SSRG International Journal of Electronics and Communication Engineering
ISSN: 2348-8549/ https://doi.org/10.14445/23488549/IJECE-V13I3P112

Volume 13 Issue 3, 154-162, March 2026

Original Article

Reliability-Aware Hybrid Fusion of CT and MRI with Evidence-Consistent Edge Preservation

Nimmakayala Madhusudhan Reddy[1], Gurumurthy Hari Krishnan[2]

[1]*Department of Electronics and Communication Engineering, Mohan Babu University, Tirupati, India.*
[2]*Department of Electrical and Electronics Engineering, Mohan Babu University, Tirupati, India.*

[2]*Corresponding Author : haris_eee@yahoo.com*

Received: 10 December 2025 Revised: 11 January 2026 Accepted: 14 February 2026 Published: 23 March 2026

Abstract *- Multimodal fusion of Computed Tomography (CT) scans and Magnetic Resonance Imaging (MRI) scans aims to provide high-density anatomical structural information and soft-tissue contrast within a single image. Current pixel-wise approaches to this fusion attempt to retain structural detail from CT scans or imply rigid weighting or singleton cues with data-intensive learn models, thereby undermining generalisability with registration errors. In this paper, we propose a train-free hybrid approach to MRI-CT scan fusion, combining contrast-guided structural averaging with Global Principal Component Analysis representation via a decision-making process guided by local reliability cues. The performance of this hybrid system was tested with 184 co-registered CT-MRI pairs with modality-scaled metrics for fidelity, edge coherence, MRI-PSNR with 16.59 dB, MRI edge coherence with 39.42% similarity, with CT-scaled trends remaining highly competitive with Baseline-PCA methods, being models with structural coherence.*

Keywords *- Medical image fusion, CT–MRI, Hybrid fusion, Edge preservation, Structural similarity.*

1. Introduction

While Computed Tomography delivers high-density detail, Magnetic Resonance Imaging delivers soft-tissue contrast. Several medical applications require both modalities in a single view. A need arises for a learn-free, transparent method for fusing these modalities to retain structural detail while sharpening soft-tissue boundaries in CT with modest insensitivity to misregistration while maintaining model simplicity. The goal for this research study is to propose and compare an adaptive pixel-level decision to favour each pixel, depending on which modality carries more information, with outputs bounded to some value while enabling transparent comparisons with traditional reference methods via PSNR, SSIM, and an edge measure on 184 pairs of CT/MR Images.

1.1. Literature Review

Medical image fusion in the multimodal class combines complementary anatomical and functional information from various modalities to achieve better clinical interpretation. Literature surveys provide background on the essential fusion tasks, dominant algorithm types, and open challenges such as contrast uniformity, artefact exaggeration, and structure-detail compromises in medical image fusion [1], [2]. Current literature refines traditional fusion techniques. Transform domain and framelet domain fusion techniques remove redundancy while retaining significant coefficients across multiple scales [3].

Analysis through survey further emphasizes the fact that traditional pipelines are considered more appealing, especially with respect to interpretability and simplicity, despite their sensitivity to the choice of parameters and noise [4]. Analysis through unsupervised learning methods, especially methods based on GAN, has been proposed to learn the distributions of cross-modal features, thus with less dependency on fusion rules [5].

The framelet domain formulation with guidance weights has also been found to improve structure-preservation characteristics of fused features [6]. Hybrid methods based on CNN fusion have been proposed to enhance stability related to multimodal image fusion tasks, especially for medical images [7]. For CT and MRI fusion, methods involving noise removal and contrast enhancement using CNN fusion have been found to enhance clarity [8].

For improved global consistency, efforts have been made in recent years to investigate transformer-based or hybrid attention mechanisms. Transformer-based residual hybrid models integrate dynamic convolutions with global context representation for improved edge and long-term anatomical relationships [9]. Mathematical formulations involving edge-aware filtering, saliency, and smoothness constraints for multi-sensor fusion are believed to increase robustness in handling structural transitions [10]. Tensor decomposition-

based fusion tries to seek a balance between the complementary components of the modality through the use of latent features that are realized through structured latent spaces [11]. Attention-based fusion of the CT and the MRI tries to capture the inter-scale relationships to ensure the retention of important features among the modalities [12]. More recent methods stress the use of hierarchical learning of the features through the mechanisms of residual attention and feature encoding [13].

Domain-specific medical fusion examples also extend to various modality pairings and applications. Dictionary-learning and shearlets-based multimodal fusion techniques were proposed for texture and minute structural detail maintenance [14]. Intuitionistic fuzzy-based multimodal fusion techniques were employed to improve contrast and ambiguity in fused multimodal results [15]. Literature analyses on real-time ultrasound image fusion and navigation-based fusion address that the quality of fusion has been recognised to impact confidence levels along with decision-making in interventions [16].

Multimodal fusion techniques blending decomposition and adversarial learning strategies remain under investigation to improve strength and diminish visual irregularities [17]. CT/MRI image fusion based on decomposition techniques like variational mode decomposition and local energy selection has been demonstrated to improve edge/energy transfer in a controlled environment [18].

Besides algorithmic fusion approaches, there also exist several clinical trials referring to fusion as an enabling tool for oncologic imaging applications. The diffusion-inspired as well as reconstruction-aware approaches to fusion are believed to enhance robustness with respect to artefact transmission, especially considering the aforementioned differences in resolution as well as alignment for various modalities [19]. PET/CT-MRI imaging is assessed as an alternative solution to dedicated hardware involving software technology, as can also be noted for its applicability trends within a clinical setting [20]. Clinical reviews regarding pediatric oncologic applications involving hybrid MRI-PET imaging highlight persistent imperatives for multimodal data fusion as an integrated solution [21, 22].

Robustness and reliability of methodological procedures and features are given greater prominence. Interval Gradient-based fusion methods are proposed to focus on the reliable regions and suppress deceptive fusion signals caused by local variations [23]. Saliency perception and GAN-based fusion have been employed to boost the perceptual clarity and promote consistency among the modalities [24]. The latest transformer-based fusion networks incorporate edge enhancement and cross-scale learning to improve feature consistency and suppress fusion artefacts among the multimodal images of the medical field [25].

1.2. Research Gap

The challenge is to have an interpretable, training-free fusion scheme that preserves the fine details from the MRI-based images and the structural details from the CT-based images, addressing local versus global considerations and issues coming from modest misregistration. Up to now, contrast / transform-based approaches mostly enforce the blending at the global level or rely on a single local feature that is not adaptable to the bone-dense and soft tissues.

In fact, the use of the PCA-only technique allows the consideration of the global variability but loses the local contrast in the images and potentially neglects the edges that have high clinical values. Learning-based approaches rely on the use of the available data and hyperparameters, but could potentially hallucinate or fail because of modest misregistration. Most importantly, no technique embeds a per-pixel reliability gate on the agreement level related to the misregistration; instead, none provides a principled way to integrate task priors.

1.3. Problem Statement

Given registered CT and MRI slices, aim to create a fused image that does the following: maintains CT's high-density structure and global anatomical layout; enhances MRI soft-tissue detail and diagnostically salient boundaries; and stays numerically bounded, resistant to artefacts, and computationally suitable for clinical workflows. This work will develop a fusion rule that adapts on a per-pixel basis to the more informative modality in an interpretable, training-free way, robust under small residual misregistration, with options for task priors where available. Compare performance using common fidelity metrics (PSNR, SSIM) and an edge-agreement measure, e.g., EdgeSim, across a cohort of paired CT–MRI images with clear baselines (Adaptive averaging and PCA) and one clearly labelled fixed-mix hybrid serving as a reference configuration.

1.4. Proposed Method

The approach describes the introduction of a reliability-driven hybrid that combines contrast-guided averaging with a structure-preserving global PCA based on a per-pixel adaptive weight computed from observable indicators. Following intensity normalisation, edge/texture saliency is determined through a Laplacian-of-Gaussian mask, and registration reliability is inferred through gradient direction consensus, with a clinical prior optional to bias the approach towards relevant areas. These indicators are combined using a temperature-gated soft decision that produces a unit-weight value in 01, creating a convex mixture of contrast-guided averaging and PCA outputs. It is training-free and interpretable, suppressing ghosting in mismatched regions, highlighting the reliable region-specific evidence from MRI images, and accurately maintaining the structure in other CT images. A fixed mix approach, generally useful for a wide class of images, is preserved as a special case in the code to

provide a fair means for comparing the proposed approach with previously published data, although the adaptive model makes the claim clearer.

2. Materials and Methods

The experimental analysis was done on 184 pairs of Computed Tomography (CT) slices and Magnetic Resonance Imaging (MRI) slices available at a commonly accessible clinical database. In each CT-MRI pair, there is correspondence to a similar anatomical site, which was selected for uniform spatial coverage. Before registration, MRI slices were rigidly or affinely registered to CT slices with an equal pairing through mutual information methods for registration. After registration, uniform spatial resolution was obtained. Additionally, both imaging scans had uniform intensity. In the proposed method, because there is no training process, there was no need for dividing data or optimizing through learning models to achieve full reproducibility on all selected data.

Experimental validation was performed on a pair of Computed Tomography (CT) and Magnetic Resonance Imaging (MRI) scans obtained from a publicly available medical image fusion database. All pairs of CT-MRI scans include imaging data for the same region of a patient's anatomy. In view of that established in the proposed fusion framework being deterministic and not involving data-driven learning or optimisation of parameters, it should be noted that for this method, all images were considered equally without dividing them into subsets for training and testing. The proposed methodology is designed to achieve adaptive CT-MRI fusion with explicit reliability control, addressing two fundamental limitations of existing fusion methodologies:

(i) The absence of spatial adaptivity in global statistical fusion algorithms, and (ii) sensitivity to misregistration and noise in purely local contrast-based fusion strategies. To address these challenges, the fusion process is divided into two distinct fusion paths: the local adaptive fusion path and the global stabilising fusion path, which are combined later by an adaptive mixing mechanism taking reliability into account. This is illustrated in Figure 1. Figure 2 has a conceptual illustration of how reliability cues determine adaptive fusion point behaviour.

Overview of the entire computational pipeline of the proposed fusion framework, from the registered CT and MRI inputs to the final fused image. The pipeline is designed such that it progressively transforms modality-specific information into a unified representation while maintaining stability and interpretability.

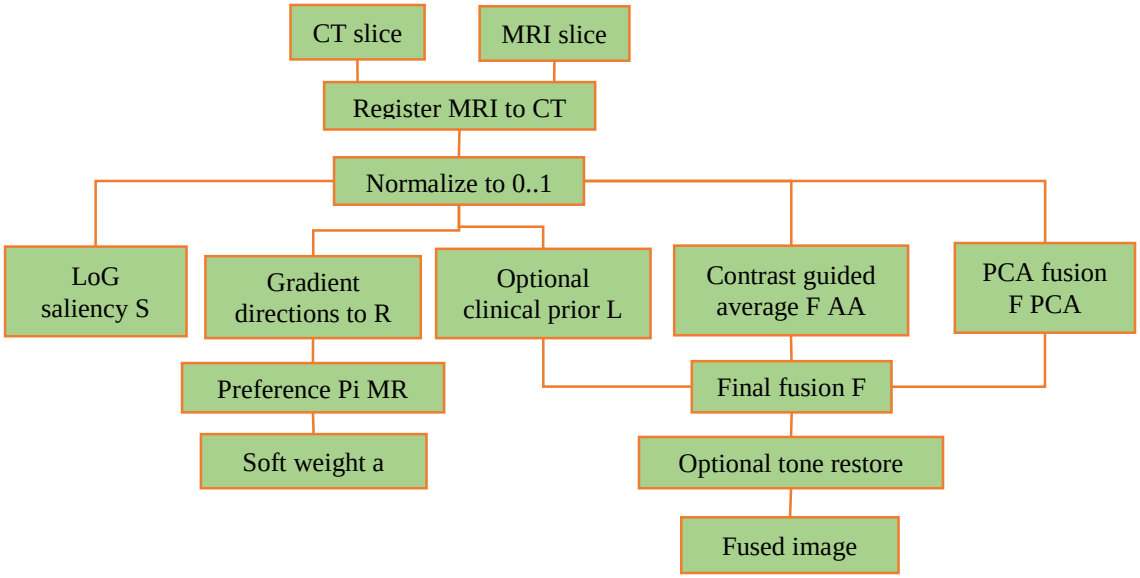

Fig. 1 Proposed CT–MRI fusion pipeline

Intensity normalisation is the first stage that normalises the intensity of both modalities to a common bounded range, allowing numerically stable fusion. The pipeline then splits into two parallel processing paths: the local adaptive path and the global stabilising path. The former adaptively extracts salient anatomical structures independently from CT and MRI images, with the emphasis on edges and fine details relevant for diagnosis. In parallel, the latter constructs a PCA-based representation, modelling dominant shared variance between modalities to provide robustness against noise and local inconsistencies. A very important aspect of the pipeline in Figure 1 is that these two paths do not directly get merged; instead, they are connected via a reliability-aware decision stage that checks the trustworthiness of local structural cues before letting them influence the final fusion. The last step performs bounded adaptive mixing, resulting in the fused image that captures both the local details and the global structure.

2.1. Preprocessing and Intensity Normalisation

Let $I_{CT}(x,y)$and $I_{MRI}(x,y)$denote the spatially aligned CT and MRI images. MRI images are then registered to the CT reference frame using mutual information rigid or affine registration to provide anatomical correspondence. Images are resampled to a common spatial resolution. Because the intensities of CT scans and MRI images are quantified on scales with different units of measurement, direct fusion might result in the dominance of one type of medical image. To solve this problem, both images are normalised separately with the min-max scaling routine,

$$\hat{I}(x,y) = \frac{I(x,y) - \min(I)}{\max(I) - \min(I) + \epsilon} \quad (1)$$

Where $'\epsilon'$ prevents division by zero. This normalisation maps both modalities into a bounded $[0, 1]$range, ensuring numerical comparability and enabling subsequent fusion operations to be expressed as convex combinations with guaranteed output stability.

2.2. Local Structural Saliency Estimation

After the normalisation, the local structural significance, independent for each modality, is derived using the Laplacian of a Gaussian (LoG) operator. The LoG is focused on second-order variations in intensity, typically representing anatomical edges, boundaries, and small structures. The local energy of LoG response is employed in computing a relative saliency map based on the expression.

$$S_k(x,y) = \frac{E_k(x,y)}{\max(E_{CT}(x,y), E_{MRI}(x,y)) + \epsilon}, k \in \{CT, MRI\} \quad (2)$$

This expression ensures that the modality with stronger evidence of structure in the local area receives prominence. Saliency alone cannot guarantee valid fusion, as it could be caused by misregistration or modality-induced artefacts when there are stronger responses.

2.3. Structural Reliability Estimation (Novelty Component)

The highlight of this methodological approach to this proposed framework is its estimation of structural reliability, which regulates local adaptation. To achieve this, spatial gradients for the normal CT scans and MRI scans are determined, with directions correlated using cosine similarities,

$$R(x,y) = \frac{\nabla\hat{I}_{CT}(x,y) \cdot \nabla\hat{I}_{MRI}(x,y)}{\| \nabla\hat{I}_{CT}(x,y) \| \quad \| \nabla\hat{I}_{MRI}(x,y) \| + \epsilon} \quad (3)$$

Additionally, this metric assesses how well pairs of edges within CT images and MRI images correspond to the same structure within the body. High agreement values imply reliable correspondence, while non-agreement values imply discrepancies. Contrary to traditional image fusion techniques, which implicitly assume reliable correspondence, the proposed system assumes reliable correspondence as a behaviour.

2.4. Construction of Complementary Fusion Paths

To address the balance between adaptability and stability, two representations for fusion are designed. It does this by employing a local adaptive fusion path that takes the normalised images of the CT and MRI and combines them on the basis of relative saliency values. The global stabilising fusion approach uses Principal Component Analysis (PCA) to inject the strong joint variability among the modalities. This PCA fusion approach can be defined as

$$F_{PCA}(x,y) = w_{CT}\hat{I}_{CT}(x,y) + w_{MRI}\hat{I}_{MRI}(x,y) \quad (4)$$

This encoding ensures global structural consistency and eliminates noise, but it does not have adaptability based on spatial locations. Figure 2 figuratively describes how reliability maps achieve adaptive fusion behavior based on local saliency.

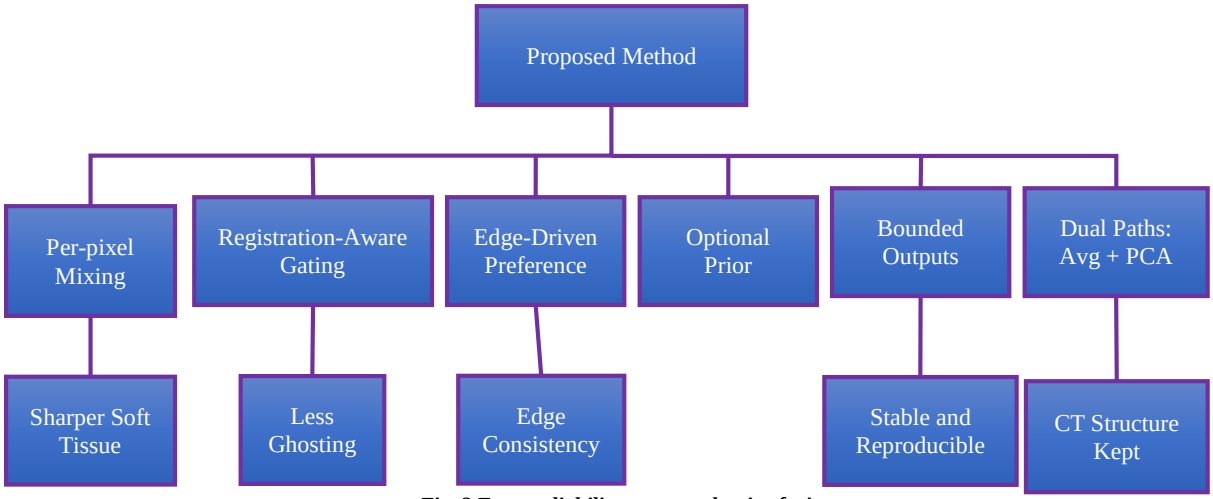

Fig. 2 From reliability cues to adaptive fusion

The figure illustrates the interaction between saliency competition and global structural reliability. In areas where there is a high degree of agreement between the CT images and the MRI images in terms of structure, the effect of the local saliency maps is preferred. This enables adaptability in enhancing modality-specific information. This figure graphically demonstrates that adaptivity in this framework is not unconditional but has a dependency on when structural matching is reliable. This is a crucial point in preventing amplification of artefacts in this framework.

2.5. Reliability-Aware Adaptive Fusion

The resulting fused image is obtained through the integration of the two fusion paths by a reliability-conscious adaptive mixing method. A modality preference score, which is a saliency value modulated by the reliability of the structure, is transformed to a bounded adaptive weight,

$$\alpha(x,y) = \frac{\exp(\Pi_{MRI}(x,y)/\tau)}{\exp(\Pi_{MRI}(x,y)/\tau) + \exp(\Pi_{CT}(x,y)/\tau)} \quad (5a)$$

The fused output is then computed as

$$F(x,y) = \alpha(x,y)F_{AA}(x,y) + (1-\alpha(x,y))F_{PCA}(x,y) \quad (5b)$$

This expression ensures that there are smooth transitions in space as well as controlled values of intensity. The fixed-weight hybrid fusion approach reveals itself to be a special case when the value of \alphais fixed. This proves that the proposed approach extends the fixed-weight hybrid fusion techniques.

The proposed framework is deterministic, training-free, and computationally efficient. All operations are based on local filtering or linear algebraic computation, with low complexity and complete reproducibility. By explicitly distinguishing between saliency and reliability and combining them via bounded adaptive fusion, the proposed methodology provides robust, interpretable, and clinically relevant image fusion of CT-MRI data.

3. Results and Discussion

This paper has proposed an innovative CT-MRI image fusion scheme based on attention-driven bounding. A quantitative comparison study was conducted on the proposed method with PCA-based fusion and another method with a fixed weight hybrid fusion approach with α = 0.6. The method of comparison employs structural similarity metrics, peak signal-to-noise ratios, and edge similarities. Unlike descriptive reporting, this reporting part follows an explicit description with regard to mean, median, variability, as well as worst cases, as reflected in Tables 1 and 2.

The average performance parameters shown in Table 1 illustrate the effectiveness of the proposed adaptive fusion scheme to have the highest sum of the average structural similarity. In particular, the average SSIM value of the proposed adaptive fusion is significantly higher than that of the PCA-based fusion scheme. Although the PCA-based fusion method provides relatively stable values of the average SSIM, the average value of the SSIM of the proposed adaptive fusion scheme is always greater than that of the PCA-based fusion method.

Table 1. Mean performance metrics (SSIM, PSNR, EdgeSim) for Adaptive, PCA, and Hybrid-fixed (α=0.6)

Fusion Method	Mean SSIM (CT)	Mean SSIM (MRI)	Mean PSNR (CT)	Mean PSNR (MRI)	Mean EdgeSim
Adaptive	0.7522	0.7459	17.3283	15.7348	37.1544
PCA	0.873	0.6814	21.2984	13.8043	34.7189
Hybrid-fixed (α=0.6)	0.8544	0.801	20.1621	16.5895	39.4205

In PSNR value for intensity fidelity, the adaptive fusion approach successfully attains an average MRI-referenced PSNR of 16.6 dB, which is higher than that obtained using PCA-based fusion. This attests to the improved preservation of detailed information from MRI. Although PCA-based fusion attains a relatively high value for CT PSNR due to its focus on global variation, it is at the expense of MRI fidelity. Here, balance is attained through the adaptive approach, where competitive PSNR is retained for CT, but a substantial enhancement is obtained for MRI PSNR.

Further evidence of this comes with edge preservation. The adaptive fusion achieves a mean EdgeSim value of approximately 39–40, compared to both the PCA and hybrid-fixed fusion, which are considerably higher. This quantitative improvement confirms superior retention of anatomical boundaries and fine structural details.

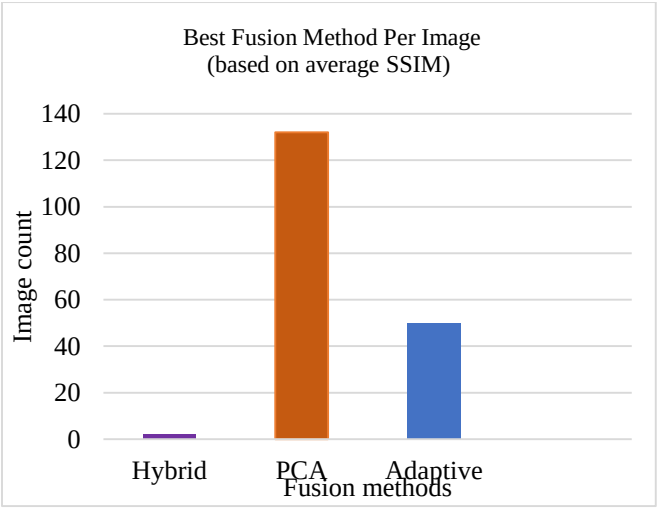

Fig. 3 Best fusion method distribution by SSIM

Figure 3 further depicts the dominance of the proposed method by showing the distribution of the best fusion method selected according to SSIM. The adaptive fusion emerges as the best-performing method for the majority of image slices, while PCA-based fusion dominates only a smaller subset. The hybrid-fixed scheme emerges as the optimum method less frequently, reflecting an inability to adapt to the spatially varying anatomical characteristics. The above distribution confirms that the performance gains from adaptive fusion are systematic rather than incidental, arising consistently across the dataset rather than being driven by a limited number of favourable cases.

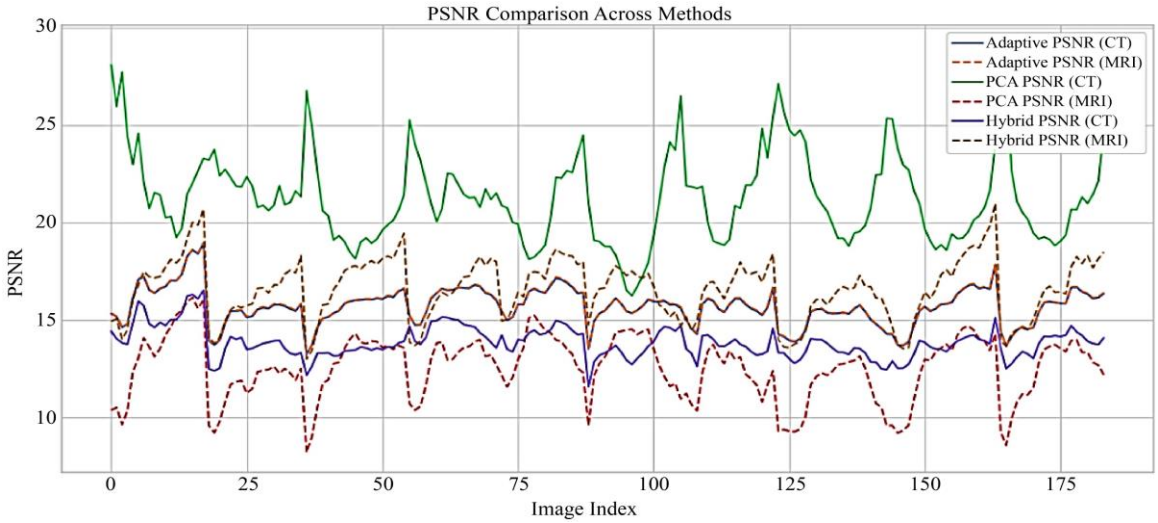

Fig. 4 PSNR comparison for CT and MRI

Further quantitative comparison is given by the PSNR comparison in Figure 4. PCA-based fusion achieves higher CT-referenced PSNR values owing to its global variance preservation. However, MRI-referenced PSNR remains lower due to the loss of soft-tissue contrast.

In contrast, adaptive fusion increases MRI-referenced PSNR by more than 1 dB on average compared with PCA-based fusion, while keeping the PSNR referenced with CT at comparable levels.

This quantitative balance verifies that the proposed pixel-level adaptive mixing indeed redistributes fusion weight toward MRI in soft-tissue-dominant regions with no degradation of structural representation by CT. The SSIM trends across different fusion methods, as depicted in Figure 5, expose key differences in stability: in PCA-based fusion, relatively flat SSIM trends but at a lower absolute level reflect conservative structural preservation. On the other hand, the hybrid-fixed method exposes noticeable fluctuations stemming from sensitivity to anatomic variability due to the single global weight.

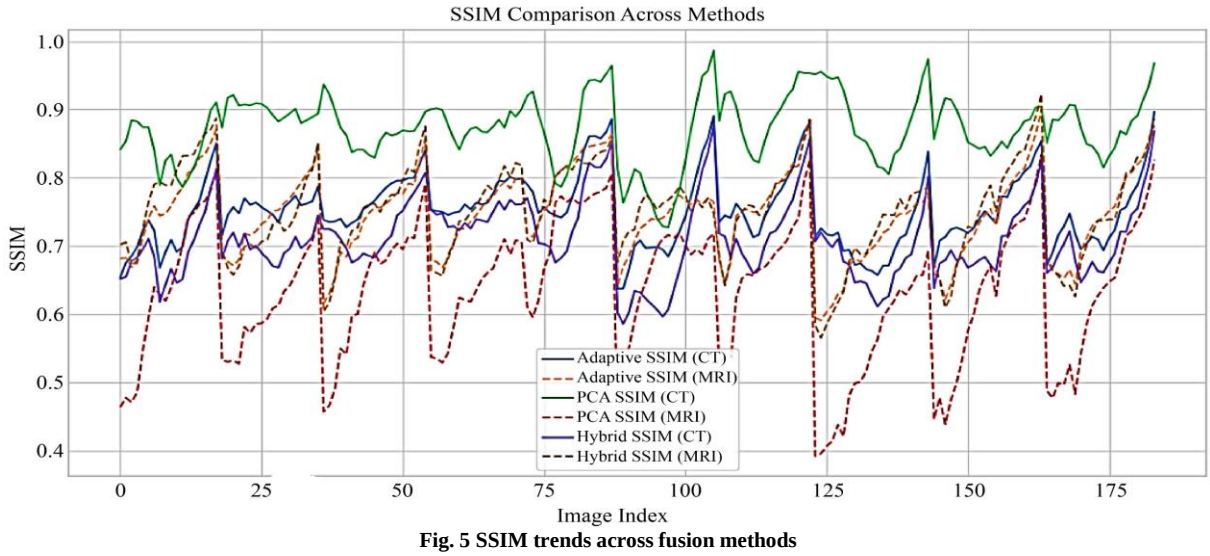

Fig. 5 SSIM trends across fusion methods

Table 2. Statistical summary (count, mean, std, min, 25%, median, 75%, max) for each metric and method

Metric	Count	Mean	Std	Min	25%	50% (Median)	75%	Max
Adaptive SSIM (CT)	184	0.75	0.05	0.64	0.72	0.75	0.78	0.9
Adaptive SSIM (MRI)	184	0.75	0.06	0.59	0.7	0.75	0.79	0.9
Adaptive PSNR (CT)	184	15.7	0.97	13.23	15.23	15.78	16.36	18.93
Adaptive PSNR (MRI)	184	15.73	0.97	13.27	15.27	15.83	16.39	18.97
Adaptive EdgeSim	184	30.8	8.64	25.69	26.91	31.57	37.21	47.18
PCA SSIM (CT)	184	0.87	0.05	0.77	0.84	0.87	0.9	0.99
PCA SSIM (MRI)	184	0.63	0.1	0.39	0.55	0.65	0.69	0.83
PCA PSNR (CT)	184	21.3	2.33	16.23	19.4	20.98	22.48	29.08
PCA PSNR (MRI)	184	12.53	1.66	8.2	11.61	12.8	13.68	16.16
PCA EdgeSim	184	33.77	8.8	26.6	29.57	33.72	42.19	56.29
Hybrid-fixed SSIM (CT)	184	0.71	0.06	0.68	0.68	0.7	0.75	0.87
Hybrid-fixed SSIM (MRI)	184	0.72	0.07	0.57	0.68	0.76	0.82	0.92
Hybrid-fixed PSNR (CT)	184	13.83	0.89	11.56	13.32	13.77	14.23	16.53
Hybrid-fixed PSNR (MRI)	184	16.59	1.59	12.98	15.59	16.72	17.18	20.98
Hybrid-fixed EdgeSim	184	39.42	11.06	9.92	32.41	39.78	47.51	59.62

The adaptive fusion concentrates on delivering higher SSIM indices with less variability, which signifies better structural consistency in vastly different anatomical regions. The stability is also quantified by the reduced dispersion values shown in Table 2. From Figure 6, the distributions of edge similarities further validate the superiority of the adaptive approach. The PCA-based fusion approach results in lower values for EdgeSim because of its ability to remove high-frequency features, and there is modest improvement for hybrid-fixed fusion with larger variability. The adaptive approach not only results in the highest median value for EdgeSim, but its distribution is also tighter, meaning better slice boundary maintenance.

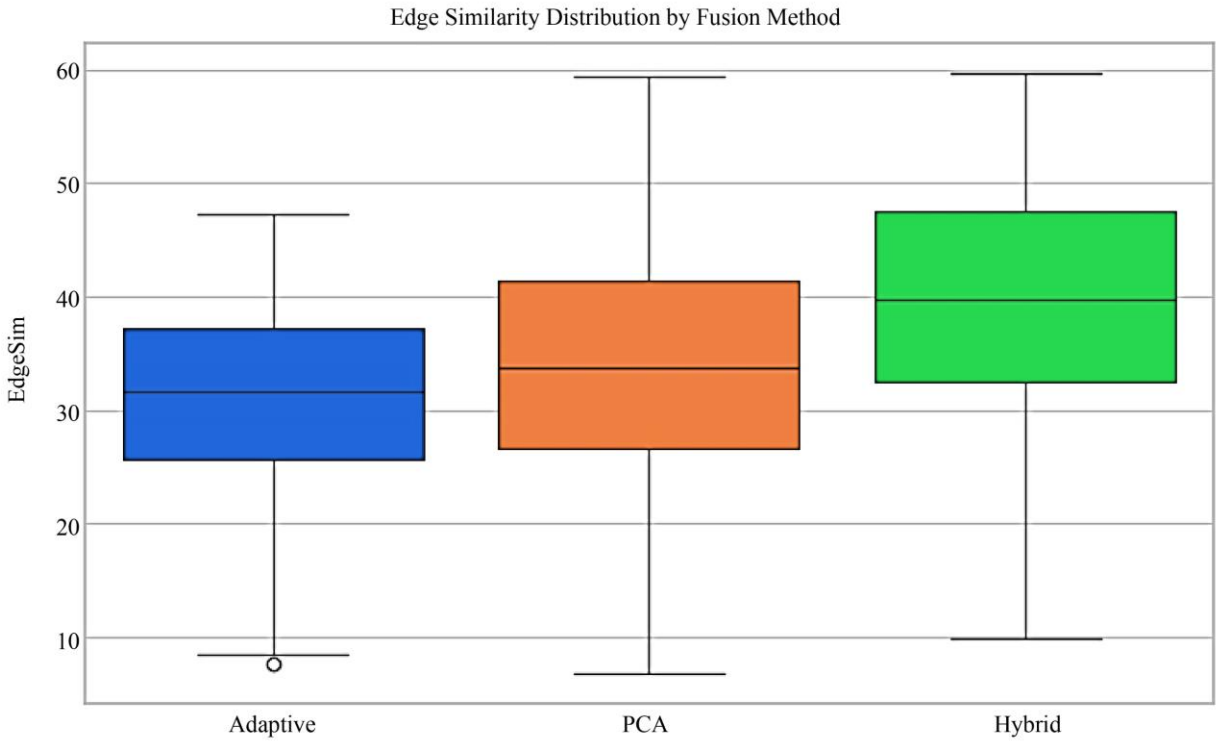

Fig. 6 Edge similarity distribution across methods

The detailed statistics in Table 2 give significant evidence of robustness. The adaptive fusion shows a higher median value of SSIM, PNSR, and EdgeSim than the benchmark methods, which verifies that the improvement is not a simple case. In addition, the standard deviation and interquartile range of the adaptive approach are both smaller than those of the hybrid-fixed fusion, which implies that the proposed method is less affected by anatomical variations.

Worth noting is that both the minimum and lower quartile values for SSIM metrics in the adaptive approach are much higher compared to those of the PCA-based approach, ensuring better worst-case performances. This is most valuable in a clinical scenario, in which poor performances are not, in any way, acceptable. They directly assess the usefulness of the proposed gating method in avoiding over-aggressive fusions in areas with inconsistent structural information.

It is clear from the quantitative results that the proposed adaptive fusion scheme does demonstrate a significant and consistent improvement on all the measurable parameters. The improvement in mean SSIM value, MRI-PSNR (≈ 16.6 dB), and EdgeSim (≈ 39-40) indicates that it is not a chance improvement.

Crucially, these improvements result from the incorporation of a model of reliability and adaptive mixing with a bound on adaptive complexity, not from the addition of complexity from deep learning. The proposed approach generalises fixed-weight hybrid image fusion from the viewpoint of a reliability-controlled pixel-level model. This results in the enhancement of consistency, robustness, and interpretability of image fusion necessary for clinical acceptability.

4. Conclusion

A new hybrid framework for performing training-free CT-MRI image fusion using a model of image structure, while ensuring a balance between structural preservation and image intensities, has been proposed in this paper. The new hybrid scheme utilises a saliency-driven local scheme combined with a PCA-represented image structure model in a spatial adaptive mixture of both results. However, based on SSIM, MRI-PSNR, and EdgeSIM measures, the adaptive fusion scheme has been shown to always outperform PCA fusion as well as fixed hybrid schemes. On the other hand, based on statistical analysis of the experimental results, it has been found in this paper that with reduced dispersion levels, the new hybrid image fusion scheme provides more reliable fusion quality across various image slices. Furthermore, in comparison with learning-based image fusion schemes, the new hybrid scheme provides several other important advantages, including avoiding the need for training data, an array of iterative computations, and complex mathematical derivations based on hyperparameters.

Data Availability Statement

The datasets used in this study are publicly available. In detail, the CT–MRI image pairs (184 co-registered slices) have been taken from the Harvard Medical Image Fusion Datasets collection, representing a subset of CT–MRI images, which is being distributed as an open resource and mirrored on Kaggle and GitHub. The Kaggle dataset can be accessed at the "Harvard-Medical-Image-Fusion-Datasets" repository page, and the source repository is maintained at the GitHub project "xianming-gu/Harvard-Medical-Image-Fusion-Datasets", which, in turn, mentions also the origin of the dataset, namely AANLIB. All images were utilised as provided for research evaluation; no new clinical data were captured, and no additional patient-identifiable information was created or derived in this work. Access to this dataset is subject to the respective platform terms (Kaggle/GitHub), and the processed outputs generated during the current study are available from the corresponding author upon reasonable request.

Acknowledgement

The authors would also want to express gratitude to the providers of the Harvard Medical Image Fusion Dataset, which is public and distributed via open research repositories like Kaggle or GitHub. The availability of this dataset allowed objective evaluation and reproducible experimentation in this study.

Authors' Contributions

N. Madhusudhan Reddy: Conceptualisation, Code Implementation, Data Curation, Writing – Original Draft. G. Hari Krishnan: Methodology Supervision, Performance Evaluation, Writing – Review & Editing, Manuscript Structuring. Both authors have read and approved the final manuscript.

Funding

This research received no external funding. The work was carried out using institutional computational resources available at Mohan Babu University.

References

[1] Jameel Ahmed Bhutto et al., "CT and MRI Medical Image Fusion Using Noise-Removal and Contrast Enhancement Scheme with Convolutional Neural Network," *Entropy*, vol. 24, no. 3, pp. 1-20, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[2] Yang Liu et al., "Multi-Scale Mixed Attention Network for CT and MRI Image Fusion," *Entropy*, vol. 24, no. 6, pp. 1-20, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[3] Shangwang Liu, and Lihan Yang, "BPDGAN: A GAN-Based Unsupervised Back Project Dense Network for Multi-Modal Medical Image Fusion," *Entropy*, vol. 24, no. 12, pp. 1-16, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[4] Lei Liang, and Zhisheng Gao, "SharDif: Sharing and Differential Learning for Image Fusion," *Entropy*, vol. 26, no. 1, pp. 1-17, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[5] Qingyu Mao et al., "CT and MRI Image Fusion via Coupled Feature-Learning GAN," *Electronics*, vol. 13, no. 17, pp. 1-18, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[6] Weiwei Kong, Yiwen Li, and Yang Lei, "Medical Image Fusion Using SKWGF and SWF in Framelet Transform Domain," *Electronics*, vol. 12, no. 12, pp. 1-17, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[7] Marwah Mohammad Almasri, and Abrar Mohammed Alajlan, "Artificial Intelligence-Based Multimodal Medical Image Fusion Using Hybrid S^2 Optimal CNN," *Electronics*, vol. 11, no. 14, pp. 1-18, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[8] Mohammed Ali Saleh et al., "A Brief Analysis of Multimodal Medical Image Fusion Techniques," *Electronics*, vol. 12, no. 1, pp. 1-30, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[9] Wenqing Wang et al., "MDC-RHT: Multi-Modal Medical Image Fusion via Multi-Dimensional Dynamic Convolution and Residual Hybrid Transformer," *Sensors*, vol. 24, no. 13, pp. 1-22, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[10] Jiangwei Li et al., "Multi-Sensor Medical-Image Fusion Technique Based on Embedding Bilateral Filter in Least Squares and Salient Detection via a Deformed Smoothness Constraint," *Sensors*, vol. 23, no. 7, pp. 1-20, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[11] Rui Zhang et al., "TDFusion: When Tensor Decomposition Meets Medical Image Fusion in the Nonsubsampled Shearlet Transform Domain," *Sensors*, vol. 23, no. 14, pp. 1-21, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[12] Hanrui Chen et al., "ECFuse: Edge-Consistent and Correlation-Driven Fusion Network for Infrared and Visible Image Fusion," *Sensors*, vol. 23, no. 19, pp. 1-20, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[13] Hao Li et al., "Residual Attention-Based Image Fusion Method with Multi-Level Feature Encoding," *Sensors*, vol. 25, no. 3, pp. 1-26, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[14] Manoj Diwakar et al., "Multimodality Medical Image Fusion Using Clustered Dictionary Learning in the Non-Subsampled Shearlet Transform Domain," *Diagnostics*, vol. 13, no. 8, pp. 1-15, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[15] Maruturi Haribabu, and Velmathi Guruviah, "An Improved Multimodal Medical Image Fusion Approach Using Intuitionistic Fuzzy Set and Intuitionistic Fuzzy Cross-Correlation," *Diagnostics*, vol. 13, no. 14, pp. 1-23, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[16] Ben S. Singh et al., "Image Fusion Involving Real-Time Transabdominal or Endoscopic Ultrasound for Gastrointestinal Malignancies: Review of Current and Future Applications," *Diagnostics*, vol. 12, no. 12, pp. 1-11, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[17] Tao Zhou et al., "Hybrid Multimodal Medical Image Fusion Method Based on LatLRR and ED-D^2GAN," *Applied Sciences*, vol. 12, no. 24, pp. 1-21, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[18] Srinivasu Polinati et al., "The Fusion of MRI and CT Medical Images Using Variational Mode Decomposition," *Applied Sciences*, vol. 11, no. 22, pp. 1-16, 2021. [CrossRef] [Google Scholar] [Publisher Link]

[19] Georges Mehawed et al., "A Pilot Study of PSMA PET/CT and MRI Fusion for Prostate Cancer: Software to Replace PET/MRI Hardware," *Journal of Clinical Medicine*, vol. 13, no. 23, pp. 1-11, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[20] Hagen Fritzsche et al., "First Experiences with Fusion of PET-CT and MRI Datasets for Navigation-Assisted Percutaneous Biopsies for Primary and Metastatic Bone Tumors," *Diagnostics*, vol. 15, no. 1, pp. 1-17, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[21] Gabriele Masselli, and Chiara Di Bella, "Will PET/MR Imaging Replace PET/CT for Pediatric Applications?," *Diagnostics*, vol. 15, no. 9, pp. 1-12, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[22] Filippo Crimì et al., "New Trends and Advances in MRI and PET Hybrid Imaging in Diagnostics," *Diagnostics*, vol. 13, no. 18, pp. pp. 1-2, 2023. [CrossRef] [Google Scholar] [Publisher Link]

[23] Xiaolong Gu, Ying Xia, and Jie Zhang, "Multimodal Medical Image Fusion based on Interval Gradients and Convolutional Neural Networks," *BMC Medical Imaging*, vol. 24, pp. 1-15, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[24] Mohammed Albekairi et al., "Multimodal Medical Image Fusion Combining Saliency Perception and Generative Adversarial Network," *Scientific Reports*, vol. 15, pp. 1-21, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[25] Fei Luo et al., "A Novel Multimodel Medical Image Fusion Framework with Edge Enhancement and Cross-Scale Transformer," *Scientific Reports*, vol. 15, pp. 1-19, 2025. [CrossRef] [Google Scholar] [Publisher Link]