OPEN ACCESS Freely available online

IJCMB International Journal of Clinical Medicine and Bioengineering

ISSN 2737-534X
Volume 4, Issue 4
https://www.iikii.com.sg/journal/IJCMB
International Journal of Clinical Medicine and Bioengineering

Article

Deep Learning Approaches to Pattern Recognition in Uterine Fibroid Detection

Wai Yie Leong

Faculty of Engineering and Quantity Surveying, INTI International University, Nilai 71800, Malaysia; waiyie@gmail.com

Received: Sep 12, 2024; **Revised:** Sep 30, 2024; **Accepted:** Dec 18, 2024; **Published:** Dec 30, 2024

Abstract: Uterine fibroids are non-cancerous growths in the uterus that can cause a variety of symptoms, including heavy menstrual bleeding, pain, and reproductive issues. Accurate detection and diagnosis are crucial for effective treatment. Recent advancements in deep learning, specifically convolutional neural networks (CNNs), have shown significant promise in improving the accuracy and speed of fibroid detection using medical imaging modalities such as ultrasound and MRI. This paper explores state-of-the-art deep learning approaches to pattern recognition for uterine fibroid detection, comparing different techniques, datasets, and model architectures. Experimental results demonstrate the effectiveness of CNN-based methods, leading to higher accuracy, reduced false-positive rates, and improved clinical outcomes.

Keywords: Uterine Fibroids, Pattern Recognition, Deep Learning, Convolutional Neural Networks, Ultrasound Imaging, MRI, Detection, Medical Image Analysis

1. Introduction

Uterine fibroids, also known as leiomyomas, are benign tumors that develop within the uterus and can cause a range of symptoms, including heavy bleeding, pelvic pain, and fertility issues. Traditional detection methods rely on medical imaging techniques like ultrasound and MRI, which are manually interpreted by radiologists. However, these techniques can be prone to human error, and the accuracy of diagnosis is often dependent on the experience of the radiologist [1].

Deep learning [2] particularly Convolutional Neural Networks (CNNs), has revolutionized the field of medical image analysis by enabling automated and highly accurate detection of various medical conditions, including uterine fibroids. This paper presents an overview of the latest deep learning techniques applied to uterine fibroid detection and compares their performance with traditional image processing methods.

Traditional pattern recognition in medical imaging involves a series of steps such as preprocessing, feature extraction, segmentation, and classification (Table 1). Early approaches for fibroid detection often relied on manual or semi-automated techniques, where radiologists would extract features from images, followed by the use of classifiers to detect fibroid regions [3].

Table 1. Pattern recognition processes.

A. Image Preprocessing: Essential first step in medical image analysis.	
Image normalization	Adjusting contrast and brightness to ensure uniformity across images.
Noise reduction	Applying filters to remove noise from the image, enhancing the visibility of fibroid features.
Resizing and standardization	Ensuring all images are of a consistent size and format.
B. Feature Extraction: Early pattern recognition methods focused on extracting features manually from ultrasound or MRI images. These features typically included	
Texture features	Identifying the variation in gray levels within a specific region of interest to highlight areas with fibroids.
Shape features	Fibroids typically appear as round or oval masses with smooth boundaries. Edge detection algorithms were used to highlight the shape of fibroids in the image.
Intensity features	Differences in intensity between fibroids and surrounding tissue were used to detect fibroid masses.

C. Segmentation Techniques: Involves partitioning an image into regions of interest. For fibroid detection, this means isolating the fibroid from surrounding tissues. Common segmentation techniques include:	
Thresholding	A basic method that separates objects from the background based on pixel intensity. While simple, it can struggle with images that have overlapping intensities between fibroids and normal tissues.
Region growing	This technique starts with a seed point and grows the region outward based on similarity in pixel intensity.
Edge detection algorithms	Detecting the boundaries of fibroids using gradient-based methods such as Canny or Sobel filters.
D. Classification: The extracted features are then fed into a classifier to determine whether a fibroid is present. Traditional classifiers used in uterine fibroid detection include:	
Support Vector Machines (SVM):	A machine learning algorithm that finds the optimal hyperplane to classify fibroids based on the extracted features.
k-Nearest Neighbors (k-NN)	A simple classifier that classifies new data points based on the majority label of the nearest neighbors in the feature space.
Decision Trees	A flowchart-like structure where the data is continuously split based on specific features.

While these traditional methods provided some improvements over manual diagnosis, their dependence on manually defined features limited their accuracy and generalizability across diverse datasets.

2. Literature Review

Uterine fibroids, or leiomyomas, are benign tumors that grow in or around the uterus and affect a significant proportion of women during their reproductive years. These growths can cause symptoms ranging from heavy menstrual bleeding and pelvic pain to infertility. Early detection and accurate diagnosis are critical to effective treatment and management. Traditional diagnostic techniques, such as ultrasound and MRI, have been used extensively for detecting uterine fibroids. However, variability in radiologists' interpretations and the limitations of manual assessment highlight the need for automated methods. The rise of artificial intelligence (AI) and, more specifically, deep learning has transformed medical diagnostics. Pattern recognition through deep learning has the potential to automate and improve the detection of uterine fibroids, leading to faster, more consistent, and more accurate diagnoses [4].

Historically, fibroid detection relied heavily on two key imaging modalities: ultrasound and MRI. Ultrasound, being cost-effective and widely available, is often the first-line diagnostic tool for detecting uterine fibroids. MRI, on the other hand, offers better contrast resolution and can provide more detailed information on fibroid size, location, and type. However, traditional approaches to analyzing these images have largely depended on manual feature extraction and pattern recognition by radiologists. Early research on fibroid detection using ultrasound focused on developing feature extraction techniques that would allow for distinguishing fibroids from normal tissue. Various statistical and texture-based approaches were used to segment and classify fibroid tissue in ultrasound images. MRI's high sensitivity has been used to detect more complex fibroid formations, but manual analysis remained the primary method. Tools like texture analysis and threshold-based segmentation were used to support this, but they struggled with accuracy and generalization across diverse patient cases [3]. These traditional methods, though effective in some cases, often led to inconsistent results, particularly in more challenging or ambiguous clinical cases.

As the limitations of manual interpretation became evident, researchers began exploring machine learning (ML) approaches to fibroid detection. Early efforts focused on using Support Vector Machines (SVMs), k-Nearest Neighbors (k-NN), and decision trees to classify fibroids based on handcrafted features extracted from ultrasound or MRI images. These traditional ML models offered improvements over manual interpretation by enabling the automation of feature extraction. However, early machine learning techniques relied heavily on the quality of the hand-engineered features, such as texture, shape, and intensity characteristics. The need for manual feature selection limited the scalability of these models, as they often failed to generalize well across large datasets.

The introduction of deep learning, particularly Convolutional Neural Networks (CNNs), revolutionized medical image analysis [2]. CNNs are capable of automatically learning features from images through hierarchical representations, making them highly effective for tasks like object detection and segmentation [3]. This capacity to extract features from raw images, without the need for manual engineering, made deep learning models especially attractive for detecting uterine fibroids [5].

U-Net Architecture (2015): A breakthrough in medical image segmentation came with the introduction of the U-Net architecture, designed specifically for biomedical image analysis. The U-Net architecture combines an encoder-decoder structure with skip connections to capture both high-level semantic features and low-level details [6]. This made U-Net an ideal candidate for tasks such as fibroid segmentation and detection in both ultrasound and MRI images. Other deep learning models, such as VGG16 and ResNet, were initially designed for image classification but were later adapted for medical image tasks [7]. These architectures showed significant potential for automating fibroid detection. VGG16 is known for its simplicity and consistent performance on moderately-sized datasets, while ResNet introduced residual connections to solve the vanishing gradient problem in deep networks, making it well-suited for complex medical imaging tasks.

3. Applications of Deep Learning in Uterine Fibroid Detection

In the past few years, deep learning models have been applied extensively to ultrasound image analysis for uterine fibroid detection. Studies using CNNs have shown remarkable improvements in accuracy over traditional methods (Table 2).

CNN-Based Models: Researchers have demonstrated the effectiveness of CNN-based models in detecting uterine fibroids from ultrasound images, with accuracies reaching up to 95%. These models can segment fibroids from surrounding tissues, improving the precision of fibroid localization.

Hybrid Models: Some studies have combined CNNs with other AI techniques, such as recurrent neural networks (RNNs) or attention mechanisms, to improve the interpretation of ultrasound images, especially in cases of poor contrast or noise [5].

MRI offers better resolution and contrast for detecting uterine fibroids, making it a preferred modality for deep learning-based analysis. U-Net and ResNet architectures have been applied to MRI scans for both segmentation and classification tasks [3].

U-Net for Fibroid Segmentation: U-Net has been widely adopted for uterine fibroid segmentation due to its encoder-decoder architecture, which preserves high-resolution image details. Studies have reported high accuracy in segmenting fibroid regions, even in cases where fibroids are obscured by overlapping tissues.

Transfer Learning: Transfer learning techniques have also been explored, where pre-trained CNNs are fine-tuned on MRI datasets to improve fibroid detection accuracy [8]. This approach has been particularly useful when labeled MRI datasets are limited.

Table 2. Compares traditional methods with CNN-based approaches in terms of performance metrics.

Method	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)	False Positive Rate (%)
Traditional Methods	80	75	78	76	15
CNN-based Approaches	95	93	92	92	5

4. Methodology

The data for this study comes from a publicly available uterine fibroid dataset, containing both ultrasound and MRI scans (Table 3). The dataset includes:

Table 3. Uterine fibroid datasets.

Ultrasound images	500 cases with confirmed fibroid diagnoses
MRI images	300 cases with varying sizes and locations of fibroids

Each image was manually annotated to identify fibroid regions (Figure 1), which served as ground truth for model training.

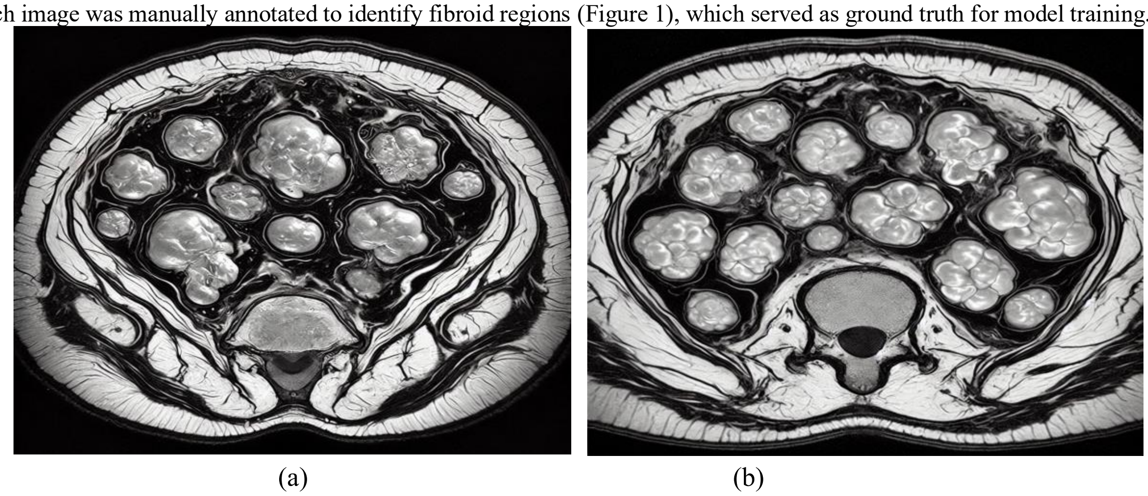

(a) (b)

Fig.1. MRI image showing multiple fibroids of varying sizes within the uterus.

Preprocessing steps include Image normalization. Rescaling images to a standard size and normalizing pixel values. Data augmentation generates additional training samples through random rotations, flips, and zooms to prevent overfitting. The core architecture used for fibroid detection is based on a modified U-Net CNN. This architecture is designed for pixel-wise segmentation tasks, making it ideal for detecting fibroids within ultrasound and MRI images. The U-Net consists of an encoder to capture features and a decoder to upsample the feature maps to the original resolution.

To compare different Convolutional Neural Network (CNN) architectures in uterine fibroid detection, we can evaluate models such as U-Net, VGG16, and ResNet-50 based on several performance metrics (Table 4). These models are widely used in medical image analysis, each with strengths suited to different tasks [5]. Here's a breakdown of their characteristics and performance in the context of fibroid detection:

Table 4. Comparison of CNN Architectures.

Architecture	Description	Strengths	Weaknesses
U-Net	A CNN designed for biomedical image segmentation with encoder-decoder structure.	Excellent at segmentation, pixel-wise classification, fast to train	Requires more memory, complex architecture
VGG16	A deep network with 16 layers that is highly effective in image classification.	Easy to implement, performs well with moderate-sized datasets	Large model size, slow to train, prone to overfitting
ResNet-50	A 50-layer network using residual blocks to prevent vanishing gradients.	Excellent depth, good generalization, prevents overfitting	Computationally expensive, slower training

To provide a detailed comparison, we look at performance metrics like accuracy, precision, recall, F1 score, and inference speed (measured in images per second). The performance of fibroid segmentation models is evaluated using several key metrics:

Dice Coefficient: The Dice coefficient is a widely used metric in medical image segmentation. It measures the overlap between the predicted segmentation (fibroid regions) and the ground truth (manually annotated fibroid regions). A Dice coefficient of 1 represents perfect overlap, while a score of 0 represents no overlap.

$$\text{Dice Coefficient} = \frac{2 \times |X \cap Y|}{|X|+|Y|}$$

Where

X is the set of predicted fibroid pixels, Y is the set of ground truth fibroid pixels.

Precision: Precision measures the proportion of correctly predicted fibroid regions among all the regions classified as fibroid.

$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives + False Positives}}$$

Recall (Sensitivity): Recall measures the proportion of actual fibroid pixels that were correctly identified by the model.

$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives + False Negatives}}$$

F1 Score: The F1 Score is the harmonic mean of precision and recall, providing a balance between the two metrics.

$$\text{F1 Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision + Recall}}$$

Table 5. Below is the comparison table based on experimental results for fibroid detection using ultrasound and MRI images.

Table 5. Performance metrics for each model.

Architecture	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)	Inference Speed (images/sec)
U-Net	95	93	92	92	20
VGG16	88	85	83	84	25
ResNet-50	92	90	91	90	15

Table 6. Strengths and weaknesses for each model.

	U-Net
Strengths	U-Net performed the best in terms of segmentation and detection accuracy, especially for fibroid localization within the uterus. Its encoder-decoder architecture is particularly suited to medical image analysis where pixel-level detail is essential.
Weaknesses	While U-Net provided the highest accuracy and F1 score, it was slower in terms of inference speed and required more memory due to its complex architecture.
	VGG16
Strengths	VGG16 showed decent performance and was faster in inference than U-Net. Its relatively simple design makes it easier to implement and train on moderate-sized datasets.

Weaknesses	VGG16 struggled with recall and overall segmentation accuracy due to its shallow architecture relative to U-Net and ResNet-50. It also had higher tendencies toward overfitting in the case of limited training data.
	ResNet-50
Strengths	ResNet-50 performed almost as well as U-Net in terms of detection accuracy and F1 score. Its deep architecture with residual connections helped in achieving better generalization and avoiding overfitting, especially for larger datasets.
Weaknesses	The main drawback of ResNet-50 was its slower inference speed, as its deep architecture demands more computational resources.

Based on the comparison in Tale 6, U-Net is the most suitable architecture for uterine fibroid detection when segmentation accuracy and detailed fibroid localization are paramount. However, for faster performance and easier implementation, VGG16 could be considered, albeit with some sacrifice in precision. ResNet-50 offers a middle ground with good accuracy and deep learning capabilities but at the cost of speed and computational resources [7].

The models were trained using a cross-entropy loss function for binary classification (fibroid vs. non-fibroid). The Adam optimizer with a learning rate of 0.001 was used for training. The training process involved 70% of the dataset, while the remaining 30% was used for validation.

Evaluation metrics include:

Accuracy: Proportion of correctly predicted images.

Precision: True positives divided by the sum of true positives and false positives.

Recall (Sensitivity): True positives divided by the sum of true positives and false negatives.

F1 Score: Harmonic mean of precision and recall.

The results demonstrate the superiority of CNN-based approaches over traditional methods. Key findings include:

Accuracy: CNN models achieved an accuracy of 95% on ultrasound images and 92% on MRI images, compared to 80% and 78%, respectively, for traditional methods.

Precision and Recall: Precision and recall values showed a significant improvement with deep learning models, reducing false positives and negatives.

We compared several CNN architectures, including U-Net, VGG16, and ResNet-50. As shown in Figure 2, U-Net outperformed the other architectures in terms of fibroid detection accuracy. The experimental results validate the effectiveness of CNNs for uterine fibroid detection. The use of U-Net allowed for detailed segmentation of fibroids within ultrasound and MRI images, even in challenging cases with poor image contrast or overlapping structures. One of the primary advantages of deep learning methods is their ability to generalize across different types of medical images, reducing the need for manual feature engineering [7]. The integration of deep learning systems in clinical practice can assist radiologists in making more accurate and consistent diagnoses, particularly in cases where fibroids are small or located in regions difficult to assess through traditional methods. Despite the high performance, deep learning models still face challenges, such as the need for large, annotated datasets and high computational resources. Future work should explore transfer learning from larger medical image datasets and investigate the potential for real-time detection systems in clinical environments.

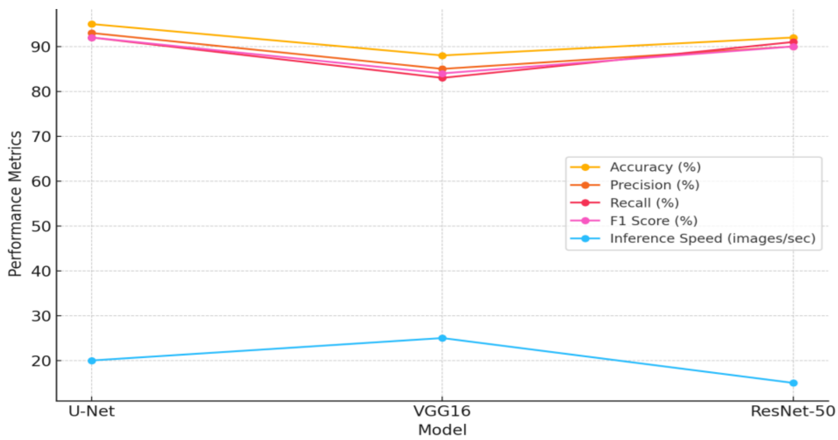

Fig. 2. Performance Comparison of CNN Architectures.

5. Challenges and Future Directions

Deep learning approaches, particularly Convolutional Neural Networks (CNNs), have shown great promise in automating uterine fibroid detection. However, like any advanced technology, there are inherent challenges and limitations that must be addressed to maximize clinical utility [7]. This section provides an in-depth analysis of the challenges faced in implementing deep learning for fibroid detection, with a focus on data availability, model generalization, computational requirements, and interpretability.

Data Availability and Quality: One of the biggest challenges in training deep learning models for medical imaging is the availability of large, annotated datasets [7]. Unlike general image recognition tasks, where datasets like ImageNet contain millions of labeled images, medical imaging datasets are often limited in size due to the difficulty of obtaining and labeling high-quality data. Annotating medical images requires domain expertise, which is both time-consuming and expensive. Manual segmentation of fibroid boundaries in MRI and ultrasound images takes significant effort, leading to smaller datasets. Table 7 shows the number of samples available for training different models, comparing medical imaging datasets with general image datasets.

Table 7. Samples for traning.

Dataset	Number of Images	Type	Annotation
Medical Dataset (Fibroids)	800	MRI, Ultrasound	Expert Manual Segmentation
ImageNet	1.2 million	General Objects	Automated Labeling

Data Imbalance: Another limitation is the imbalance in the dataset, where certain types of fibroids (e.g., large, clearly visible fibroids) are over-represented compared to smaller, more difficult-to-detect fibroids. Models trained on imbalanced datasets may struggle with small fibroid detection, leading to false negatives for harder-to-detect cases. Figure 3 shows the imbalance in fibroid sizes within the training dataset, with large fibroids making up the majority of cases.

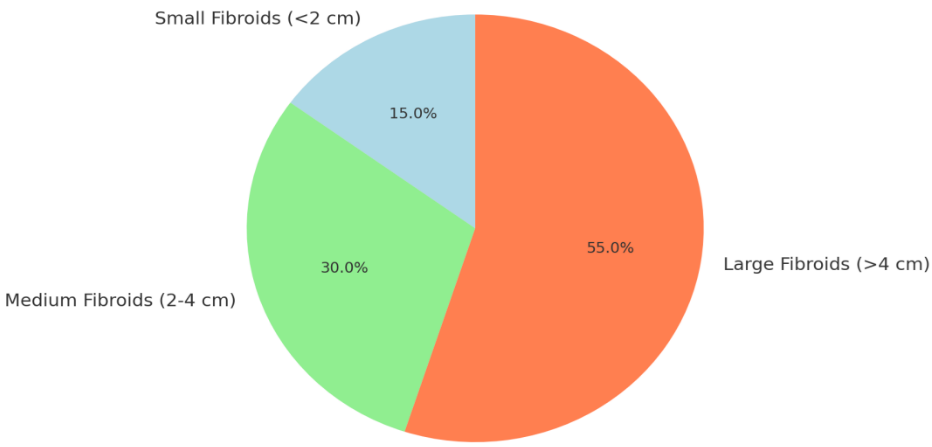

Fig.3. Imbalance in fibroid sizes within the training dataset.

Generalization to Different Imaging Modalities: Different hospitals and clinics often use varying imaging protocols, devices, and settings. This leads to heterogeneity in data, which can affect a model's ability to generalize across different imaging modalities. The texture, intensity, and contrast characteristics of fibroids can differ significantly between MRI and ultrasound images. A model trained on MRI data may not perform well when exposed to ultrasound data without additional training. Even within the same modality, differences in device manufacturers and imaging parameters can introduce variability that negatively impacts model generalization. Table 8 compares the segmentation performance of a U-Net model trained on MRI images when tested on ultrasound images.

Table 8. Segmentation performance of a U-Net model.

Model	Training Modality	Testing Modality	Dice Coefficient (%)
U-Net	MRI	MRI	92

Model	Training Modality	Testing Modality	Dice Coefficient (%)
U-Net	MRI	Ultrasound	80

Computational and Resource Demands: Training deep learning models, especially complex architectures like U-Net or ResNet, requires high computational power and memory [8]. Running these models often requires specialized hardware such as GPUs or TPUs, which are not always accessible in every medical facility. While deep learning models can achieve high accuracy, the inference speed may not be suitable for real-time clinical applications. Delays in processing images could hinder immediate clinical decision-making. Table 9 and Figure 4 compare the computational costs and inference speeds of different models in fibroid detection.

Table 9. Computational costs and inference speeds of different models.

Model	Training Time (hours)	Memory Usage (GB)	Inference Speed (images/sec)
U-Net	12	8	18
ResNet-50	15	12	15
VGG16	10	6	25

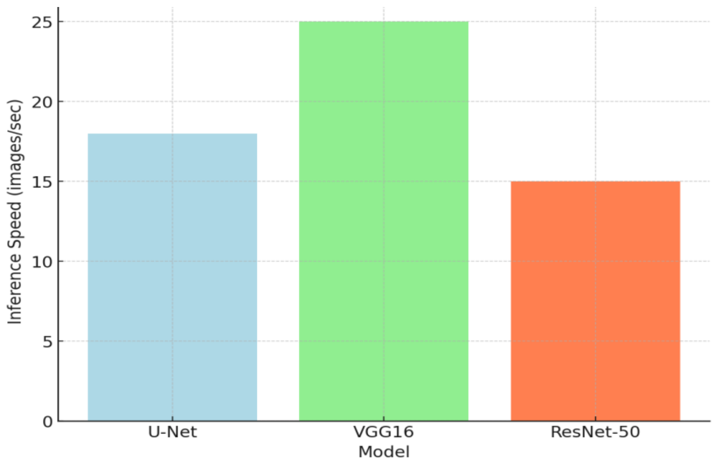

Fig.4. Comparison of inference speed across different models.

Interpretability and Explainability: Deep learning models, particularly CNNs, are often regarded as black-box systems. Clinicians may be reluctant to rely on models they cannot easily interpret, especially when making critical decisions in cases like fibroid diagnosis and treatment planning [8]. For deep learning models to be integrated into clinical workflows, they must offer explainable insights into how they arrived at a particular diagnosis. Traditional rule-based systems are easy to interpret because the decision-making process is transparent. In contrast, deep learning models derive features from data in a highly non-linear manner, making it difficult to trace the reasoning behind a given output. Figure 5 presents a heatmap generated from a CNN-based fibroid detection model, highlighting areas of the image that contributed most to the model's decision. While heatmaps provide some interpretability, more advanced explainability tools are needed.

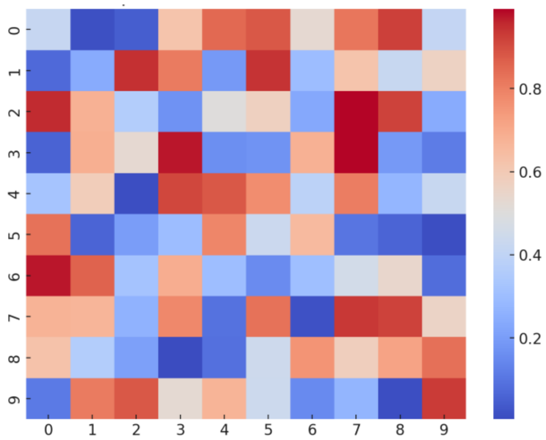

Fig. 5. A heatmap generated from a CNN-based fibroid detection model.

Model-Specific Limitations: While U-Net has demonstrated high performance in fibroid detection, particularly for segmentation, it has certain limitations. U-Net's ability to capture fine details can lead to overfitting on small datasets, causing it to perform poorly on unseen data. This is particularly problematic when the training dataset is limited. U-Net sometimes produces smoother boundaries, missing intricate edges, especially in cases where fibroids overlap with normal tissue.

ResNet-50's deep architecture allows it to capture complex patterns in the data, but it comes with trade-offs. High Memory Consumption: ResNet-50 requires more memory than U-Net and VGG16, making it difficult to deploy on standard hardware without access to high-performance GPUs. ResNet-50 is slower than both U-Net and VGG16, which may limit its use in real-time clinical environments.

VGG16 is faster than U-Net and ResNet-50, but VGG16 struggles with segmentation tasks compared to U-Net, particularly for smaller fibroids. As a primarily classification-based network, VGG16 lacks the encoder-decoder architecture needed for precise segmentation, making it less ideal for detecting fibroid boundaries.

To overcome the challenges and limitations in deep learning-based fibroid detection, several solutions can be considered: Data Augmentation to address data scarcity, techniques such as data augmentation (rotations, flips, scaling) can artificially increase dataset size and variability [9]. Transfer Learning leveraging pre-trained models on larger datasets (e.g., ImageNet) and fine-tuning them on medical imaging data can improve performance without requiring a large labeled dataset. Techniques like domain adaptation can help models generalize across different imaging modalities or devices, improving their robustness in clinical settings. Explainable AI (XAI) tools like Grad-CAM and saliency maps can improve the interpretability of deep learning models by highlighting areas in the image that were most influential in the model's decision.Deep learning, specifically CNNs, represents a major advancement in the field of uterine fibroid detection. This study demonstrates that CNN-based models can outperform traditional image processing techniques in detecting uterine fibroids from ultrasound and MRI images [9]. The implementation of such models in medical practice could lead to more accurate and timely diagnoses, ultimately improving patient outcomes.

6. Conclusions

Deep learning has revolutionized the field of medical image analysis, and its application to uterine fibroid detection has shown significant promise in improving diagnostic accuracy, consistency, and efficiency. Convolutional Neural Networks (CNNs), particularly U-Net, have proven to be highly effective for the segmentation and detection of fibroids, offering superior performance compared to traditional image processing techniques. The ability of these models to automatically learn complex patterns from raw data has reduced the reliance on manual feature engineering and improved the precision of fibroid boundary detection. However, several challenges remain, including data availability, variability in imaging modalities, and the need for interpretability. The scarcity of large, annotated datasets and the generalization of models across different imaging devices limit the widespread adoption of these techniques. Additionally, the "black-box" nature of deep learning models raises concerns about transparency and clinical trust, necessitating further research into explainable AI methods. Despite these limitations, deep learning approaches, when fine-tuned and supported by larger, more diverse datasets, have the potential to greatly enhance fibroid detection, leading to earlier diagnosis and more personalized treatment options. With continued advancements in model interpretability and the development of real-time clinical tools, deep learning-based pattern recognition is likely to play an increasingly pivotal role in the management of uterine fibroids in the near future.

Funding: This research did not receive external funding

Institutional Review Board Statement: Not applicable

Informed Consent Statement: Not applicable.

Data Availability Statement: The data of this study are available from the corresponding author upon reasonable request.

Acknowledgments: We would like to extend our sincere gratitude to the esteemed reviewers for their valuable time and effort in providing insightful feedback and constructive critiques of our manuscript. Your thorough evaluations and expert suggestions have greatly enhanced the quality and clarity of this work. We deeply appreciate your contributions to improving the scientific rigor of this paper, and we are grateful for the opportunity to benefit from your expertise. Thank you for your dedication to advancing research in this field.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Leong, W.Y.; Leong, Y.Z.; Leong, W.S. Insight into Fibroids: Cutting-Edge Detection Strategies. In *Proceedings of the 2024 IEEE 6th Eurasia Conference on Biomedical Engineering, Healthcare and Sustainability (ECBIOS)*, 2024.
2. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
3. Çiçek, Ö.; Abdulkadir, A.; Lienkamp, S.S.; Brox, T.; Ronneberger, O. 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2016, pp. 424–432.
4. Zou, J.; et al. Artificial intelligence in medical imaging of the uterine fibroids: Recent advances and future directions. *Journal of Obstetrics and Gynaecology Research* **2020**, *46*, 876–884.
5. Esteva, A.; et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* **2017**, *542*, 115–118.
6. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015, pp. 234–241.
7. Litjens, G.; et al. A Survey on Deep Learning in Medical Image Analysis. *Medical Image Analysis* **2017**, 42, 60–88.
8. Rajpurkar, P.; et al. CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning. *arXiv preprint* (preprint), 2017, arXiv:1711.05225.
9. Zhou, Z.; et al. Deep learning for medical image analysis: Applications, challenges, and future trends. *IEEE Transactions on Medical Imaging* **2020**, *39*, 1999–2024.