

PLUTO-4: Frontier Pathology Foundation Models

Harshith Padigela¹ Shima Nofallah² Atchuth Naveen Chilaparasetti² Ryun Han² Andrew Walker³
 Judy Shen³ Chintan Shah³ Blake Martin³ Aashish Sood³ Elliot Miller³ Ben Glass⁴ Andy Beck⁴
 Harsha Pokkalla⁴ Syed Ashar Javed⁴
 PathAI

Abstract

Foundation models trained on large-scale pathology image corpora have demonstrated strong transfer capabilities across diverse histopathology tasks. Building on this progress, we introduce **PLUTO-4**, our next generation of pathology foundation models that extend the Pathology-Universal Transformer (PLUTO) to frontier scale. We share two complementary Vision Transformer architectures in the PLUTO-4 family: a compact and efficient **PLUTO-4S** model optimized for multi-scale deployment using a FlexiViT setup with 2D-RoPE embeddings, and a frontier-scale **PLUTO-4G** model trained with a single patch size to maximize representation capacity and stability. Both models are pretrained using a self-supervised objective derived from DINOv2 on a large multi-institutional corpus containing 551,164 WSIs from 137,144 patients across over 50 institutions, spanning over 60 disease types and over 100 stains. Comprehensive evaluation across public and internal benchmarks demonstrates that PLUTO-4 achieves state-of-the-art performance on tasks requiring varying spatial and biological context, including patch-level classification, segmentation, and slide-level diagnosis. The compact PLUTO-4S provides high-throughput and robust performance for practical deployment, while PLUTO-4G establishes new performance frontiers across multiple pathology benchmarks, including 11% improvement in dermatopathology diagnosis. These diverse improvements underscore PLUTO-4’s potential to transform real-world applications as a backbone for translational research and diagnostic use cases.

1. Introduction

Pathology is the study of microscopic tissue morphology and remains the clinical gold standard for diagnosing disease. The digitization of histopathology slides into whole slide images (WSIs) has enabled large-scale quantitative analysis and the development of computational pathology systems that aim to assist pathologists in diagnosis, prognosis, and biomarker assessment (Campanella et al., 2019; Bulten et al., 2020). However, the intrinsic complexity of WSIs — including gigapixel scale, variations in staining protocols and scanner systems, and heterogeneous biological content — poses unique challenges for machine learning methods.

Recent advances in self-supervised learning have shown that *foundation models* (FMs), trained on large and diverse image corpora, can generate transferable visual representations across a wide range of tasks (Oquab et al., 2024; Caron et al., 2021). In pathology, several models such as PLUTO (Juyal et al., 2024), H-Optimus (Saillard et al., 2024), Virchow2 (Zimmermann et al., 2024), Atlas (Alber et al., 2025), RudolfV (Dippel et al., 2024), and UNI (Chen et al., 2024) have demonstrated the potential of this paradigm, providing encoders that generalize across tissue types, stains, and magnifications. Despite these advances, scaling pathology FMs remains constrained by three key factors: (1) limited availability of large and heterogeneous training datasets, (2) training instability and compute bottlenecks when training large scale models, and (3) high computational demands that limit deployment in diagnostic workflows at scale.

To address these challenges, we developed **PLUTO-4**, a new generation of pathology foundation models designed for both scalability, efficiency and performance. PLUTO-4 introduces two complementary encoders:

- **PLUTO-4S**, a compact, high-throughput model incorporating a **FlexiViT** backbone (Beyer et al., 2023) with Rotary Positional Embeddings (**RoPE**) (Su et al., 2021) for robust performance and deployment at scale.
- **PLUTO-4G**, a frontier-scale model trained with a single patch-token size and positional embedding scheme,

¹Project Lead and First Author ²Core Contributor ³Contributor
⁴Senior Contributor. Correspondence to: Harshith Padigela
 <harshith.padigela@pathai.com>.

Table 1. Performance comparison of PLUTO-4 with existing pathology foundation models.

PLUTO-4G achieves best-in-class performance for the majority of benchmarks and task categories.

Dataset / Metric	PLUTO-4G	H-Optimus-0	Atlas	Virchow-2	UNI2-H	Prov-Gigapath	Lunit-S	H-Optimus-1
<i>Spatial Transcriptomics</i> (Pearson r)								
HEST	0.427	0.413	0.399	0.396	0.414	0.386	–	0.422
<i>Patch-level Classification; EVA</i> (Balanced Accuracy %)								
MHIST	87.5 (0.3)	84.3	85.2	86.1	82.4	82.9	78.1	83.5
BreakHIS	81.5 (0.4)	80.1	–	82.1	85.9	82.7	74.2	–
BACH	93.8 (0.5)	75.9	93.1	88.3	91.5	75.9	78.3	–
Gleason (Arvaniti)	79.3 (0.7)	77.0	–	78.3	77.5	72.4	75.0	–
PCAM (test)	95.1 (0.1)	94.3	94.9	93.8	95.0	94.5	89.7	–
CRC	96.4 (0.2)	95.5	97.0	96.7	96.5	95.1	94.0	–
<i>Slide-level Classification</i> (Balanced Accuracy %)								
PANDA-Small (test)	66.8 (1.6)	67.1	70.0	64.6	65.7	65.3	61.0	–
Derm 2K* (Macro F1 %)	67.1	62.8	–	–	–	–	–	–
<i>Nuclear Segmentation; EVA</i> DICE								
MoNuSAC	70.4 (0.3)	68.5	–	66.9	64.2	68.0	62.9	–
CoNSep	65.0 (0.1)	64.4	–	64.0	63.0	62.6	60.2	–

Best score per dataset in **bold**. The std across multiple runs is shown in parantheses. Results on EVA for external models are reported from their leaderboard or the model’s paper. HEST results for external models are from (Bioptimus, 2025) or the model’s paper.

*Derm-2K is a proprietary dataset.

designed to maximize representation capacity and model performance.

Both models were trained on a large-scale, multi-institutional corpus capturing the broad morphological and technical diversity across disease areas, staining modalities, and scanner systems, and which represents one of the most comprehensive collections used for pathology foundation model training.

The training methodology builds upon the original PLUTO workflow (Juyal et al., 2024), maintaining a self-supervised pretraining framework derived from DINOv2 (Oquab et al., 2024) with multi-resolution image sampling. PLUTO-4 extends this framework with improved distributed training stability, higher-precision optimization, architectural improvements and extended compute scaling, enabling consistent convergence for both compact and frontier-scale architectures.

Through extensive benchmarking in patch, slide and segmentation-level tasks, PLUTO-4G shows state-of-the-art performance in multiple public benchmarks, while PLUTO-4S provides a strong and computationally efficient alternative suitable for real-world deployment. Together, these models advance the goal of creating general-purpose pathology foundation models capable of supporting a broad range of diagnostic and research applications.

2. Training Data

2.1. Dataset Overview

The **PLUTO-4** dataset represents one of the largest and most diverse corpora assembled for pathology foundation model training to date. It comprises approximately **551,164** de-identified whole slide images (WSIs) drawn from 137,144 patients across more than 50 contributing institutions, spanning a broad spectrum of disease areas, organ systems, stains, and scanners. In total, the dataset includes slides from more than 40 distinct tissue and organ categories spanning over **60 diseases** covering over **100 stain variants**. Slides were digitized on multiple scanner systems encompassing more than ten hardware models. The dataset reflects both diagnostic and research use cases, incorporating routine hematoxylin and eosin (H&E) stains, immunohistochemistry (IHC) panels, and special stains from frozen and formalin-fixed paraffin-embedded (FFPE) preparations.

The large-scale diversity of PLUTO-4 is a critical enabler of robust and generalizable pathology foundation models. By integrating data across multiple healthcare systems and scanner vendors, the dataset captures a wide distribution of staining, scanning, and morphological variability encountered in clinical practice.

PLUTO-4: Frontier Pathology Foundation Models




Figure 1. Distribution of the PLUTO-4 dataset across organs, diseases, stain groups and scanners.

2.2. Stain and Modality Diversity

As shown in Figure 1 (bottom left), the dataset is anchored by a large cohort of approximately 396K H&E stained WSIs, which is complemented by a diverse collection of special stains and Immunohistochemistry (IHC) slides. This IHC subset features over 50 unique targets, encompasses membranous, cytoplasmic, and nuclear staining patterns, and includes a wide range of biomarkers (e.g., PD-L1, HER2, Ki-67, CD8, CD20, AE1/AE3, ER, PR) from common to less frequent ones used in hematopathology and research applications. Special stains (e.g., Masson’s trichrome, PAS, reticulin, iron) and frozen-section preparations provide additional morphological variety beyond routine diagnostic material. This stain-level diversity supports the development of encoders capable of recognizing cellular and structural features under heterogeneous imaging conditions.

2.3. Scanner and Site Heterogeneity

Slides were digitized using a variety of whole-slide scanners, including systems from Aperio, Philips, Ventana and Hamamatsu. Figure 1 (bottom right) illustrates the broad vendor representation, with the largest contributions from Aperio and Philips scanners. The dataset also includes a subset of slides converted from generic TIFF formats, representing legacy or research-origin data. This diversity of scanner hardware and color calibration profiles introduces natural variation in image tone, contrast, and sharpness, which can improve feature generalization in large-scale representation learning.

2.4. Organ and Disease Coverage

The dataset has broad coverage across oncology and non-oncology domains, spanning over 40 organs. High-volume tissues such as Colon, Skin, Prostate, Breast, and Liver collectively contribute about 240K WSIs, with substantial representation from additional gastrointestinal, respiratory, and lymphoid organs, among others (Figure 1, top). This broad anatomic coverage ensures inclusion of both common and rare organ systems encountered in routine clinical practice.

Across these organs, the dataset encompasses more than **60 disease entities**, spanning malignant, benign, inflammatory, and normal tissue categories. Frequent malignancies include *Invasive Breast Carcinoma*, *Prostate Adenocarcinoma*, *NSCLC Adenocarcinoma*, *Colorectal Adenocarcinoma*, *Melanoma*, and *Urothelial Carcinoma*, while non-neoplastic conditions such as *Ulcerative and Microscopic Colitis*, *Gastritis*, and *Esophagitis* are also represented. Benign and pre-cancerous lesions (*Adenomas*, *Keratosis*), inflammatory disorders, and unremarkable tissues further expand diagnostic diversity.

At the disease-type level (Figure 1, second row), malignant neoplasms account for approximately 212K WSIs, inflammatory gastrointestinal and dermatologic disorders for 86K WSIs, and benign or pre-cancerous conditions for 72K WSIs, with the remainder corresponding to normal or unclassified tissue. This composition provides balanced representation across major histologic categories, capturing variations in staining, scanning, and disease morphology observed in clinical practice.

2.5. Data Preprocessing and Sampling

To ensure consistent quality across magnifications, we applied a multi-stage preprocessing pipeline. Usable tissue regions were extracted using the latest version of our Artifact Detect algorithm (Le et al., 2025), which segments usable tissue from background and excludes artifact regions like pen markings, folds, or scanning blur. We sampled 165M regions of usable tissue from the slides and generated around 640M image tiles across 4 magnifications (0.25, 0.5, 1.0, and 2.0 $\mu\text{m}/\text{px}$) for training with a mix of tile sizes 275px, 550px which are cropped to the global and local crops views during training.

3. Training Methodology

3.1. Training Architecture Design

The training architecture for PLUTO-4 extends the self-supervised framework introduced in the original PLUTO (Juyal et al., 2024), while incorporating architectural and optimization refinements that enable stable scaling from compact to frontier-sized vision transformers.

PLUTO-4S. The compact encoder PLUTO-4S is designed to support multiple levels of representational granularity across diverse pathology tasks. Different problems, such as cell segmentation versus slide-level classification, require varying receptive field sizes and granular context to capture relevant morphologic context. To accommodate this, we adopt the **FlexiViT** setup (Beyer et al., 2023) similar to PLUTO (Juyal et al., 2024), allowing the same backbone to operate on variable patch-token sizes. During pretraining, patch-token sizes from [8, 16, 32] are sampled dynamically across training iterations, exposing the model to multi-scale input structure without separate backbone training. Unlike the absolute positional embeddings used in PLUTO (Juyal et al., 2024), PLUTO-4S employs two-dimensional rotary positional embeddings (**2D-RoPE**) (Su et al., 2021), which provide relative positional encoding in both spatial dimensions and improve stability at large sequence lengths. This modification is consistent with the positional encoding formulation used in (Siméoni et al., 2025) and other large-scale vision transformers.

PLUTO-4G. For the frontier-scale model, PLUTO-4G, we observed that larger embeddings inherently possess sufficient representational capacity, and the difference between various patch-token-sizes diminishes as model scale increases. Empirically, we saw the difference in downstream performance between patch-sizes 8, 16 reduce across tasks as model size grows, while the computational cost and memory scales quadratically when patch-token size halves (or sequence length doubles). An example is seen in Figure 2, where we see ViT-G-14 is almost 3.5X faster than ViT-G-8. Therefore, we train PLUTO-4G using a single patch size (14), a choice consistent with other FMs and significantly reduces our training cost while improving throughput.

3.2. Stabilizing Self-Supervised ViT Training

Self-Supervised training of Vision transformers with DINOv2 is unstable and challenging. The iBOT and DINO projection heads are prone to large activations and overflows, and training suffers from noisy losses, large gradient norms, and many of these become prominent as the model size grows. Similar observations were also noted by (Zimmermann et al., 2024). We have found the following training choices helpful in stabilizing training in addition to the recommendations in DINOv2 like gradient clipping.

- **bfloat16-mixed precision:** Even with gradient clipping we found activations in the projection heads grow and overflow in float16 which has a limited range of $[-65504, 65504]$ leading to nan losses in training. bfloat16 trades off precision for range and has a much larger range of $[-3.39 \times 10^{38}, 3.39 \times 10^{38}]$ similar to float32. While float16 may still be critical in other workloads where precision is critical like RL training (Qi et al., 2025) and inference, for DINOv2 we found range more crucial than precision and thus use bfloat16 for model forward and backward passes. Key computations which need higher precision like loss computation, momentum center updates are performed in float32.
- **Adding register tokens:** We add in register tokens (Darcet et al., 2024) which can capture any high norm tokens as the model size increases and better capture spatial features.
- **Using large batch-sizes:** We found it helpful to have batch sizes of at least 1024 to ensure stable losses and gradients.

To ensure stable training, we use a local batch size of 32 tiles per GPU, corresponding to an effective global batch size of 1024 for PLUTO-4G. For PLUTO-4S we use an effective batch-size of 1536. We employ the AdamW optimizer with a fixed weight decay and linear warm-up of




Figure 2. Training throughput scaling across architectures and hardware. ViT-B shows near-linear scaling across both A40 and H200 clusters, while ViT-G throughput degrades beyond two nodes due to communication bottlenecks in DDP. Additionally we can also see, ViT-G with patch size 8 is approximately 3.5 $\times$ slower than ViT-G with patch size 14.




Figure 3. Optimizing ViT-G training throughput. Enabling GPUDirect RDMA and tuning DDP parameters (bucket_cap_mb, gradient_as_bucket_view) restores near-linear throughput scaling and saturates InfiniBand bandwidth.

learning rate followed by a cosine annealing decay for both models. Temperature and momentum schedules follow a similar approach to DINOv2, though we adjusted epoch lengths to ensure good coverage of the dataset. Gradient clipping is applied at a global norm of 3.0 and we use 4 register tokens for PLUTO-4G.

3.3. Scaling Model Training

We systematically analyze scaling efficiency when training PLUTO models across multiple GPUs and nodes. Training is performed on four NVIDIA H200 nodes, each containing eight GPUs interconnected via NVLink for intra-node communication and InfiniBand with GPUDirect RDMA for inter-node communication. Figure 2 summarizes the observed throughput trends for ViT-B and ViT-G architectures under varying hardware and node configurations.

For smaller architectures such as ViT-B, throughput scales

almost linearly with the number of nodes on both A40 and H200 clusters. However, for larger architectures such as ViT-G, throughput drops significantly when going beyond one node, with overall throughput even decreasing at higher GPU counts. The main bottleneck arises from all-reduce synchronization during Distributed Data Parallel (DDP) training. As model size grows, the gradient tensors exchanged between GPUs increase substantially, causing network congestion and underutilization of available compute resources. This effect is particularly evident in models with large sizes (ViT-G) where per-layer gradient aggregation dominates total step time.

Mitigating cross-node bottlenecks. To alleviate these communication bottlenecks, we enabled GPUDirect RDMA to allow direct GPU-to-GPU data transfer across nodes, bypassing CPU memory and reducing PCIe latency. While this provided an immediate improvement, scaling efficiency remained suboptimal and bandwidth utilization was still low, suggesting that the DDP communication pipeline itself required further optimization.

We conducted an extensive sweep of DDP parameters to tune gradient synchronization behavior. Two adjustments proved critical: increasing `bucket_cap_mb` from the default 25 MB to 360 MB, and enabling `gradient_as_bucket_view=True`. The larger communication bucket size reduces the number of all-reduce calls per step, while `gradient_as_bucket_view` minimizes tensor copies during gradient aggregation. Combined with RDMA, these modifications significantly improved communication-computation overlap, achieving near-linear throughput scaling with the number of GPUs and full utilization of the InfiniBand fabric (Fig. 3).

FSDP vs. DDP trade-offs. We also evaluated Fully Sharded Data Parallel (FSDP) as an alternative for scaling large models. FSDP shards parameters, gradients, and optimizer states across GPUs, reducing memory and communication per device. This approach is advantageous when models exceed single-GPU VRAM capacity or when network bandwidth is constrained. However, when bandwidth is ample and models fit in GPU memory, the additional sharding and all-gather overhead reduces throughput. On H200, tuned DDP achieved roughly twice the throughput of FSDP for ViT-G. FSDP could be useful for future PLUTO variants where model scale or interconnect limits make full replication infeasible.

4. Results and Evaluation

4.1. Evaluation Methodology

We evaluate the PLUTO-4 models on a broad set of public and proprietary benchmarks spanning multiple biological

Table 2. Performance results using CLS + Mean(Patch-Token) concatenated embeddings. PLUTO-4G maintains state-of-the-art performance across most datasets, demonstrating strong generalization with CLS + Patch-Token embeddings.

Dataset	PLUTO-4G	H-Optimus-0	Virchow-2	Atlas
HEST	0.432	0.424	0.398	0.421
MHIST	87.9	85.0	85.9	86.4
BACH	93.2	74.2	88.7	92.5
PCAM (Test)	95.2	94.3	93.9	94.8
CRC	96.8	96.2	96.7	97.1
PANDA-Small (Test)	66.6	68.0	66.4	70.5

Best result per dataset in **bold**. Results for H-Optimus-0, Virchow-2, and Atlas are from the Table 5 in (Alber et al., 2025). The metric for HEST is pearson r and the metric for rest of the datasets is balanced accuracy.

scales, including patch-level classification and regression, segmentation, and slide-level prediction. The evaluations were selected to assess the generality of the learned representations across tasks requiring different contextual granularity, ranging from subcellular detail to whole-slide inference.

Evaluation Framework. We leverage the open-source EVA benchmarking framework (kaiko.ai et al., 2024) for standardized evaluation of pathology foundation models. EVA provides unified loaders, preprocessing, and evaluation metrics across a wide collection of pathology datasets, allowing consistent comparison to publicly available FMs such as Virchow2 (Zimmermann et al., 2024), UNI2-H (Chen et al., 2024), Prov-Gigapath (Xu et al., 2024), Lunit (Kang et al., 2023) and Atlas (Alber et al., 2025). For HEST we use the benchmarking setup and code as described in (Jaume et al., 2024). Details regarding specific datasets are covered in subsequent section.

Embedding Selection. Unless otherwise specified, all reported results use the **CLS token embedding** extracted from the frozen encoder as the image representation. This design enables fair comparison to prior works that primarily rely on CLS embeddings for linear probing and downstream task evaluation. Results using the concatenation of CLS, Mean of patch-token embeddings for classification tasks are provided in the table 2 for completeness and to quantify the effect of incorporating spatial context.

4.2. Evaluation Datasets and Reporting Protocols

We evaluate PLUTO-4 models across four broad task categories representing different biological and spatial scales in pathology: patch-level classification, segmentation, spatial transcriptomics, and slide-level prediction. For all benchmarks, we follow standard linear-probing or ABMIL protocols as implemented in EVA (kaiko.ai et al., 2024) and HEST (Jaume et al., 2024).

Patch-level Classification. We evaluate on MHIST, BreakHis, BACH, and Gleason benchmarks from the EVA suite, which span epithelial tissue classification and breast cancer subtyping tasks requiring moderate contextual resolution. Metrics are reported as Balanced Accuracy (BA) averaged over 5 runs using the EVA evaluation protocol. For comparison with external models, results for H-Optimus-0, Virchow-2, UNI2-H, Prov-Gigapath, and Lunit are taken directly from the EVA leaderboard, while Atlas and H-Optimus-1 results are taken from their respective publications (Alber et al., 2025; Bioptimus, 2025).

Segmentation. We benchmark on CoNSep and MoNuSAC datasets for instance-level nuclei segmentation and classification across organs and stains. We use the *ConvDecoderWithImage* decoder from EVA which upsamples the input feature maps before concatenating them with the original input image, followed by a series of convolutional layers. Performance is reported as the MONAI-DICE score averaged over 5 EVA evaluation runs. Results for H-Optimus-0, Virchow-2, UNI2-H, Prov-Gigapath, and Lunit are taken from the EVA leaderboard. Segmentation metrics for Atlas and H-Optimus-1 are not reported in their respective works.

Spatial Transcriptomics (HEST). We use the HEST-1k benchmark (Jaume et al., 2024) to evaluate spatially-resolved gene expression prediction from H&E morphology across nine tissue types. We follow the official Ridge Regression with PCA protocol and report the average Pearson correlation across all tissue types. Results for all external foundation models are drawn from the unified comparison presented in the H-Optimus-1 blog (Bioptimus, 2025), and results for Atlas are cited from its publication (Alber et al., 2025). Indication-wise results are provided in Table 5.

Slide-level Prediction. Two slide-level tasks are used to assess global contextual reasoning:

- (i) **PANDA-Small:** A subset of the PANDA challenge, for predicting prostate cancer Gleason grades, using the EVA ABMIL protocol. We report **Balanced Accuracy** averaged over 20 runs using EVA.
- (ii) **Derm-2K:** A proprietary dermatopathology diagnosis benchmark comprising over 2,000 whole-slide images across 17 skin lesion categories. The dataset covers major lesion types, including actinic keratosis, basal cell carcinoma, benign nevus, cyst, dermatitis, dysplastic nevus, invasive melanoma, lichenoid keratosis, melanoma in situ, scar, seborrheic keratosis, squamous cell carcinoma, squamous cell carcinoma in situ, vascular lesion, verruca vulgaris, other benign non-melanocytic lesions, and normal skin. We evaluate

models using our internal AB-MIL setup and report **Macro-F1** score.

For PANDA-Small, results for H-Optimus-0, Virchow-2, UNI2-H, Prov-Gigapath, and Lunit are taken from the EVA leaderboard; Atlas results are reported from (Alber et al., 2025). For Derm-2K, results are compared between PLUTO-4G and H-Optimus-0.

4.3. Benchmark Results

4.3.1. PATCH-LEVEL CLASSIFICATION

PLUTO-4G demonstrates strong generalization on a diverse set of patch-level classification tasks (Table 1). On MHIST, PCAM-test and BACH, **PLUTO-4G** achieves the highest balanced accuracy outperforming other comparable models like Virchow-2, Atlas, H-Optimus-0 and H-Optimus-1. For fine-grained glandular structure classification (Gleason Arvaniti), PLUTO-4G attains the highest accuracy, indicating its ability to model large-scale context. Across general-purpose benchmarks such as BreakHis and CRC, PLUTO-4G achieves top-tier accuracy, establishing a new performance standard for frozen representation evaluation.

4.3.2. SEGMENTATION

On dense nuclear segmentation and classification tasks (MoNuSAC and CoNSep), PLUTO-4G attains the highest Dice and AJI scores (70.4 and 65.0, respectively), outperforming all previous pathology foundation models, including Virchow-2 (0.669 / 0.640) and UNI2-H (64.2 / 63.0). These gains demonstrate that self-supervised representations from PLUTO-4 maintain spatial coherence, enabling effective adaptation for dense instance segmentation.

4.3.3. SPATIAL TRANSCRIPTOMICS (HEST)

The HEST benchmark evaluates morphological correlates of spatially resolved gene expression across nine tumor types. As shown in Table 5, PLUTO-4G achieves the highest mean Pearson correlation ($r = 0.427$), surpassing all prior models including H-Optimus-1 (0.422), UNI2-H (0.413), Virchow-2 (0.396), and Atlas (0.399). PLUTO-4G shows particularly strong improvements in the clear cell renal cell carcinoma (CCRCC), pancreatic ductal adenocarcinoma (PAAD), and skin cutaneous melanoma (SKCM) datasets, indicating enhanced sensitivity to morpho-molecular associations across distinct tumor morphologies.

4.3.4. SLIDE-LEVEL PREDICTION

We evaluate PLUTO-4G on two slide-level tasks requiring aggregation of regional features into whole-slide predictions.

On the PANDA-Small Gleason grading benchmark,

PLUTO-4G demonstrates strong performance, surpassing models such as Virchow-2, UNI2-H, and Prov-Gigapath, performing on par with H-Optimus-0, and slightly below Atlas.(Table 1).

For Derm-2K, we evaluate PLUTO-4G and H-Optimus-0 on a proprietary 17-class dataset of 2K WSIs covering benign, inflammatory, and malignant skin diseases. PLUTO-4G achieves a macro-F1 of 67.1, surpassing H-Optimus-0 (0.628).

These results demonstrate robust transfer of learned morphological features to slide-level prediction tasks.

4.3.5. CLS + MEAN(PATCH-TOKEN) CONCATENATION

To ensure completeness and comparability with prior works, we additionally report results using CLS + Mean(Patch-Token) embeddings (Table 2). This evaluation mirrors the setup used in Virchow2 and Atlas. Since not all prior FM papers reported these evaluations, we use results reported in Atlas for our comparison.

Across eight benchmark datasets, PLUTO-4G achieves or matches state-of-the-art performance on nearly every task, outperforming or equaling other external models such as Atlas, Virchow-2, and H-Optimus-0. Specifically, PLUTO-4G attains the highest correlation on HEST (0.432) and top accuracy on MHIST (87.9), BACH (93.2), and PCAM (95.2). For the slide-level Gleason grading PANDA-Small, PLUTO-4G achieves comparable performance to other models, and slightly below Atlas. Overall, these results demonstrate that PLUTO-4G’s representations remain highly competitive even when evaluated under alternative embedding strategies used by other pathology foundation models.

4.4. Comparison of PLUTO Models

We next analyze the impact of **patch size** and **embedding type** across PLUTO variants (Table 3). This comparison highlights differences between the smaller PLUTO-4S configurations (trained with multiple patch sizes using FlexiViT) and the larger PLUTO-4G model (trained with a single patch size of 14).

Effect of patch size (ps = 8 vs ps = 16). Within the ViT-S family, performance differences across patch sizes are modest but consistent. Models trained with smaller patch tokens (ps = 8) slightly outperform ps = 16 on fine-grained, texture-rich datasets such as MHIST, BreakHIS, and BACH, with both CLS, CLS+Mean(Patch-Token) settings, reflecting improved sensitivity to local morphological cues. Conversely, ps = 16 performs on par or better for larger-context tasks such as Gleason, CRC and PCAM. The gap however increases in nuclear segmentation tasks CoN-Sep and MoNuSAC, where smaller patch-sizes can better

capture granular local features with PLUTO-4S-8 performing comparably or better than many larger external FMs like Virchow-2, UNI-2H and closing the gap with PLUTO-4G. These results confirm that flexible patch-token training effectively balances global and local context while providing flexibility to select the best patch-token size for modelling depending on the task.

Scaling to ViT-G. The larger PLUTO-4G model delivers substantial gains across all task categories. It surpasses both ViT-S configurations on all patch-level benchmarks and achieves the best results on segmentation, with Dice/AJI scores of 70.4/65.0, indicating stronger spatial consistency in feature maps. This scaling trend mirrors the improvements observed in other large-scale vision transformers, where increased capacity enhances representational generality across tasks.

CLS + Mean(Patch-Token) Concatenation. To assess whether incorporating broader spatial context improves classification performance, we evaluate concatenated embeddings combining the CLS token with the mean of all patch-tokens (Table 3). This formulation allows the model to jointly leverage global and local contextual features during downstream evaluation.

For PLUTO-4G, we observe moderate gains on several classification tasks, including HEST (+0.005 r), MHIST (+0.4 pp), PCAM, and CRC (+0.4 pp), while datasets such as BACH and Gleason show marginal declines. Across model variants, the improvements are not uniform but indicate a consistent trend: performance increases on 7 of 8 datasets for PLUTO-4S-8, 5 of 8 for PLUTO-4S-16, and 5 of 8 for PLUTO-4G. These results suggest that incorporating mean patch-token information can enhance discriminative ability for certain tasks, particularly those requiring finer-grained morphological context.

For segmentation, CLS+Mean Patch-Token features is not directly applicable, as downstream decoders directly operate on the entire Patch-Token features.

Overall, the addition of patch-level context yields modest but measurable benefits for classification tasks, particularly in smaller ViT-S models, while larger ViT-G models appear less sensitive to the choice of embedding aggregation strategy.

4.5. Case Study: Impact on PathAssist Derm Product

PLUTO-4 represents a significant leap in performance from our prior versions and we see these improvements across a range of our product offerings. As an example, we present results on our proprietary dermatopathology diagnosis dataset, comprising 17 lesion classes. The prior version of PLUTO corresponds to the model deployed in our PathAs-

Table 3. **Performance results for PLUTO-4: CLS vs. CLS + Mean(Patch-Token) across patch sizes and models.** ViT-S models trained with FlexiViT achieve strong performance while allowing configuration of the best patch-token size [8 / 16] for the task. Patch-Token size 8 significantly improves 4S performance on tasks needing fine-grained context. Using CLS + Patch embeddings yields modest improvements across most classification tasks.

Dataset / Metric	CLS token only			CLS + Mean(Patch-Token)		
	4S-8	4S-16	4G	4S-8	4S-16	4G
<i>Spatial Transcriptomics</i>						
<i>(Pearson r)</i>						
HEST	0.365	0.362	0.427	0.369	0.364	0.432
<i>Patch-level Classification</i>						
<i>(Balanced Accuracy %)</i>						
MHIST	83.7	83.5	87.5	84.2	83.4	87.9
BreakHis	81.3	79.2	81.5	80.8	76.8	81.8
BACH	82.7	79.6	93.8	85.1	79.8	93.2
Gleason (Arvaniti)	76.2	76.3	79.3	76.3	76.6	78.5
PCAM (Test)	90.7	90.9	95.1	91.4	91.6	95.2
CRC	95.2	95.0	96.4	95.4	95.5	96.8
<i>Slide-level Classification</i>						
<i>(Balanced Accuracy %)</i>						
PANDA-Small (Test)	61.8	61.5	66.8	63.1	63.0	66.6
<i>Nuclear Segmentation (DICE)</i>						
MoNuSAC	67.8	64.0	70.4	—	—	—
CoNSep	64.9	62.1	65.0	—	—	—

Benchmark datasets are setup and evaluated using EVA, with the exception of HEST. For segmentation datasets, CLS + Mean(Patch-Token) features are not applicable, as segmentation decoders utilize the entire set of Patch-Token features rather than pooled representations.

Table 4. **Derm-2K Slide-Level Prediction results.** PLUTO-4G achieves a 11% relative improvement in macro-F1 over the previous PLUTO-3 series, with PLUTO-4S yielding strong improvements

Model	Architecture / Size	Macro F1
PLUTO-3S-16	ViT-S (22M)	0.606
PLUTO-4S-16	ViT-S (22M)	0.628
PLUTO-4G	ViT-G (1.1B)	0.671
H-Optimus-0	ViT-G (1.1B)	0.628

sist Derm product (Billaut et al., 2025). As shown in Table 4, the new PLUTO-4 models achieve marked improvements, with PLUTO-4G delivering an 11% performance boost and PLUTO-4S achieving competitive results comparable to H-Optimus-0, despite being substantially smaller in scale.

5. Discussion

The PLUTO-4 series demonstrates the next step in scaling and adapting foundation models for digital pathology. Our results establish **PLUTO-4G** as a new performance standard across diverse pathology benchmarks, achieving state-of-the-art results in patch-level, segmentation, and slide-level prediction tasks. Trained at scale with a large, diverse corpus of histopathology data, PLUTO-4G exhibits

strong generalization across datasets and clinical contexts, capturing both cellular- and tissue-level morphology within a unified representation. In parallel, **PLUTO-4S** provides an efficient and versatile alternative—offering competitive performance with substantially reduced compute requirements and configurable patch-token sizes that make it well suited for a wide range of downstream problems and deployment settings.

Foundation models like PLUTO-4 represent a powerful substrate for computational pathology, providing general-purpose visual representations which are useful for a wide-variety of downstream tasks. However, they are only one part of an end-to-end computational pathology system. Building useful real-world applications requires additional task-specific adaptation layers that translate general embeddings into downstream applications. Developing robust and interpretable adapters remains a key step toward making these models practically deployable.

There also exists an inherent trade-off between deployability and performance. Larger architectures such as PLUTO-4G deliver the higher accuracy and generalization but require greater computational resources. Smaller models such as PLUTO-4S offer faster inference, lower memory footprint, and easier fine-tuning while maintaining competitive accuracy. Together, the two variants define a scalable family of models that balance performance and accessibility,

Table 5. HEST Benchmark: Spatial Transcriptomics Prediction Across Tumor Types.

Dataset	PLUTO-4G	H-Optimus-1	H-Optimus-0	UNI2-H	Atlas	Virchow-2
CCRCC	0.289 (0.042)	0.245 (0.125)	0.255 (0.135)	0.261 (0.132)	0.278 (0.036)	0.257 (0.123)
COAD	0.316 (0.016)	0.320 (0.016)	0.309 (0.000)	0.301 (0.004)	0.259 (0.031)	0.259 (0.016)
IDC	0.606 (0.085)	0.602 (0.081)	0.598 (0.085)	0.590 (0.081)	0.596 (0.081)	0.592 (0.080)
LUNG	0.569 (0.028)	0.578 (0.012)	0.559 (0.032)	0.558 (0.014)	0.570 (0.017)	0.553 (0.017)
LYMPH	0.273 (0.048)	0.277 (0.039)	0.259 (0.040)	0.272 (0.040)	0.257 (0.047)	0.255 (0.026)
PAAD	0.511 (0.049)	0.496 (0.051)	0.491 (0.040)	0.500 (0.040)	0.507 (0.072)	0.472 (0.065)
PRAD	0.374 (0.034)	0.378 (0.012)	0.385 (0.000)	0.357 (0.049)	0.353 (0.032)	0.348 (0.031)
READ	0.233 (0.033)	0.242 (0.015)	0.222 (0.048)	0.223 (0.038)	0.213 (0.029)	0.209 (0.050)
SKCM	0.670 (0.045)	0.659 (0.048)	0.645 (0.062)	0.659 (0.017)	0.562 (0.050)	0.619 (0.028)
AVG	0.427	0.422	0.413	0.413	0.399	0.396

Performance reported as mean Pearson correlation (r) across folds with standard deviation in parentheses. Bold values denote best performance per tumor type. Numbers for Atlas are extracted from their publication (Alber et al., 2025) and other external models are extracted from (Bioptimus, 2025)

supporting both high-throughput research applications and real-world deployment.

Beyond research benchmarks, foundation models such as PLUTO-4 can directly elevate upstream performance across PathAI’s diagnostic and biopharma product lines.¹ When integrated into products such as Explore product line for biomarker discovery (e.g., PathExplore, IHCExplore), Detect and Assist product line for workflow efficiencies (e.g., TumorDetect, PathAssist Derm), and AIM products for automated, reproducible biomarker quantification (e.g., AIM-PD-L1, AIM-HER2, AIM-TumorCellularity), these models provide stronger feature representations and significantly higher starting points for task-specific supervised fine-tuning. PathAssist Derm for reference, saw a 11% improvement in accuracy and robustness when upgrading to PLUTO-4G. By serving as the representational backbone across these product lines, PLUTO-4 enables the development of more powerful diagnostic and research applications that accelerate biomarker discovery, enhance translational insights, and expand the scope of computational pathology-driven discovery.

Acknowledgements

The authors would like to thank Ylaine Gerardin and Jacqueline Brosnan-Cashman for their valuable feedback on the manuscript. The authors also acknowledge Chris Burke, Geetika Singh, and Darren Fahy for their support with data onboarding and management; Jon Ross and Braden Senst for their MLOps and infrastructure support; and Benjamin Chen for his pathology expertise and guidance.

¹PathExplore, IHCExplore, TumorDetect, AIM-PD-L1, AIM-HER2, AIM-TumorCellularity, and PathAssist Derm are For Research Use Only. Not for use in diagnostic procedures.

References

- Alber, M., Tietz, S., Dippel, J., Milbich, T., Lesort, T., Korfiatis, P., Krügener, M., Perez Cancer, B., Shah, N., Möllers, A., and et al. Atlas: A novel pathology foundation model. *arXiv preprint arXiv:2501.05409*, 2025. URL <https://arxiv.org/abs/2501.05409>.
- Beyer, L., Izmailov, P., Kolesnikov, A., Caron, M., Kornblith, S., Zhai, X., Minderer, M., Tschanen, M., Alabdulmohsin, I., Pavetić, F., and et al. Flexivit: One model for all patch sizes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14496–14506, 2023. doi: 10.1109/CVPR52729.2023.01404. URL https://openaccess.thecvf.com/content/CVPR2023/html/Beyer_FlexiViT_One_Model_for_All_Patch_Sizes_CVPR_2023_paper.html.
- Billaut, V., Nofallah, S., Yu, L., Walker, A., Fahy, D., Zhu, S., Agrawal, N., Brosnan-Cashman, J. A., Pokkalla, H., Abel, J. H., Glass, B., Balasubramanian, S., Walk, E., and Beck, A. Ai-based triage of dermatopathology specimens for pathologist workload prioritization. *Laboratory Investigation*, 105(3), 2025. URL [https://www.laboratoryinvestigation.org/article/S0023-6837\(24\)02325-0/fulltext](https://www.laboratoryinvestigation.org/article/S0023-6837(24)02325-0/fulltext).
- Biopimus. Biopimus launches h-optimus-1: A state-of-the-art foundation model for pathology. <https://www.biopimus.com/news/biopimus-launches-h-optimus-1>, April 2025. URL <https://www.biopimus.com/news/biopimus-launches-h-optimus-1>. Blog post.
- Bulten, W., Pinckaers, H., van Boven, H., Vink, R., de Bel, T., van Ginneken, B., van der Laak, J., Hulsbergen-van de Kaa, C., and Litjens, G. Automated deep-learning system for gleason grading of prostate cancer using biopsies: a diagnostic study. *The Lancet Oncology*, 21(2):233–241, 2020. doi: 10.1016/S1470-2045(19)30739-9. URL [https://www.thelancet.com/journals/lanonc/article/PIIS1470-2045\(19\)30739-9/fulltext](https://www.thelancet.com/journals/lanonc/article/PIIS1470-2045(19)30739-9/fulltext).
- Campanella, G., Hanna, M. G., Geneslaw, L., Mirafior, A., Werneck Krauss Silva, V., Busam, K. J., Brogi, E., Reuter, V. E., Klimstra, D. S., and Fuchs, T. J. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature Medicine*, 25(8):1301–1309, 2019. doi: 10.1038/s41591-019-0508-1. URL <https://www.nature.com/articles/s41591-019-0508-1>.
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., and Joulin, A. Emerging properties in self-supervised vision transformers. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 9650–9660, 2021. doi: 10.1109/ICCV48922.2021.00953. URL https://openaccess.thecvf.com/content/ICCV2021/html/Caron_Emerging_Properties_in_Self-Supervised_Vision_Transformers_ICCV_2021_paper.html.
- 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., and et al. Towards a general-purpose foundation model for computational pathology. *Nature Medicine*, 30(3):850–862, 2024. doi: 10.1038/s41591-023-02770-9. URL <https://www.nature.com/articles/s41591-023-02770-9>.
- Darcet, T., Oquab, M., Mairal, J., and Bojanowski, P. Vision transformers need registers, 2024. URL <https://arxiv.org/abs/2309.16588>.
- Dippel, J., Feulner, B., Winterhoff, T., Milbich, T., Tietz, S., Schallenberg, S., Dernbach, G., Kunft, A., Heinke, S., Eich, M.-L., and et al. Rudolfv: A foundation model by pathologists for pathologists. *arXiv preprint arXiv:2401.04079*, 2024. URL <https://arxiv.org/abs/2401.04079>.
- Jaume, G., Doucet, P., Song, A. H., Lu, M. Y., Almagro-Perez, C., Wagner, S. J., Vaidya, A. J., Chen, R. J., Williamson, D. F. K., Kim, A., and Mahmood, F. Hest-1k: A dataset for spatial transcriptomics and histology image analysis. In *Advances in Neural Information Processing Systems*, December 2024.
- Juyal, D., Padigela, H., Shah, C., Shenker, D., Harguindeguy, N., Liu, Y., Martin, B., Zhang, Y., Nercessian, M., Markey, M., and et al. Pluto: Pathology-universal transformer. *ICML Workshop on Efficient and Accessible Foundation Models for Biological Discovery*, 2024. URL <https://arxiv.org/abs/2405.07905>.
- kaiko.ai, Gatopoulos, I., Känzig, N., Moser, R., and Otálora, S. eva: Evaluation framework for pathology foundation models. In *Medical Imaging with Deep Learning*, 2024. URL <https://openreview.net/forum?id=FNBQOPj18N>.
- Kang, M., Song, H., Park, S., Yoo, D., and Pereira, S. Benchmarking self-supervised learning on diverse pathology datasets, 2023. URL <https://arxiv.org/abs/2212.04690>.
- Le, Q., Bhargava, R., Trotter, B., Rose, K., Thibault, M., Shanis, Z., Padigela, H., Ross, J., Billaut, V., Brosnan-Cashman, J. A., Pokkalla, H., Abel, J. H., Glass, B., Walk,

- E., and Beck, A. Ai-powered detection and classification of artifacts in whole-slide images of pathology specimens. *Laboratory Investigation*, 105(3):1368–1379, 2025. doi: 10.1038/s0023-6837(24)03284-7. URL [https://www.laboratoryinvestigation.org/article/S0023-6837\(24\)03284-7/fulltext](https://www.laboratoryinvestigation.org/article/S0023-6837(24)03284-7/fulltext).
- Oquab, M., Darcet, T., Moutakanni, T., Vo, H. V., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Bojanowski, P., and Joulin, A. Dinov2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. URL <https://openreview.net/forum?id=a8TH3U2EKk>.
- Qi, P., Liu, Z., Zhou, X., Pang, T., Du, C., Lee, W. S., and Lin, M. Defeating the training-inference mismatch via fp16, 2025. URL <https://arxiv.org/abs/2510.26788>.
- Saillard, C., Jenatton, R., Llinares-López, F., Mariet, Z., Cahané, D., Durand, E., and Vert, J.-P. H-optimus-0, 2024. URL <https://github.com/bioptimus/releases/tree/main/models/h-optimus/v0>.
- Siméoni, O., Vo, H. V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., Massa, F., Haziza, D., Wehrstedt, L., Wang, J., Darcet, T., Moutakanni, T., Sentana, L., Roberts, C., Vedaldi, A., Tolan, J., Brandt, J., Couprie, C., Mairal, J., Jégou, H., Labatut, P., and Bojanowski, P. Dinov3, 2025. URL <https://arxiv.org/abs/2508.10104>.
- Su, J., Lu, Y., Pan, S., Wen, B., and Liu, Y. Roformer: Enhanced transformer with rotary position embedding. *arXiv preprint arXiv:2104.09864*, 2021. URL <https://arxiv.org/abs/2104.09864>.
- 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., and Poon, H. A whole-slide foundation model for digital pathology from real-world data. *Nature*, 630: 181–188, 2024. doi: 10.1038/s41586-024-07441-w.
- Zimmermann, E., Vorontsov, E., Viret, J., Casson, A., Zelechowski, M., Shaikovski, G., Tenenholtz, N., Hall, J., Fuchs, T. J., Fusi, N., Liu, S., and Severson, K. Virchow 2: Scaling self-supervised mixed magnification models in pathology. *arXiv preprint arXiv:2408.00738*, 2024. URL <https://arxiv.org/abs/2408.00738>.