

No Masks Needed: Explainable AI for Deriving Segmentation from Classification

Mosong Ma¹, Tania Stathaki¹, and Michalis Lazarou²

¹Dept. Electrical and Electronic Engineering, Imperial College London,
London, UK

²Centre for Vision, Speech and Signal Processing, University of Surrey,
Guildford, UK

Abstract. *Medical image segmentation is vital for modern healthcare and is a key element of computer-aided diagnosis. While recent advancements in computer vision have explored unsupervised segmentation using pre-trained models, these methods have not been translated well to the medical imaging domain. In this work, we introduce a novel approach that fine-tunes pre-trained models specifically for medical images, achieving accurate segmentation with extensive processing. Our method integrates Explainable AI to generate relevance scores, enhancing the segmentation process. Unlike traditional methods that excel in standard benchmarks but falter in medical applications, our approach achieves improved results on datasets like CBIS-DDSM, NulnsSeg and Kvasir-SEG.*

Keywords: *Medical Image Segmentation, Explainable AI, Transfer Learning.*

1 Introduction

Medical image segmentation has become a focal point in research due to its potential to revolutionize healthcare by enabling faster and more accurate diagnoses [1]. This task, however, is challenging, particularly because deep learning methods rely heavily on large, well-annotated datasets. In the medical domain, such datasets are difficult to obtain because of privacy concerns and the need for expert-level annotations. These annotations require precise pixel-level labelling, which is time-consuming and costly, often demanding physicians to divert time from patient care.

To overcome these challenges, recent approaches have explored using weakly supervised learning [2] with weak annotations or applying transfer learning [3] from pre-trained models. Our work focuses on the latter, building on the idea of fine-tuning pre-trained classification models specifically for medical tasks. While prior studies [4,5,6] have used pre-trained models for segmentation without fine-tuning, achieving notable results in standard image datasets, these methods struggle or fail in medical image segmentation.

Our approach addresses these limitations by fine-tuning a linear classifier on the medical dataset and utilizing Explainable Artificial Intelligence (XAI) [7] to interpret the model's decision-making process. This interpretation is converted into segmentation masks, providing a more human-readable and clinically useful output. By doing so, we demonstrate that




Fig. 1: General Pipeline for ExplainSeg

fine-tuning, combined with XAI, can enhance the performance of pre-trained models in the medical imaging domain, leading to more effective and interpretable segmentation results.

Our contributions are summarized below:

- We propose a novel method for medical image segmentation (ExplainSeg).
- We show the efficacy of XAI in medical segmentation by simply fine-tuning a linear probing classifier on top of a pre-trained model
- ExplainSeg achieves state-of-the-art performance and significantly outperforms other segmentation methods using pre-trained medical image segmentation models.

2 Related work

2.1 Medical image Segmentation

Medical image segmentation [8] has been widely studied, with various methods developed to address its challenges. Early approaches relied on manual or semi-automated techniques, which require significant human input, making them time-consuming and prone to variability between observers, limiting their use in large-scale studies. Segmentation methods [9] can be broadly categorized into traditional techniques, such as edge-based, region-based, and thresholding [10], along with machine learning methods, which include supervised, semi-supervised, and unsupervised approaches.

2.2 Transfer Learning Segmentation

Transfer learning [3] involves leveraging knowledge from a pre-trained model, typically trained on a large dataset, and applying it to a different but related task. This approach has become popular in image segmentation tasks, particularly when labelled data is scarce or models must generalize across different domains.

Several methods have successfully applied transfer learning to segmentation tasks. TokenCut [4] is a notable example, where a Vision Transformer (ViT) [11] is pre-trained through the self-supervised learning method of DINO [12] and used to segment objects by treating image patches as tokens and optimizing a normalized cut criterion. Cut-and-LEaRn (CutLER) [5] and FOUND [6] also employ pre-trained networks to analyze background regions and generate coarse object proposals.

2.3 Explainable Artificial Intelligence (XAI)

Machine learning has significantly advanced medical image segmentation, with deep learning models like U-Net becoming the standard for their ability to learn complex features. However, these models are often viewed as "black boxes" due to their opaque decision-making processes.

Explainable AI (XAI) is increasingly vital in addressing this issue. Techniques like integrated gradients [13] and Grad-CAM [14] provide visual explanations, helping users understand which parts of an image influenced the model's predictions. While these tools enhance the balance between performance and interpretability, challenges remain in ensuring that these explanations are accurate and relevant.

2.4 XAI to Segmentation

Our approach leverages deep learning and Explainable AI (XAI) techniques to enhance interpretability and reduce reliance on pixel-level annotations in medical image segmentation. Using transfer learning, we adapt a pre-trained model to our dataset and employ XAI to generate segmentation masks, ensuring both accuracy and transparency.

Previous works, such as [15] and [16], leverage XAI to create segmentation masks without pixel-level labels in generic domains like infrastructure imagery. In contrast, our work focuses on the medical imaging domain, addressing its higher precision and robustness requirements. By tailoring XAI methods to weakly annotated medical data, we demonstrate their potential to deliver accurate and meaningful segmentation results, highlighting the novelty of our approach.

3 Method

3.1 Problem formulation

At the fine-tuning stage, we assume access to a labelled dataset, $D_{\text{train}} = \{\mathbf{x}_i, y_i\}_{i=1}^N$, of N image-label pairs, where every image $\mathbf{x}_i$ has a label $y_i \in Y$, that belongs to the label set Y . We fine-tune a classification network F on D_{train} .

This fine-tuned network F is then trained with an extra classification head G to generate and interpret classification results with the help of XAI. The generated relevance map is refined into segmentation masks with extensive post-processing including morphology and other segmentation methods.

3.2 Optional Fine Tuning & Training

When working with a new dataset, fine-tuning [17] is necessary if a pre-trained classification network struggles to classify accurately. To minimize computational demands, gradients are typically frozen for most of the hidden layers. A task-specific classification layer can also be added on top of the deep learning model, and its parameters are trainable allowing the model to adapt to the specific medical image dataset.

3.3 Integrated Gradients (IG)

Integrated gradients [13] trace how the network determines an image's class by assigning relevance scores to each pixel that indicate its contribution to the classification result, and is chosen due to its model agnostic feature making it compatible to most models.

The integrated gradient along the i -th dimension for an input $\mathbf{x}$ and baseline $\mathbf{x}'$ is defined as follows:

$$\text{IG}_i(\mathbf{x}) = (\mathbf{x}_i - \mathbf{x}'_i) \times \int_{\alpha=0}^1 \frac{\partial F(\mathbf{x}' + \alpha \times (\mathbf{x} - \mathbf{x}'))}{\partial \mathbf{x}_i} d\alpha \quad (1)$$

where α is a scaling factor that interpolate between the baseline $\mathbf{x}'$ and the input $\mathbf{x}$, $\frac{\partial F(\mathbf{x})}{\partial x_i}$ is the gradient of $F(\mathbf{x})$ along the i -th dimension.

The Noise Tunnel function in Captum [18] is an extension of existing attribution methods used for explainable AI (XAI) by adding stochastic noise to input samples to enhance the robustness and reliability of the explanations. It applies Gaussian noise to the input data in multiple iterations, averaging the resulting attribution maps. This process helps reduce the model sensitivity to small perturbations and highlights consistent patterns in the attributions. As a result, Noise Tunnel produces more stable and interpretable explanations.

We did not employ other XAI techniques such as Grad-CAM, as preliminary experiments have consistently shown that the coarse heatmaps produced by Grad-CAM offer significantly inferior spatial localization compared to IG, particularly for fine-grained structures in medical images. In our experiments, Grad-CAM failed to highlight diagnostically relevant regions with sufficient resolution or fidelity. Given these limitations and the superior performance of IG in our setting, further exploration of Grad-CAM and related methods was deemed unnecessary.

3.4 Feature Map Extraction

Alongside integrated gradients for relevance scores, we also explored extracting and visualizing feature maps from the ViT network [19].

Feature map extraction is a key process in deep learning. It captures the essential patterns from an input image. As the image passes through the layers of a network, each layer applies filters that detect specific features and identify patterns relevant to tasks like object recognition or segmentation.

The result of this filtering process is the feature map, which highlights the areas of the image where the features were detected. Feature maps play a crucial role in helping the network understand the content of an image and perform tasks like classification or localization.

3.5 Image Segmentation

The relevance maps generated by XAI methods are insufficient for accurate segmentation mask generation. Therefore, additional post-processing is necessary to refine them into precise segmentation masks.

Otsu Thresholding The simplest approach involves applying thresholding to the IG results. Otsu thresholding [20] is a method used to automatically find the optimal threshold for separating an image into two classes: foreground and background. It works by analyzing the histogram of pixel intensities and selecting the threshold that minimizes the intra-class variance, which measures the spread of pixel values within each class. This approach ensures that the selected threshold effectively distinguishes between the two groups of pixels.

In our application, we use the relevance map generated from the Integrated Gradients method as input for Otsu thresholding. The relevance map highlights the importance of each pixel in the input image for the prediction of the model. By applying Otsu thresholding to this map, we can filter out less relevant pixels and focus on the most significant areas that contributed to the model's decision. This helps reduce noise and emphasize the most pertinent regions in the relevance map.

Morphological Processing We applied a series of image processing and morphological operations to refine the pixel relevance maps generated by the Integrated Gradients method. First, the relevance maps were converted to grayscale and enhanced through histogram equalization [21]. Thresholding was then used to create binary masks that highlight significant regions. We labelled connected components, removed border artefacts, and used dilation to connect fragmented areas. Erosion was applied to clean the boundaries, and small holes were filled to produce smooth, continuous masks.

Normalized Cut The Normalized Cut method [22] models an image as a weighted graph $G = (V, E)$, where each pixel corresponds to a node $v \in V$ and edges $e \in E$ connect pairs of nodes. The weight w_{ij} of an edge between nodes i and j represents the similarity between the corresponding pixels.

The goal is to partition the graph into disjoint sets A and B by minimizing the Ncut value and therefore separating foreground and background, which can be defined as:

$$\text{NCut}(A, B) = \frac{\text{Cut}(A, B)}{\text{Assoc}(A, V)} + \frac{\text{Cut}(A, B)}{\text{Assoc}(B, V)} \quad (2)$$

where:

$$\text{Cut}(A, B) = \sum_{i \in A, j \in B} w_{ij} \quad (3)$$

$$\text{Assoc}(A, V) = \sum_{i \in A} \sum_{j \in V} w_{ij}, \text{Assoc}(B, V) = \sum_{i \in B} \sum_{j \in V} w_{ij} \quad (4)$$

Fusion of Feature and Relevance Map Extracting the feature map from the Vision Transformer network enhances the XAI segmentation by revealing how the network perceives medical images. By integrating pixel-level importance with the feature space representation, we can reduce the influence of outlier pixels in the integrated gradient map. We combine the relevance map and feature map by multiplying them, effectively using the feature map as a weighted map to denoise the relevance map.

3.6 DenseCRF Edge Refinement

Dense Conditional Random Field (DenseCRF) [23] is a powerful post-processing technique used to refine coarse segmentation masks, particularly by leveraging the information from the original image. They work by incorporating both the initial segmentation predictions and the pixel-level details of the original image to produce more accurate and coherent results.

DenseCRF models combine unary potentials, which come from the initial segmentation output and indicate the likelihood of each pixel belonging to a specific class, with pairwise potentials that consider the relationships between pixels based on their spatial location and appearance (colour, texture). The key advantage of DenseCRF is its ability to enforce pixels with similar appearances and close proximity in the original image are more likely to share the same label in the refined segmentation.

The energy function in DenseCRF is defined as:

$$E(\mathbf{x}) = \sum_i \phi_u(\mathbf{x}_i) + \sum_{i < j} \phi_p(\mathbf{x}_i, \mathbf{x}_j) \quad (5)$$

where $\phi_u(\mathbf{x}_i)$ is the unary potential that captures the likelihood of pixel $\mathbf{x}_i$ belonging to a particular class, and $\phi_p(\mathbf{x}_i, \mathbf{x}_j)$ is the pairwise potential that models the interactions between all neighbouring pixel pairs $\mathbf{x}_i$ and $\mathbf{x}_j$ without repetition.

Algorithm 1 summarizes our method.

4 Experiments

4.1 Setup

Datasets

Algorithm 1: General Pipeline for ExplainSeg

input : Pre-trained classification network F
 & trainable classification head G
input : Training dataset, D_{train} which contains:
 Training image for optional training, $\mathbf{x}_{\text{train}}$
 Ground truth class for optional training, y_{train}
input : Input image, $\mathbf{x}$
output : Predicted class, y
output : Predicted segmentation mask $\mathbf{y}$ for $\mathbf{x}$

- 1 Optionally finetune F and train G with $\mathbf{x}_{\text{train}}$ and y_{train}
- 2 Predict the class y of input image $\mathbf{x}$ by using F
- 3 Obtain relevance map by applying XAI to result y
- 4 Process the relevance map to obtain a more detailed mask
- 5 Refine mask with DenseCRF to acquire final segmentation mask $\mathbf{y}$

In this research, we conducted experiments on three publicly available medical imaging datasets, each representing a different domain: breast cancer screening, histopathology, and gastrointestinal endoscopy. These datasets include the CBIS-DDSM (Curated Breast Imaging Subset of the Digital Database for Screening Mammography) [24], NuInsSeg (Nuclei Instance Segmentation dataset) [25], and Kvasir-SEG (Polyp Segmentation dataset) [26].

The CBIS-DDSM dataset is an updated and standardized subset of the DDSM, designed for computer-aided detection and diagnosis in mammography. It contains over 3,000 annotated mammography images across two primary classes: benign and malignant, along with normal cases in some settings. Each case is provided with a corresponding pathology-confirmed label. For our classification training, we used 1,276 images for training and 234 for validation.

The NuInsSeg dataset is a fully annotated benchmark for instance-level nuclei segmentation in H&E-stained histopathology slides. It consists of 665 high-resolution image tiles with annotations for over 30,000 individual nuclei instances across a variety of tissue types. For our classification-based training, we used 600 images for training and 65 for validation. Although the dataset supports multi-class segmentation, we focused on binary instance segmentation for each nuclei class. All images were normalized using the mean and standard deviation computed from the training set to ensure consistency across staining variations.

The Kvasir-SEG dataset is a polyp segmentation dataset sourced from real-world colonoscopy examinations. It consists of 1,000 colonoscopy images annotated with corresponding binary masks indicating polyp regions. As the dataset contains only positive samples, we split each image into smaller patches and labeled them as positive or negative based on the presence of polyp pixels in the corresponding mask. This approach enabled us to generate a balanced dataset for classification training. We created a training set of 1,852 image patches and a validation set of 254 patches. We computed the mean and standard

deviation values from the training images and used these statistics to standardize inputs for both training and inference.

For all datasets, we applied dataset-specific normalization using their respective channel-wise mean and standard deviation values, computed solely from the training subsets to avoid data leakage. Additionally, data augmentation strategies such as horizontal flipping, color jittering, and rotation were applied to improve model robustness and generalization.

Due to the unique structure of the ExplainSeg framework, no information from the ground truth segmentation masks is used during the classification network finetuning stage. Consequently, for the final segmentation mask generation step, both the training and validation images are passed through the trained model to produce predictions, which are then evaluated against the ground truth masks to assess segmentation performance.

Method Variations

The segmentation pipeline proposed in this study consists of two main stages: the **explanation stage** and the **post-processing stage**. For each stage, we introduce two distinct strategies, resulting in four unique combinations: Fusion & Morphology (Variant A), Fusion & NCut (Variant B), XAI & Morphology (Variant C), and XAI & NCut (XNCut).

In the **explanation stage**, we explore two alternative approaches to generate meaningful interpretability maps:

XAI-only: Employs explainable AI techniques (Integrated Gradients with Noise Tunnel) on the classification network to directly produce heatmaps that indicate regions contributing most to the classification outcome.

Fusion-based: Combines the XAI-generated heatmap with intermediate feature maps from the Vision Transformer (ViT) backbone, aiming to enhance spatial precision while retaining semantic relevance.

In the **post-processing stage**, the relevance maps are transformed into binary segmentation masks using one of the following two techniques:

NCut (Normalized Cut): Treats the relevance map as a weighted graph and applies spectral clustering to segment semantically coherent regions.

Morphological Operations: Applies classical image processing operations, such as thresholding, dilation, and erosion, to extract connected, object-like structures.

Each combination is designed to investigate how different forms of interpretability and refinement influence segmentation performance. We evaluate the effectiveness of these methods across three different datasets, as shown in Table 1.

Implementation details and hyperparameters

Our implementation is based on Python using the PyTorch framework [27]. We benefit from using the Distillation with No Label (DINO) [12] Vision Transformer model since it operates in a self-supervised manner making it ideal when labelling is expensive or impractical.

METHOD VARIANT	COMPONENTS			NUINSSEG		CBIS-DDSM		Kvasir-SEG	
	XAI	Fusion	Morph NCut	mIoU%	Dice%	mIoU%	Dice%	mIoU%	Dice%
Variant A		✓	✓	13.4	22.4	19.0	32.3	31.1	45.6
Variant B		✓		12.5	20.1	11.5	25.6	32.5	45.9
Variant C	✓		✓	9.2	16.0	25.7	34.3	30.2	44.6
XNCut	✓		✓	13.1	20.6	31.2	43.7	28.6	41.4

Table 1: Ablation study of ExplainSeg variants across three datasets (mIoU% / Dice%). Bold indicates best performance per dataset and metric.

The 'deitsmall16' pre-trained weights were used to fine-tune the Vision Transformer. Stochastic Gradient Descent (SGD) [28] as the optimizer with a learning rate of 0.005, combined with momentum for stability. Additionally, learning rate decay was applied, reducing the learning rate by a factor of 0.1 every 50 epochs to ensure smoother convergence during fine-tuning. To apply XAI to the classification results, both Integrated Gradients and noise tunnelling (with 5 samples) [18] were implemented to mitigate the impact of noise in both the IG results and the original image.

4.2 Experimental results

Baselines Selection

The core objective of our project is to generate segmentation masks without relying on ground truth annotations. This design choice significantly limits the pool of directly comparable baseline models, as most established segmentation methods require extensive pixel-wise labeled datasets—an approach we intentionally avoid.

While some prior work has explored explainable AI (XAI) methods for segmentation, these are primarily developed for natural image domains and often lack publicly available implementations. As such, they are neither reproducible nor directly applicable to the medical imaging setting of our study, and thus were excluded from experimental comparison.

To assess the performance of our method, we selected two representative baselines: **TokenCut**[4] and **MICRA-Net**[29]. We also incorporated **MaskCut**, an unsupervised segmentation model hosted within the CutLER repository, which forms the core inference engine in our own pipeline.

TokenCut is an unsupervised segmentation technique that leverages Vision Transformer (ViT) feature tokens and applies spectral clustering to segment images without requiring annotations. Though not specifically designed for medical data, its reliance on high-level semantic features aligns well with our aim of generating segmentation masks driven by meaningful visual cues. TokenCut serves as a strong benchmark for evaluating performance under minimal supervision.

MICRA-Net, by contrast, is a model explicitly designed for medical image analysis. It was pretrained on a domain-specific dataset consisting of F-actin fluorescence microscopy

images of hippocampal neurons, as described in the publication: “*Neuronal activity remodels the F-actin based submembrane lattice in dendrites but not axons of hippocampal neurons.*”[30] Although MICRA-Net employs supervised learning, we include it in our comparison to evaluate how a specialized, pretrained model performs in annotation-scarce segmentation tasks.

MaskCut, incorporated as part of the CutLER[5] codebase, is a recent unsupervised segmentation method that combines self-supervised ViT features with bipartition-based object discovery. Unlike CutLER’s full pipeline—which includes a clustering and mask refinement stage—MaskCut performs direct bipartition sampling to generate multiple object proposals per image. Despite being developed for general-purpose segmentation, MaskCut proves highly relevant to our study, as it provides multiple candidate masks without requiring pixel-level labels. Its use of CRF-based post-processing and ViT-based semantic representations makes it especially suitable for extracting coarse object boundaries in medical images.

By benchmarking our XAI-driven segmentation approach against TokenCut, MICRA-Net, and MaskCut, we highlight its ability to deliver informative segmentation outputs under limited supervision. This comparative evaluation underscores the novelty and practical relevance of our method in domains—like medical imaging—where annotated data is scarce or expensive to obtain.

Ablation Study

Table 1 compares the four ExplainSeg variants—Fusion & Morphology(Variant A), Fusion & NCut(Variant B), XAI & Morphology(Variant C), and XAI & NCut(XNCut)—across NuInsSeg, CBIS-DDSM, and Kvasir-SEG using mIoU and Dice metrics.

XNCut achieves the highest mIoU (31.2%) and Dice (43.7%) on CBIS-DDSM, the dataset with the highest clinical relevance. It also performs competitively on NuInsSeg and Kvasir-SEG, demonstrating a strong balance between explainability and spatial coherence. This indicates that the integration of explainability-driven pixel relevance with NCut segmentation effectively captures meaningful regions in complex medical images.

Fusion & Morphology shows solid performance on NuInsSeg and Kvasir-SEG but falls short on CBIS-DDSM, reflecting its limitations in handling highly domain-specific imagery. Fusion & NCut excels on Kvasir-SEG, achieving the highest Dice (45.9%) and mIoU (32.5%), suggesting that NCut’s spatial grouping benefits datasets with clearer anatomical structures and consistent textures. Meanwhile, XAI & Morphology performs reasonably but does not match the consistency or peak performance of XAI & NCut(XNCut).

Baseline Comparison

Table 2 presents a comparison between ExplainSeg (XAI & NCut) and other segmentation methods—TokenCut, MICRA-Net, and MaskCut—across the three datasets.

ExplainSeg (XAI & NCut) consistently achieves the best or near-best segmentation performance across all three datasets. It yields the highest mIoU (31.2%) and Dice score (43.7%) on CBIS-DDSM, the mammography dataset that poses substantial challenges due to low contrast and subtle boundaries. On NuInsSeg, which contains microscopy images

METHOD	NUINSEg	CBIS-DDSM	KVASIR-SEG
mIoU			
TokenCut [4]	5.8	1.9	23.8
MICRA-Net [29]	7.8	12.7	15.2
MaskCut [5]	9.3	5.4	41.3
ExplainSeg (XNCut)	13.1	31.2	28.6
MEAN DICE			
TokenCut [4]	10.2	3.1	28.4
MICRA-Net [29]	14.1	20.8	25.3
MaskCut [5]	14.3	8.5	48.4
ExplainSeg (XNCut)	20.6	43.7	41.4

Table 2: Comparison of segmentation methods on three datasets using mIoU and mean Dice metrics (%). Bold indicates best performance per dataset and method.

with densely packed nuclei, ExplainSeg also surpasses TokenCut and MICRA-Net by a clear margin. While it does not attain the top score on Kvasir-SEG, its performance remains highly competitive, underscoring its strong generalizability across diverse imaging modalities.

MaskCut, by contrast, achieves the highest scores on Kvasir-SEG (mIoU 41.3%, Dice 48.4%), outperforming all other methods. This is likely attributable to the nature of the dataset, which comprises RGB endoscopic images resembling natural photographs. As MaskCut is pretrained on general-domain image data, its superior performance on Kvasir-SEG is consistent with its inductive bias toward natural image statistics. However, it performs considerably worse on CBIS-DDSM and NuInsSeg, suggesting limited applicability to more specialized or clinical imaging domains.

MICRA-Net performs moderately across all datasets but is consistently outperformed by ExplainSeg, indicating that while it offers a degree of medical-domain awareness, it lacks the fine-grained localization and spatial reasoning facilitated by the integration of XAI and NCut.

TokenCut exhibits the weakest performance overall, particularly on CBIS-DDSM and NuInsSeg, further reinforcing the limitations of general image segmentation techniques when applied to structurally and semantically distinct medical imagery.

These results affirm the effectiveness of ExplainSeg’s XAI-driven relevance maps combined with spectral clustering for producing spatially coherent and clinically meaningful segmentations. Its ability to perform well across varying modalities highlights its robustness and suitability for real-world medical imaging applications.

5 Conclusions

This study presents an integrated approach that combines state-of-the-art deep learning techniques with Explainable AI (XAI) to address key challenges in medical image segmen-




Fig. 2: Comparison of results.

Top to Bottom: CBIS-DDSM, NuInsSeg, and Kvasir-SEG.

Left to Right: Original, Ground Truth, ExplainSeg, TokenCut, MICRA-Net, MaskCut

tation. Leveraging transfer learning, we fine-tuned a DINO model [12] pre-trained with deitsmall16 weights on domain-specific medical imaging data, effectively mitigating the limitations posed by scarce annotated samples. The incorporation of XAI methods enhanced the interpretability of the model, providing critical insights into its decision-making process and fostering greater trust in its predictions.

Our experimental results demonstrate that it is possible to achieve high segmentation accuracy while maintaining clinical interpretability. This balance between performance and transparency is essential for the deployment of AI systems in healthcare, where understanding the rationale behind model outputs is as important as the outputs themselves.

Future research will focus on refining these techniques by investigating more advanced and task-specific XAI methods, expanding the scope to additional medical imaging modalities, and conducting clinical validation. Ultimately, our aim is to develop robust, interpretable AI tools that can assist clinicians in making informed decisions and contribute to improved patient care outcomes.

References

1. Md Eshmam Rayed, SM Sajibul Islam, Sadia Islam Niha, Jamin Rahman Jim, Md Mohsin Kabir, and MF Mridha, "Deep learning for medical image segmentation: State-of-the-art advancements and challenges," *Informatics in Medicine Unlocked*, p. 101504, 2024.
2. Zhi-Hua Zhou, "A brief introduction to weakly supervised learning," *National science review*, vol. 5, no. 1, pp. 44–53, 2018.

3. Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He, "A comprehensive survey on transfer learning," *Proceedings of the IEEE*, vol. 109, no. 1, pp. 43–76, 2021.
4. Yangtao Wang, Xi Shen, Shell Xu Hu, Yuan Yuan, James L Crowley, and Dominique Vaufreydaz, "Self-supervised transformers for unsupervised object discovery using normalized cut," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 14543–14553.
5. Xudong Wang, Rohit Girdhar, Stella X Yu, and Ishan Misra, "Cut and learn for unsupervised object detection and instance segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 3124–3134.
6. Oriane Siméoni, Chloé Sekkat, Gilles Puy, Antonín Vobecký, Éloi Zablocki, and Patrick Pérez, "Unsupervised object localization: Observing the background to discover objects," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
7. Sheikh Rabiul Islam, William Eberle, Sheikh Khaled Ghafoor, and Mohiuddin Ahmed, "Explainable artificial intelligence approaches: A survey," 2021.
8. Dinesh D Patil and Sonal G Deore, "Medical image segmentation: a review," *International Journal of Computer Science and Mobile Computing*, vol. 2, no. 1, pp. 22–27, 2013.
9. KKD Ramesh, G Kiran Kumar, K Swapna, Debabrata Datta, and S Suman Rajest, "A review of medical image segmentation algorithms," *EAI Endorsed Transactions on Pervasive Health and Technology*, vol. 7, no. 27, pp. e6–e6, 2021.
10. Milan Sonka, J Michael Fitzpatrick, et al., "Handbook of medical imaging: Volume 2, medical image processing and analysis," SPIE, 2000.
11. Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," *ICLR*, 2021.
12. Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin, "Emerging properties in self-supervised vision transformers," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 9650–9660.
13. Mukund Sundararajan, Ankur Taly, and Qiqi Yan, "Axiomatic attribution for deep networks," 2017.
14. Ramprasaath R Selvaraju, Abhishek Das, Ramakrishna Vedantam, Michael Cogswell, Devi Parikh, and Dhruv Batra, "Grad-cam: Why did you say that?," 2017.
15. Clemens Seibold, Johannes Künzel, Anna Hilsmann, and Peter Eisert, "From explanations to segmentation: Using explainable ai for image segmentation," in *Proceedings of the 17th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*. 2022, p. 616–626, SCITEPRESS - Science and Technology Publications.
16. Florent Forest, Hugo Porta, Devis Tuia, and Olga Fink, "From classification to segmentation with explainable ai: A study on crack detection and growth monitoring," 2024.
17. Yu-Xiong Wang, Deva Ramanan, and Martial Hebert, "Growing a brain: Fine-tuning by increasing model capacity," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 2471–2480.
18. Narine Kokhlikyan, Vivek Miglani, Miguel Martin, Edward Wang, Bilal Alsallakh, Jonathan Reynolds, Alexander Melnikov, Natalia Kliushkina, Carlos Araya, Siqi Yan, and Orion Reblitz-Richardson, "Captum: A unified and generic model interpretability library for pytorch," 2020.
19. Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel, "Deep vit features as dense visual descriptors," *arXiv preprint arXiv:2112.05814*, vol. 2, no. 3, pp. 4, 2021.
20. Mehmet Sezgin and Bülent Sankur, "Survey over image thresholding techniques and quantitative performance evaluation," *J. Electronic Imaging*, vol. 13, pp. 146–168, 2004.
21. Stephen M Pizer, E Philip Amburn, John D Austin, Robert Cromartie, Ari Geselowitz, Trey Greer, Bart ter Haar Romeny, John B Zimmerman, and Karel Zuiderveld, "Adaptive histogram equalization and its variations," *Computer vision, graphics, and image processing*, vol. 39, no. 3, pp. 355–368, 1987.
22. Jianbo Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 888–905, 2000.

23. Philipp Krähenbühl and Vladlen Koltun, "Efficient inference in fully connected crfs with gaussian edge potentials," *Advances in neural information processing systems*, vol. 24, 2011.
24. Rebecca Sawyer Lee, Francisco Gimenez, Assaf Hoogi, Kanae Kawai Miyake, Mia Gorovoy, and Daniel L Rubin, "A curated mammography data set for use in computer-aided detection and diagnosis research," *Scientific Data*, vol. 4, pp. 170177, 2017.
25. Amirreza Mahbod, Christine Polak, Katharina Feldmann, Rumsha Khan, Katharina Gelles, Georg Dorffner, Ramona Woitek, Sepideh Hatamikia, and Isabella Ellinger, "Nuinsseg: A fully annotated dataset for nuclei instance segmentation in h&e-stained histological images," 2023.
26. Debesh Jha, Pia H. Smedsrud, Michael A. Riegler, Pål Halvorsen, Thomas de Lange, Dag Johansen, and Håvard D. Johansen, "Kvasir-seg: A segmented polyp dataset," 2019.
27. Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer, "Automatic differentiation in pytorch," 2017.
28. Herbert Robbins and Sutton Monro, "A stochastic approximation method," *The annals of mathematical statistics*, 1951.
29. Anthony Bilodeau, Constantin VL Delmas, Martin Parent, Paul De Koninck, Audrey Durand, and Flavie Lavoie-Cardinal, "Microscopy analysis neural network to solve detection, enumeration and segmentation from image-level annotations," *Nature Machine Intelligence*, vol. 4, no. 5, pp. 455–466, 2022.
30. Flavie Lavoie-Cardinal, Anthony Bilodeau, Mado Lemieux, Marc-André Gardner, Theresa Wiesner, Gabrielle Laramée, Christian Gagné, and Paul De Koninck, "Neuronal activity remodels the f-actin based submembrane lattice in dendrites but not axons of hippocampal neurons," *Scientific reports*, vol. 10, no. 1, pp. 11960, 2020.