

GaraMoSt: Parallel Multi-Granularity Motion and Structural Modeling for Efficient Multi-Frame Interpolation in DSA Images

Ziyang Xu¹, Huangxuan Zhao^{2,3,*}, Wenyu Liu¹, Xinggong Wang^{1,*}

¹Institute of AI, School of EIC, Huazhong University of Science and Technology

²School of Computer Science, Wuhan University

³Union Hospital, Tongji Medical College, Huazhong University of Science and Technology
{xuzyoung, liuwuy, xgwang}@hust.edu.cn, zhao_huangxuan@sina.com

Abstract

The rapid and accurate direct multi-frame interpolation method for Digital Subtraction Angiography (DSA) images is crucial for reducing radiation and providing real-time assistance to physicians for precise diagnostics and treatment. DSA images contain complex vascular structures and various motions. Applying natural scene Video Frame Interpolation (VFI) methods results in motion artifacts, structural dissipation, and blurriness. Recently, MoSt-DSA has specifically addressed these issues for the first time and achieved SOTA results. However, MoSt-DSA’s focus on real-time performance leads to insufficient suppression of high-frequency noise and incomplete filtering of low-frequency noise in the generated images. To address these issues within the same computational time scale, we propose GaraMoSt. Specifically, we optimize the network pipeline with a parallel design and propose a module named MG-MSFE. MG-MSFE extracts frame-relative motion and structural features at various granularities in a fully convolutional parallel manner and supports independent, flexible adjustment of context-aware granularity at different scales, thus enhancing computational efficiency and accuracy. Extensive experiments demonstrate that GaraMoSt achieves the SOTA performance in accuracy, robustness, visual effects, and noise suppression, comprehensively surpassing MoSt-DSA and other natural scene VFI methods.

Code — <https://github.com/ZyoungXu/GaraMoSt>

Introduction

4D Digital Subtraction Angiography (DSA) is an advanced medical imaging technology critical for diagnosing and treating various vascular diseases, including in the brain and heart (2022). It is widely utilized in hospital interventional surgeries. DSA operates by injecting a contrast agent, usually iodine-based, into the patient’s body, then capturing images from different angles in space at a fixed rotational speed, alongside the dynamic blood flow changes over time (2024). Due to the radiative nature of the imaging environment, the radiation dose received by patients and physicians is directly proportional to the image count, posing a threat to health.

Decreasing the number of images through frame interpolation is a straightforward and effective method to reduce





Figure 1: SSIM-Time comparison of various methods for interpolating 1 to 3 frames. Our GaraMoSt- \mathcal{L}_1 achieves SOTA accuracy, while time cost is almost the same as MoSt-DSA. Details in Table 1,2,3.

radiation. However, simply achieving interpolation is insufficient for practical application; the interpolated images must be accurate to assist in diagnostics effectively, and the time cost must be low to aid surgeries in real-time and save valuable patient time. Due to data scarcity and cross-disciplinary complexities, these issues received little attention and lacked targeted solutions. Only recently, MoSt-DSA (2024a) was the first to address these issues by proposing a direct multi-frame interpolation (MFI) method. MFI is a video frame interpolation (VFI) task and a significant research topic in computer vision (2023; 2018; 2023; 2024). Conventional MFI methods using recursion are time-consuming and unsuitable for real-time support (2022; 2020; 2021). In contrast, MoSt-DSA’s direct MFI method extracts general motion and structural features, maps them at different time steps, and directly infers any number of frames in once forward computation, achieving state-of-the-art (SOTA) performance in terms of image quality, robustness, inference time, and computational cost.

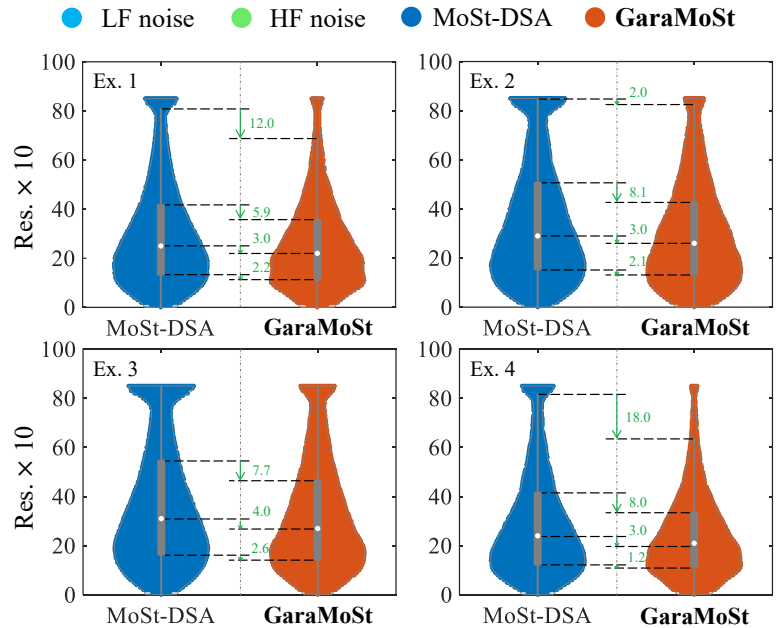
The core of the direct MFI method is the precise extraction of motion and structural features (2023). Existing VFI

*Corresponding Author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



(a) Qualitative comparison



(b) Quantitative comparison of residual distribution

Figure 2: Qualitative and quantitative comparison of GaraMoSt vs. MoSt-DSA. (a) Qualitatively, MoSt-DSA shows insufficient suppression of **high** and **low** frequency noise, whereas GaraMoSt significantly improves these issues. (b) Quantitatively, GaraMoSt enhances noise suppression, with notable reductions in the first quartile, median, third quartile, and upper adjacent of the noise distribution. Additionally, the reduction in width demonstrates a decrease in noise quantity.

methods, typically designed for natural scenes, inadequately extract motion and structural features of DSA images or do so at a coarse granularity. Common approaches are categorized into three types, see Figure 3 (a), (b), (c): (a) blending motion and structural feature extraction in a single module, leading to ambiguity in both features (2023; 2022; 2022; 2019; 2021); (b) sequentially extracting structural features of each frame and inter-frame motion features in multiple modules, although clearly defining motion features, lack correspondence in inter-frame structure (2023; 2023; 2023; 2022; 2022; 2018; 2021; 2022; 2021; 2019; 2020); (c) simultaneously extracting relative motion and structural features in a single module, but due to the coarse granularity of context, it fails to adapt to the fine-grained, complex structures of DSA images (2023). Ultimately, these methods produce motion artifacts, structural dissipation, and blurring issues defined by MoSt-DSA (2024a). MoSt-DSA introduced a fine-grained context interaction module for simultaneously extracting motion and structural features that enable flexible adjustment of context granularity, significantly improving these issues, as shown in Figure 3 (d). However, MoSt-DSA focuses on real-time performance, and although the feature granularity is flexible, the feature sampling scale is limited, leading to insufficient suppression of high-frequency (HF) and low-frequency (LF) noise. As illustrated in Figure 2, MoSt-DSA’s inference residuals primarily concentrate on the edges of the tiny blood vessels (HF noise), as well as the sparse background of the vessels and areas with consistent pixel (LF noise).

In summary, a significant challenge is how to further en-

hance the suppression of HF and LF noise in interpolated images while maintaining low time costs, thereby assisting physicians in more accurate real-time diagnostics and treatment. To address this challenge, we introduce GaraMoSt. Unlike conventional approaches that increase network depth via serial connections for deeper downsampling (2023; 2020), our method optimizes the network pipeline with a parallel design, using a wider network width to achieve deeper downsampling, maintaining inference time at the same level as MoSt-DSA. Specifically, we first obtain four different scales of features through sequential convolution and downsampling; second, we apply multi-scale dilated convolutions for cross-scale fusion, merging features from layers 1-3 and 2-4 to form two different scales of fused features (1/8 and 1/16); third, we introduce a module named MG-MSFE that extracts relative motion and structural features at various granularities in a fully convolutional parallel manner, enabling flexible adjustment of context-aware granularity across various scales, and transforming optimal context into linear functions, moving away from reliance on expensive attention maps, thereby enhancing computational efficiency and accuracy, as depicted in Figure 3 (e); finally, combining different time steps t and the two scales of motion and structural features, we predict dual-layer flows and masks and refine them using a simplified UNet (2015) to produce the final intermediate frame I_t .

In conclusion, our work makes the following contributions:

(1) We introduce GaraMoSt. Compared to MoSt-DSA, GaraMoSt further enhances the suppression of HF and LF noise in interpolated images, with inference time maintained

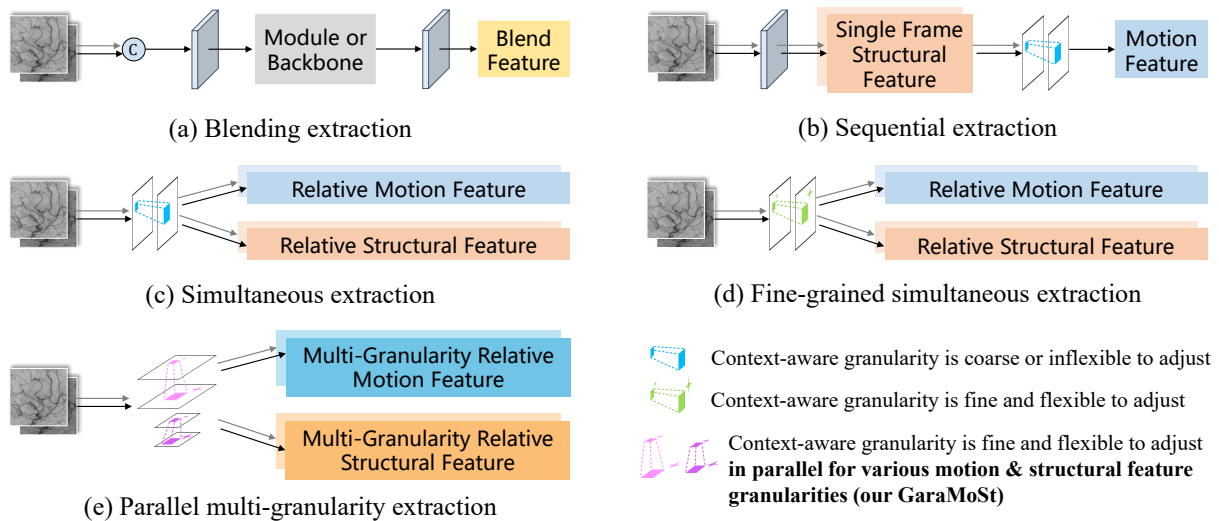


Figure 3: Illustration of our proposed parallel multi-granularity extraction and other methods of extracting motion and structural features. The approach proposed by our GaraMoSt is shown in (e). Notably, "simultaneous" should not be confused with "parallel". While (c) and (d) achieve simultaneous output of motion and structural features outside the module, the internal computation process is still highly sequential.

at the same level (for interpolating 3 frames, only increasing by 0.005s), enabling real-time assistance for more precise diagnostic and therapeutic procedures by physicians.

(2) We propose a module named MG-MSFE that extracts frame-relative motion and structural features at various granularities in a fully convolutional parallel manner, supporting independent and flexible adjustment of context-aware granularity across various scales, and transforming optimal context into linear functions, moving away from reliance on expensive attention maps, thereby improving computational efficiency and accuracy.

(3) Extensive experiments demonstrate that GaraMoSt achieves SOTA performance in accuracy, robustness, visual effects, and noise suppression, comprehensively surpassing MoSt-DSA and other natural scene VFI methods.

Related Work

DSA frame interpolation. Characteristics of DSA Images: (1) Filled with microscopic vessels, the fine-grained structure and motion complexity are much higher than in macroscopic natural images; (2) Composed of numerous tiny vessels, whereas natural images are usually composed of countable large instances (e.g., people, cars); (3) Sufficient structural features can usually be obtained with less downsampling; (4) Background pixels are highly similar and uniform. As described in Intro, MoSt-DSA is the first real-time solution but has insufficient suppression of HF and LF noise.

Method

Improvements Overview. The pipeline improvements of GaraMoSt over MoSt-DSA include: **Feature Extraction:** Added a parallel 1/8 scale feature downsampling for both I0 and I1. **Cross-Scale Feature Fusion:** Added a parallel fusion path (Path 1) for 1/2, 1/4, and 1/8 scale features in

both I0 and I1. Path 1 runs parallel to Path 2, which fuses 1/1, 1/2, and 1/4 scale features. **Motion and Structural Feature Extraction:** Added parallel extraction of relative motion and structural features from Path 1, alongside extraction from Path 2. Moreover, different granularities are adopted for Path 1 and Path 2. **Flow and Mask Prediction:** Modified to predict dual-layer flow and mask, with the output as the sum of both layers. **Refiner:** Incorporated relative structural features from Path 2 into the lowest-level feature concatenation.

Multi-Granularity Motion-Structure Feature Extractor (MG-MSFE)

As outlined in the introduction, DSA frame interpolation faces challenges in enhancing noise suppression while maintaining low time costs, crucial for accurate real-time diagnostics and treatment. To address this, deepening the feature extraction downsampling is effective but traditional methods (2021; 2022; 2020; 2023) via serial feature extractor connections are too slow for real-time use. Furthermore, DSA interpolation requires both structural and motion features rich in geometric details. Since deeper layers often lose geometric precision for abstract semantics, using these layers for feature computation and refinement may degrade performance, as our ablation study shows. Hence, we introduce MG-MSFE, parallel extracting multi-granularity relative motion and structural features in shallow layers, enhancing accuracy and reducing noise effectively without sacrificing speed.

Module design and calculation. The complete structure of MG-MSFE is shown in Figure 4. The input consists of four feature sets: I_0^0 and I_1^0 are the cross-scale fusion features of I_0 and I_1 from the 1-3 layers in Multi-Scale Feature Extractor, and I_0^1 and I_1^1 are from the 2-4 layers. After encoding these features with Multi-Granularity MoSt Attention, we obtain two different scale motion features, M_0 and M_0'

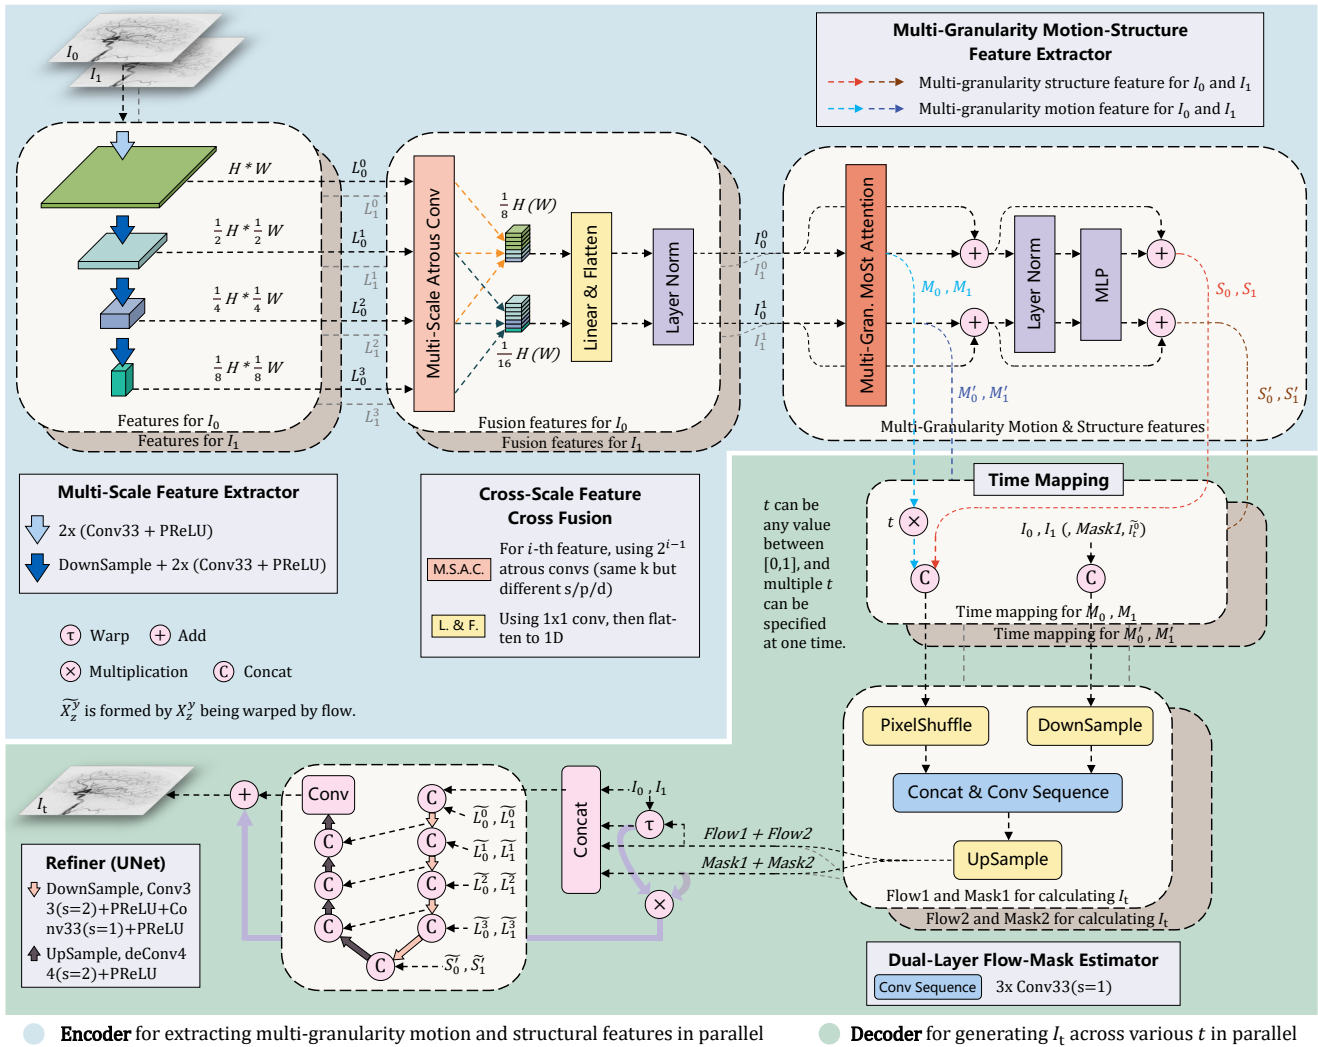


Figure 4: Overall pipeline of our GaraMoSt. The **encoder** includes the multi-scale feature extractor (MSFE), cross-scale feature cross fusion (CSFCF), and the multi-granularity motion-structure feature extractor (MG-MSFE) for extracting general multi-granularity motion and structural features in parallel. The **decoder** consists of the time mapping (TM), dual-layer flow-mask estimator (DL-FME), and Refiner, tailored to parallel decode and refine general multi-granularity features to generate I_t .

for I_0 relative to I_1 , and M_1 and M'_1 for I_1 relative to I_0 . Simultaneously, Multi-Granularity MoSt Attention encodes the enhanced inter-frame structural features. These initial structural features then pass through residual connections, layer norm, and MLP to obtain the final structural features: S_0, S'_0 for I_0 relative to I_1 and S_1, S'_1 for I_1 relative to I_0 . The sizes of M_0 and M_1 match those of S_0 and S_1 , while the sizes of M'_0 and M'_1 match those of S'_0 and S'_1 .

The calculation process of Multi-Granularity MoSt Attention is shown in Figure 5. In the left image pyramid, yellow and blue blocks represent the cross-scale fusion features of I_0 and I_1 respectively. Different scales correspond to s values of 1 or 2. First, we employ a Lambda Layer (2021) to simulate content-based and position-based contextual interactions in a fully convolutional manner. Specifically, we denote the depths of the query and value as $|k|$ and $|v|$, respectively,

and the position information as $P_0^s \in \mathbb{R}^{|n| \times d}$ (each pixel corresponds to a value between $[-1.0, 1.0]$, increasing along rows or columns across different channels). Second, we perform parallel relative attention calculations on the cross-scale fusion features $I_0^s \in \mathbb{R}^{|n| \times d}$ and $I_1^s \in \mathbb{R}^{|n| \times d}$ for scale s . Mathematically, the calculation of I_0^s relative to I_1^s is as follows, and the calculation of I_1^s relative to I_0^s is similar:

$$\begin{cases} Q_0^s = I_0^s W_{Q_0^s} \in \mathbb{R}^{|n| \times |k|} \\ K_0^s = I_1^s W_{K_0^s} \in \mathbb{R}^{|n| \times |k|} \\ V_0^s = I_1^s W_{V_0^s} \in \mathbb{R}^{|n| \times |v|} \end{cases} \quad (1)$$

Then we represent relative position embeddings as $E_0^s \in \mathbb{R}^{|n| \times |k|}$ (hyperparameters of a 3D-Conv). By normalizing the keys, we obtain $\bar{K}_0^s = \text{softmax}(K_0^s, \text{axis} = n)$. Next, we compute the content-based contextual interactions $\lambda_{c_0}^s$

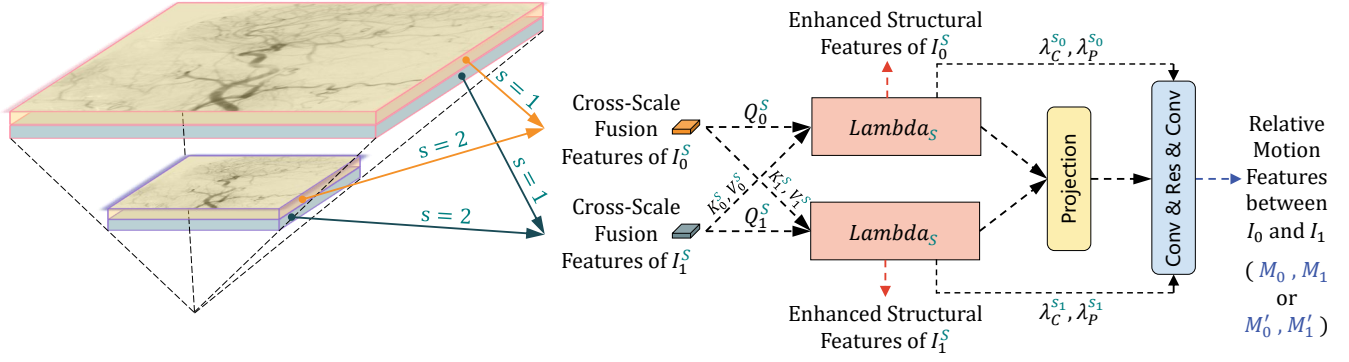


Figure 5: Multi-Granularity MoSt-Attention for parallel extracting of multi-granularity motion and structural features. Different scales correspond to different values of s , with s being either 1 or 2. In the left image pyramid, the yellow and blue blocks represent the cross-scale fusion features of I_0 and I_1 respectively. Enhanced structural features are used in MG-MSFE for subsequent calculations to derive the final structural features, see Figure 4 for details.




Figure 6: Using various granularities in Lambda for feature extraction across scales: for regions at similar positions across scales, $Lambda_{s1}$ and $Lambda_{s2}$ summarize different context within different scope r (i.e., granularity) into a fixed-size linear function (i.e., a matrix) applied to the corresponding query, thus bypassing the need for memory-intensive attention maps while maintaining flexibility.

and position-based contextual interactions $\lambda_{p_0}^s$, as:

$$\begin{cases} \lambda_{c_0}^s = \bar{K}_0^s T V_0^s \in \mathbb{R}^{|k| \times |v|} \\ \lambda_{p_0}^s = E_0^s T V_0^s \in \mathbb{R}^{|k| \times |v|} \end{cases} \quad (2)$$

Finally, by applying contextual interactions to the queries as well as P_0^s , we obtain the general multi-granularity motion and structural features. Specifically, when s takes values of 1 and 2, we obtain S_0 and M_0 , and S'_0 and M'_0 , as:

$$\begin{cases} S_0 = Q_0^1 (\lambda_{c_0}^1 + \lambda_{p_0}^1), S'_0 = Q_0^2 (\lambda_{c_0}^2 + \lambda_{p_0}^2) \\ M_0 = P_0^1 (\lambda_{c_0}^1 + \lambda_{p_0}^1), M'_0 = P_0^2 (\lambda_{c_0}^2 + \lambda_{p_0}^2) \end{cases} \quad (3)$$

All these processes are similar for the calculation of I_1^s relative to I_0^s . In the actual code, we concatenate I_0^s and I_1^s along the channel dimension to enable parallel computation. Figure 6 provides a detailed illustration of how to flexibly extract features at different scales with varying granularities.

Parallel Extraction of General Multi-Granularity Motion and Structural Features

Multi-Scale Feature Extractor (MSFE). We use MSFE to obtain vascular features at four different scales. For I_0 and I_1 , we first compute the first layer low-level features L_0^0 and L_1^0 using a 3x3 convolution followed by PReLU (2015). Then, through downsampling and the same configuration, we compute the second layer low-level features L_0^1 and L_1^1 . Similarly, we obtain the third and fourth layer low-level features L_0^2, L_1^2 and L_0^3, L_1^3 . Mathematically,

$$\begin{cases} L_j^0 = H(I_j), & L_j^1 = D(L_j^0) \\ L_j^2 = D(L_j^1), & L_j^3 = D(L_j^2) \end{cases}, \quad (4)$$

where H is a stack of convolution and activation functions, and D represents the integration of H and downsampling operations, with j being 0 or 1.

Cross-Scale Feature Cross Fusion (CSFCF). We use CSFCF to fuse vascular features. Specifically, we first use multi-scale atrous convolution to fuse the 1-3 layer features and the 2-4 layer features from MSFE. For the i -th layer low-level features L_0^i and L_1^i , we use 2^{i-1} atrous convolutions (2017) (fixed kernel size of 3, stride of 2^i , and for the n -th atrous convolution, padding/dilation is n). Mathematically,

$$\mathcal{F}(L_j^i) = (A_1(L_j^i), \dots, A_n(L_j^i)), \quad (5)$$

where \mathcal{F} signifies feature fusion, A indicates atrous convolution. The variable n , representing the number of A , takes a value of 2^{i-1} (i equal to 0, 1, or 2). By merging the fused features and applying a linear mapping, we obtain four different cross-scale fusion features: I_0^0 and I_1^0 from the 1-3 layers, and I_0^1 and I_1^1 from the 2-4 layers, as:

$$\begin{cases} I_j^0 = \mathcal{T}[\mathcal{C}(\mathcal{F}(L_j^0), \mathcal{F}(L_j^1), \mathcal{F}(L_j^2))] \\ I_j^1 = \mathcal{T}[\mathcal{C}(\mathcal{F}(L_j^1), \mathcal{F}(L_j^2), \mathcal{F}(L_j^3))] \end{cases}, \quad (6)$$

where j being 0 or 1. \mathcal{T} represents the linear mapping, and \mathcal{C} indicates the concatenation operation. Finally, we flatten and normalize I_j^0 and I_j^1 , preparing them for subsequent processing by MG-MSFE.

Method	SSIM (%)		PSNR		Time (s)
	Mean \uparrow	STD \downarrow	Mean \uparrow	STD \downarrow	1 frame \downarrow
ABME (2021)	94.02	2.28	39.83	3.39	0.383
FILM- \mathcal{L}_1 (2022)	94.11	2.37	39.86	3.43	0.201
EMA-small (2023)	94.19	2.21	40.07	3.47	0.027
EMA (2023)	94.33	2.13	40.13	3.40	0.056
MoSt-DSA- \mathcal{L}_1 (2024a)	94.62	2.12	40.32	3.35	0.024
GaraMoSt- \mathcal{L}_1 †	95.05	1.89	40.69	3.29	0.029
SoftSplat- \mathcal{L}_F (2020)	91.78	3.07	38.59	3.51	0.035
FILM- \mathcal{L}_{VGG} (2022)	93.10	2.67	39.27	3.45	0.201
FILM- \mathcal{L}_{Style} (2022)	93.05	2.72	39.25	3.48	0.201
MoSt-DSA (2024a)	93.65	2.61	39.55	3.45	0.024
GaraMoSt†	94.37	2.22	40.19	3.31	0.029

Table 1: Quantitative comparison with SOTA methods on single-frame interpolation. Best scores for color losses in blue, and for perceptually-sensitive losses in red. “†” obtained by ourselves, the rest are copied from (2024a).

Parallel Decoding and Refinement for Generating Multiple Frames

Time Mapping (TM). We use TM to map the general multi-granularity motion features to the specified time steps in parallel and concatenate them with the corresponding scale general structural features as part of the input (α) to the DL-FME. The other part of the input is β , which varies across different scales. At higher scales, β_h consists only of the concatenation of I_0 and I_1 , while at lower scales, β_l includes I_0 , I_1 , \tilde{I}_t^0 , and μ_t^0 . Here, μ_t^0 is the first layer mask calculated by the DL-FME, and \tilde{I}_t^0 is the first layer flow-warped result of I_0 and I_1 calculated by the DL-FME.

Dual-Layer Flow-Mask Estimator (DL-FME). As shown in Figure 4, DL-FME (denoted by \mathcal{F}) applies PixelShuffle (2016) upsampling to α , and downsampling to β . Subsequently, α and β merge and undergo continuous convolution operations, eventually generating bidirectional optical flow ϕ_t^s and mask μ_t^s for the specific t through upsampling in different layers, as:

$$\begin{cases} \phi_t^0, \mu_t^0 = \mathcal{F}(\mathcal{C}(M_{0 \rightarrow t}^0, M_{1 \rightarrow t}^0, S_0^0, S_1^0), \beta_h) \\ \phi_t^1, \mu_t^1 = \mathcal{F}(\mathcal{C}(M_{0 \rightarrow t}^1, M_{1 \rightarrow t}^1, S_0^1, S_1^1), \beta_l) \end{cases} \quad (7)$$

Next, we sum ϕ_t and μ_t across different layers, using the summed ϕ_t to warp I_0 , I_1 , the low-level features L_j^i , and S_0^i and S_1^i . For instance, for X_z^y , the result after warping is denoted as \tilde{X}_z^y . Subsequently, we concatenate I_0 , I_1 , \tilde{I}_0 , \tilde{I}_1 , ϕ_t , and μ_t together, referred to as \mathcal{O}_t for Refiner to process.

Refiner. Follow the design of MoSt-DSA, except incorporated \tilde{S}_0^i and \tilde{S}_1^i into the lowest-level feature concatenation.

Experiments

Model Configuration. For the best accuracy-time trade-off, context-aware granularities (r) are [7, 7], [7, 29], [7, 29]. See the ablation study for proof.

Training Details. We followed the settings of MoSt-DSA (2024a), except for the following differences. For interpolating 1 to 3 frames, we set the batch sizes to 10, 6, and 4, with 1000 warm-up steps. We use the AdamW (2017) optimizer

Method	SSIM (%)		PSNR		Time (s)
	Mean \uparrow	STD \downarrow	Mean \uparrow	STD \downarrow	2 frame \downarrow
FILM- \mathcal{L}_1 (2022)	92.62	3.13	37.93	3.82	0.388
EMA-small (2023)	91.82	3.68	37.37	4.07	0.074
EMA (2023)	91.90	3.63	37.41	4.08	0.112
MoSt-DSA- \mathcal{L}_1 (2024a)	94.35	2.29	39.78	3.44	0.070
GaraMoSt- \mathcal{L}_1 †	94.70	2.08	40.16	3.36	0.076
SoftSplat- \mathcal{L}_F (2020)	90.86	3.49	37.84	3.43	0.084
FILM- \mathcal{L}_{VGG} (2022)	91.42	3.42	37.47	3.67	0.388
FILM- \mathcal{L}_{Style} (2022)	91.31	3.50	37.39	3.74	0.388
MoSt-DSA (2024a)	93.14	2.84	38.94	3.38	0.070
GaraMoSt†	94.01	2.42	39.66	3.37	0.076

Table 2: Quantitative comparison with SOTA methods on two frames interpolation. Blue, red, † same to Table 1.

Method	SSIM (%)		PSNR		Time (s)
	Mean \uparrow	STD \downarrow	Mean \uparrow	STD \downarrow	3 frame \downarrow
FILM- \mathcal{L}_1 (2022)	91.94	3.52	37.21	3.91	0.548
EMA-small (2023)	90.48	4.52	36.31	4.24	0.122
EMA (2023)	90.57	4.50	36.35	4.24	0.165
MoSt-DSA- \mathcal{L}_1 (2024a)	93.58	2.69	38.85	3.56	0.117
GaraMoSt- \mathcal{L}_1 †	94.22	2.34	39.53	3.51	0.122
SoftSplat- \mathcal{L}_F (2020)	90.07	3.87	37.24	3.53	0.137
FILM- \mathcal{L}_{VGG} (2022)	90.63	3.83	36.75	3.74	0.548
FILM- \mathcal{L}_{Style} (2022)	90.54	3.90	36.66	3.82	0.548
MoSt-DSA (2024a)	93.03	2.94	38.66	3.59	0.117
GaraMoSt†	93.49	2.71	39.06	3.48	0.122

Table 3: Quantitative comparison with SOTA methods on three frames interpolation. Blue, red, † same to Table 1.

with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and a weight decay of $6e - 5$. The learning rate is warmed up to $6e - 5$ and then decays to $6e - 6$ over 100 epochs following a cosine schedule (2016).

Comparison Details. For a fair comparison, we followed MoSt-DSA’s settings, including trained two versions on their DSA dataset: one (GaraMoSt- \mathcal{L}_1) using only the \mathcal{L}_1 loss, which achieves higher test scores, and the other (GaraMoSt) using the combined loss \mathcal{L} (2024a), which benefits image quality. MoSt-DSA’s dataset (2024a) contains 470 head DSA image sequences (329 for training and 141 for testing; 489*489 resolution) from 8 hospitals, which are made in the data form of interpolating 1 to 3 frames. Our results are fairly derived from a single NVIDIA RTX 3090 GPU.

Single-Frame Interpolation (SFI)

The quantitative evaluation results are shown in Table 1. EMA is short for EMA-VFI (2023). Our method achieved significant accuracy improvements whether compared with methods using color losses or perceptually-sensitive losses. Moreover, the substantial reduction in the STD of SSIM (S.SSIM) demonstrates our model’s significant robustness enhancement. Despite leading in accuracy, our inference time is nearly identical to MoSt-DSA, with only a 0.005s difference. Specifically, compared to methods using color losses, GaraMoSt- \mathcal{L}_1 achieved a Mean of SSIM (M.SSIM) of 95.05 (+0.33), an S.SSIM of 1.89 (-0.23), a Mean of PSNR (M.PSNR) of 40.69 (+0.37), and a STD of PSNR (S.PSNR) of 3.29 (-0.06). Compared to methods using perceptually-sensitive losses, GaraMoSt performed even better, achieving

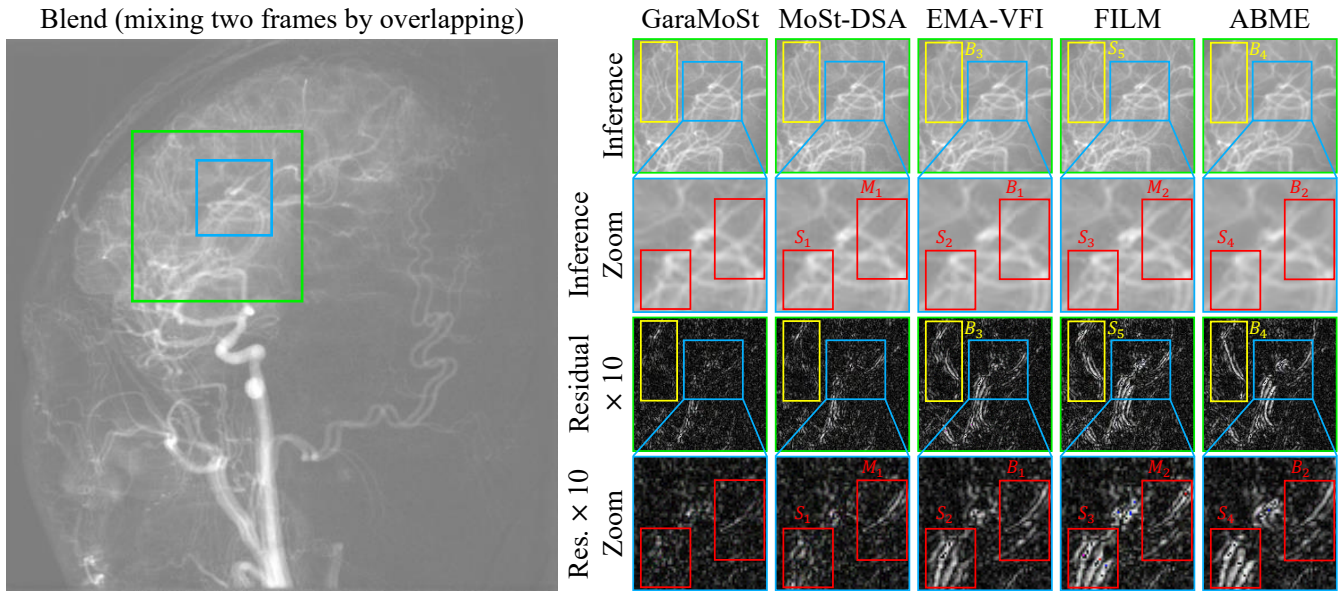


Figure 7: Visual comparison for interpolating one frame: our GaraMoSt vs. other SOTA methods. The first and third rows correspond to the green box in the "blend", while the second and fourth rows correspond to the blue box in the "blend". M stands for motion artifact, S for structural dissipation, and B for blurring. The residual levels of our GaraMoSt are significantly lower than those of all other methods. This demonstrates that we effectively suppressed noise and further improved issues related to motion artifacts, structural dissipation, and blurring.

an M.SSIM of 94.37 (+0.72), an S.SSIM of 2.22 (-0.39), an M.PSNR of 40.19 (+0.64), and an S.PSNR of 3.31 (-0.14). Notably, the score differences in SSIM and PSNR among SOTA VFI methods are minimal. For instance, on the UCF101 (2012), the SOTA EMA-VFI surpasses the second-best (2023) by only 0.01% in SSIM and 0.01 in PSNR, and the third-best (2023) by 0.04% in SSIM and 0.01 in PSNR. Thus, the superiority of our method is significant.

The qualitative comparison results are shown in Figure 7. The residual levels of our GaraMoSt are significantly lower than those of all other methods. This demonstrates that we effectively suppressed noise and further improved issues related to motion blur, structural dissipation, and blurring.

Multi-Frame Interpolation (MFI)

The quantitative results for interpolating 2 and 3 frames are presented in Table 2, 3. We continue to show significant improvements in accuracy and robustness. The inference time for interpolating 2 and 3 frames remains almost the same as MoSt-DSA, with differences of 0.006s and 0.005s. Notably, when interpolating 2 frames, our GaraMoSt achieves an M.SSIM of 94.01 (+0.87), an S.SSIM of 2.42 (-0.42), and an M.PSNR of 39.66 (+0.72). When interpolating 3 frames, our GaraMoSt- \mathcal{L}_1 achieves an M.SSIM of 94.22 (+0.64), an S.SSIM of 2.34 (-0.35), and an M.PSNR of 39.53 (+0.68).

Ablation Study

Impact of various context-aware granularity in SFI. As shown in Supplementary Material’s Table 4 (2024b), for the best accuracy-time trade-off, [7, 7], [15, 15], [7, 29], and [15, 29] are the best 4 options, while [29, 15] is the worst.

Among various schemes, the time differences up to 0.016s, M.SSIM differences up to 0.19, and S.SSIM differences up to 0.07. This demonstrates the importance of choosing the appropriate granularity.

Impact of various context-aware granularity in MFI.

We further compared the SFI’s best 4 options on MFI, as shown in Supplementary Material’s Table 5 (2024b). The results indicate that [7, 29] is more suitable for MFI.

Influence of structural features with different depths on refinement.

We upsampled \tilde{S}_0^i and \tilde{S}_1^i to match the size of \tilde{L}_0^3 and \tilde{L}_1^3 and replaced \tilde{L}_0^3 and \tilde{L}_1^3 in the Refiner with the upsampled features, resulting in a decrease of 0.10 in M.SSIM, and a decrease of 0.09 in M.PSNR. This indicates the importance of using shallow structural features for refinement. See details in Supplementary Material’s Table 6 (2024b).

Conclusion

We proposed GaraMoSt, which further enhances HF and LF noise suppression in interpolated images while maintaining low time costs, enabling real-time assistance for more precise diagnostic and therapeutic procedures by physicians. In particular, we devised a general module MG-MSFE that efficiently extracts relative motion and structural features at various granularities in a fully convolutional parallel manner, supporting independent and flexible adjustment of context-aware granularity across various scales. Extensive experiments show that our GaraMoSt achieves SOTA performance in accuracy, robustness, visual effects, and noise suppression, comprehensively surpassing MoSt-DSA and other natural scene VFI methods.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (No. 62376102, No. 82472070).

References

- Bao, W.; Lai, W.-S.; Ma, C.; Zhang, X.; Gao, Z.; and Yang, M.-H. 2019. Depth-aware video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3703–3712.
- Bello, I. 2021. Lambdanetworks: Modeling long-range interactions without attention. *arXiv preprint arXiv:2102.08602*.
- Chen, L.-C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2017. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4): 834–848.
- Danier, D.; Zhang, F.; and Bull, D. 2022. ST-MFNet: A Spatio-Temporal Multi-Flow Network for Frame Interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3521–3531.
- Ding, T.; Liang, L.; Zhu, Z.; and Zharkov, I. 2021. Cdfi: Compression-driven network design for frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8001–8011.
- Fang, Y.; Sun, Q.; Wang, X.; Huang, T.; Wang, X.; and Cao, Y. 2024. Eva-02: A visual representation for neon genesis. *Image and Vision Computing*, 105171.
- Gao, P.; Tian, H.; and Qin, J. 2023. Video Frame Interpolation with Flow Transformer. *arXiv preprint arXiv:2307.16144*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Huang, Z.; Zhang, T.; Heng, W.; Shi, B.; and Zhou, S. 2020. Rife: Real-time intermediate flow estimation for video frame interpolation. *arXiv preprint arXiv:2011.06294*.
- Jia, Z.; Lu, Y.; and Li, H. 2022. Neighbor Correspondence Matching for Flow-based Video Frame Synthesis. In *Proceedings of the 30th ACM International Conference on Multimedia*, 5389–5397.
- Jiang, H.; Sun, D.; Jampani, V.; Yang, M.-H.; Learned-Miller, E.; and Kautz, J. 2018. Super slo-mo: High quality estimation of multiple intermediate frames for video interpolation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 9000–9008.
- Jin, X.; Wu, L.; Chen, J.; Chen, Y.; Koo, J.; and Hahm, C.-h. 2023. A Unified Pyramid Recurrent Network for Video Frame Interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1578–1587.
- Kalluri, T.; Pathak, D.; Chandraker, M.; and Tran, D. 2023. FLAVR: Flow-Agnostic Video Representations for Fast Frame Interpolation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2071–2082.
- Kong, L.; Jiang, B.; Luo, D.; Chu, W.; Huang, X.; Tai, Y.; Wang, C.; and Yang, J. 2022. IFRNet: Intermediate Feature Refine Network for Efficient Frame Interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1969–1978.
- Loshchilov, I.; and Hutter, F. 2016. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*.
- Loshchilov, I.; Hutter, F.; et al. 2017. Fixing weight decay regularization in adam. *arXiv preprint arXiv:1711.05101*, 5.
- Lu, L.; Wu, R.; Lin, H.; Lu, J.; and Jia, J. 2022. Video Frame Interpolation with Transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3532–3542.
- Niklaus, S.; and Liu, F. 2018. Context-aware synthesis for video frame interpolation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1701–1710.
- Niklaus, S.; and Liu, F. 2020. Softmax splatting for video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5437–5446.
- Park, J.; Lee, C.; and Kim, C.-S. 2021. Asymmetric bilateral motion estimation for video frame interpolation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14539–14548.
- Posa, A.; Tanzilli, A.; Barbieri, P.; Steri, L.; Arbia, F.; Mazza, G.; Longo, V.; and Iezzi, R. 2022. Digital Subtraction Angiography (DSA) Technical and Diagnostic Aspects in the Study of Lower Limb Arteries. *Radiation*.
- Reda, F.; Kontkanen, J.; Tabellion, E.; Sun, D.; Pantofaru, C.; and Curless, B. 2022. FILM: Frame Interpolation for Large Motion. *arXiv preprint arXiv:2202.04901*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241. Springer.
- Shang, W.; Ren, D.; Yang, Y.; Zhang, H.; Ma, K.; and Zuo, W. 2023. Joint Video Multi-Frame Interpolation and Deblurring Under Unknown Exposure Time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13935–13944.
- Shi, W.; Caballero, J.; Huszár, F.; Totz, J.; Aitken, A. P.; Bishop, R.; Rueckert, D.; and Wang, Z. 2016. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1874–1883.
- Sim, H.; Oh, J.; and Kim, M. 2021. Xvfi: Extreme video frame interpolation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14489–14498.
- Soomro, K.; Zamir, A. R.; and Shah, M. 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*.
- Xu, Z.; Zhao, H.; Cui, Z.; Liu, W.; Zheng, C.; and Wang, X. 2024a. MoSt-DNA: Modeling Motion and Structural Interactions for Direct Multi-Frame Interpolation in DSA Images. In *European Conference on Artificial Intelligence*, 537–544.

- Xu, Z.; Zhao, H.; Liu, W.; and Wang, X. 2024b. GaraMoSt: Parallel Multi-Granularity Motion and Structural Modeling for Efficient Multi-Frame Interpolation in DSA Images. *arXiv preprint*.
- Xue, T.; Chen, B.; Wu, J.; Wei, D.; and Freeman, W. T. 2019. Video enhancement with task-oriented flow. *International Journal of Computer Vision*.
- Yu, Z.; Zhang, Y.; Zou, D.; Chen, X.; Ren, J. S.; and Ren, S. 2023. Range-Nullspace Video Frame Interpolation With Focalized Motion Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22159–22168.
- Zhang, G.; Zhu, Y.; Wang, H.; Chen, Y.; Wu, G.; and Wang, L. 2023. Extracting motion and appearance via inter-frame attention for efficient video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5682–5692.
- Zhao, H.; Xu, Z.; Chen, L.; Wu, L.; Cui, Z.; Ma, J.; Sun, T.; Lei, Y.; Wang, N.; Hu, H.; et al. 2024. Large-scale pre-trained frame generative model enables real-time low-dose DSA imaging: An AI system development and multi-center validation study. *Med*.
- Zhou, C.; Liu, J.; Tang, J.; and Wu, G. 2023. Video Frame Interpolation with Densely Queried Bilateral Correlation. *arXiv preprint arXiv:2304.13596*.