3,6,7,18,19,22,25,28]

<sup>1</sup> <https://huggingface.co/biopimus/H0-mini>

<sup>2</sup> Available at <https://github.com/owkin/plism-benchmark>.

(FM) to map small patches of tissues to a lower dimensional space by computing features. Due to label scarcity at the tile-level, recent pathology feature extractors are pretrained using self-supervised learning (SSL) methods such as DINO, iBOT or DINOv2 [2,21,27] and have become a cornerstone of modern CPath frameworks. Recent studies have mostly focused on scaling the pretraining dataset and the model size, resulting in large models whose size can reach more than one billion parameters, pretrained in up to a few million WSIs using the DINOv2 framework [22,25,28]. This yielded spectacular improvements in performance on a variety of downstream tasks, yet introduced an increased computational cost. Besides, the robustness of FMs to variations in sample preparation or digitization has not attracted the same attention as their performance, which is critical to the clinical deployment of CPath workflows [8]. Some datasets and benchmarks have been proposed to investigate this aspect [16,20,24].

To jointly tackle the additional computational cost while promoting the robustness of the models, this study proposes to investigate the distillation of large FMs. Starting from a recent FM, H-Optimus-0 [22], a Vision Transformer (ViT)-giant [4] with more than one billion parameters, we investigate its distillation into a smaller ViT-Base of 86 million parameters. On several public benchmarks, the distilled model, H0-mini, demonstrates competitive performance with significantly larger state-of-the-art models. Additionally, leveraging the PLISM dataset [20] and a private multi-scanner, multi-staining dataset for breast cancer biomarkers prediction, we show that the distilled model significantly outperforms all other foundation models reported in the literature in its robustness to variations in staining and scanning conditions. This uncovers new perspectives to design robust CPath models in view of their adoption in clinical practice.

**Related work.** Distillation is a well known machine learning technique that consists in supervising a student model by a teacher model [10]. In the context of building image foundation models, it has been shown to be more efficient to distill a large model into a smaller one rather than training the small model from scratch [21]. In [5], a simple method is proposed to perform the distillation, and variations corresponding to the various SSL frameworks are investigated such as RoB-DINO and RoB-iBOT. While distillation techniques are well established in computer vision in general, their application to FMs for digital pathology remains relatively scarce. In GPFM [18], 3 foundation models (CONCH [17], Phikon and UNI) are simultaneously distilled into a student model in addition to DINOv2 losses. In [28], a ViT-Small model Virchow2G-Mini is introduced, resulting from the distillation of Virchow2G (a ViT-Giant) on one billion tiles.

## 2 Material and methods

### 2.1 Pre-training setup

**Distillation setup.** In this study, H-Optimus-0 (H0, [22]) is considered as the teacher model. We then follow the general methodology described in [21].

More precisely, for an image  $x$ , let  $x_1$  and  $x_2$  denote two augmented views of  $x$ . Additionally, for  $i \in \llbracket 1, 2 \rrbracket$ , we denote by  $z_i^{(t)}$  (resp.  $z_i^{(s)}$ ) the class tokens output by the teacher (resp. student) model for image  $x_i$ . For  $i \in \llbracket 1, 2 \rrbracket$  and for a patch  $p$ , we also denote by  $z_{i,p}^{(t)}$  (resp.  $z_{i,p}^{(s)}$ ) the patch tokens output by the teacher (resp. student) for the image  $x_i$ . We pass the teacher and student class tokens (resp. patch tokens) through the corresponding DINO (resp. iBOT) head.  $h_i$  denotes the vector of prototype scores output from the head projection of  $z_i$ . Taking inspiration from [5], distillation is performed by combining two objectives.

*DINO objective.* In this setting, only the class scores are used to perform the distillation. Let  $H$  denote the cross-entropy loss. The corresponding loss function  $L_{\text{dino}}$  is defined as

$$L_{\text{dino}} := \left( H(h_1^{(t)}, h_2^{(s)}) + H(h_2^{(t)}, h_1^{(s)}) \right) / 2. \quad (1)$$

*iBOT objective.* The iBOT objective extends the distillation by incorporating patch scores supervision. This loss is defined as the following

$$L_{\text{ibot}} := \frac{1}{2P} \sum_{p=1}^P \sum_{j=1}^2 H(h_{j,p}^{(t)}, h_{j,p}^{(s)}), \quad (2)$$

where  $P$  denotes the total number of patches. Unlike [5] we do not apply masking of the patches and the iBOT loss is applied for all patch scores.

Finally, for all experiments, we keep a sparse exponential moving average (EMA) of the student as in [5] and remove the stochastic depth as in [21]. Following recent studies such as [28, 14] we also removed the Kolo regularization. To speed up the distillation, mixed precision is used as in [21]. H0-mini was trained for 105,000 iterations and a batch size of 2,048 using 128 Nvidia V100 32Go for a total of 4,350 gpu hours. For all downstream experiments, we used the sparse EMA of the student as the feature extractor.

**Pre-training datasets.** We perform the distillation on a dataset of 43M tiles (224x224 at 20x magnification) extracted from 6,093 TCGA slides covering 16 cancer sites. This allows for a fair comparison with Phikon, a ViT-B feature extractor pretrained from scratch on the same dataset with iBOT. Compared to the typical sizes of pre-training datasets, it stands out by its relatively small size. In particular, we note the difference with the pretraining setup proposed in [28], where the distillation is performed on a 1B-tile dataset.

## 2.2 Evaluation setup

**Performance benchmarks.** To rigorously evaluate model performance, we use two publicly available benchmarks<sup>3</sup>. The EVA benchmark [9] consists in 4

<sup>3</sup> The HEST benchmark is available at <https://github.com/mahmoodlab/HEST>, and the EVA benchmark is available at <https://kaiko-ai.github.io/eva/main/>.

patch-level classification tasks, 2 patch-level segmentation tasks, and 2 slide-level classification tasks. For each category of tasks, EVA’s evaluation protocol is fixed to allow for a fair evaluation across feature extractors (*e.g.*, learning rate, batch size or schedulers). We also leverage the HEST-Benchmark [12] which collects gene expression prediction tasks for 9 different indications. For each task, a subset of 50 highly variable genes is considered. To account for various embedding dimensions between feature extractors, we follow the recommended evaluation procedure from HEST-Benchmark, fitting a ridge regression on top of a PCA reduction with 256 components. Additionally, and following the recommended methodology in [28], we use the concatenation of the class token and the mean over all patch tokens as embeddings of our distilled model.

**PLISM robustness dataset.** To evaluate model robustness, we first use the public PLISM-WSI [20] dataset which consists of 46 human tissue types stained using 13 different H&E conditions, and captured using 7 WSI scanners (Figure 1). This results in 91 WSIs later registered using the Elastix software [15]. A matter detection step was then applied, followed by a tiling step to extract 16,278 tissue tiles for each slide, at pre-defined coordinates. We make this processed dataset publicly available<sup>4</sup>.

Cosine similarity and top- $k$  accuracy are used to evaluate the robustness of the feature extractors across scanner and staining combinations. Cosine similarity is assessed on features extracted from tissue tiles averaged over all matching tile pairs. We note that this metric was proposed in [24] to quantify the robustness of feature extractors to stain normalization. Top- $k$  accuracy computes the percentage of tiles from one slide whose matching tile on the other slide ranks among the  $k$  closest tiles (by cosine similarity) when compared to all other tiles from both slides. Finally, metrics are aggregated across slide pairs. Specific robustness to scanner (resp. staining) is assessed on all fixed-staining (resp. scanner) cross-scanner (resp. staining) pairs. We also report overall robustness on cross-stainings, cross-scanners slide pairs.

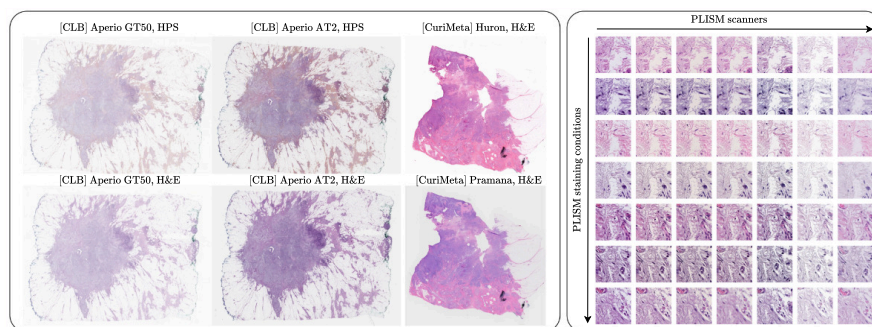
### **BreastBiomarker (BreastBm) performance & robustness benchmark.**

We also evaluate the downstream robustness of H0-Mini to variations in scanner and staining in the context of breast cancer biomarker prediction. Slide-level tasks aim at predicting the positivity of estrogen receptor status (ER), progesterone receptor status (PR), human epidermal growth factor receptor 2 status (HER2) and germline BRCA1/BRCA2 mutations (gBRCA). We leverage 3 private cohorts for robustness evaluation (see Figure 1). The “IGR” (resp. “CLB”) cohorts consists of 671 (resp. 353) tumor blocks providing 2 consecutive slides: one stained with Hematoxylin Eosin Saffron (HES) (resp. archival Hematoxylin Phloxine Saffron, HPS) and one with H&E. All IGR (resp. CLB) slides are scanned using an Olympus VS200 (resp. Aperio GT450 and Aperio AT2). IGR

<sup>4</sup> The processed PLISM-WSI dataset is made available at <https://huggingface.co/datasets/owkin/plism-dataset>

H&E slides are also scanned with Aperio ScanScope and Hamamatsu S60. This yields 6 possible subcohort pairs for IGR and for CLB. The CuriMeta cohort consists of 208 H&E slides scanned with Pramana and Huron scanners (1 pair).

For all endpoints and to best isolate the impact of feature extractors, all downstream models use a mean pooling architecture: a logistic regression is trained on the average of tile-level feature representations. Performance is assessed using AUC while robustness is evaluated via the Concordance Correlation Coefficient (CCC) between WSI-level predictions from different stainings and scanners for the same tumor blocks. For ER and PR (resp. HER2) endpoints, training was performed on 1,109 WSIs from TCGA-BRCA and evaluation on IGR, CLB and CuriMeta cohorts (resp. IGR cohort). Due to labels unavailability in TCGA, predictive models for the gBRCA endpoint were trained on IGR and validated on CLB and CuriMeta.



**Fig. 1.** Visualizations of BreastBm (left) and PLISM-WSI (right) datasets. For PLISM, we only display 7 of the 13 different stainings on the y-axis.

### 3 Results

#### 3.1 Performance evaluation

In Table 1, we report the results on the EVA benchmark. As suggested by its authors, we report the average results without the BACH task, as the spatial resolution of the images differs from the other tasks and therefore tend to favor mixed-magnification models such as Virchow2 or the Kaiko models [1]. Even though the results tend to saturate (6 models show a mean performance between 0.78 and 0.79), H0-mini is competitive with the other state-of-the-art models while being much smaller. A noticeable exception is the Kaiko-B/8 model reaching equivalent performance (but with a smaller patch size which results in more computations). This shows that the distilled model can perform well on a variety of tasks, including patch classification, WSI classification or segmentation.

**Table 1.** EVA results. For all models, input embeddings are only the [CLS] tokens. Balanced accuracy (resp. “MonaiDiceScore”) is reported for classification (resp. segmentation) tasks. (\*) results may be over-estimated due to an overlap between downstream and pre-training datasets.

Model	Size	Patch-level classification				Slide-level classification		Segmentation		Mean
		Bach	Crc	Mhist	Pcam	Cam16	Panda	Consep	Monusac	
Virchow2	632M	<u>0.880</u>	<b>0.966</b>	<b>0.858</b>	0.936	<b>0.864</b>	0.642	0.630	0.663	<b>0.794</b>
UNI2-h	682M	<b>0.914</b>	<u>0.965</u>	0.820	<u>0.949</u>	<u>0.855</u>	<b>0.672</b>	0.632	0.642	<u>0.791</u>
H0	1,100M	0.758	0.958	0.839	0.942	0.820	0.645	0.637	<b>0.679</b>	0.789
Gigapath	1,100M	0.761	0.952	0.829	0.945	0.814	0.664	0.621	0.672	0.785
GPFM	307M	0.830*	0.952	0.811	0.945	0.851*	0.647*	0.639	0.640	0.784
UNI	307M	0.797	0.947	<u>0.844</u>	0.936	0.834	0.656	0.628	0.638	0.783
<b>H0-mini</b>	86M	0.774	0.961	0.790	0.942	0.842	<u>0.667</u>	0.629	0.643	0.782
Hibou L <sub>16</sub>	307M	0.816	0.931	0.826	<b>0.951</b>	0.832	0.633	<u>0.642</u>	0.658	0.782
Kaiko B <sub>8</sub>	86M	0.858	0.957	0.823	0.918	0.818	0.638	<b>0.645</b>	<u>0.675</u>	0.782
Phikon	86M	0.722	0.936	0.799	0.922	0.797	0.640	0.629	0.644	0.767
PhikonV2	307M	0.727	0.939	0.775	0.893	0.808	0.635	0.630	0.639	0.760

In Table 2, we report the relative inference time to extract  $10^4$  tiles (with H0-mini taken as the reference, using 1 NVIDIA T4 GPU equipped with 16Gb vRAM, 32 CPUs and 128 GB RAM) and the results on the HEST benchmark. H0-mini also performs well on this benchmark highlighting the efficiency of distillation. The model significantly outperforms Phikon, its ViT-B equivalent trained from scratch with iBOT on the same pre-training data. Moreover, it surpasses on average much bigger foundation models such as Virchow2 or Gigapath.

**Table 2.** Relative inference times to extract  $10^4$  tiles (taking H0-mini as the reference) and HEST results (Pearson correlation). For all models, [CLS] tokens are concatenated to the mean over all patch tokens to form the input embeddings.

Model	$t/10^4$ tiles	Idc	Prad	Paad	Skcm	Coad	Read	Ccrcc	Luad	L. Idc	Mean
UNI2-h	5.7	<u>0.605</u>	0.375	<b>0.523</b>	<b>0.683</b>	<b>0.332</b>	<u>0.226</u>	0.266	0.574	<b>0.274</b>	<b>0.429</b>
H0	10.1	<b>0.611</b>	0.362	<u>0.511</u>	<u>0.661</u>	0.309	<b>0.240</b>	<u>0.267</u>	<u>0.575</u>	0.266	<u>0.422</u>
<b>H0-mini</b>	1	0.591	0.363	0.507	0.612	0.270	0.205	0.264	0.563	0.264	0.404
Virchow2	6.1	0.597	0.353	0.478	0.640	0.258	0.207	0.260	0.568	0.257	0.402
Hibou L <sub>16</sub>	2.8	0.594	0.323	0.476	0.606	<u>0.313</u>	0.182	<b>0.278</b>	0.572	0.249	0.399
Gigapath	7.0	0.571	<b>0.384</b>	0.492	0.582	0.308	0.186	0.228	0.558	0.250	0.395
GPFM	2.7	0.580	0.373	0.469	0.584	0.280	0.177	0.251	0.552	0.239	0.394
Kaiko B <sub>8</sub>	2.9	0.571	<u>0.383</u>	0.473	0.590	0.310	0.173	0.266	<b>0.588</b>	0.236	0.391
UNI	2.1	0.585	0.327	0.488	0.623	0.258	0.176	0.246	0.556	0.258	0.391
PhikonV2	2.1	0.568	0.379	0.477	0.584	0.256	0.186	0.261	0.550	0.248	0.390
Phikon	1	0.548	0.345	0.464	0.555	0.267	0.167	0.250	0.568	0.239	0.378

Overall ranking and comparison between models is performed through a one-sided Wilcoxon sign rank test with Holm correction (at  $\alpha = 0.05$ ) on HEST and EVA combined (17 tasks). This places H0-mini at 4th position, outperformed by UNI2-h ( $p < 10^{-4}$ ), H-Optimus-0 ( $p = 0.039$ ) and Virchow2 (n.s.,  $p = 0.29$ ). H0-mini outperforms UNI ( $p = 0.044$ ) and the Phikon models ( $p < 10^{-3}$ ).

### 3.2 Robustness evaluation

Table 3 presents the robustness metrics evaluated on the PLISM-WSI dataset. H0-mini achieves high scores, outperforming all other state-of-the-art models. We also note that multi-modal extractors, such as CONCH, demonstrate high robustness. While there is a correlation between the cosine similarity metric and the retrieval metric, cosine similarity alone is not enough to assess the robustness of a FM. Models such as PLIP [11] and Phikon present high values of cosine similarity, while failing at the retrieval task. It should also be noted that robustness to staining is harder to achieve compared to scanner as staining variations may have an impact on biological morphology while scanning variations mostly impact color tints. Looking at top-10 accuracy, H0-Mini presents the highest robustness to scanning variations, as well as joint scanning and staining variations.

**Table 3.** PLISM results (cosine similarity and top-10 accuracy). For each metric, we report median and inter-quartile range over the corresponding slide pairs.

Model	Fixed-staining, cross-scanner		Fixed-scanner, cross-staining		Cross-staining, cross-scanner	
	Cosine	top-10	Cosine	top-10	Cosine	top-10 ↓
<b>H0-mini</b>	0.92 (0.06)	<b>0.86 (0.28)</b>	0.83 (0.09)	0.32 (0.33)	0.79 (0.11)	<b>0.18 (0.25)</b>
H0	0.85 (0.08)	0.74 (0.32)	0.73 (0.11)	<b>0.33 (0.29)</b>	0.67 (0.12)	0.17 (0.20)
Virchow2	0.88 (0.06)	0.61 (0.31)	0.81 (0.09)	0.31 (0.25)	0.77 (0.10)	0.16 (0.18)
CONCH	<b>0.93 (0.04)</b>	0.75 (0.29)	0.86 (0.05)	0.24 (0.21)	0.84 (0.06)	0.16 (0.16)
GigaPath	0.79 (0.11)	0.59 (0.40)	0.61 (0.13)	0.12 (0.18)	0.56 (0.14)	0.05 (0.08)
UNI	0.76 (0.09)	0.53 (0.30)	0.62 (0.13)	0.17 (0.22)	0.53 (0.13)	0.05 (0.08)
UNI2-h	0.76 (0.17)	0.50 (0.56)	0.68 (0.13)	0.19 (0.25)	0.56 (0.17)	0.05 (0.10)
Kaiko B <sub>8</sub>	0.87 (0.07)	0.35 (0.46)	0.81 (0.09)	0.15 (0.20)	0.75 (0.11)	0.05 (0.09)
GPFM	0.80 (0.08)	0.36 (0.34)	0.69 (0.17)	0.09 (0.18)	0.58 (0.14)	0.02 (0.04)
Hibou L <sub>16</sub>	0.65 (0.10)	0.06 (0.09)	0.61 (0.16)	0.03 (0.06)	0.47 (0.14)	0.008 (0.01)
PLIP	0.92 (0.04)	0.05 (0.20)	<b>0.91 (0.05)</b>	0.04 (0.10)	<b>0.87 (0.06)</b>	0.004 (0.01)
Phikon	0.81 (0.08)	0.13 (0.24)	0.69 (0.19)	0.02 (0.07)	0.60 (0.18)	0.004 (0.01)
PhikonV2	0.72 (0.09)	0.06 (0.10)	0.66 (0.18)	0.03 (0.07)	0.54 (0.14)	0.003 (0.008)

Finally, table 4 presents the downstream results on the BreastBm dataset. H0-mini consistently and significantly outperforms H-Optimus-0 on each of the performance and robustness metrics ( $p < 10^{-4}$ ) as assessed by multiple bootstrap-based comparison tests combined with harmonic mean p-value across subcohorts (resp. subcohort pairs) for AUC (resp. CCC). Considering that the fea-



tures aggregation module is a simple average, BreastBm results suggest that robustness in the feature space translates to downstream robustness.

**Table 4.** Performance (AUC) and robustness (CCC) results on the BreastBM dataset.

Task	Model	Performance metric: AUC			Robustness metric: CCC		
		Clb	Igr	Curimeta	Clb	Igr	Curimeta
gBrca	H0-mini	<b>0.69</b> (0.02)	-	<b>0.69</b> (0.01)	<b>0.92</b> (0.03)	-	<b>0.88</b> (0.84-0.91)
	H0	0.66 (0.02)	-	0.64 (0.05)	0.42 (0.25)	-	0.71 (0.65-0.78)
ER	H0-mini	<b>0.95</b> (0.00)	<b>0.76</b> (0.02)	<b>0.88</b> (0.00)	0.97 (0.02)	<b>0.82</b> (0.08)	<b>0.98</b> (0.97-0.98)
	H0	0.95 (0.01)	0.75 (0.03)	<b>0.88</b> (0.00)	<b>0.97</b> (0.01)	0.72 (0.12)	0.90 (0.87-0.95)
PR	H0-mini	<b>0.89</b> (0.01)	0.87 (0.00)	<b>0.69</b> (0.01)	<b>0.96</b> (0.03)	<b>0.78</b> (0.10)	<b>0.98</b> (0.97-0.98)
	H0	0.88 (0.01)	<b>0.88</b> (0.02)	<b>0.69</b> (0.01)	0.92 (0.03)	0.63 (0.18)	0.94 (0.92-0.96)
HER2	H0-mini	-	<b>0.66</b> (0.02)	-	-	<b>0.74</b> (0.14)	-
	H0	-	0.65 (0.02)	-	-	0.64 (0.18)	-

## 4 Discussion

While the results presented in this work are promising, several limitations and perspectives can be noted. First, we note that the final projection head of the teacher model is required to perform the distillation. Whether the distillation can be performed without the projection head (for instance, by learning them during the pre-training as in [18]) with equivalent downstream performance and robustness should be explored. Second, even though distilled models will be much cheaper to use for downstream tasks (both in time, cost or carbon footprint), they remain computationally demanding to train compared to training a similar model from scratch (1.7x longer for a ViT-Base) as it requires the teacher model. However, the resulting performance and robustness is greatly improved.

Finally, we mention some open questions. The successful distillation of large models into smaller ones raises a question on the intrinsic dimension of FMs in digital pathology. Besides, we did not perform an extensive hyperparameter optimization. Whether some could be tailored for distillation remains to be explored (*e.g.*, adding more global crops, reducing the student patch size, etc.). Additionally, we proposed several metrics to assess the robustness of FMs, yet future work should be carried out to have a more comprehensive understanding of a model’s behavior [13]. Note that this may require high-quality multi-scanned, multi-stained datasets. We hope that releasing our processed version of PLISM-WSI will encourage further research in this direction.

## 5 Conclusion

This study shows that distilling a large foundation model results in a competitive smaller model. We show that it can efficiently be leveraged for downstream applications closing the performance gap with larger foundation models. By releasing



the model publicly, we hope our work will democratize the use of pathology foundation models for researchers with limited computational resources. Besides, we show that distilled models present an additional robustness property, with their features having a better invariance to changes in scanners or staining conditions. Self-supervised learning is known to be an efficient way to address domain generalization in CPath [26] and, to the best of our knowledge, this study is the first to highlight that distilling a large foundation model further improves the robustness of the resulting model. This enhanced robustness property lays the groundwork for future research, supporting the broader adoption of CPath models in clinical practice.

**Acknowledgments.** This work was granted access to the high performance Computing (HPC) resources of IDRIS under the allocations 2023-A0141012519 and 2024-A0161012519 made by GENCI. The results presented here are in part based upon data generated by the TCGA Research Network: <https://www.cancer.gov/tcga>. Results on the gBRCA dataset are based on data generated by Gustave Roussy (GR) and Centre Léon Bérard (CLB) through PortrAI (Owkin, Tribun Health, Cypath, GR, CLB, Unicancer), a French consortium advancing precision medicine. PortrAI was funded by the French government as part of the France 2030 program, and the European Union – NextGenerationEU through the France Relance program. An additional external data set from the United States was accessed through CuriMeta.

**Disclosure of Interests.** Persons affiliated with Owkin own stocks in the company (A.F., N.D., O.T., A.R., R.D., G.R., A.O.). Persons affiliated with Biopitimus own stocks in the company (T.P., D.V., M.S., C.S.).

## References

1. kaiko. ai, Aben, N., de Jong, E.D., Gatopoulos, I., Känzig, N., Karasikov, M., Lagré, A., Moser, R., van Doorn, J., Tang, F.: Towards large-scale training of pathology foundation models (2024)
2. Caron, M.: Emerging properties in self-supervised vision transformers (Apr 2021), <https://arxiv.org/abs/2104.14294>
3. Chen, R.J., Ding, T., Lu, M.Y., Williamson, D.F.K., Jaume, G., Song, A.H., Chen, B., Zhang, A., Shao, D., Shaban, M., Williams, M., Oldenburg, L., Weishaupt, L.L., Wang, J.J., Vaidya, A., Le, L.P., Gerber, G., Sahai, S., Williams, W., Mahmood, F.: Towards a general-purpose foundation model for computational pathology. *Nature Medicine* **30**(3), 850–862 (Mar 2024). <https://doi.org/10.1038/s41591-024-02857-3>
4. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale (arXiv:2010.11929) (Jun 2021), <http://arxiv.org/abs/2010.11929>, arXiv:2010.11929 [cs]
5. Duval, Q., Misra, I., Ballas, N.: A simple recipe for competitive low-compute self supervised vision models (2023). <https://doi.org/10.48550/ARXIV.2301.09451>
6. Filiot, A., Ghermi, R., Olivier, A., Jacob, P., Fidon, L., Kain, A.M., Saillard, C., Schiratti, J.B.: Scaling self-supervised learning for histopathology with masked

- image modeling p. 2023.07.21.23292757 (Jul 2023). <https://doi.org/10.1101/2023.07.21.23292757>
7. Filiot, A., Jacob, P., Mac Kain, A., Saillard, C.: Phikon-v2, a large and public feature extractor for biomarker prediction (2024). <https://doi.org/10.48550/ARXIV.2409.09173>
  8. Food, U., Administration, D.: Developing and labeling in vitro companion diagnostic devices for a specific group of oncology therapeutic products (Apr 2020), fDA-2018-D-3380
  9. Gatopoulos, I., Känzig, N., Moser, R., Otalora, S.: eva: Evaluation framework for pathology foundation models. Medical Imaging with Deep Learning (2024)
  10. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network (2015). <https://doi.org/10.48550/ARXIV.1503.02531>
  11. Huang, Z., Bianchi, F., Yuksekgonul, M., Montine, T.J., Zou, J.: A visual-language foundation model for pathology image analysis using medical twitter. *Nature Medicine* **29**(9), 2307–2316 (2023). <https://doi.org/10.1038/s41591-023-02504-3>
  12. Jaume, G., Doucet, P., Song, A.H., Lu, M.Y., Almagro-Pérez, C., Wagner, S.J., Vaidya, A.J., Chen, R.J., Williamson, D.F.K., Kim, A., Mahmood, F.: Hest-1k: A dataset for spatial transcriptomics and histology image analysis (arXiv:2406.16192) (Jun 2024), arXiv:2406.16192 [cs]
  13. de Jong, E.D., Marcus, E., Teuwen, J.: Current pathology foundation models are unrobust to medical center differences (2025), <https://arxiv.org/abs/2501.18055>
  14. Karasikov, M., van Doorn, J., Känzig, N., Cesur, M.E., Horlings, H.M., Berke, R., Tang, F., Otálora, S.: Training state-of-the-art pathology foundation models with orders of magnitude less data (2025), <https://arxiv.org/abs/2504.05186>
  15. Klein, S., Staring, M., Murphy, K., Viergever, M.A., Pluim, J.P.W.: elastix: a toolbox for intensity-based medical image registration. *IEEE Trans Med Imaging* **29**(1), 196–205 (Nov 2009)
  16. Lee, J., Lim, J., Byeon, K., Kwak, J.T.: Benchmarking pathology foundation models: Adaptation strategies and scenarios (2024). <https://doi.org/10.48550/ARXIV.2410.16038>
  17. Lu, M.Y., Chen, B., Williamson, D.F.K., Chen, R.J., Liang, I., Ding, T., Jaume, G., Odintsov, I., Le, L.P., Gerber, G., Parwani, A.V., Zhang, A., Mahmood, F.: A visual-language foundation model for computational pathology. *Nature Medicine* **30**(3), 863–874 (2024). <https://doi.org/10.1038/s41591-024-02856-4>
  18. Ma, J., Guo, Z., Zhou, F., Wang, Y., Xu, Y., Cai, Y., Zhu, Z., Jin, C., Lin, Y., Jiang, X., Han, A., Liang, L., Chan, R.C.K., Wang, J., Cheng, K.T., Chen, H.: Towards a generalizable pathology foundation model via unified knowledge distillation (2024). <https://doi.org/10.48550/ARXIV.2407.18449>
  19. Nechaev, D., Pchelnikov, A., Ivanova, E.: Hibou: A family of foundational vision transformers for pathology (2024). <https://doi.org/10.48550/ARXIV.2406.05074>
  20. Ochi, M., Komura, D., Onoyama, T., Shinbo, K., Endo, H., Odaka, H., Kakiuchi, M., Katoh, H., Ushiku, T., Ishikawa, S.: Registered multi-device/staining histology image dataset for domain-agnostic machine learning models. *Scientific Data* **11**(1), 330 (Apr 2024). <https://doi.org/10.1038/s41597-024-03122-5>
  21. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Assran, M., Ballas, N., Galuba, W., Howes, R., Huang, P.Y., Li, S.W., Misra, I., Rabbat, M., Sharma, V., Synnaeve,

- G., Xu, H., Jegou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.: Dinov2: Learning robust visual features without supervision (arXiv:2304.07193) (Apr 2023). <https://doi.org/10.48550/arXiv.2304.07193>, arXiv:2304.07193 [cs]
22. Saillard, C., Jenatton, R., Llinares-López, F., Mariet, Z., Cahané, D., Durand, E., Vert, J.P.: H-optimus-0 (2024), <https://github.com/bioptimus/releases/tree/main/models/h-optimus/v0>
  23. Song, A.H., Jaume, G., Williamson, D.F.K., Lu, M.Y., Vaidya, A., Miller, T.R., Mahmood, F.: Artificial intelligence for digital and computational pathology. *Nature Reviews Bioengineering* **1**(12), 930–949 (2023). <https://doi.org/10.1038/s44222-023-00096-8>
  24. Wölflein, G., Ferber, D., Meneghetti, A.R., Nahhas, O.S.M.E., Truhn, D., Carrero, Z.I., Harrison, D.J., Arandjelović, O., Kather, J.N.: Benchmarking pathology feature extractors for whole slide image classification (2023). <https://doi.org/10.48550/ARXIV.2311.11772>
  25. Xu, H., Usuyama, N., Bagga, J., Zhang, S., Rao, R., Naumann, T., Wong, C., Gero, Z., González, J., Gu, Y., Xu, Y., Wei, M., Wang, W., Ma, S., Wei, F., Yang, J., Li, C., Gao, J., Rosemon, J., Bower, T., Lee, S., Weerasinghe, R., Wright, B.J., Robicsek, A., Piening, B., Bifulco, C., Wang, S., Poon, H.: A whole-slide foundation model for digital pathology from real-world data. *Nature* **630**(8015), 181–188 (Jun 2024). <https://doi.org/10.1038/s41586-024-07441-w>
  26. Zamanitajeddin, N., Jahanifar, M., Xu, K., Siraj, F., Rajpoot, N.: Benchmarking domain generalization algorithms in computational pathology (arXiv:2409.17063) (Sep 2024), <http://arxiv.org/abs/2409.17063>, arXiv:2409.17063 [cs]
  27. Zhou, J., Wei, C., Wang, H., Shen, W., Xie, C., Yuille, A., Kong, T.: ibot: Image bert pre-training with online tokenizer (arXiv:2111.07832) (Jan 2022). <https://doi.org/10.48550/arXiv.2111.07832>, arXiv:2111.07832 [cs]
  28. Zimmermann, E., Vorontsov, E., Viret, J., Casson, A., Zelechowski, M., Shaikovski, G., Tenenholtz, N., Hall, J., Klimstra, D., Yousfi, R., Fuchs, T., Fusi, N., Liu, S., Severson, K.: Virchow2: Scaling self-supervised mixed magnification models in pathology (2024). <https://doi.org/10.48550/ARXIV.2408.00738>