

# SportsSloMo: A New Benchmark and Baselines for Human-centric Video Frame Interpolation

Jiabben Chen\*  
UC San Diego  
jic088@ucsd.edu

Huaizu Jiang  
Northeastern University  
h.jiang@northeastern.edu

## Abstract

Human-centric video frame interpolation has great potential for enhancing entertainment experiences and finding commercial applications in sports analysis industry, e.g., synthesizing slow-motion videos. Although there are multiple benchmark datasets available for video frame interpolation in the community, none of them is dedicated to human-centric scenarios. To bridge this gap, we introduce **SportsSloMo**, a benchmark featuring over 130K high-resolution ( $\geq 720p$ ) slow-motion sports video clips, totaling over 1M video frames, sourced from YouTube. We re-train several state-of-the-art methods on our benchmark, and we observed a noticeable decrease in their accuracy compared to other datasets. This highlights the difficulty of our benchmark and suggests that it poses significant challenges even for the best-performing methods, as human bodies are highly deformable and occlusions are frequent in sports videos. To tackle these challenges, we propose human-aware loss terms, where we add auxiliary supervision for human segmentation in panoptic settings and keypoints detection. These loss terms are model-agnostic and can be easily plugged into any video frame interpolation approach. Experimental results validate the effectiveness of our proposed human-aware loss terms, leading to consistent performance improvement over existing models. The dataset and code can be found at: <https://neu-vi.github.io/SportsSloMo/>.

## 1. Introduction

Video frame interpolation (VFI) is a technique that synthesizes intermediate frames from input images, enhancing the clarity of content that may be difficult to see otherwise. This technique finds wide-ranging applications, including slow-motion video generation [27], novel view synthesis [39], video compression [74], cartoon and rendered content gen-

\*Work mainly done when Jiabben Chen was an intern at Northeastern University.




Figure 1. **Human-centric video frame interpolation results.** We propose human-aware auxiliary losses to improve interpolation accuracy at motion boundaries. From left to right: (a): Overlaid inputs, (b): Ground truth, (c): Interpolation results, (d): Interpolation results with our proposed human-aware loss terms.

eration [4, 37], etc. In recent years, we have witnessed significant advances in this field, due in large part to the development of various benchmarks [1, 13, 22, 57, 64–66, 80].

Humans feature prominently in most contemporary videos. With the widespread use of mobile devices, people can easily record and share their daily experiences with family, friends, and colleagues. Meanwhile, live broadcasts of sporting events attract a large audience. Automatically generated slow-motion videos can create a more immersive and engaging experience for users by highlighting details and valuable moments of their lives that may be missed in real-time. Therefore, improving video frame interpolation results for human-centric videos has great potential for enhancing user experiences in entertainment. Human-centric video frame interpolation approaches can also be beneficial in various industries. Athletes and coaches, for example, can use slow-motion synthesis to identify flaws in techniques, highlight areas for improvement, and gain a more detailed understanding of how different factors may contribute to the success.

**A new benchmark.** Despite the availability of various benchmarks for video frame interpolation, a notable gap exists in datasets specifically tailored for human-centric sce-

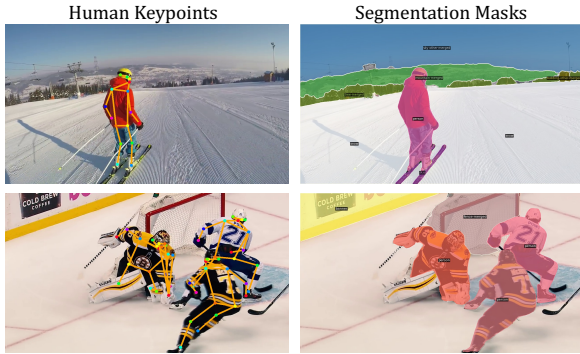


Figure 2. Visualization of human keypoints [78] and panoptic segmentation masks [9].

narios. To bridge this gap and to foster the research in this important direction, we introduce **SportsSloMo**, a new dataset comprising high-resolution ( $\geq 720p$ ) slow-motion sports videos crawled from YouTube under the Common Creative Licence. This dataset encompasses a diverse range of sports, such as football, basketball, baseball, hockey, etc. Since a video may contain advertisement, transition frames, changes of shot, and non-slow-motion content, we carefully curate the data to remove such unwanted content and finally split each long video into a set of short slow-motion clips of 9 frames. The first and last frames are used as input and the rest 7 intermediate frames are reserved as ground truths for training and evaluating VFI models. In total, our benchmark has 130K video clips and more than 1M video frames. Compared with existing datasets, as shown in Table 1, our proposed SportsSloMo benchmark is the largest one so far, with high resolution and focus on human-centric scenarios.

While primarily designed for human-centric VFI in this paper, we believe SportsSloMo dataset may also has the potential to aid research in other tasks such as video super-resolution [6, 38], group activity recognition [67, 82, 86], and dynamic view synthesis [16, 39, 72, 84]. By releasing this dataset to the entire community, we hope to encourage technical advancement in human-centric video frame interpolation and empower researchers to explore innovative applications in other adjacent fields.

**Benchmarking existing approaches.** To facilitate development and evaluation of human-centric video frame interpolation methods, we re-train several state-of-the-art approaches [23, 25, 27–29, 36, 83] using their publicly released code on our SportsSloMo dataset. As the human bodies are highly deformable and occlusions are frequent in sports videos, the accuracy of all the methods decrease compared to their performance on other datasets. For instance, EBME [29] and EMA-VFI [83], two of the top-performing approaches on SportsSloMo, produce PSNR scores of 30.15 and 30.70, respectively - markedly lower than their scores of 36.19 and 36.64 on Vimeo90K [80], as well as 30.64 and 30.94 on the hard split of the SNU-FILM [13] benchmarks.

Table 1. Comparisons of different benchmark datasets for video frame interpolation. (#inter. frames indicates the number of intermediate frames to synthesize.)

Dataset	#clips	#images	#inter. frames	resolution	human-centric
UCF101 [65]	0.4K	1K	1	256×256	✓
Adobe240fps [66]	0.1K	80K	7	1280×720	✗
Vimeo90K [80]	70K	220K	1	448×256	✗
SNU-FILM [13]	1.2K	3.6K	1	1280×720	✗
X4K1000FPS [64]	4.4K	286K	7	4096×2160	✗
SportsSloMo	130K	1183K	7	1280×720	✓

It highlights the difficulty of our benchmark and suggests significant challenges need to be addressed.

**Enhancing models for human-centric VFI.** To improve the existing VFI models on our benchmark, we introduce human-aware priors to enhance the model training. Specifically, we propose loss terms based on human segmentation in the panoptic setting [9] and human keypoints estimation [78] as extra supervision for intermediate frame synthesis. Fig. 2 shows a visualization of detected human keypoints and segmentation masks in our dataset. Human segmentation masks help delineate human body boundaries, which are helpful for reducing ghost effects around the motion boundaries. At the same time, human keypoints estimation can also indicate where each body part is, enforcing coherent motion trajectories in the synthesized video frames. Specifically, we compare the output from pre-trained panoptic segmentation and keypoints detection models taking input as a synthesized and the ground-truth intermediate frame, respectively, and use the consistency as supervision. As shown in Fig. 1, by supervising with our proposed human-aware loss terms, we improve the interpolation quality at motion boundaries with less blurry results in scenarios with large motion and occlusion. Both of these human-aware loss terms are model agnostic and can be easily integrated into any video frame interpolation approach. Experimental results show that they can consistently improve the accuracy of seven existing approaches, leading to strong baselines on our benchmark.

To sum up, this paper makes the following contributions:

- We introduce SportsSloMo, a new benchmark dataset consisting of a large amount of slow-motion sports videos. To the best of our knowledge, this is the first high-resolution dataset tailored for human-centric video frame interpolation, supporting synthesis of multiple intermediate frames.
- We benchmark state-of-the-art approaches on the new benchmark, highlighting the challenges of the human-centric video frame interpolation task.
- We propose two human-aware loss terms, which take the

human priors into account for video frame interpolation and can easily be plugged into existing video frame interpolation approaches. Experimental results validate that they can consistently improve existing models, yielding strong baseline models on our new benchmark.

## 2. Related Work

### 2.1. Benchmark datasets

Existing publicly available datasets already provide a valuable resource for developing and evaluating video frame interpolation methods. We would briefly introduce these datasets and reveal some limitations in this section. Table. 1 shows a comparison between existing datasets and our proposed SportsSloMo dataset.

The SNU-FILM [13] dataset is a widely-used benchmark for VFI evaluation, containing 1240 frame triplets of  $1280 \times 720$  resolution. And it is divided into four different parts, namely, Easy, Medium, Hard, and Extreme according to motion magnitude. The Middlebury benchmark [1] is another widely used dataset, image resolution in this dataset is around  $640 \times 480$ . However, it is commonly only used to evaluate VFI methods for 8 sequences. UCF101 [65] is originally a dataset for human action recognition, containing a variety of human actions. With the test set constructed by [42], it is also used to evaluate VFI methods, containing 379 triplets of  $256 \times 256$  frame size. Nevertheless, its scale is rather small and its resolution is also low. The Adobe240fps dataset [66], originally for video deblurring, is another widely used dataset for VFI. It is consisted of high frame-rate videos (240 fps) with a resolution of  $1280 \times 720$ , yet the videos are from only 118 clips.

The mostly used dataset for training and evaluating VFI methods is the Vimeo90K dataset [80]. It contains 73171 frame triplets from 14777 video clips extracted from real-life video clips with a  $\text{fps} \leq 30$ . Nevertheless, one of its main drawback is the low resolution of  $448 \times 256$  obtained by downscaling the original high resolution frames. Moreover, as VFI methods have been rapidly improving in recent years, their performance on the widely-used Vimeo90K dataset has approached saturation. To further advance the state-of-the-art in video frame interpolation, a new dataset with bigger scale, higher resolution and more challenging scenarios is necessary. X4K1000FPS is a recently released high frame rate (1000 fps) with 4K spatial resolution to promote the study of VFI for very high resolution videos.

As a result, none of existing datasets contains rich human-centric data with high resolution and big scale. To bridge this gap, we introduce the SportsSloMo dataset, a human-centric VFI dataset with 130K video clips and 1M video frames at 240 fps with a resolution of  $1280 \times 720$ , aiming to foster research in human-centric VFI.

### 2.2. Video frame interpolation methods

Existing VFI methods could be generally classified into flow-agnostic and flow-based methods.

**Flow-agnostic approaches** model VFI without explicit intermediate motion representation. *Phase-based* methods [47, 48] directly predict the phase decomposition of the intermediate frame, but can only handle motion within a limited range. *Kernel-based* methods are the mainstream approach in this category, which typically aims to estimate intermediate frames by learning adaptive kernels to convolve input frames [53, 54]. Over the years, numerous improvements are proposed in this field, including using deformable convolution [10, 11], formulating interpolated motion estimation as classification [56], blending deep features [18], introducing dual-frame adversarial loss [36], performing channel attention [13] and utilizing 3D space-time convolutions [32]. Recently, Shi *et al.* [63] introduced a Transformer-based framework to model long-range dependencies with the aid of attention mechanisms. By directly hallucinating pixel values, these methods tend to generate blurry results and artifacts, especially in fast-moving scenes [44].

**Flow-based approaches** currently serve as a promising direction in VFI. Generally speaking, flow-based method takes a paradigm of a two-stages pipeline: (1) flow estimation, and (2) frame synthesis. They first estimate optical flow between input frames, and then synthesize intermediate frames using image warping [26]. As a representative work, SuperSloMo [27] by Jiang *et al.* adopted a skip-connected U-Net to estimate bi-directional optical flows under the assumption of linear motion. Quadratic [77] and cubic [12, 70] trajectory assumptions have also been made to approximate intermediate motion. Recent work has explored various techniques to improve intermediate flow estimation and interpolation accuracy, including forward-warping via softmax splatting [23, 52], voxel flow [42], cycle consistency loss [41, 61], task-oriented flow distillation loss [35], Gram matrix loss [60], implicit neural function [8], occlusion mask [3], anchor points alignment [64], privileged distillation [25], and pyramid recurrent flow estimation [28]. Considering additional information like contextual maps [51], depth maps [2] and auxiliary visual information from event cameras [20, 69] can also further improve interpolation accuracy. Park *et al.* employed symmetric bilateral motion field estimation, and further improved intermediate motion estimation accuracy through asymmetric bilateral motion field [55]. Lu *et al.* [44] leveraged Transformer architecture [71] to model long-term dependency. Jin *et al.* [29] proposed a novel bi-directional motion estimator in a pyramid structure. Zhang *et al.* [83] proposed a novel feature extraction strategy to combine motion and appearance information via a hybrid CNN and Transformer architecture.





Figure 3. **The SportsSloMo dataset.** (a) Sampled frames, covering various sports categories and challenging human-centric content for VFI; (b) Histogram of flow magnitude of all pixels in the dataset; (c) Histogram of mean flow magnitude of all images in the dataset.




Figure 4. **Distribution of different sports categories** in our SportsSloMo benchmark.

In a nutshell, although previous methods have proposed successful designs to handle complex motion and occlusion, none of them is carefully designed for human-centric scenes. As discussed in Sec. 1, human-centric VFI is confronted with various challenges including dynamic pose variation, complex human motion patterns and occlusion in crowded scenes. Additionally, accurately synthesizing fine details such as facial expressions and hand gestures can be challenging. To this end, we propose to consider human-aware loss terms through incorporating extra supervision to both human keypoints detection and segmentation in the panoptic setting.

### 3. SportsSloMo Benchmark

A myriad of benchmark datasets for video frame interpolation are available, including Middlebury [1], GoPro [50], UCF101 [65], DAVIS [57], Adobe240fps [66], Vimeo90K [80], SNU-FILM [13] and X4K1000FPS [64]. But none of these them focuses on human-centric VFI, *e.g.*, in sports scenes. This limits the study of VFI methods targeted for human-centric applications such as enhanced entertainment experiences, commercial deployment, etc.

To bridge this gap and to foster future research, we propose **SportsSloMo**, a challenging dataset consisting of high-resolution ( $\geq 720p$ ) sports videos crawled from YouTube under the Common Creative License. Careful curation has been taken to remove the unwanted content in the videos, including advertisement, transition frames, changes of shot, flashing lights, and non-slow-motion content. Specifically, we first conduct human detection on all the videos by utilizing Yolov3 [62] and remove video frames without detected humans. Second, we remove frames with flashing lights by setting a threshold about the brightness change between consecutive frames (large brightness change indicates the existence of a flashing light). Third, we use RAFT [68] to measure motion magnitude and set a threshold to discard non-slow-motion video segments. Finally, we carefully curate the the extracted clips manually to keep only high-quality self-consistent videos. With such semi-automatic curation, we end up having a set of short slow-motion clips of 9 frames, where each frame has a spatial resolution of at least  $1280 \times 720$ . The first and last frames are used as input for a VFI method and the 7 intermediate frames as ground truths for model training and evaluation, approximately corresponding to converting a video of 30-fps (frames per second) to 240-fps.

In total, we collect 131,464 video clips and 1.2M individual video frames from 259 raw YouTube videos, which covers 22 various sports categories in our dataset with different content and motion patterns, including hockey, baseball, skating, basketball, volleyball, etc. Fig. 4 shows the proportion of different sports categories. Visualizations of randomly sampled video frames are available in Fig. 3. We refer the readers to the supplemental video for more results.

As shown in Table 1, compared with existing datasets,

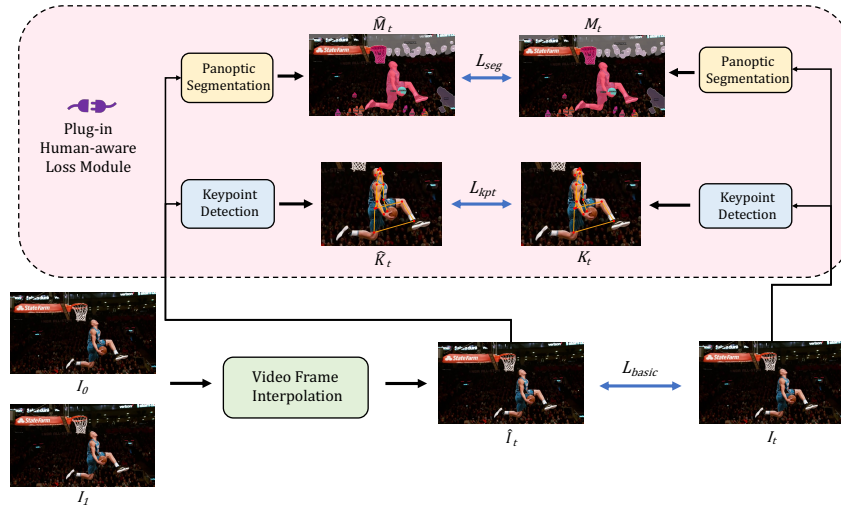


Figure 5. **Illustration of our proposed human-aware loss terms**, consisting of losses about human segmentation in the panoptic setting and keypoint detection.

not only our dataset is tailored for human-centric scenarios, but also surpasses others in terms of scale, resolution, and frame rate. Fig. 3 demonstrates the histogram of flow magnitude of all pixels and mean flow magnitude of all images in our dataset (calculated using GMFlow [76]), with a comparison with the widely-used SNU-FILM [13] dataset (which has the same spatial resolution as ours) and the recently introduced X4k1000FPS [64] dataset. As can be seen, our SportsSloMo benchmark contains more large-displacement motion compared to the SNU-FILM dataset. For instance, for the average flow magnitude on the image-level (Fig. 3(c)), our dataset has way more images, whose mean flow magnitude is greater than 20 pixels. And we can also observe that both our SportsSloMo and X4K1000FPS datasets contain large motions.

We split our proposed SportsSloMo dataset into *train* and *test*, containing 115,421 and 16,043 video clips, respectively. For each sports category, videos are split into *train* and *test without intersection*, so that the test videos are completely unseen during training.

## 4. Human-aware Video Frame Interpolation

### 4.1. Overview

Given two input frames  $I_0$  and  $I_1$ , the goal of VFI is to synthesize the intermediate frame  $I_t$  at an intermediate time step  $t \in (0, 1)$ . A VFI method needs to find the pixel-wise correspondences between  $I_0$  and  $I_1$  and adaptively fuse corresponding pixels to synthesize each pixel in  $I_t$ . Flow-based approaches, such as SuperSloMo [27] and EBME [29], usually explicitly model the correspondences as optical flow, whereas flow-agnostic methods, such as AdaCof [36], use learned kernels to process the visual correspondence implicitly.

To supervise the network training, the reconstruction error of the intermediate frame is used as the loss. For instance, in [29], the weighted sum of the Charbonnier loss [7]  $L_{char}$  and census loss [46]  $L_{cen}$  between the ground truth intermediate frame  $I_t$  and synthesized frame  $\hat{I}_t$  are adopted

$$L_{basic} = L_{char}(I_t - \hat{I}_t) + \lambda_{cen} \cdot L_{cen}(I_t, \hat{I}_t), \quad (1)$$

where  $L_{char}(x) = (x^2 + \epsilon^2)^\alpha$ ,  $\alpha = 0.5$ ,  $\epsilon = 10^{-6}$ .

Such reconstruction loss, however, may not provide enough supervision for human-centric VFI due to the challenges of highly deformable human motion and frequent occlusion. To address these challenges, we propose to incorporate  $L_{seg}$  and  $L_{kpt}$  into the network supervision based on human segmentation in the panoptic setting and keypoint detection, respectively, enforcing the model to produce high-quality synthesis results over human boundaries and along the keypoint trajectories. The final loss is

$$L = L_{basic} + \lambda_{seg} L_{seg} + \lambda_{kpt} L_{kpt}. \quad (2)$$

It is worth noting that our proposed human-aware loss terms are flexible, which can be plugged into other flow-based (e.g., SuperSloMo [27] and EBME [29]) or flow-agnostic VFI models (e.g., AdaCof [36]), as shown in Section 5.

### 4.2. Human Segmentation Loss

In human-centric scenarios, accurate estimation of human body boundaries is crucial. In regions where the body movement is complex or there are heavy occlusions, inaccuracies in the estimated body boundaries may lead to visible artifacts in the synthesized video frame. To this end, we propose to incorporate human segmentation as extra supervision to improve the synthesis of intermediate video frames. Human segmentation masks directly tell where human body




Figure 6. **VFI results with human-aware loss (HL) terms.** We can see better human segmentation, keypoint detection, and interpolation result can be obtained enhanced by the human-aware loss.

boundaries are, as shown in Fig. 2, helping to reduce ghost effects. Though instance and panoptic segmentation could both serve for this purpose, we empirically find that the panoptic segmentation models tend to yield better interpolation results than instance segmentation counterparts.

Specifically, we adopt state-of-the-art Mask2Former [9] trained on COCO panoptic dataset [33] with the Swin-L backbone [43] from Detectron2 [75] to generate panoptic segmentation masks  $M_t$  from the ground-truth intermediate frame  $I_t$ . During training, the panoptic segmentation results  $\hat{M}_t$  of the synthesized intermediate frame  $\hat{I}_t$  obtained using the same Mask2Former model is compared with  $M_t$ . Ideally, if the synthesized intermediate frame  $\hat{I}_t$  is identical to the ground truth  $I_t$ ,  $\hat{M}_t$  should be the same as  $M_t$ . By comparing their difference, the human segmentation loss enhances the consistency of the human body boundaries between  $I_t$  and  $\hat{I}_t$ , which helps reduce the ghost effects. We follow the loss function of Mask2Former and use the binary cross-entropy loss  $L_{ce}$  and the dice loss  $L_{dice}$  [49] for the human segmentation loss

$$L_{seg} = \lambda_{ce} L_{ce}(\hat{M}_t, M_t) + \lambda_{dice} L_{dice}(\hat{M}_t, M_t), \quad (3)$$

where we set  $\lambda_{ce} = 5.0$  and  $\lambda_{dice} = 5.0$ . The panoptic segmentation model is frozen during training and the gradient of this loss will be backpropagated to a VFI model only in a way similar to the perceptual loss [31], enforcing the model to produce high-quality synthesis results.

### 4.3. Human Keypoint Detection Loss

In addition to segmentation, we also consider the human keypoints estimation as an extra supervision. Accurate estimation of human keypoints is also crucial for human-centric VFI, which are indicative about the position of each body part, providing additional cues for motion estimation. Hence, by incorporating the supervision of human keypoints, we can better preserve human motions by enforcing coherent motion trajectories and guide the interpolation process to generate more plausible intermediate frames.

Specifically, we adopt the state-of-the-art human pose estimation model ViTPose [78] trained on COCO [40], with the ViT-L backbone [15] initialized with MAE [19] pre-trained weights. We first use an off-the-shelf human detector YOLOv8 [30] to detect person instances in the ground-truth intermediate frame  $I_t$ . We then employ ViTPose to estimate the keypoints heatmap  $K_t$  containing the locations of each joint. Similar to the human segmentation loss, we estimate the keypoints heatmap  $\hat{K}_t$  using the same ViTPose model on the synthesized intermediate frame  $\hat{I}_t$ . The differences between  $K_t$  and  $\hat{K}_t$  are then used as the supervision to improve VFI results. We follow the loss function of ViTPose and use the MSE (mean square error) loss over heatmaps as the human keypoint loss  $L_{kpt}$ .

$$L_{kpt} = MSE(K_t, \hat{K}_t). \quad (4)$$

The ViTPose model is frozen during training and the gradient of this loss will be backpropagated to a VFI model to encourage high-quality intermediate frame synthesis so that the estimated human keypoints are close to the ones obtained from the true intermediate frame.

## 5. Experiments

### 5.1. Setup

**VFI models.** We benchmark existing VFI approaches by re-training several state-of-the-art methods on our SportsSloMo dataset using their publicly available code, including SuperSloMo [27], AdaCoF [36], M2MVFI [23], RIFE [25], EBME [29], UPR-Net [28], and EMA-VFI [83]. AdaCoF is a flow-agnostic VFI model and the rest are flow-based. All these models are trained from scratch on our SportsSloMo dataset.

We also incorporate our proposed human-aware loss terms into these models to validate the effectiveness of explicitly taking human priors into account to improve human-centric VFI.

**Evaluation metrics.** Following previous VFI methods [24, 27], we adopt the widely-used signal-to-noise ratio (PSNR), structural similarity (SSIM) [73], and interpolation error (IE) [1] to evaluate the interpolation results. Additionally, we employ the keypoint detector [78, 79] to measure






Figure 7. **Qualitative comparisons on SportsSloMo dataset.** With the enhancement of the human-aware (HL) losses, EBME achieves better interpolation results.

Table 2. **Quantitative results of different approaches on the SportsSloMo dataset.** As can be seen, our proposed human-aware loss (HL) terms consistently improve all VFI models’ performance.

Method	PSNR $\uparrow$	SSIM $\uparrow$	IE $\downarrow$	LPIPS $\downarrow$	AP $\uparrow$
SuperSlomo [27]	29.77	0.910	9.14	0.082	87.2
SuperSloMo + HL	<b>30.24</b>	<b>0.917</b>	<b>8.46</b>	<b>0.079</b>	<b>89.8</b>
AdaCoF [36]	28.79	0.926	5.84	0.099	87.0
AdaCoF + HL	<b>28.94</b>	0.926	<b>5.72</b>	<b>0.096</b>	<b>88.6</b>
M2MVFI [23]	29.03	0.935	5.16	0.091	87.6
M2MVFI + HL	<b>29.29</b>	<b>0.936</b>	<b>5.07</b>	<b>0.088</b>	<b>89.3</b>
RIFE [25]	29.69	0.931	5.25	0.078	87.4
RIFE + HL	<b>29.87</b>	<b>0.933</b>	<b>5.16</b>	<b>0.076</b>	<b>89.2</b>
EBME [29]	30.15	0.941	4.66	0.065	90.2
EBME + HL	<b>30.48</b>	<b>0.944</b>	<b>4.40</b>	<b>0.059</b>	<b>90.7</b>
UPR-Net [28]	30.25	0.945	4.56	0.065	90.2
UPR-Net + HL	<b>30.50</b>	0.945	<b>4.38</b>	<b>0.064</b>	<b>90.6</b>
EMA-VFI [83]	30.70	0.949	4.28	0.052	90.4
EMA-VFI + HL	<b>30.75</b>	<b>0.952</b>	<b>4.19</b>	<b>0.051</b>	<b>90.6</b>

estimated keypoint errors and report standard Average Precision (AP). Moreover, we also adopt a perceptual image quality metric LPIPS [85], which is widely used for image generation and correlates better with human perception. For PSNR, SSIM and AP, higher indicates better performance. For IE and LPIPS, the lower the better.

**Implementation details.** For the human segmentation and keypoint loss terms, to avoid the influence of Batch Normalization (BN) layers on the training phase, we replaced the BN layers with frozen Batch Normalization (frozen BN) layers. We use the default hyperparameters provided by

each VFI model. Training and evaluation are both conducted on 8 NVIDIA RTX A6000 GPUs.

During training, we randomly sample a frame between the first and the ninth (last) frame for VFI methods supporting arbitrary time frame interpolation like SuperSlomo [27], M2MVFI [23], RIFE [25], EBME [29], UPR-Net [28], and EMA-VFI [83]. For VFI methods that can only synthesize the middle frame like AdaCoF [36], we randomly generate triplets within the nine frames such that the target frame is always in the middle of the input two frames. During evaluation, arbitrary-time-supported VFI methods take the time step as input and interpolate every intermediate frame, while single-frame methods recursively generate each intermediate frame.

## 5.2. Benchmarking Existing Methods

Table 2 shows quantitative comparisons between existing VFI methods on our proposed SportsSloMo dataset. We can see that EMA-VFI [83] performs the best in terms of all five evaluation metrics. It is worth noting that the performance for each model decreases compared with results on popular datasets. For instance, two top-performing approaches on SportsSloMo, EBME [29] and EMA-VFI [83] produce PSNR scores of 30.15 and 30.70, respectively, compared with 36.19 and 36.64 on the Vimeo90K [80] as well as 30.64 and 30.94 on the hard split of the SNU-FILM [13] benchmarks. It highlights the difficulty of our benchmark and suggests significant challenges need to be addressed. We provide more details in the supplementary material.

## 5.3. Enhanced Human-aware VFI

Table 2 also demonstrates the effectiveness of our proposed human-aware losses. Specifically, we incorporate the proposed human segmentation and keypoint losses into each of the VFI models. As can be seen, they consistently im-

Table 3. Effectiveness of different auxiliary loss terms.

Model	segmentation loss	keypoint loss	PSNR $\uparrow$	SSIM $\uparrow$	IE $\downarrow$
SuperSloMo	$\times$	$\times$	29.77	0.910	9.14
	$\checkmark$	$\times$	30.19	0.916	8.56
	$\times$	$\checkmark$	30.22	0.916	8.49
	$\checkmark$	$\checkmark$	<b>30.24</b>	<b>0.917</b>	<b>8.46</b>
EBME	$\times$	$\times$	30.15	0.941	4.66
	$\checkmark$	$\times$	30.26	0.942	4.61
	$\times$	$\checkmark$	30.34	0.943	4.46
	$\checkmark$	$\checkmark$	<b>30.48</b>	<b>0.944</b>	<b>4.40</b>

prove the performance of every single VFI method for all three evaluation metrics. Specifically, in terms of PSNR, our proposed human-aware losses lead to improvement of 1.6% and 1.1% for SuperSloMo [27] and EBME [29], respectively. In terms of IE, the improvement for these two methods are 7.4% and 5.6%, respectively. Qualitatively, we can see from Fig. 6 and Fig. 7, our proposed human-aware loss terms can improve interpolation results for the highly deformable arms of the athletes and the background under occlusions.

#### 5.4. Ablation Study

In this section, we present ablation study on SuperSloMo [27] and EBME [29] to analyze the design choices of our proposed human-aware losses in Table 3. As can be seen, both human segmentation and keypoint detection losses can successfully improve the performance of SuperSloMo and EBME. This justifies our motivation to leverage human-aware priors to improve the human-centric VFI. By combining these two losses, the largest performance gain can be obtained for both methods. Furthermore, from Fig. 6, we can clearly see that better interpolation results can be obtained around the elbow and hand.

#### 5.5. Limitations and Discussions

While we have introduced human-aware loss terms to handle challenging human-centric scenarios in our proposed dataset, there still remains noteworthy limitations for future exploration.

First of all, for human-centric scenes, the presence of large, complex, and highly deformable human motions and occlusion by crowds make finding correspondences of humans (*i.e.*, optical flow) between input frames a challenging task. Recent studies [25, 35] have demonstrated the effectiveness of knowledge distillation [21] to improve the intermediate optical flow estimation. We also explored this loss to incorporate extra supervision for flow-based VFI methods (*e.g.*, SuperSloMo [27], EBME [29]). Specifically, we use the optical flow results obtained from state-of-the-art

model GMFlow [76], which is pre-trained on the optical flow benchmarks [5, 14, 45], as the teacher network to perform flow distillation. We found, however, mixed results on our SportsSloMo benchmark. While it indeed improves the performance for SuperSloMo [27] (PSNR: 29.77 vs. 30.18), it hurts the performance for EBME [29] (PSNR: 30.15 vs. 29.85). We also fine-tune the GMFlow model on the human flow dataset [59], consisting of synthetic human figures overlaid on top of real world images as background. However, we observed significant domain gap between the synthetic training images and real-world test images, which lead to poor optical flow estimation results. How to improve human-centric optical flow to improve VFI is still a challenging problem.

Second, we only utilize 2D cues of human bodies in our proposed loss terms to enhance VFI. Recently, there has been significant advance made in 3D human body reconstruction from videos, for instance, [17, 34, 58, 81], considering both the spatial and temporal information. By lifting 2D videos into the 3D space, occlusions may be better solved. More advanced loss terms involving 3D human reconstruction is worth exploring. We leave this as future work to explore.

Finally, in the SportsSloMo dataset, various sports involve fast-moving balls and sports equipment. In this paper, we only consider human boundary and body parts for motion estimation. Fast-moving objects, however, in these scenarios are also challenging for VFI. More effort is needed for such a problem.

## 6. Conclusion

In this paper, we introduced a new benchmark, **SportsSloMo**, focusing on human-centric video frame interpolation. Our benchmark contains 130K video clips and more than 1M video frames obtained from high-resolution ( $\geq 720p$ ) slow-motion sports videos crawled from YouTube with careful curation. Due to the complex, highly deformation human motion and frequent occlusion, this benchmark imposes significant challenges to existing VFI models. To enhance their accuracy, we introduce human-aware loss terms to improve existing methods, where the supervision of human segmentation and keypoints detection are incorporated. Our loss terms are model agnostic and have been successfully applied to seven existing VFI methods, leading to better accuracy consistently. Our benchmark dataset and code will be publicly released to foster future research in the new exciting direction of human-centric VFI.

## Acknowledgement

We thank Thuy-Tien Bui and Devroop Kar for their help on data collection. This work was partially supported by the National Science Foundation under Award IIS-2310254.



## References

- [1] Simon Baker, Daniel Scharstein, JP Lewis, Stefan Roth, Michael J Black, and Richard Szeliski. A database and evaluation methodology for optical flow. *International journal of computer vision*, 92(1):1–31, 2011. 1, 3, 4, 6
- [2] Wenbo Bao, Wei-Sheng Lai, Chao Ma, Xiaoyun Zhang, Zhiyong Gao, and Ming-Hsuan Yang. Depth-aware video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3703–3712, 2019. 3
- [3] 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*, 43(3):933–948, 2019. 3
- [4] Karlis Martins Briedis, Abdelaziz Djelouah, Mark Meyer, Ian McGonigal, Markus H. Gross, and Christopher Schroers. Neural frame interpolation for rendered content. *ACM Trans. Graph.*, 40(6):239:1–239:13, 2021. 1
- [5] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, 2012. 8
- [6] Kelvin C.K. Chan, Shangchen Zhou, Xiangyu Xu, and Chen Change Loy. Investigating tradeoffs in real-world video super-resolution. In *CVPR*, 2022. 2
- [7] Pierre Charbonnier, Laure Blanc-Feraud, Gilles Aubert, and Michel Barlaud. Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proceedings of 1st international conference on image processing*, pages 168–172. IEEE, 1994. 5
- [8] Zeyuan Chen, Yinbo Chen, Jingwen Liu, Xingqian Xu, Vidit Goel, Zhangyang Wang, Humphrey Shi, and Xiaolong Wang. Videoir: Learning video implicit neural representation for continuous space-time super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2047–2057, 2022. 3
- [9] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1290–1299, 2022. 2, 6
- [10] Xianhang Cheng and Zhenzhong Chen. Video frame interpolation via deformable separable convolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10607–10614, 2020. 3
- [11] Xianhang Cheng and Zhenzhong Chen. Multiple video frame interpolation via enhanced deformable separable convolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):7029–7045, 2021. 3
- [12] 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 *European Conference on Computer Vision*, pages 107–123. Springer, 2020. 3
- [13] 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*, pages 10663–10671, 2020. 1, 2, 3, 4, 5, 7
- [14] Alexey Dosovitskiy, Philipp Fischery, Eddy Ilg, Caner Hazirbas, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, Thomas Brox, et al. FlowNet: Learning optical flow with convolutional networks. In *ICCV*, 2015. 8
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 6
- [16] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *ICCV*, 2021. 2
- [17] Shubham Goel, Georgios Pavlakos, Jathushan Rajasegaran, Angjoo Kanazawa\*, and Jitendra Malik\*. Humans in 4D: Reconstructing and tracking humans with transformers. In *International Conference on Computer Vision (ICCV)*, 2023. 8
- [18] 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*, pages 14004–14013, 2020. 3
- [19] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 6
- [20] Weihua He, Kaichao You, Zhendong Qiao, Xu Jia, Ziyang Zhang, Wenhui Wang, Huchuan Lu, Yaoyuan Wang, and Jianxing Liao. Timereplayer: Unlocking the potential of event cameras for video interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17804–17813, 2022. 3
- [21] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 8
- [22] Hsuan-I Ho, Xu Chen, Jie Song, and Otmar Hilliges. Render in-between: Motion guided video synthesis for action interpolation. *arXiv preprint arXiv:2111.01029*, 2021. 1
- [23] Ping Hu, Simon Niklaus, Stan Sclaroff, and Kate Saenko. Many-to-many splatting for efficient video frame interpolation. In *CVPR*, 2022. 2, 3, 6, 7
- [24] Zhewei Huang, Tianyuan Zhang, Wen Heng, Boxin Shi, and Shuchang Zhou. Rife: Real-time intermediate flow estimation for video frame interpolation. *arXiv preprint arXiv:2011.06294*, 2020. 6
- [25] Zhewei Huang, Tianyuan Zhang, Wen Heng, Boxin Shi, and Shuchang Zhou. Real-time intermediate flow estimation for video frame interpolation. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIV*, pages 624–642. Springer, 2022. 2, 3, 6, 7, 8
- [26] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. *Advances in neural information processing systems*, 28, 2015. 3

- [27] Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik G. Learned-Miller, and Jan Kautz. Super slo-mo: High quality estimation of multiple intermediate frames for video interpolation. In *CVPR*, 2018. 1, 2, 3, 5, 6, 7, 8
- [28] Xin Jin, Longhai Wu, Jie Chen, Youxin Chen, Jayoon Koo, and Cheul-hee Hahm. A unified pyramid recurrent network for video frame interpolation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2023. 3, 6, 7
- [29] Xin Jin, Longhai Wu, Guotao Shen, Youxin Chen, Jie Chen, Jayoon Koo, and Cheul-hee Hahm. Enhanced bi-directional motion estimation for video frame interpolation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5049–5057, 2023. 2, 3, 5, 6, 7, 8
- [30] Glenn Jocher, Ayush Chaurasia, and Jing Qiu. YOLO by Ultralytics, 2023. 6
- [31] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. 6
- [32] Tarun Kalluri, Deepak Pathak, Manmohan Chandraker, and Du Tran. Flavr: Flow-agnostic video representations for fast frame interpolation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2071–2082, 2023. 3
- [33] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9404–9413, 2019. 6
- [34] Muhammed Kocabas, Chun-Hao P. Huang, Otmar Hilliges, and Michael J. Black. PARE: Part attention regressor for 3D human body estimation. In *ICCV*, 2021. 8
- [35] Lingtong Kong, Boyuan Jiang, Donghao Luo, Wenqing Chu, Xiaoming Huang, Ying Tai, Chengjie Wang, and Jie Yang. Ifrnet: Intermediate feature refine network for efficient frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1969–1978, 2022. 3, 8
- [36] Hyeongmin Lee, Taeh 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*, pages 5316–5325, 2020. 2, 3, 5, 6, 7
- [37] Siyao Li, Shiyu Zhao, Weijiang Yu, Wenxiu Sun, Dimitris N. Metaxas, Chen Change Loy, and Ziwei Liu. Deep animation video interpolation in the wild. In *CVPR*, 2021. 1
- [38] Yinxiao Li, Pengchong Jin, Feng Yang, Ce Liu, Ming-Hsuan Yang, and Peyman Milanfar. Comisr: Compression-informed video super-resolution. In *ICCV*, 2021. 2
- [39] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *CVPR*, 2021. 1, 2
- [40] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 6
- [41] 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*, pages 8794–8802, 2019. 3
- [42] Ziwei Liu, Raymond A Yeh, Xiaoou Tang, Yiming Liu, and Aseem Agarwala. Video frame synthesis using deep voxel flow. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4463–4471, 2017. 3
- [43] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 6
- [44] Liying Lu, Ruizheng Wu, Huajia Lin, Jiangbo Lu, and Jiaya Jia. Video frame interpolation with transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3532–3542, 2022. 3
- [45] Nikolaus Mayer, Eddy Ilg, Philip Häusser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *CVPR*, 2016. 8
- [46] 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. 5
- [47] Simone Meyer, Oliver Wang, Henning Zimmer, Max Grosse, and Alexander Sorkine-Hornung. Phase-based frame interpolation for video. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1410–1418, 2015. 3
- [48] Simone Meyer, Abdelaziz Djelouah, Brian McWilliams, Alexander Sorkine-Hornung, Markus Gross, and Christopher Schroers. Phasenet for video frame interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 498–507, 2018. 3
- [49] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. Ieee, 2016. 6
- [50] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3883–3891, 2017. 4
- [51] Simon Niklaus and Feng Liu. Context-aware synthesis for video frame interpolation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1701–1710, 2018. 3
- [52] Simon Niklaus and Feng Liu. Softmax splatting for video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5437–5446, 2020. 3
- [53] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive convolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 670–679, 2017. 3

- [54] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive separable convolution. In *Proceedings of the IEEE international conference on computer vision*, pages 261–270, 2017. 3
- [55] 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*, pages 14539–14548, 2021. 3
- [56] Tomer Peleg, Pablo Szekely, Doron Sabo, and Omry Sendik. Im-net for high resolution video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2398–2407, 2019. 3
- [57] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 724–732, 2016. 1, 4
- [58] Jathushan Rajasegaran, Georgios Pavlakos, Angjoo Kanazawa, Christoph Feichtenhofer, and Jitendra Malik. On the benefits of 3d pose and tracking for human action recognition. In *CVPR*, 2023. 8
- [59] Anurag Ranjan, David T. Hoffmann, Dimitrios Tzionas, Siyu Tang, Javier Romero, and Michael J. Black. Learning multi-human optical flow. *Int. J. Comput. Vis.*, 128(4):873–890, 2020. 8
- [60] Fitsum Reda, Janne Kontkanen, Eric Tabellion, Deqing Sun, Caroline Pantofaru, and Brian Curless. Film: Frame interpolation for large motion. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII*, pages 250–266. Springer, 2022. 3
- [61] Fitsum A Reda, Deqing Sun, Aysegul Dundar, Mohammad Shoeybi, Guilin Liu, Kevin J Shih, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Unsupervised video interpolation using cycle consistency. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 892–900, 2019. 3
- [62] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 4
- [63] Zhihao Shi, Xiangyu Xu, Xiaohong Liu, Jun Chen, and Ming-Hsuan Yang. Video frame interpolation transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17482–17491, 2022. 3
- [64] Hyeonjun Sim, Jihyong Oh, and Munchurl Kim. Xvfi: Extreme video frame interpolation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14489–14498, 2021. 1, 2, 3, 4, 5
- [65] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 2, 3, 4
- [66] Shuochen Su, Mauricio Delbracio, Jue Wang, Guillermo Sapiro, Wolfgang Heidrich, and Oliver Wang. Deep video deblurring for hand-held cameras. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1279–1288, 2017. 1, 2, 3, 4
- [67] Masato Tamura, Rahul Vishwakarma, and Ravigopal Vennelakanti. Hunting group clues with transformers for social group activity recognition. In *ECCV*, 2022. 2
- [68] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 402–419. Springer, 2020. 4
- [69] Stepan Tulyakov, Daniel Gehrig, Stamatiios Georgoulis, Julius Erbach, Mathias Gehrig, Yuanyou Li, and Davide Scaramuzza. Time lens: Event-based video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16155–16164, 2021. 3
- [70] Stepan Tulyakov, Alfredo Bochicchio, Daniel Gehrig, Stamatiios Georgoulis, Yuanyou Li, and Davide Scaramuzza. Time lens++: Event-based frame interpolation with parametric non-linear flow and multi-scale fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17755–17764, 2022. 3
- [71] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 3
- [72] Liao Wang, Jiakai Zhang, Xinhang Liu, Fuqiang Zhao, Yanshun Zhang, Yingliang Zhang, Minye Wu, Jingyi Yu, and Lan Xu. Fourier plenotrees for dynamic radiance field rendering in real-time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13524–13534, 2022. 2
- [73] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [74] Chao-Yuan Wu, Nayan Singhal, and Philipp Krähenbühl. Video compression through image interpolation. In *ECCV*, 2018. 1
- [75] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 6
- [76] Hao-fei Xu, Jing Zhang, Jian-fei Cai, Hamid Reza Tofighi, and Dacheng Tao. Gmflow: Learning optical flow via global matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8121–8130, 2022. 5, 8
- [77] Xiangyu Xu, Li Siyao, Wenxiu Sun, Qian Yin, and Ming-Hsuan Yang. Quadratic video interpolation. *Advances in Neural Information Processing Systems*, 32, 2019. 3
- [78] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. Vit-pose: Simple vision transformer baselines for human pose estimation. *Advances in Neural Information Processing Systems*, 35:38571–38584, 2022. 2, 6
- [79] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. Vit-pose++: Vision transformer for generic body pose estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 6
- [80] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented



- flow. *International Journal of Computer Vision*, 127(8): 1106–1125, 2019. [1](#), [2](#), [3](#), [4](#), [7](#)
- [81] Vickie Ye, Georgios Pavlakos, Jitendra Malik, and Angjoo Kanazawa. Decoupling human and camera motion from videos in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [8](#)
- [82] Hangjie Yuan, Dong Ni, and Mang Wang. Spatio-temporal dynamic inference network for group activity recognition. In *ICCV*, 2021. [2](#)
- [83] Guozhen Zhang, Yuhan Zhu, Haonan Wang, Youxin Chen, Gangshan Wu, and Limin Wang. Extracting motion and appearance via inter-frame attention for efficient video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5682–5692, 2023. [2](#), [3](#), [6](#), [7](#)
- [84] Jiakai Zhang, Xinhang Liu, Xinyi Ye, Fuqiang Zhao, Yanshun Zhang, Minye Wu, Yingliang Zhang, Jingyi Yu, and Lan Xu. Editable free-viewpoint video using a layered neural representation. 2021. [2](#)
- [85] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. [7](#)
- [86] Honglu Zhou, Asim Kadav, Aviv Shamsian, Shijie Geng, Farley Lai, Long Zhao, Ting Liu, Mubbasir Kapadia, and Hans Peter Graf. Composer: Compositional reasoning of group activity in videos with keypoint-only modality. In *ECCV*, 2022. [2](#)