Challenges in Deep Learning-Based Small Organ Segmentation: A Benchmarking Perspective for Medical Research with Limited Datasets

Phongsakon Mark Konrad[a], Andrei-Alexandru Popa[b], Yaser Sabzehmeidani[b], Liang Zhong[c], Elisa A. Liehn[d], Serkan Ayvaz[a,*]

[a]*Centre for Industrial Software, University of Southern Denmark, Alsion 2, Sønderborg, 6400, Denmark*
[b]*Centre for Industrial Mechanics, University of Southern Denmark, Alsion 2, Sønderborg, 6400, Denmark*
[c]*Duke-NUS, 8 College Road, Singapore, 169857, Singapore*
[d]*National Heart Center Singapore, 5 Hospital Dr, Singapore, 169609, Singapore*

Abstract

Accurate segmentation of carotid artery structures in histopathological images is vital for advancing cardiovascular disease research and diagnosis. However, deep learning model development in this domain is constrained by the scarcity of annotated cardiovascular histopathological data. This study investigates a systematic evaluation of state-of-the-art deep learning segmentation models—including convolutional neural networks (U-Net, DeepLabV3+), a Vision Transformer (SegFormer), and recent foundation models (SAM, MedSAM, MedSAM + UNet) on a limited dataset of cardiovascular histology images. Despite employing an extensive hyperparameter optimization strategy with extensive Bayesian search, our findings reveal that model performance is highly sensitive to data splits, with minor differences driven more by statistical noise than by true algorithmic superiority. This instability exposes the limitations of standard benchmarking practices in low-data clinical settings and challenges the assumption that performance rankings reflect meaningful clinical utility.

*Corresponding author

Email addresses: phon@mmmi.sdu.dk (Phongsakon Mark Konrad), andrei@sdu.dk (Andrei-Alexandru Popa), yasers@sdu.dk (Yaser Sabzehmeidani), zhong.liang@duke-nus.edu.sg (Liang Zhong), liehn.elisaanamaria@nhcs.com.sg (Elisa A. Liehn), seay@mmmi.sdu.dk (Serkan Ayvaz)

We caution that such practices may foster a false sense of certainty that could mislead model selection and downstream clinical decisions. Our findings highlight the need for rethinking model evaluation frameworks in small-data medical AI research and advocate for a more uncertainty-aware, context-sensitive assessment strategy to better align machine learning practices with the realities of clinical research.

Keywords: Medical Image Segmentation, Illusion of Control, Small-Dataset Evaluation, Statistical Stability, Histopathology, Foundation Models, Explainable AI

---

1. Introduction

Histological analysis remains a cornerstone of cardiovascular pathology, providing invaluable insights for establishing diagnoses and tracking the morphological changes that define diseases like myocardial infarction and instent restenosis (Birare et al., 2024; Dinc, 2023; Honda et al., 2023; Matsuhiro et al., 2020).

However, ensuring reproducibility of results across different laboratories remains a critical challenge. Although numerous guidelines exist for standardizing instrument calibration and operation, there is a notable deficiency in comprehensive documentation of microscopy protocols and data analysis procedures. This shortfall in detailed reporting compromises experimental integrity and hampers efforts to achieve true scientific reproducibility.

In the cardiovascular field, debates continue over the optimal methodologies for quantifying functional and morphological changes following surgical procedure in murine models, like atherosclerosis or myocardial infarction. Key questions revolve around how plaque size or infarct size should be measured, how it should be expressed (as a percentage based on a single section, averaged across multiple sections, and should focus solely on affected regions or include the entire organ). The absence of consensus has led to a heterogeneous body of published data, which complicates the translation of findings and impedes the development of standardized measurement protocols. As a result, researchers often find themselves reconstructing methodologies from scratch, slowing progress and undermining reproducibility within this vital area of cardiovascular research.

The development of computer-based systems and artificial intelligence (AI) offers promising solutions for improving standardization and consistency in

cardiovascular research. While these tools are mainly used for large datasets, smaller datasets from animal studies are often neglected. There is a growing movement in the cardiovascular community to minimize animal use (Hubrecht and Carter, 2019) and obtain maximum information from minimal samples, even though current AI applications like 3D reconstructions do not fully support this approach. The variety of methods currently used raises concerns about the reliability of tissue analysis results. Therefore, efforts are being made to compare different methods on the same limited datasets to improve consistency and standardization across laboratories, ultimately strengthening the reproducibility of cardiovascular research.

As an example, the detailed characterization of vascular lesions from histological sections is especially critical in understanding conditions like carotid artery stenosis, a primary contributor to ischemic stroke (Arasu et al., 2021). However, the translation of this vital technique from research to widespread clinical application is fraught with challenges. The current gold standard, which relies on manual analysis of tissue by highly trained personnel, is notoriously laborious and time-consuming (Roberts et al., 2020). This manual process not only risks damaging the fragile tissue samples but, more critically, it introduces significant inter-observer variability, a fundamental barrier to producing reproducible science (Pandelea-Dobrovicescu et al., 2020).

In recent years, deep learning models, from established Convolutional Neural Networks (CNNs) to massive, pre-trained Foundation Models (FMs), have shown immense promise (Ahmad et al., 2020; De Matos et al., 2021; Kong et al., 2020; Shamai et al., 2022). Yet, a critical "valley of credibility" separates the data-rich environments where these models are developed from the data-scarce reality where they are most needed, a problem especially apparent in specialized medical research (Basla et al., 2024). This leads to a difficult question: how do we select the right model when we only have a handful of patient samples to guide us? What if our methods for choosing the "best" model are themselves flawed?

This study was initiated to systematically investigate the performance of state-of-the-art deep learning models for carotid artery segmentation in the context of limited cardiovascular histopathological data. Our objective was to identify the most effective model architecture for this clinically significant yet data-constrained task. To this end, we conducted a comprehensive and transparent benchmarking study involving traditional CNNs (U-Net(Ronneberger et al., 2015), DeepLabV3+(Chen et al., 2017)), a modern lightweight transformer (SegFormer(Xie et al., 2021)), and the latest Foundation Models (SAM

(Kirillov et al., 2023), MedSAM (Ma et al., 2024), MedSAM + UNet (Yang et al., 2024)).

The evaluation methodology followed widely adopted practices in machine learning research, incorporating extensive hyperparameter tuning via a thorough Bayesian optimization strategy and multiple cross-validation schemes (Cooper et al., 2021). While such approaches aim to reduce variability and enhance reliability, they are not immune to the well-documented "illusion of control", a bias that can lead researchers to overestimate the stability of model performance in low-data regimes (Langer, 1975).

This study moves beyond its initial goal of simple model selection to test a more critical hypothesis: that standard benchmarking leaderboards are inherently fragile in low-data regimes. We contend that any observed model ranking is likely an artifact of statistical noise and specific evaluation choices, rather than a reflection of true algorithmic superiority. Our aim is to empirically demonstrate this instability and propose a more robust evaluation framework centered on statistical equivalence and practical efficiency. We therefore frame this work as a deliberate case study on an extremely small dataset (N=9), using it as a stress test to probe the fundamental limits of common benchmarking practices. The goal is to provide a clear, empirical demonstration of the underlying risks of instability that arise when data is scarce.

2. Related Work

2.1. CNN-based Architectures for Medical Segmentation

For over a decade, Convolutional Neural Networks (CNNs) have been the cornerstone of medical image segmentation. The U-Net architecture, introduced by Ronneberger et al. (2015), is arguably the most influential model in this domain. Its symmetric encoder-decoder design, coupled with skip connections that concatenate feature maps from the encoder to the decoder, allows it to preserve high-resolution spatial information that is crucial for precise boundary delineation. The success of U-Net has led to a creation of variants, such as UNet++ (Zhou et al., 2019), which uses redesigned, nested skip pathways to capture features at multiple scales, and R2U-Net (Alom et al., 2018), which integrates recurrent and residual connections to create an even more powerful feature representation. As noted by Paheding et al. (2024), these models created a vital bridge between AI research and clinical application.

Another prominent CNN architecture is DeepLabV3+ (Chen et al., 2017). Its key innovation is the Atrous Spatial Pyramid Pooling (ASPP) module, which employs parallel atrous (or dilated) convolutions with different dilation rates. This allows the model to probe an incoming feature map with filters that have multiple fields of view, capturing multi-scale context without increasing the number of parameters or computational cost. This ability to handle objects of varying sizes makes it particularly effective for complex medical scenes where pathologies can manifest at different scales. Both U-Net and DeepLabV3+ have become foundational tools, but as Yu et al. (2021) highlight, their performance is fundamentally linked to the availability of large, expertly annotated datasets, a condition often unmet in specialized clinical research.

2.2. Transformer-based Architectures in Vision and Medicine

More recently, the field has seen a shift towards Transformer-based architectures, which leverage self-attention mechanisms to model long-range dependencies within an image, overcoming the inherent local receptive field limitations of CNNs (Vaswani et al., 2017). While initial Vision Transformers (ViTs) produced single-resolution feature maps, making them less suitable for dense prediction tasks like segmentation, newer architectures have addressed this. SegFormer, proposed by Xie et al. (2021), stands out as a powerful and efficient design. It employs a hierarchical Transformer encoder that generates multi-scale features, much like a CNN, but with the global context advantage of self-attention. This is paired with a simple, lightweight All-MLP decoder that effectively fuses these multi-scale features. This design avoids the complex decoders of other models, resulting in a computationally efficient yet powerful architecture for semantic segmentation. The adoption of such modern architectures reflects a desire for more flexible and powerful AI systems (Chrysos, 2024), but also introduces new validation challenges due to their intricate training dynamics (Xia and Wang, 2023).

2.3. Foundation Models and Their Adaptation for Medical Imaging

The paradigm of large-scale pre-training resulted in foundation models, with the Segment Anything Model (SAM) representing a significant breakthrough for image segmentation (Kirillov et al., 2023). SAM is a generalist, promptable model trained on over a billion masks from natural images, enabling it to perform zero-shot segmentation of unseen objects and images. Its

architecture consists of a powerful ViT image encoder, a prompt encoder that interprets points, boxes, and text, and a fast mask decoder.

The remarkable capabilities of SAM spurred immediate efforts to adapt it for the specialized medical domain. As Ma et al. (2024) demonstrated with MedSAM, fine-tuning the model on a large corpus of medical images can significantly improve its performance on these tasks. Their strategy involved freezing the massive image encoder and only updating the mask decoder to reduce computational cost. Further innovations have come from hybrid models. For instance, combining a powerful foundation model encoder with a proven CNN-based decoder, such as in a SAM-UNet hybrid (Yang et al., 2024), aims to leverage the global feature extraction of Transformers with the robust local feature refinement of architectures like U-Net. However, the out-of-the-box performance of SAM on fine-grained medical tasks can be poor, necessitating careful fine-tuning.

2.4. The Challenge of Hyperparameter Optimization in Medical AI

Successfully applying any of these advanced models is critically dependent on finding an optimal set of hyperparameters. The performance of a given architecture can vary dramatically based on the choice of learning rate, batch size, optimizer (e.g., Adam, SGD), learning rate scheduler, and loss function (Lai, 2024). The combination of these choices creates a vast, high-dimensional search space that is computationally infeasible to explore exhaustively with traditional methods like grid search. While random search is more efficient, it provides no guarantee of finding an optimal configuration (Subaşı, 2024).

More advanced strategies like Bayesian Optimization have become essential for navigating this complex space efficiently (Agrawal et al., 2024). These methods build a probabilistic surrogate model of the objective function to intelligently select which hyperparameter combinations to evaluate next. This challenge is further compounded by the models themselves. Even Parameter-Efficient Fine-Tuning (PEFT) techniques like LoRA (Hu et al., 2022), which reduce the number of trainable *model weights*, often introduce new *hyperparameters* (e.g., the LoRA rank 'r' and alpha 'α') that must themselves be tuned. Therefore, ensuring that each model in a comparison is performing at its peak potential requires a rigorous and systematic hyperparameter optimization strategy, which is a central component of our study's methodology.

2.5. The Credibility of Benchmarking on Small Medical Datasets

Beyond the challenge of optimizing individual models, a greater methodological problem arises when comparing them, especially in the low-data regimes common in medical research (Salehi et al., 2023; Hong et al., 2023; Maier-Hein et al., 2018). The standard practice of declaring a single "state-of-the-art" model based on its rank on a benchmark leaderboard is coming under increasing scrutiny (Lin, 2019). As Lin et al. (2021) have argued, a model's rank can be a fragile construct, highly dependent on the chosen evaluation protocol and the specific data split. The well-known bias-variance trade-off becomes a central dilemma: Leave-One-Out Cross-Validation (LOOCV), while offering a nearly unbiased performance estimate, suffers from high variance, making results unstable (Ranglani, 2024).

Furthermore, the process of extensive hyperparameter optimization, while necessary for a fair comparison, carries the inherent risk of "overfitting to the validation set" (Makarova et al., 2021; Hertel et al., 2022). This can lead to optimistically biased results that are not reproducible on new data (Tetko et al., 2024). While the literature is rich with papers proposing novel model architectures, there is a distinct lack of methodological studies that address the *stability and reliability of the evaluation process itself*. Our work directly engages with this credibility crisis by empirically demonstrating how these theoretical concerns manifest as dramatic instabilities in a real-world medical AI task. We provide a cautionary tale for any researcher working with limited data, shifting the focus from the pursuit of an illusory top rank toward the critical importance of statistical stability.

3. A Systematic Model Evaluation Framework

To investigate the stability of deep learning benchmarks on small medical datasets, we designed and executed a rigorous, multi-stage evaluation framework. This process, illustrated in Figure 1, was developed to ensure a fair and comprehensive comparison across diverse model architectures by systematically controlling for common sources of variability.

The pipeline begins with data collection and expert annotation (1). Crucially, we reject the use of default model configurations. Instead, each architecture undergoes an exhaustive Bayesian hyperparameter optimization (2) to ensure it performs at its peak potential. This commitment is vital for mitigating the confounding effects of suboptimal configurations when comparing models (Agrawal et al., 2024; Lai, 2024). To generate statistically sound

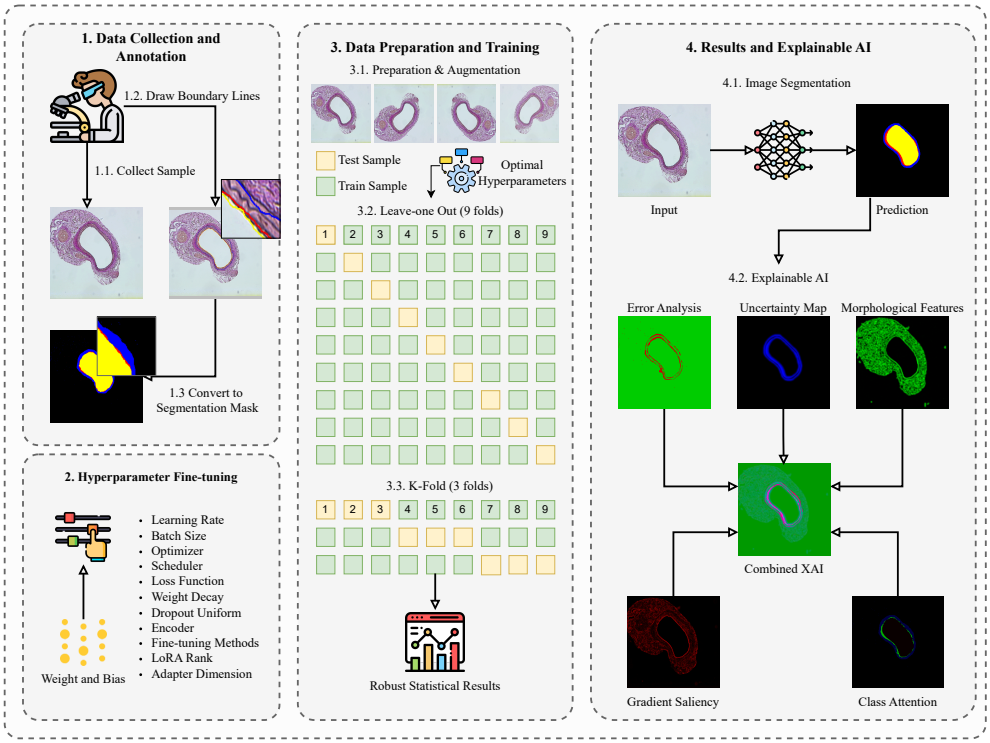

Figure 1: The systematic evaluation framework. Our process begins with (1) data collection and expert annotation. Each model architecture then undergoes (2) extensive Bayesian hyperparameter optimization to find its optimal configuration. (3) Data preparation, augmentation and rigorous cross-validation (LOOCV and 3-Fold) is used to generate robust statistical results for (4.1.) segmentation performance, which are then dissected using our (4.2.) multi-modal XAI framework for qualitative insight.

results, we then subject each optimized model to a demanding multi-strategy cross-validation (3.2, 3.3) in combination with on-the-fly image augmentation (3.1), employing both Leave-One-Out (LOOCV) and 3-Fold splits. LOOCV, in particular, has been shown to provide robust statistical estimates for small datasets (Xu et al., 2023). The quantitative results (4.1) are then dissected using our five-layer Explainable AI (XAI) framework (4.2) to understand not just *what* the model decided, but also the potential sources of performance instability across different data splits (Patil, 2024; Islam et al., 2024).

3.1. Dataset and Clinical Context

The dataset consists of nine high-resolution histological images of vascular tissue from the injured carotid artery of mice. Mice (C57/Bl6, male) undergo the carotid de-endothelialization, as described before (Curaj et al., 2016) (Permission no.31 from 27.09.2022, Directia Sanitara Veterinara si Pentru Siguranta Alimentelor Bucharest, Romania). The carotide artery was excized 4 weeks after de-endothelialization and embedded in paraffin. Serial sections (5 µm thick) were collected starting with the bifurcation between the internal carotide artery and external carotid artery. 9 sections, in 50µm appart were stained with Elastica van Gieson (EVG) and photographed. These source images are stored in BMP format with variable native resolutions (e.g., 984×792 pixels) and were prepared with Elastica van Gieson (EVG) staining. This small sample size ($N = 9$) was intentionally chosen to simulate a realistic and challenging clinical research scenario where data is scarce, forcing a focus on model robustness and data efficiency. The task is a four-class semantic segmentation problem designed to delineate key vascular structures: Lumen (class 1), Neointima (class 2), and Media (class 3), against the Background (class 0). Ground truth masks were meticulously generated from expert annotations; initial color-coded JPG annotations were processed into four-class, integer-labeled PNG masks to create clean, analyzable data for model training and evaluation. For use in our pipeline, these high-resolution source images and masks are systematically preprocessed, including resizing to standardized input dimensions (e.g., 256×256 pixels for conventional models). An example is shown in Figure 2.

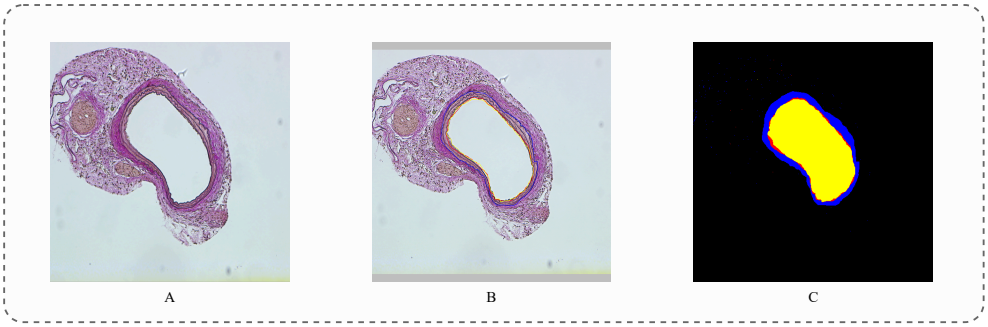

Figure 2: Example of data preparation. (A) Original histological image. (B) Expert line-art annotation. (C) Processed ground truth segmentation mask for Lumen (yellow), Neointima (red), and Media (blue).

3.2. Data Preprocessing and Augmentation

Our data preparation strategy was tailored to the specific requirements of each model family to ensure optimal performance and a fair comparison. The pipeline involved model-specific resizing followed by a unified augmentation strategy. For the conventional CNN and SegFormer architectures, all images and their corresponding masks were resized to a uniform dimension of 256×256 pixels and normalized to the range $[0, 1]$. In contrast, the foundation models, including SAM and MedSAM, required a different approach to leverage their pre-trained capabilities. Their encoders (ViT-B, ViT-L) were designed for high-resolution inputs; therefore, images destined for these models were resized to 1024×1024 pixels to fully utilize the intricate spatial features learned during their original training.

Following this model-specific resizing, we employed an *on-the-fly data augmentation* strategy during training for all models to expand the limited dataset and improve generalization (Liu et al., 2024; Xu et al., 2020; Dasari et al., 2022). This technique applied a series of random transformations to each image-mask pair at every epoch, encompassing both geometric and color-based changes. To account for variations in sample orientation, geometric augmentations included random horizontal and vertical flips, alongside affine transformations such as rotations ($\pm 20°$), scaling (from 80% to 120%), and shearing. Critically, these transforms were applied synchronously to both the image and its segmentation mask to maintain perfect spatial alignment. Complementing these spatial adjustments, color augmentations were applied exclusively to the input images. This involved color jittering to randomly alter brightness, contrast, saturation, and hue, thereby simulating variations in staining protocols and lighting conditions.

To ensure a rigorous and unbiased assessment of true performance, the validation and test sets for all models were subjected only to their respective resizing and normalization steps, without any data augmentation.

3.3. Model Architecture Selection

To robustly test benchmark stability, we chose a diverse set of models spanning three distinct architectural paradigms: classic Convolutional Neural Networks (CNNs), a modern Vision Transformer, and strong newcomers large-scale Foundation Models. This selection provides a challenging testbed, as each paradigm has different strengths, complexities, and data requirements.

More specifically, we selected U-Net and DeepLabV3+ as established standards, SegFormer due to its computational efficiency, and SAM-based

models to evaluate the impact of large-scale pre-training, including a novel hybrid architecture to test synergistic effects.

3.3.1. Traditional CNN Architectures

We included two powerful and widely-used CNN architectures, as they represent the established standard in medical image segmentation. The first, U-Net, is a foundational model in medical imaging due to its iconic encoder-decoder structure and skip connections. To ensure we evaluated its strongest possible configuration, we conducted an exhaustive search across eight different pre-trained encoders, including the ResNet (He et al., 2016), EfficientNet (Tan and Le, 2019), DenseNet (Huang et al., 2017) and MobileNet Howard et al. (2017) families. As a strong alternative, we also included DeepLabV3+. Its key feature is the Atrous Spatial Pyramid Pooling (ASPP) module, which uses parallel dilated convolutions to capture features at multiple scales, making it particularly effective for segmenting complex scenes with objects of varying sizes.

3.3.2. Modern Vision Transformer

To represent the latest advancements in efficient model design, we selected SegFormer. This architecture stands out from earlier Vision Transformers by employing a hierarchical encoder that produces multi-scale features, similar to a CNN, but without requiring complex positional encodings. This encoder is paired with a lightweight MLP (Multi-Layer Perceptron) decoder, resulting in a model that is both computationally efficient and highly effective for semantic segmentation tasks.

3.3.3. Foundation Models and Hybrid Architectures

To assess the impact of large-scale pre-training, we included the Segment Anything Model (SAM) and its medical derivative, MedSAM. The standard SAM model is a powerful generalist but requires manual prompts to guide segmentation. Before developing a fully automated method, we explored a two-stage approach where a detection model, such as a fine-tuned YOLO model, would generate bounding box prompts for SAM (Khalili et al., 2024). However, this proved to be ineffective for two main reasons. First, fine-tuning a detection model on our small dataset did not produce bounding boxes with sufficient accuracy for the fine-grained segmentation task. Second, converting our ground-truth data, which is based on boundary lines, into distinct bounding boxes was inherently difficult due to the close proximity

and occasional crossing of the lines between adjacent tissue layers (Xu et al., 2024). For these reasons, we abandoned the two-stage strategy and instead implemented a more robust *auto-prompting mechanism* to integrate SAM into our automated benchmarking framework. This system replaces manual input with a lightweight neural network that takes the image embeddings from SAM's encoder as input. The network uses a *multi-head cross-attention mechanism* to fuse these image-conditioned features with a small set of learnable prompt tokens. The resulting context-aware prompt embedding is then passed directly to SAM's mask decoder, enabling fully automated, end-to-end training without human intervention. Furthermore, we included MedSAM, a version of SAM specifically adapted for medical images, which we fine-tuned using Low-Rank Adaptation (LoRA). Finally, to explore architectural synergies, we engineered a novel MedSAM-UNet Hybrid. This model uses the powerful, frozen Vision Transformer encoder from MedSAM to extract rich, global features and feeds them into a custom-built U-Net-style decoder, thereby combining the global context awareness of a Transformer with the proven spatial precision of a U-Net.

3.4. Hyperparameter Optimization

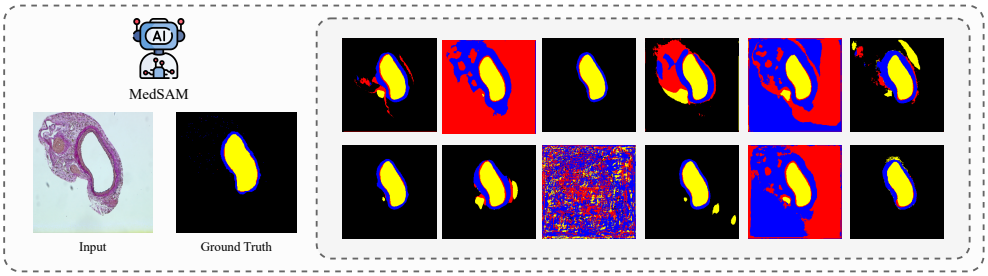

Figure 3: The critical impact of hyperparameter selection. All twelve segmentation results were generated by the same MedSAM architecture. The stark variation, from accurate predictions to catastrophic failure, is due solely to different hyperparameter configurations. This illustrates that without a rigorous search, one could incorrectly dismiss a powerful model based on a single, suboptimal trial.

Hyperparameters significantly impact model performance, particularly with limited data (Lai, 2024). To account for this, we used Bayesian optimization (explained in Appendix A) over a total of 1,000 runs, managed by

Weights & Biases Sweeps[1]. The number of runs per model varied based on the search space and prior nature; for example, UNet had 300 runs, SAM had 100, and other models had 150. Our extensive search space is detailed in Table 1. As Figure 3 clearly illustrates, this rigorous approach is crucial: without it, a strong architecture might be unfairly dismissed due to a single suboptimal configuration (Wang and Yang, 2023). This process ensures each architecture is compared at its near-optimal configuration, which is detailed in Table 2.

Table 1: Hyperparameter Search Space Summary.

Parameter	**Search Strategy**	**Range / Values**
General Parameters		
Learning Rate	Log-uniform	[1e-5, 1e-2]
Batch Size	Categorical	[2, 4, 8, 16, 32]
Optimizer	Categorical	[Adam, AdamW, SGD, RMSprop, Nadam]
Scheduler	Categorical	[Cosine, Polynomial, Step, Warmup Cosine]
Loss Function	Categorical	[Focal-Dice, Focal-Tversky, Unified Focal]
Weight Decay	Log-uniform	[1e-6, 1e-2]
Dropout	Uniform	[0.0, 0.3]
Architecture-Specific Parameters		
Encoder (UNet)	Categorical	[ResNet18/34/50, EfficientNet-B0/B1/B2, DenseNet121, MobileNetV2]
SAM Model Type	Categorical	[ViT-B, ViT-L]
Fine-tuning Method	Categorical	[SVD, LoRA, Norm Tuning]
LoRA Rank	Categorical	[8, 16, 32]
Adapter Dimension	Categorical	[128, 256, 512]

[1]https://wandb.ai/site/sweeps/

Table 2: Optimal Hyperparameters for Each Model, Identified via Bayesian Search

Parameter	**UNet**	**DLV3+**[a]	**SF**[b]	**SAM**	**MedSAM**	**M-UNet**[c]
Core Architecture						
Architecture	RN34[d]	RN50	MiT-B0	ViT-B	ViT-L	ViT-B + U
Encoder Weights	IN[d]	IN	IN	—	—	—
Loss Function	F-D[d]	F-D	F-D	F-D	F-D	F-D
Epochs	200	200	200	200	200	200
Optimization & Regularization						
Learning Rate	1×10^{-4}	5×10^{-5}	6×10^{-5}	6×10^{-4}	8.2×10^{-4}	8.7×10^{-4}
Optimizer	Adam	Adam	Adam	AdamW	Adam	Adam
LR Scheduler	1Cycle[e]	1Cycle	Cosine	W-Cos[e]	W-Cos	W-Cos
Batch Size	8	28	8	4	2	4
Weight Decay	1×10^{-4}	1×10^{-4}	1×10^{-4}	1×10^{-3}	6.7×10^{-4}	8.1×10^{-5}
Dropout	0.1	—	—	—	—	0.0
Grad. Accum.	1	1	1	1	2	1
Grad. Clipping	—	—	—	2.0	1.0	1.0
Parameter-Efficient Fine-Tuning (PEFT)						
Method	—	—	—	NT[e]	LoRA	NT
LoRA Rank	—	—	—	16	8	16
Adapter Dim.	—	—	—	512	512	—

[a] DLV3+: DeepLabV3+; [b] SF: SegFormer;
[c] M-UNet: MedSAM + UNet. [d] RN: ResNet; IN: ImageNet; F-D: Focal-Dice.
[e] 1Cycle: OneCycle; W-Cos: Warmup Cosine; NT: Norm Tuning.

3.5. Evaluation Metrics and Cross-Validation

Our primary evaluation metric is the macro-averaged Dice Similarity Coefficient (DSC) over the three tissue classes, chosen for its robustness to class imbalance (Yeung et al., 2023). We also report Intersection over Union (IoU) (Müller et al., 2022). To probe for instability, we employ two cross-validation strategies: Leave-One-Out (LOOCV), a low-bias, high-variance estimator, and 3-Fold CV, a higher-bias, lower-variance estimator (Bhagat and Bakariya, 2025).

Our loss function, identified as optimal through our hyperparameter search, is a composite of Focal and Dice loss designed to balance pixel-level accuracy

and region-level overlap (Yeung et al., 2021).

3.5.1. Focal Loss

To mitigate the impact of severe class imbalance, we employ Focal Loss ($\mathcal{L}_{\text{focal}}$). This loss function down-weights easy examples and focuses training on hard, misclassified examples:

$$\mathcal{L}_{\text{focal}} = -\alpha(1 - p_t)^{\gamma} \log(p_t)$$

Here, p_t represents the model's estimated probability for the ground truth class. We set $\alpha = 0.8$ for class balancing and $\gamma = 2.0$ to control the focusing strength on hard examples.

3.5.2. Dice Loss

The Dice Loss ($\mathcal{L}_{\text{dice}}$) directly optimizes our primary evaluation metric, promoting better overlap between predicted and ground truth masks. Its differentiable form is given by:

$$\mathcal{L}_{\text{dice}} = 1 - \frac{1}{K} \sum_{k=1}^{K} \frac{2 \sum_{i,j} \mathbf{P}_k(i,j) \cdot \mathbf{Y}_k(i,j) + \epsilon}{\sum_{i,j} \mathbf{P}_k(i,j) + \sum_{i,j} \mathbf{Y}_k(i,j) + \epsilon} \tag{1}$$

In this formulation, $\mathbf{P}_k$ denotes the predicted probabilities for class k, $\mathbf{Y}_k$ is the one-hot encoded ground truth for class k, and $\epsilon = 10^{-6}$ is a small constant added to prevent division by zero. K is the total number of classes.

3.5.3. Composite Loss Function

Our primary and most effective loss function, the Composite Loss Function ($\mathcal{L}_{\text{focal\_dice}}$), strategically combines Focal Loss and Dice Loss with adaptive weighting:

$$\mathcal{L}_{\text{focal\_dice}} = \lambda_f \mathcal{L}_{\text{focal}} + \lambda_d \mathcal{L}_{\text{dice}}$$

Here, $\lambda_f = 0.5$ and $\lambda_d = 0.5$ provide a balanced optimization of both pixel-level accuracy and region-level overlap, leading to superior segmentation results.

3.6. Multi-Modal Explainable AI (XAI) Framework

Interpretability is a prerequisite for clinical trust. We introduce a framework that synthesizes five XAI modalities (Figure 4) to provide a holistic diagnostic view. Crucially, we repurpose this framework from merely explaining a prediction to diagnosing the sources of statistical instability across

different data splits. An explainability framework is only useful if its outputs map directly to clinical needs. We designed our five layers to address distinct, practical questions that arise during pathological review, from assessing diagnostic accuracy to building confidence in the model's predictions. Table 3 provides a systematic validation of how each component serves critical clinical functions.

Mathematically, our framework generates a composite explanation map $\mathbf{E}$ from an input image $\mathbf{I}$, its ground truth $\mathbf{Y}$, and the model's prediction $\hat{\mathbf{Y}}$. This composite map is constructed as a weighted sum of five individual, distinct explanation layers $\mathbf{E}_i$:

$$\mathbf{E} = \sum_{i=1}^{5} \alpha_i \mathbf{E}_i$$

The coefficients α_i are adaptive and can be adjusted to emphasize different explanatory aspects based on the specific clinical question or the user's focus. The individual layers are detailed below:

3.6.1. Layer 1: Error Analysis

This foundational layer provides direct visual feedback on the correctness of the model's prediction. It generates a binary error map $\mathbf{E}_1$ that immediately highlights regions of agreement and disagreement between the predicted mask $\hat{\mathbf{Y}}$ and the ground truth $\mathbf{Y}$.

$$\mathbf{E}_1(x, y) = \begin{cases} \text{color for correct} & \text{if } \hat{\mathbf{Y}}(x, y) = \mathbf{Y}(x, y) \\ \text{color for error} & \text{otherwise} \end{cases}$$

Here, distinct colors represent correct and erroneous pixels, respectively.

3.6.2. Layer 2: Uncertainty Estimation

This layer quantifies the model's confidence in its pixel-wise predictions. Uncertainty is often highest at critical tissue boundaries, indicating areas where the model is less certain about its classification. We compute an uncertainty map $\mathbf{E}_2$ using the pixel-wise entropy $\mathcal{H}$ of the model's class probabilities $\mathbf{P}$.

$$\mathbf{E}_2(x, y) = \mathcal{H}(\mathbf{P}(x, y)) = -\sum_k P_k(x, y) \log_2 P_k(x, y)$$

3.6.3. Layer 3: Morphological Analysis

To ground the explanation in clinically relevant anatomical and structural features, this layer $\mathbf{E}_3$ captures tissue-specific morphological characteristics by integrating various image processing techniques, such as texture descriptors and structural filters.

3.6.4. Layer 4: Class-wise Attention

This layer $\mathbf{E}_4$ adaptively highlights regions corresponding to specific classes where the model exhibits suboptimal performance. The contribution of each class mask $\mathbf{M}_k$ to the overall attention map is weighted by its performance deficit, calculated as $(1 - DSC_k)$. This approach focuses the explanation on the most problematic tissue types.

$$\mathbf{E}_4(x, y) = \sum_{k=1}^{3} (1 - DSC_k) \cdot \mathbf{M}_k(x, y)$$

3.6.5. Layer 5: Gradient-based Saliency

The final layer $\mathbf{E}_5$ emphasizes the fine-grained structural edges that are often critical for accurate histological interpretation. This is achieved by computing the magnitude of the image gradient, highlighting areas of rapid intensity change.

$$\mathbf{E}_5(x, y) = \sqrt{\left(\frac{\partial \mathbf{I}}{\partial x}\right)^2 + \left(\frac{\partial \mathbf{I}}{\partial y}\right)^2}$$

This comprehensive, multi-modal explanation framework provides clinicians with a deeper understanding of the model's behavior, facilitating trust and enabling more informed diagnostic decisions.

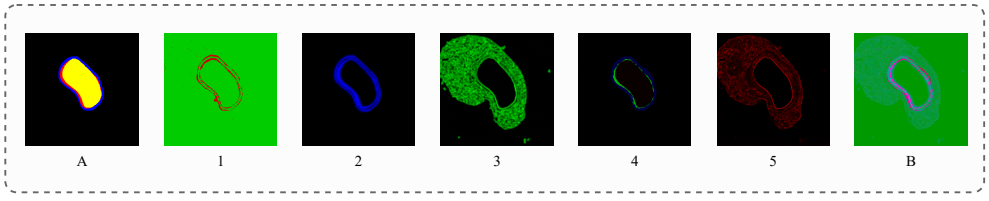

Figure 4: The five-layer XAI framework. (A) A model's prediction is dissected into (1) Error Analysis, (2) Uncertainty, (3) Morphology, (4) Attention, and (5) Saliency. (B) These are synthesized into an integrated explanation, providing a multi-faceted view of the model's behavior.

Table 3: Clinical Validation Framework for XAI Components. Column headers (1-5, B) refer to the components detailed in Figure 4: (1) Error, (2) Uncertainty, (3) Morphology, (4) Attention, (5) Saliency, and (B) Integrated.

Clinical Need	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(B)**
Diagnostic Accuracy	✓	✓	✓	✓	✓	✓
Boundary Delineation	×	✓	✓	×	✓	✓
Tissue Characterization	×	×	✓	×	×	✓
Quality Assurance	✓	✓	×	✓	×	✓
Educational Value	✓	✓	✓	✓	✓	✓
Clinical Confidence	×	✓	✓	✓	×	✓

3.7. Statistical Analysis

Given the small sample size of our dataset ($N = 9$), we cannot assume a normal distribution for performance scores. Therefore, our statistical analysis relies primarily on *non-parametric methods* to compare model performance and assess the stability of the results. All statistical analyses were conducted using Python, leveraging the *pandas* library for data manipulation and *Matplotlib* with *Seaborn* for visualization.

1. **Bootstrap Confidence Intervals**: To quantify the uncertainty around mean performance scores without assuming a normal distribution, we calculated 95% bootstrap confidence intervals (Efron, 1992). This was achieved by resampling the cross-validation fold scores with replacement 10,000 times for each model to generate a distribution of the mean, from which the percentile-based confidence intervals were derived. The core resampling was implemented using the `resample` function from *scikit-learn*[2] (Pedregosa et al., 2011).
2. **Non-Parametric Rank Comparison**: To formally compare the models against each other across all cross-validation folds, we employed the Friedman test (Friedman, 1937). This non-parametric test is highly recommended for comparing multiple classifiers over multiple datasets (or, in our case, folds) (Demšar, 2006). We then used the Nemenyi test as a post-hoc analysis to identify which specific models have statistically

[2] https://scikit-learn.org/

different average ranks. The results are visualized using Critical Difference (CD) plots. This entire analysis was performed using functions from the *SciPy*[3] library (Virtanen et al., 2020).

3. **Effect Size Measurement**: To assess the practical magnitude of the performance differences between model pairs, we calculated Cohen's d (Cohen, 2013). This provides a standardized measure of the difference between two means, helping to distinguish between statistical significance and practical importance.

By combining these methods, we created a robust statistical framework to investigate our hypothesis that model rankings are unstable and not statistically meaningful in a low-data setting.

4. Results

This section details the empirical findings of our study. We first present the quantitative performance of the models under two standard cross-validation protocols, illustrating how a conventional analysis can suggest a clear winning architecture. We then systematically deconstruct this initial finding through a series of deeper statistical analyses, revealing the profound instability of the benchmark and highlighting the methodological challenges inherent in small-dataset evaluation.

4.1. Quantitative Performance Across Validation Protocols

To establish a performance baseline, all optimally-configured models were evaluated using both Leave-One-Out (LOOCV) and 3-Fold cross-validation (CV).

Under the LOOCV protocol, which provides a nearly unbiased but high-variance estimate of performance, the **SegFormer** model emerged as the apparent leader. As shown in Table 4, **SegFormer** achieved the highest mean scores across all metrics, including the primary macro-averaged Dice Similarity Coefficient (DSC) of 0.821.

When we switched to a 3-Fold CV protocol, a method with lower variance but potentially higher bias, **SegFormer** continued to show strong performance, again achieving the highest scores in most categories, as detailed in Table 5. A surface-level interpretation of these results would typically lead to

[3] https://scipy.org/

the conclusion that **SegFormer** is the state-of-the-art model for this task. The subsequent analysis will demonstrate why this conclusion is statistically unfounded.

Table 4: LOOCV Performance Summary. Results are shown as Mean ± Std Dev. Highest values in each column are in bold. $\boldsymbol{\mu}_{\text{macro}}$ denotes the macro-averaged mean.

Model	Lumen	Neointima	Media	$\boldsymbol{\mu}_{\text{macro}}$
Dice Score (DSC) ↑				
UNet	0.969 ± 0.02	0.628 ± 0.19	0.824 ± 0.05	0.807 ± 0.07
DeepLabV3+	0.967 ± 0.01	0.632 ± 0.20	0.788 ± 0.02	0.796 ± 0.07
SegFormer	$\mathbf{0.974 \pm 0.01}$	$\mathbf{0.659 \pm 0.23}$	$\mathbf{0.829 \pm 0.04}$	$\mathbf{0.821 \pm 0.09}$
SAM	0.964 ± 0.02	0.627 ± 0.19	0.806 ± 0.06	0.799 ± 0.08
MedSAM	0.971 ± 0.02	0.638 ± 0.18	0.801 ± 0.05	0.803 ± 0.07
MedSAM+UNet	0.966 ± 0.02	0.501 ± 0.15	0.740 ± 0.06	0.735 ± 0.06
Intersection over Union (IoU) ↑				
UNet	0.941 ± 0.04	0.482 ± 0.20	0.703 ± 0.07	0.709 ± 0.09
DeepLabV3+	0.936 ± 0.02	0.489 ± 0.21	0.651 ± 0.03	0.692 ± 0.07
SegFormer	$\mathbf{0.950 \pm 0.01}$	$\mathbf{0.527 \pm 0.23}$	$\mathbf{0.710 \pm 0.06}$	$\mathbf{0.729 \pm 0.09}$
SAM	0.930 ± 0.03	0.480 ± 0.19	0.679 ± 0.08	0.696 ± 0.08
MedSAM	0.944 ± 0.03	0.492 ± 0.19	0.671 ± 0.07	0.702 ± 0.08
MedSAM+UNet	0.934 ± 0.03	0.347 ± 0.15	0.591 ± 0.08	0.624 ± 0.06

4.2. Analysis of Benchmark Instability

A closer inspection reveals that the notion of a single "best" model is an illusion created by the specifics of the evaluation protocol. We identified three primary sources of this instability.

4.2.1. The Winning Model Changes with the Evaluation Protocol

The most direct evidence of instability is that the top-ranked model changes depending on the evaluation context. Table 6 summarizes the best-performing model for each metric across three contexts: the single best run from the 1,000-run hyperparameter sweep, the mean LOOCV result, and the mean 3-Fold CV result. The "Winner" column changes constantly. For the main Macro-Dice metric, the winner is **MedSAM+UNet** in the sweep, **SegFormer** in LOOCV, and **DeepLabV3+** in 3-Fold CV. This demonstrates

Table 5: 3-Fold Cross-Validation Performance Summary. Results are shown as Mean ± Std Dev. Highest values in each column are in bold. $\boldsymbol{\mu}_{\text{macro}}$ denotes the macro-averaged mean.

Model	Lumen	Neointima	Media	$\boldsymbol{\mu}_{\text{macro}}$
Dice Score (DSC) ↑				
UNet	0.969 ± 0.01	0.628 ± 0.02	0.824 ± 0.01	0.726 ± 0.01
DeepLabV3+	0.967 ± 0.01	0.632 ± 0.02	0.788 ± 0.00	0.710 ± 0.01
SegFormer	$\mathbf{0.974 \pm 0.00}$	$\mathbf{0.659 \pm 0.02}$	$\mathbf{0.829 \pm 0.01}$	$\mathbf{0.744 \pm 0.01}$
SAM	0.964 ± 0.01	0.627 ± 0.02	0.806 ± 0.01	0.717 ± 0.01
MedSAM	0.971 ± 0.01	0.638 ± 0.02	0.801 ± 0.01	0.719 ± 0.01
MedSAM+UNet	0.966 ± 0.01	0.501 ± 0.01	0.740 ± 0.01	0.620 ± 0.01
Intersection over Union (IoU) ↑				
UNet	0.941 ± 0.01	0.482 ± 0.02	0.703 ± 0.01	0.593 ± 0.01
DeepLabV3+	0.936 ± 0.02	0.489 ± 0.02	0.651 ± 0.00	0.570 ± 0.01
SegFormer	$\mathbf{0.950 \pm 0.01}$	$\mathbf{0.527 \pm 0.03}$	$\mathbf{0.710 \pm 0.01}$	$\mathbf{0.618 \pm 0.02}$
SAM	0.930 ± 0.01	0.480 ± 0.02	0.679 ± 0.01	0.580 ± 0.01
MedSAM	0.944 ± 0.01	0.492 ± 0.02	0.671 ± 0.01	0.581 ± 0.01
MedSAM+UNet	0.934 ± 0.01	0.347 ± 0.01	0.591 ± 0.01	0.469 ± 0.01

that model leadership is not a stable property but an artifact of the evaluation choice.

4.2.2. Mean Scores Obscure High Fold-to-Fold Variance

Average performance metrics, while useful, can mask severe performance swings between individual data splits. This is visualized in Figure 5, which plots the performance rank of each model on each fold of the cross-validation. The left panel, showing 3-Fold CV, exhibits relatively consistent rankings. In stark contrast, the right panel for LOOCV reveals extreme rank volatility. For example, **UNet** ranks 2nd on fold 2 but 5th again on fold 5 and in the final fold ranks 1st. This proves that a high average score does not guarantee reliable performance on any single, unseen sample, a critical concern for clinical applications.

4.2.3. Statistical Analysis of Performance Differences

To rigorously test the stability of the observed performance, we conducted a multi-faceted statistical analysis. First, to move beyond simple mean scores, we estimated the uncertainty of each model's performance using bootstrap

Table 6: Analysis of Performance Instability Across Evaluation Contexts. Winning models are color-coded by architectural class (Green: CNN, Blue: Transformer, Red: Foundation Model). The Range (Δ) column is highlighted with a red gradient, where darker shades indicate a larger range and thus higher performance instability.

Metric	Winner	Highest	Lowest	Range (Δ)
Best Single Run (Sweep)				
Macro-Averaged Metrics (Overall Performance)				
Dice (Macro)	**MedSAM+UNet**	0.710	0.000	0.710
IoU (Macro)	**MedSAM+UNet**	0.606	0.000	0.606
Class-Specific Metrics (Tissue-Level Performance)				
Dice (Lumen)	**MedSAM**	0.979	0.000	0.979
Dice (Neointima)	**MedSAM**	0.442	0.391	0.051
Dice (Media)	**MedSAM**	0.797	0.100	0.697
IoU (Lumen)	**MedSAM**	0.959	0.000	0.959
IoU (Neointima)	**MedSAM**	0.391	0.296	0.095
IoU (Media)	**MedSAM**	0.663	0.100	0.563
Leave-One-Out CV (LOOCV)				
Macro-Averaged Metrics (Overall Performance)				
Dice (Macro)	**SegFormer**	0.890	0.580	0.310
IoU (Macro)	**SegFormer**	0.807	0.474	0.333
Class-Specific Metrics (Tissue-Level Performance)				
Dice (Lumen)	**SegFormer**	0.985	0.945	0.040
Dice (Neointima)	**DeepLabV3+**	0.876	0.202	0.674
Dice (Media)	**SegFormer**	0.881	0.479	0.402
IoU (Lumen)	**SegFormer**	0.971	0.731	0.240
IoU (Neointima)	**DeepLabV3+**	0.780	0.112	0.668
IoU (Media)	**SegFormer**	0.787	0.315	0.472
3-Fold CV (K-Fold)				
Macro-Averaged Metrics (Overall Performance)				
Dice (Macro)	**DeepLabV3+**	0.831	0.542	0.289
IoU (Macro)	**DeepLabV3+**	0.725	0.441	0.284
Class-Specific Metrics (Tissue-Level Performance)				
Dice (Lumen)	**MedSAM**	0.980	0.942	0.038
Dice (Neointima)	**DeepLabV3+**	0.770	0.184	0.586
Dice (Media)	**UNet**	0.857	0.438	0.419
IoU (Lumen)	**MedSAM**	0.960	0.891	0.069
IoU (Neointima)	**DeepLabV3+**	0.631	0.123	0.508
IoU (Media)	**UNet**	0.750	0.282	0.468

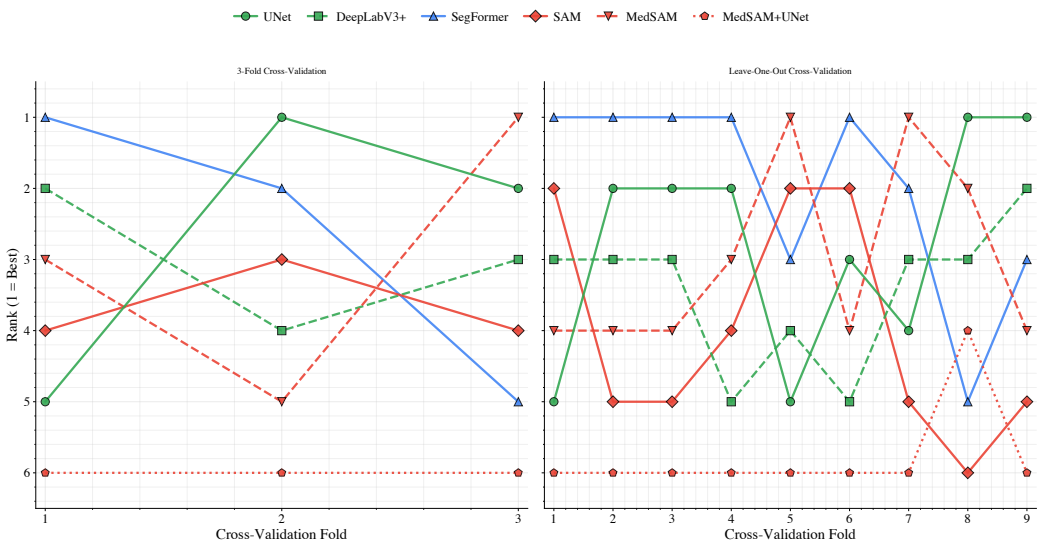

Figure 5: Model Rank Stability in 3-Fold vs. LOOCV. The plots track the performance rank of each model (1=Best) across the validation folds for 3-Fold CV (left) and LOOCV (right).

resampling. By creating 10,000 new sets of fold scores via sampling with replacement, we constructed 95% confidence intervals (CIs) for the true mean macro-Dice score. The results, shown in Table 7 and visualized in Figure 6, reveal substantial overlap between the CIs for the top-performing models across both validation strategies. For instance, under the LOOCV protocol, the CI for SegFormer ([0.764, 0.869]) almost entirely contains the CIs for UNet, MedSAM, SAM, and DeepLabV3+. This overlap indicates that we cannot be statistically confident that any one of these models is genuinely superior to the others.

Next, we performed a non-parametric rank comparison using the Friedman test and Nemenyi post-hoc test, visualized with Critical Difference plots in Figure 7. This analysis directly compares the models' rankings across all folds. While the high-variance LOOCV analysis shows some statistical separation, the more stable 3-Fold CV analysis is not statistically significant overall ($p = 0.221$). Crucially, its CD plot connects all the top-tier models, providing strong evidence that they are statistically indistinguishable.

Finally, to distinguish between statistical significance and practical importance, we calculated the effect size (Cohen's d) for all pairwise comparisons (Table 8 and Figure 8). This analysis shows that most performance differences

Table 7: Bootstrap 95% Confidence Intervals for Macro-Dice Score

	LOOCV		3-Fold CV	
Model	**Mean**	**95% CI**	**Mean**	**95% CI**
UNet	0.807	[0.760, 0.851]	0.756	[0.708, 0.850]
DeepLabV3+	0.796	[0.752, 0.836]	**0.785**	[0.708, 0.837]
SegFormer	**0.821**	[0.764, 0.869]	0.780	[0.656, 0.854]
SAM	0.799	[0.747, 0.843]	0.761	[0.672, 0.813]
MedSAM	0.803	[0.756, 0.843]	0.771	[0.709, 0.807]
MedSAM+UNet	0.735	[0.700, 0.772]	0.672	[0.632, 0.737]

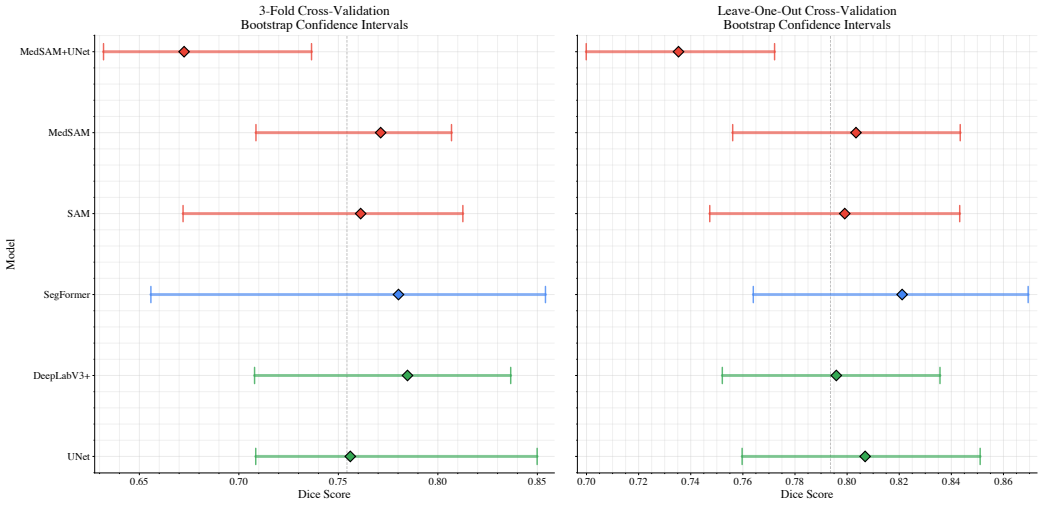

Figure 6: Bootstrap 95% Confidence Intervals for Macro-Dice Score. The plots show the mean Dice score (diamond) and the 95% confidence interval for each model under 3-Fold CV (left) and LOOCV (right). The substantial overlap between the intervals for the top-performing models indicates that their mean performance scores are not statistically different from one another.

were negligible to small, further supporting the conclusion that the observed variations are not practically meaningful. The rank volatility, detailed in Table 9, confirms this; with low Spearman correlations between fold rankings, there is no consistent, reproducible model hierarchy. Taken together, this body of evidence demonstrates that the small performance gaps between models are not statistically significant.

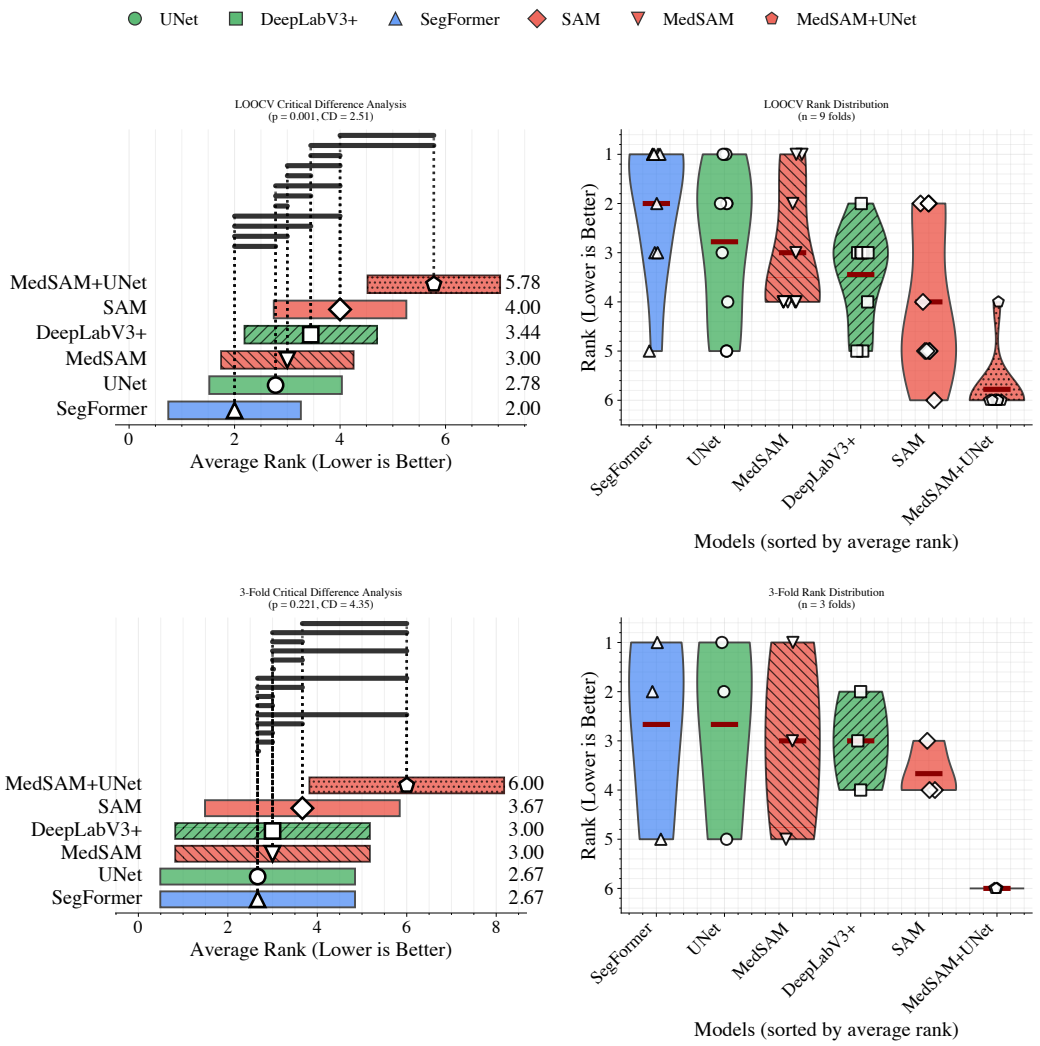

Figure 7: Critical Difference Analysis of Model Ranks. This analysis compares average model ranks under LOOCV (top) and 3-Fold CV (bottom). The plots on the left are Critical Difference (CD) diagrams from the Friedman and Nemenyi tests; models connected by a solid bar are not statistically different. The plots on the right show the distribution of ranks for each model across the validation folds. While the high-variance LOOCV shows some separation, the more stable 3-Fold CV analysis (p=0.221) finds no statistically significant difference among the top-performing models.

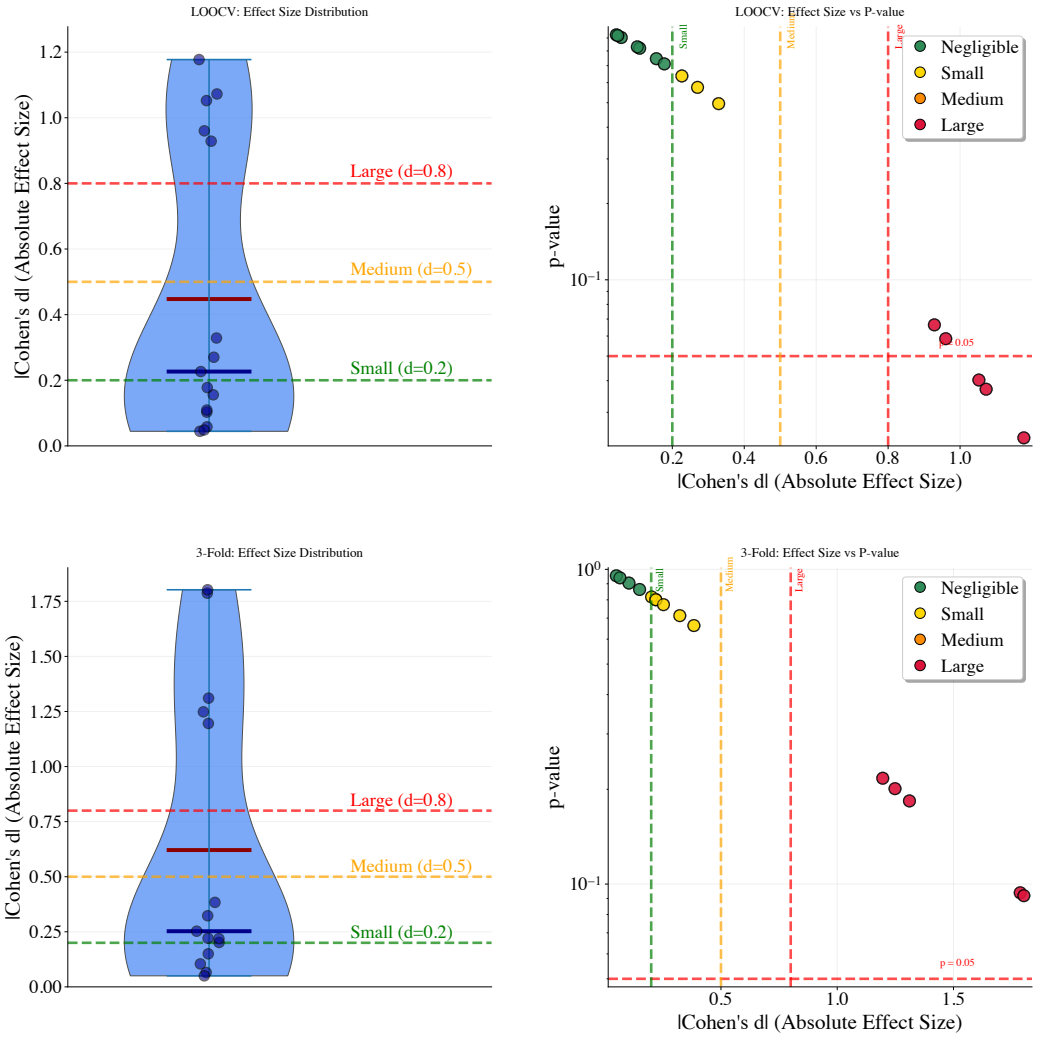

Figure 8: The relationship between practical significance (Cohen's d) and statistical significance (p-value) across two cross-validation protocols. The violin plots (left) show the distribution of absolute effect sizes for Leave-One-Out (LOOCV) and 3-Fold CV, with most comparisons yielding negligible to small effects. The scatter plots (right) show that only the few comparisons with large effect sizes are also statistically significant ($p < 0.05$). This indicates that most observed performance differences between models are neither practically meaningful nor statistically significant.

4.3. The Paradox of Competence: Similar Predictions, Unstable Metrics

A final, subtle paradox emerges when comparing the quantitative chaos with the qualitative visual results. As shown in Figure 9, despite the fluc-

Table 8: Pairwise Effect Size (Cohen's d) for LOOCV Macro-Dice Scores

vs.	DLV3+[a]	SegFormer	SAM	MedSAM	M-UNet[b]
UNet	0.37	-0.25	0.07	-0.22	1.20
DeepLabV3+[a]	–	-0.37	-0.32	-0.22	1.80
SegFormer	–	–	0.20	0.10	1.25
SAM	–	–	–	-0.15	1.31
MedSAM	–	–	–	–	1.79

[a] DLV3+: DeepLabV3+ [b] M-UNet: MedSAM+UNet

Table 9: Comparative Rank Instability Analysis Across Cross-Validation Protocols

	LOOCV (High Variance)					3-Fold CV (Lower Variance)				
Model	$\bar{R}$	R	σ	Top-1 (%)	Top-2 (%)	$\bar{R}$	R	σ	Top-1 (%)	Top-2 (%)
UNet	2.78	4.0	1.56	22.22	55.56	2.67	1.0	0.58	33.33	66.67
DeepLabV3+	3.44	3.0	1.01	0.00	11.11	3.00	2.0	1.00	33.33	33.33
SegFormer	2.00	4.0	1.41	55.56	66.67	2.67	1.0	0.58	33.33	66.67
SAM	4.00	4.0	1.58	0.00	33.33	3.67	1.0	0.58	0.00	33.33
MedSAM	3.00	3.0	1.32	22.22	33.33	2.67	1.0	0.58	0.00	66.67
MedSAM+UNet	5.78	2.0	0.67	0.00	0.00	5.67	1.0	0.58	0.00	0.00

Note: $\bar{R}$ is the mean rank, R is the rank range (Max - Min), and σ is the rank standard deviation.

tuating metrics and lack of statistical significance, the actual segmentation masks produced by the top-performing models are strikingly similar and appear highly competent. The "Model Prediction" column shows that **UNet**, **DeepLabV3+**, and **SegFormer** produce outputs that are nearly indistinguishable from each other and close to the ground truth. This suggests that while our metrics are sensitive enough to detect minor, unstable boundary disagreements, these differences may not be practically or clinically relevant. This "paradox of competence" implies that the search for a single best model based on these metrics may be an illusion of a meaningful difference.

4.4. Quantitative Analysis of XAI Stability

To move beyond a purely qualitative assessment of the "paradox of competence," we conducted a quantitative analysis of the stability of the models' reasoning. We used the generated uncertainty maps from the 9 folds of the LOOCV protocol as a proxy for the model's confidence. For each top-performing model, we calculated the pixel-wise mean and variance of these

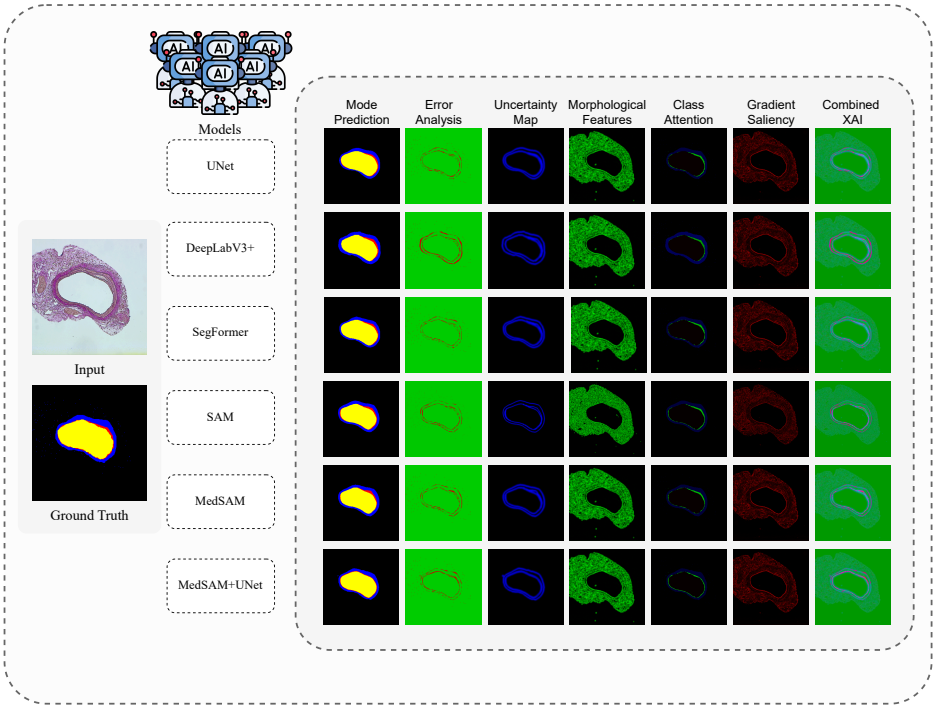

Figure 9: Qualitative Comparison of Model Predictions and Explainable AI (XAI) Outputs, Illustrating the 'Paradox of Competence'.

maps. The results for **SegFormer**, the top-ranked model in the LOOCV protocol, are shown in Figure 10.

The mean uncertainty map (Figure 10A) confirms that, on average, the model's uncertainty is highest at the boundaries between tissue types, which is clinically expected. The crucial finding is in the variance map (Figure 10B), which is nearly black across the entire image. This indicates a quantitatively minimal variance in uncertainty across the 9 independent folds. This provides strong evidence that while the quantitative *performance metrics* were unstable, the model's underlying *reasoning*—its understanding of which regions were easy or difficult—was highly stable and consistent. The instability is therefore confined to superficial metric noise at the boundaries, not a fundamental disagreement in model competence.

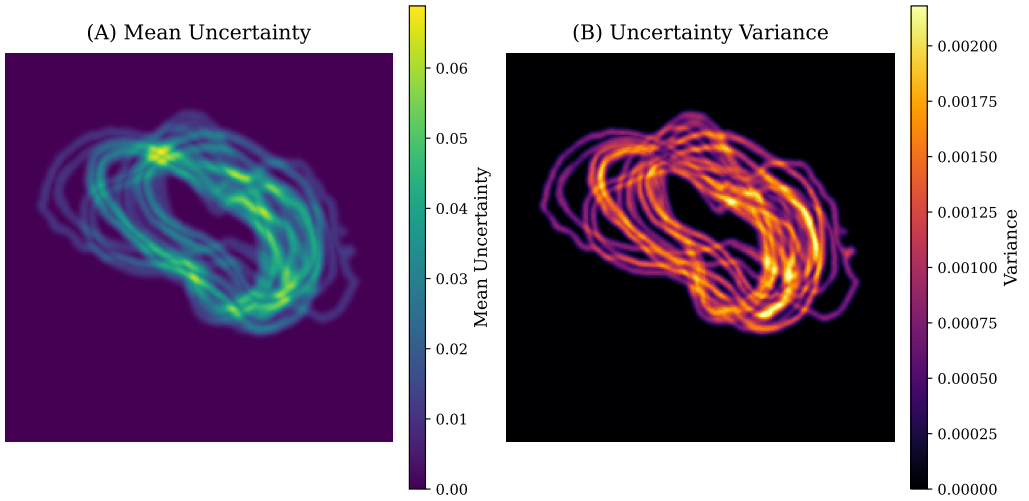

Figure 10: Quantitative XAI Stability Analysis. (A) The mean uncertainty map, averaged across all 9 LOOCV folds, showing that uncertainty is consistently highest at tissue boundaries. (B) The per-pixel variance of uncertainty across the 9 folds. The extremely low variance (dark color) provides quantitative evidence that the model's confidence and reasoning were highly stable, despite fluctuating performance metrics.

5. Discussion

Our results lead to a conclusion that is both simple and profound: on a genuinely small medical dataset, the concept of a single "best" model is a statistical illusion. The research journey, which began with the conventional goal of identifying a top-performing architecture, instead uncovered the inherent fragility of the benchmarking process itself. This section interprets these findings, discusses their broad implications for the medical AI community, and proposes a more robust path forward.

5.1. The Fragility of Leaderboards in Low-Data Regimes

The central finding of our work is that benchmark leadership is an artifact of the evaluation protocol, not an intrinsic property of a model. This instability is driven by a combination of the well-known bias-variance trade-off in validation strategies and the brittleness of hyperparameter optimization. As shown in our results (Figure 5), the high variance of LOOCV created chaotic model rankings sensitive to single data points, while the higher bias of 3-Fold CV may have unfairly penalized more complex models.

This phenomenon provides strong empirical evidence for the *illusion of control* cognitive bias in machine learning research (Langer, 1975). The immense effort invested in rigorous hyperparameter tuning creates a false sense of confidence that randomness has been controlled and a definitive outcome has been reached. However, our findings show that even with optimal configurations, performance is still beholden to the chance composition of data splits. This highlights a critical vulnerability in the standard practice of declaring a state-of-the-art model based on a single leaderboard, particularly when data is scarce.

5.2. When Metrics Mislead: The Paradox of Competence

A critical implication of our work is the "paradox of competence," where quantitative metrics suggest chaos, yet qualitative results show consensus. As seen in the visual comparisons (Figure 9), the top-performing models produce segmentation masks that are strikingly similar to each other and to the ground truth. This suggests that all well-optimized models have learned the same core anatomical features required for the task. The statistical noise detected by the metrics likely arises from minor, pixel-level disagreements at the tissue boundary-disagreements that, while mathematically significant to an overlap score like Dice, may be clinically irrelevant (Maier-Hein et al., 2018).

This paradox forces a shift in how we approach model evaluation. If performance metrics are statistically tied and volatile, we must turn to other criteria. This opens the door for a more mature, engineering-driven approach to model selection. When faced with a pool of statistically equivalent models, the decision should be guided by pragmatic factors such as computational efficiency, memory requirements, and ease of deployment. Our findings provide a clear example: the leading models were statistically indistinguishable on accuracy, but **SegFormer** achieves its top-tier performance with only a fraction of its competitors' parameters. Therefore, choosing **SegFormer** is not a claim of superior accuracy, but a responsible and justifiable engineering decision made after acknowledging the statistical dead heat.

5.3. Generalizability and the Boundary of Statistical Stability

This study was conducted in a deliberately extreme low-data environment ($N = 9$), which serves as a "stress test" for benchmarking practices. By pushing the evaluation to its limit, we establish a valuable "worst-case scenario" for benchmark stability. The catastrophic failure of ranking stability here

begs the question of where the boundary of this "low-data danger zone" lies. If leaderboard rankings are meaningless at $N = 9$, are they reliable at $N = 20$, $N = 50$, or $N = 100$?.

While answering this question definitively requires further experiments, we can theorize based on our results and established statistical principles. The chaotic model rankings, particularly under the high-variance LOOCV protocol (Figure 5), are an expected consequence of small sample sizes. As the number of samples (N) increases, the variance of performance estimates is expected to decrease, which would likely lead to more stable and reproducible model rankings.

However, we hypothesize that even as rankings stabilize with more data, the practical implications of our findings will remain. For instance, at $N = 50$, we might find that SegFormer consistently outperforms UNet. But as our effect size analysis shows (Figure 8), the magnitude of these differences between well-optimized models is often negligible to small. Therefore, even a statistically significant result at a larger N may not translate to a clinically meaningful improvement. This reinforces our central conclusion: the pursuit of a new state-of-the-art score on an unstable benchmark does not automatically equate to a more useful clinical tool.

Thus, we propose a new research agenda focused on mapping this "phase transition" from statistical instability to stability. By systematically repeating this rigorous evaluation framework on nested subsets of larger public datasets, we can establish evidence-based guidelines for the minimum sample sizes required to confidently declare a winning model. Answering this question is essential for improving the rigor and reproducibility of our field and providing clear, actionable guidance to all researchers, especially those working in rare diseases or other data-limited domains.

6. Conclusion

This work demonstrates that on a realistic, data-scarce medical dataset, the pursuit of a single "best" deep learning model is a chase after a statistical mirage. The model rankings we observed were not a reflection of true architectural superiority but were artifacts of the evaluation protocol, statistical noise, and the high variance inherent in small-sample cross-validation. Our findings provide stark, empirical evidence that the common practice of declaring a state-of-the-art model based on a fragile leaderboard, a practice driven by a cognitive illusion of control, can be profoundly misleading in low-data regimes.

We therefore advocate for a paradigm shift in our community's approach to benchmarking. The goal should not be to celebrate incremental gains on an unstable leaderboard, but to transparently report the stability of results across multiple evaluation protocols and statistical tests. Future work should focus not on crowning illusory winners, but on identifying a *set* of statistically equivalent, competent models and then selecting among them based on pragmatic criteria like efficiency and deployability. Success must be redefined: away from the pursuit of a single, unstable top rank and toward a new standard where the honest reporting of instability is considered the highest form of scientific rigor. By embracing this principle, we can build a more trustworthy and reliable foundation for the future of medical AI.

Author contributions: CRediT

PMK: Conceptualization, Data curation, Formal analysis, Methodology, Software, Visualization, Validation, Writing – original draft; **AAP:** Conceptualization, Data curation, Methodology, Validation, Writing – review and editing; **YS:** Conceptualization, Methodology, Validation, Writing – review and editing; **ZL:** Conceptualization, Methodology, Validation, Writing – review and editing; **EAL:** Conceptualization, Resources, Funding acquisition, Providing samples, Project administration, Supervision, Writing – review and editing; **SA:** Conceptualization, Funding acquisition, Methodology, Validation, Supervision, Writing – review and editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

This work is supported by the Innovation Fund Denmark for the project DIREC (9142-00001B) and CNCS—UEFISCDI, project number PN-III-P4-PCE-2021-1680. The content is solely the responsibility of the authors and does not represent the official views of funding sources.

References

Agrawal, P., et al., 2024. A survey on hyperparameter optimization of machine learning models, in: 2024 2nd International Conference on Disruptive Technologies (ICDT), IEEE. pp. 11–15.

Ahmad, H.M., Khan, M.J., Yousaf, A., Ghuffar, S., Khurshid, K., 2020. Deep learning: a breakthrough in medical imaging. Current Medical Imaging 16, 946–956.

Alom, M.Z., Hasan, M., Yakopcic, C., Taha, T.M., Asari, V.K., 2018. Recurrent residual convolutional neural network based on u-net (r2u-net) for medical image segmentation. arXiv preprint arXiv:1802.06955 .

Arasu, R., Arasu, A., Muller, J., 2021. Carotid artery stenosis: An approach to its diagnosis and management. Australian journal of general practice 50, 821–825.

Basla, R., Giulivi, L., Magri, L., Boracchi, G., 2024. An expert-driven data generation pipeline for histological images, in: 2024 IEEE International Symposium on Biomedical Imaging (ISBI), IEEE. pp. 1–5.

Bhagat, M., Bakariya, B., 2025. A comprehensive review of cross-validation techniques in machine learning. IJSAT-International Journal on Science and Technology 16.

Birare, S.D., Swami, D.V., Gaikwad, S.S., Malpani, R.R., 2024. Study of the histopathological changes in the heart in postmortem/autopsy cases in tertiary care center. Heart Views 25, 9–12.

Chen, L.C., Papandreou, G., Schroff, F., Adam, H., 2017. Rethinking atrous convolution for semantic image segmentation. arXiv preprint arXiv:1706.05587 .

Chrysos, G., 2024. Architecture design: From neural networks to foundation models, in: 2024 IEEE 11th International Conference on Data Science and Advanced Analytics (DSAA), IEEE. pp. 1–3.

Cohen, J., 2013. Statistical power analysis for the behavioral sciences. routledge.

Cooper, A.F., Lu, Y., Forde, J., De Sa, C.M., 2021. Hyperparameter optimization is deceiving us, and how to stop it. Advances in Neural Information Processing Systems 34, 3081–3095.

Curaj, A., Zhoujun, W., Staudt, M., Liehn, E.A., 2016. Induction of accelerated atherosclerosis in mice: The" wire-injury" model. MyJoVE Corporation.

Dasari, V.R., Geerhart, B.E., Alexander, D.M., 2022. Increasing image segmentation accuracy on small datasets by merging multiple inferences on augmented images, in: Disruptive Technologies in Information Sciences VI, SPIE. pp. 86–89.

De Matos, J., Ataky, S.T.M., de Souza Britto Jr, A., Soares de Oliveira, L.E., Lameiras Koerich, A., 2021. Machine learning methods for histopathological image analysis: A review. Electronics 10, 562.

Demšar, J., 2006. Statistical comparisons of classifiers over multiple data sets. Journal of Machine learning research 7, 1–30.

Dinc, R., 2023. Post-stenting restenosis in coronary artery diseases-pathophysiology and morphological features of re-stenotic tissue. Biomedical Journal of Scientific & Technical Research 53, 44836–44841.

Efron, B., 1992. Bootstrap methods: another look at the jackknife, in: Breakthroughs in statistics: Methodology and distribution. Springer, pp. 569–593.

Friedman, M., 1937. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. Journal of the american statistical association 32, 675–701.

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770–778.

Hertel, L., Baldi, P., Gillen, D.L., 2022. Reproducible hyperparameter optimization. Journal of Computational and Graphical Statistics 31, 84–99.

Honda, Y., Nishihira, K., Shibata, Y., Asada, Y., 2023. Early-phase restenosis following directional coronary atherectomy and paclitaxel-coated balloon

angioplasty: a histopathological analysis. European Heart Journal 44, 1098–1098.

Hong, G.S., Jang, M., Kyung, S., Cho, K., Jeong, J., Lee, G.Y., Shin, K., Kim, K.D., Ryu, S.M., Seo, J.B., et al., 2023. Overcoming the challenges in the development and implementation of artificial intelligence in radiology: a comprehensive review of solutions beyond supervised learning. Korean Journal of Radiology 24, 1061.

Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H., 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 .

Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al., 2022. Lora: Low-rank adaptation of large language models. ICLR 1, 3.

Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q., 2017. Densely connected convolutional networks, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 4700–4708.

Hubrecht, R.C., Carter, E., 2019. The 3rs and humane experimental technique: implementing change. Animals 9, 754.

Islam, M.A., Mridha, M., Jahin, M.A., Dey, N., 2024. A unified framework for evaluating the effectiveness and enhancing the transparency of explainable ai methods in real-world applications. arXiv preprint arXiv:2412.03884 .

Khalili, E., Priego-Torres, B., Leon-Jimenez, A., Sanchez-Morillo, D., 2024. Automatic lung segmentation in chest x-ray images using sam with prompts from yolo. IEEE Access .

Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al., 2023. Segment anything, in: Proceedings of the IEEE/CVF international conference on computer vision, pp. 4015–4026.

Kong, B., Li, Z., Zhang, S., 2020. Toward large-scale histopathological image analysis via deep learning, in: Biomedical Information Technology. Elsevier, pp. 397–414.

Lai, Y., 2024. Application and effectiveness evaluation of bayesian optimization algorithm in hyperparameter tuning of machine learning models, in: 2024 International Conference on Power, Electrical Engineering, Electronics and Control (PEEEC), IEEE. pp. 351–355.

Langer, E.J., 1975. The illusion of control. Journal of personality and social psychology 32, 311.

Lin, J., 2019. The neural hype and comparisons against weak baselines, in: Acm sigir forum, ACM New York, NY, USA. pp. 40–51.

Lin, J., Campos, D., Craswell, N., Mitra, B., Yilmaz, E., 2021. Significant improvements over the state of the art? a case study of the ms marco document ranking leaderboard, in: Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval, pp. 2283–2287.

Liu, Z., Lv, Q., Lee, C.H., Shen, L., 2024. Segmenting medical images with limited data. Neural Networks 177, 106367.

Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B., 2024. Segment anything in medical images. Nature Communications 15, 654.

Maier-Hein, L., Eisenmann, M., Reinke, A., Onogur, S., Stankovic, M., Scholz, P., Arbel, T., Bogunovic, H., Bradley, A.P., Carass, A., et al., 2018. Why rankings of biomedical image analysis competitions should be interpreted with care. Nature communications 9, 5217.

Makarova, A., Shen, H., Perrone, V., Klein, A., Faddoul, J.B., Krause, A., Seeger, M., Archambeau, C., 2021. Overfitting in bayesian optimization: an empirical study and early-stopping solution, in: 2nd Workshop on Neural Architecture Search (NAS 2021)@ ICLR 2021, NAS 2021.

Matsuhiro, Y., Shutta, R., Nakamura, H., Yasumoto, K., Yasumura, K., Tanaka, A., Matsunaga-Lee, Y., Nakamura, D., Yano, M., Yamato, M., et al., 2020. Histological findings of rapid progression of neoatherosclerosis including calcification in hemodialysis patients. Coronary Artery Disease 31, 479–480.

Müller, D., Soto-Rey, I., Kramer, F., 2022. Towards a guideline for evaluation metrics in medical image segmentation. BMC Research Notes 15, 210.

Paheding, S., Reyes-Angulo, A.A., Alam, M.S., 2024. U-net-based medical image segmentation: A comparative analysis and future trends, in: Computer Vision. Chapman and Hall/CRC, pp. 128–147.

Pandelea-Dobrovicescu, G.R., Prodana, M., Golgovici, F., Ionita, D., Sajin, M., Demetrescu, I., 2020. Surface morphology and histopathological aspects of metallic used cardiovascular cocr stents. Metals 10, 1112.

Patil, D., 2024. Explainable artificial intelligence (xai): Enhancing transparency and trust in machine learning models. Available at SSRN 5057400 .

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al., 2011. Scikit-learn: Machine learning in python. the Journal of machine Learning research 12, 2825–2830.

Ranglani, H., 2024. Empirical analysis of the bias-variance tradeoff across machine learning models. Machine Learning and Applications: An International Journal (MLAIJ) Vol 11.

Roberts, W.C., Everett, B.P., Won, V.S., Kondapalli, N., 2020. Diagnostic usefulness of histological examination of the left ventricular "core" excised to insert a left ventricular assist device in patients with severe heart failure. The American Journal of Cardiology 137, 71–76.

Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation, in: International Conference on Medical image computing and computer-assisted intervention, Springer. pp. 234–241.

Salehi, A.W., Khan, S., Gupta, G., Alabduallah, B.I., Almjally, A., Alsolai, H., Siddiqui, T., Mellit, A., 2023. A study of cnn and transfer learning in medical imaging: Advantages, challenges, future scope. Sustainability 15, 5930.

Shamai, G., Livne, A., Polónia, A., Sabo, E., Cretu, A., Bar-Sela, G., Kimmel, R., 2022. Deep learning-based image analysis predicts pd-l1 status from h&e-stained histopathology images in breast cancer. Nature Communications 13, 6753.

Subaşı, N., 2024. Comprehensive analysis of grid and randomized search on dataset performance. European Journal of Engineering and Applied Sciences 7, 77–83.

Tan, M., Le, Q., 2019. Efficientnet: Rethinking model scaling for convolutional neural networks, in: International conference on machine learning, PMLR. pp. 6105–6114.

Tetko, I.V., van Deursen, R., Godin, G., 2024. Be aware of overfitting by hyperparameter optimization! Journal of Cheminformatics 16, 139.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need. Advances in neural information processing systems 30.

Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al., 2020. Scipy 1.0: fundamental algorithms for scientific computing in python. Nature methods 17, 261–272.

Wang, H., Yang, K., 2023. Bayesian optimization, in: Many-Criteria Optimization and Decision Analysis: State-of-the-Art, Present Challenges, and Future Perspectives. Springer, pp. 271–297.

Xia, K., Wang, J., 2023. Recent advances of transformers in medical image analysis: a comprehensive review. MedComm–Future Medicine 2, e38.

Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P., 2021. Segformer: Simple and efficient design for semantic segmentation with transformers. Advances in neural information processing systems 34, 12077–12090.

Xu, J., Li, M., Zhu, Z., 2020. Automatic data augmentation for 3d medical image segmentation, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 378–387.

Xu, K., Goetz, L., Rajpoot, N., 2024. On generalisability of segment anything model for nuclear instance segmentation in histology images. arXiv preprint arXiv:2401.14248 .

Xu, P., Ji, X., Li, M., Lu, W., 2023. Small data machine learning in materials science. npj Computational Materials 9, 42.

Yang, S., Bi, H., Zhang, H., Sun, J., 2024. Sam-unet: Enhancing zero-shot segmentation of sam for universal medical images. arXiv preprint arXiv:2408.09886 10.

Yeung, M., Rundo, L., Nan, Y., Sala, E., Schönlieb, C.B., Yang, G., 2023. Calibrating the dice loss to handle neural network overconfidence for biomedical image segmentation. Journal of Digital Imaging 36, 739–752.

Yeung, M., Sala, E., Schönlieb, C.B., Rundo, L., 2021. A mixed focal loss function for handling class imbalanced medical image segmentation. arXiv preprint arXiv:2102.04525 .

Yu, H., Yang, L.T., Zhang, Q., Armstrong, D., Deen, M.J., 2021. Convolutional neural networks for medical image analysis: state-of-the-art, comparisons, improvement and perspectives. Neurocomputing 444, 92–110.

Zhou, Z., Siddiquee, M.M.R., Tajbakhsh, N., Liang, J., 2019. Unet++: Redesigning skip connections to exploit multiscale features in image segmentation. IEEE transactions on medical imaging 39, 1856–1867.

Supplementary Materials

Appendix A. Bayesian Hyperparameter Optimization Details

Our hyperparameter optimization strategy is built upon Bayesian Optimization, a powerful technique that efficiently searches the hyperparameter space. This approach is driven by two core components: a surrogate model and an acquisition function.

First, a Gaussian Process (GP) serves as the surrogate model. This GP approximates the complex relationship between hyperparameters (x) and the model's performance score ($f(x)$). After t trials (i.e., evaluations of different hyperparameter configurations), the GP provides a posterior distribution for any new, unevaluated hyperparameter set x_*. This distribution is characterized by a mean predicted score $\mu_t(x_*)$ and an associated uncertainty $\sigma_t^2(x_*)$, represented as $\mathcal{N}(\mu_t(x_*), \sigma_t^2(x_*))$.

Second, an acquisition function leverages this probabilistic model to intelligently decide which set of hyperparameters to evaluate next. We utilize the Expected Improvement (EI) function, which quantifies the expected gain in performance beyond the best score found so far, $f(x^+) = \max(f(x_1), ..., f(x_t))$, where x_t is the hyperparameter set from the t-th trial. The EI function is defined as:

$$\begin{aligned} EI(x) = (\mu_t(x) - f(x^+)) \cdot \Phi\left(\frac{\mu_t(x) - f(x^+)}{\sigma_t(x)}\right) \\ + \sigma_t(x) \cdot \phi\left(\frac{\mu_t(x) - f(x^+)}{\sigma_t(x)}\right) \end{aligned} \tag{S1}$$

In this equation, $\Phi(\cdot)$ represents the cumulative distribution function (CDF) of the standard normal distribution, and $\phi(\cdot)$ is its probability density function (PDF).

At each step of the optimization, the algorithm identifies the hyperparameters that maximize this EI value. This mechanism effectively balances two crucial aspects of optimization: *exploration* (investigating uncertain regions of the hyperparameter space that might contain better solutions) and *exploitation* (focusing on regions already known to yield high performance). This intelligent trade-off allows for efficient discovery of optimal hyperparameter configurations with fewer evaluations compared to traditional grid search or random search methods.