

Motion-aware Latent Diffusion Models for Video Frame Interpolation

Zhilin Huang
Shenzhen International Graduate
School, Tsinghua University
Pengcheng Laboratory
Shenzhen, China
zerinhwang03@pku.edu.cn

Chujun Qin
China Southern Power Grid
Guangzhou, China
chujun.qin@pku.edu.cn

Zikun Zhou*
Pengcheng Laboratory
Shenzhen, China
zhouzikunhit@gmail.com

Yijie Yu
Shenzhen International Graduate
School, Tsinghua University
Pengcheng Laboratory
Shenzhen, China
yyj23@mails.tsinghua.edu.cn

Bing Zheng
Shenzhen International Graduate
School, Tsinghua University
Pengcheng Laboratory
Shenzhen, China
zhengb21@mails.tsinghua.edu.cn

Yaowei Wang
Pengcheng Laboratory
Harbin Institute of Technology
Shenzhen, China
wangyw@pcl.ac.cn

Ling Yang
Peking University
Beijing, China
yangling0818@163.com

Xiawu Zheng
Xiamen University
Pengcheng Laboratory
Xiamen, China
zhengxiawu@xmu.edu.cn

Wenming Yang*
Shenzhen International Graduate
School, Tsinghua University
Pengcheng Laboratory
Shenzhen, China
yangelwm@163.com

ABSTRACT

With the advancement of AIGC, video frame interpolation (VFI) has become a crucial component in existing video generation frameworks, attracting widespread research interest. For the VFI task, the motion estimation between neighboring frames plays a crucial role in avoiding motion ambiguity. However, existing VFI methods always struggle to accurately predict the motion information between consecutive frames, and this imprecise estimation leads to blurred and visually incoherent interpolated frames. In this paper, we propose a novel diffusion framework, **Motion-Aware latent Diffusion models (MADIFF)**, which is specifically designed for the VFI task. By incorporating motion priors between the conditional neighboring frames with the target interpolated frame predicted throughout the diffusion sampling procedure, MADIFF progressively refines the intermediate outcomes, culminating in generating both visually smooth and realistic results. Extensive experiments conducted on benchmark datasets demonstrate that our method achieves state-of-the-art performance significantly outperforming existing approaches, especially under challenging scenarios involving dynamic textures with complex motion.

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '24, October 28-November 1, 2024, Melbourne, VIC, Australia

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0686-8/24/10

<https://doi.org/10.1145/3664647.3680846>

CCS CONCEPTS

• **Computing methodologies** → **Reconstruction**.

KEYWORDS

Generative Models, Diffusion Models, Video Frame Interpolation

ACM Reference Format:

Zhilin Huang, Yijie Yu, Ling Yang, Chujun Qin, Bing Zheng, Xiawu Zheng, Zikun Zhou, Yaowei Wang, and Wenming Yang. 2024. Motion-aware Latent Diffusion Models for Video Frame Interpolation. In *Proceedings of the 32nd ACM International Conference on Multimedia (MM '24)*, October 28-November 1, 2024, Melbourne, VIC, Australia. *Proceedings of the 32nd ACM International Conference on Multimedia (MM'24)*, October 28-November 1, 2024, Melbourne, Australia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3664647.3680846>

1 INTRODUCTION

Video frame interpolation (VFI) aims to generate intermediate frames between two consecutive video frames. It is commonly used to increase the frame resolution, e.g., to produce slow-motion content [29]. VFI has also been applied to video compression [69], video generation [21] and animation production [57].

Existing VFI methods [29, 31, 37, 48, 71] are mostly based on deep neural networks. The early deep learning methods always rely on 3D convolution [28] or RNNs [9] to model the contextual correlations of conditional neighboring frames. Inspired by the success of generative adversarial networks (GANs) in image synthesis, more recent attempts have been made to develop video frame interpolation methods by incorporating the adversarial loss. Owing to the remarkable generative capabilities of GANs, these methods exhibit significant efficacy in predicting interpolated video frames that

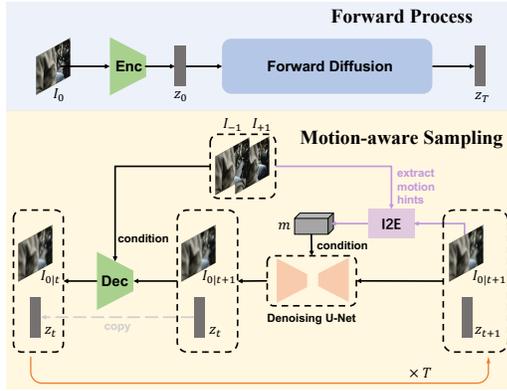


Figure 1: Overview of the diffusion processes in MADIFF. The encoder and decoder enable projection between image and latent spaces, and the diffusion processes take place in the latent space. I2E denotes image-to-event generator [78] which have capability of generating event volume by taking two continuous frames as input. And m denotes motion hints extracted from I2E.

possess natural and consistent content, holding the state-of-the-art in the literature. However, these deep learning-based methods for VFI always tend to generate unrealistic texture, artifacts and low-perceptual results. The reason is that the primary contributor to the optimization objective — and consequently the final performance of the model — remains the L1/L2-based distortion loss between their outputs and the ground-truth interpolated frames [15]. This type of loss may not accurately reflect the perceptual quality of the interpolated videos. Consequently, [14] indicates that current techniques, despite attaining elevated PSNR scores, often fall short in terms of perceptual quality, particularly in demanding situations characterized by dynamic textures and complex motion patterns.

Recently, de-noising diffusion probabilistic models (DDPMs) have been gaining widespread interest and have risen to prominence as state-of-the-art solutions across various areas [20, 21, 23, 24, 61, 72]. These diffusion models have demonstrated exceptional capabilities in creating realistic and perceptually-enhanced images and videos, reportedly outperforming other generative models. [15, 67] are among the early methods that have explored the application of diffusion models for VFI tasks. Specifically, they address the VFI tasks as a form of conditional image generation by taking neighboring frames into the de-noising network for the target interpolated frame generation.

However, these methods fail to explicitly model the inter-frame motions between the interpolated frames and the given neighboring conditional frames, a crucial factor in preventing the generation of blurred interpolated frames due to motion ambiguity. This is particularly important in complex dynamic scenes that involve intricate motions, occlusions, or abrupt changes in brightness.

To tackle these challenges, in this paper, we propose a novel latent diffusion framework, **Motion-Aware latent Diffusion model (MADIFF)** for the video frame interpolation task. Specifically, we follow [15] to build upon MADIFF by adopting recently proposed latent diffusion models (LDMs) [55]. LDMs consist of an autoencoder that maps images into a latent space and a de-noising U-net, which

carries out the reverse diffusion process within that latent space, forming the foundation of our framework. To incorporate the inter-frame motion priors between given conditional neighboring frames with the interpolated frames into MADIFF, we propose a novel vector quantized motion-aware generative adversarial network, named **VQ-MAGAN**. In particular, we initially utilize a pre-trained EventGAN [78] to predict the event volume which reflects the pixel-level intensity changes between two consecutive frames. Subsequently, the event volumes between the interpolated frame and the two neighboring conditional frames are employed as motion hints to enhance image reconstructions within the decoder of VQ-MAGAN. In this design, VQ-MAGAN possesses the capability to predict the target interpolated frame by aggregating contextual details from the given neighboring frames under the guidance of inter-frame motion hints. Furthermore, for the de-noising process in LDMs, we also incorporate motion hints between the interpolated frame and the two neighboring frames as additional conditions.

During the training process of both VQ-MAGAN and the de-noising U-net, we directly utilize the ground-truth interpolated frame for extracting inter-frame motion hints between interpolated frame and neighboring conditional frames. Since the ground-truth interpolated frame is unknown during the sampling process of LDM, extracting motion hints between the interpolated frame and the conditional neighboring frames is not feasible. To eliminate the discrepancy of motion hints extraction between the sampling phase and the training phase, making the motion hints in the sampling process available, we propose a novel motion-aware sampling procedure (**MA-SAMPLING**). Specifically, during the sampling process, we use the coarse interpolated frame predicted in the previous time step to extract inter-frame motion hints in conjunction with the conditional neighboring frames. The extracted motion hints are then fed into both VQ-MAGAN and the de-noising U-net for the prediction of the interpolated frame in the current time step. By refining the interpolated frames progressively, our MADIFF can effectively integrate the inter-frame motion hints into the sampling process, resulting in visually smooth and realistic frames.

Extensive experiments on various VFI benchmark datasets, encompassing both low and high resolution content (up to 4K), demonstrate that our MADIFF achieves the state-of-the-art performance significantly outperforming existing approaches, especially under challenging scenarios involving dynamic textures with complex motion. Our contributions are summarized as follows:

- We propose a novel vector quantized motion-aware generative adversarial network, named **VQ-MAGAN**, which fully incorporates the inter-frame motion hints between the target interpolated frame and the given neighboring conditional frames into the prediction of the interpolated frame.
- We propose a novel motion-aware sampling procedure, named **MA-SAMPLING**, to eliminate the discrepancy of motion hints extraction between the sampling phase and the training phase, making the extraction of motion hints in the sampling process feasible and refine the predicted interpolated frames progressively.
- We demonstrate, through quantitative and qualitative experiments, that the proposed method achieves the state-of-the-art performance outperforming existing approaches.

2 RELATED WORK

2.1 Image-to-Event Generation

Event cameras are a novel bio-inspired asynchronous sensor [18, 41] with advantages such as high dynamic range, high temporal resolution, and low power consumption [4, 18, 41]. Event cameras have the potential to provide solutions for a wide range of visually challenging scenarios and has already found extensive applications [2, 5, 22, 33, 38–40, 46, 54, 63, 64, 70, 77]. However, acquiring event stream data is expensive and requires specific devices. Recently, several studies have exploited the event simulation from the standard camera, i.e. simulating the event stream from continuous images or video sequences. Kaiser et al. [30] generate events by simply applying a threshold on the image difference. A positive or negative event is generated depending on the pixel’s intensity difference. Pix2NVS [3] computes per-pixel luminance from conventional video frames. The technique generates synthetic events with inaccurate timestamps clustered to frame timestamps. [76, 78] propose a GAN-based method to generate the event volume by simply taking two continuous image frames as input. By modeling the correlation between continuous frames and event volumes, image-to-event generation models can capture inter-frame motion hints, which can serve as auxiliary guidance for the VFI task. In our MADIFF, we directly utilized the pre-trained EventGAN [78] to extract motion hints for guiding the interpolation process.

2.2 Video Frame Interpolation

Current approaches to video frame interpolation (VFI) that leverage deep learning can typically be divided into two main categories: those that are based on flow estimation and those that rely on kernel methods. Existing flow-based methods [29, 35, 42, 44, 47, 48, 50, 51, 56, 71] primarily generate intermediate frames by analyzing the pixel motion between consecutive frames. Specifically, these methods use optical flow estimation to capture the movement of objects within the scene, and then synthesize intermediate frames based on these estimations. These methods can be founded on various assumptions and algorithms, such as assuming uniform motion within the scene or employing machine learning models to predict the optical flow field, thereby enhancing the accuracy and visual quality of the interpolated frames. On the other hand, Existing kernel-based video frame interpolation methods [7, 8, 11, 16, 37, 49] typically utilize kernel functions to estimate the relationships between pixels across different frames, generating intermediate frames by performing a weighted average of pixel values based on known frames. These methods can flexibly adapt to various motion patterns in the scene by selecting appropriate kernel sizes and shapes, thus effectively smoothing the visual transitions between frames while preserving details. In addition to these two main categories, there are also some methods integrate flow-based and kernel-based approaches to better synthesis performance [10, 31].

Meanwhile, several works attempt to introduce the event stream as a supplemental conditions for guiding accurate models between the neighboring frames and interpolated frames. However, almost event-based methods require taking the accurate ground-truth event streams as inputs which is inconvenient for practical application. In contrast, our MADIFF directly utilizes off-the-shelf pre-trained image-to-event models for providing the motion hints

in the subsequent interpolated frame generation process. Besides, our MADIFF refines the target interpolated frame in a progressive manner, which is quite different from previous VFI methods that predict interpolated frame in a one-shot manner.

Recently, denoising diffusion probabilistic models (DDPMs) have gained increasing attention and emerged as a new state-of-the-art in several areas of computer vision. [15, 67] are the most early methods exploiting the application of DDPMs for VFI tasks. Compared to previous VFI methods, diffusion-based methods show satisfactory perceptual performance, especially on dynamic textures with complex motions. However, existing diffusion-based methods simply define the VFI task as a conditional image generation task by considering the neighboring frames as conditions, neglecting to explicitly model inter-frame motions which are a crucial factor for generating realistic and visually smooth interpolated frames. Compared with these methods, we propose a novel diffusion framework, MADIFF, to effectively incorporate motion hints from existing motion-related models into diffusion models in a progressive manner.

3 PRELIMINARY

3.1 Representation of Event Volume

Each event \mathbf{e} can be represented by a tuple (x, y, t, p) , where x and y represent the spatial position of the event, t represents the timestamp, and $p = \pm 1$ indicates its polarity.

As described in [78], the event are presented for easily processing: events are scattered into a fixed size 3D spatio-temporal volume, where each event, (x, y, t, p) is inserted into the volume, which has $B = 9$ temporal channels, with a linear kernel:

$$t_i^* = (B - 1) \cdot \frac{t_i - t_1}{t_N - t_1} \quad (1)$$

$$V(x, y, t) = \sum_i \max(0, 1 - |t - t_i^*|) \quad (2)$$

This retains the distribution of the events in x - y - t space, and has shown success in a number of tasks [6, 53, 79].

Since the event volume generated by EventGAN [78] are the concatenation of event volumes of different polarity along the time dimension, the final event volume which is strictly non-negative. In our MADIFF, we directly utilize the event volume generated by EventGAN as the inter-frame motion hints.

3.2 Latent Diffusion Models

Latent diffusion models (LDMs) [55] is variant of de-noising diffusion probabilistic models which executes the de-noising process in the latent space of an autoencoder, namely $\mathcal{E}(\cdot)$ and $\mathcal{D}(\cdot)$, implemented as pre-trained VQ-GAN [17] or VQ-VAE [65]. Compared with executing de-noising process in the pixel-level data, LDM can reduce computational costs while preserving high visual quality.

For the training of LDM, the given latent code z for a randomly sampled training image x is converted to noise with a Markov process defined by the transition kernel:

$$q(z_t | z_{t-1}) = \mathcal{N}(z_t; \sqrt{\alpha_t} z_{t-1}, (1 - \alpha_t) \mathbf{I}) \quad (3)$$

where $t = 1, 2, \dots, T$, $z_0 = z$, and α_t is a hyper-parameter that controls the rate of noise injection. When the amount of noise is sufficiently large, z_T becomes approximately distributed according

to $\mathcal{N}(0, \mathbf{I})$. In order to convert noise back to data for sample generation, the reverse diffusion process is estimated by learning the reverse transition kernel:

$$p_{\theta}(z_{t-1} | z_t) = \mathcal{N}(\mu_{\theta}(z_t), \Sigma_{\theta}(z_t)) \quad (4)$$

and then take it as an approximation to $q(z_{t-1} | z_t)$. Following Ho et al. [20], $\mu_{\theta}(z_t)$ is parameterized with a neural network $\epsilon_{\theta}(z_t, t)$ (called the score model [59, 60]) and Σ_{θ} is fixed to be a constant. The score model can be optimized with de-noising score matching [26, 66]. During sample generation, within each time step, the de-noising U-net initially predicts the \hat{z}_0 . Finally, the decoder of VQ-GAN or VQ-VAE generates the image \hat{I}_0 from the de-noised latent representation \hat{z}_0 , disregarding any contextual information.

4 METHODS

4.1 Motion Hints Extraction

In MADIFF, we utilize the pre-trained EventGAN [78] to capture inter-frame motion hints between the interpolated frame with the conditional neighboring frames. Specifically, given the interpolated frame $I_0 \in \mathbb{R}^{H \times W \times 3}$ and two conditional neighboring frames $I_{-1}, I_{+1} \in \mathbb{R}^{H \times W \times 3}$, where I_{-1} denotes the previous frame and I_{+1} denotes the next frame. The motion hints extraction process are formulated as follows:

$$m_{-1 \rightarrow 0} = f_{2E}(I_{-1}, I_0) \quad (5)$$

$$m_{0 \rightarrow +1} = f_{2E}(I_0, I_{+1}) \quad (6)$$

where $f_{2E}(\cdot)$ is the pre-trained EventGAN, $m_{i \rightarrow j}$ denotes the extracted motion hints from the frame i to the frame j . In practice, we directly use the predicted event volume $EV_{i \rightarrow j} \in \mathbb{R}^{H \times W \times (2 \times B)}$ as $m_{i \rightarrow j}$. Besides, MADIFF we proposed is a general framework which can incorporate different motion-related models easily. More details please refer to Section 5.5.2.

4.2 VQ-MAGAN

Implementation Details. For the consideration of motion information between the interpolated frame and neighboring frames is crucial for the VFI task, we propose a novel VQ-GAN, namely VQ-MAGAN as presented in Figure 2.

The encoder \mathcal{E} produces the latent encoding $z_0 = \mathcal{E}(I_0)$ by taking the given ground-truth target frame $I_0 \in \mathbb{R}^{H \times W \times 3}$ as the input, where $z_0 \in \mathbb{R}^{\frac{H}{f} \times \frac{W}{f} \times 3}$, and f is a hyper-parameter. In practice, we set $f = 32$ following [15].

Then the decoder \mathcal{D} reconstruct target frame \hat{I}_0 by taking z_0 and the feature pyramids ϕ_{-1}, ϕ_{+1} which are extracted by \mathcal{E} from two neighboring frames I_{-1}, I_{+1} following [15]. Moreover, we utilize the motion hints extractor to capture the inter-frame motion hints $m_{-1 \rightarrow 0}$ and $m_{0 \rightarrow +1}$ between the ground-truth target frame I_0 with the neighboring frames I_{-1} and I_{+1} . And then we take the $m_{-1 \rightarrow 0}$ and $m_{0 \rightarrow +1}$ as an additional guidance for the layer-wisely contextual aggregation in the decoder \mathcal{D} through a motion-aware warp (MA-WARP) module. Specifically, for the feature $h_0^l \in \mathbb{R}^{U \times V \times C}$ of interpolated frame and the motion hints $m_{-1 \rightarrow 0}, m_{0 \rightarrow +1} \in \mathbb{R}^{H \times W \times (2 \times B)}$, the MA-WARP in l -th decoder layer firstly reshape the motion hints to the resolution of $U \times V$, obtained $m_{-1 \rightarrow 0}^l, m_{0 \rightarrow +1}^l \in \mathbb{R}^{U \times V \times (2 \times B)}$. Then for each motion hints, a 2-channel offset map $\Omega_{-1 \rightarrow 0}^l$ and

$\Omega_{0 \rightarrow +1}^l$, respectively, which reflects the pixel-level feature correlations from the neighboring frame to the target interpolated frame is generated through a learnable neural networks: $f_{\Omega}(\cdot)$

$$\Omega_{0 \leftarrow -1}^l = f_{\Omega}(h_0^l, m_{-1 \rightarrow 0}^l, \phi_{-1}^l) \quad (7)$$

$$\Omega_{0 \leftarrow +1}^l = f_{\Omega}(h_0^l, m_{0 \rightarrow +1}^l, \phi_{+1}^l) \quad (8)$$

After that, a warp function $f_{warp}(\cdot)$ proposed in [25] is introduced and serves as a aggregation mechanism:

$$h_{0 \leftarrow -1}^l = f_{warp}(\Omega_{0 \leftarrow -1}^l, \phi_{-1}^l) \quad (9)$$

$$h_{0 \leftarrow +1}^l = f_{warp}(\Omega_{0 \leftarrow +1}^l, \phi_{+1}^l) \quad (10)$$

The MA-WARP also generates a gate map $g \in [0, 1]^{U \times V \times 1}$ to account for occlusion [29], and a residual map $\delta \in \mathbb{R}^{U \times V \times C}$ to further enhance the performance:

$$\tilde{h}_0^l = g \cdot h_{0 \leftarrow -1}^l + (1 - g) \cdot h_{0 \leftarrow +1}^l + \delta, \quad (11)$$

$$g = f_g(h_{0 \leftarrow -1}^l, h_{0 \leftarrow +1}^l), \quad (12)$$

$$\delta = f_{\delta}(h_0^l). \quad (13)$$

where both $f_g(\cdot)$ and $f_{\delta}(\cdot)$ are learnable neural networks, \tilde{h}_0^l is the output of MA-WARP in l -th decoder layer. By hierarchically applying MA-WARP in the decoder layer, VQ-MAGAN is able to fully utilize the motion hints for accurately aggregating pyramid contexts from neighboring frames. Compared with VQ-FIGAN proposed in [15], our VQ-MAGAN has capability of incorporating the inter-frame motions between the target interpolated frames with the conditional neighboring frames.

Training VQ-MAGAN. For the training of VQ-MAGAN, we follow the original training settings of VQGAN in [17, 55], where the loss function consists of an LPIPS-based [75] perceptual loss, a patch-based adversarial loss [27] and a latent regularization term based on a vector quantization (VQ) layer [65]. Particularly, we use the ground-truth interpolated frame to extract motion hints with given conditional neighboring frames.

Since the ground-truth target frame are provided for extracting motion hints during the training process of VQ-MAGAN, the reconstruction task may becomes much easier for VQ-MAGAN, potentially degrading the performance of reconstruction in the inference stage. To address this issue, we only utilize the motion hints with a probability of 0.5 during the training stage to assist in the reconstruction process of VQ-MAGAN. The pseudo code is provided in the Appendix. While we replace the ground-truth interpolated frame with the predicted one from the previous time step during the sampling procedure as described in Section 4.4.

4.3 De-noising with Conditional Motion Hints

Implementation Details. The encoder of trained VQ-MAGAN allows us to access a compact latent space in which we perform forward diffusion by gradually adding Gaussian noise to the latent z_0 of the target frame I_0 according to a pre-defined noise schedule, and learn the reverse (de-noising) process to perform conditional generation. Moreover, compared with previous methods, we additionally incorporate the dynamic inter-frame motion hints into the de-noising process. To this end, we adopt the noise-prediction parameterization [20] of diffusion models and train a de-noising U-Net by

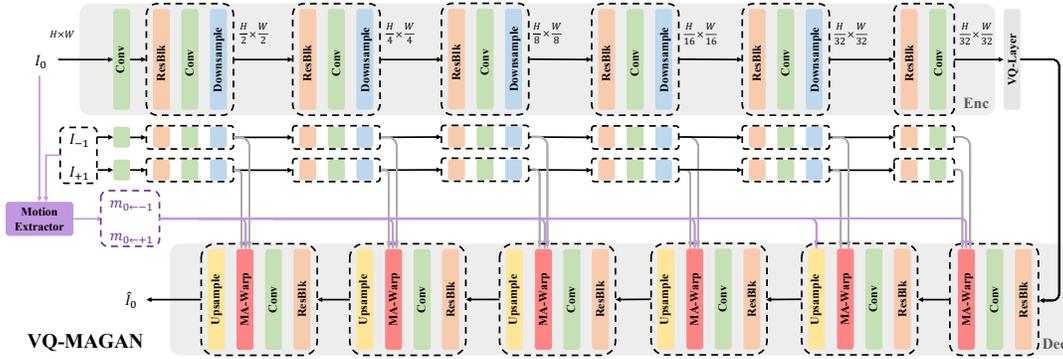


Figure 2: The architecture of the vector quantized motion-aware generative adversarial network, VQ-MAGAN. In practice, the motion extractor is image-to-event generator [78], $m_{i \rightarrow j}$ denotes the inter-frame motion hints between frame i and j .

minimizing the re-weighted variational lower bound on the conditional log-likelihood $\log p_\theta(z_0|z_{-1}, z_{+1}, m_{-1 \rightarrow 0}, m_{0 \rightarrow +1})$, where z_{-1} and z_{+1} are the latent representations of the two conditional neighboring frames, and $m_{i \rightarrow j}$ denotes the motion hints between frame i and j as described in Section 4.1.

Specifically, the de-noising U-Net ϵ_θ takes as input the noisy latent representation z_t for the target frame I_0 (sampled from the t -th step in the forward diffusion process of length T), the diffusion step t , as well as the conditioning latent representations z_{-1}, z_{+1} for the neighboring frames I_{-1}, I_{+1} . It is trained to predict the noise added to z_0 at each time step t by minimizing

$$\mathcal{L} = \mathbb{E} \left[\|\epsilon - \epsilon_\theta(z_t, t, z_{-1}, z_{+1}, m_{-1 \rightarrow 0}, m_{0 \rightarrow +1})\|^2 \right] \quad (14)$$

where $t \sim \mathcal{U}(1, T)$. The derivation and full details of the training procedure of ϵ_θ are provide in Appendix. Intuitively, the training is performed by alternately adding a random Gaussian noise to z_0 according to a pre-defined noising schedule, and having the network ϵ_θ predict the noise added given the step t , conditioning on z_{-1}, z_{+1} and $m_{i \rightarrow j}$.

Training of De-noising U-net. It is worth noting that during the training of de-noising U-net, we directly utilize the ground-truth interpolated frame for extracting inter-frame motion hints which is the same as the extraction process in the training of VQ-MAGAN.

4.4 MA-SAMPLING of MADIFF

As described above, both VQ-MAGAN and de-noising U-net are conditioned on the inter-frame motion hints extracted from the interpolated frame and the conditional neighboring frames. During the training stage of both VQ-MAGAN and de-noising U-net we directly using the ground-truth interpolated frame for the motion hints extraction in a teach-forcing manner. However, the interpolated frame is unknown in the sampling phase, rendering the extraction of inter-frame motion hints between the interpolated frame and neighboring frames infeasible. While the motion hints directly extracted from given neighboring frames are often inaccurate and cannot provide sufficient guidance, leading to sub-optimal performance as described in Table 5. For eliminating this discrepancy of motion hints extraction between the training and sampling

phase, making the motion hints in the sampling process available, we propose a novel MA-SAMPLING.

Before introducing MA-SAMPLING, we provide a review of the sampling process within traditional LDM for VFI tasks [15]: Within each time step, firstly the de-noising U-net ϵ_θ predicts the noise $\hat{\epsilon}$ conditioned on the latent representations z_{-1}, z_{+1} of neighboring frames I_{-1}, I_{+1} . Then $\hat{z}_0|t$ is obtained as follows:

$$\hat{\epsilon} = \epsilon_\theta(\hat{z}_t, t, z_{-1}, z_{+1}) \quad (15)$$

$$\hat{z}_0|t = \frac{1}{\sqrt{\alpha_t}} \left(\hat{z}_t - \frac{1 - \alpha_t}{\sqrt{1 - \alpha_t}} \hat{\epsilon} \right) \quad (16)$$

where $\epsilon_\theta(\cdot)$ is the de-noising U-net, $\hat{z}_0|t$ denotes the predicted z_0 at time step t (particularly, we denote $\hat{z}_0|1$ as \hat{z}_0 for simplification), \hat{z}_t is the noisy latent representation of the predicted $\hat{z}_0|_{t+1}$ obtained at previous time step $t+1$ during the sampling process. And \hat{z}_t can be calculated by using $\hat{\epsilon}$ and relevant parameters of the pre-defined forward process as eq (3). Finally, the decoder of VQ-GAN produces the image \hat{I}_0 from $\hat{z}_0|1$ with the help of feature pyramids ϕ_{-1}, ϕ_{+1} extracted by the encoder \mathcal{E} from I_{-1}, I_{+1} . Compared with the traditional sampling process mentioned above, our MA-SAMPLING has capability of incorporating accurate motion hints between the interpolated frame and the neighboring frames for progressively refining the predicted target frame. Specifically, at time step t , firstly the de-noising U-net ϵ_θ predicts the noise $\hat{\epsilon}$ conditioned on the latent representations z_{-1}, z_{+1} of neighboring frames and additional motion hints $\hat{m}_{0 \rightarrow +1|t+1}, \hat{m}_{-1 \rightarrow 0|t+1}$. Then $\hat{z}_0|t$ is obtained as follows:

$$\hat{\epsilon} = \epsilon_\theta(\hat{z}_t, t, z_{-1}, z_{+1}, \hat{m}_{-1 \rightarrow 0|t+1}, \hat{m}_{0 \rightarrow +1|t+1}) \quad (17)$$

$$\hat{z}_0|t = \frac{1}{\sqrt{\alpha_t}} \left(\hat{z}_t - \frac{1 - \alpha_t}{\sqrt{1 - \alpha_t}} \hat{\epsilon} \right) \quad (18)$$

where $\hat{m}_{-1 \rightarrow 0|t+1}, \hat{m}_{0 \rightarrow +1|t+1}$ are extracted from the predicted interpolated frame $\hat{I}_0|_{t+1}$ and neighboring frames I_{-1}, I_{+1} :

$$\hat{I}_0|_{t+1} = \mathcal{D}(\hat{z}_0|_{t+1}) \quad (19)$$

$$\hat{m}_{-1 \rightarrow 0|t+1} = f_{I2E}(I_{-1}, \hat{I}_0|_{t+1}) \quad (20)$$

$$\hat{m}_{0 \rightarrow +1|t+1} = f_{I2E}(\hat{I}_0|_{t+1}, I_{+1}) \quad (21)$$

And \hat{z}_{t-1} can be calculated using $\hat{\epsilon}$ and relevant parameters of the pre-defined forward process as sampling process of previous methods eq (3). Particularly, at time step T , motion hints

$\hat{m}_{-1 \rightarrow 0|T+1}$ and $\hat{m}_{0 \rightarrow +1|T+1}$ are both replaced with empty features $\mathbf{O} \in \mathbb{R}^{H \times W \times (2 \times B)}$. Finally, the decoder \mathcal{D} produces the interpolated frame $\hat{I}_{0|1}$ (for simplification we denote it as \hat{I}_0) from the de-noised latent representation $\hat{z}_{0|t+1}$ i.e. \hat{z}_0 , by fully considering feature pyramids ϕ_{-1}, ϕ_{+1} extracted by the encoder \mathcal{E} from I_{-1}, I_{+1} as contexts under the guidance of motion hints $\hat{m}_{-1 \rightarrow 0|1}, \hat{m}_{0 \rightarrow +1|1}$. Full details and pseudo code are provided in the Appendix.

5 EXPERIMENTS

5.1 Implementation Details

Regarding the diffusion processes, following [15], we adopt a linear noise schedule and a codebook size of 8192 for vector quantization in VQ-MAGAN. We sample from all baseline diffusion models with the DDIM [58] sampler for 200 steps. We also follow [15] to train the VQ-MAGAN using the Adam optimizer [34] and the de-noising U-net using the Adam-W optimizer [43], with the initial learning rates set to 10^{-5} and 10^{-6} respectively. Single NVIDIA V100 GPU were used for all training and evaluation.

5.2 Experimental Setup

5.2.1 Datasets. In our experiments, we strictly follow the dataset configuration specified in [15] for both training and evaluation. For training MADIFF, we utilize the widely adopted Vimeo90k dataset [71]. And additional samples from the BVI-DVC dataset [45] are adopted to better evaluate our VFI methods across a broader range of scenarios. The overall training set comprises 64,612 frame triplets from Vimeo90k-septuplets and 17,600 frame triplets from BVI-DVC, utilizing only the central three frames. For data augmentation, we randomly crop 256×256 patches and perform random flipping and temporal order reversing following [15]. For evaluation, MADIFF is validated on four widely recognized VFI benchmarks, including Middlebury [1], UCF-101 [62], DAVIS [52] and SNU-FILM [10]. The resolutions of these test datasets vary from 225×225 up to 4K, covering various levels of VFI complexity.

5.2.2 Evaluation. Following [15], we adopt a perceptual image quality metric LPIPS [75], FloLPIPS [12] for performance evaluation. These metrics exhibit a stronger correlation with human assessments of VFI quality in contrast to the traditionally utilized quality metrics, such as PSNR and SSIM [68]. We also evaluate FID [19] which measures the similarity between the distributions of interpolated and ground-truth frames. FID was previously used as a perceptual metric for video compression [74], enhancement [73] and colorization [32]. We also provide benchmark results based on PSNR and SSIM in the Appendix, noting that these are limited in reflecting the perceptual quality of interpolated content [14] and are therefore not the focus of this paper.

5.2.3 Baseline Methods. MADIFF was compared against 11 recent SOTA VFI methods, including BMBC [50], AdaCoF [37], CDFI [16], XVFI [56], ABME [51], IFRNet [35], VFIformer [44], ST-MFNet [13], FLAVR [31], MCVD [67] and LDMVFI [15]. It is noted that MCVD and LDMVFI are diffusion-based VFI method.

5.3 Quantitative Comparison

Table 1 shows the performance of the evaluated methods on the Middlebury, UCF-101 and DAVIS test sets. It can be observed that MADIFF consistently outperforms all the other VFI methods including both non-diffusion or diffusion-based methods. Moreover, we evaluate the performance on the four splits of the SNU-FILM dataset (the average motion magnitude increases from Easy to Extreme), as summarized in Table 2, which further demonstrates the superior perceptual quality of MADIFF, especially in the scenes including complex motions (SNU-FILM-Hard and Extreme). Notably, the performance of other diffusion-based VFI methods, namely MCVD and LDMVFI, are generally unsatisfactory, which suggests that simply defining the VFI task as a conditional image generation task, without explicitly modeling inter-frame motions, may not be adequate for generating realistic and visually smooth results. As presented in Table 1, the number of parameters in MADIFF is large. This is because we adopted the existing de-noising U-net [55] designed for generic image generation following [15] and introduced adapters for fusing motion hints into the VQ-MAGAN and de-noising U-net.

5.4 Qualitative Comparison

Figure 3 shows the comparison between example frames interpolated by MADIFF with the competing methods. We can observe that, non-diffusion-based methods (BMBC [50], VFIformer [44], IFRNet [35] and ST-MFNet [13]) tend to predict blurry results due to the L1/L2-based distortion loss. On the contrary, diffusion-based VFI method (LDMVFI [15]) is able to generate more realistic results, while it also produces several artifacts leads by motion ambiguity. By incorporating inter-frame motion hints as the guidance, our MADIFF has capability of predicting realistic results.

5.5 Ablation Study

In this section we experimentally validate and study the main components including the VQ-MAGAN and MA-SAMPLING in MADIFF.

5.5.1 Effectiveness of VQ-MAGAN and MA-SAMPLING. To validate the effectiveness of the proposed VQ-MAGAN and MA-SAMPLING, we conduct 4 experiments: (1) **Exp0**: without introducing the extracted motion hints from the pre-trained motion-related model into the VQ-MAGAN, and replacing the MA-SAMPLING with the original sampling procedure from LDMs; (2) **Exp1**: without introducing the extracted motion hints from the pre-trained motion-related model into the VQ-MAGAN, while still using MA-SAMPLING for extracting inter-frame motion hints as additional conditions of the de-nosing U-net; (3) **Exp2**: introducing the extracted motion hints from the pre-trained motion-related model into the VQ-MAGAN, while replacing MA-SAMPLING with original sampling procedure in LDMs; (4) **Exp3**: MADIFF equipped with both VQ-MAGAN and MA-SAMPLING. The experimental results are presented in Table 3. By comparing **Exp0**, **Exp1** and **Exp2**, we can observe that simply introducing MA-SAMPLING can only brings slightly performance improvements, while simply introducing motion hints to guide the contexts warping process in VQ-MAGAN can significantly improve the performance of MADIFF. Moreover, the results

	Middlebury			UCF-101			DAVIS			#P (M)
	LPIPS↓	FloLPIPS↓	FID↓	LPIPS↓	FloLPIPS↓	FID↓	LPIPS↓	FloLPIPS↓	FID↓	
BMBC	0.023	0.037	12.974	0.034	0.045	33.171	0.125	0.185	15.354	11.0
AdaCoF	0.031	0.052	15.633	0.034	0.046	32.783	0.148	0.198	17.194	21.8
CDFI	0.022	0.043	12.224	0.036	0.049	33.742	0.157	0.211	18.098	5.0
XVFI	0.036	0.070	16.959	0.038	0.050	33.868	0.129	0.185	16.163	5.6
ABME	0.027	0.040	11.393	0.058	0.069	37.066	0.151	0.209	16.931	18.1
IFRNet	0.020	0.039	12.256	0.032	0.044	28.803	0.114	0.170	14.227	5.0
VFIformer	0.031	0.065	15.634	0.039	0.051	34.112	0.191	0.242	21.702	5.0
ST-MFNet	N/A	N/A	N/A	0.036	0.049	34.475	0.125	0.181	15.626	21.0
FLAVR	N/A	N/A	N/A	0.035	0.046	31.449	0.209	0.248	22.663	42.1
MCVD	0.123	0.138	41.053	0.155	0.169	102.054	0.247	0.293	28.002	27.3
LDMVFI	0.019	0.044	16.167	0.026	0.035	26.301	0.107	0.153	12.554	439.0
MADIFF w/o MS	0.016	0.034	13.649	0.024	0.032	24.677	0.098	0.143	11.764	447.8
MADIFF	0.016	0.034	11.678	0.024	0.033	24.289	0.096	0.142	11.089	448.8

Table 1: Quantitative comparison of MADIFF ($f = 32$) and 11 tested methods on Middlebury, UCF-101 and DAVIS. Note ST-MFNet and FLAVR require four input frames so cannot be evaluated on Middlebury dataset which contains frame triplets. For each column, we highlight the best result in red and the second best in blue.

	SNU-FILM-Easy			SNU-FILM-Medium			SNU-FILM-Hard			SNU-FILM-Extreme		
	LPIPS↓	FloLPIPS↓	FID↓	LPIPS↓	FloLPIPS↓	FID↓	LPIPS↓	FloLPIPS↓	FID↓	LPIPS↓	FloLPIPS↓	FID↓
BMBC	0.020	0.031	6.162	0.034	0.059	12.272	0.068	0.118	25.773	0.145	0.237	49.519
AdaCoF	0.021	0.033	6.587	0.039	0.066	14.173	0.080	0.131	27.982	0.152	0.234	52.848
CDFI	0.019	0.031	6.133	0.036	0.066	12.906	0.081	0.141	29.087	0.163	0.255	53.916
XVFI	0.022	0.037	7.401	0.039	0.072	16.000	0.075	0.138	29.483	0.142	0.233	54.449
ABME	0.022	0.034	6.363	0.042	0.076	15.159	0.092	0.168	34.236	0.182	0.300	63.561
IFRNet	0.019	0.030	5.939	0.033	0.058	12.084	0.065	0.122	25.436	0.136	0.229	50.047
ST-MFNet	0.019	0.031	5.973	0.036	0.061	11.716	0.073	0.123	25.512	0.148	0.238	53.563
FLAVR	0.022	0.034	6.320	0.049	0.077	15.006	0.112	0.169	34.746	0.217	0.303	72.673
MCVD	0.199	0.230	32.246	0.213	0.243	37.474	0.250	0.292	51.529	0.320	0.385	83.156
LDMVFI	0.014	0.024	5.752	0.028	0.053	12.485	0.060	0.114	26.520	0.123	0.204	47.042
MADIFF w/o MS	0.013	0.021	5.157	0.025	0.048	10.919	0.058	0.110	23.143	0.125	0.210	49.435
MADIFF	0.013	0.021	5.334	0.027	0.049	11.022	0.058	0.107	22.707	0.118	0.198	44.923

Table 2: Quantitative comparison results on SNU-FILM which contains 4 subsets with different motion complexities (the average motion magnitude increases from Easy to Extreme). And VFIformer is not included because the GPU goes out of memory. For each column, we highlight the best result in red and the second best in blue.

Exps	VQ-MAGAN	MA-SAMPLING	Middlebury			UCF-101			DAVIS		
			LPIPS↓	FloLPIPS↓	FID↓	LPIPS↓	FloLPIPS↓	FID↓	LPIPS↓	FloLPIPS↓	FID↓
Exp0	×	×	0.022	0.043	17.202	0.027	0.036	25.578	0.113	0.157	12.250
Exp1	×	✓	0.020	0.037	11.632	0.028	0.037	26.885	0.109	0.153	11.558
Exp2	✓	×	0.016	0.034	13.649	0.024	0.032	24.677	0.098	0.143	11.764
Exp3	✓	✓	0.016	0.034	11.678	0.024	0.033	24.289	0.096	0.142	11.089

Table 3: Ablation study of main components of MADIFF. For each column, we highlight the best result in red and the second best in blue.

of Exp3 demonstrate that simultaneously applying VQ-MAGAN and MA-SAMPLING can help MADIFF achieve optimal performance.

5.5.2 Influence of Different Types of Motion Hints. In MADIFF, we propose a novel framework to bridge the motion-related prediction models and the diffusion-based VFI models. To demonstrate the generalizations of MADIFF, we conduct ablation study on the effect of different motion hints. Specifically, we utilize two types of motion hints for guiding the interpolated generation process

in VQ-MAGAN and MA-SAMPLING of MADIFF: (1) **Flow-based motion hints:** flow maps between the interpolated frame with the neighboring frames estimated by the pre-trained FastFlowNet [36]; (2) **Event-based motion hints:** event volumes between the interpolated frame with the neighboring frames predicted by the pre-trained EventGAN [78] as described in Section 4.1.

Experimental results are presented in Table 4, where the baseline refers to the MADIFF without incorporating motion hints in

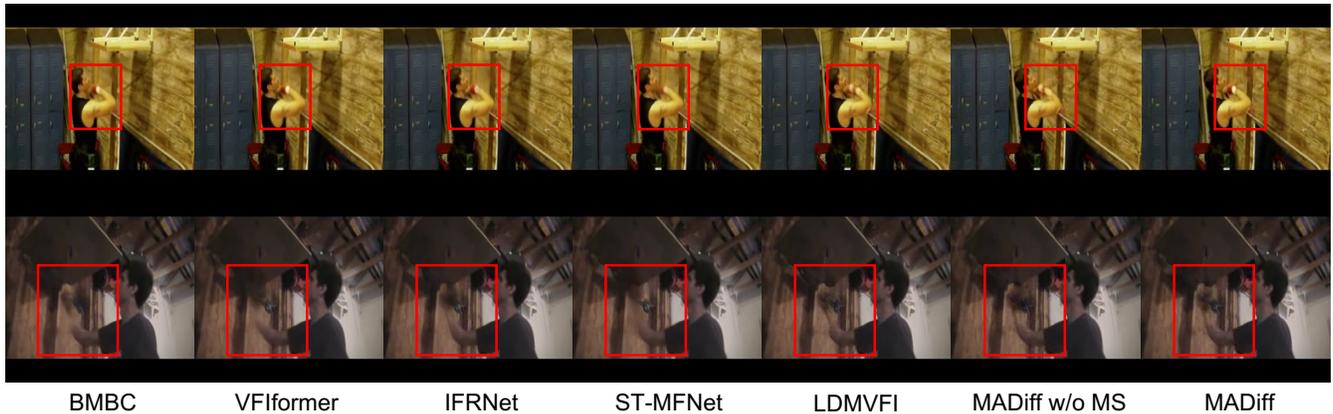


Figure 3: Visual examples of frames interpolated by the state-of-the-art methods and the proposed MADiff. Under large and complex motions, our method preserves the most high-frequency details, delivering superior perceptual quality.

Exps	Middlebury			UCF-101			DAVIS		
	LPIPS↓	FloLPIPS↓	FID↓	LPIPS↓	FloLPIPS↓	FID↓	LPIPS↓	FloLPIPS↓	FID↓
baseline	0.022	0.043	17.202	0.027	0.036	25.578	0.113	0.157	12.250
Flow-based Motion Hints	0.016	0.037	16.294	0.025	0.034	23.793	0.104	0.152	12.647
Event-based Motion Hints	0.016	0.034	11.678	0.024	0.033	24.289	0.096	0.142	11.089

Table 4: Ablation study of different type of motion hints in MADiff. For each column, we highlight the best result in red and the second best in blue.

Exps	Middlebury			UCF-101			DAVIS		
	LPIPS↓	FloLPIPS↓	FID↓	LPIPS↓	FloLPIPS↓	FID↓	LPIPS↓	FloLPIPS↓	FID↓
baseline	0.022	0.043	17.202	0.027	0.036	25.578	0.113	0.157	12.250
Global Motion Hints	0.018	0.034	11.990	0.025	0.034	23.757	0.104	0.150	11.865
Dynamic Motion Hints	0.016	0.034	11.678	0.024	0.033	24.289	0.096	0.142	11.089

Table 5: Effectiveness of extracting motion hints between interpolated frames with neighboring frames instead of directly extracting motion hints from two continuous neighboring frames. For each column, we highlight the best result in red and the second best in blue.

both VQ-MAGAN and MA-SAMPLING. The results indicate that our MADiff can effectively integrate various types of motion hints, thereby enhancing the performance of video frame interpolation. This demonstrates that our MADiff constitutes a flexible framework, allowing for the straightforward replacement of the motion hint extractor to incorporate diverse motion hints into diffusion models for the VFI task. Moreover, we conduct two additional experiments to ascertain the superiority of motion hints between the interpolated frame and the neighboring frames (referred to as **Dynamic Motion Hints** in Table 5) over motion hints between two neighboring frames (referred to as **Global Motion Hints** in Table 5). The results presented in Table 5 demonstrate that motion hints extracted between the interpolated frame and the neighboring frames offer more precise guidance and help to mitigate motion ambiguity in the VFI task, compared to motion hints directly extracted from two consecutive neighboring frames.

6 CONCLUSION

In this paper, for the VFI task, we propose a novel motion-aware latent diffusion model (MADiff) which can fully leverage rich inter-frame motion priors from readily available and pre-trained motion-related models during the generation of interpolated frames. Specifically, our MADiff consist of a vector quantized motion-aware generative adversarial network (VQ-MAGAN) and a de-noising U-net. And VQ-MAGAN is adept at aggregating contextual details under the guidance of inter-frame motion hints between the interpolated and given neighboring frames. Additionally, we propose a novel motion-aware sampling procedure (MA-SAMPLING) that progressively refines the predicted frame throughout the sampling process of the diffusion model. Comprehensive experiments conducted on benchmark datasets demonstrate that our MADiff achieves the state-of-the-art performance, significantly surpassing existing approaches, particularly in scenarios characterized by dynamic textures and complex motions.

ACKNOWLEDGMENTS

This work was partly supported by the National Natural Science Foundation of China (Nos.62171251&62311530100) and the Special Foundations for the Development of Strategic Emerging Industries of Shenzhen (No.KJZD20231023094700001) and the Major Key Research Project of PCL (No.PCL2023A08).

REFERENCES

- [1] Simon Baker, Daniel Scharstein, JP Lewis, Stefan Roth, Michael J Black, and Richard Szeliski. 2011. A database and evaluation methodology for optical flow. *Int. J. of Comput. Vis.* 92, 1 (2011), 1–31.
- [2] Patrick Bardow, Andrew J Davison, and Stefan Leutenegger. 2016. Simultaneous optical flow and intensity estimation from an event camera. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 884–892.
- [3] Yin Bi and Yiannis Andreopoulos. 2017. PIX2NVIS: Parameterized conversion of pixel-domain video frames to neuromorphic vision streams. In *2017 IEEE International Conference on Image Processing (ICIP)*. IEEE, 1990–1994.
- [4] Christian Brandli, Raphael Berner, Minhao Yang, Shih-Chii Liu, and Tobi Delbrück. 2014. A $240 \times 180 \times 130$ db $3 \mu\text{s}$ latency global shutter spatiotemporal vision sensor. *IEEE Journal of Solid-State Circuits* 49, 10 (2014), 2333–2341.
- [5] Marco Cannici, Marco Ciccone, Andrea Romanoni, and Matteo Matteucci. 2020. A differentiable recurrent surface for asynchronous event-based data. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX* 16. Springer, 136–152.
- [6] Kenneth Chaney, Alex Zihao Zhu, and Kostas Daniilidis. 2019. Learning event-based height from plane and parallax. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 0–0.
- [7] Xianhang Cheng and Zhenzhong Chen. 2020. Video frame interpolation via deformable separable convolution. In *Proc. of the AAAI Conf. on Artificial Intell.* 10607–10614.
- [8] Xianhang Cheng and Zhenzhong Chen. 2022. Multiple video frame interpolation via enhanced deformable separable convolution. *IEEE Trans. on Pattern Anal. Mach. Intell.* 44, 10 (2022), 7029–7045.
- [9] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078* (2014).
- [10] Myungsub Choi, Heewon Kim, Bohyung Han, Ning Xu, and Kyoung Mu Lee. 2020. Channel attention is all you need for video frame interpolation. In *Proc. of the AAAI Conf. on Artificial Intell.*, Vol. 34. 10663–10671.
- [11] Duolikun Danier, Fan Zhang, and David Bull. 2022. Enhancing deformable convolution based video frame interpolation with coarse-to-fine 3D CNN. In *IEEE Int. Conf. on Image Process.* 1396–1400.
- [12] Duolikun Danier, Fan Zhang, and David Bull. 2022. FloLPIPS: A Bespoke Video Quality Metric for Frame Interpolation. In *IEEE Picture Coding Symposium*. 283–287.
- [13] Duolikun Danier, Fan Zhang, and David Bull. 2022. ST-MFNet: A spatio-temporal multi-flow network for frame interpolation. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 3521–3531.
- [14] Duolikun Danier, Fan Zhang, and David Bull. 2022. A subjective quality study for video frame interpolation. In *IEEE Int. Conf. on Image Process.* 1361–1365.
- [15] Duolikun Danier, Fan Zhang, and David Bull. 2024. Ldmvfi: Video frame interpolation with latent diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 1472–1480.
- [16] Tianyu Ding, Luming Liang, Zhihui Zhu, and Ilya Zharkov. 2021. CDFI: Compression-Driven Network Design for Frame Interpolation. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 8001–8011.
- [17] Patrick Esser, Robin Rombach, and Bjorn Ommer. 2021. Taming transformers for high-resolution image synthesis. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 12873–12883.
- [18] Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J Davison, Jörg Conradt, Kostas Daniilidis, et al. 2020. Event-based vision: A survey. *IEEE transactions on pattern analysis and machine intelligence* 44, 1 (2020), 154–180.
- [19] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. GANs trained by a two time-scale update rule converge to a local nash equilibrium. *Adv. in Neural Inform. Process. Syst.* 30 (2017).
- [20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Adv. in Neural Inform. Process. Syst.* 33 (2020), 6840–6851.
- [21] Jonathan Ho, Tim Salimans, Alexey A. Gritsenko, William Chan, Mohammad Norouzi, and David J. Fleet. 2022. Video Diffusion Models. In *Adv. in Neural Inform. Process. Syst.*, Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (Eds.). https://openreview.net/forum?id=f3zNgKga_ep
- [22] Zhilin Huang, Quanmin Liang, Yijie Yu, Chujun Qin, Xiwu Zheng, Kai Huang, Zikun Zhou, and Wenming Yang. 2024. Bilateral Event Mining and Complementary for Event Stream Super-Resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 34–43.
- [23] Zhilin Huang, Ling Yang, Zaixi Zhang, Xiangxin Zhou, Yu Bao, Xiwu Zheng, Yuwei Yang, Yu Wang, and Wenming Yang. 2024. Binding-Adaptive Diffusion Models for Structure-Based Drug Design. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 12671–12679.
- [24] Zhilin Huang, Ling Yang, Xiangxin Zhou, Zhilong Zhang, Wentao Zhang, Xiwu Zheng, Jie Chen, Yu Wang, CUI Bin, and Wenming Yang. 2024. Protein-ligand interaction prior for binding-aware 3d molecule diffusion models. In *The Twelfth International Conference on Learning Representations*.
- [25] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy. 2018. Liteflownet: A lightweight convolutional neural network for optical flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 8981–8989.
- [26] Aapo Hyvärinen. 2005. Estimation of Non-Normalized Statistical Models by Score Matching. *J. Mach. Learn. Res.* 6 (2005), 695–709.
- [27] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 1125–1134.
- [28] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 2012. 3D convolutional neural networks for human action recognition. *IEEE transactions on pattern analysis and machine intelligence* 35, 1 (2012), 221–231.
- [29] Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik Learned-Miller, and Jan Kautz. 2018. Super slo-mo: High quality estimation of multiple intermediate frames for video interpolation. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 9000–9008.
- [30] Jacques Kaiser, J Camilo Vasquez Tieck, Christian Hubschneider, Peter Wolf, Michael Weber, Michael Hoff, Alexander Friedrich, Konrad Wojtasik, Arne Roennau, Ralf Kohlhaas, et al. 2016. Towards a framework for end-to-end control of a simulated vehicle with spiking neural networks. In *2016 IEEE International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPACT)*. IEEE, 127–134.
- [31] Tarun Kalluri, Deepak Pathak, Manmohan Chandraker, and Du Tran. 2023. FLAVR: Flow-agnostic video representations for fast frame interpolation. In *Proc. of the IEEE Winter Conf. on Applications of Comput. Vis.* 2071–2082.
- [32] Xiaoyang Kang, Xianhui Lin, Kai Zhang, Zheng Hui, Wangmeng Xiang, Jun-Yan He, Xiaoming Li, Peiran Ren, Xuansong Xie, Radu Timofte, et al. 2023. NTIRE 2023 video colorization challenge. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 1570–1581.
- [33] Junho Kim, Jaehyeok Bae, Gangin Park, Dongsu Zhang, and Young Min Kim. 2021. N-imagenet: Towards robust, fine-grained object recognition with event cameras. In *Proceedings of the IEEE/CVF international conference on computer vision*. 2146–2156.
- [34] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Int. Conf. on Learn. Represent.*
- [35] Lingtong Kong, Boyuan Jiang, Donghao Luo, Wenqing Chu, Xiaoming Huang, Ying Tai, Chengjie Wang, and Jie Yang. 2022. IFRNet: Intermediate feature refine network for efficient frame interpolation. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 1969–1978.
- [36] Lingtong Kong, Chunhua Shen, and Jie Yang. 2021. Fastflownet: A lightweight network for fast optical flow estimation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 10310–10316.
- [37] Hyeonmin Lee, Taeh Kim, Tae-young Chung, Daehyun Pak, Yuseok Ban, and Sangyoun Lee. 2020. AdaCoF: Adaptive collaboration of flows for video frame interpolation. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 5316–5325.
- [38] Diance Li, Jianing Li, and Yonghong Tian. 2023. SODFormer: Streaming Object Detection with Transformer Using Events and Frames. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023).
- [39] Quanmin Liang, Zhilin Huang, Xiwu Zheng, Feidiao Yang, Jun Peng, Kai Huang, and Yonghong Tian. 2024. Efficient Event Stream Super-Resolution with Recursive Multi-Branch Fusion. *arXiv preprint arXiv:2406.19640* (2024).
- [40] Quanmin Liang, Xiwu Zheng, Kai Huang, Yan Zhang, Jie Chen, and Yonghong Tian. 2023. Event-Diffusion: Event-Based Image Reconstruction and Restoration with Diffusion Models. In *Proceedings of the 31st ACM International Conference on Multimedia*. 3837–3846.
- [41] Patrick Lichtsteiner, Christoph Posch, and Tobi Delbrück. 2008. A $128 \times 128 \times 120$ dB $15 \mu\text{s}$ Latency Asynchronous Temporal Contrast Vision Sensor. *IEEE J. Solid State Circuits* 43, 2 (2008), 566–576. <https://doi.org/10.1109/JSSC.2007.914337>
- [42] Ziwei Liu, Raymond A Yeh, Xiaoou Tang, Yiming Liu, and Aseem Agarwala. 2017. Video frame synthesis using deep voxel flow. In *Proc. of the IEEE Int. Conf. on Comput. Vis.* 4463–4471.
- [43] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *Int. Conf. on Learn. Represent.* <https://openreview.net/forum?id=Bkg6RiCqY7>
- [44] Liying Lu, Ruizheng Wu, Huaajia Lin, Jiangbo Lu, and Jiaya Jia. 2022. Video frame interpolation with transformer. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 3532–3542.

- [45] Di Ma, Fan Zhang, and David Bull. 2021. BVI-DVC: A Training Database for Deep Video Compression. *IEEE Trans. on Multimedia* (2021), 1–1. <https://doi.org/10.1109/TMM.2021.3108943>
- [46] Nico Messikommer, Carter Fang, Mathias Gehrig, and Davide Scaramuzza. 2023. Data-driven feature tracking for event cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5642–5651.
- [47] Simon Niklaus and Feng Liu. 2018. Context-aware synthesis for video frame interpolation. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 1701–1710.
- [48] Simon Niklaus and Feng Liu. 2020. Softmax splatting for video frame interpolation. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 5437–5446.
- [49] Simon Niklaus, Long Mai, and Feng Liu. 2017. Video frame interpolation via adaptive convolution. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 670–679.
- [50] Junheum Park, Keunsoo Ko, Chul Lee, and Chang-Su Kim. 2020. BMBC: Bilateral motion estimation with bilateral cost volume for video interpolation. In *Comput. Vis. –ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV* 16. Springer, 109–125.
- [51] Junheum Park, Chul Lee, and Chang-Su Kim. 2021. Asymmetric Bilateral Motion Estimation for Video Frame Interpolation. In *Proc. of IEEE Int. Conf. on Comput. Vis.* 14539–14548.
- [52] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. 2016. A benchmark dataset and evaluation methodology for video object segmentation. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 724–732.
- [53] Henri Rebecq, René Ranfl, Vladlen Koltun, and Davide Scaramuzza. 2019. Events-to-video: Bringing modern computer vision to event cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3857–3866.
- [54] Henri Rebecq, René Ranfl, Vladlen Koltun, and Davide Scaramuzza. 2019. High speed and high dynamic range video with an event camera. *IEEE transactions on pattern analysis and machine intelligence* 43, 6 (2019), 1964–1980.
- [55] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 10684–10695.
- [56] Hyeonjun Sim, Jihyong Oh, and Munchurl Kim. 2021. XVFI: eXtreme Video Frame Interpolation. In *Proc. of the IEEE Int. Conf. on Comput. Vis.* 14489–14498.
- [57] Li Siyao, Shiyu Zhao, Weijiang Yu, Wenxiu Sun, Dimitris Metaxas, Chen Change Loy, and Ziwei Liu. 2021. Deep animation video interpolation in the wild. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 6587–6595.
- [58] Jiaming Song, Chenlin Meng, and Stefano Ermon. 2021. Denoising Diffusion Implicit Models. In *Int. Conf. on Learn. Represent.*
- [59] Yang Song and Stefano Ermon. 2019. Generative modeling by estimating gradients of the data distribution, Vol. 32.
- [60] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2020. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.
- [61] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *Int. Conf. on Learn. Represent.* <https://openreview.net/forum?id=PXTIG12RRHS>
- [62] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402* (2012).
- [63] Gemma Taverni, Diederik Paul Moeys, Chenghan Li, Celso Cavaco, Vasy1Motsnyi, David San Segundo Bello, and Tobi Delbruck. 2018. Front and back illuminated dynamic and active pixel vision sensors comparison. *IEEE Transactions on Circuits and Systems II: Express Briefs* 65, 5 (2018), 677–681.
- [64] Stepan Tulyakov, Alfredo Bochicchio, Daniel Gehrig, Stamatios Georgoulis, Yuanyou Li, and Davide Scaramuzza. 2022. 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*. 17755–17764.
- [65] Aaron Van Den Oord, Oriol Vinyals, et al. 2017. Neural discrete representation learning. *Adv. in Neural Inform. Process. Syst.* 30 (2017).
- [66] Pascal Vincent. 2011. A connection between score matching and denoising autoencoders. *Neural computation* 23, 7 (2011), 1661–1674.
- [67] Vikram Voleti, Alexia Jolicœur-Martineau, and Christopher Pal. 2022. MCVD - Masked Conditional Video Diffusion for Prediction, Generation, and Interpolation. In *Adv. in Neural Inform. Process. Syst.*, Vol. 35. 23371–23385.
- [68] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. on Image Process.* 13, 4 (2004), 600–612.
- [69] Chao-Yuan Wu, Nayan Singhal, and Philipp Krahenbuhl. 2018. Video compression through image interpolation. In *Proc. of the European Conf. on Comput. Vis.* 416–431.
- [70] Lan Xu, Weipeng Xu, Vladislav Golyanik, Marc Habermann, Lu Fang, and Christian Theobalt. 2020. Eventcap: Monocular 3d capture of high-speed human motions using an event camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4968–4978.
- [71] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. 2019. Video enhancement with task-oriented flow. *Int. J. of Comput. Vis.* 127, 8 (2019), 1106–1125.
- [72] Ling Yang, Zhilin Huang, Yang Song, Shenda Hong, Guohao Li, Wentao Zhang, Bin Cui, Bernard Ghanem, and Ming-Hsuan Yang. 2022. Diffusion-based scene graph to image generation with masked contrastive pre-training. *arXiv preprint arXiv:2211.11138* (2022).
- [73] Ren Yang. 2021. NTIRE 2021 challenge on quality enhancement of compressed video: Dataset and study. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 667–676.
- [74] Ren Yang, Radu Timofte, and Luc Van Gool. 2022. Perceptual Learned Video Compression with Recurrent Conditional GAN. In *Processings of the Int. Joint Conf. on Artificial Intell.* 1537–1544.
- [75] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recog.* 586–595.
- [76] Zhongyang Zhang, Shuyang Cui, Kaidong Chai, Haowen Yu, Subhasis Dasgupta, Upal Mahbub, and Tauhidur Rahman. 2023. V2CE: Video to Continuous Events Simulator. *arXiv preprint arXiv:2309.08891* (2023).
- [77] Yi Zhou, Guillermo Gallego, Henri Rebecq, Laurent Kneip, Hongdong Li, and Davide Scaramuzza. 2018. Semi-dense 3D reconstruction with a stereo event camera. In *Proceedings of the European conference on computer vision (ECCV)*. 235–251.
- [78] Alex Zihao Zhu, Ziyun Wang, Kaung Khant, and Kostas Daniilidis. 2021. Eventgan: Leveraging large scale image datasets for event cameras. In *2021 IEEE international conference on computational photography (ICCP)*. IEEE, 1–11.
- [79] Alex Zihao Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. 2019. Unsupervised event-based learning of optical flow, depth, and egomotion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 989–997.