

# MEGAN: Memory Enhanced Graph Attention Network for Space-Time Video Super-Resolution

Chenyu You<sup>1</sup> Lianyi Han<sup>2</sup> Aosong Feng<sup>1</sup> Ruihan Zhao<sup>3</sup>  
Hui Tang<sup>2</sup> Wei Fan<sup>2</sup>  
<sup>1</sup>Yale University <sup>2</sup>Tencent Hippocrates Research Lab  
<sup>3</sup>The University of Texas at Austin

## Abstract

*Space-time video super-resolution (STVSR) aims to construct a high space-time resolution video sequence from the corresponding low-frame-rate, low-resolution video sequence. Inspired by the recent success to consider spatial-temporal information for space-time super-resolution, our main goal in this work is to take full considerations of spatial and temporal correlations within the video sequences of fast dynamic events. To this end, we propose a novel one-stage memory enhanced graph attention network (MEGAN) for space-time video super-resolution. Specifically, we build a novel long-range memory graph aggregation (LMGA) module to dynamically capture correlations along the channel dimensions of the feature maps and adaptively aggregate channel features to enhance the feature representations. We introduce a non-local residual block, which enables each channel-wise feature to attend global spatial hierarchical features. In addition, we adopt a progressive fusion module to further enhance the representation ability by extensively exploiting spatial-temporal correlations from multiple frames. Experiment results demonstrate that our method achieves better results compared with the state-of-the-art methods quantitatively and visually.*

## 1. Introduction

Space-time video super-resolution (STVSR) [43] aims at reconstructing a photo-realistic video sequence with high spatial and temporal resolutions from its corresponding low-frame-rate (LFR) and low-resolution (LR) counterparts. High-resolution (HR) slow-motion video sequences provide more visually appealing details in the space-time domain, which finds a wide range of applications [74, 71, 73, 25, 33, 24], such as sports video, movie making, and medical imaging analysis.

In recent years, there is a surge of research interests in upscaling the spatial-temporal resolution of a video simul-

taneously. Most of traditional space-time super-resolution (SR) methods [42, 6, 26, 36] propose to increase the resolution both in the time and space domain by combining information from multiple low-resolution video recordings in a dynamic scene, where they address the space-time SR problem as an optimization problem. For example, Shechtman *et al.* [42] formulates the problem of the space-time SR as linear least-square minimization with suitable regularization priors. Specifically, they demonstrate that it is possible to obtain an improved HR video sequence in both spatial and temporal dimensions. However, the complex nature of the real-world video sequences limits the effectiveness of these methods since they can only super-resolve video sequences containing slow-varying motions with a simple analytic form. Furthermore, the high computational cost limits their practical use in real-world scenarios.

Deep convolutional neural networks (CNNs) have shown promising results in various machine learning tasks [45, 72, 75, 21, 20, 76, 66, 68, 3, 70, 51, 50, 48, 53, 67, 69], such as video frame interpolation (VFI) [39, 64, 13], video super-resolution (VSR) [57, 8, 55], and image recognition [11, 52, 49]. To jointly up-scaling videos both in space and time, the most common way is to perform learning-based VFI and VSR methods alternately and independently. For example, it first constructs in-between frames by VFI, and then generates high-resolution video sequences by VSR. However, due to the strong space-time relationship among video sequences, such two-stage methods may fail to capture correlations between space and time, which is essential for increasing spatial-temporal resolution. Moreover, these methods consist of two independent CNN-based networks, which are computationally expensive.

To enhance feature representations used for space-time video SR purposes, it is desirable to integrate the spatial and temporal information. Existing one-stage space-time SR approaches enhance feature representations by jointly aggregating the representations both in spatial and temporal dimensions. However, these methods have two major limitations. First, these learned feature representations only

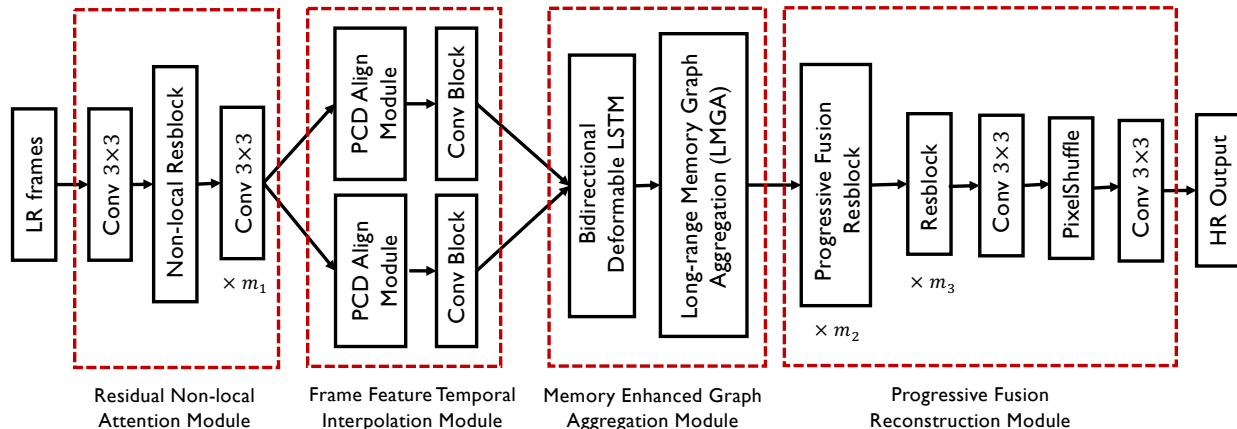


Figure 1: Framework overview of Memory Enhanced Graph Attention Networks (MEGAN). Our approach uses the residual non-local attention module to extensively exploit information both in temporal and spatial dimensions. Then we apply frame feature temporal interpolation and long-range memory graph aggregation (LMGA) module to handle complex motions in a dynamic scene. By capturing global temporal features, it enables progressive function reconstruction to generate HR output.

contain local contexts, while lacking global structural information. Features from multiple frames and different locations can help the network to model long-term dependencies among the video sequences more efficiently. Thus, only considering spatial information is an insufficient approximation of the global and local influence. Second, space-time SR is a dynamic process with complex scenes. Most CNN-based methods indistinguishably extract different types of features to super-solve in space-time domain, which lack discriminative representation ability to attend to different types of information of fast dynamic events.

In this paper, we work toward simultaneously reconstructing a video sequence of high spatial and temporal resolution given a low-frame-rate, low-resolution video sequence. The objective is achieved by the proposed memory enhanced graph attention network (MEGAN), which comprehensively leverages the spatial-temporal information to enhance the feature representation capability of CNNs. The motivation comes from the fact that features in spatial domain correlate with the contexts in temporal domain. In implementation, we use the non-local residual block to capture the global dependencies between channel-wise features and encourage the model to capture more informative features in the spatial dimension. Considering the non-local block requires high computational cost, which is practically challenging to use for the task of STVSR. Therefore, we improve the non-local residual block to boost convergence. In addition, we design a novel long-range memory graph aggregation (LMGA) module to adaptively learn a robust spatial-temporal feature representation between the current key frame and the neighboring frames. Specifically, we cache the extracted hierarchical features in LMGA and combine embedded features in the collection to propagate contextual information towards the improved feature repre-

sentations by constructing the adaptive graph. The embedded spatial features and temporal contexts provide a more clear definition of the feature aspects. To further promote the capability of feature representation, we apply progressive fusion residual dense blocks (PFRDBs) to progressively extract both intra-frame spatial correlations and inter-frame temporal correlations. Experimental results show that the proposed method achieves better results compared with the state-of-the-art methods. Our contributions are summarized as follows:

- We present a unified graph attention network for space-time SR problems. Our method utilizes the non-local residual blocks to better learn different kinds of information (*e.g.*, low-and high-frequency information) and long-range spatial dependencies.
- We devise a novel long-range memory graph aggregation (LMGA) module which enables the key frames to adaptively model temporal dependency among the video sequences of fast dynamic events. Both local and global information of the rapid dynamic space-time scene can allow the network to obtain an improved video sequence of high spatial-temporal resolution.
- We utilize progressive fusion residual dense blocks (PFRDBs) to combine spatial features and temporal information for better reconstruction performance.
- We validate the effectiveness of our MEGAN with extensive experiments. MEGAN achieves superior performance compared with state-of-the-art methods on three benchmark datasets.

## 2. Related Work

In this section, we overview the progress on three related topics: video super-resolution, video frame interpolation,




Figure 2: Overview of Long-range Memory Graph Aggregation (LMGA) module. The LMGA includes two operations: frame feature aggregation and graph construction. We randomly select  $\tau$  frame features from the shuffled ordered-index frame feature maps  $\{1, \dots, 2n + 1\}$ , and then cache the selected features to form the global memory block. By leveraging these cached information, the global-local feature aggregation function  $N_g(\cdot)$  enables the key frame to access more complex spatial-temporal information. We also dynamically construct  $\mathcal{G}_k$  by Graph Construction to exploit the spatial-temporal relationship between input feature frames in a discriminative manner. We then construct the output features by the small network containing two convolutional layers.

and space-time video super-resolution.

**Video Super-Resolution** Video super-resolution aims to recover HR video frames from the corresponding corrupted video frames. Several works [2, 64, 54, 40] adopted optical flow for explicit temporal alignment in video sequences, which can effectively capture image details among multiple frames and inter-frame motion to improve reconstruction accuracy. Other works [17, 47, 77] utilized deep residual learning to reconstruct HR output. Some works [14, 55, 57] attempted to address the implicit temporal alignment without motion estimation between video sequences at subpixel and subframe accuracy. For example, Jo *et al.* [14] generated dynamic upsampling filters and residuals to recover HR frames by preserving temporal consistency. Tian *et al.* [55] and Wang *et al.* [57] adaptively learned the temporal information without computing optical flow. Considering that recurrent neural networks (RNNs) can model long-term temporal information, some recent efforts [12, 40, 9] have been made to adopt RNN-based networks to encourage temporal dependency modeling without explicit temporal alignment, and therefore can steer very complex scene dynamics.

**Video Frame Interpolation** Video frame interpolation aims at synthesizing the in-between frames from two consecutive video frames. Video frame interpolation approach can usually be categorized as kernel-based [1, 39, 38], phase-based [35, 34], and flow-based [37, 31, 13, 64]. Some recent kernel-based works [39, 38] employed CNNs to estimate adaptive filters for each pixel and then convolved with input frames to synthesize an intermediate frame. More recently, some efforts [35, 34] proposed to synthesize an intermediate frame by utilizing a neural network decoder to estimate a per-pixel phase-based motion representation. Most

recently, the common approach is to guide the synthesis of an intermediate frame by predicting optical flow between two input frames. For example, Context-Aware Synthesis (CtxSyn) [37] computed the optical flow between two input frames and then forward-warped the images with temporal contextual information extracted from input images corresponding to the optical flow for high-quality frame synthesis.

**Space-Time Video Super-Resolution** The pioneer space-time super-resolution (SR) work [43] addressed the problem of simultaneous spatial and temporal super-resolution by solving linear function. The following work [42] achieved space-time SR from a single video by combing information from multiple space-time patches at sub-frame accuracy. More recently, the success of CNNs has motivated the development of STVSR methods [60, 10, 18, 15, 62, 61, 56]. Xiang *et al.* [60] proposed a one-stage STVSR method which employed a deformable convolutional LSTMs (ConvLSTM) to align and aggregate temporal contexts for better utilization of global information. Haris *et al.* [10] introduced a method to super-solve in spatial and temporal dimensions simultaneously and jointly. However, most methods have limited capacity to capture global dependencies and handle space-time visual patterns in complex dynamic scenes.

### 3. Method

Given a LR space-time video sequence ( $I^{LR} = \{I_{2t-1}^{LR}\}_{t=1}^{n+1}$ ), our method aims at reconstructing the corresponding HR video sequence ( $I^{HR} = \{I_t^{HR}\}_{t=1}^{2n+1}$ ). The size of input LR video frame is  $H \times W$ , where  $H$  and  $W$  refer to height and width, respectively. The size of HR output

is  $rH \times rW$  with the spatial upscaling factor  $r$ . An overview of MEGAN for space-time video super-resolution is shown in Figure 1. The framework consists of four main parts: residual non-local attention module, frame feature temporal interpolation module, memory enhanced graph aggregation module, and progressive fusion reconstruction module.

### 3.1. Residual Non-local Attention Module

Inspired by [29, 30, 58, 65], we first use a  $3 \times 3$  convolutional (Conv) layer as a shallow feature extractor. Then, we incorporate a residual non-local attention block [58, 65] to capture long-range dependencies between channel-wise features and allow the model to discriminatively learn different types of information (e.g., low and high frequency information). Mathematically, the non-local operation is defined as follows:

$$y_i = \left( \sum_{\forall j} f(x_i, x_j)g(x_j) \right) / \sum_{\forall j} f(x_i, x_j), \quad (1)$$

where  $x$  denotes the input of non-local operation and  $y$  denotes the output with the same size as  $x$ .  $i$  is the output position index, and  $j$  is the index of all possible positions. The pairwise function  $f(\cdot)$  computes the attention between two inputs, while the function  $g(\cdot)$  computes the feature representation at a certain position. Non-local neural networks [58] provide several versions of  $f(\cdot)$ , such as Gaussian function, dot product similarity, and feature concatenation. Different from [65], we use the dot product similarity to evaluate the pairwise attention:

$$f(x_i, x_j) = u(x_i)v(x_j) = (W_u x_i)W_v x_j, \quad (2)$$

where  $W_u$  and  $W_v$  are weight matrices, and we adopt a linear embedding for  $g : g(x_j) = W_g x_j$  with the weight matrix  $W_g$ . The reason why we use dot product similarity rather than the Gaussian function is to boost convergence, which allows us to train very deep networks, being more suitable for STVSR. Thus, we have the output  $z$  at position  $i$  of the non-local residual block:

$$z_i = W_z y_i + x_i = W_z \text{softmax}(f(x_i, x_j))g(x_j) + x_i, \quad (3)$$

where  $W_z$  is a weight matrix. After the  $m_1$  Conv layers, the hierarchical feature maps ( $F^{LR} = \{F_{2t-1}^{LR}\}_{t=1}^{n+1}$ ) serve as inputs for the following frame feature interpolation module.

### 3.2. Frame Feature Temporal Interpolation

Frame feature temporal interpolation module has been proven to have superior performance in [60], thus we use it in our network to for frame feature interpolation. Given two input feature maps  $F_1^{LR}$  and  $F_3^{LR}$ , our goal is to learn a feature temporal interpolation function  $f(\cdot)$  to estimate the

in-between feature map  $F_2^{LR}$ . To accurately capture complex intra-sequence motions [55], we adopt the modulated deformable sampling module [60] for frame feature sampling. Since the intermediate frame  $F_2^{LR}$  is not available, we use the deformable convolution to approximate the forward and backward motion conditions between  $F_1^{LR}$  and  $F_3^{LR}$ . We first use  $g_1$  to predict the learnable offset  $\Delta p_1$  for the feature  $F_1^{LR}$ . With  $\Delta p_1$  and  $F_1^{LR}$ , we obtain the sampled features  $F_1^{LR'}$  by the deformable convolution [5, 78]. The feature sampling procedure can be defined as follows:

$$\Delta p_1 = g_1([F_1^{LR}, F_3^{LR}]), \quad (4)$$

$$F_1^{LR'} = DCConv(F_1^{LR}, \Delta p_1), \quad (5)$$

where  $g_1$  denotes the general function including several convolution layers. We follow the same procedure in Equation 4 and 5 to compute the sampled feature  $F_3^{LR'}$  from  $\Delta p_3$  and  $F_3^{LR}$ . Finally, we utilize a simple linear blending function to synthesize the LR feature map. Overall, the general form of the interpolation function is formulated as follows:

$$F_2^{LR} = f(F_1^{LR}, F_3^{LR}) = C_1(F_1^{LR'}) + C_3(F_3^{LR'}), \quad (6)$$

where  $C_1$  and  $C_3$  are  $1 \times 1$  Conv layer. Similarly, we apply the feature temporal interpolation function  $f(\cdot)$  to  $\{F_{2t-1}^{LR}\}_{t=1}^{n+1}$ . As a result, we can obtain the in-between frame feature maps  $\{F_{2t}^{LR}\}_{t=1}^n$ .

### 3.3. Memory Enhanced Graph Aggregation

We now have the consecutive frame feature maps  $\{F_t^{LR}\}_{t=1}^{2n+1}$  as the input. The module mainly consists of two components: bidirectional deformable ConvLSTM, and long-range memory graph aggregation block (See in Figure 2). The input is the consecutive frame feature maps  $\{F_t^{LR}\}_{t=1}^{2n+1}$ .

**Deformable ConvLSTM** We adopt the bidirectional deformable ConvLSTM [60] to super-solve video with the complex motion in dynamic scenes. Comparing to naive ConvLSTM, deformable ConvLSTM can better handle lager motion in complex dynamic scenes by aligning hidden and cell states to the reference feature map.

$$\Delta p_t^h = g^h([h_{t-1}, F_t^{LR}]), \quad (7)$$

$$\Delta p_t^c = g^c([c_{t-1}, F_t^{LR}]), \quad (8)$$

$$h'_{t-1} = DCConv(h_{t-1}, \Delta p_t^h), \quad (9)$$

$$c'_{t-1} = DCConv(c_{t-1}, \Delta p_t^c), \quad (10)$$

$$h_t, c_t = ConvLSTM(h'_{t-1}, c'_{t-1}, F_t^{LR}), \quad (11)$$

where  $\Delta p_t^h$  and  $\Delta p_t^c$  refer to the estimated offset of the hidden and cell states.  $h'_{t-1}$  and  $c'_{t-1}$  are aligned hidden and cell states.  $g^h$  and  $g^c$  denote the embedding function which

VFI Method	VSR Method	Vid4		Vimeo-Fast		Vimeo-Medium		Vimeo-Slow		Adobe240		Parameters (Million)
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	
SuperSloMo [13]	Bicubic	23.38	0.6154	34.94	0.8848	30.35	0.8570	28.82	0.8230	26.34	0.7361	19.8
SuperSloMo [13]	RCAN [77]	23.81	0.6401	34.54	0.9076	32.50	0.8884	30.69	0.8624	26.43	0.7398	19.8+16.0
SuperSloMo [13]	RBPN [9]	24.42	0.7105	34.94	0.9174	33.23	0.9059	31.04	0.8753	28.22	0.8170	19.8+12.7
SuperSloMo [13]	EDVR [57]	25.33	0.7417	35.19	0.9201	33.51	0.9096	31.42	0.8832	28.72	0.8320	19.8+20.7
SepCov [39]	Bicubic	23.62	0.6322	32.42	0.8901	30.77	0.8650	29.18	0.8310	26.61	0.7457	21.7
SepCov [39]	RCAN [77]	24.92	0.7236	34.97	0.9195	33.59	0.9125	32.13	0.8967	26.84	0.7488	21.7+16.0
SepCov [39]	RBPN [9]	25.63	0.7546	35.21	0.9244	34.19	0.9231	32.86	0.9090	30.18	0.8706	21.7+12.7
SepCov [39]	EDVR [57]	25.90	0.7785	35.24	0.9252	34.23	0.9240	32.97	0.9112	30.28	0.8745	21.7+20.7
DAIN [1]	Bicubic	23.64	0.6315	32.56	0.8922	30.82	0.8654	29.20	0.8309	26.61	0.7453	24.0
DAIN [1]	RCAN [77]	25.07	0.7281	35.27	0.9242	33.82	0.9146	32.26	0.8974	26.86	0.7489	24.0+16.0
DAIN [1]	RBPN [9]	25.81	0.7576	35.67	0.9304	34.53	0.9262	33.01	0.9097	30.29	0.8709	24.0+12.7
DAIN [1]	EDVR [57]	26.10	0.7828	35.81	0.9323	34.66	0.9280	33.11	0.9119	<b>30.40</b>	0.8749	24.0+20.7
AdaCoF [23]	Bicubic	23.60	0.6287	32.57	0.8902	30.78	0.8641	29.20	0.8308	26.59	0.7449	21.8
AdaCoF [23]	RCAN [77]	24.94	0.6807	35.08	0.9068	33.73	0.9165	31.93	0.8915	26.83	0.7479	21.8+16.0
AdaCoF [23]	RBPN [9]	25.78	0.7546	35.55	0.9266	34.35	0.9236	33.08	0.9108	30.39	<b>0.8751</b>	21.8+12.7
AdaCoF [23]	EDVR [57]	26.02	0.7813	35.81	0.9294	34.56	0.9262	33.21	0.9130	30.26	0.8709	21.8+20.7
STARnet [10]		26.03	0.7818	36.04	0.9342	34.84	0.9298	33.01	0.9086	29.92	0.8589	111.6
Zooming-SloMo [60]		<b>26.31</b>	<b>0.7976</b>	<b>36.81</b>	<b>0.9415</b>	<b>35.41</b>	<b>0.9361</b>	<b>33.36</b>	<b>0.9138</b>	30.34	0.8713	<b>11.1</b>
MEGAN		<b>26.57</b>	<b>0.8044</b>	<b>37.18</b>	<b>0.9446</b>	<b>35.71</b>	<b>0.9389</b>	<b>33.62</b>	<b>0.9171</b>	<b>30.56</b>	<b>0.8753</b>	<b>10.7</b>

Table 1: Quantitative Evaluation of state-of-the-art methods. **Red** and **blue** indicate the best and the second best performance, respectively.

deploys the Pyramid, Cascading and Deformable (PCD) architecture in [57].

**Long-range Memory Graph Aggregation** Based on the inter-connected property between space and time domains, we propose a novel long-range memory graph aggregation (LMGA) module for STVSR, drawing inspiration from the context video object detection [4]. Different from [4], we utilize Graph Convolutional Network (GCN) to learn appropriate temporal correlations among whole videos. This is because GCN well exploits frame interactions by constructing prior temporal graph for better improving the semantic information among different subsets of frame features. Another major difference is that we use enlarged (interpolated) feature framesets to build our global and local pool, while [4] uses Region Proposal Network. Since the enlarged features are often of low quality, our LMGA greatly improves feature quality by using both short and long-range space-time information. In implementations, LMGA includes two operations: frame feature aggregation and graph construction. An example of frame feature aggregation in feature space is shown in Figure 2. We group the single frame features from the deformable ConvLSTM to form the local pool as  $\mathbf{L} = \{E_t^{LR}\}_{t=1}^{2n+1}$ . For long-term global modeling, we randomly shuffle and pick  $\tau$  indices from the ordered index sequence  $\{1, \dots, 2n+1\}$  to create the global pool  $\mathbf{G} = \{E_t^G\}_{t=1}^\tau$ . We denote the enriched feature set as  $\mathbf{M} = \{E_t^M\}_{t=1}^{\varphi=\tau+1}$  by concatenating each local feature and randomly selected global contexts from  $\mathbf{G}$ . To be specific, the aggregation function can be summarized as:

$$\mathbf{M} = N_g(\mathbf{L}, \mathbf{G}), \quad (12)$$

where  $N_g(\cdot)$  is the global-local feature aggregation function containing three convolutional layers. The spatial-temporal

correlations among a collection of frame features can be achieved via a well constructed graph  $\mathcal{G}$ , where every frame feature is a vertex and the corresponding edge feature is a similarity-weighted interrelation of two vertices. To exploit the global and local information of space-time domain for STVSR, we propose to perform graph construction in feature domain iteratively.

The graph construction is composed of  $K$  weighted graph  $\{\mathcal{G}^k(\mathcal{V}, \mathcal{E})\}_{k=1}^K$ . The graph vertices  $\mathcal{V}$  are the enriched feature set  $\mathbf{M}$  with the  $\varphi$  nodes, and the correlation (edges) set  $\mathcal{E}$  is with the size  $|\mathcal{E}| = |\mathcal{V}| \times (\varphi - 1)$ . Following the message-passing algorithms [16, 7], we define the learned edge feature  $A_{p,q}^k$  as the non-linear combination of the absolute difference between any two node features ( $p$  and  $q$ ), e.g.,  $A_{p,q}^k = \phi_\theta(|E_p^k - E_q^k|)$ , where  $\phi$  is a transformation function (e.g. a neural network) with learnable parameters  $\theta$ .  $E_p^k$  is the feature embedding for the node  $p$  of  $\mathcal{G}^k$ . In this work, we use stacked convolutional layers to embed the node features to get the weighted adjacency  $A_{p,q}^k$  between nodes  $p$  and  $q$ . The adjacent matrix is normalized to a stochastic matrix by using a softmax along each row.

Inspired by GCN [41], we aggregate  $\varphi$  node embeddings with its neighbours using *adjacency operator*  $A : E \mapsto A(E)$  where  $A(E_p) := \sum_{q \sim p} A_{p,q} E_q$ , with  $p \sim q$  if and only if  $(p, q) \in \mathcal{E}$ . Let us denote  $E_{p,r}^k$  as the  $r$ -th channel of the embedding of node  $p$  in  $\mathcal{G}^k$ . Consequently, the  $k$ -th GCN layer  $\text{Gc}(\cdot)$  is conducted as:

$$E_{p,r'}^{k+1} = \text{Gc}(E^k) = \rho \left( \sum_{r=1}^{r_k} \sum_{v \in \mathcal{N}_p} A_{p,v} E_{v,r}^k * \theta_{r',r}^k \right) \quad \forall r', p \quad (13)$$

where  $\theta \in \mathbb{R}^{r_k \times r_{k+1} \times b \times b}$  are learned parameters of the




Figure 3: Visual comparison of Space-time SR results produced by different methods. Our MEGAN achieves more visually pleasant results by handling motion blur and motion aliasing in complex dynamic space-time scenes (**Zoomed for visual clarity**).

non-linear transformation.  $r_k$ ,  $r_{k+1}$ , and  $b$  denotes the size of the input features and the output features, 2D convolutional kernel size, respectively.  $\mathcal{N}(p)$  is the neighbour node set of the node  $p$ .  $*$  denotes the convolution operation.  $\rho$  is the operation of Leaky ReLU [63]. To refine the enriched intermediate features set  $\{E_t^{K+1}\}_{t=1}^{2n+1}$ , we use a small network containing two convolutional layers for embedding features, and then concatenate them to obtain  $\{E_t^{K+1}\}_{t=1}^{2n+1}$ . The refined features appears to consistently enable the network with more expressive power, which provides a better starting point for training a model for downstream tasks. As a result, we use LMGA-refined features  $F^M = \{E_t^{K+1}\}_{t=1}^{2n+1}$  as the input for the subsequent layers of the network. Please see Figure 2 for more details.

### 3.4. Progressive Fusion Reconstruction

In the reconstruction process  $\Psi$ , we first utilize  $m_2$  PFRDBs [65] to progressively learn intra-frame spatial representations and inter-frame temporal correlations. After

that, we concatenate  $m_3$  residual blocks [27] to further enable very deep network for video super resolution tasks. For up-sampling, we perform  $\times 2$  upsample *PixelShuffle* operations [44, 27]. The last Conv layer fuses all the feature maps to reconstruct HR video sequence  $I^{HR} = \{I_t^{HR}\}_{t=1}^{2n+1}$ . We obtain the final HR output via the reconstruction module, e.g.,  $I^{HR} = \Psi(F^M)$ .

### 3.5. Loss Function

In the training process, our proposed network is optimized by the following objective function in an end-to-end manner:

$$L_{\text{MEGAN}} = \sqrt{\|I^{GT} - I^{HR}\|^2 + \epsilon^2}, \quad (14)$$

where  $I^{GT}$  denotes the ground-truth HR video sequence. We adopt Charbonnier penalty function [22] as the loss term and empirically set  $\epsilon = 1e^{-3}$ .




Figure 4: Visual comparison. Our MEGAN recovers more detailed textures with sharper edges (**Zoomed for visual clarity**).

### 3.6. Implementation Details

In the proposed MEGAN, we utilize  $m_1 = 3$  Conv layers in the residual non-local attention module. For the progressive fusion reconstruction network, we set  $m_2, m_3 = 5, 20$ . In the LMGA, we set  $\tau$  as 4, and  $k$  as 2. During the training, we select the odd-indexed 4 frames and the corresponding 7 consecutive frames as the LR input and ground-truth HR, respectively. We augment the training data in these ways, such as rotation, flipping, and randomly reversing the direction of video sequences. On the image level, each video frame is down-sampled by 4, and a random crop of  $32 \times 32$  on LR images is performed. We train our network on 2 NVIDIA 1080Ti GPUs using the Adam optimizer [19]. Using a cosine annealing strategy [32], the initial learning rate starts with  $1e^{-4}$ , and is annealed to  $1e^{-7}$  after 20000 iterations. In the training process, we set batch size as 16. For fair comparisons, we evaluate all the methods on the same single Nvidia 1080Ti GPU. The number of parameters of MEGAN is 10.7 Million.

## 4. Experiments

**Experimental Setup** In this study, we utilize the Vimeo-90K dataset [64] for training to demonstrate the fidelity and robustness of our proposed method. The dataset consists of over 60,000 7-frame training video sequences with the resolution of  $256 \times 448$ . The resolution of down-scaled LR frames is  $64 \times 112$ . To evaluate the performance of different methods, we select several datasets, including Vid4 [28], Vimeo [64], and Adobe240 [46] testsets. Following the setting in [8], we split Vimeo testset into fast motion, medium motion, and slow motion subsets. For evaluation, we utilize two widely-used image quality metrics to quantitatively validate the STVSR performance, including peak signal-to-noise ratio (PSNR) and structural similarity index metrics (SSIM) [59].

**Comparison to State-of-the-art** In this study, we compare the proposed MEGAN with the state-of-the-art methods. For clarity, we categorize the methods into one-stage

and two-stage. For one-stage methods, we choose Zooming SlowMo [60] and STARnet [10] for comparison. For two-stage methods, we pick SepConv [39], Super-SloMo [13], DAIN [1], AdaCoF [23] as Time SR and RCAN [77], RBPN [9], EDVR [57], PFRB [65] as Space SR.

We evaluate the proposed method against the state-of-the-art methods on the Vid4 and Vimeo datasets. The typical results are given in Figures 3 and 4. To further evaluate the robustness of our method, the recovered features are zoomed in Figures 3 and 4. We can clearly observe that our MEGAN provides more visually pleasant results with finer details and sharper edges than the competing methods. The quantitative results are summarized in Table 1. We can see that our model obtains 0.26, 0.37, 0.30, 0.26, and 0.22 dB PSNR gains over the best one-stage-based model Zooming-SloMo on Vid4, Vimeo-Fast, Vimeo-Medium, Vimeo-Slow, and Adobe240, and also outperforms the best two-stage method DAIN+EDVR by 0.47, 1.37, 1.05, 0.51 and 0.16 dB on Vid4, Vimeo-Fast, Vimeo-Medium, Vimeo-Slow, and Adobe240 in terms of PSNR, respectively. To better illustrate the robustness for space-time upsampling, we also compare our proposed method with 18 state-of-the-art methods on Adobe240 [28] dataset. As shown in Table 1, the proposed MEGAN consistently outperforms all the evaluated methods. This suggests that the proposed MEGAN is capable of recovering better results with finer details by capturing long-term spatial-temporal information. These comparisons intuitively show that it is essential to access more global and local information from the space-time domain to handle motion blur and motion aliasing in complex dynamic space-time scenes. In other words, the results demonstrate that the proposed MEGAN reconstructs more detailed results with more sharper contexts, as shown in Figures 3 and 4.

## 5. Ablation Study

We conduct a thorough ablation analysis to demonstrate the effectiveness of different key components of MEGAN. Detailed results are found in the supplementary material.




Figure 5: Illustration of effectiveness of LMGA. As shown in the yellow boxes, we can observe that w/ LMGA module recovers more richer details (e.g., headphone)

Method	Vid4		Vimeo-Fast	
	PSNR	SSIM	PSNR	SSIM
w/o LMGA	26.40	0.8001	36.89	0.9419
w/o NLRB	26.44	0.8017	37.01	0.9430
Ours	<b>26.57</b>	<b>0.8044</b>	<b>37.18</b>	<b>0.9446</b>

Table 2: Ablation study on LMGA and NLRB over Vid4 and Vimeo-Fast dataset.

**Effectiveness of LMGA** We first explore the effectiveness of the proposed LMGA block as it is the key component in our model. We establish our baseline network (w/o LMGA model) by removing the LMGA block from our model. The quantitative results are provided in Table 2. The increased performance in terms of PSNR value over the the Vimeo-Fast dataset, from 36.89 to 37.18 dB, demonstrates that the LMGA module can better capture fast motions by aggregating global temporal information. The visual comparisons are depicted in Figure 5. From the visual results, we can see that the LMGA block enables the model to accurately reconstruct frames with more details. In particular, the “w/o LMGA” lower bound fails to realize the headphone on the woman’s head, but our model provides better visualization of the headphone. These observations demonstrate that adding the LMGA block to better utilize the global and local temporal contexts can further encourage our model to obtain more accurate spatial-temporal representations from fast dynamic video sequences.

**Effectiveness of NLRB** We then investigate the effectiveness of the NLRB block. Our baseline network adopts the proposed network w/o NLRB block, i.e., removing two NLRB blocks when training the network. The results in Table 2 show that the deep non-local residual learning strategy helps improve the performances on two benchmark datasets. In Figure 6, it is observed that our model preserves




Figure 6: Illustration of robustness of NLRB. As we see, w/ NLRB is able to contain richer and sharper details from the overlaid LR image especially in the regions of surrounding areas.

better textural details of grasses (shown at left), and captures the human leg motion (shown at right). In comparison, the “w/o NLRB” model has limited ability to recover fine details of grasses in the central area, but the surrounding area is still blurry. These results show that the NLRB block enables the model to capture long-range dependencies in the spatial dimension, while the channel-wise attention enables the model to attend different types of feature representations. NLRB can further encourage the network to retain high-frequency details by incorporating more spatial information. Its support is further illustrated in Figure 6. We can see that NLRB is capable of retraining low- and high-frequency information to reconstruct more realistic video sequences. These observations demonstrate the effectiveness of NLRB in providing more detailed textures.

## 6. Conclusion

In this work, we present a unified framework for space-time video super-resolution. Aided by temporal and spatial information, our approach learns complex spatial-temporal features to reconstruct a video sequence of high space-time resolution. To this end, we propose a novel long-range memory graph aggregation (LMGA) module to enhance the features of key frames with global and local information in the dynamic space-time scene. We also introduce the deep non-local residual learning to enable our model to adaptively learn mixed attention hierarchical features from multiple video sequences. Furthermore, we adopt the progressive fusion reconstruction network to capture long-range spatial-temporal dependencies. In general, the proposed MEGAN has achieved the superior performance among all the involved methods in both static and dynamic scenes.



## References

- [1] Wenbo Bao, Wei-Sheng Lai, Chao Ma, Xiaoyun Zhang, Zhiyong Gao, and Ming-Hsuan Yang. Depth-aware video frame interpolation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3703–3712, 2019.
- [2] Jose Caballero, Christian Ledig, Andrew Aitken, Alejandro Acosta, Johannes Totz, Zehan Wang, and Wenzhe Shi. Real-time video super-resolution with spatio-temporal networks and motion compensation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4778–4787, 2017.
- [3] Nuo Chen, Chenyu You, and Yuexian Zou. Self-supervised dialogue learning for spoken conversational question answering. In *INTERSPEECH*, 2021.
- [4] Yihong Chen, Yue Cao, Han Hu, and Liwei Wang. Memory enhanced global-local aggregation for video object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10337–10346, 2020.
- [5] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Int. Conf. Comput. Vis.*, pages 764–773, 2017.
- [6] Esmaeil Faramarzi, Dinesh Rajan, and Marc P Christensen. Space-time super-resolution from multiple-videos. In *International Conference on Information Science, Signal Processing and their Applications*, pages 23–28. IEEE, 2012.
- [7] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. *arXiv preprint arXiv:1704.01212*, 2017.
- [8] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. Recurrent back-projection network for video super-resolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019.
- [9] Muhammad Haris, Gregory Shakhnarovich, and Norimichi Ukita. Recurrent back-projection network for video super-resolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3897–3906, 2019.
- [10] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. Space-time-aware multi-resolution video enhancement. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2859–2868, 2020.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016.
- [12] Yan Huang, Wei Wang, and Liang Wang. Bidirectional recurrent convolutional networks for multi-frame super-resolution. In *Adv. Neural Inform. Process. Syst.*, pages 235–243, 2015.
- [13] Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik Learned-Miller, and Jan Kautz. Super slomo: High quality estimation of multiple intermediate frames for video interpolation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9000–9008, 2018.
- [14] Younghyun Jo, Seoung Wug Oh, Jaeyeon Kang, and Seon Joo Kim. Deep video super-resolution network using dynamic upsampling filters without explicit motion compensation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3224–3232, 2018.
- [15] Jaeyeon Kang, Younghyun Jo, Seoung Wug Oh, Peter Vajda, and Seon Joo Kim. Deep space-time video upsampling networks. *arXiv preprint arXiv:2004.02432*, 2020.
- [16] Steven Kearnes, Kevin McCloskey, Marc Berndl, Vijay Pande, and Patrick Riley. Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design*, 30(8):595–608, 2016.
- [17] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1646–1654, 2016.
- [18] Soo Ye Kim, Jihyong Oh, and Munchurl Kim. Fisr: Deep joint frame interpolation and super-resolution with a multi-scale temporal loss. *arXiv preprint arXiv:1912.07213*, 2019.
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [20] Deying Kong, Yifei Chen, Haoyu Ma, Xiangyi Yan, and Xiaohui Xie. Adaptive graphical model network for 2d hand-pose estimation. *arXiv preprint arXiv:1909.08205*, 2019.
- [21] Deying Kong, Haoyu Ma, and Xiaohui Xie. Sia-gcn: A spatial information aware graph neural network with 2d convolutions for hand pose estimation. *arXiv preprint arXiv:2009.12473*, 2020.
- [22] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 624–632, 2017.
- [23] Hyeongmin Lee, Taehoh Kim, Tae-young Chung, Daehyun Pak, Yuseok Ban, and Sangyoun Lee. Adacof: Adaptive collaboration of flows for video frame interpolation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5316–5325, 2020.
- [24] Guang Li, Shouhua Luo, Chenyu You, Matthew Getzin, Liang Zheng, Ge Wang, and Ning Gu. A novel calibration method incorporating nonlinear optimization and ball-bearing markers for cone-beam ct with a parameterized trajectory. *Medical physics*, 46(1):152–164, 2019.
- [25] Kaiyan Li, Weimin Zhou, Hua Li, and Mark A Anastasio. Assessing the impact of deep neural network-based image denoising on binary signal detection tasks. *IEEE Transactions on Medical Imaging*, 2021.
- [26] Tao Li, Xiaohai He, Qizhi Teng, Zhengyong Wang, and Chao Ren. Space-time super-resolution with patch group cuts prior. *Signal Processing: Image Communication*, 30:147–165, 2015.
- [27] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 136–144, 2017.
- [28] Ce Liu and Deqing Sun. A bayesian approach to adaptive video super resolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 209–216. IEEE, 2011.
- [29] Ding Liu, Zhaowen Wang, Yuchen Fan, Xianming Liu, Zhangyang Wang, Shiyu Chang, and Thomas Huang. Robust video super-resolution with learned temporal dynamics. In *Int. Conf. Comput. Vis.*, 2017.

- [30] Ding Liu, Bihan Wen, Yuchen Fan, Chen Change Loy, and Thomas S Huang. Non-local recurrent network for image restoration. In *Adv. Neural Inform. Process. Syst.*, 2018.
- [31] Ziwei Liu, Raymond A Yeh, Xiaoou Tang, Yiming Liu, and Aseem Agarwala. Video frame synthesis using deep voxel flow. In *Int. Conf. Comput. Vis.*, pages 4463–4471, 2017.
- [32] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [33] Qing Lyu, Chenyu You, Hongming Shan, and Ge Wang. Super-resolution mri through deep learning. *arXiv preprint arXiv:1810.06776*, 2018.
- [34] Simone Meyer, Abdelaziz Djelouah, Brian McWilliams, Alexander Sorkine-Hornung, Markus Gross, and Christopher Schroers. Phasenet for video frame interpolation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 498–507, 2018.
- [35] Simone Meyer, Oliver Wang, Henning Zimmer, Max Grosse, and Alexander Sorkine-Hornung. Phase-based frame interpolation for video. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1410–1418, 2015.
- [36] Uma Mudenagudi, Subhashis Banerjee, and Prem Kumar Kalra. Space-time super-resolution using graph-cut optimization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(5):995–1008, 2010.
- [37] Simon Niklaus and Feng Liu. Context-aware synthesis for video frame interpolation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1701–1710, 2018.
- [38] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive convolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 670–679, 2017.
- [39] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive separable convolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 261–270, 2017.
- [40] Mehdi SM Sajjadi, Raviteja Vemulapalli, and Matthew Brown. Frame-recurrent video super-resolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6626–6634, 2018.
- [41] Victor Garcia Satorras and Joan Bruna Estrach. Few-shot learning with graph neural networks. In *Int. Conf. Learn. Represent.*, 2018.
- [42] O Shahar, A Faktor, and M Irani. Space-time super-resolution from a single video. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3353–3360, 2011.
- [43] Eli Shechtman, Yaron Caspi, and Michal Irani. Space-time super-resolution. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(4):531–545, 2005.
- [44] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1874–1883, 2016.
- [45] Kun Su, Xiulong Liu, and Eli Shlizerman. Audeo: Audio generation for a silent performance video. In *Adv. Neural Inform. Process. Syst.*, volume 33, 2020.
- [46] Shuochen Su, Mauricio Delbracio, Jue Wang, Guillermo Sapiro, Wolfgang Heidrich, and Oliver Wang. Deep video deblurring for hand-held cameras. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1279–1288, 2017.
- [47] Ying Tai, Jian Yang, and Xiaoming Liu. Image super-resolution via deep recursive residual network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3147–3155, 2017.
- [48] Hao Tang, Xuming Chen, Yang Liu, Zhipeng Lu, Junhua You, Mingzhou Yang, Shengyu Yao, Guoqi Zhao, Yi Xu, Tingfeng Chen, et al. Clinically applicable deep learning framework for organs at risk delineation in ct images. *Nature Machine Intelligence*, 1(10):480–491, 2019.
- [49] Hao Tang, Daniel R Kim, and Xiaohui Xie. Automated pulmonary nodule detection using 3d deep convolutional neural networks. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 523–526. IEEE, 2018.
- [50] Hao Tang, Xingwei Liu, Kun Han, Xiaohui Xie, Xuming Chen, Huang Qian, Yong Liu, Shanlin Sun, and Narisu Bai. Spatial context-aware self-attention model for multi-organ segmentation. In *IEEE Winter Conf. Appl. Comput. Vis.*, pages 939–949, 2021.
- [51] Hao Tang, Xingwei Liu, Shanlin Sun, Xiangyi Yan, and Xiaohui Xie. Recurrent mask refinement for few-shot medical image segmentation. *arXiv preprint arXiv:2108.00622*, 2021.
- [52] Hao Tang, Xingwei Liu, and Xiaohui Xie. An end-to-end framework for integrated pulmonary nodule detection and false positive reduction. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pages 859–862. IEEE, 2019.
- [53] Hao Tang, Chupeng Zhang, and Xiaohui Xie. Nodulenet: Decoupled false positive reduction for pulmonary nodule detection and segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 266–274. Springer, 2019.
- [54] Xin Tao, Hongyun Gao, Renjie Liao, Jue Wang, and Jiaya Jia. Detail-revealing deep video super-resolution. In *Int. Conf. Comput. Vis.*, pages 4472–4480, 2017.
- [55] Yapeng Tian, Yulun Zhang, Yun Fu, and Chenliang Xu. Tdan: Temporally deformable alignment network for video super-resolution. *arXiv preprint arXiv:1812.02898*, 2018.
- [56] Min-Yuan Tseng, Yen-Chung Chen, Yi-Lun Lee, Wei-Sheng Lai, Yi-Hsuan Tsai, and Wei-Chen Chiu. Dual-stream fusion network for spatiotemporal video super-resolution. In *IEEE Winter Conf. Appl. Comput. Vis.*, pages 2684–2693, 2021.
- [57] Xintao Wang, Kelvin C.K. Chan, Ke Yu, Chao Dong, and Chen Change Loy. Edvr: Video restoration with enhanced deformable convolutional networks. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, June 2019.
- [58] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7794–7803, 2018.
- [59] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4):600–612, 2004.
- [60] Xiaoyu Xiang, Yapeng Tian, Yulun Zhang, Yun Fu, Jan P Allebach, and Chenliang Xu. Zooming slow-mo: Fast and

- accurate one-stage space-time video super-resolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3370–3379, 2020.
- [61] Zeyu Xiao, Xueyang Fu, Jie Huang, Zhen Cheng, and Zhiwei Xiong. Space-time distillation for video super-resolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2113–2122, 2021.
- [62] Zeyu Xiao, Zhiwei Xiong, Xueyang Fu, Dong Liu, and Zheng-Jun Zha. Space-time video super-resolution using temporal profiles. In *ACM Int. Conf. Multimedia*, pages 664–672, 2020.
- [63] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*, 2015.
- [64] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *Int. J. Comput. Vis.*, 127(8):1106–1125, 2019.
- [65] Peng Yi, Zhongyuan Wang, Kui Jiang, Junjun Jiang, and Jiayi Ma. Progressive fusion video super-resolution network via exploiting non-local spatio-temporal correlations. In *Int. Conf. Comput. Vis.*, pages 3106–3115, 2019.
- [66] Chenyu You, Nuo Chen, Fenglin Liu, Dongchao Yang, and Yuexian Zou. Towards data distillation for end-to-end spoken conversational question answering. *arXiv preprint arXiv:2010.08923*, 2020.
- [67] Chenyu You, Nuo Chen, and Yuexian Zou. Contextualized attention-based knowledge transfer for spoken conversational question answering. In *INTERSPEECH*, 2021.
- [68] Chenyu You, Nuo Chen, and Yuexian Zou. Knowledge distillation for improved accuracy in spoken question answering. In *IEEE Int. Conf. Acoust. Speech Signal Process.*, 2021.
- [69] Chenyu You, Nuo Chen, and Yuexian Zou. MRD-Net: Multi-Modal Residual Knowledge Distillation for Spoken Question Answering. In *IJCAI*, 2021.
- [70] Chenyu You, Nuo Chen, and Yuexian Zou. Self-supervised contrastive cross-modality representation learning for spoken question answering. *arXiv preprint arXiv:2109.03381*, 2021.
- [71] Chenyu You, Guang Li, Yi Zhang, Xiaoliu Zhang, Hongming Shan, Mengzhou Li, Shenghong Ju, Zhen Zhao, Zhuiyang Zhang, Wenxiang Cong, et al. CT super-resolution GAN constrained by the identical, residual, and cycle learning ensemble (gan-circle). *IEEE Trans. Med. Imag.*, 39(1):188–203, 2019.
- [72] Chenyu You, Junlin Yang, Julius Chapiro, and James S. Duncan. Unsupervised wasserstein distance guided domain adaptation for 3d multi-domain liver segmentation. In *Interpretable and Annotation-Efficient Learning for Medical Image Computing*, pages 155–163. Springer International Publishing, 2020.
- [73] Chenyu You, Linfeng Yang, Yi Zhang, and Ge Wang. Low-Dose CT via Deep CNN with Skip Connection and Network in Network. In *Developments in X-Ray Tomography XII*, volume 11113, page 111131W. International Society for Optics and Photonics, 2019.
- [74] Chenyu You, Qingsong Yang, Lars Gjestebj, Guang Li, Shenghong Ju, Zhuiyang Zhang, Zhen Zhao, Yi Zhang, Wenxiang Cong, Ge Wang, et al. Structurally-sensitive multi-scale deep neural network for low-dose CT denoising. *IEEE Access*, 6:41839–41855, 2018.
- [75] Chenyu You, Ruihan Zhao, Lawrence Staib, and James S Duncan. Momentum contrastive voxel-wise representation learning for semi-supervised volumetric medical image segmentation. *arXiv preprint arXiv:2105.07059*, 2021.
- [76] Chenyu You, Yuan Zhou, Ruihan Zhao, Lawrence Staib, and James S. Duncan. Simcvd: Simple contrastive voxel-wise representation distillation for semi-supervised medical image segmentation. *arXiv preprint arXiv:2108.06227*, 2021.
- [77] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2472–2481, 2018.
- [78] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9308–9316, 2019.