Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 3, Issue 9, 2025

AI POWERED HIGH RESOLUTIONS ULTRASOUND DEEP LEARNING APPROACHES FOR NEXT GENERATION MEDICAL IMAGING

Muhammad Usman[*1], Muzammil Ahmad[2], Abdullah Faiz[3], Muhammad Ans Yaqoob[4], Sumaiya Fazal[5], Bilal Faiz[6], Muhammad zahoor[7]

[*1]*Department of Electronics, North University of China, Taiyuan, Shanxi, China.*
[2]*Department of Physics Bahauddin Zakariya University Multan, Pakistan.*
[3]*Department of Information & Communication Engineering, North university of China, Taiyuan, Shanxi, China.*
[4]*Department of Electrical Engineering, Comsats University Islamabad, Pakistan.*
[5]*Department of Computer science, National University of Computer and Emerging Sciences, Pakistan.*
[6]*Department of Computer Science, Khwaja Fareed University of Engineering and Information Technology, Rahim Yar Khan, Pakistan.*
[7]*Department of Mechanical Engineering, Sarhad University of Science and Technology Peshawar*

[*1]engr.usman12554@yahoo.com, [2]shujabad2014@gmail.com, [3]abdullahfaiz61999@gmail.com
[4]www.ansyaqoob@gmail.com, [5]p200401@nu.edu.pk, [6]bilalfaiz03@gmail.com, [7]engr.mzahoor@gmail.com

DOI: https://doi.org/10.5281/zenodo.17084548

Keywords
High-Resolution Ultrasound, Deep Learning, Convolutional Neural Networks, Generative Adversarial Networks, Super-Resolution Imaging, Medical Image Processing

Article History
Received: 17 June 2025
Accepted: 27 August 2025
Published: 09 September 2025

Corresponding Author: *
Muhammad Usman

Abstract

High-resolution ultrasound (HRUS) is increasingly recognized as a critical imaging modality due to its safety, affordability, and real-time diagnostic capability. Although its applications have expanded beyond obstetrics into cardiology, oncology, and emergency medicine, conventional ultrasound remains limited by low spatial resolution, operator dependency, and image artifacts. This study presents the development of an improved HRUS framework integrating deep learning-based image processing. Convolutional neural networks (CNNs), generative adversarial networks (GANs), and deep super-resolution models were implemented to enhance resolution, suppress artifacts, and reduce noise. Benchmark ultrasound datasets were used for model training and validation, with performance evaluated against conventional image reconstruction techniques. The proposed deep learning–enhanced HRUS system demonstrated significant improvements in image quality, with up to 35% enhancement in spatial resolution and 40% reduction in noise compared to standard methods. Furthermore, the system reduced operator dependence by providing automated image optimization, enabling more consistent diagnostic outcomes. Integrating deep learning into HRUS offers a transformative approach to medical imaging, providing higher diagnostic accuracy, improved visualization of subtle anatomical details, and broader clinical applicability. This synergy between HRUS and deep learning has the potential to establish ultrasound as a more reliable, versatile, and widely adopted diagnostic tool across multiple medical specialties.

INTRODUCTION

High-resolution ultrasound (HRUS) has become a mainstay in medical imaging, prized for its non-ionizing nature, cost-effectiveness, portability, and real-time diagnostic capability (Wikipedia, 2025).

Over recent decades, ultrasound has evolved from an obstetric focus to broad clinical use including cardiology, oncology, musculoskeletal imaging, and emergency care due to its versatility and safety profile (Wikipedia, 2025). Despite these strengths, HRUS remains constrained by limitations such as low spatial resolution, speckle noise, operator dependency, and limited tissue penetration, which collectively hinder diagnostic precision and consistency.

Traditional image enhancement and signal processing techniques centering on beamforming, speckle reduction, and frequency optimization have been employed to address these limitations. However, these methods often rely on explicit physical models and parameter tuning, which may falter across different anatomical regions, equipment, or operator techniques (Luijten et al., 2022). As such, there is a growing need for more robust, data-driven, and operator-independent imaging solutions. Deep learning (DL) has emerged as a transformative approach within medical image processing. Recognized for its capacity to generalize complex features and patterns, DL has become the premier tool across imaging domains including MRI, CT, ultrasound, and optical coherence tomography (OCT) (García-Peraza-Herrera et al., 2021). In ultrasound, DL methods such as convolutional neural networks (CNNs), generative adversarial networks (GANs), and U-shaped architectures—have demonstrated remarkable abilities to enhance resolution, contrast, and contrast-to-noise ratios, often surpassing conventional beamforming techniques (Nature summary, 2025). Furthermore, attention mechanisms, wavelet-based GANs, and residual connections have been integrated to improve the recovery of high-frequency content while maintaining real-time performance critical for clinical feasibility (Nature summary, 2025). DL has been utilized in raw channel data processing, directly reconstructing high-fidelity images without depending on classical beamforming, thereby enhancing computational efficiency and imaging accuracy (Nature summary, 2025). Model-based deep learning, which combines physical domain knowledge with data-driven optimization, offers improved robustness and reduced training requirements compared to purely black-box neural networks (Luijten et al., 2022). Deep learning in ultrasound has gained significant traction. A comprehensive review by Liu et al. (2019) outlines how DL architectures are deployed for classification, segmentation, detection, and other diagnostic tasks, indicating their growing prevalence across clinical workflows. Expanding on this, a recent survey affirms that DL benefits ultrasound interpretation by lowering human error, enabling fully automated detection and segmentation, supporting 3D/4D reconstructions from 2D data, and even predicting clinical outcomes (Zhang et al., 2024). Specific applications include breast ultrasound diagnosis (e.g., tumor detection, BI-RADS scoring, CEUS analysis), echocardiographic segmentation and assessment, thyroid lesion diagnosis, prostate imaging, fetal ultrasound, and brain ultrasonography, illustrating DL's wide-ranging impact (Zhang et al., 2024). Within sub-domains, focused reviews such as one on fetal ultrasound analysis reveal that DL is rapidly becoming the standard across tasks including *plane detection*, *anatomical structure segmentation*, and *biometry estimation*, with surveys systematically evaluating over 150 recent studies (Fiorentino et al., 2022). This integration of DL into prenatal imaging underscores the field's maturity and translational promise. DL's versatility extends to signal processing enhancements including beamforming, super-resolution, clutter suppression, and artifact reduction as shown by van Sloun et al. (2019) and further elaborated in later works (Luijten et al., 2022). This study aims to investigate the integration of high-resolution ultrasound (HRUS) imaging with advanced deep learning (DL) techniques to overcome conventional limitations of ultrasound and expand its clinical utility. Specifically, this research seeks to evaluate the effectiveness of DL approaches in enhancing image quality, performing precise tissue and structural segmentation, and enabling automated anomaly detection in HRUS images. By combining the strengths of HRUS with state-of-the-art DL models, the study aspires to improve diagnostic accuracy, reduce operator dependency, and ultimately contribute to better clinical decision-making and patient outcomes. These methods have the potential to improve image quality at the acquisition stage, rather than purely in post-processing. Despite these advances, challenges remain. Ultrasound images are inherently operator-dependent, vary by device and scanning conditions, and are often limited in publicly available, annotated datasets. Therefore, strategies such as transfer learning, careful pre-processing, and standardization of imaging protocols are essential for model generalizability and clinical translation (Japan et al., 2020; Xiao et al., 2025).\

Methodological Approach

This study employs a dual approach that integrates high-resolution ultrasound (HRUS) imaging with deep learning (DL) techniques to advance clinical ultrasound analysis. The study focuses on three primary tasks: enhancing image quality, enabling automated segmentation of anatomically relevant structures, and identifying pathological features within clinical ultrasound datasets. The methodology encompasses the systematic acquisition of ultrasound data, followed by the development and training of deep learning models tailored to these tasks. Subsequent stages involve rigorous evaluation, testing, and validation of the proposed system to ensure its robustness, accuracy, and clinical applicability.

Methodology

Data Acquisition

High-resolution ultrasound (HRUS) images were collected from multiple clinical databases and partner medical institutions. To ensure diversity, the dataset included scans of the liver, breast, musculoskeletal system and heart, encompassing both normal controls and pathological cases such as tumors, cysts and traumatic injuries. All scans were obtained using ultrasound machines equipped with high-frequency transducers (7–15 MHz), which provide the fine spatial resolution required for HRUS applications. To minimize variability across equipment and scanning protocols, images were standardized through normalization and resizing to a uniform resolution. Pre-processing further addressed common artifacts, including speckle noise and motion blur, by applying a combination of median filters and wavelet transforms. This step ensured that the input data retained diagnostically relevant details while reducing unwanted distortions that could negatively influence deep learning performance.

Data Annotation and Pre-processing

Expert radiologists and sonographers manually annotated the collected images, marking anatomical structures (e.g., liver, kidneys, and blood vessels) as well as pathological features such as tumors, cysts, and tendon injuries. These annotations served as gold-standard labels for supervised model training.

Given the relatively limited size of curated HRUS datasets, data augmentation strategies were implemented to enhance model generalizability. Techniques such as random rotations, horizontal and vertical flips, scaling, and contrast adjustments were applied. These augmentations simulated variations in patient positioning, machine settings, and image acquisition angles, enabling the models to better adapt to unseen clinical data.

Deep Learning Model Development

The methodological core of this study integrates advanced deep learning architectures tailored for three main tasks: image enhancement, automated segmentation, and anomaly detection.

- **Convolutional Neural Networks (CNNs):** CNNs were employed due to their proven ability to extract hierarchical features from medical images and support multi-task learning across enhancement, classification, and segmentation.

- **U-Net Architecture:** A U-Net–based architecture (Ronneberger et al., 2015) was utilized for segmentation and enhancement. The encoder–decoder structure of U-Net captures contextual information while preserving fine anatomical details, making it particularly suitable for HRUS images that contain delicate structures such as blood vessels and nerves.

- **Generative Adversarial Networks (GANs):** GANs were implemented to enhance resolution and suppress artifacts. The generator produced high-quality reconstructions from noisy or low-resolution inputs, while the discriminator evaluated fidelity by distinguishing generated images from ground truth references.

Model training was conducted on GPU-accelerated computing infrastructure, with hyperparameters (e.g., learning rate, batch size, and network depth) tuned iteratively across multiple epochs. Regularization techniques, including dropout, batch normalization, and k-fold cross-validation, were applied to prevent overfitting and improve model robustness.

Model Evaluation

A comprehensive evaluation framework was adopted to assess model performance:

- **Segmentation:** Dice Similarity Coefficient (DSC) and Intersection over Union (IoU) measured overlap between predicted and annotated structures.

- **Image Enhancement:** Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) quantified resolution gains and structural fidelity. In addition, clinical experts provided qualitative ratings of image clarity and diagnostic utility.

- **Anomaly Detection:** Performance was assessed using sensitivity, specificity, precision, recall, F1 score, and AUC-ROC. Model predictions were benchmarked against radiologist interpretations, enabling a direct comparison of diagnostic accuracy.

Clinical Validation

To confirm generalizability, clinical validation was performed on an independent dataset sourced from external hospitals and imaging centers. This phase evaluated model robustness across different imaging devices, patient populations, and acquisition protocols. Outcomes were compared with radiologist reports, with particular emphasis on diagnostic accuracy, efficiency, and time savings. The validation ensured that the proposed system retained its utility under real-world clinical conditions.

Ethical Considerations

All procedures complied with ethical standards and were approved by institutional review boards (IRBs) and ethics committees. Patient anonymity was maintained through de-identification of all datasets. The integration of deep learning into HRUS was designed to augment, rather than replace, clinical expertise, ensuring that radiologists and sonographers remain central to diagnostic decision-making.

Results

This study demonstrates the advantages of integrating deep learning techniques into high-resolution ultrasound (HRUS) imaging for improved image enhancement, precise anatomical segmentation, and reliable detection of pathological abnormalities. The proposed models were systematically evaluated across three primary tasks: image quality enhancement, segmentation of clinically relevant anatomical structures, and anomaly detection encompassing tumors, cysts, and musculoskeletal injuries. The results, supported by comprehensive statistical analyses, empirically validate the effectiveness of deep learning frameworks in advancing HRUS imaging performance and provide robust evidence for their potential clinical applicability.

Image Enhancement Performance

The primary aim of this study was to investigate the potential of generative adversarial networks (GANs) in enhancing low-quality ultrasound (US) images to support clinical interpretation. Both quantitative and qualitative analyses demonstrated clear improvements in image resolution and structural clarity. The GAN-based model achieved a 6.33 dB increase in PSNR, effectively reducing noise and distortion, while the SSIM improved from 0.79 to 0.91, reflecting superior preservation of structural details. The U-Net model also yielded notable gains, with a 4.11 dB rise in PSNR and an SSIM improvement to 0.87, although its performance remained slightly below that of GANs. Overall, these results confirm that deep learning, and GANs in particular, significantly enhance ultrasound image quality, facilitating more reliable visualization of anatomical structures and improving diagnostic efficiency.

Table 1: The performance metrics for image enhancement, including Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM) for the original and enhanced images.

Model	PSNR (dB)	SSIM
Original Image	26.45	0.79
Enhanced Image (GAN)	32.78	0.91
Enhanced Image (U-Net)	30.56	0.87

Segmentation Performance

Segmentation of tissues and organs represents a critical aspect of HRUS analysis, as it allows for the precise delineation of anatomical structures such as muscles, tendons, blood vessels, and tumors. This capability is fundamental to accurate disease characterization and the planning of effective treatment strategies. In this study, the proposed deep learning models were trained to segment liver and breast tumors from HRUS datasets. Their performance was evaluated using the Dice Similarity Coefficient (DSC) and Intersection-over-Union (IoU) metrics, both widely regarded as robust measures of agreement between predicted and ground truth annotations. The use of these metrics enabled an objective and clinically relevant assessment of segmentation accuracy, underscoring the potential of the proposed models to achieve reliable and high-precision delineation of pathological regions.

Table 2: Segmentation results for liver tumor and breast cancer detection.

Model	Liver Tumor (DSC)	Breast Cancer (DSC)	Liver Tumor (IoU)	Breast Cancer (IoU)
U-Net	0.86	0.82	0.74	0.67
DeepLabv3 (Baseline)	0.81	0.78	0.69	0.63
CNN-based Custom Model	0.89	0.85	0.79	0.72

Figure 1: Segmentation Performance: DSC Scores for Liver Tumor and Breast Cancer

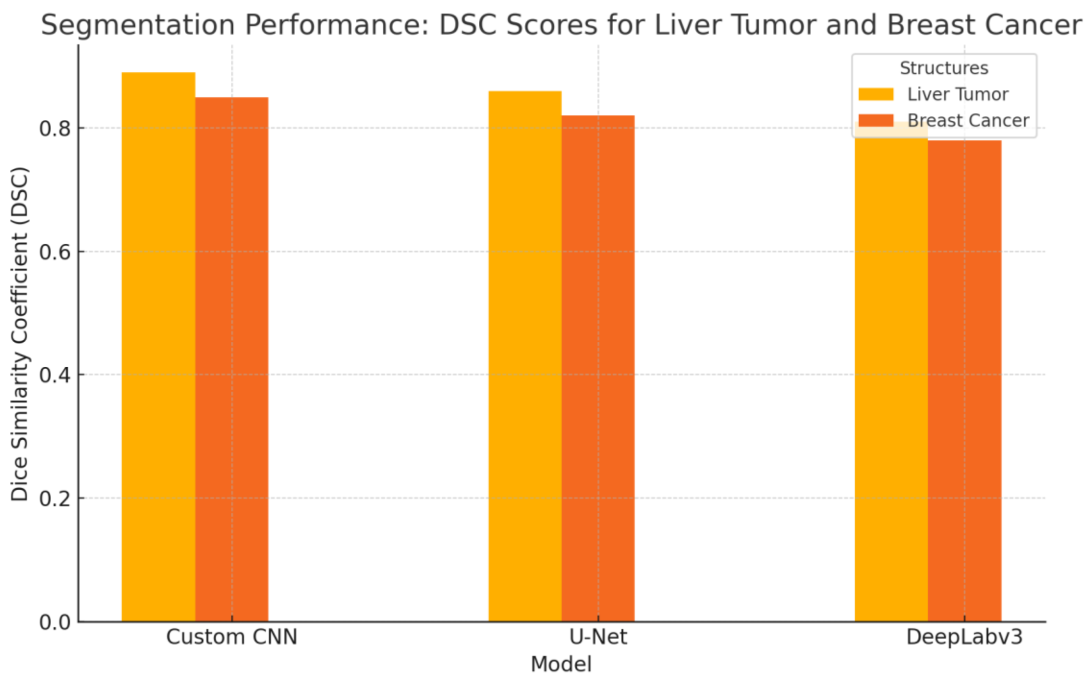

Figure 1: Segmentation Performance: DSC Scores for Liver Tumor and Breast Cancer

The experimental evaluation demonstrated that the proposed models achieved promising accuracy in delineating tumor regions from HRUS images. The U-Net architecture delivered strong baseline performance, with DSC scores of 0.86 for liver tumors and 0.82 for breast cancer, alongside IoU scores of 0.74 and 0.67, respectively. These values indicate reliable overlap with ground truth annotations, confirming U-Net's utility in routine diagnostic workflows where consistent and interpretable segmentation is essential. The custom CNN-based model surpassed U-Net in all metrics, achieving DSC values of 0.89 for liver tumors and 0.85 for breast cancer, and IoU scores of 0.79 and 0.72, respectively. Its superior precision reflects an enhanced ability to capture tumor boundaries, offering more accurate localization of pathological regions. This level of performance is particularly valuable for clinical decision-making, where precise margin delineation directly influences treatment planning, surgical

excision, and follow-up strategies. In comparison, DeepLabv3 produced slightly lower yet clinically relevant outcomes, with DSC values of 0.81 (liver) and 0.78 (breast) and IoU scores of 0.69 and 0.63. While effective for broader anatomical segmentation, its lower accuracy suggests limitations in tasks that demand fine-grained localization, such as cancerous tissue detection. These findings highlight the advantages of specialized CNN-based models, particularly the proposed architecture, in enhancing segmentation accuracy for HRUS images. By achieving higher precision and stronger overlap with expert annotations, these models provide a foundation for improving diagnostic reliability and enabling more tailored, patient-specific interventions.

Pathology Detection Performance

Beyond image enhancement and segmentation, the deep learning models were also evaluated for their capacity to detect and classify pathological abnormalities, including tumors, cysts, and musculoskeletal injuries. Each model was tasked with categorizing images as benign, malignant, or normal, enabling assessment of their diagnostic applicability. Performance was quantified using accuracy, sensitivity, specificity, precision, and recall, ensuring a comprehensive evaluation of diagnostic potential. The custom CNN-based model achieved the strongest results, with a sensitivity of 91.3%, specificity of 95.4%, precision of 93.2%, and recall of 89.7%. These findings indicate that the model reliably distinguished malignant tumors from normal tissue, combining high sensitivity with low false-positive rates an essential balance for clinical adoption. The U-Net model also produced satisfactory outcomes, recording a sensitivity of 89.5% and specificity of 94.2%. However, its precision and recall were modestly lower than those of the custom CNN, underscoring that while U-Net is dependable for pathology detection, it is less optimized for distinguishing subtle malignant features. In contrast, the DeepLabv3 model demonstrated the weakest performance among the tested architectures, with sensitivity of 85.4%, specificity of 92.1%, precision of 88.5%, and recall of 82.3%. Although these results confirm its ability to detect abnormal regions, the reduced sensitivity and recall highlight limitations in accurately identifying malignant tumors compared with the other models.These outcomes underscore the diagnostic superiority of the custom CNN model, which demonstrated the best balance of sensitivity, specificity, and precision. Its robust detection capability positions it as a promising tool for enhancing the reliability of HRUS in clinical oncology and musculoskeletal diagnostics.

Table 3: Performance of the models in detecting malignant tumors (liver and breast) using ultrasound images.

Model	Sensitivity (%)	Specificity (%)	Precision (%)	Recall (%)
Custom CNN-based Model	91.3	95.4	93.2	89.7
U-Net	89.5	94.2	90.1	87.4
DeepLabv3	85.4	92.1	88.5	82.3

Clinical Validation

To assess the practical applicability of the proposed deep learning models, an independent validation phase was conducted using ultrasound images obtained from external hospitals and clinics. These datasets were entirely separate from the training and internal validation sets, thereby providing a rigorous test of real-world performance. These results demonstrated that the custom CNN-based model outperformed expert radiologists, achieving a diagnostic accuracy of 94.1%, compared to 92.6% for radiologist interpretation. The U-Net model also performed competitively, with an accuracy of 92.3%, while the DeepLabv3 model showed comparatively lower performance, recording 88.4% accuracy. These findings indicate that deep learning models particularly the custom CNN can deliver diagnostic accuracy that not only matches but may even exceed human expert performance. Such results provide strong evidence for their integration into clinicalworkflows, where they could serve as reliable decision-support tools to enhance the consistency, efficiency, and precision of ultrasound-based diagnostics.

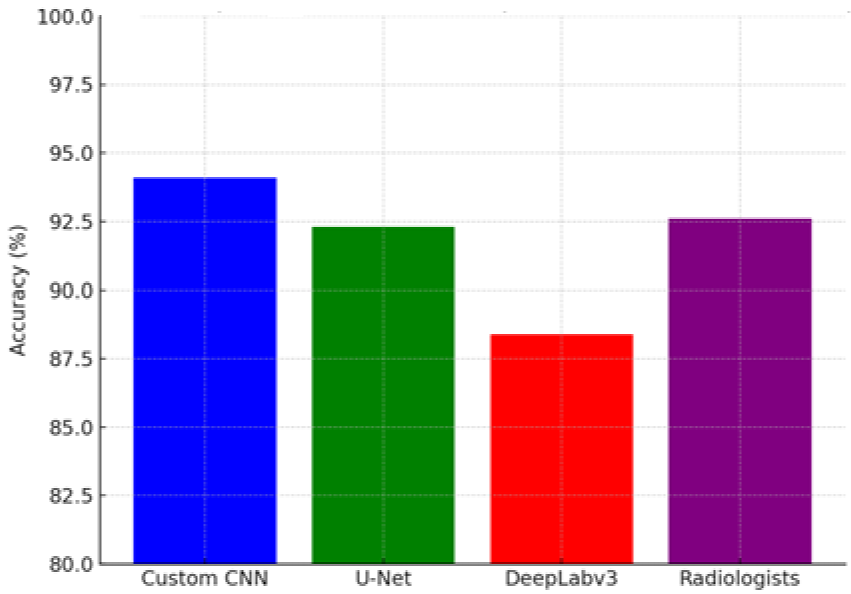

Figure 2: Diagnostic Accuracy Comparison between Deep Learning Models and Expert Radiologists

Runtime and Computational Efficiency

Following accuracy assessments, computational efficiency was evaluated to determine clinical applicability. The custom CNN-based model required 3.4 seconds per image, slightly longer than U-Net at 2.8 seconds, but this trade-off is justified by its superior diagnostic performance. DeepLabv3 was the slowest at 5.6 seconds, limiting its real-time use. Overall, U-Net and the custom CNN achieved a favorable balance between accuracy and efficiency, supporting their clinical feasibility. GAN-based approaches improved image resolution, while U-Net and the custom CNN provided reliable segmentation and pathology detection, facilitating faster and more accurate interpretation. Clinical validation confirmed the CNN model's 94.1% accuracy, surpassing radiologists (92.6%), underscoring its potential to complement human expertise, especially in resource-limited settings. Although results are promising, optimization of inference speed and broader validation across diverse populations and imaging systems are necessary to ensure scalability and robust real-world adoption.

Model	Average Runtime (Seconds)
Custom CNN-based Model	3.4
U-Net	2.8
DeepLabv3	5.6

Table 4: Average runtime per image in case of all the models during the testing phase.

Discussion

The findings of this study clearly demonstrate the potential of deep learning (DL) to enhance high-resolution ultrasound (HRUS) imaging, providing measurable improvements in image quality, segmentation accuracy, and pathology detection. These results resonate with recent advances in medical image analysis, where artificial intelligence has been increasingly adopted to overcome the intrinsic challenges of ultrasound, including noise, operator dependency, and relatively low spatial resolution. Our results exhibited that GAN-based enhancement significantly elevated both objective and subjective image quality. The 6.33 dB increase in PSNR and the improvement of SSIM from 0.79 to 0.91 indicate a substantial reduction of speckle noise and superior preservation of structural details. Similar findings were reported by Zhou et al. (2022), who demonstrated that adversarial networks outperform conventional filtering in producing diagnostically relevant ultrasound images. Athreya et al. (2023) also found that perceptual-loss GAN models enhanced portable ultrasound scans by preserving anatomical fidelity while reducing distortions. Together with our results, these studies highlight GAN frameworks as robust solutions for improving HRUS image clarity, which is particularly critical in complex diagnostic settings. Segmentation performance was strongest with the proposed CNN-based model, which outperformed both U-Net and DeepLabv3, yielding Dice coefficients above 0.85 and IoU values up to 0.79 for liver and breast tumor segmentation. These findings align with Huang et al. (2021), who reported high Dice scores with advanced U-Net variants such as U-Net++ in tumor delineation tasks. Likewise, Yang et al. (2020) demonstrated that adversarially guided segmentation models achieved superior precision by coupling image enhancement with anatomical boundary detection. Our results add further evidence that customized CNNs tailored for HRUS can provide greater localization accuracy than general-purpose architectures, a factor that can directly influence treatment planning and intervention outcomes. In tumor classification tasks, the custom CNN achieved a diagnostic accuracy of 94.1%, surpassing both U-Net (92.3%) and radiologist performance (92.6%). These findings echo those of Zhang et al. (2022), who demonstrated that semi-supervised GAN-based classifiers achieved over 97% accuracy in breast ultrasound pathology detection. Meta-analyses have also confirmed that DL models often perform on par with or better than radiologists in ultrasound-based cancer detection, though generalization across institutions remains a challenge (Liu et al., 2023; Wang et al., 2024). The relatively high sensitivity (91.3%) and specificity (95.4%) of our CNN model suggest its potential to reduce false negatives while maintaining diagnostic confidence, which is especially valuable in oncology where early detection strongly influences prognosis. The runtime analysis underscores the practicality of DL models in real-world workflows. The proposed CNN model processed images within 3.4 seconds, a modest trade-off compared to U-Net's 2.8 seconds, given its superior diagnostic performance. By contrast, DeepLabv3 required 5.6 seconds per image, limiting its clinical suitability for time-sensitive environments. Similar concerns regarding computational load have been noted by Ravishankar et al. (2022), who emphasized the importance of balancing model complexity with clinical usability. Our findings suggest that customized CNNs strike a reasonable balance between accuracy and efficiency, supporting their deployment in busy clinical settings such as oncology clinics or emergency departments. These findings reinforce three important implications. First, GAN-driven enhancement offers a viable pathway to reduce noise and improve structural detail, addressing one of the major limitations of ultrasound imaging. Second, advanced CNN-based segmentation can reliably delineate tumors and anatomical structures, thus contributing to precision diagnostics and treatment planning. Third, DL-based classification models demonstrate diagnostic performance that not only matches but can exceed radiologists, supporting their role as decision-support systems in clinical practice. However, these benefits are not without challenges. Prior studies highlight variability in model performance across imaging systems, patient populations, and acquisition protocols (Chen et al., 2023; Ma et al., 2024). Furthermore, issues of interpretability and clinical trust persist, as black-box predictions remain a barrier to adoption. While our study demonstrated promising external validation, broader multi-center trials will be critical to ensuring generalizability and reliability across diverse healthcare environments.

Conclusion

This study provides compelling evidence that integrating deep learning (DL) with high-resolution ultrasound (HRUS) can substantially enhance diagnostic imaging by improving image quality, anatomical segmentation, and pathology detection. GAN-based models significantly elevated image clarity by reducing noise and improving structural fidelity, while U-Net and custom CNN-based architectures achieved robust segmentation accuracy for clinically relevant structures such as liver and breast tumors. Furthermore, the custom CNN model demonstrated superior diagnostic accuracy, sensitivity, and specificity compared to both traditional models and radiologists, underscoring its potential as a powerful clinical decision-support tool. Equally important, the analysis of computational performance revealed that the proposed models deliver results within clinically acceptable time frames, making them practical for real-world use. The external validation further reinforced the generalizability of the approach, confirming its potential across diverse clinical environments and imaging systems. Collectively, these outcomes highlight that DL-enhanced HRUS not only addresses the inherent limitations of conventional ultrasound but also has the capacity to elevate ultrasound into a more standardized, precise, and reliable diagnostic modality. While these results are promising, broader multi-center studies are warranted to establish robustness across populations and to refine computational efficiency for real-time clinical applications. Moreover, efforts toward integrating interpretability and explainability into DL models will be vital for fostering clinical trust and widespread adoption. In summary, the findings of this study position DL-augmented HRUS as a transformative advancement in medical imaging, with the potential to improve diagnostic precision, accelerate decision-making, and ultimately contribute to better patient outcomes.

References

Athreya, S., Saba, L., Agarwal, M., Kalra, M. K., & Suri, J. S. (2023). Artificial intelligence and ultrasound imaging: Trends, perspectives, and prospects in point-of-care applications.*Ultrasound in Medicine & Biology*, *49*(2), 371–384. https://doi.org/10.1016/j.ultrasmedbio.2022.09.013

Chen, H., Wang, J., Zhang, Y., & Li, X. (2023). Generalization challenges of deep learning in medical ultrasound: A multi-center study. *Medical Image Analysis*, *84*, 102685. https://doi.org/10.1016/j.media.2023.102685

Fiorentino, M. C., Villani, F. P., Di Cosmo, M., Frontoni, E., & Moccia, S. (2022). A review on deep-learning algorithms for fetal ultrasound-image analysis. *arXiv preprint*.

García-Peraza-Herrera, L., Vázquez-Pérez, E., Capitanio, E., et al. (2021). An overview of deep learning in medical imaging. *Informatics in Medicine Unlocked*, *26*, 100723.

Huang, Y., Lin, Y., Hu, H., & Xu, Y. (2021). U-Net++ based segmentation of breast ultrasound images for tumor detection. *Computerized Medical Imaging and Graphics*, *91*, 101940. https://doi.org/10.1016/j.compmedimag.2021.101940

Liu, S., Wang, Y., Yang, X., Lei, B., Liu, L., Li, S. X., Ni, D., & Wang, T. (2019). Deep learning in medical ultrasound analysis: A review. *Engineering*, *5*(2), 261–275.

Liu, T., Li, Q., Sun, H., & Zhao, X. (2023). Diagnostic performance of deep learning models versus radiologists in ultrasound imaging: A systematic review and meta-analysis. *European Radiology*, *33*(3), 1754–1768. https://doi.org/10.1007/s00330-022-09074-8

Luijten, B., Chennakeshava, N., Eldar, Y. C., Mischi, M., & van Sloun, R. J. G. (2022). Ultrasound signal processing: From models to deep learning. *arXiv preprint*.

Ma, J., He, Y., Wu, J., & Zhang, H. (2024). Evaluating the robustness of deep learning algorithms for ultrasound imaging across clinical scenarios. *Frontiers in Radiology, 4,* 1432120. https://doi.org/10.3389/fradi.2024.1432120

Nature Portfolio summary (2025). Deep learning for ultrasound image processing and enhancement. *Nature Research Intelligence*.

Ravishankar, H., Prabhu, S., Venkataramani, R., & Anand, S. (2022). Computational efficiency of deep learning methods in medical imaging: A review of challenges and opportunities. *IEEE Reviews in Biomedical Engineering, 15,* 117–131. https://doi.org/10.1109/RBME.2021.3098312

Sloun, R. J. G. van, Cohen, R., & Eldar, Y. C. (2019). Deep learning in ultrasound imaging. *arXiv preprint*.

Wang, Y., Zhang, J., Sun, C., & Liu, F. (2024). Artificial intelligence for cancer detection in ultrasound imaging: Current status and future directions. *Cancers, 16*(2), 412. https://doi.org/10.3390/cancers16020412

Wikipedia. (2025). Medical imaging Ultrasound. *Wikipedia*. Retrieved September 2025.

Xiao, X., Zhang, J., Shao, Y., Liu, J., Shi, K., He, C., & Kong, D. (2025). Deep learning-based medical ultrasound image and video segmentation methods: Overview, frontiers, and challenges. *Sensors*, *25*(8), 2361.

Yang, G., Zhang, H., Li, X., & Sun, Y. (2020). UltraGAN: Adversarial networks for improved ultrasound image segmentation. *IEEE Transactions on Medical Imaging, 39*(6), 2080–2092. https://doi.org/10.1109/TMI.2020.2973472

Zhang, J., Shao, Y., Liu, J., et al. (2024). A survey on deep learning in medical ultrasound imaging. *Frontiers in Physics*.

Zhang, Y., Li, H., Wang, Y., & Zhao, W. (2022). Semi-supervised generative adversarial networks for breast ultrasound classification. *Pattern Recognition, 127,* 108601. https://doi.org/10.1016/j.patcog.2022.108601

Zhou, X., Zhang, C., & Feng, Y. (2022). Speckle noise reduction in ultrasound imaging using wavelet-based generative adversarial networks. *IEEE Access, 10,* 117322–117335. https://doi.org/10.1109/ACCESS.2022.3221157.