

IFRNet: Intermediate Feature Refine Network for Efficient Frame Interpolation

Lingtong Kong^{1*}, Boyuan Jiang^{2*}, Donghao Luo², Wenqing Chu², Xiaoming Huang²,
Ying Tai², Chengjie Wang², Jie Yang^{1†}

¹Shanghai Jiao Tong University, China, ²Youtu Lab, Tencent

{ltkong, jieyang}@sjtu.edu.cn

{byronjiang, michaello, wenqingchu, skyhuang, yingtai, jasoncjwang}@tencent.com

Abstract

Prevailing video frame interpolation algorithms, that generate the intermediate frames from consecutive inputs, typically rely on complex model architectures with heavy parameters or large delay, hindering them from diverse real-time applications. In this work, we devise an efficient encoder-decoder based network, termed IFRNet, for fast intermediate frame synthesizing. It first extracts pyramid features from given inputs, and then refines the bilateral intermediate flow fields together with a powerful intermediate feature until generating the desired output. The gradually refined intermediate feature can not only facilitate intermediate flow estimation, but also compensate for contextual details, making IFRNet do not need additional synthesis or refinement module. To fully release its potential, we further propose a novel task-oriented optical flow distillation loss to focus on learning the useful teacher knowledge towards frame synthesizing. Meanwhile, a new geometry consistency regularization term is imposed on the gradually refined intermediate features to keep better structure layout. Experiments on various benchmarks demonstrate the excellent performance and fast inference speed of proposed approaches. Code is available at <https://github.com/ltkong218/IFRNet>.

1. Introduction

Video frame interpolation (VFI), that converts low frame rate (LFR) image sequences to high frame rate (HFR) videos is an important low-level computer vision task. Related techniques are widely applied to various practical applications, such as slow-motion generation [22], novel view synthesis [55] and cartoon creation [42]. Although it has been studied by a large number of researches, there are still

* Equal contribution. This work was done when Lingtong Kong was an intern at Tencent Youtu Lab.

† Corresponding author: Jie Yang (jieyang@sjtu.edu.cn). This research is partly supported by NSFC, China (No: 61876107, U1803261).




Figure 1. **Speed, accuracy and parameters comparison.** Proposed IFRNet achieves state-of-the-art frame interpolation accuracy with fast inference speed and lightweight model size.

great challenges when dealing with complicated dynamic scenes, which include large displacement, severe occlusion, motion blur and abrupt brightness change.

Recently, with the development of optical flow networks [13, 24, 45, 46], significant progress has been made by flow-based VFI approaches [22, 33, 37, 49], since optical flow can provide an explicit correspondence to register frames in a video sequence. Successful flow-based approaches usually follow a three-step pipeline: **1)** Estimate optical flow between target frame and input frames. **2)** Warp input frames or context features by predicted flow fields for spatial alignment. **3)** Refine warped frames or features and generate the target frame by a synthesis network. Denoting input frames and target frame to be I_0, I_1 and I_t ($0 < t < 1$), existing methods either first estimate optical flow $F_{0 \rightarrow 1}, F_{1 \rightarrow 0}$ [3, 22, 32, 33, 36], and then approximate or refine bilateral intermediate flow $F_{t \rightarrow 0}, F_{t \rightarrow 1}$ [9, 22, 40, 49] as shown in Figure 2 (a), or throw the intractable intermediate flow estimation sub-task to a learnable flow network for end-to-end training [20, 50, 54] as depicted in Figure 2 (b). Their common step is to further employ an image synthesis network to encode spatial aligned context feature [32] for target frame generation or refinement.




Figure 2. **Different flow-based VFI paradigms.** We roughly classify existing flow-based VFI methods based on encoder-decoders with specific function. In (a) [3, 22, 32, 33, 36, 37, 40, 49], FlowNet estimates conventional optical flow $F_{0 \rightarrow 1}, F_{1 \rightarrow 0}$, the middle part approximates or further refines flow fields $F_{t \rightarrow 0}, F_{t \rightarrow 1}$. In (b) [20, 50, 54], the Intermediate FlowNet directly predicts intermediate flow of $F_{t \rightarrow 0}, F_{t \rightarrow 1}$. Both (a) and (b) contain a separate synthesis network for target frame generation. In (c), proposed IFRNet jointly refines the intermediate flow $F_{t \rightarrow 0}, F_{t \rightarrow 1}$ together with a powerful intermediate feature ϕ_t to generate the target frame in a single encoder-decoder.

Although above pipeline that first estimates intermediate flow and then context feature has become the most popular paradigm for flow-based VFI approaches [9, 32, 33, 37, 40], it suffers from several defects. First, they divide intermediate flow and context feature refinement into separate encoder-decoders, which ignores the mutual promotion of these two crucial elements for frame interpolation. Second, their cascaded architecture based on above design concept can substantially increase the inference delay and model parameters, blocking them from mobile and real-time applications.

In this paper, we propose a novel Intermediate Feature Refine Network (IFRNet) for VFI to overcome the above limitations. For the first time, we merge above separated flow estimation and feature refinement into a *single encoder-decoder based model for compactness and fast inference*, abstracted in Figure 2 (c). It first extracts pyramid features from given inputs by the encoder, and then jointly refines the bilateral intermediate flow fields together with a powerful intermediate feature through coarse-to-fine decoders. The improved architecture can benefit intermediate flow and intermediate feature with each other, endowing our model with the ability to not only generate sharper moving objects but also capture better texture details.

For better supervision, we propose task-oriented flow distillation loss and feature space geometry consistency loss to effectively guide the multi-scale motion estimation and intermediate feature refinement. Specifically, our flow distillation approach adjusts the robustness of distillation loss adaptively in space and focuses on learning the useful teacher knowledge for frame synthesizing. Besides, proposed geometry consistency loss can employ the extracted intermediate features from ground truth to constrain the reconstructed intermediate features for keeping better structure layout. Figure 1 gives a speed, accuracy and parameters comparison among advanced VFI methods, demonstrating the state-of-the-art performance of our approaches. In summary, our main contributions are listed as follows:

- We devise a novel IFRNet to jointly perform intermediate flow estimation and intermediate feature refinement for efficient video frame interpolation.
- Task-oriented flow distillation loss and feature space geometry consistency loss are newly proposed to promote intermediate motion estimation and intermediate feature reconstruction of IFRNet, respectively.

- Benchmark results demonstrate that our IFRNet not only achieves state-of-the-art VFI accuracy, but also enjoys fast inference speed and lightweight model size.

2. Related Work

Video Frame Interpolation. The mainstream VFI methods can be classified into flow-based [3, 22, 29, 32, 33, 36, 37, 40, 49–51, 54], kernel-based [7, 8, 12, 25, 34, 35, 38] and hallucination-based approaches [10, 16, 23]. Different VFI paradigms have their own merits and flaws due to the substantial frame synthesizing manner. For example, kernel-based methods are good at handling motion blur by convolving over local patches [34, 35], successive works mainly extend it to deal with high resolution videos [38], increase the degrees of freedom for convolution kernel [7, 8, 25], or combine them with other paradigms for compensation [4, 12]. However, they are typically computationally expensive and short of dealing with occlusion. In another way, hallucination-based methods directly synthesize frames from the feature domain by blending field-of-view features generated by deformable convolution [11] or PixelShuffle operations [10]. They can naturally generate complex contextual details, while the predicted frames tend to be blurry when fast-moving objects exist.

Recently, significant progress has been made by flow-based VFI approaches, since optical flow can provide an explicit correspondence for frame registration. These solutions either employ an off-the-shelf flow model [32, 49] or estimate task-specific flow [22, 29, 37, 40, 50] as a guidance for pixel-level motion. Common subsequent step is to forward [14] or backward [48] warp input images to target frame, and finally refine warped frames by an image synthesis network [12, 32, 33, 37], often instantiated as a GridNet [15]. For achieving better image interpolation quality, more complicated deep models are devised to estimate intermediate flow fields [9, 49] and refine the generated target frame [22, 33, 36, 37]. However, the heavy computation cost and large inference delay make them unsuitable for resource limited devices. To take a breath from above module cascading competition, and reconsider the improvement of prior efficient flow-based VFI paradigm, *e.g.* DVF [29], we propose a novel single encoder-decoder based IFRNet, that can perform real-time inference with excellent accuracy.




Figure 3. **Architecture overview and loss functions of IFRNet.** Our model is an efficient encoder-decoder based network, which first extracts pyramid context features from input frames with a shared encoder, and then gradually refines bilateral intermediate flow fields $F_{t \rightarrow 0}^k, F_{t \rightarrow 1}^k$ together with reconstructed intermediate feature $\hat{\phi}_t^k$ through coarse-to-fine decoders, until yielding the final output. Besides the common image reconstruction loss \mathcal{L}_r , task-oriented flow distillation loss \mathcal{L}_d and feature space geometry consistency loss \mathcal{L}_g are newly devised to guide the feature alignment procedure more effectively towards intermediate frame synthesizing.

Optical Flow Estimation. Finding dense correspondence between adjacent frames, namely optical flow estimation [19], has been studied for decades for its fundamental role in many downstream video processing tasks [5, 52]. FlowNet [13] is the first attempt to apply deep learning for optical flow estimation based on the encoder-decoder U-shape network. Inspired by traditional coarse-to-fine paradigm, SPyNet [39], PWC-Net [45] and FastFlowNet [24] integrate pyramid feature, backward warping and achieve impressive real-time performance. Knowledge distillation [18] also plays an important role in optical flow prediction, usually embodied as generating pseudo label in unsupervised optical flow learning [26, 27] or related tasks [1, 41]. A recent VFI method [20] also uses a distillation strategy to promote motion prediction. Beyond the difference of architecture design, our distillation approach can focus on the useful knowledge for intermediate frame synthesizing in a task adaptive manner.

3. Proposed Approach

In this section, we first introduce the IFRNet architecture built on the principle of joint refinement of intermediate flow and intermediate feature, to obtain an efficient encoder-decoder based framework for VFI. Then two novel objective functions, *i.e.*, task-oriented flow distillation loss and feature space geometry consistency loss are introduced to help our model achieve excellent performance.

3.1. IFRNet

Given two input frames I_0 and I_1 at adjacent time instances, video frame interpolation aims to synthesize an intermediate frame I_t , where $0 < t < 1$. To achieve this goal, proposed model performs a first extraction phase so as to retrieve a pyramid of features from each frame, then in a coarse-to-fine manner it progressively refines bilateral intermediate flow fields together with reconstructed intermediate feature until reaching the highest level of the pyramid to obtain the final output. Figure 3 sketches the overall architecture of proposed IFRNet.

Pyramid Encoder. To obtain contextual representation from each input frame, we design a compact encoder \mathcal{E} to extract a pyramid of features. Purposely, the parameter shared encoder is built of a block of two 3×3 convolutions in each pyramid level, respectively with strides 2 and 1. As shown in Figure 3, IFRNet extracts 4 levels of pyramid features, counting 8 convolution layers, each followed by a PReLU activation [17]. By gradually decimating the spatial size, it increases the feature channels to 32, 48, 72 and 96, generating pyramid features ϕ_0^k, ϕ_1^k in level k ($k \in \{1, 2, 3, 4\}$) for frames I_0 and I_1 , respectively.

Coarse-to-Fine Decoders. After extracting meaningful hierarchical representations, we then gradually refine intermediate flow fields through multiple decoders by backward warping pyramid features ϕ_0^k, ϕ_1^k to generate $\tilde{\phi}_0^k, \tilde{\phi}_1^k$ according to $F_{t \rightarrow 0}^k$ and $F_{t \rightarrow 1}^k$, respectively. The main advantage of




Figure 4. Details of the decoder in each pyramid level.

coarse-to-fine warping strategy consists of computing easier residual flow at each scale. Different from previous VFI approaches containing post-refinement [12, 20, 33, 37], we explore to improve the bilateral flow prediction during its coarse-to-fine procedure for higher efficiency. Specifically, we make each decoder \mathcal{D}^{k+1} output a higher level reconstructed intermediate feature $\hat{\phi}_t^k$ besides bilateral flow fields $F_{t \rightarrow 0}^k, F_{t \rightarrow 1}^k$, which can fill up the missing reference information to facilitate motion estimation. On the other hand, better predicted flow fields $F_{t \rightarrow 0}^k, F_{t \rightarrow 1}^k$ will align source pyramid features to the target position more precisely, thus, generating better $\tilde{\phi}_0^k, \tilde{\phi}_1^k$, which can in turn improve higher level intermediate feature reconstruction. Therefore, decoders in proposed IFRNet can jointly refine bilateral intermediate flow fields together with reconstructed intermediate feature, benefitting each other until reaching desired output. Moreover, the gradually refined intermediate feature, containing bilateral occlusion and global context information, can finally generate fusion mask and compensate for motion details, that are often missing by flow-based methods, enabling IFRNet a powerful encoder-decoder VFI architecture without additional refinement [33, 37].

Concretely, in each pyramid level, we stack corresponding input features into a holistic volume that is forwarded by a compact decoder network \mathcal{D}^k , consisting of a block of six 3×3 convolutions and one 4×4 deconvolution, with strides 1 and $1/2$, respectively. A PReLU [17] follows each convolution layer. Details of each decoder is shown in Figure 4. In order to keep relative large receptive field and channel numbers for motion estimation and feature encoding while maintaining efficiency, we modify the third and the fifth convolution to update only partial channels of previous output tensor. Furthermore, residual connection and interlaced placement can promote information propagation and joint refinement. More details are shown in supplementary. Note that inputs of \mathcal{D}^4 and outputs of \mathcal{D}^1 are different from other decoders due to the task-related characteristics. In summary, features among decoders can be computed by

$$[F_{t \rightarrow 0}^3, F_{t \rightarrow 1}^3, \hat{\phi}_t^3] = \mathcal{D}^4([\phi_0^4, \phi_1^4, T]), \quad (1)$$

$$[F_{t \rightarrow 0}^{k-1}, F_{t \rightarrow 1}^{k-1}, \hat{\phi}_t^{k-1}] = \mathcal{D}^k([F_{t \rightarrow 0}^k, F_{t \rightarrow 1}^k, \hat{\phi}_t^k, \tilde{\phi}_0^k, \tilde{\phi}_1^k]), \quad (2)$$

$$[F_{t \rightarrow 0}, F_{t \rightarrow 1}, M, R] = \mathcal{D}^1([F_{t \rightarrow 0}^1, F_{t \rightarrow 1}^1, \hat{\phi}_t^1, \tilde{\phi}_0^1, \tilde{\phi}_1^1]), \quad (3)$$

where \mathcal{D}^k ($k = 2, 3$) stand for decoders at middle pyramid levels, $[\cdot]$ denotes concatenation operation. T is a one-channel conditional input for arbitrary time interpolation,

whose values are all the same and set to t . M is a one-channel merge mask exported by a sigmoid layer whose elements range from 0 to 1, and R is a three-channel image residual that can compensate for details. Finally, we can synthesize the desired frame \hat{I}_t by following formulation

$$\hat{I}_t = M \odot \tilde{I}_0 + (1 - M) \odot \tilde{I}_1 + R, \quad (4)$$

$$\tilde{I}_0 = w(I_0, F_{t \rightarrow 0}), \tilde{I}_1 = w(I_1, F_{t \rightarrow 1}), \quad (5)$$

where w means backward warping, \odot is element-wise multiplication. M adjusts the mixing ratio according to bidirectional occlusion information, while R compensates for some details when flow-based generation is unreliable, such as regions of target frame are occluded in both views.

Discussion with Optical Flow Networks. Different from the coarse-to-fine pipeline in real-time optical flow [24, 45] which mainly deals with large displacement matching challenge, in video interpolation, since the target frame is missing, its motion estimation becomes a “chicken-and-egg” problem. Therefore, decoders of IFRNet reconstruct intermediate feature besides intermediate flow fields, performing spatio-temporal feature aggregation and intermediate motion refinement jointly to benefit from each other.

Image Reconstruction Loss. According to above analysis, an efficient IFRNet has been designed for VFI, which is end-to-end trainable. For the purpose of generating intermediate frame, we employ the same image reconstruction loss \mathcal{L}_r as [37] between network output \hat{I}_t and ground truth frame I_t^{gt} , which is the sum of two terms and denoted by

$$\mathcal{L}_r = \rho(\hat{I}_t - I_t^{gt}) + \mathcal{L}_{cen}(\hat{I}_t, I_t^{gt}), \quad (6)$$

where $\rho(x) = (x^2 + \epsilon^2)^\alpha$ with $\alpha = 0.5, \epsilon = 10^{-3}$ is the Charbonnier loss [6] severing as a surrogate for the \mathcal{L}_1 loss. While \mathcal{L}_{cen} is the census loss, which calculates the soft Hamming distance between census-transformed [31] image patches of size 7×7 .

3.2. Task-Oriented Flow Distillation Loss

Training IFRNet with above reconstruction loss \mathcal{L}_r can already perform intermediate frame synthesizing. However, the simple optimization target usually drops into local minimum, since illuminance cases are often challenging, *i.e.*, extreme brightness and repetitive texture regions. To deal with this problem, we try to adopt the knowledge distillation [18] strategy to guide multi-scale intermediate flow estimation of IFRNet by an off-the-shelf teacher flow network, that helps to align multi-scale pyramid features explicitly. In practice, the pre-trained teacher is only used during training, and we calculate its flow prediction as pseudo label $F_{t \rightarrow 0}^p, F_{t \rightarrow 1}^p$ in advance for efficiency. Note that RIFE [20] also uses flow distillation. However, their indiscriminate distillation manner usually learns undesired noise existed in pseudo label. Even if ground truth is available, optical flow itself is often

a sub-optimal representation for specific video task [50]. To overcome above limitations, we propose task-oriented flow distillation loss that can decrease the adverse impacts while focusing on the useful knowledge for better VFI.

Observing that $F_{t \rightarrow 0}, F_{t \rightarrow 1}$ which directly control frame synthesis are sensitive to harmful information in pseudo label. Therefore, we impose multi-scale flow distillation except for the decoder \mathcal{D}^1 , and leave its flow prediction totally constrained by the reconstruction loss \mathcal{L}_r in a task-oriented manner [50]. Furthermore, we can compare above relaxed flow prediction $F_{t \rightarrow 0}, F_{t \rightarrow 1}$ with pseudo label $F_{t \rightarrow 0}^p, F_{t \rightarrow 1}^p$ to calculate robustness masks P_0, P_1 , and use them to adjust the robustness of distillation loss spatially in lower multiple scales for better task-oriented flow distillation, whose procedure is depicted in Figure 3. Specifically, we can obtain $P_l (l \in \{0, 1\})$ by the following formulation

$$P_l = \exp(-\beta |F_{t \rightarrow l} - F_{t \rightarrow l}^p|_{epe}), \quad (7)$$

where $|\cdot|_{epe}$ calculates per-pixel end-point-error, the coefficient β controlling sensibility for robustness is set to 0.3 according to grid search. Foundation of above operations is based on the assumption that task-oriented flow generally agrees with true optical flow but differs in some details.

Following previous experience [21, 44], our task-oriented flow distillation employs the generalized Charbonnier loss $\rho(x) = (x^2 + \epsilon^2)^\alpha$ for better robust learning of intermediate flow, where parameters ϵ and α control the robustness of this loss. Formally, it can be written as

$$\mathcal{L}_d = \sum_{k=1}^3 \sum_{l=0}^1 \rho(\mathcal{U}_{2^k}(F_{t \rightarrow l}^k) - F_{t \rightarrow l}^p), \quad (8)$$

where \mathcal{U}_s is the bilinear upsampling operation with scale factor s . However, different from the fixed format like previous methods [21, 44], we make it adjustable about VFI task by letting ϵ and α be functions of the robustness parameter p , where $p \in (0, 1]$ means the robustness value of any position in aforementioned robustness masks P_0, P_1 . In general, we employ the linear and exponential linear functions to generate α and ϵ separately as follows

$$\alpha = p/2, \quad \epsilon = 10^{-(10p-1)/3}. \quad (9)$$

The coefficients are selected based on two typical cases. For example, when $p = 1.0$, $\rho(x)$ becomes the surrogate \mathcal{L}_1 loss in Eq. 6. And when $p = 0.4$, it turns to be the robust loss used in LiteFlowNet [21]. Figure 5 gives some intuitive examples of this adaptive robust loss. Comprehensively speaking, in each spatial location, if the task-oriented flow prediction of decoder \mathcal{D}^1 is consistent with that in pseudo label, the gradient of the adaptive distillation loss is relatively steep, which tends to distill this helpful information to the bottom three decoders by common gradient descent optimizer. On the other hand, the loss will become more robust to downgrade this relatively harmful flow knowledge.




Figure 5. **Task-oriented flow distillation loss.** It takes the format of generalized Charbonnier loss, while the concrete form in each location is controlled by the corresponding robustness parameter p , which is determined by Eq. 7 to acquire task adaptive ability.

3.3. Feature Space Geometry Consistency Loss

Besides above task-oriented flow distillation loss for facilitating multi-scale intermediate flow estimation, better supervision of intermediate feature is preferred for further improvement. Observing that extracted pyramid features ϕ_0^k, ϕ_1^k by the encoder \mathcal{E} , in a sense, play an equivalent role as the reconstructed intermediate feature $\hat{\phi}_t^k$ from the decoder \mathcal{D}^{k+1} , we try to employ the same parameter shared encoder \mathcal{E} to extract a pyramid of features ϕ_t^k from ground truth frame I_t^{gt} , and use ϕ_t^k to regularize the reconstructed intermediate feature $\hat{\phi}_t^k$ in multi-scale feature domain.

Intuitively, we can adopt the commonly used \mathcal{L}_1 loss to restrict $\hat{\phi}_t^k$ to be close to ϕ_t^k . However, the overtighten constraint will harm the global context and occlusion information contained in reconstructed intermediate feature $\hat{\phi}_t^k$. To relax it and inspired by the local geometry alignment property of census transform [53], we extend the census loss \mathcal{L}_{cen} [31] into multi-scale feature space for progressive supervision, where the soft Hamming distance is calculated between census-transformed corresponding feature maps with 3×3 patches in a channel-by-channel manner. Formally, this loss can be written as

$$\mathcal{L}_g = \sum_{k=1}^3 \mathcal{L}_{cen}(\hat{\phi}_t^k, \phi_t^k). \quad (10)$$

Our motivation is that the extracted pyramid feature, containing useful low-level structure information for frame synthesizing, can regularize the reconstructed intermediate feature to keep better geometry layout. For each spatial location, \mathcal{L}_g only constrain the geometry of its neighbor local patch in every feature map. Consequently, there is no restriction on the channel-wise representation for $\hat{\phi}_t^k$ to encode bilateral occlusion and residual information.

Based on above analysis, our final loss function, containing three parts for joint optimization, is formulated as

$$\mathcal{L} = \mathcal{L}_r + \lambda \mathcal{L}_d + \eta \mathcal{L}_g, \quad (11)$$

where weighting parameters are set to $\lambda = 0.01, \eta = 0.01$.

Method	Vimeo90K	UCF101	SNU-FILM				Time (s)	Params (M)	FLOPs (T)
			Easy	Medium	Hard	Extreme			
SepConv [35]	33.79/0.9702	34.78/0.9669	39.41/0.9900	34.97/0.9762	29.36/0.9253	24.31/0.8448	0.065	21.7	0.36
CAIN [10]	34.65/0.9730	34.91/0.9690	39.89/0.9900	35.61/0.9776	29.90/0.9292	24.78/0.8507	0.069	42.8	1.29
AdaCoF [25]	34.47/0.9730	34.90/0.9680	39.80/0.9900	35.05/0.9754	29.46/0.9244	24.31/0.8439	0.054	21.8	0.36
RIFE [20]	35.62/0.9780	35.28/0.9690	40.06/ 0.9907	35.75/0.9789	30.10/0.9330	24.84/0.8534	0.026	9.8	0.20
IFRNet	35.80/0.9794	35.29/ <u>0.9693</u>	40.03/0.9905	<u>35.94/0.9793</u>	30.41/0.9358	25.05/0.8587	0.025	5.0	0.21
IFRNet small	35.59/0.9786	35.28/0.9691	39.96/0.9905	35.92/0.9792	30.36/0.9357	25.05/0.8582	0.019	<u>2.8</u>	0.12
ToFlow [50]	33.73/0.9682	34.58/0.9667	39.08/0.9890	34.39/0.9740	28.44/0.9180	23.39/0.8310	0.152	1.4	0.62
CyclicGen [28]	32.09/0.9490	35.11/0.9684	37.72/0.9840	32.47/0.9554	26.95/0.8871	22.70/0.8083	0.161	19.8	1.77
DAIN [3]	34.71/0.9756	34.99/0.9683	39.73/0.9902	35.46/0.9780	30.17/0.9335	25.09/0.8584	1.033	24.0	5.51
SoftSplat [33]	36.10/0.9700	<u>35.39/0.9520</u>	-	-	-	-	0.195	12.2	0.90
BMBC [36]	35.01/0.9764	35.15/0.9689	39.90/0.9902	35.31/0.9774	29.33/0.9270	23.92/0.8432	3.845	11.0	2.50
CDFI full [12]	35.17/0.9640	35.21/0.9500	40.12/0.9906	35.51/0.9778	29.73/0.9277	24.53/0.8476	0.380	5.0	0.82
ABME [37]	<u>36.18/0.9805</u>	35.38/ 0.9698	39.59/0.9901	35.77/0.9789	<u>30.58/0.9364</u>	25.42/0.8639	0.905	18.1	1.30
IFRNet large	36.20/0.9808	35.42/0.9698	<u>40.10/0.9906</u>	36.12/0.9797	30.63/0.9368	<u>25.27/0.8609</u>	0.079	19.7	0.79

Table 1. **Quantitative comparison (PSNR/SSIM) of VFI results on the Vimeo90K, UCF101 and SNU-FILM datasets.** For each item, the best result is **boldfaced**, and the second best is underlined. Top and bottom parts are divided by running time.

4. Experiments

In this section, we first introduce implementation details and datasets used in this paper. Then, we quantitatively and qualitatively compare IFRNet with recent state-of-the-arts on various benchmarks. Finally, ablation studies are carried out to analyze the contribution of proposed approaches. Experiments in the main paper follow a common practice of $t = 0.5$, that is synthesizing the single middle frame. IFRNet also supports multi-frame interpolation with temporal encoding T , whose results are presented in supplementary.

4.1. Implementation Details

We implement proposed algorithm in PyTorch, and use Vimeo90K [50] training set to train IFRNet from scratch. Our model is optimized by AdamW [30] algorithm for 300 epochs with total batch size 24 on four NVIDIA Tesla V100 GPUs. The learning rate is initially set to 1×10^{-4} , and gradually decays to 1×10^{-5} following a cosine attenuation schedule. During training, we augment the samples by random flipping, rotating, reversing sequence order and random cropping patches with size 224×224 . For optical flow distillation, we extract pseudo label of bilateral intermediate flow fields with the pre-trained LiteFlowNet [21] in advance, and perform consistent augmentation operations with frame triplets during the whole training process.

4.2. Evaluation Metrics and Datasets

We evaluate our method on various datasets covering diverse motion scenes for comprehensive comparison. Common metrics, such as PSNR and SSIM [47] are adopted for quantitative evaluation. For Middlebury, we use the official IE and NIE indices. Now, we briefly introduce the used test datasets to assess our approaches.

Vimeo90K [50]: It contains frame triplets of 448×256 resolution. There are 3,782 triplets consisted in the test part.

UCF101 [43]: We adopt the test set selected in DVF [29], which includes 379 triplets of 256×256 frame size.

SNU-FILM [10]: SNU-FILM contains 1,240 frame triplets of approximate 1280×720 resolution. According to motion magnitude, it is divided into four different parts, namely, Easy, Medium, Hard, and Extreme for detailed comparison.

Middlebury [2]: The Middlebury benchmark is a widely used dataset to evaluate optical flow and VFI methods. Image resolution in this dataset is around 640×480 . In this paper, we test on the Evaluation set without using Other set.

4.3. Comparison with the State-of-the-Arts

We compare IFRNet with state-of-the-art VFI methods, including kernel-based SepConv [35], AdaCoF [25] and CDFI [12], flow-based ToFlow [50], DAIN [3], SoftSplat [33], BMBC [36], RIFE [20] and ABME [37], and hallucination-based CAIN [10] and FeFlow [16]. For results on SNU-FILM, we execute the released codes of CDFI and RIFE and refer to the other results tested in ABME. For Middlebury, we directly test on the Evaluation part and submit interpolation results to the online benchmark. To measure the inference speed and computation complexity, we run all methods on one Tesla V100 GPU under 1280×720 resolution and average the running time with 100 iterations. For fair comparison, we further build a large and a small version of IFRNet by scaling feature channels with 2.0 and 0.75, respectively, and separate above methods into two classes, *i.e.*, *fast* and *slow*, according to their inference time.

Quantitative Evaluation. Table 1 and Table 2 summarize quantitative results on diverse benchmarks. On Vimeo90K and UCF101 test datasets, IFRNet large achieves the best results on both PSNR and SSIM metrics. A recent method ABME [37] also gets similar accuracy. However, our model runs $11.5 \times$ faster with similar amount of parameters due to the efficiency of single encoder-decoder based architecture. Our large model also obtains leading results on the Easy,




Figure 6. **Qualitative comparison of different VFI methods on SNU-FILM (Hard) dataset.** Proposed IFRNet algorithm can synthesize fast moving objects with sharp boundary while maintaining distinct contextual details. Zoom in for best view.

Method	Average		Mequon		Schefflera		Urban		Teddy		Backyard		Basketball		Dumprtruck		Evergreen	
	IE	NIE	IE	NIE	IE	NIE	IE	NIE	IE	NIE	IE	NIE	IE	NIE	IE	NIE	IE	NIE
SuperSloMo [22]	5.310	0.778	2.51	0.59	3.66	0.72	2.91	0.74	5.05	0.98	9.56	0.94	5.37	0.96	6.69	0.60	6.73	0.69
ToFlow [50]	5.490	0.840	2.54	0.55	3.70	0.72	3.43	0.92	5.05	0.96	9.84	0.97	5.34	0.98	6.88	0.72	7.14	0.90
DAIN [3]	4.856	0.713	2.38	0.58	3.28	0.60	3.32	0.69	4.65	0.86	7.88	0.87	4.73	0.85	6.36	0.59	6.25	0.66
FeFlow [16]	4.820	0.719	2.28	0.51	3.50	0.66	2.82	0.70	4.75	0.87	<u>7.62</u>	0.84	4.74	0.86	6.07	0.64	6.78	0.67
AdaCoF [25]	4.751	0.730	2.41	0.60	3.10	0.59	3.48	0.84	4.84	0.92	8.68	0.90	4.13	0.84	5.77	<u>0.58</u>	5.60	<u>0.57</u>
BMBC [36]	4.479	0.696	2.30	0.57	3.07	0.58	3.17	0.77	4.24	0.84	7.79	<u>0.85</u>	4.08	<u>0.82</u>	5.63	<u>0.58</u>	<u>5.55</u>	0.56
SoftSplat [33]	<u>4.223</u>	<u>0.645</u>	2.06	<u>0.53</u>	<u>2.80</u>	<u>0.52</u>	<u>1.99</u>	<u>0.52</u>	3.84	0.80	8.10	<u>0.85</u>	<u>4.10</u>	0.81	5.49	0.56	5.40	<u>0.57</u>
IFRNet large	4.216	0.644	<u>2.08</u>	<u>0.53</u>	2.78	0.51	1.74	0.43	<u>3.96</u>	<u>0.83</u>	7.55	0.87	4.42	0.84	<u>5.56</u>	0.56	5.64	0.58

Table 2. **Evaluation results on the Middlebury benchmark.** For each item, the best result is **boldfaced**, and the second best is underlined.

Medium and Hard parts of SNU-FILM datasets, while only falls behind ABME on the Extreme part. We attribute the reason to be that the bilateral cost volume constructed by ABME is good at estimating large displacement motion. In Table 2, IFRNet large achieves top-performing VFI accuracy in most of the eight Middlebury test sequences, and outperforms the previous state-of-the-art SoftSplat [33] on both average IE and NIE metrics. Although the improvement is limited, our approach runs $2.5 \times$ faster than SoftSplat which takes cascaded VFI architecture. For FLOPs in convolution layers, IFRNet large also consumes significantly less computation than other VFI architectures.

In regard to real-time and lightweight VFI approaches, IFRNet yields about 0.2 dB better result than RIFE [20] on Vimeo90K, and the margin is more distinct on large motion cases in SNU-FILM dataset. It is worth noting that IFRNet only contains **half** parameters to achieve better results than RIFE thanks to the superiority of joint refinement of intermediate flow and context feature. Compared with CDFI full [12], IFRNet has the same 5M parameters, while achieving **0.63** dB higher PSNR on Vimeo90K with $15.2 \times$ faster inference speed. Moreover, IFRNet small can further

improve speed by **31%** and reduce parameters and computation complexity by **44%** than IFRNet while with only slight frame interpolation accuracy decrease.

Qualitative Evaluation. Figure 6 visually compares well-behaved VFI methods on SNU-FILM (Hard) dataset which contains large and complex motion scenes. It can be seen that kernel-based [12, 25, 35] and hallucination-based [10] methods fail to synthesize sharp motion boundary, containing ghost and blur artifacts. Compared with flow-based algorithms [3, 37], our approach can generate texture details faithfully thanks to the powerfulness of gradually refined intermediate feature. In short, IFRNet can synthesize pleasing target frame with more comfortable visual experience. More qualitative results can be found in our supplementary.

4.4. Ablation Study

To verify the effectiveness of proposed approaches, we carry out ablation study in terms of network architecture and loss function on Vimeo90K and SNU-FILM Hard datasets.

Intermediate Feature. To ablate the effectiveness of intermediate feature $\hat{\phi}_t^k$ in IFRNet, we build a model by removing $\hat{\phi}_t^k$ from the input and output of multiple decoders, while

Architecture		Vimeo90K	Hard
IF	R	PSNR	PSNR
✗	✗	34.83	29.96
✓	✗	35.22	30.22
✗	✓	35.11	30.06
✓	✓	35.51	30.27

Table 3. **Ablation study on different architecture variants.** ‘IF’ means intermediate feature $\hat{\phi}_t^k$ and ‘R’ stands for residual R .




Figure 7. **Visual comparison of intermediate flow and predicted frame of IFRNet w/o and w/ intermediate feature.**

keeping feature channels of middle parts of decoders unchanged. Also, we selectively remove residual R in Eq. 4 to isolate the improvement from intermediate flow and residual. We train them with only the reconstruction loss \mathcal{L}_r under the same learning schedule as before. As listed in Table 3, from the first two rows, we can observe that intermediate feature can provide reference anchor information to promote intermediate flow estimation. Figure 7 also presents some visual examples to confirm the conclusion. Compared with the last and the second rows in Table 3, it demonstrates that gradually refined intermediate feature, containing global context information, can compensate better scene details. Conclusively, residual compensation from the intermediate context feature is necessary for IFRNet to achieve advanced VFI performance, since intermediate flow prediction is substantively unreliable. Overall, the two-fold benefits from intermediate feature greatly improves VFI accuracy of IFRNet with relatively small additional cost.

Task-Oriented Flow Distillation. Table 4 compares VFI accuracy under different combinations of proposed loss functions quantitatively. It can be seen that adding task-oriented flow distillation loss \mathcal{L}_d consistently improves PSNR of 0.2 dB on Vimeo90K. To verify the superiority of its task adaptive ability, we also perform flow distillation with generalized Charbonnier loss under different robustness shown in Figure 5, whose results are summarized in Figure 8. It turns out that robustness parameter $p = 0.3$ achieves best VFI accuracy in the fixed robustness setting. On the other hand, flow distillation can damage frame quality when p approaches to 1.0 due to the harmful knowledge in pseudo label. In a word, proposed task-oriented approach achieves the best accuracy thanks to its spatial adaptive abil-

Loss Function			Vimeo90K	Hard
\mathcal{L}_r	\mathcal{L}_d	\mathcal{L}_g	PSNR	PSNR
✓	✗	✗	35.51	30.27
✓	✓	✗	35.72	30.38
✓	✗	✓	35.61	30.30
✓	✓	✓	35.80	30.41

Table 4. **Ablation study on different loss functions.**




Figure 8. **Ablation study on different flow distillation losses.**




Figure 9. **Visual comparison of mean feature map of intermediate feature $\hat{\phi}_t^1$ w/o and w/ \mathcal{L}_g .** Leftmost is the ground truth.

ity for adjusting robustness loss during flow distillation.

Feature Space Geometry Consistency. As shown in Table 4, adding proposed feature space geometry consistency loss \mathcal{L}_g based on above contributions, we can obtain a further improvement, that confirms the complementary effect of \mathcal{L}_g in regard to \mathcal{L}_d . Figure 9 visually compares mean feature maps of intermediate feature $\hat{\phi}_t^1$ w/o and w/ \mathcal{L}_g . It shows that \mathcal{L}_g can regularize the reconstructed intermediate feature to keep better geometry layout in multi-scale feature space, resulting in better VFI performance.

5. Conclusion

In this paper, we have devised an efficient deep architecture, termed IFRNet, for video frame interpolation, without any cascaded synthesis or refinement module. It gradually refines intermediate flow together with a powerful intermediate feature, that can not only boost intermediate flow estimation to synthesize sharp motion boundary but also provide global context representation to generate vivid motion details. Moreover, we have presented task-oriented flow distillation loss and feature space geometry consistency loss to fully release its potential. Experiments on various benchmarks demonstrate the state-of-the-art performance and fast inference speed of proposed approaches. We expect proposed single encoder-decoder joint refinement based IFRNet to be a useful component for many frame rate up-conversion and intermediate view synthesis systems.

References

- [1] Filippo Aleotti, Matteo Poggi, Fabio Tosi, and Stefano Mattoccia. Learning end-to-end scene flow by distilling single tasks knowledge. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020. 3
- [2] Simon Baker, Daniel Scharstein, J. P. Lewis, Stefan Roth, Michael J. Black, and Richard Szeliski. A database and evaluation methodology for optical flow. *International Journal of Computer Vision*, 2011. 6
- [3] Wenbo Bao, Wei-Sheng Lai, Chao Ma, Xiaoyun Zhang, Zhiyong Gao, and Ming-Hsuan Yang. Depth-aware video frame interpolation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 6, 7
- [4] Wenbo Bao, Wei-Sheng Lai, Xiaoyun Zhang, Zhiyong Gao, and Ming-Hsuan Yang. Memc-net: Motion estimation and motion compensation driven neural network for video interpolation and enhancement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2
- [5] Kelvin C.K. Chan, Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Basicvsr: The search for essential components in video super-resolution and beyond. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [6] P. Charbonnier, L. Blanc-Feraud, G. Aubert, and M. Barlaud. Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proceedings of 1st International Conference on Image Processing*, 1994. 4
- [7] Xianhang Cheng and Zhenzhong Chen. Video frame interpolation via deformable separable convolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020. 2
- [8] Xianhang Cheng and Zhenzhong Chen. Multiple video frame interpolation via enhanced deformable separable convolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2
- [9] Zhixiang Chi, Rasoul Mohammadi Nasiri, Zheng Liu, Juwei Lu, Jin Tang, and Konstantinos N. Plataniotis. All at once: Temporally adaptive multi-frame interpolation with advanced motion modeling. In *Computer Vision – ECCV 2020*, 2020. 1, 2
- [10] Myungsub Choi, Heewon Kim, Bohyung Han, Ning Xu, and Kyoung Mu Lee. Channel attention is all you need for video frame interpolation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020. 2, 6, 7
- [11] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017. 2
- [12] Tianyu Ding, Luming Liang, Zhihui Zhu, and Ilya Zharkov. Cdffi: Compression-driven network design for frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 2, 4, 6, 7
- [13] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Häusser, Caner Hazirbas, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, 2015. 1, 3
- [14] Karl M. Fant. A nonaliasing, real-time spatial transform technique. *IEEE Computer Graphics and Applications*, 1986. 2
- [15] Damien Fourure, Rémi Emonet, Elisa Fromont, Damien Muselet, Alain Trémeau, and Christian Wolf. Residual conv-deconv grid network for semantic segmentation. In *Proceedings of the British Machine Vision Conference, 2017*, 2017. 2
- [16] Shurui Gui, Chaoyue Wang, Qihua Chen, and Dacheng Tao. Featureflow: Robust video interpolation via structure-to-texture generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 6, 7
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *2015 IEEE International Conference on Computer Vision (ICCV)*, 2015. 3, 4
- [18] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015. 3, 4
- [19] Berthold K.P. Horn and Brian G. Schunck. Determining optical flow. *Artificial Intelligence*, 1981. 3
- [20] Zhewei Huang, Tianyuan Zhang, Wen Heng, Boxin Shi, and Shuchang Zhou. Rife: Real-time intermediate flow estimation for video frame interpolation. *CoRR*, 2021. 1, 2, 3, 4, 6, 7
- [21] Tak-Wai Hui, Xiaou Tang, and Chen Change Loy. Liteflownet: A lightweight convolutional neural network for optical flow estimation. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018. 5, 6
- [22] Huaizu Jiang, Deqing Sun, Varan Jampani, Ming-Hsuan Yang, Erik Learned-Miller, and Jan Kautz. Super slo-mo: High quality estimation of multiple intermediate frames for video interpolation. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018. 1, 2, 7
- [23] Soo Kim, Jihyong Oh, and Munchurl Kim. Fisr: Deep joint frame interpolation and super-resolution with a multi-scale temporal loss. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020. 2
- [24] Lingtong Kong, Chunhua Shen, and Jie Yang. FastflowNet: A lightweight network for fast optical flow estimation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021. 1, 3, 4
- [25] Hyeongmin Lee, Taeoh Kim, Tae-young Chung, Daehyun Pak, Yuseok Ban, and Sangyoun Lee. Adacof: Adaptive collaboration of flows for video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 6, 7
- [26] Pengpeng Liu, Irwin King, Michael R. Lyu, and Jia Xu. Ddflow: Learning optical flow with unlabeled data distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019. 3
- [27] Pengpeng Liu, Michael R. Lyu, Irwin King, and Jia Xu. Selfflow: Self-supervised learning of optical flow. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. 3

- [28] Yu-Lun Liu, Yi-Tung Liao, Yen-Yu Lin, and Yung-Yu Chuang. Deep video frame interpolation using cyclic frame generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019. 6
- [29] Ziwei Liu, Raymond A. Yeh, Xiaoou Tang, Yiming Liu, and Aseem Agarwala. Video frame synthesis using deep voxel flow. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017. 2, 6
- [30] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *7th International Conference on Learning Representations*, 2019. 6
- [31] Simon Meister, Junhwa Hur, and Stefan Roth. UnFlow: Un-supervised learning of optical flow with a bidirectional census loss. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018. 4, 5
- [32] Simon Niklaus and Feng Liu. Context-aware synthesis for video frame interpolation. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018. 1, 2
- [33] Simon Niklaus and Feng Liu. Softmax splatting for video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2, 4, 6, 7
- [34] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive convolution. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [35] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive separable convolution. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017. 2, 6, 7
- [36] Junheum Park, Keunsoo Ko, Chul Lee, and Chang-Su Kim. Bmbc: Bilateral motion estimation with bilateral cost volume for video interpolation. In *European Conference on Computer Vision*, 2020. 1, 2, 6, 7
- [37] Junheum Park, Chul Lee, and Chang-Su Kim. Asymmetric bilateral motion estimation for video frame interpolation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 1, 2, 4, 6, 7
- [38] Tomer Peleg, Pablo Szekely, Doron Sabo, and Omry Sendik. Im-net for high resolution video frame interpolation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [39] Anurag Ranjan and Michael J. Black. Optical flow estimation using a spatial pyramid network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [40] Hyeonjun Sim, Jihyong Oh, and Munchurl Kim. Xvfi: extreme video frame interpolation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 1, 2
- [41] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems*, 2014. 3
- [42] Li Siyao, Shiyu Zhao, Weijiang Yu, Wenxiu Sun, Dimitris Metaxas, Chen Change Loy, and Ziwei Liu. Deep animation video interpolation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1
- [43] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *CoRR*, 2012. 6
- [44] Deqing Sun, Stefan Roth, and Michael J. Black. A quantitative analysis of current practices in optical flow estimation and the principles behind them. *International Journal of Computer Vision*, 2014. 5
- [45] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018. 1, 3, 4
- [46] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision – ECCV 2020*, 2020. 1
- [47] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 2004. 6
- [48] George Wolberg, H. M. Sueyllam, M. A. Ismail, and K. M. Ahmed. One-dimensional resampling with inverse and forward mapping functions. *Journal of Graphics Tools*, 2000. 2
- [49] Xiangyu Xu, Li Siyao, Wenxiu Sun, Qian Yin, and Ming-Hsuan Yang. Quadratic video interpolation. In *Advances in Neural Information Processing Systems*, 2019. 1, 2
- [50] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision (IJCV)*, 2019. 1, 2, 5, 6, 7
- [51] Liangzhe Yuan, Yibo Chen, Hantian Liu, Tao Kong, and Jianbo Shi. Zoom-in-to-check: Boosting video interpolation via instance-level discrimination. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [52] Yuan Yuan, Wei Su, and Dandan Ma. Efficient dynamic scene deblurring using spatially variant deconvolution network with optical flow guided training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [53] Ramin Zabih and John Woodfill. Non-parametric local transforms for computing visual correspondence. In *Computer Vision – ECCV '94*, 1994. 5
- [54] Haoxian Zhang, Yang Zhao, and Ronggang Wang. A flexible recurrent residual pyramid network for video frame interpolation. In *Computer Vision – ECCV 2020*, 2020. 1, 2
- [55] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A. Efros. View synthesis by appearance flow. In *Computer Vision – ECCV 2016*, 2016. 1