MemoryFusion: Cross-Temporal Memory Learning for Multimodal Video Fusion

Meiqi Gong
Electronic Information School
Wuhan University
Wuhan, China
meiqigong@whu.edu.cn

Hao Zhang*
Electronic Information School
Wuhan University
Wuhan, China
zhpersonalbox@gmail.com

Jiayi Ma*
School of Robotics
Wuhan University
Wuhan, China
jyma2010@gmail.com

Abstract

Multimodal videos contain rich temporal cues that can facilitate faithful fusion. However, existing video fusion methods mainly rely on short-range temporal interactions, leaving valuable long-term information underexplored. Moreover, inconsistent fusion representations can cause flickering artifacts, undermining temporal coherence. To address these coupled challenges, we propose MemoryFusion, a two-stage framework integrating cross-temporal memory learning with progressive temporal refinement. Considering the complementary roles of adjacent and distant frames, the first stage combines a Short-Term Memory Bank for local temporal guidance with a Long-Term Memory Bank that preserves reliable historical information through a fusion-oriented update strategy. To mitigate temporal fluctuations, a Temporal Smoothing Module leverages soft motion cues for coarse background stabilization. Building upon the resulting fusion prior, The second stage further employs a lightweight 3D residual Temporal Refinement Module to correct remaining temporal inconsistencies while enhancing structural details. Extensive experiments demonstrate that MemoryFusion achieves superior fusion quality and temporal consistency over existing frame-based and video-based fusion methods. The source code is publicly available at https://github.com/Meiqi-Gong/MemoryFusion.

1 Introduction

Due to the inherent limitations of sensing mechanisms and hardware, single-modal images often fail to provide reliable scene understanding, especially in complex environments [4, 6, 15, 26]. To address this issue, multimodal image fusion has emerged as a crucial technique that integrates complementary information from different modalities to produce more comprehensive and reliable representations [20, 8, 22]. Benefiting from its enhanced representational capacity, multimodal fusion has been successfully applied in a wide range of fields, such as autonomous driving [19, 9], intelligent surveillance [25], and medical diagnosis [2], among others [10, 11].

Early research on multi-modal fusion has primarily focused on frame-based design, where frame-based modality-specific features are typically combined to enhance detail and contrast. For instance, Ma *et al.* introduced SwinFusion [12], which employs a Swin Transformer backbone to jointly capture intra-modal and inter-modal spatial dependencies within a single frame, enabling more adaptive and content-aware frame-based fusion. After that, CDDFuse [27] employs Restormer blocks to extract shallow cross-modal frame-wise features, followed by a dual-branch Transformer-CNN extractor that captures global and local representations within each frame independently. Further, Text-DiFuse [24] introduces a text-modulated diffusion model into the frame-based fusion paradigm.

* Corresponding authors.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

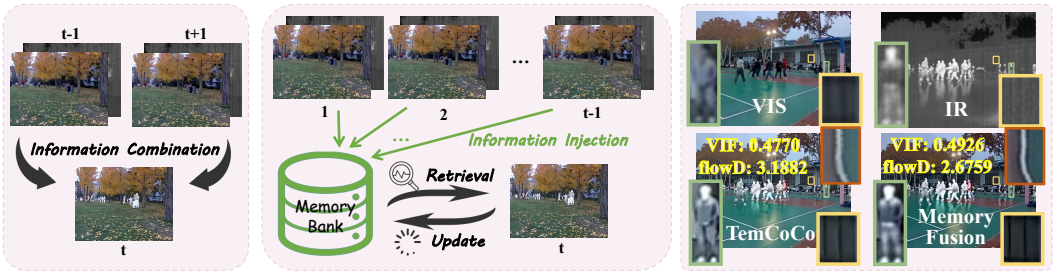

(a) Existing video fusion methods (b) Our proposed MemoryFusion (c) Results Comparison with TemCoCo

Figure 1: Comparison of existing video fusion methods and our proposed MemoryFusion. (a) and (b) illustrate the differences in how temporal information is exploited. (c) shows the corresponding visual comparison results, where the top-right region displays the results of stacked 40 consecutive frames in the same row for temporal consistency evaluation. Higher VIF means better fusion quality and lower flowD means better temporal consistency.

It embeds feature-level integration directly into the single-frame generative diffusion process, enabling simultaneous degradation removal and cross-modal information fusion within each frame. Recently, FC-Former [17] uses a fully-connected self-attention mechanism to capture local and non-local correlations within each single frame. It embeds a multi-source image representation as a non-local prior to enable effective single-frame feature fusion across modalities. Despite their remarkable success in image fusion, they are designed under single frame, whereas in practical scenarios, multimodal data are typically acquired in video form. This discrepancy prevents such methods from effectively modeling and leveraging inter-frame dependencies, thereby resulting in limited fusion performance and undesired temporal flickering when applied to practical video sequences.

Due to the growing recognition of the limitations of frame-based fusion, researchers have recently shifted their attention toward video fusion, which explicitly models and exploits temporal information across consecutive frames. By leveraging temporal correlations, video fusion enables a more comprehensive and stable fusion of multimodal data, resulting in significantly improved visual quality and reduced inter-frame flickering. For example, Gong *et al.* proposed TemCoCo [5], which aligns features from adjacent frames through deformable convolution (DCN) to capture local short-term correlations, while incorporating a temporal modeling module and a dedicated temporal consistency loss to ensure temporal consistency. In addition, Tang *et al.* introduced VideoFusion [16], which employs an attention mechanism to associate neighboring frames with the current one and maintains temporal coherence through explicit consistency constraints. Besides that, Zhao *et al.* proposed UniVF [29], leveraging neighboring-frame learning and optical flow-based feature warping to fuse adjacent frames. Also, they integrate a Transformer-based encoder-decoder with a temporal consistency loss to enhance temporal coherence. Overall, the development of video fusion methods is still at a preliminary stage. Existing video fusion methods primarily focus on short-range correlations between adjacent frames, as concluded in Figure 1 (a), where information of adjacent frames is transferred through various strategies to assist current-frame fusion. This short-term dependency limits the full exploitation of accessible temporal context, thus constraining fusion fidelity. Moreover, their temporal consistency modeling is confined to complex networks with merely simple temporal constraints and no incorporation of temporal or spatial prior, which often results in limited smoothness and the loss of fine spatial details.

Beyond the video fusion literature, memory mechanisms have also been investigated in broader video processing tasks. For example, memory banks are commonly used in video object segmentation to establish long-term correspondence for mask propagation [30, 3], while motion-adaptive temporal filtering and 3D convolution have been explored to suppress temporal fluctuations and capture local spatio-temporal dependencies [13, 23]. These studies demonstrate the potential of temporal memory and temporal filtering for video understanding. Inspired by these temporal modeling paradigms, we redesign them under the unique characteristics of multimodal video fusion.

To address the above limitation, we propose MemoryFusion, a two-stage multimodal video fusion method to explicitly tackle the challenges of underutilized temporal context and inadequate temporal consistency. On the one hand, we fully exploit cross-temporal learning through dual memory

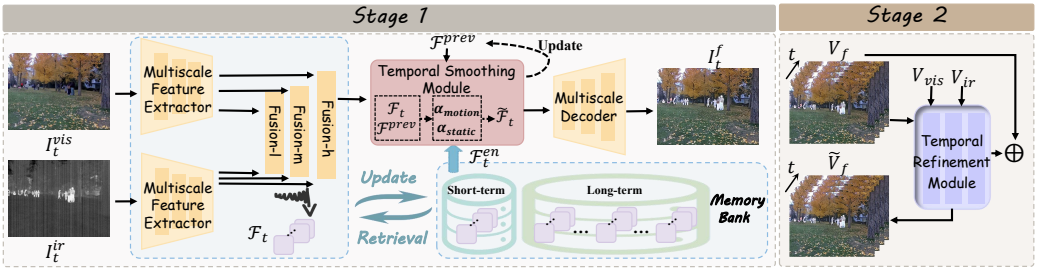

Figure 2: The overall framework of the proposed MemoryFusion.

banks in the first stage. Specifically, the Short-Term Memory Bank (STMB) aggregates informative cues from recent frames to provide local temporal guidance, while the Long-Term Memory Bank (LTMB) adaptively captures and preserves key representations from long-term frames to supply global temporal guidance. As shown in Figure 1 (b), unlike conventional methods that rely solely on short-range temporal cues, our MemoryFusion also explicitly maintains long-term temporal information through dedicated memory banks, enabling the network to retrieve more comprehensive information. On the other hand, to alleviate temporal fluctuations introduced during cross-frame feature aggregation, a Temporal Smoothing Module (TSM) is incorporated at the end of the first stage to perform region-adaptive coarse temporal smoothing by selectively preserving background information from the preceding frame. Building upon the resulting fusion prior, the second stage incorporates a Temporal Refinement Module (TRM) to guarantee fine temporal coherence. Treating the first-stage output as the initial results, the TRM leverages a 3D convolutional network to explicitly model and enforce temporal constraints, achieving a holistic optimization. Through this progressive coarse-to-fine design, MemoryFusion effectively integrates historical information while maintaining temporal stability. As illustrated in Figure 1(c), our MemoryFusion not only restores structural details that remain blurred in results from other video fusion methods such as TemCoCo, but also preserves temporal coherence across consecutive frames, effectively suppressing flickering artifacts and ensuring stable visual dynamics.

Our main contributions of this work can be summarized as:

- We propose MemoryFusion, a cross-temporal memory learning for multimodal video fusion that jointly enhances visual quality and temporal consistency.
- Dual memory banks are designed with a redundancy-aware and fusion-quality-aware LTMB update strategy to preserve diverse and reliable historical information.
- We develop a two-fold progressive temporal modeling approach, a TSM alleviating coarse background transitions, and a TRM enforcing fine consistency while refining details.
- Extensive experiments on public datasets demonstrate that MemoryFusion achieves superior performance in terms of fusion quality and temporal consistency.

2 Method

This section details the proposed two-stage MemoryFusion, as illustrated in Figure 2. The first stage performs memory-guided fusion with coarse temporal consistency modeling, while the second stage refines temporal consistency based on the fusion prior from the first stage.

2.1 Notation and Overview

In the first stage, dual memory banks are designed to store information from multiple temporal scales, and the framework primarily focuses on exploiting historical information stored in the memory bank to facilitate enhanced video fusion. Taking the t-th frame as an example, the visible frame $I_t^{\text{vis}} \in \mathbb{R}^{H\times W\times 3}$ and the infrared frame $I_t^{\text{ir}} \in \mathbb{R}^{H\times W\times 1}$ are first fed into two Multiscale Feature Encoders $\mathcal{E}_{\text{vis}}, \mathcal{E}_{\text{ir}}$ to extract hierarchical representations by:

$$\mathcal{F}_{t,l}^{\text{vis}}, \mathcal{F}_{t,m}^{\text{vis}}, \mathcal{F}_{t,h}^{\text{vis}} = \mathcal{E}_{\text{vis}}(I_t^{\text{vis}}), \quad \mathcal{F}_{t,l}^{\text{ir}}, \mathcal{F}_{t,m}^{\text{ir}}, \mathcal{F}_{t,h}^{\text{ir}} = \mathcal{E}_{\text{ir}}(I_t^{\text{ir}}), \tag{1}$$

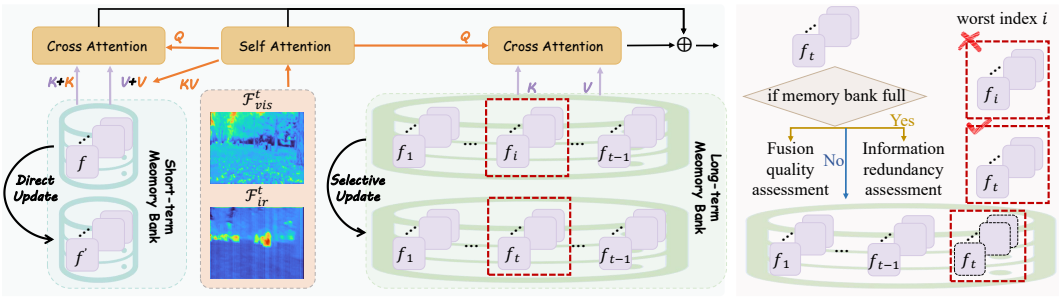

(a) Memory-based Enhancement (b) Update Rules of LTMB

Figure 3: Implementation of memory banks.

where $\mathcal{F}_{t,l}^{\text{vis}}, \mathcal{F}_{t,l}^{\text{ir}} \in \mathbb{R}^{\frac{H}{4}\times\frac{W}{4}\times C_1}$, $\mathcal{F}_{t,m}^{\text{vis}}, \mathcal{F}_{t,m}^{\text{ir}} \in \mathbb{R}^{\frac{H}{2}\times\frac{W}{2}\times C_2}$ and $\mathcal{F}_{t,h}^{\text{vis}}, \mathcal{F}_{t,h}^{\text{ir}} \in \mathbb{R}^{H\times W\times C_3}$. The corresponding multi-modal features are then fused through multiple feature fusion layers to obtain preliminary fused features $\mathcal{F}_{t,l}^{f}, \mathcal{F}_{t,m}^{f}$ and $\mathcal{F}_{t,h}^{f}$. These fused features interact with the memory banks through a retrieval process, where the network selectively extracts the relevant information from the memory banks to support the current fusion. The memory banks are dynamically updated following distinct strategies: the Short-term Memory Bank (STMB) is refreshed at every frame, while the Long-term Memory Bank (LTMB) is updated only when its capacity limit is reached, following a specific update rule that will be detailed in Section 2.2. Subsequently, the features are passed to the Temporal Smoothing Module (TSM), which integrates temporal information by incorporating the previous features $\mathcal{F}^{\text{prev}}$ at corresponding network layers to get the coarse temporal-smoothed feature $\mathcal{F}_t^{sm}$. Finally, the memory-enhanced features $\mathcal{F}_t^{\text{en}}$ and the smoothed fusion features $\mathcal{F}_t^{\text{sm}}$ are fed into a Multiscale Decoder $\mathcal{D}$ to reconstruct the fused frame I_t^f:

$$I_t^f = \mathcal{D}(\mathcal{F}_t^{\text{en}}, \mathcal{F}_t^{\text{sm}}). \tag{2}$$

Except for the first frame ($\mathcal{F}^{prev}$ is absent), all subsequent frames follow the same processing pipeline. After the first stage, the entire fused video sequence V_f, together with the original input video sequences V_{vis} and V_{ir}, is fed into the second-stage Temporal Refinement Module (TRM) to further enhance global temporal consistency across the video.

2.2 Memory Banks

For each frame t, after obtaining the preliminary fused feature $\mathcal{F}_t$ by fusion layers, we retrieve relevant information from a Short-Term Memory Bank (STMB) and a Long-Term Memory Bank (LTMB) to implement cross-temporal enhancement, as illustrated in Figure 3 (a).

Short-term Memory Bank (STMB) The STMB stores enhanced features from recent frames to provide short-term temporal context for the current fused feature. Since recent frames usually contain locally correlated structures and textures, the short-term memory can provide complementary cues for enhancing dynamic details. Different from directly adding past-frame features to the current representation, our STMB performs current-anchored memory retrieval. Specifically, the current fused feature $\mathcal{F}_t$ provides the query $\mathbf{Q}_c$, as well as the current key and value $(\mathbf{K}_c, \mathbf{V}_c)$, while the STMB provides memory keys and values $(\mathbf{K}_s, \mathbf{V}_s)$. Then the short-term enhanced output is computed via multi-head attention:

$$\mathbf{tgt}_s = \text{Attention}(\mathbf{Q}_c, \mathbf{K}_c + \mathbf{K}_s, \mathbf{V}_c + \mathbf{V}_s). \tag{3}$$

Long-term Memory Bank (LTMB) The LTMB stores representative features from wider-range frames to provide long-term contextual guidance for the current frame. Similar to STMB, the LTMB also adopts current-anchored memory retrieval, where the current fused feature provides the query $\mathbf{Q}_c$, and the LTMB provides memory keys and values $(\mathbf{K}_l, \mathbf{V}_l)$. Different from STMB, the LTMB focuses on retrieving representative long-range contextual cues from temporally distant frames, and is updated with a selective strategy to avoid being dominated by redundant or low-quality historical features. The long-term enhanced output is computed as:

$$\mathbf{tgt}_l = \text{Attention}(\mathbf{Q}_c, \mathbf{K}_l, \mathbf{V}_l). \tag{4}$$

After both short-term and long-term enhancement, the final enhanced feature is $\mathcal{F}_t^{\text{en}} = \mathcal{F}_t + \mathbf{tgt}_s + \mathbf{tgt}_l$. The enhanced feature $\mathcal{F}_t^{\text{en}}$ is stored in STMB for the next frame enhancement. For LTMB, new features are directly added until the maximum length is reached, and we follow previous work to set the maximum length as 8 [30]. Otherwise, LTMB follows update rules based on two criteria when receiving a new frame, detailed as follows.

Update Rules Although the LTMB effectively stores historical cues for enhancement, its capacity must be restricted to avoid redundant memory accumulation and outdated information. As illustrated in Figure 3 (b), we introduce a dual-criterion pruning mechanism that evaluates each memory entry from two complementary perspectives: *information redundancy*, which measures whether a memory entry can be well reconstructed by the remaining entries, and *fusion quality*, which reflects its ability to perform enhancement and fusion. This design encourages the LTMB to retain diverse and reliable memories, while removing entries that are either redundant or associated with poor fusion quality.

Specifically, when the number of long-term memory entries exceeds the predefined capacity, we perform memory pruning for each batch independently. Given N long-term memory entries $\{\mathbf{M}_i\}_{i=1}^{N}$, we first reshape and normalize them along the feature dimension to obtain $\mathbf{M} \in \mathbb{R}^{N\times d}$. For each memory entry $\mathbf{M}_i$, we use the remaining entries $\mathbf{M}_{-i}$ to linearly reconstruct it by solving a least-squares problem:

$$\mathbf{w}_i^* = \arg\min_{\mathbf{w}_i} \left\|\mathbf{M}_{-i}^\top \mathbf{w}_i - \mathbf{M}_i\right\|_2^2, \tag{5}$$

where $\mathbf{M}_{-i}$ denotes all memory entries except $\mathbf{M}_i$. The reconstructed memory is then obtained as $\hat{\mathbf{M}}_i = \mathbf{M}_{-i}^\top \mathbf{w}_i^*$. The reconstruction error is calculated by $e_i = \left\|\hat{\mathbf{M}}_i - \mathbf{M}_i\right\|_2^2$. A smaller reconstruction error indicates that $\mathbf{M}_i$ can be well represented by other memory entries and is therefore more redundant. Accordingly, we define the reconstruction-based redundancy score as:

$$r_i = 1 - \frac{e_i}{\sum_{j=1}^{N} e_j + \epsilon}. \tag{6}$$

In addition to redundancy, we also consider the enhancement and fusion quality. Given the image-level loss $\mathcal{L}_i$ corresponding to the frame stored in the LTMB, we define the quality score as:

$$q_i = \frac{\mathcal{L}_i}{\sum_{j=1}^{N} \mathcal{L}_j + \epsilon}, \tag{7}$$

where a larger q_i indicates poorer fusion quality and thus a higher priority for removal. Finally, the overall removal score is computed as

$$s_i = r_i + q_i. \tag{8}$$

The memory entry with the maximum score $i^* = \arg\max_i s_i$ is removed, and the LTMB is updated by discarding $\mathbf{M}_{i^*}$ while keeping the remaining memory entries. In this way, the LTMB preserves memory entries that are both unique and beneficial to fusion quality, thereby maintaining a compact yet informative long-term representation.

2.3 Temporal Smoothing Module (TSM)

To prevent obvious background variations between consecutive frames caused by the integration of differing memory bank features, we introduce a soft adaptive Temporal Smoothing Module (TSM) to implement coarse temporal consistency, which adaptively balances the reliance on current and previous frame information across motion and static regions. Specifically, TSM first estimates a soft motion mask by analyzing inter-frame variations in the original video sequence. Since the infrared modality is often contaminated by noise, which may be erroneously interpreted as motion, the motion estimation is instead performed using the visible video features. Given two consecutive visible-frame features, the difference map is computed as: $D_t = \|\mathcal{F}_t^{\text{vis}} - \mathcal{F}_{t-1}^{\text{vis}}\|_1$, where $\|\cdot\|_1$ denotes the l_1 norm. After normalization, a soft motion probability map M_t is obtained via a sigmoid activation: $M_t = \sigma(norm(D_t))$. To further suppress boundary artifacts and enhance spatial smoothness, the motion map is refined through mean-kernel blurring: $\tilde{M}_t = K * M_t$, where K denotes a mean filter.

Using the resulting soft mask $\tilde{M}_t$, the adaptive smoothing coefficient is constructed as:

$$\alpha_t = \alpha_{\text{motion}} \cdot \tilde{M}_t + \alpha_{\text{static}} \cdot (1 - \tilde{M}_t),$$

where α_{motion} and α_{static} control the degree of temporal filtering for motion and static regions, respectively. Finally, the feature output of the current frame is given by:

$$\tilde{\mathcal{F}}_t = \alpha_t \odot \mathcal{F}_t + (1 - \alpha_t) \odot \mathcal{F}^{\text{prev}},$$

where $\tilde{\mathcal{F}}_t$ will serve as the previous frame feature $\mathcal{F}^{\text{prev}}$ in the computation of the next frame.

The whole first stage is constrained by temporal-level consistency loss $\mathcal{L}_{\text{t}}$, and image-level quality loss $\mathcal{L}_{\text{grad}}$, $\mathcal{L}_{\text{color}}$ and $\mathcal{L}_{\text{content}}$ following [5], concluded as:

$$\mathcal{L} = \mathcal{L}_{\text{t}} + \lambda_1 \mathcal{L}_{\text{grad}} + \lambda_2 \mathcal{L}_{\text{color}} + \lambda_3 \mathcal{L}_{\text{content}}. \tag{9}$$

2.4 Temporal Refinement Module (TRM)

Although the previous TSM coarsely enhances temporal consistency, the overall temporal coherence of the network remains largely constrained by the 2D convolution-based encoder-decoder and other spatially oriented structures, which are inherently limited in modeling temporal dynamics. Therefore, we design the second-stage Temporal Refinement Module (TRM) to further refine temporal consistency in a fine-grained and efficient manner.

We assume that the model has already learned high-quality fused feature representations in the first stage. With this strong fusion prior, the TRM is deliberately kept lightweight, consisting of only several stacked 3D separable convolutions to capture subtle temporal variations without excessive computational overhead. To further uncover the informative details overlooked in the first stage, both the fused outputs V_f and the original visible and infrared sequences (V_{vis}, V_{ir}) are fed into this module. A residual learning strategy is employed to inherit and refine the performance of the first stage, leading to a stable and progressive improvement in temporal coherence.

This process can be mathematically formulated as:

$$\tilde{V}_f = TRM(V_f, V_{vis}, V_{ir}) + V_f, \tag{10}$$

where $TRM(\cdot)$ denotes the 3D convolution-based refinement network. In this stage, we adopt an optimization objective composed of a gradient loss $\mathcal{L}_{\text{grad}}$ to preserve fine spatial textures and a temporal loss $\mathcal{L}_{\text{t}}$ to enforce inter-frame consistency. This supervision enables the TRM to recover spatio-temporal dynamics overlooked by the first-stage fusion process.

3 Experiments

In this section, we first introduce the experimental setup, including the comparison methods, evaluation metrics, hyperparameters, and other implementation details. Subsequently, we present the quantitative and qualitative results on various benchmark datasets. Finally, we perform ablation studies to demonstrate the effectiveness of each designed module in our method.

3.1 Implementation Details

Datasets and Metrics We conduct experiments on two infrared and visible video datasets, namely M3SVD [16] and HDO [18], both of which contain pre-registered video pairs. The M3SVD dataset consists of 220 pairs of video sequences with frame counts ranging from 100 to 1000, while the HDO dataset includes 26 pairs of sequences, each containing 150 to 200 frames. Our method is trained only on the M3SVD dataset and tested on both datasets. To comprehensively evaluate the model's performance, we employ MI, CC, SCD and VIF to assess the fusion quality, while flowD and feaCD [5] are used to measure the temporal consistency of the fused videos.

Comparison Methods Since video fusion methods are still in their early stage of development, we select five state-of-the-art image fusion algorithms and three video fusion methods for comparison. The image fusion algorithms include LRRNet [7], Diff-IF [21], MMIF-EMMA [28], Text-DiFuse [24], and TC-MoA [31], while the video fusion methods include TemCoCo [5], VideoFusion [16] and UniVF [29]. To ensure a fair comparison, for those fusion models lacking integrated de-degradation modules, we apply DSTNet [14] to deblur the visible videos and MDIVDNet [1] to denoise the infrared videos before fusion. Both preprocessing works have been fully retrained on our datasets. All comparative video fusion methods are evaluated using the standard number of frames specified in their official implementations.

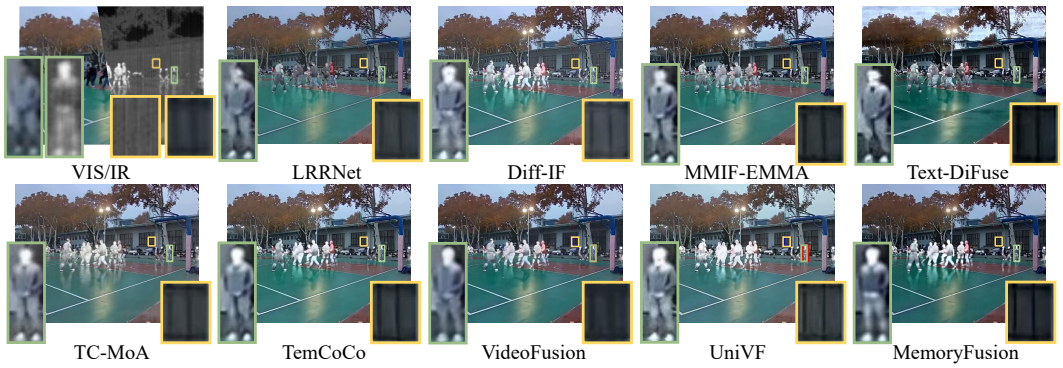

Figure 4: Qualitative results on M3SVD dataset under degradation.

Table 1: Quantitative results on the M3SVD and HDO datasets. **Bold** indicates the best result.

Method	M3SVD						HDO					
	MI↑	CC↑	SCD↑	VIF↑	flowD↓	feaCD↓	MI↑	CC↑	SCD↑	VIF↑	flowD↓	feaCD↓
LRRNet	1.7889	0.4584	1.4884	0.3906	5.5065	1.9309	1.0616	0.6725	0.7250	0.2416	7.4598	1.9523
Diff-IF	2.0692	0.4912	1.6068	0.4847	4.0890	1.8210	1.1387	0.6851	1.1483	0.3039	8.0450	1.9587
MMIF-EMMA	2.1175	0.4758	1.6443	0.4758	4.6592	1.8412	1.1668	0.6945	1.2582	0.3071	7.9598	1.9565
Text-DiFuse	1.6253	0.4576	1.4989	0.3514	8.6578	1.9744	1.1682	0.7024	1.4096	0.2940	9.7262	1.9732
TC-MoA	2.0178	0.4660	1.5276	0.4054	4.1954	1.8118	1.0942	0.7023	1.1639	0.3181	7.8125	1.9570
TemCoCo	2.1717	0.4800	1.5961	0.4770	3.1882	1.8058	1.3212	0.7192	1.3689	0.3823	5.1147	1.7336
VideoFusion	2.0353	0.4612	1.5299	0.4126	3.4161	1.8154	1.1428	0.7106	1.3152	0.3238	6.0564	1.8643
UniVF	2.1316	0.4848	1.6122	0.4677	4.0271	1.8138	1.1854	0.7205	1.3266	0.3427	6.7987	1.8680
Ours-stage1	2.1631	0.4872	1.5975	0.4877	3.6135	1.8088	1.3261	0.7319	1.4002	0.3733	6.6011	1.7747
MemoryFusion	**2.2350**	**0.4952**	**1.6152**	**0.4926**	**2.6759**	**1.7640**	**1.3264**	**0.7347**	**1.4450**	**0.3947**	**4.9188**	**1.7293**

Experimental Settings We empirically set the hyperparameters as $\alpha_{motion} = 0.7, \alpha_{static} = 0.3, \lambda_1 = 5, \lambda_2 = 10$ and $\lambda_3 = 5$. During training, the batch size is set to 2, and each batch contains 30 frames. During testing, 200 frames are used for each video sequence. All experiments are conducted on the PyTorch platform using a 2.0 GHz Intel Xeon Gold 6330 CPU and an NVIDIA A800 80 GB PCIe GPU.

3.2 Video Fusion

In this part, we present the experimental results on M3SVD and HDO datasets, including both qualitative and quantitative results, to comprehensively demonstrate the effectiveness of MemoryFusion.

Results on M3SVD Figure 4 compares different methods on the M3SVD dataset. As shown, the proposed MemoryFusion exhibits clear advantages in both detail restoration and contrast preservation. Specifically, frame-based fusion methods fail to preserve complete and realistic human structures or suffer from noticeable detail blurring. Among video-based methods, TemCoCo and VideoFusion remain limited by low contrast or misalignment artifacts, while UniVF, despite sharper structures, is still inferior. In the yellow box at the lower right, our method reconstructs window structures with the highest contrast among all competitors. Table 1 reports the quantitative results on M3SVD. Our method achieves the best performance on all image quality metrics (MI, CC, SCD, VIF) as well as temporal consistency metrics (flowD, feaCD), demonstrating superior spatial fidelity and temporal coherence. While TemCoCo and VideoFusion show relatively strong temporal consistency due to video-sequence modeling, UniVF exhibits weaker temporal coherence without such mechanisms. Notably, our first-stage model already surpasses image-based methods due to temporal consistency loss and TSM, though frame-wise processing limits its temporal performance. With the lightweight second stage, temporal coherence is effectively restored and image quality is further improved.

Results on HDO Figure 5 presents the qualitative results on the HDO dataset. In the enlarged yellow box at the upper right, similar to the previous results, all compared methods except TemCoCo and our

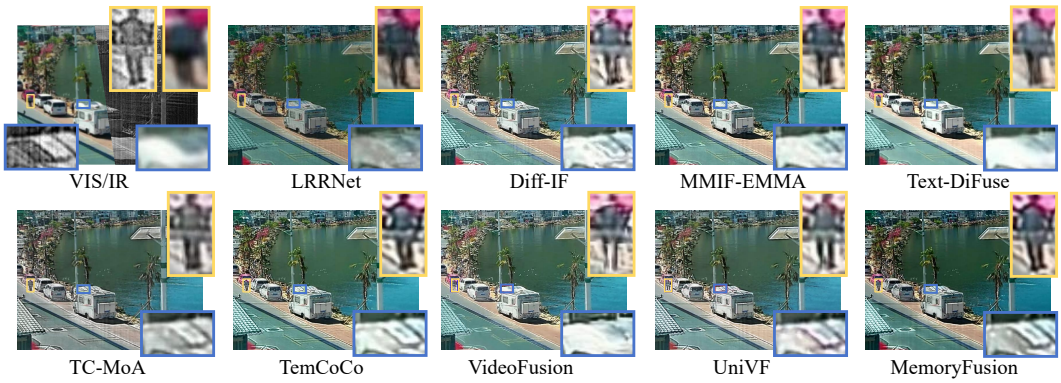

Figure 5: Qualitative results on HDO dataset under degradation.

proposed MemoryFusion fail to preserve the complete human structure, where TC-MoA and UniVF exhibit noticeable color distortion, and Diff-IF, MMIF-EMMA and VideoFusion produce blurred or fragmented contours. In contrast, both TemCoCo and our approach maintain more coherent human outlines, yet TemCoCo still suffers from residual blur and local detail loss. Beyond the structural regions, the enlarged blue box at the lower right highlights differences in fine texture reconstruction. Our MemoryFusion yields sharper edges and more faithful contrast, effectively restoring spatial details that are heavily degraded in other methods. The quantitative results summarized in Table 1 further confirm these observations, showing that MemoryFusion consistently achieves superior scores across all metrics, validating its robustness and effectiveness under severe degradation conditions.

Temporal Consistency Comparison To evaluate temporal consistency, we concatenate the same pixel row from 60 consecutive frames for visualization, as shown in Figure 6. Existing methods suffer from temporal artifacts. For instance, LRRNet and Text-DiFuse show ghosting and tearing effects, while Diff-IF and MMIF-EMMA exhibit structural discontinuities. Although TemCoCo, VideoFusion, and UniVF reduce some distortions via frame alignment, they still struggle with smooth temporal transitions. In contrast, MemoryFusion produces temporally coherent and visually stable results with well-preserved structures, demonstrating its effectiveness in modeling temporal dependencies and maintaining spatial-temporal stability.

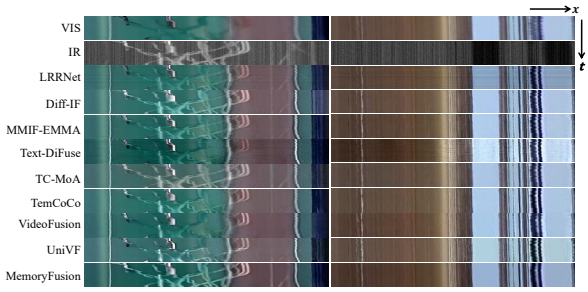

Figure 6: Temporal consistency comparison.

3.3 Ablation Studies

In this section, we conduct comprehensive ablation studies to validate the effectiveness of each designed component in MemoryFusion. Specifically, we provide qualitative evaluations from both image quality and temporal consistency perspectives, as shown in Figure 7, and report the corresponding quantitative results in Table 2. It is worth noting that, except for the two variants related to the second refinement stage (*i.e.*, w/o stage2 tloss and w/o stage2 Conv3D), all other ablation settings are evaluated only within the first stage. Therefore, these variants are compared with the stage-1 baseline rather than the complete MemoryFusion model, so as to more accurately reflect the contribution of each stage-1 component.

w/o TSM This setting removes the Temporal Smoothing Module (TSM). The TSM leverages the previous frame to smooth the current frame, reducing obvious inconsistency between consecutive frames. As shown in Figure 7 and Table 2, omitting this module has little effect on image quality but results in temporal flickering, confirming its role in enhancing temporal smoothness. This indicates that TSM mainly contributes to inter-frame stability.

w/o STMB This experiment disables the Short-Term Memory Bank (STMB). Without it, the model cannot effectively reuse short-term contextual cues from recent frames to guide current-frame fusion.

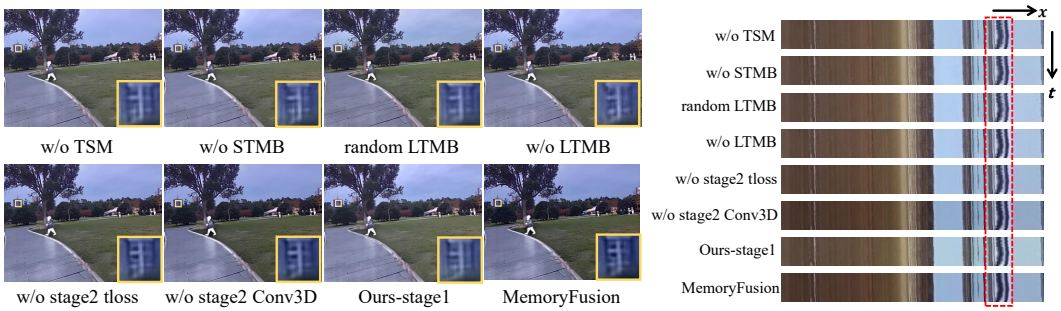

Figure 7: Qualitative results of ablation experiments.

As shown in the figure, the result shows blurred textures and weakened local contrast. These results confirm that STMB plays a vital role in reinforcing details during enhancement and fusion. This suggests that recent temporal information is beneficial for preserving local structures and fine textures.

Random LTMB In this configuration, the Long-Term Memory Bank (LTMB) is updated randomly rather than by the proposed dual-criterion rule. Such stochastic replacement disrupts the accumulation of informative long-term cues, leading to unstable global guidance. The results highlight the importance of an effective update mechanism for maintaining meaningful long-term memory. This further shows that the benefit of LTMB comes not only from memory storage, but also from selective memory updating.

w/o LTMB This variant removes the LTMB to examine the contribution of global memory aggregation. Without access to long-term frame information, the model is limited to short-range temporal cues from the STMB and fails to capture global contextual relations and scene structures. As a result, the fused results exhibit degraded structural integrity and weakened detail representation, especially in complex or spatially diverse regions, demonstrating the effectiveness of LTMB. This verifies that long-term memory provides complementary global cues beyond short-term temporal memory.

w/o stage2 tloss & w/o stage2 Conv3D These two variants evaluate the temporal loss and 3D convolution in the second refinement stage. Removing the temporal loss weakens the explicit supervision on inter-frame consistency, leading to less stable transitions. For the w/o stage2 Conv3D variant, we replace the 3D convolution with 2D convolution, which preserves spatial refinement but removes explicit temporal modeling across adjacent frames. As shown in Table 2, 3D convolution plays a more decisive role in maintaining temporal stability, demonstrating that the joint spatio-temporal refinement in the second stage is crucial for stabilizing video fusion results.

Table 2: Quantitative results of ablation experiments.

M3SVD	MI↑	CC↑	SCD↑	VIF↑	flowD↓	feaCD↓
w/o TSM	2.0173	0.4821	1.5734	**0.4943**	4.2172	1.8038
w/o STMB	2.0239	0.4767	1.5538	0.4817	3.6586	1.8061
random LTMB	2.0199	0.4755	1.5580	0.4825	3.7010	1.8036
w/o LTMB	2.0031	0.4646	1.5509	0.4770	3.7768	1.8078
w/o stage2 tloss	2.1987	0.4884	1.5961	0.4932	3.2065	1.7985
w/o stage2 Conv3D	2.1007	0.4942	1.5994	0.4938	3.5290	1.7956
Ours-stage1	2.1631	0.4872	1.5975	0.4877	3.6135	1.8088
MemoryFusion	**2.2350**	**0.4952**	**1.6152**	0.4926	**2.6759**	**1.7640**

3.4 Plug-and-Play Spatiotemporal Refinement with TRM

To further investigate the generalization and effectiveness of the proposed TRM, we apply it to two existing fusion methods, MMIF-EMMA and UniVF. Specifically, the original fusion models are kept fixed, and their fused outputs are directly used as fusion priors for TRM. This allows TRM to perform spatiotemporal refinement without modifying the original fusion networks, demonstrating its flexibility in enhancing existing fusion methods.

As shown in Table 3, introducing TRM significantly improves the temporal consistency of both methods, achieving relative reductions in flowD of 36.6% and 30.4% for MMIF-EMMA and UniVF, respectively. In addition to these substantial improvements in temporal consistency, TRM also enhances fusion quality in terms of CC and SCD for both methods, while maintaining comparable performance on other fusion metrics. These results demonstrate that TRM can effectively enhance

temporal consistency while preserving or improving spatial fusion quality, even when applied to fusion models without memory mechanisms.

Furthermore, the complete MemoryFusion achieves the best temporal consistency and superior overall fusion performance, suggesting that the memory-enhanced Stage-1 fusion prior provides a favorable foundation for subsequent refinement. More importantly, TRM achieves substantial improvements when directly applied to other fusion methods without requiring additional training or finetuning, highlighting its plug-and-play capability. These results demonstrate the general applicability of TRM as a flexible spatiotemporal refinement module that can effectively enhance fusion quality and, more importantly, temporal consistency, with the potential for further improvements through finetuning.

Table 3: Quantitative comparison of applying TRM to different fusion methods. The best results are highlighted in bold.

Method	MI↑	CC↑	SCD↑	VIF↑	flowD↓	feaCD↓
MMIF-EMMA	2.1175	0.4758	1.6443	0.4758	4.6592	1.8412
MMIF-EMMA+TRM	2.0641	0.4932	**1.6624**	0.4762	2.9549	1.8125
UniVF	2.1316	0.4848	1.6122	0.4677	4.0271	1.8138
UniVF+TRM	2.0337	0.4921	1.6325	0.4673	2.8025	1.8122
Ours-stage1	2.1631	0.4872	1.5975	0.4877	3.6135	1.8088
MemoryFusion	**2.2350**	**0.4952**	1.6152	**0.4926**	**2.6759**	**1.7640**

3.5 Computation Efficiency Comparison

We compare the computational efficiency in terms of FLOPs, parameter count, and runtime in Table 4. Although MemoryFusion introduces an additional second stage, the overall computational cost remains moderate. Specifically, the first-stage model achieves competitive efficiency compared with existing video-based fusion methods, while the lightweight second stage incurs only negligible overhead. For preprocessing-based methods, the reported values only account for the computational cost of the fusion networks, excluding additional pretrained deblurring or denoising modules. Therefore, the comparison reflects the intrinsic efficiency of the fusion networks rather than the computational cost of the complete preprocessing-and-fusion pipelines. Overall, our full model strikes a favorable balance between fusion quality and computational efficiency, outperforming heavy diffusion-based or large-scale temporal models while maintaining comparable inference speed.

Table 4: Comparison of computational efficiency.

	LRRNet	Diff-IF	MMIF-EMMA	Text-DiFuse	TC-MoA	TemCoCo	VideoFusion	UniVF	Stage1	Stage2
FLOPs (G)	14.2	86.9	41.5	1268	4833	505.6	301.1	740.7	213.6	1.65
Param (M)	0.047	22.36	1.52	119.4	340	19.18	6.74	9.15	17.2	0.033
Time (s)	0.162	4.23	0.056	24.2	0.185	0.109	0.105	0.245	0.099	0.001

4 Conclusion

This paper proposes MemoryFusion, a two-stage cross-temporal multimodal video fusion framework designed to achieve high-fidelity fusion and superior temporal consistency. In the first stage, dual memory banks dynamically update and effectively retrieve informative cues from historical frames to enhance current-frame fusion. Additionally, a TSM is incorporated to adaptively mitigate coarse frame variations and establish a stable temporal foundation. In the second stage, guided by the powerful fusion prior from the first stage, a lightweight TRM further enforces fine-grained temporal coherence and enhances spatial fidelity through residual refinement. Extensive experiments on multiple public datasets demonstrate that MemoryFusion not only achieves state-of-the-art performance in both fusion quality and temporal consistency but also exhibits strong generalization across diverse scenes.

Acknowledgments and Disclosure of Funding

This work was supported by the National Natural Science Foundation of China under Grant no. 62676289, and the China Postdoctoral Science Foundation under no. 2026T190426. The numerical calculations in this paper have been done on the supercomputing system in the Supercomputing Center of Wuhan University. The authors declare no competing interests.

References

[1] Lijing Cai, Xiangyu Dong, Kailai Zhou, and Xun Cao. Exploring video denoising in thermal infrared imaging: Physics-inspired noise generator, dataset and model. *IEEE Transactions on Image Processing*, 33:3839–3854, 2024.

[2] Wanyi Chen, Zihua Zhao, Jiangchao Yao, Ya Zhang, Jiajun Bu, and Haishuai Wang. Multi-modal medical diagnosis via large-small model collaboration. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 30763–30773, 2025.

[3] Ho Kei Cheng and Alexander G Schwing. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *European conference on computer vision*, pages 640–658. Springer, 2022.

[4] Meiqi Gong, Hao Zhang, Hebaixu Wang, Jun Chen, Jun Huang, Xin Tian, and Jiayi Ma. Cross-scale domain adaptation with comprehensive information for pansharpening. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 785–793, 2024.

[5] Meiqi Gong, Hao Zhang, Xunpeng Yi, Linfeng Tang, and Jiayi Ma. Temcoco: Temporally consistent multi-modal video fusion with visual-semantic collaboration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14326–14335, 2025.

[6] Meiqi Gong, Hao Zhang, Bingwei Hui, and Jiayi Ma. Transfm: Visible-to-infrared image translation via flow matching. *IEEE/CAA Journal of Automatica Sinica*, 13(5):1239–1241, 2026.

[7] Hui Li, Tianyang Xu, Xiao-Jun Wu, Jiwen Lu, and Josef Kittler. Lrrnet: A novel representation learning guided fusion network for infrared and visible images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):11040–11052, 2023.

[8] Xuan Li, Rongfu Chen, Jie Wang, Lei Ma, Li Cheng, and Haiwen Yuan. Dstcfuse: A method based on dual-cycled cross-awareness of structure tensor for semantic segmentation via infrared and visible image fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5558–5567, 2024.

[9] Jinyuan Liu, Zhu Liu, Guanyao Wu, Long Ma, Risheng Liu, Wei Zhong, Zhongxuan Luo, and Xin Fan. Multi-interactive feature learning and a full-time multi-modality benchmark for image fusion and segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8115–8124, 2023.

[10] Jinyuan Liu, Runjia Lin, Guanyao Wu, Risheng Liu, Zhongxuan Luo, and Xin Fan. Coconet: Coupled contrastive learning network with multi-level feature ensemble for multi-modality image fusion. *International Journal of Computer Vision*, 132(5):1748–1775, 2024.

[11] Fuya Luo, Yunhan Li, Guang Zeng, Peng Peng, Gang Wang, and Yongjie Li. Thermal infrared image colorization for nighttime driving scenes with top-down guided attention. *IEEE Transactions on Intelligent Transportation Systems*, 23(9):15808–15823, 2022.

[12] 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.

[13] Zhong-Zheng Peng, Yi-Xin Yang, Jin-Hui Tang, and Jin-Shan Pan. Video colorization: A survey. *Journal of Computer Science and Technology*, 39(3):487–508, 2024.

[14] Chen Rao, Guangyuan Li, Zehua Lan, Jiakai Sun, Junsheng Luan, Wei Xing, Lei Zhao, Huaizhong Lin, Jianfeng Dong, and Dalong Zhang. Rethinking video deblurring with wavelet-aware dynamic transformer and diffusion model. In *Proceedings of the European Conference on Computer Vision*, pages 421–437, 2024.

[15] Linfeng Tang, Chunyu Li, Yeda Wang, Guoqing Wang, Yixuan Yuan, and Jiayi Ma. Dspfusion: Image fusion via degradation and semantic dual-prior guidance. *IEEE Transactions on Image Processing*, 35: 6331–6345, 2026. doi: 10.1109/TIP.2026.3700938.

[16] Linfeng Tang, Yeda Wang, Meiqi Gong, Zizhuo Li, Yuxin Deng, Xunpeng Yi, Chunyu Li, Han Xu, Hao Zhang, and Jiayi Ma. Videofusion: A spatio-temporal collaborative network for multi-modal video fusion and restoration. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2026.

[17] Xiao Wu, Zi-Han Cao, Ting-Zhu Huang, Liang-Jian Deng, Jocelyn Chanussot, and Gemine Vivone. Fully-connected transformer for multi-source image fusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(3):2071–2088, 2025.

[18] Housheng Xie, Meng Sang, Yukuan Zhang, Yang Yang, Shan Zhao, and Jianbo Zhong. Rcvs: A unified registration and fusion framework for video streams. *IEEE Transactions on Multimedia*, 26:11031–11043, 2024.

[19] Jun Xie, Jiazhen Dou, Liyun Zhong, Jianglei Di, and Yuwen Qin. A dual-mode intensity and polarized imaging system for assisting autonomous driving. *IEEE Transactions on Instrumentation and Measurement*, 73:1–13, 2024.

[20] Han Xu, Meiqi Gong, Xin Tian, Jun Huang, and Jiayi Ma. Cufd: An encoder–decoder network for visible and infrared image fusion based on common and unique feature decomposition. *Computer Vision and Image Understanding*, 218:103407, 2022.

[21] Xunpeng Yi, Linfeng Tang, Hao Zhang, Han Xu, and Jiayi Ma. Diff-if: Multi-modality image fusion via diffusion model with fusion knowledge prior. *Information Fusion*, 110:102450, 2024.

[22] Xunpeng Yi, Qinglong Yan, Yibing Zhang, Zongrong Li, Hao Zhang, Zhanchuan Cai, and Jiayi Ma. Ll-refiner: Learning adaptive refinement for ultra-high-definition low-light image enhancement. *IEEE/CAA Journal of Automatica Sinica*, 13(6):1288–1300, 2026.

[23] Xinyi Ying, Longguang Wang, Yingqian Wang, Weidong Sheng, Wei An, and Yulan Guo. Deformable 3d convolution for video super-resolution. *IEEE Signal Processing Letters*, 27:1500–1504, 2020.

[24] Hao Zhang, Lei Cao, and Jiayi Ma. Text-difuse: An interactive multi-modal image fusion framework based on text-modulated diffusion model. *Advances in Neural Information Processing Systems*, 37:39552–39572, 2024.

[25] Hao Zhang, Linfeng Tang, Xinyu Xiang, Xuhui Zuo, and Jiayi Ma. Dispel darkness for better fusion: A controllable visual enhancer based on cross-modal conditional adversarial learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26487–26496, 2024.

[26] Hao Zhang, Lei Cao, Xuhui Zuo, Zhenfeng Shao, and Jiayi Ma. Omnifuse: Composite degradation-robust image fusion with language-driven semantics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(9):7577–7595, 2025.

[27] Zixiang Zhao, Haowen Bai, Jiangshe Zhang, Yulun Zhang, Shuang Xu, Zudi Lin, Radu Timofte, and Luc Van Gool. Cddfuse: Correlation-driven dual-branch feature decomposition for multi-modality image fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5906–5916, 2023.

[28] Zixiang Zhao, Haowen Bai, Jiangshe Zhang, Yulun Zhang, Kai Zhang, Shuang Xu, Dongdong Chen, Radu Timofte, and Luc Van Gool. Equivariant multi-modality image fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25912–25921, 2024.

[29] Zixiang Zhao, Haowen Bai, Bingxin Ke, Yukun Cui, Lilun Deng, Yulun Zhang, Kai Zhang, and Konrad Schindler. A unified solution to video fusion: From multi-frame learning to benchmarking. *Advances in Neural Information Processing Systems*, 2025.

[30] Junbao Zhou, Ziqi Pang, and Yu-Xiong Wang. Rmem: Restricted memory banks improve video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18602–18611, 2024.

[31] Pengfei Zhu, Yang Sun, Bing Cao, and Qinghua Hu. Task-customized mixture of adapters for general image fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7099–7108, 2024.

A Supplemental Material

A.1 Experimental Details

MemoryFusion directly takes the raw infrared-visible video sequences as input without using external de-degradation preprocessing. For baseline methods that do not contain degradation-removal components, we additionally apply DSTNet/MDIVDNet as preprocessing only for robustness evaluation under degraded inputs. Unless otherwise stated, the FLOPs, parameters, and inference time reported for these baseline fusion networks do not include the additional computational cost of DSTNet/MDIVDNet. Therefore, the efficiency comparison is favorable to the baselines, while MemoryFusion remains competitive without relying on external preprocessing.

A.2 Extension to Medical Scenes

We further extend the proposed model to medical sequence fusion to evaluate its generalization capability to medical scenes. Specifically, experiments are conducted on the Harvard medical imaging dataset[0], where the MRI and PET sequences consist of a series of consecutive slice images acquired from cross-sectional scans of the same organ. Owing to the strong anatomical correlation between adjacent slices, the spatial continuity among them is analogous to the temporal coherence between consecutive frames in a video. Hence, we treat each MRI or PET slice as a video frame to evaluate the model's performance in medical sequence fusion. Since medical scenes are relatively homogeneous, this experiment primarily focuses on assessing the fusion performance. As shown in Figure s1, methods such as LRRNet, Text-DiFuse and VideoFusion tend to blur fine structural boundaries or over-smooth functional details, while Diff-IF and MMIF-EMMA exhibit noticeable inconsistencies between structural and metabolic regions. Despite their competence in detail preservation, TC-MoA and UniVF fail to retain the color information of PET images, resulting in severe spectral distortion. In contrast, our MemoryFusion framework effectively integrates structural information from MRI with metabolic patterns from PET, achieving both high structural fidelity and functional contrast. The highlighted regions further demonstrate that our approach preserves sharp anatomical edges and detailed texture in gray matter areas while maintaining accurate PET intensity distributions, thereby providing more clinically interpretable fused results.

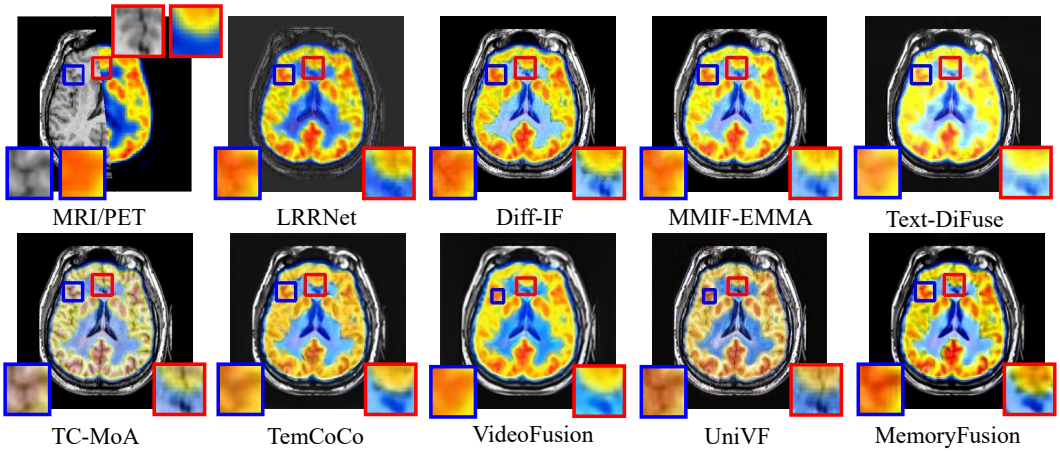

Figure s1: Results of medical sequence fusion on Harvard dataset.

A.3 Results of Non-degraded Scenes

Figure s2 shows the qualitative results on non-degraded scenes. MemoryFusion enhances visual contrast and salient structural cues. In the enlarged yellow box at the upper left, most methods fail to clearly separate the human from the background, while TemCoCo and VideoFusion introduce boundary artifacts due to over-enhancement. In contrast, MemoryFusion effectively incorporates thermal cues from the infrared modality, producing clearer human boundaries with well-preserved

[0] http://www.med.harvard.edu/AANLIB/

texture continuity. The enlarged green box at the lower right further compares texture reconstruction around the window region. Obviously, our MemoryFusion delivers sharper and more consistent structures with higher contrast. These results demonstrate that the proposed memory-enhanced modeling not only preserves global scene semantics but also improves fine-detail reconstruction, even under non-degraded conditions where other methods tend to over-smooth or distort details.

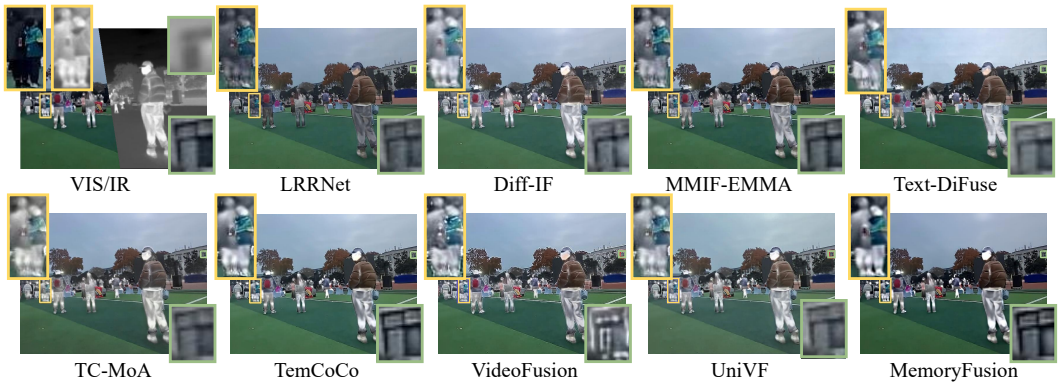

Figure s2: Qualitative results of non-degraded scene.

A.4 Limitation

Although MemoryFusion achieves strong fusion quality and temporal consistency, it is not explicitly designed as a degradation-specific restoration framework. Its robustness thus depends partly on the degradation coverage of the training data, and may decrease under severely unseen degradations. Moreover, since TSM uses visible-frame differences only as a lightweight temporal smoothing cue, when motion is mainly reflected in the infrared modality, the first-stage smoothing may provide limited guidance, placing greater pressure on the second-stage Temporal Refinement Module to restore temporal consistency. Nevertheless, these limitations are orthogonal to our core contribution, and degradation-aware modeling or reliability-aware bi-modal motion estimation can be naturally incorporated into MemoryFusion as additional modules in future work.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction state the main contributions and scope of the work, including cross-temporal memory learning and progressive temporal modeling. The claims are aligned with the reported quantitative, qualitative, and ablation results on the evaluated video fusion datasets.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper includes a separate limitation section discussing the scope of the proposed method, noting that MemoryFusion focuses on cross-temporal memory learning and temporal consistency modeling rather than degradation-specific restoration. It also clarifies that robustness to complex degradations depends on the diversity of degradation patterns covered in the training data.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

   Justification: This paper does not present theoretical results, theorems, or formal proofs. The proposed method is an empirical framework whose effectiveness is validated through quantitative comparisons, qualitative results, and ablation studies.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: The paper reports the key information needed to reproduce the main results, including the network architecture, memory update strategy, datasets, evaluation metrics, baselines, training hyperparameters, and implementation environment. The source code and pretrained models will be released to further support reproducibility.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: The code is included in the supplemental material with instructions for reproducing the main experimental results. The paper also describes the datasets, evaluation metrics, baselines, and implementation details needed for faithful reproduction.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: The paper specifies the main experimental settings, including the training and testing datasets, evaluation metrics, comparison methods, hyperparameters, batch size, training/testing frame settings, and implementation environment. Additional details are provided in the supplemental material and released code.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [N/A]

Justification: The paper reports deterministic quantitative results following the standard evaluation protocol in multimodal video fusion, but does not include error bars or statistical significance tests. However, the reported results are averaged over a large number of test frames/sequences from public datasets, which reduces sample-level variability and provides a stable empirical comparison.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: The paper reports the compute resources and execution time required for the experiments, including the CPU, GPU type, GPU memory, implementation platform, and running time for the proposed method. These details provide sufficient information to assess the computational cost and reproduce the reported experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research conforms to the NeurIPS Code of Ethics. It uses public datasets for multimodal video fusion and does not involve human subjects, private or sensitive personal data, or applications that raise specific ethical concerns.

   Guidelines:

- The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: The paper discusses the broader impacts of the proposed multimodal video fusion method, mainly including its potential benefits for robust perception in autonomous driving, intelligent monitoring, and robotic systems. Since this work focuses on low-level multimodal fusion using public benchmark datasets and does not involve personal data, decision-making, or user-facing deployment, we do not identify direct negative societal impacts.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: This work does not involve the release of high-risk models or datasets, such as large language models, image generators, or scraped datasets. The released code and pretrained models are for low-level multimodal video fusion on public benchmark datasets, so no additional safeguards are required.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.

- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: We properly cite the original papers of the datasets and baseline methods used in this work. All public datasets and existing assets are used only for academic research and are handled in accordance with their original licenses and terms of use.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: The paper introduces and releases new assets, including the implementation code and pretrained models of MemoryFusion. These assets are documented in the supplemental material with instructions for usage, reproduction, and evaluation.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: This work does not involve crowdsourcing experiments, human-subject studies, participant instructions, or compensation.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: This work does not involve crowdsourcing experiments, human-subject studies, study participants, or personally identifiable data. Therefore, IRB approval or equivalent review is not applicable.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: The core method development in this work does not involve LLMs as important, original, or non-standard components. Any possible use of LLMs is limited to writing, editing, or formatting assistance and does not affect the methodology, experiments, scientific rigor, or originality of the research.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.