Edge-Enhanced Dilated Residual Attention Network for Multimodal Medical Image Fusion

Meng Zhou,* **Yuxuan Zhang,*** **Xiaolan Xu**
Department of Computer Science
University of Toronto
Toronto, Canada
{simonzhou,yuxuan,landyxu}@cs.toronto.edu

Jiayi Wang
Desautels Faculty of Management
McGill University
Montreal, Canada
jiayi.wang10@mail.mcgill.ca

Farzad Khalvati
Department of Computer Science, Department of Medical Imaging
University of Toronto
Toronto, Canada
farzad.khalvati@utoronto.ca

Abstract

Multimodal medical image fusion is a crucial task that combines complementary information from different imaging modalities into a unified representation, thereby enhancing diagnostic accuracy and treatment planning. While deep learning methods, particularly Convolutional Neural Networks (CNNs) and Transformers, have significantly advanced fusion performance, some of the existing CNN-based methods fall short in capturing fine-grained multiscale and edge features, leading to suboptimal feature integration. Transformer-based models, on the other hand, are computationally intensive in both the training and fusion stages, making them impractical for real-time clinical use. Moreover, the clinical application of fused images remains unexplored. In this paper, we propose a novel CNN-based architecture that addresses these limitations by introducing a Dilated Residual Attention Network Module for effective multiscale feature extraction, coupled with a gradient operator to enhance edge detail learning. To ensure fast and efficient fusion, we present a parameter-free fusion strategy based on the weighted nuclear norm of softmax, which requires no additional computations during training or inference. Extensive experiments, including a downstream brain tumor classification task, demonstrate that our approach outperforms various baseline methods in terms of visual quality, texture preservation, and fusion speed, making it a possible practical solution for real-world clinical applications. The code will be released at https://github.com/simonZhou86/en_dran.

1 Introduction

Medical imaging plays an increasingly prominent role in clinical diagnosis, it aims to aggregate common and complementary information from different image modalities as well as integrate the information to generate more clearer images (Xie et al., 2023). Medical image fusion can enhance crucial details of anatomy and tissue information from different image modalities and hence helps physicians and radiologists in accurate diagnosis of diseases, e.g., precise localization of tumor boundaries and tissues (Chen et al., 2024) and effective radiotherapy treatments (Safari et al., 2023; Xie et al., 2023).

Multimodal medical images provide organizations and structures from various aspects. For instance, Magnetic Resonance Imaging (MRI) offers high-resolution soft-tissue anatomical details, while computerized tomography (CT) scans reveal high-density information like bone structures

---

*Equal Contribution, corresponding to M.Zhou

1

and implants (Xie et al., 2023). Positron Emission Tomography (PET) and Single-Photon Emission Computed Tomography (SPECT) highlight metabolic activity and blood flow in tissues. Due to hardware limitations, single-modal images cannot provide accurate and comprehensive information (Xie et al., 2023). Physicians must analyze multiple images to make informed decisions, a process that is both time-consuming and laborious. Multimodal medical image fusion (MMIF) addresses this challenge by integrating prominent and complementary information into a single, more visually perceptive image, thereby supporting more accurate diagnoses (Wang et al., 2022).

Recent advances in deep learning have significantly improved multimodality fusion performance due to its powerful representation capabilities. Several works have utilized convolutional neural networks (CNNs) for various image fusion tasks. For example, Li & Wu (2019) proposed DenseFuse, an infrared and visible image fusion framework using dense blocks. To address the limitation of single-scale feature extraction in Li & Wu (2019), Song et al. (2019) introduced a multi-scale DenseNet (MSDNet), capturing features at different scales with various convolutional kernel sizes. Additionally, Zhang et al. (2020) proposed a general CNN-based image fusion framework for multifocus, infrared-visible, and multimodal medical image fusion. They introduced elementwise fusion rules to combine feature maps directly. There are also many efforts on solely multi-modal medical image fusion tasks. Fu et al. (2021) introduced the residual pyramid attention network capable of better deep feature extraction capabilities for MRI-CT, MRI-PET, and MRI-SPECT fusion tasks. They also proposed a Feature Energy Ratio Strategy to fuse two feature maps in the latent space. Similarly, Li et al. (2022) proposed a double residual attention network to capture detailed features while avoiding gradient vanishing or explosion. However, both Fu et al. (2021) and Li et al. (2022) suffer from losing the structural information and edge details, which are crucial for medical images (Meng et al., 2019). Recently, Transformer-based methods (Vaswani et al., 2017) have also garnered attention in computer vision tasks (Dosovitskiy et al., 2020), with some applications in medical image fusion. For example, Ma et al. (2022) proposed SwinFusion, combining a CNN feature extractor with a cross-domain transformer model to fuse local and global information, achieving superior fusion performance on MRI-CT and MRI-PET tasks. However, the high computational cost of computing global interactions hinders its clinical applications. Xie et al. (2023) proposed a multiscale CNN model and applied residual Swin Transformer layers in the fusion strategy. Although their method further improved the fusion performance on all MRI-CT, MRI-PET, and MRI-SPECT tasks, the computational cost is still high and hinders the clinical application. Recently, Xie et al. (2024b) proposed a lightweight cross-modality transformer using window and grid attention, minimizing computational costs while maintaining superior fusion performance. More recently, the emergency of the improved selective structured state space models (Mamba) (Gu & Dao, 2023), provides a novel solution to the problem above. Mamba has been demonstrated to outperform Transformer models in tasks requiring long-term dependency modeling, due to its selective global information modeling capabilities while maintaining linear complexity. This design not only reduces the computational costs but also enhances the inference speed. Notably, several studies have already explored the application of Mamba in various multimodal image fusion tasks (Peng et al., 2024; Li et al., 2024; Xie et al., 2024a). Despite these advancements, many of the existing fusion methods lack validation on downstream tasks and remain unexplored. Therefore, a computationally friendly model that preserves both soft-tissue and detailed structural information (e.g., edges) is crucial for practical clinical applications. To this end, we propose a novel end-to-end feature fusion framework for multimodal anatomical and functional image fusion. We introduce the Dilated Residual Attention Network (DRAN) to extract multi-scale features and a family of *parameter-free* fusion strategies to fuse feature maps in the latent space. Our approach combines input images to generate a fused image with more detailed textures and less information loss, surpassing several state-of-the-art fusion methods both qualitatively and quantitatively. We also take the clinical applicability into account and validate our approach on a downstream region-of-interest(ROI)-based brain tumor classification task between high-grade gliomas (HGG) and low-grade gliomas (LGG) following a similar procedure done in Zhou & Khalvati (2024). To summarize, our contributions are as follows:

1. We propose a novel Dilated Residual Attention Network (DRAN) for extracting multi-scale fine-grained features effectively. We also incorporate a learnable dense residual gradient operator (DRGO) to enhance edge feature representations.

2. We introduce a family of *parameter-free* fusion strategy, Softmax Feature Weighted Strategy, that outperforms other existing parameter-free fusion strategies and takes a step further to real-time fusion.

3. To the best of our knowledge, we are the first to provide preliminary evaluations of the proposed fusion framework on a downstream brain LGG/HGG tumor type classification task.

4. Experiments show our proposed method outperforms several baselines in both objective fusion metrics and subjective image quality, as well as the performance on the downstream classification task.
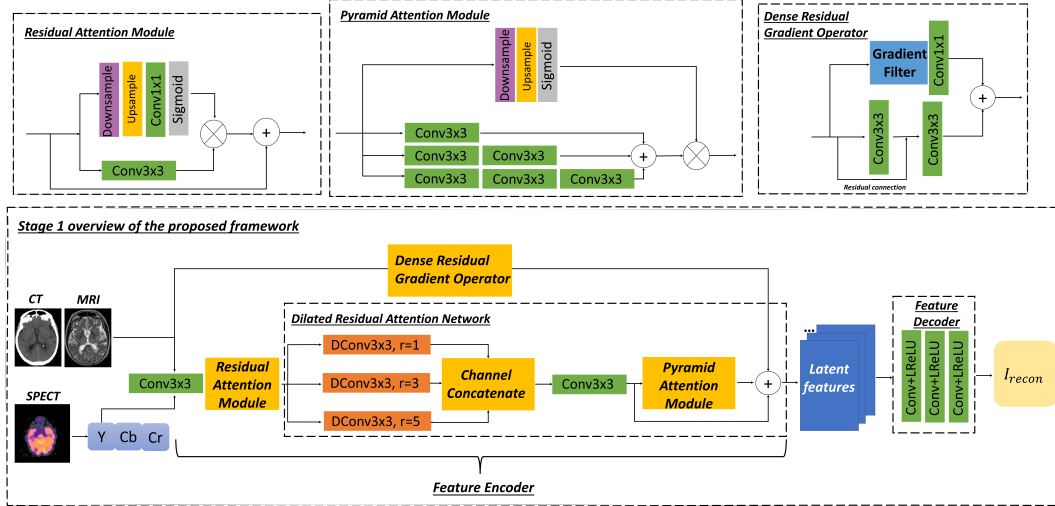
2 MATERIALS AND METHOD

Figure 1: An overview of Stage 1 model in the proposed framework. `DConv3x3, r=c` represents dilated convolution with kernel size 3×3 and dilation rate equals to c. All `Conv+LReLU` layers in the decoder have 3×3 kernel followed by Leaky-ReLU. Note that the YCbCr conversion only applies to SPECT images.

The proposed fusion framework consists of three main components: a feature encoder, a fusion module, and a feature decoder. Specifically, we designed an **asymmetric autoencoder** comprising a deep feature encoder and a lightweight feature decoder. The autoencoder extracts multi-scale features from the input image into the latent space using the proposed DRAN and an edge enhancer, then reconstructs the image back to the original space, as illustrated in Figure 1. Since there is no golden ground truth for the fusion task, we disentangle the training process of the proposed fusion framework into two stages following recent works (Xie et al., 2023; Fu et al., 2021; Li et al., 2022). In the first stage, we train the autoencoder *solely* to extract multi-scale deep features and reconstruct images through a general reconstruction task. In the second stage depicted in Figure 2, we introduce a family of *parameter-free* fusion strategy utilizing softmax weights and the nuclear norm to fuse two feature maps, f^a and f^b, into a single feature map f^c. Using the trained decoder, this fused feature map f^c is then decoded back to the image space to obtain the final fused image. Below, we provide a detailed explanation of each component.

2.1 FEATURE EXTRACTION AND IMAGE RECONSTRUCTION

Feature Encoder. The core module in the feature encoder is the **Dilated Residual Attention Network** (DRAN), which is inspired by two state-of-the-art mechanisms: residual attention (Wang et al., 2017) and pyramid attention (Li et al., 2018). The residual attention mechanism includes a sequence of convolutional layers for feature processing and, in parallel, a downsample-upsample block with a sigmoid function for learning weights based on feature importance. The residual attention network enhances the feature expression capability through the attention mechanism and aids in faster model convergence by preventing gradient vanishing or explosion via the residual connection. A graphical illustration of our residual attention module is shown at the top-left of Figure 1. However, residual attention alone cannot effectively extract and learn multi-scale semantic features. To

address this, we incorporate an additional pyramid attention network (Li et al., 2018). The convolution block in pyramid attention usually contains multiple convolutional layers that capture features at different scales and receptive fields. We adopt this approach by replacing convolutions with larger kernel filters with a sequence of smaller kernel filters (Szegedy et al., 2016). As shown at the top-middle of Figure 1, we use a single 3×3 convolution to represent a 3×3 receptive field, *two* 3×3 convolutions to represent a 5×5 receptive field, and *three* 3×3 convolutions to represent a 7×7 receptive field. Different from Fu et al. (2021), we leverage the $\{1, 3, 5\}$-dilated convolution (Yu & Koltun, 2015) on shallow features of the original input image to further enhance the learning ability of local multi-scale information and fine details without downsampling the feature map. The receptive field is expanded using three different dilated convolutions to improve the *discriminative* multi-scale feature extraction ability of the model. Once the multi-scale features are extracted, they are concatenated channel-wise. The residual-pyramid attention paradigm is then applied to further extract deep features. These deep features are the output of the feature extraction module and are used in both the fusion and reconstruction modules.

Edge information in MRI and CT is also crucial for accurate feature representation. To precisely capture these fine-grained edge features, we introduce a learnable dense residual gradient operator (DRGO) to enhance edge feature representation. The proposed module aggregates learnable convolutional features with gradient magnitude information on shallow features, and these are directly added element-wise to the features after the DRAN block, as shown at the bottom of Figure 1. Our DRGO module uses two convolutional layers with 3×3 kernel sizes with residual connections to extract features. Additionally, a Sobel gradient operator (Kanopoulos et al., 1988), followed by a 1×1 convolutional layer, is used to learn the gradient information. This combination enhances the edge features by integrating both convolutional and gradient-based information. A graphical illustration of the DRGO module is shown at the top-right of Figure 1.

Asymmetric Autoencoder. The detailed architecture for the feature encoder is shown in Figure 1. After we obtained the latent features of input images from the feature encoder, we passed them into a lightweight decoder which only contains three convolutional layers and leaky ReLU activation function. The rationale behind this asymmetric design is that if the latent feature maps contain a wealth of useful information, they can be easily reconstructed back to the original image. Therefore, we employ a lightweight decoder to maximize the encoder's capability to extract rich semantic information in the latent space. Note that the latent features will be directly used in the subsequent fusion process, making the informative latent features crucial to the quality of the overall fused image.

2.2 FEATURE FUSION

Softmax Weighted Fusion Strategy. The fusion strategy is used to fuse the extracted features of input images into a single feature map. In this work, we introduce a novel parameter-free fusion strategy termed Softmax Feature Weighted Strategy. First, we obtained two output feature maps f^a, f^b from the extraction module for input images I_a, I_b, respectively. These feature maps can be used to generate the corresponding weight maps that indicate the amount of contribution of each pixel to the final fused feature map (Lahoud & Süsstrunk, 2019). then, to get the weight map, we take the channel-wise softmax operation to the feature map which can be realized by Equation equation 1, where x_i is the ith channel of the output feature map x.

$$S(x_i) = \frac{\exp(x_i)}{\sum_i \exp(x_i)} \tag{1}$$

After we obtained the softmax output, we computed the matrix nuclear norm ($\| \cdot \|_*$), which is the summation of its singular values. Finally, we obtain the weights for the output feature map by taking the weighted average of the *maximum value* of the nuclear norm. The formula is given in Equation equation 2.

$$W_k = \frac{\phi(\|S(x_i)^k\|_*)}{\sum_{k=1}^{C} \phi(\|S(x_i)^k\|_*)} \tag{2}$$

Where C is the number of multimodality images ($C = 2$ in our work), $S(x_i)^k, k \in [a, b]$ is the weight map after softmax operation for the feature map $f^k, k \in [a, b]$. As mentioned, we selected $\phi(\cdot)$ to be `max()` in this work, but it can also be `mean()`, `sum()` or `identity()`. The final fused feature map is then given by $f = \sum W_k * f_k, k \in [a, b]$.
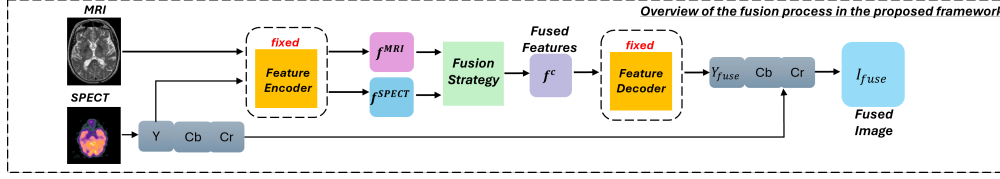

Figure 2: A sample illustration of the fusion process, we take MRI-SPECT fusion as an example. The fusion process is the same for MRI-CT fusion except the RGB to YCbCr conversion is ignored.

2.3 LOSS FUNCTION

Since only the autoencoder in Stage 1 involves training and optimization, we use the reconstruction objective which combines pixel loss, image gradient loss (Ma et al., 2020; Fabbri et al., 2018), and perceptual loss (Johnson et al., 2016) to optimize the network. The gradient loss (Ma et al., 2020; Fabbri et al., 2018) is added to model the fine details of textures in the reconstructed image, and the perceptual loss (Johnson et al., 2016) is added to model the high-level semantic similarity between reconstructed and input images. In detail, our loss function is defined as follows:

$$\mathcal{L}_{pixel} = \|x - \hat{x}\|_2^2, \; \mathcal{L}_{grad} = \|\nabla x - \nabla \hat{x}\|_2^2,$$
$$\mathcal{L}_{perp} = \sum_{k=1}^{C} \|f_i^k(x) - f_i^k(\hat{x})\|_2^2 \tag{3}$$

$$\mathcal{L}(\theta) = \mathcal{L}_{pixel} + \lambda_1 * \mathcal{L}_{grad} + \lambda_2 * \mathcal{L}_{perp} \tag{4}$$

Equation equation 3 shows the three losses utilized. We use a standard $\mathcal{L}_2$ distance for the pixel loss, where M is the number of input images, I_o is the output image. The image gradient loss is also realized by the $\mathcal{L}_2$ norm of the image gradient in both x and y-direction. Finally, the perceptual loss (Johnson et al., 2016), where $f_i^k(x)$ is the the k-th channel in i-th layer (with size $W_i \times H_i$) from the pre-trained VGG16 network (Simonyan & Zisserman, 2014) with input image x, and C is the number of channels. We prefer i to be large, i.e., the deeper layer of the VGG network. The total loss function is given in Equation equation 4, θ is the set of network weights to be optimized, λ_1, λ_2 are weight balancing factors of the gradient and perceptual loss, respectively. We empirically set λ_1, λ_2 equal to 0.5.

2.4 DATA AND PREPROCESSING

In this work, we use three datasets to validate the effectiveness of our proposed approach: MRI-CT (184 pairs) and MRI-SPECT (357 pairs) multi-modality fusion from The Harvard Whole Brain Atlas dataset[1], and additional FLAIR and T2 sequence data from the BraTS 2019 dataset (Bakas et al., 2017; 2018; Menze et al., 2014) (335 patients) to evaluate our method on the downstream brain tumor type classification task. For MRI-SPECT fusion, we converted SPECT images from the RGB color space to the YCbCr space following (Fu et al., 2021; Xie et al., 2023; Li et al., 2022), using only the Y-channel images to train the model. For MRI-CT fusion, we used the original images as they are all single-channel grayscale images. All MRI-CT and MRI-SPECT image pairs were co-registered and preprocessed beforehand so that each pixel intensity is in the range of [0,255]. We further normalized the pixel intensity to [0,1].

For the T2 and FLAIR sequence data from BraTS, we first obtained the ROIs by multiplying the images with masks, then reshaped the data from 240×240×155 to 128×128×128, and normalized

---

[1]https://www.med.harvard.edu/aanlib/

all pixel intensities to the range [0,1]. We converted the 3D data to 2D by slicing over the *Axial plane* for each patient and only considered slices with at least 10% non-zero pixels.

3 Experiments

All programs were implemented in PyTorch and were trained on Google Colab and Compute Canada. For both MRI-CT and MRI-SPECT pairs, we trained the autoencoder for 100 epochs with an initial learning rate of 0.0001 and cosine decay to 3e-7, a mini-batch size of 4, and with the Adam optimizer (Kingma & Ba, 2014). We randomly held out 30 image pairs from the MRI-CT dataset and 50 pairs from the MRI-SPECT dataset as the standalone test set. To ensure the robustness of our model, we repeated our experiments three times and ensured that we had different test sets in each run.

To assess the usability of our fusion framework, we conducted an ROI-based brain tumor classification task between LGG and HGG using the BraTS 2019 data as we discussed previously. First, we randomly held out 40 *patients* (20 LGG and 20 HGG patients, 1152 slices in total) as a standalone test set. The rest of the data is used to train our model. We trained our fusion autoencoder for 25 epochs with a constant learning rate of 0.0001, a mini-batch size of 4, and with the Adam optimizer. We used our proposed SFNN-max fusion strategy to fuse T2 and FLAIR images. For the classification model, we used ResNet-50 for all experiments and trained with focal loss (Lin et al., 2017) followed (Zhou & Khalvati, 2024). We trained the model for 50 epochs with a constant learning rate of 0.001, a mini-batch size of 8, and an Adam optimizer. We ran the classification experiment for three trials with different train-validation splits to ensure the robustness and reliability of our findings.

Baseline Model & Comparison. For comparison of image fusion results on MRI-CT and MRI-SPECT data, we considered five open-sourced state-of-the-art methods: IFCNN (Zhang et al., 2020), MSRPAN (Fu et al., 2021), MSDRA (Li et al., 2022), SwinFusion (Ma et al., 2022) and MRSCFusion (Xie et al., 2023), where the first three are CNN-based methods and the last two are Transformer-based methods. We rerun all models except MRSCFusion due to memory limitations, and we used the pre-trained weights of MRSCFusion from their official repository[2]. For quantitative comparisons, we select five commonly used metrics in previous works (Xie et al., 2023; Li et al., 2022; Fu et al., 2021; Chen et al., 2024; Safari et al., 2023): Peak signal-to-noise ratio (PSNR), Structural Similarity (SSIM) (Wang et al., 2004), Feature Mutual Information (Haghighat et al., 2011), Feature SSIM (FSIM) (Zhang et al., 2011), and Information Entropy (EN). The downstream classification performance is evaluated using AUC, F1-Score, and Accuracy.

4 Results and Discussions

4.1 Image Fusion Results

Main Results. In Figure 3, we compare the fusion results of different methods across three randomly selected MRI-CT and MRI-SPECT test pairs. For the MRI-CT fusion task, we focus on the area with more tissue information from MRI images and bone structures from CT images (highlighted and zoomed in red). We expect that the dense information (e.g., bone structures) in CT images and soft tissues in MRI images should be simultaneously retained in the fused images. We observed that the MRSCFusion had undesirable pixel intensities and some of the details from the MRI image were missing, such as the brain contour line. SwinFusion preserves tissue information from MRI images well but fails to maintain the boundary details for both CT and MRI images. IFCNN retains MRI tissue details well but slightly weakens dense structures from CT images. MSRPAN struggles to distinguish tissue boundaries between CT and MRI, losing fine-grained details (e.g., the second image in the MSRPAN column), resulting in overly sharp edges. MSDRA produces unclear boundaries and lacks significant intensity differences, making it difficult to differentiate between dense boundaries and soft tissues. In contrast, our proposed method provides a clear, bright brain contour from MRI and better edge contrast from CT. Our fused images offer superior contrast between edges and inner tissues, preserving more edge and fine-grained information from both modalities. Visually, our results appear more natural and the overall contrast is more promising.

---

[2]https://github.com/millieXie/MRSCFusion

For the MRI-SPECT fusion task, we focus on the area with more morphological (color) information from the SPECT and texture information from MRI images (highlighted and zoomed in red box) for a better comparison. MRSCFusion introduces noticeable intensity distortions and loses texture details from MRI images. SwinFusion achieves satisfactory results, but some color information blurs the MRI texture. IFCNN effectively represents the functional information from SPECT images but still misses some fine details from MRI images. MSRPAN and MSDRA preserve color information well but blur some MRI texture details. Our fused images retain the appropriate color information from SPECT while preserving more structure and tissue details from MRI.
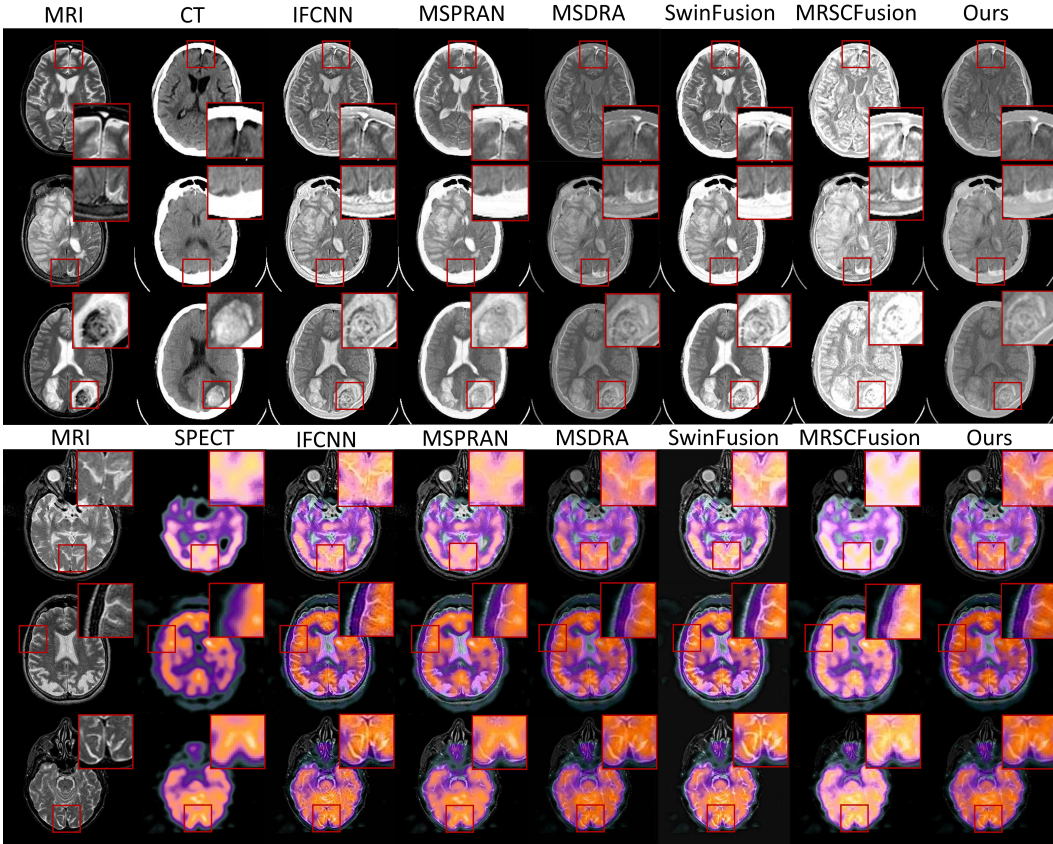

Figure 3: Qualitative results for MRI-CT (top three rows) and MRI-SPECT (bottom three rows) fusion task. We randomly select three sample pairs from both test sets and show the fusion results across different methods. Zoom in for a better view.

The quantitative metrics, computed over three distinct test sets, are reported with mean values and standard deviations in Table 1. Our EH-DRAN method achieves the best performance in terms of PSNR, FMI, FSIM, and Information Entropy for the MRI-CT fusion task. The high FMI, FSIM, and Information Entropy scores indicate that our fused images maintain superior structural similarity and contain richer information. Although our method shows a slightly lower SSIM score compared to SwinFusion, this may be attributed to SwinFusion's tendency to emphasize MRI features. In contrast, our approach balances the contributions from both MRI and CT images.

For the MRI-SPECT fusion task, our method consistently surpasses baseline methods in PSNR, SSIM, FMI, and FSIM. Despite slightly lower Entropy values than SwinFusion, all other metrics demonstrate that our approach effectively preserves more functional and morphological information from MRI and SPECT images. This outcome aligns with the qualitative fusion results discussed earlier.

Selection of Fusion Strategies. As discussed in Section 2.2, our proposed fusion strategy offers several variants, including $\phi(\cdot) = \text{mean}()$, $\text{sum}()$, and $\text{max}()$. We conducted an extensive qual-

Table 1: Comparison between different methods on two test datasets, **bold** and <u>underline</u> numbers represent best and second-best values in each dataset, respectively.

Dataset	Method	PSNR	SSIM	FMI	FSIM	EN
MRI-CT	IFCNN (Zhang et al., 2020)	<u>15.594±0.112</u>	0.700±0.015	0.870±0.012	0.801±0.001	8.968±0.227
	MSRPAN (Fu et al., 2021)	14.790±0.233	0.749±0.003	0.744±0.001	0.804±0.001	7.773±0.273
	MSDRA (Li et al., 2022)	15.308±0.437	0.742±0.037	0.872±0.002	0.788±0.005	<u>9.554±0.767</u>
	SwinFusion (Ma et al., 2022)	14.962±0.173	**0.768±0.007**	0.882±0.002	<u>0.810±0.001</u>	8.445±0.078
	MRSCFusion (Xie et al., 2023)	14.476±0.205	0.713±0.012	0.877±0.006	0.791±0.010	7.544±0.232
	EH-DRAN(Ours)	**16.830±0.490**	<u>0.753±0.007</u>	**0.883±0.005**	**0.820±0.003**	**10.727±0.531**
MRI-SPECT	IFCNN (Zhang et al., 2020)	<u>19.728±0.228</u>	0.721±0.025	<u>0.846±0.062</u>	0.783±0.027	10.167±0.429
	MSRPAN (Fu et al., 2021)	19.174±0.046	0.732±0.002	0.838±0.003	0.793±0.002	9.737±0.202
	MSDRA (Li et al., 2022)	19.662±0.165	0.725±0.003	0.839±0.003	0.794±0.003	10.784±0.447
	SwinFusion (Ma et al., 2022)	17.557±0.021	0.728±0.004	0.808±0.007	<u>0.819±0.011</u>	**13.066±0.428**
	MRSCFusion (Xie et al., 2023)	18.412±0.211	<u>0.734±0.012</u>	0.827±0.009	0.814±0.006	9.87±0.600
	EH-DRAN(Ours)	**21.455±0.071**	**0.736±0.002**	**0.876±0.004**	**0.843±0.003**	<u>11.970±0.538</u>

itative and quantitative analysis of these strategies for both the MRI-CT and MRI-SPECT fusion tasks, with the results summarized in Table 2. For comparison, we used FER (Fu et al., 2021) and FL1N (Li et al., 2022), two parameter-free fusion strategies from prior research, as baselines.

Our results demonstrated that the proposed fusion strategy consistently outperformed the baseline methods across all quantitative metrics. Notably, the robustness of our approach was evident, as the performance varied only slightly across different $\phi(\cdot)$ functions. This robustness underscores the stability and reliability of our strategy. Detailed qualitative comparisons of different fusion strategies are provided in Appendix A.

Table 2: Comparison between different fusion strategies on two test datasets.

Dataset	Fusion Strategy	PSNR	SSIM	FMI	FSIM	EN
MRI-CT	FER (Fu et al., 2021)	14.718±0.549	0.743±0.005	0.874±0.002	0.798±0.008	8.675±0.182
	FL1N (Li et al., 2022)	15.620±0.279	0.737±0.002	0.878±0.003	0.804±0.008	9.001±0.583
	SFNN-mean	15.631±0.290	0.736±0.003	0.877±0.002	0.820±0.003	9.013±0.351
	SFNN-max	16.830±0.490	0.753±0.007	0.883±0.005	0.820±0.003	10.727±0.531
	SFNN-sum	15.590±0.370	0.735±0.005	0.876±0.003	0.810±0.007	9.397±0.732
MRI-SPECT	FER (Fu et al., 2021)	19.635±0.039	0.832±0.004	0.832±0.005	0.796±0.003	9.751±0.112
	FL1N (Li et al., 2022)	20.337±0.058	0.833±0.002	0.842±0.006	0.800±0.010	10.562±0.482
	SFNN-mean	20.337±0.043	0.734±0.003	0.841±0.003	0.836±0.006	10.115±0.574
	SFNN-max	21.455±0.071	0.736±0.002	0.876±0.004	0.843±0.003	11.970±0.538
	SFNN-sum	20.336±0.136	0.734±0.002	0.842±0.003	0.838±0.010	10.937±0.451

Fusion Time. Next, we assess the model complexity by examining the total number of trainable parameters and the image fusion time for each image pair using the MRI-SPECT dataset, as detailed in Table 3. The MRI-SPECT dataset is selected due to its larger number of image pairs and its representation of a more complex task, which closely mirrors real-world scenarios. Compared to other baselines, our proposed method has a reasonable number of parameters and achieves image fusion in only 1 second. The proposed parameter-free fusion strategy suggests great potential for **real-time** fusion in clinical settings. Furthermore, we expect the fusion time for MRI-CT pairs to be even shorter than that for MRI-SPECT, primarily due to the absence of color space conversion between RGB and YCbCr.

Table 3: Comparison between different methods in average inference time on MRI-SPECT dataset using a T4 GPU.

	IFCNN	MSRPAN	MSDRA	SwinFusion	MRSCFusion	Ours
Params(M)	0.08	0.10	0.20	0.97	23.00	0.50
Time(s)	0.89	0.79	0.81	1.31	2.85	1.26

Ablation Study. We performed our ablation studies in two folds. (1) to assess the effectiveness of the proposed Dense Residual Gradient Operator (DRGO) module to learn edge details, (2) to evaluate the impact of incorporating the gradient loss during model optimization. The ablation experiments were performed using both the MRI-CT and MRI-SPECT datasets. We hypothesize that (1) the edge enhancer aids the autoencoder in extracting meaningful edge features, thereby reinforc-

ing edge details from both source images in the fused images, and (2) the gradient loss assists the model in learning and reproducing fine texture details in the reconstructed output. Table 4 presents the results of our ablation study. We began with a baseline model that excluded both the DRGO module and the gradient loss ($\mathcal{L}_{grad}$). Adding $\mathcal{L}_{grad}$ to the loss function led to an improvement across all evaluation metrics compared to the baseline, confirming the benefit of gradient loss for optimizing texture detail learning. When we further integrated the DRGO module into the model, we observed additional improvements across all metrics. The notable increase in SSIM and FSIM scores demonstrates the DRGO module's effectiveness in preserving edge information from both source images, thus validating its contribution to overall fusion quality.

Table 4: Ablations on the edge enhancer and loss function components of our proposed framework on both datasets. Base Model represents the model that is trained without using the DRGO module and the loss function only contains pixel and perceptual loss.

Dataset	Method	PSNR	SSIM	FMI	FSIM	Entropy
MRI-CT	Base Model	15.623±0.032	0.745±0.013	0.878±0.003	0.802±0.003	9.122±0.706
	Base Model+$\mathcal{L}_{grad}$	16.355±0.038	0.749±0.010	0.881±0.002	0.818±0.002	9.771±0.528
	Base Model+$\mathcal{L}_{grad}$+DRGO	**16.830±0.490**	**0.753±0.007**	**0.883±0.005**	**0.820±0.003**	**10.727±0.531**
MRI-SPECT	Base Model	20.698±0.002	**0.743±0.008**	0.833±0.006	0.836±0.005	10.010±0.563
	Base Model+$\mathcal{L}_{grad}$	20.738±0.026	0.740±0.011	0.837±0.002	0.838±0.004	10.454±0.426
	Base Model+$\mathcal{L}_{grad}$+DRGO	**21.455±0.071**	0.736±0.002	**0.876±0.004**	**0.843±0.003**	**11.970±0.538**

4.2 CLASSIFICATION RESULTS

To validate the efficacy of the proposed fusion framework in clinical multi-modality brain tumor images, we trained a classification model to validate the efficacy of the proposed fusion framework in distinguishing between HGG and LGG brain tumor types. Following (Xie et al., 2024b), we used T2 and FLAIR images for the classification task. For comparison, we evaluated three approaches: single-modality images (using either T2 or FLAIR), dual-modality images (using a channel-wise concatenation of T2 and FLAIR), and T2-FLAIR fused images. The classification results, presented in Table 5, demonstrate that utilizing the fused image generated by our proposed fusion framework significantly improves performance in terms of AUC and F1-Score compared to the single-modality and dual-modality baselines. This improvement highlights our method's ability to enhance details in the brain tumor ROIs while preserving overall image contrast, corroborating the findings from the previous section. These results further suggest the potential of leveraging fused images for improved accuracy in real clinical diagnosis.

Table 5: Comparison of brain tumor type LGG/HGG classification performance between single-modality, dual-modality, and T2-FLAIR fused images. Values are reported as mean±standard deviation.

	AUC	F1-Score	Accuracy
T2 (1-channel)	0.722±0.021	0.703±0.018	0.604±0.037
FLAIR (1-channel)	0.727±0.024	0.701±0.008	0.611±0.017
T2+FLAIR (2-channel)	0.723±0.028	0.717±0.012	0.640±0.015
Fused (1-channel)	**0.769±0.003**	**0.723±0.006**	**0.640±0.011**

5 CONCLUSIONS

In this work, we proposed a novel asymmetric autoencoder architecture incorporating a Dilated Residual Attention Network (DRAN) for effective multi-scale feature extraction. Additionally, we integrated a Dense Residual Gradient Operator (DRGO) as an edge enhancer to capture fine-grained edge details. We introduced a family of *parameter-free* softmax-weighted fusion strategies for multimodal image fusion, designed to operate without requiring parameter computation during both training and inference phases, pushing us a step further to achieve real-time image fusion. Our extensive evaluations demonstrate that the proposed method outperforms several baseline approaches in subjective visual quality and objective fusion metrics. The improved performance in the downstream brain tumor classification task further highlights the effectiveness of our fusion framework.

We envision our approach being applied to subsequent disease localization tasks, radiotherapy treatment planning, and surgical navigation in real-world clinical settings.

In the future, we plan to extend our method from 2D to 3D medical image fusion, as 3D data is more common in practical medical imaging. Furthermore, we aim to explore the integration of Transformers or Selective State Spaces Models (Gu & Dao, 2023) to enhance feature extraction and image fusion capabilities.

REFERENCES

Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S Kirby, John B Freymann, Keyvan Farahani, and Christos Davatzikos. Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. *Scientific data*, 4 (1):1–13, 2017.

Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv preprint arXiv:1811.02629*, 2018.

Wei Chen, Qixuan Li, Heng Zhang, Kangkang Sun, Wei Sun, Zhuqing Jiao, and Xinye Ni. Mr–ct image fusion method of intracranial tumors based on res2net. *BMC Medical Imaging*, 24(1):169, 2024.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

Cameron Fabbri, Md Jahidul Islam, and Junaed Sattar. Enhancing underwater imagery using generative adversarial networks. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7159–7165. IEEE, 2018.

Jun Fu, Weisheng Li, Jiao Du, and Yuping Huang. A multiscale residual pyramid attention network for medical image fusion. *Biomedical Signal Processing and Control*, 66:102488, 2021.

Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023.

Mohammad Bagher Akbari Haghighat, Ali Aghagolzadeh, and Hadi Seyedarabi. A non-reference image fusion metric based on mutual information of image features. *Computers & Electrical Engineering*, 37(5):744–756, 2011.

Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pp. 694–711. Springer, 2016.

Nick Kanopoulos, Nagesh Vasanthavada, and Robert L Baker. Design of an image edge detection filter using the sobel operator. *IEEE Journal of solid-state circuits*, 23(2):358–367, 1988.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Fayez Lahoud and Sabine Süsstrunk. Zero-learning fast medical image fusion. In *2019 22th International Conference on Information Fusion (FUSION)*, pp. 1–8. IEEE, 2019.

Hanchao Li, Pengfei Xiong, Jie An, and Lingxue Wang. Pyramid attention network for semantic segmentation. *arXiv preprint arXiv:1805.10180*, 2018.

Hui Li and Xiao-Jun Wu. Densefuse: A fusion approach to infrared and visible images. *IEEE Transactions on Image Processing*, 28(5):2614–2623, 2019. doi: 10.1109/TIP.2018.2887342.

Weisheng Li, Xiuxiu Peng, Jun Fu, Guofen Wang, Yuping Huang, and Feifei Chao. A multiscale double-branch residual attention network for anatomical–functional medical image fusion. *Computers in Biology and Medicine*, 141:105005, 2022.

Zhe Li, Haiwei Pan, Kejia Zhang, Yuhua Wang, and Fengming Yu. Mambadfuse: A mamba-based dual-phase model for multi-modality image fusion. *arXiv preprint arXiv:2404.08406*, 2024.

Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pp. 2980–2988, 2017.

Cheng Ma, Yongming Rao, Yean Cheng, Ce Chen, Jiwen Lu, and Jie Zhou. Structure-preserving super resolution with gradient guidance. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 7769–7778, 2020.

Jiayi Ma, Linfeng Tang, Fan Fan, Jun Huang, Xiaoguang Mei, and Yong Ma. Swinfusion: Cross-domain long-range learning for general image fusion via swin transformer. *IEEE/CAA Journal of Automatica Sinica*, 9(7):1200–1217, 2022.

Lingyu Meng, Xiaopeng Guo, and Huaguang Li. Mri/ct fusion based on latent low rank representation and gradient transfer. *Biomedical Signal Processing and Control*, 53:101536, 2019.

Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34 (10):1993–2024, 2014.

Siran Peng, Xiangyu Zhu, Haoyu Deng, Zhen Lei, and Liang-Jian Deng. Fusionmamba: Efficient image fusion with state space model. *arXiv preprint arXiv:2404.07932*, 2024.

Mojtaba Safari, Ali Fatemi, and Louis Archambault. Medfusiongan: multimodal medical image fusion using an unsupervised deep generative adversarial network. *BMC Medical Imaging*, 23(1): 203, 2023.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

Xu Song, Xiao-Jun Wu, and Hui Li. Msdnet for medical image fusion. In *International conference on image and graphics*, pp. 278–288. Springer, 2019.

Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3156–3164, 2017.

Guofen Wang, Weisheng Li, Xinbo Gao, Bin Xiao, and Jiao Du. Functional and anatomical image fusion based on gradient enhanced decomposition model. *IEEE Transactions on Instrumentation and Measurement*, 71:1–14, 2022.

Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

Xinyu Xie, Xiaozhi Zhang, Shengcheng Ye, Dongping Xiong, Lijun Ouyang, Bin Yang, Hong Zhou, and Yaping Wan. Mrscfusion: Joint residual swin transformer and multiscale cnn for unsupervised multimodal medical image fusion. *IEEE Transactions on Instrumentation and Measurement*, 72: 1–17, 2023. doi: 10.1109/TIM.2023.3317470.

Xinyu Xie, Yawen Cui, Chio-In Ieong, Tao Tan, Xiaozhi Zhang, Xubin Zheng, and Zitong Yu. Fusionmamba: Dynamic feature enhancement for multimodal image fusion with mamba. *arXiv preprint arXiv:2404.09498*, 2024a.

Xinyu Xie, Xiaozhi Zhang, Xinglong Tang, Jiaxi Zhao, Dongping Xiong, Lijun Ouyang, Bin Yang, Hong Zhou, Bingo Wing-Kuen Ling, and Kok Lay Teo. Mactfusion: Lightweight cross transformer for adaptive multimodal medical image fusion. *IEEE Journal of Biomedical and Health Informatics*, 2024b.

Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.

Lin Zhang, Lei Zhang, Xuanqin Mou, and David Zhang. Fsim: A feature similarity index for image quality assessment. *IEEE transactions on Image Processing*, 20(8):2378–2386, 2011.

Yu Zhang, Yu Liu, Peng Sun, Han Yan, Xiaolin Zhao, and Li Zhang. Ifcnn: A general image fusion framework based on convolutional neural network. *Information Fusion*, 54:99–118, 2020.

Meng Zhou and Farzad Khalvati. Conditional generation of 3d brain tumor regions via VQGAN and temporal-agnostic masked transformer. In *Medical Imaging with Deep Learning*, 2024. URL `https://openreview.net/forum?id=LLoSHPorlM`.

A MORE ON FUSION STRATEGIES

Qualitative of different fusion strategies: We selected two variants of our proposed SFNN methods, `sum()` and `max()`, due to their superior visual quality and performance, as illustrated in Figure 4. For the MRI-CT fusion task, our method demonstrated clear advantages over the FER strategy (Fu et al., 2021), achieving enhanced fidelity with sharper delineation of inner tissue structures from MRI images and well-preserved boundaries from CT images. Furthermore, compared to the FL1N method (Li et al., 2022), our approach produced brighter and more defined edges, thereby improving the contrast between edges and tissues. In the case of MRI-SPECT fusion, our method retained more detailed features from the MRI image compared to the FER method. While the visual distinctions between our method and the FL1N strategy were subtle, the quantitative analysis further demonstrates the superiority of our approach.
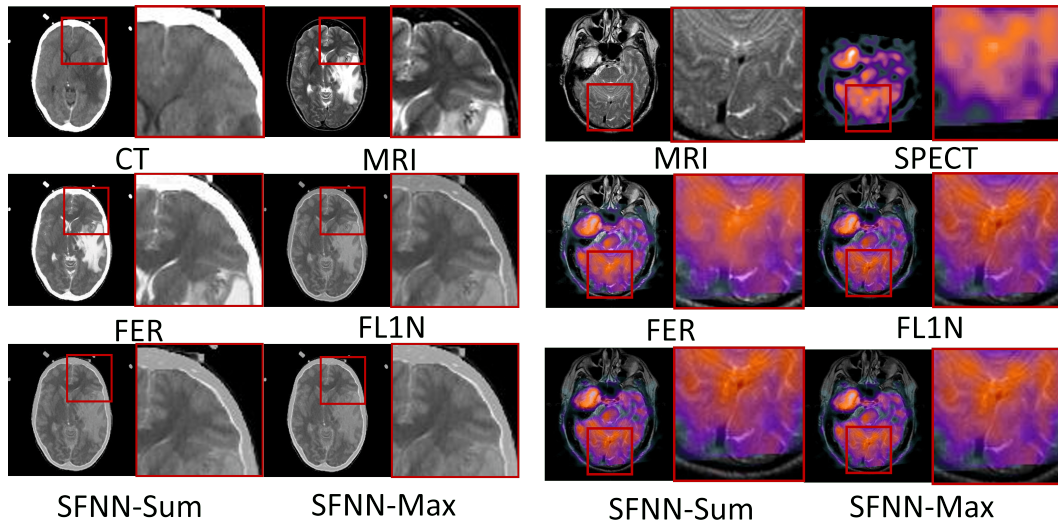
Figure 4: Qualitative comparison between different fusion strategies on MRI-CT and SPECT dataset. *Left panel:* visualization on MRI-CT dataset, *right panel:* visualization on the MRI-SPECT dataset. For both datasets, we randomly sample a pair from the test set and zoom in on a selected region for a better view.