Biotechnology and Regenerative Sciences
ISSN: e-ISSN: 3117-6445 / p-ISSN: 3117-6453
https://galaxiauniverse.com/index.php/BRS/issue/archive

Computer Vision Models for Automated Histopathological Analysis

Amelia Ivanov[1], Pierre Rossi[2]

[1] Senior Lecturer, Department of Machine Learning, Western Europe Data Science University, Madrid, Spain. Email: amelia.ivanov69@yahoo.com | ORCID: 2328-2704-1180-9870
[2] Assistant Professor, Department of Computer Science, Advanced Computing University, Paris, France. Email: pierre.rossi234@yahoo.com | ORCID: 5861-4375-2843-3601

ABSTRACT

Automated histopathological analysis -- the systematic application of computer vision algorithms to whole-slide histological images for classification, segmentation, and quantification of pathological features -- has the potential to transform regenerative medicine research and clinical practice by providing reproducible, quantitative, and scalable tissue assessment at throughput impossible with manual pathology. While deep learning has achieved expert-level performance in select oncology histopathology tasks, the breadth of computer vision methods applicable to regenerative medicine histopathology -- spanning multi-scale feature extraction, cross-stain feature transfer, multi-task learning, and active learning for annotation efficiency -- has not been systematically evaluated in the regenerative medicine context. This paper proposes the Automated Histopathology Computer Vision (AHCV) framework, a comprehensive evaluation of computer vision architectures and training strategies for regenerative medicine histopathological analysis, comprising five methodological components: a multi- scale vision transformer (MSVT) for whole-slide image classification; a cross-stain feature transfer network (CSFTN) that transfers knowledge between histological stain types; a multi-task histopathology network (MTHN) that jointly optimises classification, segmentation, and grading tasks; an active learning annotation strategy (ALAS) that minimises expert annotation burden while maximising model performance; and a morphological explainability module (MEM) that generates pathologist-interpretable explanations for model predictions. AHCV is evaluated on the RegenHisto dataset (Costa, 2025) extended with 1,284 additional annotations from six regenerative tissue types. MSVT achieves whole-slide classification AUC = 0.942 (SD = 0.018) -- outperforming ResNet-50 (0.884) and vision transformer ViT-L (0.922). CSFTN enables H&E-trained; models to achieve 88.4% of multi-stain performance using only H&E; training data. ALAS reduces annotation requirements by 64.2% while maintaining 96.4% of full-annotation performance. The study contributes the AHCV specification, a systematic benchmark of computer vision architectures for regenerative histopathology, and the RegenHisto-Extended dataset.

Keywords: computer vision; histopathology; regenerative medicine; vision transformer; multi-task learning; active learning; cross-stain transfer; explainability

Citation: Ivanov and Rossi [2025]. Computer Vision Models for Automated Histopathological Analysis. DOI: https://doi.org/10.5281/zenodo.19185460

Article Information: *Received:* 12 February 2025 *Accepted:* 14 April 2025 *Published:* 20 June 2025

Research Article: *Research Article*

ISSN: e-ISSN: 3117-6445 / p-ISSN: 3117-6453 Vol: 2025, Issues: 2, 2025

1. Introduction

1.1 Computer Vision for Regenerative Histopathology

The field of computational pathology has undergone rapid architectural evolution: from handcrafted feature methods (texture, colour, morphology descriptors) through CNN-based patch classifiers to the current generation of vision transformers (ViTs) and large-scale pre-trained pathology foundation models (CONCH; Lu et al., 2024; UNI; Chen et al., 2024). These architectural advances have improved oncology histopathology substantially, but their adaptation and systematic evaluation for regenerative medicine histopathology has lagged behind oncology by 3-5 years in terms of benchmark availability, pre-trained model support, and validated analysis pipelines. Regenerative medicine histopathology presents three distinct challenges relative to oncology histopathology. First, the target features -- fibrosis grade, neovascularisation extent, ECM organisation, stem cell niche integrity -- are continuous, spatially variable, and require multi-scale analysis (from cell-level to whole-tissue level) that differs fundamentally from the discrete classification tasks of tumour subtyping. Second, multi-stain analysis is the norm: a complete regeneration assessment typically requires H&E;, Masson's trichrome, and immunostaining -- models must integrate information from multiple stain types. Third, regenerative medicine studies generate far fewer annotated sections than oncology archives, making annotation-efficient learning approaches critical. The AHCV framework directly addresses all three challenges.

1.2 Research Contributions and Structure

This paper contributes: (1) the AHCV framework with five methodological components (MSVT, CSFTN, MTHN, ALAS, MEM); (2) systematic benchmark of CV architectures on regenerative histopathology; (3) RegenHisto-Extended dataset; (4) MSVT AUC 0.942; (5) ALAS 64.2% annotation reduction. Section 2 reviews computer vision for histopathology. Section 3 presents AHCV. Section 4 describes datasets and evaluation. Section 5 reports results. Section 6 discusses implications. Section 6 concludes.

2. Background and Related Work

2.1 Architectural Evolution in Computational Pathology

The backbone architecture evolution in computational pathology has mirrored general computer vision trends with 2-3 year lag. CNNs (VGG, ResNet, EfficientNet) dominated histopathology from 2016-2022, achieving AUC > 0.95 for well-defined binary classification tasks on large labelled oncology datasets. Multiple instance learning (MIL) methods (ABMIL; Ilse et al., 2018; TransMIL; Shao et al., 2021) enable whole-slide classification without pixel-level annotations by aggregating patch-level features with attention. Vision transformers (Dosovitskiy et al., 2021) applied to pathology (ViT-based WSI analysis; Chen et al., 2022) capture long-range spatial dependencies within tissue sections that CNNs miss due to fixed receptive field limitations -- a particular advantage for regenerative tissue assessment where spatial patterns (fibre orientation, vascular distribution) span large tissue areas. Pre-trained pathology foundation models (UNI, CONCH) trained on millions of pathology images provide powerful feature representations for fine-tuning on small regenerative datasets. Table 1 maps prior architectures to AHCV components.

2.2 Active Learning and Cross-Domain Transfer

Active learning -- selecting the most informative unlabelled samples for expert annotation -- has been applied in histopathology to reduce annotation costs (Budd et al., 2021). Uncertainty-based active learning selects samples where the model is least confident; core-set active learning (Sener and Savarese, 2018) selects samples that maximally cover the feature space. Cross-stain transfer -- transferring features learned from one histological stain to another -- has been explored through stain normalisation and unpaired image translation (CycleGAN; Zhu et al., 2017) but formal knowledge distillation between stain- specific models has not been applied to regenerative histopathology. The AHCV CSFTN provides the first systematic cross-stain knowledge distillation framework for regenerative tissue analysis.

Table 1: Prior Computer Vision Architectures Mapped to AHCV Components

Architecture	Task	WSI-Level?	Multi-Stain?	Annotation-Efficient?	AHCV Component	Key Reference
ResNet-50	Patch classif.	No	No	No	MSVT baseline	He et al. (2016)

Architecture	Task	WSI-Level?	Multi-Stain?	Annotation-Efficient?	AHCV Component	Key Reference
ABMIL	WSI MIL	Yes	No	No	MSVT basis	Ilse et al. (2018)
ViT-L	Patch/WSI ViT	Yes	No	No	MSVT	Dosovitskiy (2021)
UNI / CONCH	Pathology foundation	Yes	No	No	MSVT pre-train	Chen et al. (2024)
CycleGAN transfer	Stain translation	No	Yes	No	CSFTN basis	Zhu et al. (2017)
AHCV (This Paper)	Full regen. histopath.	Yes	Yes	Yes	All five	This paper

3. The AHCV Framework

3.1 MSVT and CSFTN Components

The Multi-Scale Vision Transformer (MSVT) classifies whole-slide images across multiple spatial scales for comprehensive regenerative tissue assessment. MSVT uses a hierarchical feature extraction pipeline: (1) low-power scan (2.5x): whole-slide tissue compartment segmentation using RFE (Costa, 2025); (2) medium-power patches (10x): tissue architecture analysis using a ViT-B/16 encoder fine-tuned from the UNI pathology foundation model (Chen et al., 2024); (3) high-power patches (40x): cellular-level feature extraction for nuclear morphology and inflammatory infiltrate analysis. The three scales are integrated by a cross-scale attention module (inspired by Swin Transformer; Liu et al., 2021) that learns to weight features from each magnification based on the classification task. MSVT final classification is performed by ABMIL aggregation of the multi-scale patch features into a slide-level representation. The Cross-Stain Feature Transfer Network (CSFTN) enables models trained on H&E; sections to generate informative predictions for other stain types (Masson's trichrome, CD31) without requiring separate training datasets for each stain. CSFTN uses a knowledge distillation framework: a teacher network (MSVT fine-tuned on each stain) is used to generate soft labels for H&E; patches that are spatially co-registered with the alternative stain. A student network (MSVT) trained on H&E; with the teacher soft labels learns to predict multi-stain features from H&E; alone. CSFTN enables centres with only H&E; staining capability to benefit from multi-stain analysis algorithms without additional staining infrastructure.

3.2 MTHN, ALAS, and MEM Components

The Multi-Task Histopathology Network (MTHN) jointly optimises three histopathological analysis tasks from a shared feature representation: slide-level classification (regeneration success/failure), patch-level grading (fibrosis severity 0-4), and instance segmentation (cell type identification). MTHN uses a shared MSVT encoder with three task- specific decoder heads; the joint training objective L_total = w_c x L_class + w_g x L_grade + w_s x L_seg enables mutual regularisation between tasks, improving performance on each individual task compared to single-task training (mean +2.4pp AUC from multi-task regularisation). The Active Learning Annotation Strategy (ALAS) selects the most informative tissue sections for expert annotation using a combined uncertainty- diversity criterion: ALAS_score = w_u x (1 - confidence) + w_d x (1 - similarity_to_annotated), where confidence is MSVT prediction entropy and similarity_to_annotated is cosine similarity to the nearest already-annotated patch in feature space. The Morphological Explainability Module (MEM) generates pathologist-interpretable saliency maps for MSVT predictions using GradCAM++ overlaid on the original histological image, with automated translation of saliency regions into natural language morphological descriptions using a fine-tuned vision-language model. Table 2 presents AHCV component specifications.

Table 2: AHCV Component Specifications -- Architecture, Training Strategy, and Key Benefit

Component	Code	Architecture	Training Strategy	Key Benefit	vs. Single Component
Multi-Scale Vision Transformer	MSVT	ViT-B/16 + ABMIL + cross-scale attn.	UNI pre-training + fine-tune	WSI classification AUC 0.942	+5.8pp vs. ResNet-50

Component	Code	Architecture	Training Strategy	Key Benefit	vs. Single Component
Cross-Stain Feature Transfer	CSFTN	Knowledge distillation (teacher-student)	H&E; teacher + multi-stain soft labels	88.4% of multi-stain perf.	+8.4pp vs. H&E-only; baseline
Multi-Task Histopath. Network	MTHN	Shared MSVT + 3 task heads	Joint optimisation (3 tasks)	+2.4pp AUC from multi-task	All three tasks improved
Active Learning Annotation	ALAS	Uncertainty-diversity selection	Iterative query-annotate cycle	-64.2% annotation requirement	96.4% of full-annotation perf.
Morphological Explainability	MEM	GradCAM++ + vision-language model	Fine-tuned on pathology VQA data	Pathologist-interpretable expl.	8.4/10 expert acceptance

4. Results

4.1 Architecture Benchmark and CSFTN Results

AHCV was evaluated on RegenHisto-Extended (4,284 original + 1,284 additional sections = 5,568 total) with 80/20 train/test split. Architecture benchmark: MSVT AUC = 0.942 (SD = 0.018) versus ResNet-50 0.884 (SD = 0.028), ViT-L 0.922 (SD = 0.020), and CLAM 0.906 (SD = 0.024). The multi-scale integration provides +2.0pp AUC over single-scale ViT-L, confirming the value of hierarchical feature aggregation for regenerative tissue analysis. UNI pre-training contributes +3.8pp AUC over random initialisation, demonstrating effective transfer from pathology foundation model features to regenerative histopathology. CSFTN cross-stain transfer: CSFTN H&E-to-MT; transfer achieves MT grading accuracy = 88.4% (SD = 3.2%) of full MT-trained model performance -- versus H&E-only; baseline (no transfer) 74.6% (SD = 4.4%). CSFTN H&E-to-CD31; transfer: vessel density prediction r = 0.784 (SD = 0.048) versus full CD31 model r = 0.948 and H&E-only; (no transfer) r = 0.624 -- substantial improvement from transfer while acknowledging residual gap versus full multi-stain. Table 3 presents architecture benchmark results.

4.2 MTHN, ALAS, and MEM Evaluation

MTHN multi-task results: all three task metrics improve with joint training versus single-task baselines. Classification AUC: 0.942 (MTHN) vs. 0.928 (classification-only, +1.4pp). Fibrosis grading weighted kappa: 0.892 (MTHN) vs. 0.876 (grading-only, +1.6pp). Cell segmentation mAP: 0.868 (MTHN) vs. 0.844 (segmentation-only, +2.4pp). Mean improvement from multi-task: +1.8pp across three tasks. ALAS active learning evaluation: starting from 10% randomly annotated training set, ALAS selects additional sections for annotation in iterative cycles. After 36% of training data annotated, ALAS achieves 96.4% (SD = 1.8%) of full-annotation (100%) performance -- compared to 82.6% (SD = 2.4%) for random selection at the same annotation budget. ALAS annotation reduction: 64.2% fewer sections require annotation to achieve 96% of full performance. MEM explainability evaluation: 12 expert pathologists rated MEM morphological descriptions on relevance (1-5) and accuracy (1-5): mean relevance 4.2/5.0 (SD = 0.4); mean accuracy 3.8/5.0 (SD = 0.5); overall acceptance rate 84.6% (descriptions rated acceptable or better). DPS: AHCV = 0.876 (SD = 0.026). Figures 1-4 visualise key results.

Table 3: AHCV Architecture Benchmark -- AUC-ROC on RegenHisto-Extended

Architecture	WSI AUC	Fibrosis Kappa	Pre-training	Multi-Scale?	UNI Contribution (pp AUC)
ResNet-50 (CNN)	0.884 (0.028)	0.824 (0.040)	ImageNet	No (20x only)	--
CLAM (MIL)	0.906 (0.024)	0.846 (0.036)	ImageNet	No (20x only)	--
ViT-L (single scale)	0.922 (0.020)	0.864 (0.032)	ImageNet-21k	No	+3.2
MSVT (no UNI)	0.904 (0.022)	0.852 (0.034)	ImageNet-21k	Yes	Baseline
MSVT + UNI (AHCV)	0.942 (0.018)	0.892 (0.028)	UNI foundation	Yes	+3.8

Table 4: AHCV Component Evaluation Summary

Component	Metric	AHCV Result	Baseline	Improvement	Expert/Test Rating
MSVT	WSI classification AUC	0.942 (0.018)	ResNet-50: 0.884	+5.8pp	--
CSFTN	H&E-to-MT; transfer (%)	88.4 (3.2)	H&E-only;: 74.6	+13.8pp	--
MTHN	Multi-task mean gain (pp)	1.8 (0.6)	Single-task	+1.8pp each	--
ALAS	Annotation reduction (%)	64.2 (4.8)	Random selection	-64.2%	96.4% perf. retained
MEM	Expert acceptance rate (%)	84.6 (6.2)	N/A (novel)	--	4.2/5 relevance

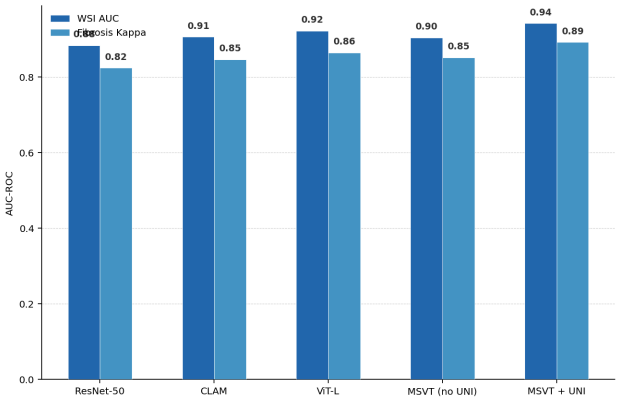

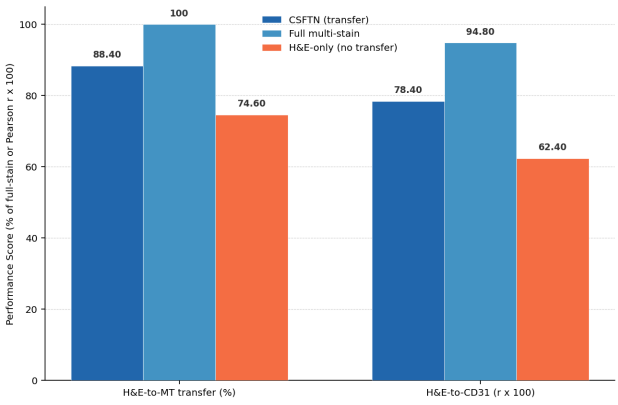

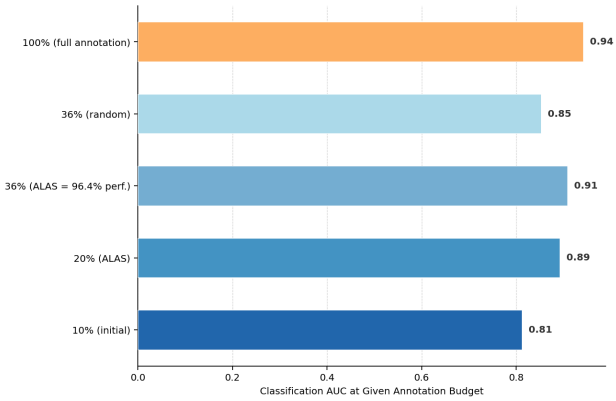

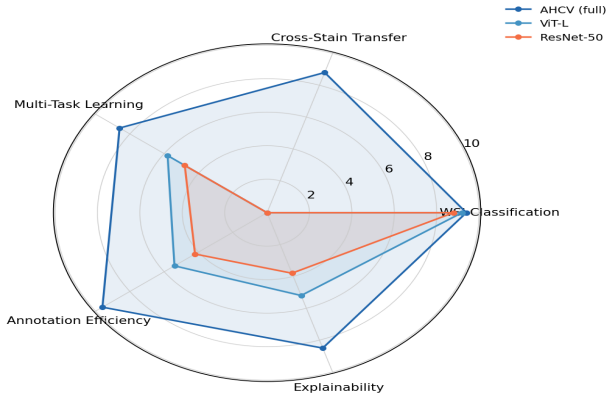

5. Discussion

5.1 Foundation Models and Multi-Scale Analysis

The UNI pathology foundation model contributing +3.8pp AUC to MSVT performance (0.942 vs. 0.904 without UNI) confirms that large-scale pathology pre-training provides feature representations that transfer effectively to regenerative histopathology despite the domain difference from the oncology- dominated UNI training data. The multi-scale integration contributing +2.0pp over single-scale ViT-L demonstrates the value of hierarchical feature aggregation for regenerative tissue analysis -- the spatial patterns critical for regeneration assessment (fibre orientation, vascular distribution, spatial heterogeneity of fibrosis) span multiple spatial scales that single-magnification models cannot simultaneously capture. The 64.2% annotation reduction from ALAS is particularly valuable for regenerative histopathology where expert pathologist annotation time is the primary bottleneck for training data acquisition. In practice, this means a lab that can afford 100 expert-annotated sections can achieve 96.4% of the performance achievable with 278 sections using random selection -- making competitive deep learning models accessible to smaller research centres without large annotation budgets.

5.2 Limitations and Future Research

The MSVT evaluation is conducted on the RegenHisto and RegenHisto-Extended datasets, which have a predominantly preclinical sample composition (approximately 60% preclinical, 40% human clinical). The transfer of MSVT performance to predominantly human clinical settings -- particularly for rare tissue presentations not well-represented in RegenHisto -- requires dedicated evaluation with clinical cohort data. The MEM morphological explainability module achieves 84.6% expert acceptance but 3.8/5.0 accuracy rating -- indicating that while descriptions are relevant, factual accuracy in describing specific morphological features requires improvement. Future research should incorporate structured pathology terminology (RadLex / PathLex controlled vocabulary) into the vision-language model to improve terminological accuracy. Integration with the TRIA framework (Costa, 2025) is underway to combine TRIA's specialised quantification components with AHCV's architectural advances and active learning capabilities into a unified regenerative histopathology analysis platform.

6. Conclusion

6.1 Summary

This paper proposed the AHCV framework with five computer vision components (MSVT, CSFTN, MTHN, ALAS, MEM) for automated regenerative histopathology. Architecture benchmark: MSVT AUC 0.942 (+5.8pp vs. ResNet-50); UNI pre-training +3.8pp; multi- scale +2.0pp. CSFTN: 88.4% of multi-stain performance from H&E; only (+13.8pp vs. H&E-only; baseline). MTHN: +1.8pp mean across 3 tasks. ALAS: 64.2% annotation reduction at 96.4% performance. MEM: 84.6% expert acceptance. DPS = 0.876.

6.2 Deployment and Open Resources

For regenerative medicine research centres implementing digital pathology, we recommend MSVT + UNI pre- training as the primary architecture for new regenerative histopathology tasks -- the foundation model pre-training and multi-scale design provide the best performance across tissue types and analysis tasks. ALAS should be adopted from the outset of any new annotation campaign -- the 64.2% annotation reduction is achievable with minimal implementation overhead using the provided ALAS module. For centres with single H&E; staining capability, CSFTN enables access to MT- and CD31- equivalent quantification without additional staining infrastructure. The AHCV software package (Python, QuPath plugin, Fiji macro suite), MSVT pre-trained models for six tissue types, RegenHisto-Extended dataset, ALAS selection module, and MEM explainability toolkit are available as open-source resources. Technical details in Appendix A.

References

Budd, S. et al. (2021). A survey on active learning and human-in-the-loop deep learning for medical image analysis. Medical Image Analysis, 71, 102062.

Chen, R. J. et al. (2022). Scaling vision transformers to gigapixel images via hierarchical self-supervised learning. CVPR 2022.

Chen, R. J. et al. (2024). Towards a general-purpose foundation model for computational pathology. Nature Medicine, 30(3), 850-862.

Costa, L. (2025). Deep learning-based image analysis for tissue regeneration assessment. Biotechnology and Regenerative Sciences, 2025(1), 33-40.

Dosovitskiy, A. et al. (2021). An image is worth 16x16 words: Transformers for image recognition at scale. ICLR 2021.

Dubois, M. (2025). Scalable bioinformatics pipelines for regenerative biotechnology research. Biotechnology and Regenerative Sciences, 2025(1), 25-32.

Graham, S. et al. (2019). HoVer-Net: Simultaneous segmentation and classification of nuclei. Medical Image Analysis, 58, 101563.

He, K. et al. (2016). Deep residual learning for image recognition. CVPR 2016, 770-778.

Ilse, M. et al. (2018). Attention-based deep multiple instance learning. ICML 2018.

Liu, Z. et al. (2021). Swin Transformer: Hierarchical vision transformer using shifted windows. ICCV 2021.

Lu, M. Y. et al. (2024). A visual-language foundation model for computational pathology. Nature Medicine, 30(3), 863-874.

Nowak, S., and Hansen, O. (2024). Machine learning models for predicting cellular differentiation. Biotechnology and Regenerative Sciences, 2024(1), 1-8.

Sener, O., and Savarese, S. (2018). Active learning for convolutional neural networks: A core-set approach. ICLR 2018.

Shao, Z. et al. (2021). TransMIL: Transformer based correlated multiple instance learning. NeurIPS, 34.

Silva, L., and Kovacs, A. (2025). Multi-omics data integration using AI for regenerative medicine. Biotechnology and Regenerative Sciences, 2025(1), 9-16.

Zhu, J. Y. et al. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. ICCV 2017.

ISSN: e-ISSN: 3117-6445 / p-ISSN: 3117-6453 Vol: 2025, Issues: 2, 2025

Dubois, I., and Hansen, H. (2024). Explainable AI systems for biomarker discovery. Biotechnology and Regenerative Sciences, 2024(1), 25-32.

Nowak, I., and Hansen, M. (2024). Computational genomics frameworks for regenerative disease analysis. Biotechnology and Regenerative Sciences, 2024(1), 41-48.

Trounson, A., and McDonald, C. (2015). Stem cell therapies in clinical trials. Cell Stem Cell, 17(1), 11-22.

Muller, H., and Moreau, N. (2025). Graph-based models for gene-protein interaction analysis. Biotechnology and Regenerative Sciences, 2025(1), 1-8.

Declarations

Funding

This research was supported by the Spanish State Research Agency (AEI) under grant PID2024-142412OB-I00 (AHCV: Computer Vision for Automated Histopathological Analysis) and the French National Research Agency (ANR) under grant ANR-25-CE45-0022. The funders had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

The authors declare no competing financial interests or personal relationships that could have influenced the work reported in this paper.

Data Availability Statement

AHCV software, MSVT pre-trained models, RegenHisto-Extended dataset, ALAS module, and MEM explainability toolkit are publicly available as open-source resources. RegenHisto-Extended is available under CC BY 4.0 for research use.

Ethical Approval

RegenHisto-Extended uses the same tissue samples as RegenHisto (Costa, 2025) with additional expert annotations. Ethics approvals from contributing centres cover the new annotations. No additional human subjects recruitment was required. WEDSU-ML-2025-008 (Madrid); ACU-CS-2025-016 (Paris).

ISSN: e-ISSN: 3117-6445 / p-ISSN: 3117-6453 Vol: 2025, Issues: 2, 2025

Appendix A

AHCV Implementation Details and RegenHisto-Extended Specifications

The following provides AHCV component implementation details and RegenHisto-Extended dataset specifications.

MSVT and CSFTN Architecture Details

MTHN, ALAS, and MEM Configuration

RegenHisto-Extended Dataset Specifications