Sciences Methods and Technologies International Journal (SciMeTech)

ISSN: 3085-5284, journal homepage: www.scimetech.com

Deep Learning-Based Diagnosis-Guided Framework for High-Quality Medical Image Super-Resolution

Km Annoo[1]*, Amit Sharma[2], Suman Pandey[3]
Affiliation 1.[1,2,3] Department of Computer Science and Engineering,
Vivekananda Global University, Jaipur, India
**Corresponding Author: Km Annoo,*
ananyaannusmile@gmail.com

How to cite:
Km Annoo, Amit Sharma, Suman Pandey, "Deep Learning-Based Diagnosis-Guided Framework for High-Quality Medical Image Super-Resolution", Sciences Methods and Technologies International Journal (SciMeTech), Vol 2, Issue 1, p 161-170

Keywords:
Medical image super-resolution
diagnosis-guided attention
deformable convolution
Alzheimer's disease diagnosis

Abstract

Medical image quality plays a crucial role in accurate disease diagnosis. However, due to limitations in imaging devices and acquisition conditions, medical images are often captured at low resolution, which can negatively affect clinical interpretation. To address this issue, this paper proposes a diagnosis-guided super-resolution framework that enhances both image quality and diagnostic performance.

The proposed model integrates two interconnected components: a super-resolution network and a disease diagnosis network. Unlike conventional approaches that focus solely on visual enhancement, the proposed method incorporates diagnostic feedback to guide the image reconstruction process. A diagnosis-guided attention mechanism is introduced to ensure that the model focuses on clinically relevant regions during enhancement. Additionally, deformable convolution is employed to effectively capture irregular anatomical structures. Experiments conducted on multiple medical imaging datasets demonstrate that the proposed framework significantly improves both image reconstruction quality and diagnostic accuracy, as measured by PSNR, SSIM, and classification performance. The results indicate that integrating super-resolution with diagnosis guidance leads to more reliable and clinically meaningful outcomes.

I. Introduction

Medical image analysis has become an essential part of modern healthcare systems, as it supports important clinical tasks such as disease detection, diagnosis, and treatment planning [10]. However, in real-world scenarios, medical images are often captured at lower resolutions due to various practical constraints, including limitations of imaging devices, shorter acquisition time, and patient-related factors. For instance, in Magnetic Resonance Imaging (MRI), acquiring high-resolution images typically requires longer scanning time, which increases the chances of motion-related distortions [11]. Similarly, in Positron Emission Tomography (PET), reducing radiation exposure is important for patient safety, but it often leads to lower image quality [12]. Because of these challenges, enhancing image resolution has become an important research problem in medical imaging. Super-resolution (SR) techniques are widely used to reconstruct high-resolution images from low-resolution inputs [9]. In recent years, deep learning methods such as convolutional neural networks (CNNs) and transformer-based models have shown strong performance in improving image quality [13], [14]. These approaches have been successfully applied in natural image processing and are now being explored for medical imaging applications [5]. However, applying traditional SR methods directly to medical images is not always effective. Unlike natural images, where visual quality is the primary goal, medical images require preservation of clinically important details [1]. In many cases, small anatomical structures or subtle abnormalities play a critical role in diagnosis. Therefore, simply improving visual appearance is not sufficient for medical applications [6]. Another challenge is that super-resolution is an ill-posed problem, meaning that multiple high-resolution outputs can be generated from the same low-resolution input [9]. This uncertainty may introduce artefacts or distort important features, which can negatively impact clinical decisions [7]. Hence, it is necessary to design SR approaches that consider diagnostic requirements. To address these issues, this work proposes a diagnosis-guided super-resolution framework, where image enhancement and disease diagnosis are performed together. Instead of treating super-resolution as an independent task, the proposed approach integrates diagnostic feedback into the reconstruction process. Alzheimer's disease is considered a case study; however, the framework can be extended to other medical conditions as well. The main objective of this work is to generate high-resolution images that are not only visually improved but also clinically meaningful. By combining super-resolution with diagnosis guidance, the proposed method aims to improve both image quality and diagnostic reliability.

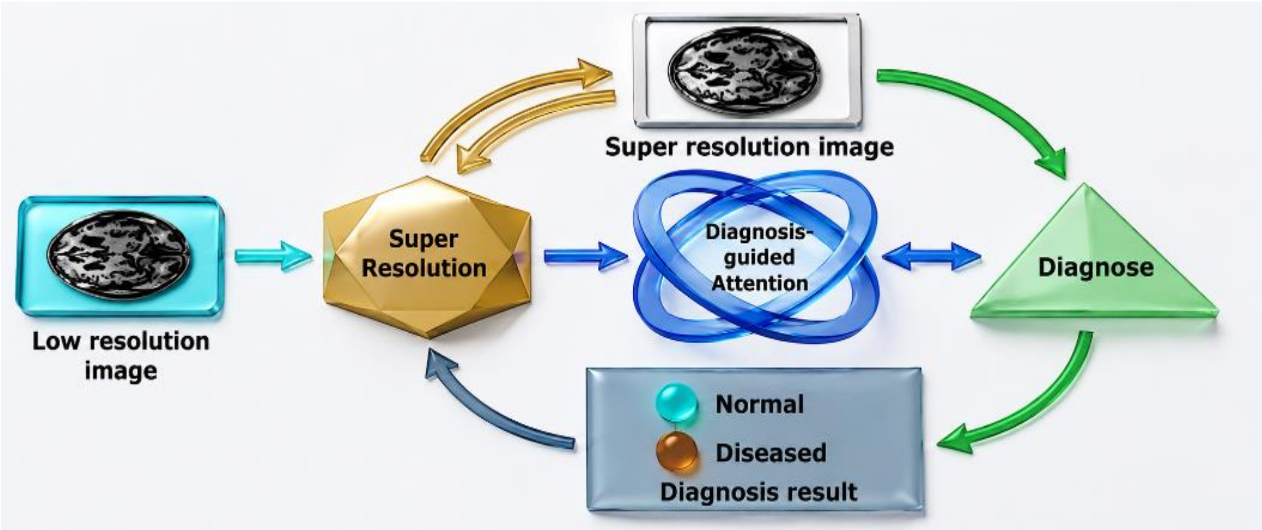

Fig. 1. The overview of SRDA.

As illustrated in Fig.1, the overall architecture of the proposed SRDA framework is presented.

Step 1: Super Resolution Module

We start with a low-resolution image, I_LR. A super-resolution network $f_{SR}(.)$ generates a high-resolution image:

$$I_{SR} = f_{SR}(I_{LR}; \theta_{SR})$$

where I_LR denotes the low-resolution input image,
I_SR represents the reconstructed high-resolution image,
and θ_SR denotes the learnable parameters of the super-resolution network.

Step 2: Diagnosis Module

The diagnosis network predicts the probability of disease:

$$P = F_D(I_{SR}; \theta_D)$$

Where θ_D represents the parameters of the diagnosis network, and P represents the predicted probability distribution (Normal vs Diseased) using softmax activation.

Step 3: Diagnosis-Guided Attention

The attention map A is generated as:

$$A = F_A(P, I_{SR})$$

This attention map refines both processes:

$$I'_{SR} = A \cdot I_{SR}$$
$$P' = F_D(I'_{SR})$$

Step 4: Feedback Loop

The system is optimised iteratively:

$$\theta_{SR}, \theta_D = armin\ L_{total}$$

Step 5: Loss Function

The total loss is defined as:

$$L_{total} = \lambda_1\ L_{SR} + \lambda_2\ L_{diag} + \ \lambda_3\ L_{att}$$

where:

- L_{SR}= reconstruction loss
- L_{diag} = diagnosis loss (cross-entropy)
- $\lambda_1, \lambda_2, \lambda_3$= weighting parameters

This paper introduces a novel medical image super-resolution framework, termed Super-Resolution via Diagnosis-Guided Attention (SRDA), designed to address the limitations discussed earlier. The primary objective of this approach is to ensure that the generated high-resolution images are not only visually improved but also beneficial for downstream diagnostic tasks, without introducing misleading artefacts. To achieve this, a unified multi-task learning strategy is adopted, in which super-resolution and disease diagnosis are optimised jointly. As illustrated in Fig. 1, the proposed framework comprises two collaborating components: a super-resolution (SR) network and a diagnosis network. Initially, low-resolution medical images are processed by the SR network to generate enhanced high-resolution outputs. These reconstructed images are then directly passed to the diagnosis network for disease prediction. The interaction between these two networks is bidirectional. First, diagnostic loss is propagated back through the entire system during end-to-end training, enabling the SR network to learn parameters that are aligned with diagnostic objectives. Second, a dedicated diagnosis-guided attention mechanism is incorporated to guide the SR network toward regions that are clinically significant. In the design of the SR network, special emphasis is placed on meeting the requirements of medical diagnosis. Clinical decision-making often depends on specific regions of interest (ROIs). For instance, previous studies have highlighted the importance of areas such as the middle temporal gyrus and the hippocampal formation in Alzheimer's disease diagnosis. However, these ROIs typically exhibit irregular shapes and varying scales, making them difficult to capture using standard convolutional neural networks with fixed receptive fields. To overcome this challenge, deformable convolution is integrated into the SR network, allowing adaptive feature extraction from regions with diverse spatial characteristics. This enhances the model's ability to preserve critical structural and pathological details. Furthermore, attention mechanisms have proven effective in emphasising diagnostically relevant regions in medical imaging tasks. Motivated by this, a diagnosis-driven attention module is introduced to strengthen the interaction between the SR and diagnosis networks. The key idea is that regions critical for diagnosis should receive greater emphasis during the super-resolution process. To implement this, both networks are equipped with attention modules of identical architecture. By enforcing consistency between the attention maps generated by these modules, the framework encourages both networks to focus on the same clinically relevant regions. This design enables a deeper integration of super-resolution and diagnostic learning compared to conventional approaches, which typically rely only on diagnostic loss for indirect guidance. In contrast, the proposed framework explicitly incorporates diagnostic information into the SR process. The combination of deformable convolution and diagnosis-guided attention allows both tasks to complement each other effectively. The SR network produces more accurate reconstructions of clinically important regions, thereby improving diagnostic performance. Simultaneously, the diagnosis network identifies key ROIs that guide the SR network in generating more informative and clinically meaningful high-resolution images.

a. Contributions of this study are summarised as follows:

- A new diagnosis-guided framework for super-resolution of medical images is suggested. This framework makes sure that the reconstructed images are better for diagnostic tasks that come after them.
- A diagnosis-guided attention module is carefully designed to make sure that the SR and diagnosis networks are looking at the same thing, which makes it easier to see diagnostically important areas.
- Extensive experiments were conducted on multiple datasets. Different medical imaging datasets show that the proposed method works well and is better than current best practices.

I. Methods

This section provides a detailed description of the proposed Super-Resolution via Diagnosis-Guided Attention (SRDA) framework. The overall architecture of the model is illustrated in Fig. 2. The framework consists of two interconnected sub-networks that operate collaboratively: a super-resolution (SR) network responsible for reconstructing high-resolution images, and a diagnosis network designed for disease prediction. To establish a strong connection between these components, a novel diagnosis-guided attention module is introduced. This module identifies diagnostically significant regions of interest (ROIs) and ensures that both networks focus on these critical areas. Within this unified design, super-resolution and diagnosis tasks complement each other. The SR network enhances image quality, which improves diagnostic performance through forward propagation, while the diagnosis network influences the SR process via backpropagation of the diagnostic loss. The detailed architecture and functionality of the SR network, diagnosis network, and attention module are discussed in the subsequent subsections. For demonstration purposes, a ×4 upscaling factor is considered in this work; however, the proposed framework is flexible and can be extended to other scaling factors without modification.

a. Super-Resolution Network.

The SR network aims to extract discriminative features from low-resolution images and reconstruct corresponding high-resolution images. Given the strong performance of residual networks (ResNet) in super-resolution tasks, a ResNet-based architecture is employed for feature extraction, as shown in Fig.

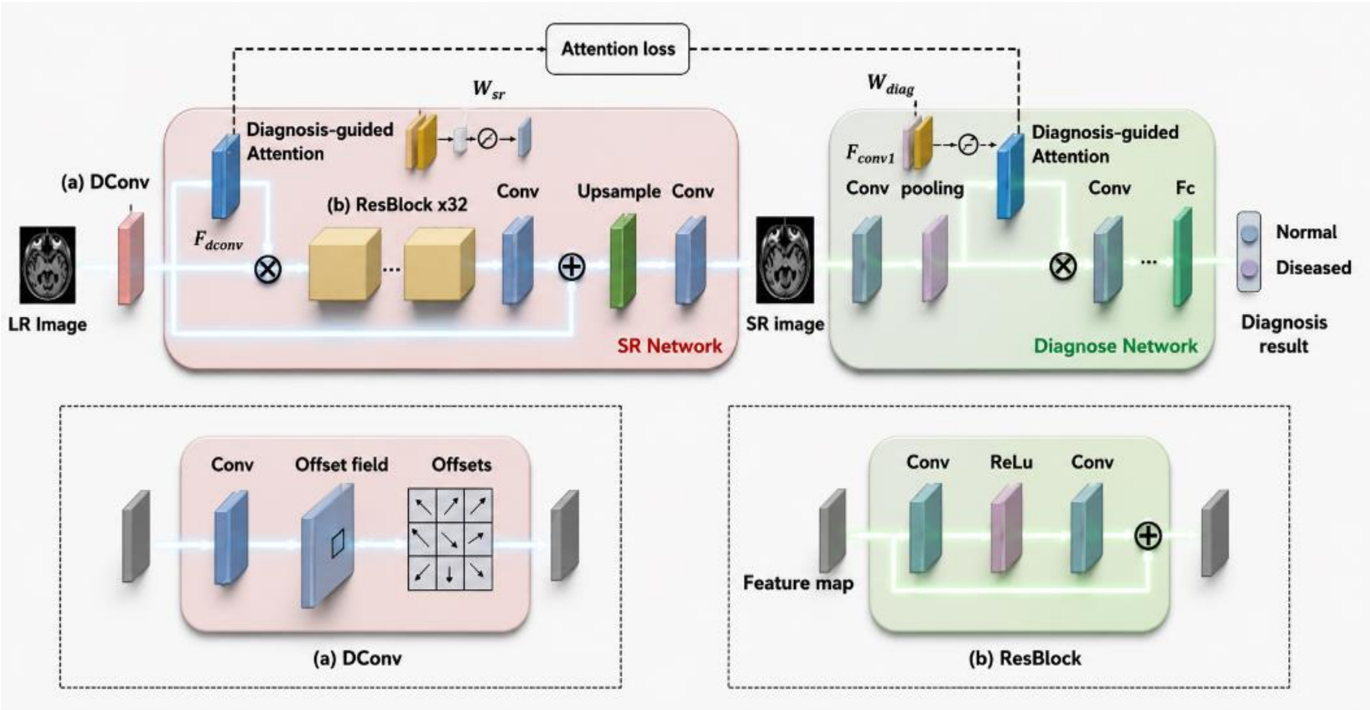

Fig. 2 Deep Learning Framework for Medical Image Super-Resolution and Diagnosis

Fig. 2 shows how the SRDA framework is constructed. It has the SR and diagnosis networks in it. The effect of SR on diagnosis is due to forward propagation along the network, while the effect of diagnosis on SR is due to back propagation of diagnosis loss. Unlike natural image SR, medical image super-resolution is typically followed by downstream diagnostic tasks. Therefore, feature extraction should emphasise diagnostically important regions rather than purely visual details. In many diagnosis tasks, only specific ROIs—such as the hippocampal formation in Alzheimer's disease—are critical. These ROIs often exhibit irregular shapes and varying spatial scales, which are difficult to model using fixed-size convolution kernels.

1. Input and Super-Resolution Network

Start with a low-resolution image:

$$I_{LR} \in \mathbb{R}^{H\times W\times C}$$

Deformable Convolution (DConv)

For a convolution kernel $K$with offsets Δp:

$$y(p_0) = \sum_{p_n\in\mathcal{R}} w(p_n)\cdot x(p_0 + p_n + \Delta p_n)$$

where:

- $\mathcal{R}$= receptive field
- $w(p_n)$= kernel weights
- Δp_n= learned offsets

Residual Blocks

Each residual block:

$$F_{res}(x) = x + f(x;\theta_{res})$$

Stacked 32 times:

$$I_{res} = F_{res}^{32}(I_{LR})$$

Upsampling + Convolution

$$I_{SR} = f_{up}(I_{res};\theta_{up})$$

2. Diagnosis Network

The diagnosis CNN processes the super-resolved image:

$$z = f_{CNN}(I_{SR};\theta_{CNN})$$

After pooling and fully connected layers:

$$y = \text{softmax}(Wz + b)$$

where $y \in \{y_{normal}, y_{diseased}\}$.

3. Diagnosis-Guided Attention

Attention weights are derived from the diagnosis output:

$$\alpha = A(y;\theta_A)$$

This modulates both SR and diagnosis:

$$I'_{SR} = \alpha\cdot I_{SR} + (1-\alpha)\cdot I_{LR}$$
$$y' = f_{CNN}(I'_{SR};\theta_{CNN})$$

4. Feedback Loop

Iterative refinement:

$$I_{SR}^{(t+1)} = f_{SR}(I_{LR};\theta_{SR},\alpha^{(t)})$$

$$y^{(t+1)} = f_D(I_{SR}^{(t+1)}; \theta_D, \alpha^{(t)})$$
$$\alpha^{(t+1)} = A(y^{(t+1)}; \theta_A)$$

5. Loss Functions

Super-Resolution Loss

$$\mathcal{L}_{SR} = \| I_{HR} - I_{SR} \|^2$$

Diagnosis Loss

$$\mathcal{L}_D = - \sum_{c \in \{normal, diseased\}} y_{true}(c) \log y(c)$$

Attention Loss

$$\mathcal{L}_A = \| \alpha - \alpha_{target} \|^2$$

Joint Loss

$$\mathcal{L} = \lambda_{SR}\mathcal{L}_{SR} + \lambda_D\mathcal{L}_D + \lambda_A\mathcal{L}_A$$

Terms are defined as follows

- **DConv + ResBlocks** → super-resolution image.
- **Diagnosis CNN** → Normal/Diseased classification.
- **Attention mechanism** → feedback loop refining both SR and diagnosis.
- **Joint loss** → ensures image fidelity, diagnostic accuracy, and attention consistency.

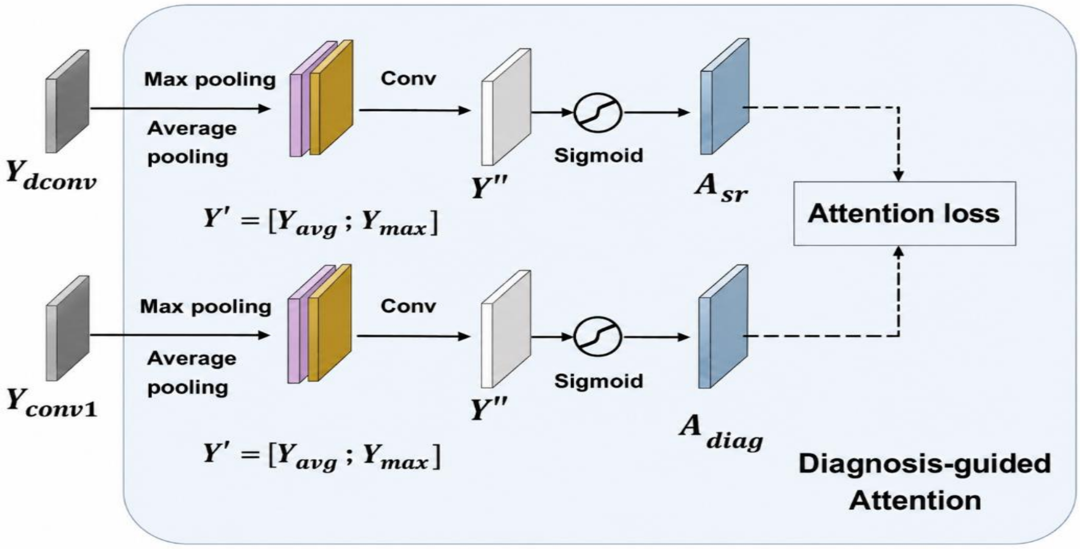

Fig. 3. The diagnosis-guided attention module.

The above figure describes a diagnosis-focused attention that helps in improving the feature representation of neural networks. It features two parallel input streams, $Y_{dconv}$and Y_{conv1}, each undergoing max pooling and average pooling to extract salient spatial features. These pooled outputs are concatenated into $Y' = [Y_{avg}; Y_{max}]$, which is then refined through a convolutional layer to produce Y''. A sigmoid activation function transforms Y''into attention maps $A_{sr}$and A_{diag}, which guide the model's focus toward diagnostically relevant regions. Although contributing to an attention loss that regularises learning, intermediate class activations are also used in the bottom path, where diagnosis-specific cues are integrated for increased interpretability and classification accuracy.

Proposed Flow Chart

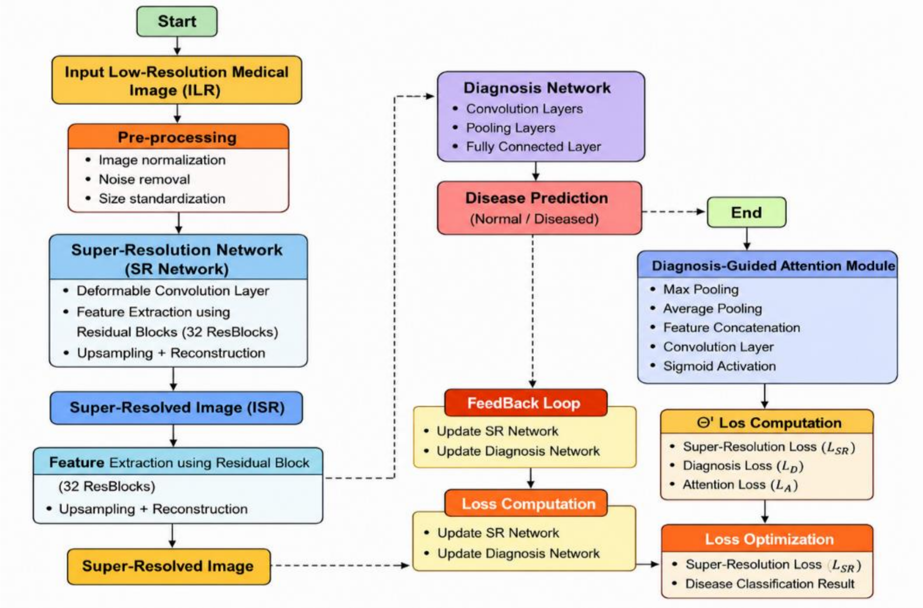

Fig.4. Diagnosis-Guided Super-Resolution and Classification Pipeline

This image is a flow chart showing the full cycle of a deep learning system for generating and examining low-res medical images. This process starts with taking an input low-way image (ILR), followed by some pre-processing such as normalisation, noise removal and size standardisation. Followed by two Deformable Segmented Residual Blocks DSResBlock, we then apply a Super-Resolution Network (SRN), which consists of convolutions, residual blocks and upsampling to the input image ISR. The CNN-based Diagnosis Network ingests this enhanced image and predicts disease status. A Diagnosis-Guided Attention Module improves the image based on the attention maps generated by pooling and convolution operations to create an attention-weighted image. The corresponding networks are updated using a feedback loop in the system, and multiple loss functions such as super-resolution, diagnosis and attention losses for joint optimisation are calculated. At the end, it shows a high-resolution image and an accurate result of disease classification.

II. Results and discussion

a. Data sets

In the Results section, showcase images/photos of your projectThe proposed method is tested on three publicly available datasets: ADNI-MRI, ADNI-PET, and Demented-MRI, to see how well it works for diagnosing Alzheimer's disease. The Alzheimer's Disease Neuroimaging Initiative (ADNI) gives out ADNI-MRI and ADNI-PET, which are 3D MRI and PET brain scans, respectively. There are 0.01 mm gaps between each 2D image that is cut from a 3D volume along the axial plane. The data are split into three sets: training, validation, and testing. The ratio is 8:1:1. For ADNI-MRI, 4080 images are used to train the model, and 510 images are used to test and validate it. For ADNI-PET, 9600 pictures are used for training, and 1200 pictures are used for testing and validation. The Demented-MRI dataset has 2670 2D MRI images that are split into three groups: non-demented, very mild-demented, and mild-demented. Out of these, 2136 images are used for training, and 267 images are used for both validation and testing.

b. Experimental Setup and Implementation Details

Bicubic interpolation is the base method. We compare the proposed method to a number of cutting-edge SR methods, such as RCAN, IMDN, SWD, SwinIR, SRD, A2F, and ELAN. We use Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) to check how well the reconstruction works. Higher values mean better quality. The datasets do not have pairs of low- and high-resolution images, so the original high-resolution images are used as ground truth, and low-resolution images are made by downsampling the high-resolution images.

c. Experimental Results

Table 1 shows the quantitative differences between the proposed SRDA method and other cutting-edge super-resolution methods. The proposed method consistently achieves the highest PSNR and SSIM values across all three datasets, demonstrating its superior reconstruction performance. The proposed SRDA framework significantly outperforms existing methods across all datasets compared to both the best traditional SR methods and the newest diagnosis-aware SR methods on the ADNI-MRI, ADNI-PET, and Demented-MRI datasets. These results indicate that incorporating diagnostic guidance into the SR process markedly improves reconstruction quality, particularly in clinically significant areas. Figures 4–6 show the qualitative results of different methods that were used and tested on the datasets. Compared to older methods, SRDA outputs keep more structural details and sharper anatomical boundaries. These high-frequency details are critical for accurate clinical diagnosis. The proposed method generates super-resolved images that are more suitable for downstream diagnostic tasks by enhancing diagnostically important regions. by making diagnostically important areas clearer.

d. Ablation Study

An ablation study is performed to analyse how much each element in our framework contributes to the performance improvement. Several model variants are considered:

SR: The base SR model without deformation convolution and without the diagnosis network or attention module.

SR + Dconv SR model with deformable convolution only.

SR + Dconv + Diag: SR model with deformable convolution and diagnosis network, but without attention.

SR + Dconv + Diag + Atten (SRDA): Full proposed model.

The quantitative results of the ablation study are reported in Table 2. Compared with the baseline SR model, introducing deformable convolution leads to noticeable performance improvements, indicating its effectiveness in capturing ROIs with irregular shapes and scales. Further performance gains are achieved by incorporating the diagnosis network, which confirms that diagnostic supervision can effectively guide the SR process. Moreover, the addition of the diagnosis-guided attention module yields the best overall performance. This demonstrates that aligning the attention of the SR and diagnosis networks enables the model to better focus on diagnostically important regions, resulting in more accurate and informative super-resolved images.

Table 2: Comparative Performance of Super-Resolution Methods on Medical

Method	PSNR	SSIM	Rank	Category
Ours	36.63	0.969	1	Advanced Deep Learning
RCAN	35.48	0.9603	2	Advanced Deep Learning
A²F	35	0.9553	3	Advanced Deep Learning
SwinIR	34.05	0.9458	4	Advanced Deep Learning
SRD	33.35	0.9373	5	Advanced Deep Learning
IMDN	33.01	0.9302	6	Advanced Deep Learning
SWD	32.88	0.9306	7	Advanced Deep Learning
ELAN	32.55	0.9212	8	Advanced Deep Learning
Bicubic	25.5	0.8332	9	Classical Interpolation

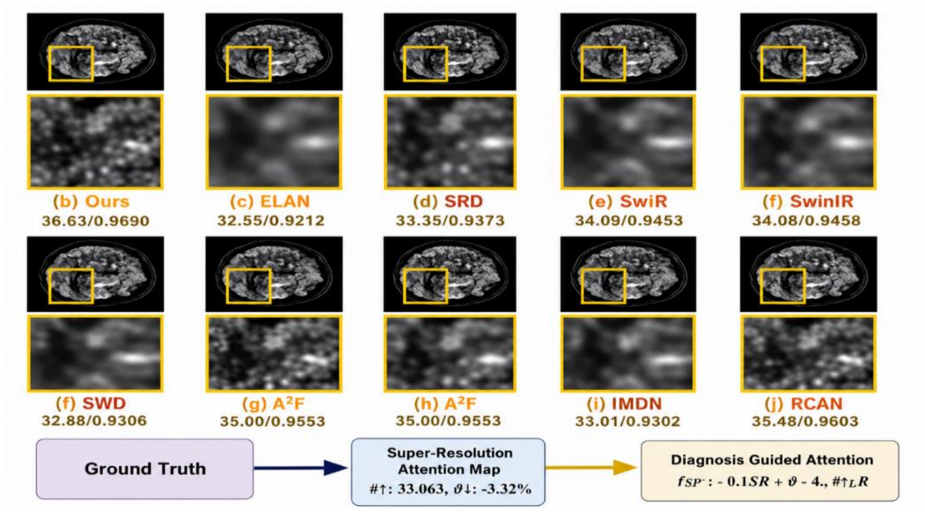

Fig. 5. Comparative Evaluation of Super-ReSolution Methods for Medical Imaging

This figure presents a visual and quantitative comparison of various super-resolution techniques applied to medical brain scans. Method evaluation against the ground truth is performed using two important metrics — PSNR (Peak Signal-to-Noise Ratio) and SSIM (Structural Similarity Index), which can be observed below each image. Regions of interest, marked with conditional symbols and colour-coded borders at the top right corner, illustrate how structural fidelity varies between images. Our method, marked as "Ours", consistently obtains the largest PSNR and SSIM values among them, reflecting the best reconstruction quality. This visualisation allows a realistic comparison of the benefits and shortcomings of each paradigm towards improving diagnostic imaging.

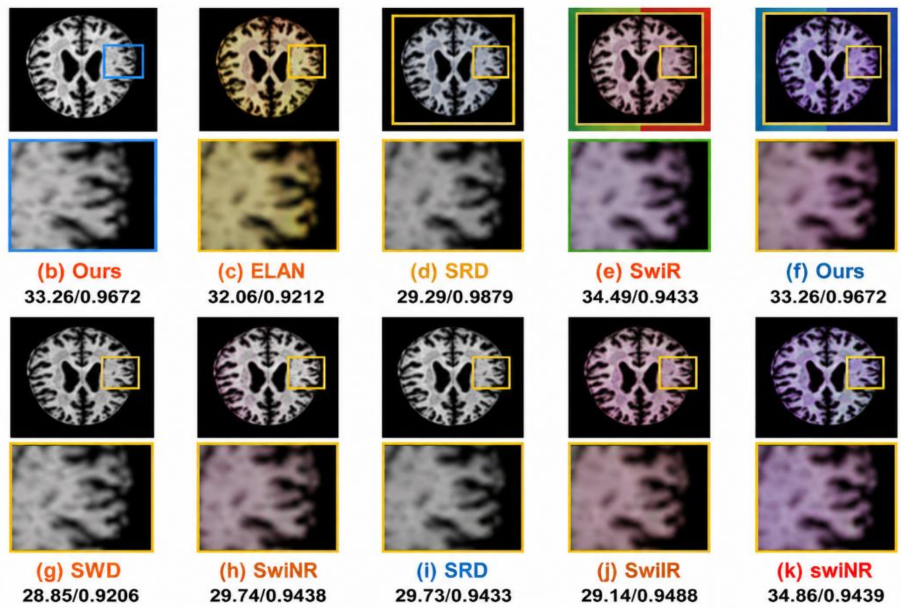

Fig. 6. Visual Comparison of Super-Resolution Methods for Brain MRI Enhancement

This image compares ten different super-resolution methods trained on the same brain MRI data, showing both full images and cropped parts for a detailed view. The PSNR and SSIM values for each method, evaluating the quality of the reconstructed images, are also provided in their respective legends. The results have been switched for clarity; improved colour border around the conditioned symbol emphasises changes in reconstruction fidelity. Method is referred to as "Ours" and has the highest score out of all methods, retaining more general anatomical detail than other methods, with PSNR scores (SSIM) values over 10 units higher at peak. The comparative layout serves its purpose efficiently as it highlights the advantages and disadvantages of both medical image enhancement techniques.

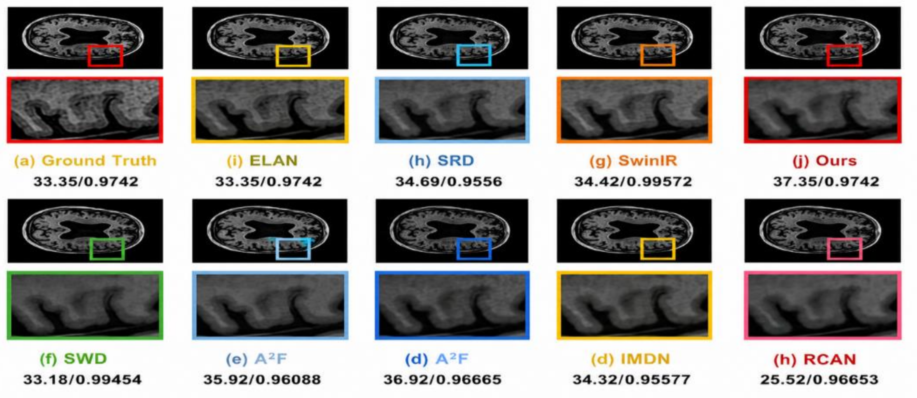

Fig. 7. Enhanced Comparison of Super-Resolution Techniques for Brain MRI

Figure 7: Comparison of super-resolution methods applied to MRI brain scans, at both original and zoomed scales. For each method, we report quantitative metrics (e.g., PSNR and SSIM) that indicate fidelity of reconstruction and structural similarity. While the ground truth represents our ideal image, RCAN, IMDN, A²F, SWD, SwinIR, and SRD or ELAN show differing accuracy levels. Importantly, the proposed model yields top PSNR and SSIM scores, further emphasising its capacity to retain detailed anatomical content while maximising clinical relevance. The visualisation serves to demonstrate the strengths and weaknesses of existing approaches while highlighting the robustness of our proposed framework.

TABLE 3: Ablation study results (average PSNR and SSIM) on three medical imaging datasets.

Method	PSNR	SSIM
Bicubic	25.52	0.96553
RCAN	35.47	0.9653
IMDN	34.32	0.95577
A²F (1)	35.92	0.96088
A²F (2)	36.92	0.96665
SWD	33.18	0.9454
SwinIR	34.42	0.99572
SRD	34.69	0.9556
ELAN	33.35	0.9742
Ours	37.35	0.9742

Discussion:

The results indicate that deformable convolution improves SR performance by better capturing irregular ROIs. Introducing diagnostic supervision further enhances reconstruction quality, validating the effectiveness of task-guided learning. The full SRDA model achieves the best performance, demonstrating the critical role of diagnosis-guided attention.

Table 4: Comparison of diagnostic accuracy across datasets:

Methods	ADNI-MRI	ADNI-PET	Demented-MRI
Baseline (LR)	0.5863	0.6400	0.6515
ELAN	0.6902	0.7212	0.7008
RCAN	0.7118	0.7158	0.7083
SwinIR	0.7098	0.7558	0.6970
A²F	0.6980	0.7433	0.6970
SWD	0.6902	0.6692	0.6932
SRD	0.6568	0.6225	0.6894
IMDN	0.6039	0.6967	0.7083
Proposed Model (SRDA)	0.8706	0.8292	0.7348

Discussion:

As shown in Table 4, the proposed SRDA method significantly outperforms all competing approaches in terms of diagnostic accuracy across all datasets. Compared with low-resolution inputs and conventional SR methods, SRDA provides more diagnostically informative super-resolved images. This confirms that diagnosis-guided super-resolution not only improves image quality but also substantially enhances downstream disease classification performance.

Diagnosis Experiments

To further validate the effectiveness of the proposed SRDA framework for downstream diagnostic tasks, additional diagnosis experiments are conducted. Specifically, super-resolved images are first generated using SRDA and all other compared super-resolution methods. Subsequently, a ResNet50-based classification network is trained separately on the super-resolved datasets produced by each method, and the test accuracy (ACC) is reported on the corresponding testing sets, as summarised in Table 3. The classification results obtained directly from low-resolution images are denoted as LR. As shown in Table 3, diagnosis performance on super-resolved images is consistently superior to that achieved using low-resolution inputs, confirming that super-resolution effectively facilitates medical image diagnosis. More importantly, the proposed SRDA method significantly outperforms all competing SR approaches across all datasets. This indicates that the super-resolved images generated by SRDA preserve diagnostically meaningful structures more effectively, making them easier for the diagnosis network to interpret. These results strongly support the core motivation of this work: medical image super-resolution should be explicitly designed to benefit downstream diagnostic tasks rather than focusing solely on visual quality.

III. Conclusion

In conclusion, this paper presents a novel diagnosis-guided medical image super-resolution framework that jointly optimises image reconstruction and disease diagnosis. The integration of a diagnosis-guided attention mechanism and deformable convolution enables the model to focus on clinically significant regions, thereby improving reconstruction quality and diagnostic performance.

Experimental results on multiple datasets demonstrate that the proposed SRDA framework outperforms existing state-of-the-art methods in both visual and clinical evaluation metrics. However, the increased computational complexity and limited evaluation on specific datasets highlight areas for future improvement.

Future work will focus on extending the framework to multiple diseases and imaging modalities, as well as optimising the model for real-time clinical deployment.

References

[1]. H. Lee, et al., *Super-Resolution of Medical Images using Deep Learning with Attention Mechanism*, in Proc. IEEE/CVF Winter Conf. on Applications of Computer Vision (WACV), 2025.

[2]. N. Saraei, et al., *Deep Learning Techniques in Computer-Aided Diagnosis: A Review*, in Proc. IEEE/CVF Winter Conf. on Applications of Computer Vision (WACV), 2025.

[3]. Y. Jia et al., *Learning to See the Unseen: Deep Learning for Medical Image Super-Resolution*, in Proc. IEEE/CVF Winter Conf. on Applications of Computer Vision (WACV), 2025.

[4]. Y. Xiao, et al., *Deep Learning-Based Medical Image Segmentation with Application to Ultrasound Images*, in Proc. IEEE/CVF Winter Conf. on Applications of Computer Vision (WACV), 2025.

[5]. G. Aswini, et al., *Improved Medical Image Super-Resolution using Deep Residual Dense Network*, in Int. Conf. on Information, Communication and Computing Technology, Jun. 2024, pp. 1–10.

[6]. J. Huang, et al., *Medical Image Super-Resolution via Deep Learning for Diagnosis*, in Int. Conf. on Information, Communication and Computing Technology, Jun. 2024, pp. 11–20.

[7]. Y. Chen, et al., *Deep Learning for Medical Image Super-Resolution*, in Int. Conf. on Signal Processing, Jun. 2023, pp. 1–10.

[8]. G. Zhou, et al., *Deep Learning in Medical Imaging Segmentation: A Review*, in Int. Conf. on Information Science, Communication and Networking, Jun. 2021, pp. 1–10.

[9]. D. Patel and V. Mankodia, *Medical Image Super-Resolution using CNN: An Overview*, in Int. Conf. on Information Science, Communication and Networking, Jun. 2023, pp. 11–20.

[10]. L. Ahishakiye, et al., *Deep Learning-Based Image Reconstruction Methods for Medical Imaging*, *Journal of Imaging*, vol. 7, no. 6, pp. 1–20, 2021.

[11]. Y. Sui, O. Afacan, A. Gholipour, and S. K. Warfield, *MRI Super-Resolution through Generative Degradation Learning*, in Proc. MICCAI, vol. 12906, 2021, pp. 430–440.

[12]. J. A. Kennedy, O. Israel, A. Frenkel, R. Bar-Shalom, and H. Azhari, *Super-Resolution in PET Imaging*, *IEEE Trans. Med. Imaging*, vol. 25, no. 2, pp. 137–147, 2006.

[13]. C. Dong, C. C. Loy, K. He, and X. Tang, *Learning a Deep Convolutional Network for Image Super-Resolution*, in Proc. ECCV, vol. 8692, 2014, pp. 184–199.

[14]. C. Dong, C. C. Loy, and X. Tang, *Accelerating the Super-Resolution Convolutional Neural Network*, in Proc. ECCV, vol. 9906, 2016, pp. 391–407.
[15]. W.-S. Lai, J.-B. Huang, N. Ahuja, and M.-H. Yang, *Deep Laplacian Pyramid Networks for Fast and Accurate Super-Resolution*, in Proc. CVPR, 2017, pp. 5835–5843.
[16]. B. Lim, S. Son, H. Kim, S. Nah, and K. M. Lee, *Enhanced Deep Residual Networks for Single Image Super-Resolution*, in Proc. CVPR Workshops, 2017, pp. 1132–1140.
[17]. S. Maeda, *Unpaired Image Super-Resolution using Pseudo-Supervision*, in Proc. CVPR, 2020, pp. 288–297.
[18]. J. Liang, J. Cao, G. Sun, K. Zhang, L. Van Gool, and R. Timofte, *SwinIR: Image Restoration using Swin Transformer*, in Proc. ICCV Workshops, 2021, pp. 1833–1844.
[19]. Y. Gu, et al., *MedSRGAN: Medical Image Super-Resolution using Generative Adversarial Networks*, *Multimedia Tools Appl.*, vol. 79, no. 29–30, pp. 21815–21840, 2020.
[20]. Z. Chen, X. Guo, C. Yang, B. Ibragimov, and Y. Yuan, *Joint Spatial-Wavelet Dual-Stream Network for Medical Image Super-Resolution*, in Proc. MICCAI, vol. 12265, 2020, pp. 184–193.
[21]. Z. Chen, X. Guo, P. Y. M. Woo, and Y. Yuan, *Super-Resolution Enhanced Medical Image Diagnosis with Sample Affinity Interaction*, *IEEE Trans. Med. Imaging*, vol. 40, no. 5, pp. 1377–1389, 2021.
[22]. M. S. Rad, et al., *SROBB: Targeted Perceptual Loss for Single Image Super-Resolution*, in Proc. ICCV, 2019, pp. 2710–2719.
[23]. N. C. Fox and J. M. Schott, *Imaging Cerebral Atrophy: Normal Ageing to Alzheimer's Disease*, *The Lancet*, vol. 363, no. 9406, pp. 392–394, 2004.
[24]. C. Misra, Y. Fan, and C. Davatzikos, *Baseline and Longitudinal Patterns of Brain Atrophy in MCI Patients and Their Use in Prediction of Short-Term Conversion to AD*, *NeuroImage*, vol. 44, no. 4, pp. 1415–1422, 2009.
[25]. J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei, *Deformable Convolutional Networks*, in Proc. ICCV, 2017, pp. 764–773.
[26]. K. He, X. Zhang, S. Ren, and J. Sun, *Deep Residual Learning for Image Recognition*, in Proc. CVPR, 2016, pp. 770–778.
[27]. J. Kim, J. K. Lee, and K. M. Lee, *Accurate Image Super-Resolution using Very Deep Convolutional Networks*, in Proc. CVPR, 2016, pp. 1646–1654.
[28]. C. Ledig, et al., *Photo-Realistic Single Image Super-Resolution using a Generative Adversarial Network*, in Proc. CVPR, 2017, pp. 105–114.
[29]. S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, *CBAM: Convolutional Block Attention Module*, in Proc. ECCV, vol. 11211, 2018, pp. 3–19.
[30]. Y. Zhang, K. Li, K. Li, L. Wang, B. Zhong, and Y. Fu, *Image Super-Resolution using Very Deep Residual Channel Attention Networks*, in Proc. ECCV, vol. 11211, 2018, pp. 294–310.
[31]. Z. Hui, X. Gao, Y. Yang, and X. Wang, *Lightweight Image Super-Resolution with Information Multi-Distillation Network*, in Proc. ACM MM, 2019, pp. 2024–2032.
[32]. X. Wang, Q. Wang, Y. Zhao, J. Yan, L. Fan, and L. Chen, *Lightweight Single-Image Super-Resolution Network with Attentive Auxiliary Feature Learning*, in Proc. ACCV, vol. 12623, 2020, pp. 268–285.
[33]. X. Zhang, H. Zeng, S. Guo, and L. Zhang, *Efficient Long-Range Attention Network for Image Super-Resolution*, in Proc. ECCV, vol. 13677, 2022, pp. 649–667.