

TLB-VFI: Temporal-Aware Latent Brownian Bridge Diffusion for Video Frame Interpolation

Zonglin Lyu Chen Chen

Center for Research in Computer Vision, University of Central Florida

zonglin.lyu@ucf.edu chen.chen@crcv.ucf.edu

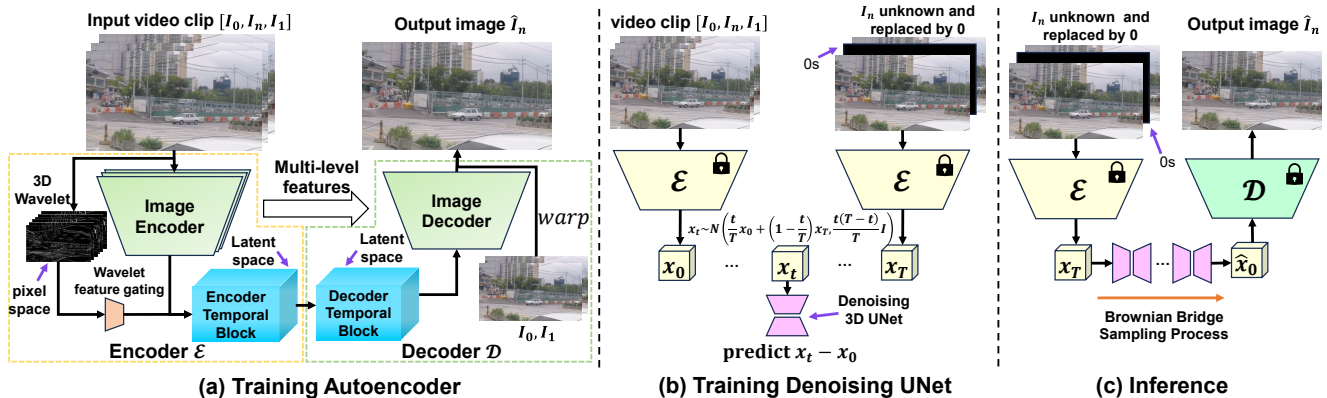


Figure 1. **Overview of the proposed method.** (a) **Training autoencoder.** The autoencoder is trained with video clip $V = [I_0, I_n, I_1]$ and aims to reconstruct I_n . It contains an image encoder (shared for all frames) and an image decoder, where multi-level encoder features from I_0, I_1 are passed to the decoder. Temporal blocks extract temporal information in the latent space and aggregate video features into a single image feature for the image decoder, and 3D Wavelet extracts temporal information in the pixel space. (b) **Training Denoising UNet.** The video clip V is encoded to \mathbf{x}_0 by Encoder \mathcal{E} (spatial + temporal). Since I_n is unknown, we replace it by 0 and obtain another video clip $\tilde{V} = [I_0, 0, I_1]$, which is encoded to \mathbf{x}_T . With the Brownian Bridge Diffusion Process, \mathbf{x}_t is computed and sent to denoising UNet to predict $\mathbf{x}_t - \mathbf{x}_0$. (c) **Inference.** During inference, we encode \tilde{V} to \mathbf{x}_T and sample with the Brownian Bridge Sampling Process to get $\hat{\mathbf{x}}_0$, which is decoded to the output frame \hat{I}_n .

Abstract

Video Frame Interpolation (VFI) aims to predict the intermediate frame I_n (we use n to denote time in videos to avoid notation overload with the timestep t in diffusion models) based on two consecutive neighboring frames I_0 and I_1 . Recent approaches apply diffusion models (both image-based and video-based) in this task and achieve strong performance. However, image-based diffusion models are unable to extract temporal information and are relatively inefficient compared to non-diffusion methods. Video-based diffusion models can extract temporal information, but they are too large in terms of training scale, model size, and inference time. To mitigate the above issues, we propose Temporal-Aware Latent Brownian Bridge Diffusion for Video Frame Interpolation (TLB-VFI), an efficient video-based diffusion model. By extracting rich temporal information from video inputs through our proposed 3D-wavelet gating and temporal-aware autoencoder, our method achieves 20% improvement in FID on the most challenging datasets over recent SOTA of image-based diffusion

models. Meanwhile, due to the existence of rich temporal information, our method achieves strong performance while having $3\times$ fewer parameters. Such a parameter reduction results in $2.3\times$ speed up. By incorporating optical flow guidance, our method requires $9000\times$ less training data and achieves over $20\times$ fewer parameters than video-based diffusion models. Codes and results are available at our [Project Page](#).

1. Introduction

Video Frame Interpolation (VFI) is a crucial task in computer vision that aims to predict the intermediate frame I_n between two known neighboring frames (past frame I_0 and future frame I_1). Good quality interpolation results benefit a wide range of applications, including novel-view synthesis [12] and video compression [44].

Recent VFI methods generally fall into two categories: diffusion-based and traditional (a.k.a non-diffusion-based) methods. Traditional methods use either kernel-based methods to generate convolution kernels that predict in-

intermediate frames from pixels of neighboring images [5, 20, 30, 31, 38] or flow-based methods that estimate optical flows [1, 6, 11, 16, 22, 25, 32, 34, 45, 49, 50]. Advances in deep learning make researchers prefer flow-based methods over kernel-based ones because the motion estimation in flow-based methods benefit from deep architecture, whereas kernel-based methods do not model motion explicitly. Diffusion-based methods can be split into image-based diffusion models [9, 26] and video-based diffusion models [17, 37, 42, 48], both achieving strong performances. However, they encounter several limitations:

- Traditional methods and image-based diffusion methods (generally less efficient than traditional methods) [9, 26], which incorporate kernel- or flow-based methods to explicitly guide frame interpolation, only extract spatial information in I_0, I_1 . Therefore, they lack explicit temporal information extraction to enhance temporal consistency.
- Video-based diffusion models [17, 37, 42, 48] directly generate pixels. Though they successfully extract temporal information by taking video clips as inputs, the training and inference costs of these methods are extremely high. They require more than ten million videos to train for a strong performance due to the lack of explicit pixel-level information from I_0, I_1 guided by optical flow, and the inference time is also extremely long.

Our goal is to extract temporal information while keeping training scale, model size, and inference time reasonable. To achieve so, we notice that a video-based diffusion model is necessary. Video-based diffusion models can gradually build up temporal information in their sampling process with 3D UNet, and the autoencoder can further extract temporal information. This is not feasible in traditional methods as I_n is not available as input. Therefore, we propose Temporal-Aware Latent Brownian Bridge Diffusion for Video Frame Interpolation (TLB-VFI), shown in Fig. 1. Starting with the selection of diffusion models, Consec. BB [26] points out that Brownian Bridge Diffusion Model (BBDM) [21] better fits the VFI task than traditional Diffusion Models [14] since the sampling variance of BBDM is lower. This is because VFI expects deterministic interpolation results rather than diverse ones. However, in the setup of Consec. BB [26], Brownian Bridge is applied between adjacent frames, but the latent features are almost identical for adjacent frames. Therefore, the Brownian Bridge becomes approximately an identity mapping and loses its functionality. To mitigate this problem, given input video clip $V = [I_0, I_n, I_1]$, we construct $\tilde{V} = [I_0, 0, I_1]$ and apply Brownian Bridge Diffusion between their latent features, shown in Fig. 1 (b). The latent features of \tilde{V} largely differ from that of V , resolving the issue of identity mapping. Details and justifications are in Sec. 4.3 and Fig. 6.

Next, it is important to construct an autoencoder so that diffusion models can efficiently run in latent space. To

overcome the limitation on lack of temporal information, a simple approach is to use 3D Convolution and Spatiotemporal attention. However, since I_n is replaced with 0 during inference (see Fig. 1 (c)), the multi-level encoder features, which benefits the decoding stage [9], in Fig. 1 (a) will contain incomplete information due to zero replacement, which harms the performance (see Sec. 4.3). We mitigate this issue by splitting the temporal information extraction into pixel space and latent space as shown in Fig. 1 (a). We include image-level encoder and decoder for spatial features only and add temporal feature extraction between the image-level encoder and decoder to extract temporal information in the latent space. This enables us to utilize the multi-level encoder features of I_0, I_1 during the decoding phase as *they are not impacted by zero replacement.* To extract temporal information in the pixel space, we propose our 3D-wavelet feature gating mechanism to extract high-frequency temporal information, since high-frequency information represents the changes along temporal dimension. **To address the limitation on efficiency**, instead of directly predicting pixel values, we utilize optical flow estimation to warp neighboring frames and refine the output. Under the guidance of optical flow at the pixel level in I_0 and I_1 , our method achieves better efficiency than directly generating pixels from scratch. Importantly, due to rich temporal features extracted, our method achieves the state-of-the-art performance while requiring less parameters and inference time than image-based diffusion models [9, 26]. Our contributions are summarized as:

- Our method requires **9000** \times fewer training data, has over **20** \times fewer the number of parameters, and is over **10** \times faster than video-based diffusion models [17, 37, 48]. Comparing to Image-based diffusion methods, our method has **3** \times fewer parameters and over **2** \times faster.
- We introduce our temporal design of autoencoder to extract temporal information in the latent space and 3D wavelet feature gating to extract temporal information in the pixel space, serving as complement with each other.
- We propose a theoretical constraint on *when Brownian Bridge Diffusion is effective* in Sec. 3.3. Our temporal design of autoencoder adheres to our proposed constraint, validating the effectiveness of the Brownian Bridge.
- Through extensive experiments, our method achieves state-of-the-art performance in various datasets. Specifically, our method achieves around **20%** improvement in FID in the most challenging datasets: SNU-FILM extreme subset [7] and Xiph-4K [29] over recent SOTAs.

2. Related Work

2.1. Traditional Methods in VFI

Video Frame Interpolation (VFI) is a task to predict the intermediate frame I_n given its neighboring frames I_0 and I_1 .

Traditional (non-diffusion) VFI methods fall into two categories: flow-based methods [1, 6, 11, 16, 18, 22, 25, 32, 34, 49] and kernel-based methods [5, 20, 30, 31, 38]. Flow-based methods utilize optical flow estimations via deep neural networks. Some estimate optical flows from the I_n to I_0 and I_1 and apply backward warping [1, 6, 11, 16, 22, 25, 32, 34, 49]. Others estimate bidirectional flows from I_0, I_1 to each other, and apply forward splatting [18]. In addition to the basic framework, various architectures like transformers [41] are introduced together with techniques such as multi-resolution recurrence refinement [18], 4D-correlations extraction [22], and asymmetric flow blending [45] to enhance interpolation quality. Kernel-based methods, initially proposed by [30], try to estimate local convolution kernels that are applied to pixels of neighboring frames to generate pixels in the intermediate frame. Several improvements, such as adaptive and deformable convolutions [5, 20], are proposed to improve performance under large motion change.

2.2. Diffusion Models in VFI

Denosing Diffusion Probabilistic Models (DDPM) [14] are proposed to generate realistic and diverse images. In DDPM, an image is transformed into standard Gaussian noise with a predefined diffusion process, and the noise is transformed back to an image with a sampling (denoising) process. The sampling process contains T steps of iteration, where T is set to 1000 experimentally. The following works [23, 24, 39] improve the efficiency of sampling by reducing the number of steps for iterative sampling while keeping high-quality generation. Latent Diffusion Models [35] introduce autoencoders with vector-quantized or KL-divergence regularization to compress images into latent space, where diffusion models operate more efficiently. Beyond image generation, BBDM [21] proposes Brownian Bridge Diffusion for image-to-image translation, and video diffusion models [3, 15] are developed to generate high-quality videos.

LDMVFI [9] incorporates the idea of Latent Diffusion Models (LDMs) [35] and kernel-based methods, proposing an autoencoder that utilizes kernel-based methods to reconstruct I_n . Recent work [26] claims that VFI prefers deterministic prediction of the intermediate frames and introduces the Consecutive Brownian Bridge, which has a small sampling variance and achieves state-of-the-art performance. Other diffusion-based works [17, 37, 42, 48] do not employ kernel- or flow-based methods but instead generate raw pixels. This approach allows these diffusion methods to extract explicit temporal information by taking videos as inputs to autoencoder and build up temporal information during the sampling process from a random noise. However, these methods require large-scale training as they lack pixel guidance from neighboring frames

I_0, I_1 : most works [17, 37, 48] require over ten million videos to train, while the common datasets for VFI only contain 51K triplets [47]. MCVD [42], a diffusion approach directly generating pixels, is trained with a small-scale dataset, resulting in unsatisfactory performance in VFI [9]. Our method addresses these challenges by extracting temporal information from video inputs and taking advantage of optical flow estimation.

3. Methodology

3.1. Preliminary

Diffusion Models. DDPM [14] defines a diffusion process that converts images into standard Gaussian noise:

$$q(\mathbf{x}_t|x_0) = \mathcal{N}(x_t; \sqrt{\alpha_t}\mathbf{x}_0, (1 - \alpha_t)\mathbf{I}), \quad (1)$$

$$\text{where } \alpha_t = \prod_{s=1}^t (1 - \beta_s).$$

$\{\beta_t\}$ are small pre-defined constants. The sampling process (i.e. how to denoise \mathbf{x}_{t-1} from \mathbf{x}_t) is given by [14]:

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(x_{t-1}; \tilde{\mu}_t, \tilde{\beta}_t), \quad (2)$$

$$\text{where } \tilde{\mu}_t = \frac{1}{1 - \beta_t} \left(\mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}} \epsilon \right), \quad (3)$$

$$\tilde{\beta}_t = \frac{1 - \alpha_{t-1}}{1 - \alpha_t} \beta_t, \text{ and } \epsilon \sim \mathcal{N}(0, \mathbf{I}). \quad (4)$$

Based on Eq. (3), the only unknown term is ϵ , which is estimated with a deep neural network $\epsilon_\theta(\mathbf{x}_t, t)$.

Brownian Bridge Diffusion Models. Similar to DDPM, Brownian Bridge Diffusion Models (BBDM) [21] replace its diffusion and sampling process with Brownian Bridge Diffusion [36] for image-to-image translation tasks such as image inpainting and image colorization. Consecutive Brownian Bridge [26] introduces a cleaner formulation of BBDM and suggests that it fits the VFI task since the sampling process has low variance, which VFI prefers. The diffusion process is:

$$q(\mathbf{x}_t|\mathbf{x}_0, \mathbf{x}_T) = \mathcal{N} \left(\frac{t}{T} \mathbf{x}_0 + (1 - \frac{t}{T}) \mathbf{x}_T, \frac{t(T-t)}{T} \mathbf{I} \right). \quad (5)$$

\mathbf{x}_0 and \mathbf{x}_T are the latent features of two images to be translated. T controls the maximum variance of the Brownian Bridge. The sampling process is given as [26]:

$$p_\theta(\mathbf{x}_s|\mathbf{x}_t, \mathbf{x}_T) = \mathcal{N} \left(\mathbf{x}_t - \frac{\Delta_t}{t} (\mathbf{x}_t - \mathbf{x}_0), \frac{s\Delta_t}{t} \mathbf{I} \right). \quad (6)$$

Here s is an arbitrary time before t , and $\Delta_t = t - s$. Analogous to DDPM [14], the only unknown term is $\mathbf{x}_t - \mathbf{x}_0$, which can be estimated with a deep neural network.

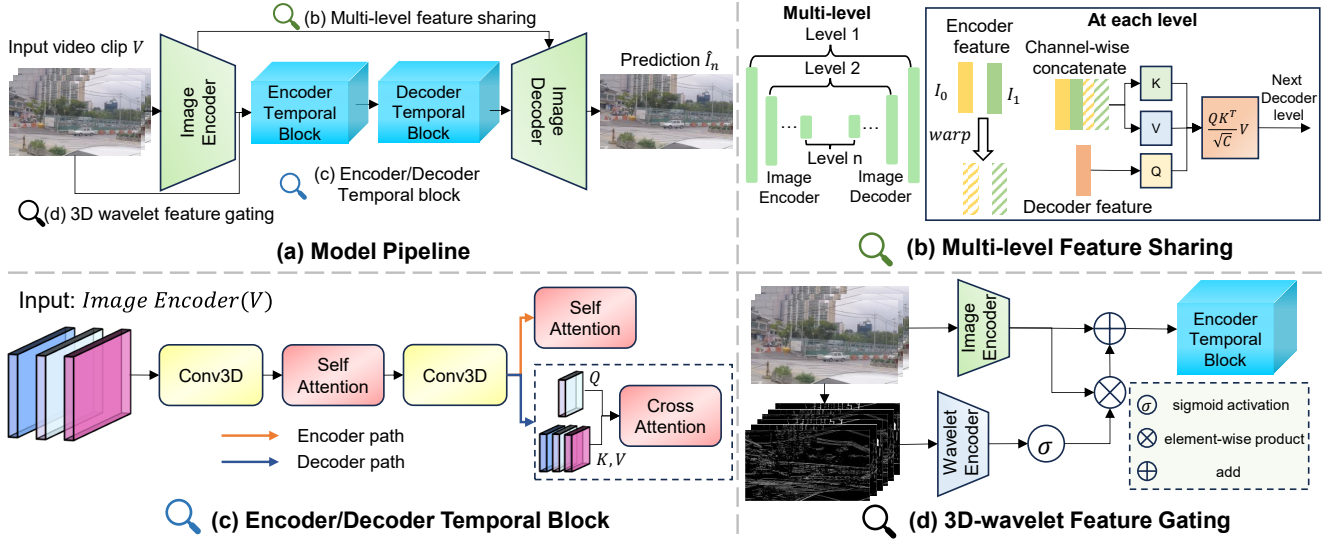


Figure 2. **(a) Model Pipeline.** The Image Encoder is shared across all frames, and temporal blocks extract temporal information in the latent space. **(b) Multi-level Feature Sharing.** The Image Encoder and Decoder consist of several levels of resolution due to downsampling/upsampling latent features. At the i^{th} level of the encoder and the decoder, features from I_0 and I_1 in the encoder are warped and concatenated with the original copy (when the downsampling rate is larger than 8, warped features are excluded). The concatenated features are used as keys and values in cross attention [41] where the decoder feature at the same level is the query. **(c) Encoder/Decoder Temporal Block.** Each temporal block consists of two sets of 3D convolution + attention. In the decoder, the second attention is cross-attention between the intermediate frame (query) and all frames (key and value) to aggregate the video feature into one feature map. **(d) 3D-wavelet Feature Gating.** Wavelet information is extracted from the input video clip and encoded by CNNs. A sigmoid activation is applied, and the result is element-wise multiplied by the output of the Image Encoder with a skip connection.

3.2. Problem Definition

Video Frame Interpolation aims to predict $I_n \in \mathbb{R}^{3 \times H \times W}$, which is the intermediate frames between I_0 and I_1 . Following the common strategy of recent VFI works [26, 45], the goal is to predict \hat{I}_n that predicts I_n :

$$\hat{I}_n = M \odot \text{warp}(I_0) + (1 - M) \odot \text{warp}(I_1) + \Delta \quad (7)$$

M is a mask with values between 0 and 1. warp denotes warping based on estimated optical flows. Δ is the residual term with values between -1 and 1. Therefore, our model will predict the mask M and residual Δ , with flow estimation embedded in the model. The basics of optical flow and warping are included in the [Supplementary Material](#).

Method Pipeline. Our method is divided into two main components: an autoencoder and a Brownian Bridge Diffusion Model (referred to as ‘diffusion model’ later), shown in Fig. 1 (a) and (b), respectively. The autoencoder contains an encoder \mathcal{E} and a decoder \mathcal{D} , shown in yellow and green dashes in Fig. 1 (a) respectively. The encoder \mathcal{E} contains an image encoder and temporal blocks, and the decoder contains temporal blocks and an image decoder. The reason for this design is discussed in Sec. 3.3. The objective of the autoencoder is to predict M, Δ given input video clip $V = [I_0, I_n, I_1]$: $M, \Delta = \mathcal{D}(\mathcal{E}(V))$. During inference, I_n is replaced with a zero matrix, resulting in $\tilde{V} = [I_0, 0, I_1]$. The goal of the diffusion model (denoted as BB) is to estimate $\mathcal{E}(V)$ from $\mathcal{E}(\tilde{V})$. In summary, our pipeline to predict I_n , with $n = 0.5$, involves the following steps:

1. $x = BB(\mathcal{E}(\tilde{V}))$.
2. $M, \Delta = \mathcal{D}(x)$.
3. $\hat{I}_n = M \odot \text{warp}(I_0) + (1 - M) \odot \text{warp}(I_1) + \Delta$

3.3. Temporal Aware Latent Brownian Bridge

Temporal Feature Extraction in Latent Space. Our goal is to extract temporal information in the autoencoder, but it is important to provide multi-level encoder features of I_0, I_1 to guide the decoder [9, 26]. Details of how these multi-level features can provide guidance are shown in Fig. 2 (b). At i^{th} level of the encoder and decoder, the features are downsampled to $2^{-i} \times$ of the image’s height and width. We denote encoder features of I_0, I_1 at i^{th} level as f_0, f_1 . Then $\text{concat}(f_0, f_1, \text{warp}(f_0), f_1, \text{warp}(f_1))$ is sent to the decoder at the i^{th} level via cross-attention [41] (concat is channel-wise concatenation). To ensure multi-level encoder features of I_0, I_1 are not affected by changing V to \tilde{V} during inference, we use shared image encoders to encode each frame into low-resolution latent features and apply temporal feature extraction of those latent features. The temporal blocks for the encoder, shown in Fig. 2 (c), utilize 3D convolution and spatiotemporal attention.

The decoder is mostly symmetric with the encoder, beginning with 3D convolution and spatiotemporal attention followed by an image decoder. To convert the video latent representation into the image latent representation, we include a spatiotemporal cross-attention aggregation in the decoder, illustrated in Fig. 2 (c). Given the video latent representation $F \in \mathbb{R}^{C \times T \times H \times W}$ where C, T, H, W rep-

resents the channel, temporal, height, and weight. The spatiotemporal cross-attention aggregation is defined as:

$$V_{out} = softmax\left(\frac{QK^T}{\sqrt{C}}\right)V$$

$$Q = W_Q F_Q, K = W_K F_{KV}, V = W_V F_{KV}$$

$$F_Q = F[t].flatten(1), F_{KV} = F.flatten(1)$$

F_Q and F_{KV} are in PyTorch style notation. This approach allows us to remove the temporal dimension, retaining only the image-level feature.

Temporal Feature Extraction in Pixel Space with 3D Wavelet Transform. Since we use a shared image encoder to encode images, the encoder and decoder only extract temporal information in latent space. However, extracting temporal information at the pixel level is essential as the latent space is highly compressed. Therefore, we apply 3D wavelet transform [27] to extract spatial and temporal frequency information with low-pass filter $[\frac{1}{\sqrt{2}} \frac{1}{\sqrt{2}}]$ and high-pass filter $[-\frac{1}{\sqrt{2}} \frac{1}{\sqrt{2}}]$. These filters are applied on height, width, and temporal dimensions, respectively. The reason to choose 3D wavelet transform is that it can apply different combinations of high pass and low pass filters across different dimensions, resulting in rich pixel-level information. Then, the frequency information is encoded with convolution layers. The 3D wavelet transform can be applied twice. With the first pass, we can extract temporal information between I_0 and I_n and between I_n and I_1 , which is not achievable with input only containing I_0 and I_1 . The second pass extracts the temporal information throughout all frames. Such pixel-level temporal information tells the model *where the motion changes more drastically*, and with this intuition, we design our 3D wavelet feature gating, shown in Fig. 2 (d). If f_w is the encoded feature of frequency information, and if f_i is the latent features encoded by the shared image encoder, then the 3D wavelet gating is:

$$f = \sigma(f_w) \odot f_i + f_i, \quad (8)$$

where σ is the sigmoid activation, and \odot is element-wise product. This implicitly guides the model to learn which parts of the video clip contain more changes in the temporal dimension. The details about the wavelet transform are included in the [Supplementary Material](#).

Temporal Information Restoration with Brownian Bridge Diffusion. During inference, V (original video clip) is replaced with \hat{V} (I_n replaced by zero matrix), leading to a distribution shift of encoded features $\mathcal{E}(\hat{V})$ from $\mathcal{E}(V)$ due to loss of temporal information. We employ the Brownian Bridge Diffusion to align the distribution of $\mathcal{E}(v)$ and $\mathcal{E}(\hat{V})$ to restore such information because Brownian Bridge diffusion demonstrates efficacy in VFI [26] by achieving low sampling variance. In addition, restoring the temporal

information is conceptually similar to spatial information restoration in image inpainting in BBDM [21], where the Brownian Bridge Diffusion is originally proposed. To effectively restore temporal information, temporal features need to be extracted. Therefore, we replace the convolution in the denoising U-Net with 3D convolution, and self-attentions are performed in spatial and temporal dimensions.

Notably, by Eq. (6), the denoising UNet ϵ_θ aims to predict $\mathbf{x}_t - \mathbf{x}_0$. When $t = T$, the ϵ_θ aims to predict $\mathbf{x}_T - \mathbf{x}_0$. If $\mathbf{x}_T = \mathbf{x}_0$, then by Eq. (5), $\mathbb{E}(\mathbf{x}_t) = 0$ at any time step. The training objective of ϵ_θ becomes 0 on expectation based on Eq. (6). The variance is negligible because $\Delta_t = \frac{1}{1000}$ and $T = 2$ during training [26]. In this case, when ϵ_θ learns to predict 0, we can easily compute that the sampling process defined by Eq. (6) is an identity mapping with a small variance. Therefore, we need a significantly large distribution shift between \mathbf{x}_0 and \mathbf{x}_T to prevent such a problem of identity mapping. Specifically, a big shift in the mean is required, which is described in the following proposition:

Proposition 1. *If the Brownian Bridge Diffusion is applied to translate between two distributions \mathbf{x}_0 and \mathbf{x}_T , there should be a large shift between $\mathbb{E}(\mathbf{x}_0)$ and $\mathbb{E}(\mathbf{x}_T)$. A sufficient constraint is to reject the Null Hypothesis $H_0 : \mathbb{E}(\mathbf{x}_0 - \mathbf{x}_T) = 0$ at significance level α .*

The proof of this proposition is given in the [Supplementary Material](#). We include experiments in Sec. 4.3 to show that our method satisfies this constraint, while the Consec. BB [26] does not.

4. Experiments

4.1. Datasets, Evaluation Metrics, and Baselines

Datasets. Following recent works [26, 49], our autoencoder and diffusion models are trained on the Vimeo 90K triplets dataset [47], which comprises 51,312 training triplets. We employ random flipping, cropping, rotation, and temporal order reversing as data augmentation. To evaluate our methods, we include Xiph [29], DAVIS [33], and SNU-FILM [7]. SNU-FILM contains four subsets based on the magnitude of motion changes: easy, medium, hard, and extreme, and Xiph contains 4K and 2K resolution for evaluation. These datasets contain various resolutions (480p to 4K) and motion changes. Among these datasets, Xiph and SNU-FILM-extreme are the most challenging and important due to high resolution and extremely large motion changes respectively, and strong performance on these two datasets indicates strong real-world applicability.

Evaluation Metric and Baselines. Recent works [9, 26, 45] identify that pixel-based metrics such as PSNR and SSIM [43] are less consistent with visual quality and humans' evaluation than learning-based metrics such as FID [13], LPIPS [51], and FloLPIPS [8]. There are exam-

Table 1. Quantitative results in LPIPS \downarrow /FloLPIPS \downarrow /FID \downarrow , the lower the better, on evaluation datasets. The best performances are **blue-boldfaced**, and the second-best performances are **red-underlined**. Results from PerVFI are in **gray background** and excluded from the ranking because their training scale is significantly larger. Results for baselines, except for PerVFI, are adopted from Consec.BB [26]. OOM indicates that the inference with one image exceeds the 24GB GPU memory of an Nvidia RTX A5000 GPU. Runtime measures seconds per frame to interpolate a 480×720 image with A5000 GPU. To ensure a fair comparison, we report the runtime of Consec. BB [26] and LDMVFI [9] in 10 step sampling (our setup) and also report the runtime of their setup in **orange** and parenthesis. Note that results from [17, 37, 48] are not included as their training scales are millions of videos.

Methods	Xiph-4K		Xiph-2K		DAVIS		SNU-FILM				Runtime Seconds					
	LPIPS/FloLPIPS/FID		LPIPS/FloLPIPS/FID		LPIPS/FloLPIPS/FID		easy		medium			hard		extreme		
	LPIPS	FloLPIPS	FID	LPIPS	FloLPIPS	FID	LPIPS	FloLPIPS	FID	LPIPS		FloLPIPS	FID	LPIPS	FloLPIPS	FID
MCVD'22 [42]	OOM		OOM		0.247/0.293/28.002		0.199/0.230/32.246		0.213/0.243/37.474		0.250/0.292/51.529		0.320/0.385/83.156		52.55	
VFIformer [25]	OOM		OOM		0.127/0.184/14.407		0.018/0.029/5.918		0.033/0.053/11.271		0.061/0.100/22.775		0.119/0.185/40.586		4.34	
IFRNet'22 [19]	0.136/0.164/23.647	0.068/0.093/11.465	0.106/0.156/12.422	0.106/0.145/13.018	0.022/0.034/6.139	0.035/0.055/11.039	0.060/0.092/20.810	0.112/0.177/40.075	0.110/0.132/12.167	0.132/0.166/15.186	0.019/0.038/5.882	0.033/0.053/11.051	0.060/0.091/20.679	0.114/0.170/39.051	0.72	
AMT'23 [22]	0.199/0.230/29.183	0.089/0.126/13.100	0.109/0.145/13.018	0.022/0.034/6.139	0.035/0.055/11.039	0.060/0.092/20.810	0.112/0.177/40.075	0.110/0.132/12.167	0.132/0.166/15.186	0.019/0.038/5.882	0.033/0.053/11.051	0.060/0.091/20.679	0.114/0.170/39.051	0.72		
UPR-Net'23 [18]	0.230/0.269/31.043	0.103/0.144/12.909	0.134/0.172/15.002	0.018/0.029/5.669	0.034/0.052/10.983	0.062/0.097/22.127	0.112/0.176/40.098	0.110/0.132/12.167	0.132/0.166/15.186	0.019/0.038/5.882	0.033/0.053/11.051	0.060/0.091/20.679	0.114/0.170/39.051	0.72		
EMA-VFI'23 [49]	0.241/0.260/28.695	0.110/0.132/12.167	0.107/0.153/12.554	0.014/0.024/5.752	0.028/0.053/12.485	0.060/0.114/26.520	0.123/0.204/47.042	0.110/0.132/12.167	0.132/0.166/15.186	0.014/0.024/5.752	0.028/0.053/12.485	0.060/0.114/26.520	0.123/0.204/47.042	2.48 (22.32)		
LDMVFI'24 [9]	OOM		OOM		0.107/0.153/12.554		0.014/0.024/5.752		0.028/0.053/12.485		0.060/0.114/26.520		0.123/0.204/47.042		2.48 (22.32)	
PerVFI'24 [45]	0.086/0.128/18.852	0.038/0.069/10.078	0.081/0.122/8.217	0.014/0.022/5.917	0.024/0.040/10.395	0.046/0.077/18.887	0.090/0.151/32.372	1.52	0.097/0.135/24.424	0.042/0.080/12.011	0.092/0.136/9.220	0.012/0.019/4.791	0.022/0.039/9.039	0.047/0.091/18.589	0.104/0.184/36.631	1.62 (2.60)
Consec. BB'24 [26]	0.097/0.135/24.424	0.042/0.080/12.011	0.092/0.136/9.220	0.012/0.019/4.791	0.022/0.039/9.039	0.047/0.091/18.589	0.104/0.184/36.631	1.62 (2.60)	0.077/0.113/19.114	0.032/0.067/9.901	0.086/0.126/8.299	0.012/0.018/4.658	0.021/0.036/8.518	0.044/0.085/17.470	0.095/0.151/29.868	0.69
Ours	0.077/0.113/19.114	0.032/0.067/9.901	0.086/0.126/8.299	0.012/0.018/4.658	0.021/0.036/8.518	0.044/0.085/17.470	0.095/0.151/29.868	0.69								

ples shown in [26] indicating that PSNR/SSIM are inconsistent with visual qualities in VFI task, and we also include examples in Sec. 4.2. Thus, FID, LPIPS, and FloLPIPS are selected as the evaluation metrics. FID and LPIPS measure perceptual similarity by Fréchet distance or normalized distances between features extracted by deep learning models. FloLPIPS, developed on top of LPIPS, accounts for motion consistency. Recent state-of-the-art methods including PerVFI (2024) [45], Consec. BB (2024) [26], LDMVFI (2024) [9], EMA-VFI (2023) [49], AMT (2023) [22], UPR-Net (2023) [18], IFRNet (2022) [19], VFIformer (2022) [25], and MCVD (2022) [42] are selected as our baselines. For models with different versions on the number of parameters, the largest version is chosen. For completeness of our experiment, we include PSNR/SSIM and implementation details in our [Supplementary Material](#).

4.2. Experimental Results

Quantitative Evaluation. The quantitative results are shown in Tab. 1, with 10 sampling steps in diffusion. Notably, PerVFI uses a combination of flow estimators, RAFT [40] and GMflow [46], to guide their model. These flow estimators are trained with three datasets: FlyingThings3D [28], FlyingChairs [10], and Sintel [4], where Sintel and Flyingthings3D contain high-resolution images. They additionally train their model on the Vimeo 90K Triplets dataset [47]. In contrast, all other methods are trained purely on the Vimeo 90K Triplets dataset without high-resolution images. Therefore, **the training scale of PerVFI is almost doubled with high-resolution data included** compared to other methods. To ensure a fair comparison, we display the results of PerVFI in **gray background** and exclude them from ranking. However, we still include it as a reference to assess our method’s capability. Even with a much smaller training scale than PerVFI, our method still achieves a better performance than PerVFI in most datasets and metrics.

Under a fair comparison, our method achieves the best performance in all metrics and datasets, as shown in Tab. 1.

In the most challenging datasets: SNU-FILM extreme [7] and Xiph-4K [29], our proposed methods achieve approximately **20%** improvements in FID and FloLPIPS over the second-best results. In relatively easier datasets like Xiph-2K, our method achieves around **20%** improvements in all metrics over the second-best result. In some datasets like DAVIS and Xiph-2K, we notice that newer methods can sometimes underperform older methods (like EMAVFI’23 vs IFRNet’22), but our method consistently outperforms others, indicating a **strong generalization capability** over different datasets. Moreover, we observe that the improvement margin increases when the task gets harder (such as SNU-FILM), indicating **our strong capability for challenging cases**, which is more crucial than incrementing quantitative scores on near-perfect results such as SNU-FILM-easy and medium.

Runtime Analysis. Other than interpolation quality, our method achieves good efficiency, shown in the last column of Tab. 1. Under the same number of sampling steps (10 in our setup), our method achieves **2.3** \times speedup over Consec. BB and **4.3** \times speedup over LDMVFI. Our method also achieves **2.2** \times speedup than recent non-diffusion SOTA: PerVFI. This is an important step toward efficiency of diffusion-based methods, as previous diffusion-based methods are inefficient. Moreover, even PerVFI contains a larger training scale, our method achieves comparable results and much faster inference speed.

Qualitative Evaluation. We select Consec. BB [26], PerVFI [45], EMAVFI [49], and IFRNet [19] to visually compare the interpolation results since they achieve strong results in challenging cases like SNU-FILM-extreme. The qualitative results are shown in Fig. 3, with examples selected from challenging cases in SNU-FILM-extreme, Xiph-4K, and DAVIS. Overlaid Input means the blended image of I_0 and I_1 . In the first row, results from IFRNet, EMAVFI, and Consec. BB are largely distorted, and the result in PerVFI contains a distorted eye. Our method aligns with the ground truth. In the second row, our method is the only one that does not contain “the third leg” in the middle.



Figure 3. **Qualitative Comparison between our method and recent SOTAs.** The leftmost image is the overlaid image of I_0 and I_1 (blended image). Areas with drastic motion changes are cropped with blue boxes to better visualize the results. Red circles and boxes indicate the area where we perform significantly better. Our method achieves better visual quality than recent SOTAs. Additional qualitative results are included in the [Supplementary Material](#).



Figure 4. **Inconsistency between PSNR/SSIM and Visual Quality.** Red circles and arrows indicate where the results from EMAVFI are distorted.

The person’s right leg in I_0 and left leg in I_1 are at almost the same position, resulting in an artificial “third leg” generated by other methods. In the third row, IFRNet, EMAVFI, and Consec. BB generate distorted results, and the vehicle generated by PerVFI misses some parts. However, our method generates a high-quality image. In the fourth row, EMAVFI generates blurred results, and IFRNet and Consec. BB generate distorted results. PerVFI, though without large distortion, contains artifacts shown in the red circles, but our result is realistic. We additionally include $8\times$ interpolation results in our [Supplementary Material](#).

Inconsistency Between PSNR/SSIM and Visual Qual-

ity. Examples in Fig. 4 indicate that PSNR/SSIM are inconsistent with visual quality. In the first row, the female’s hand in the red circle is blurred and distorted in the result of EMAVFI, but our result is consistent with ground truth. In the second row, there is distortion (pointed by arrow) on hairs in the result of EMAVFI, and ours is also consistent with the ground truth. However, in both examples, the PSNR/SSIM in EMAVFI is better, and our LPIPS is better, indicating that LPIPS is consistent with visual quality while PSNR/SSIM are not.

Training Cost of Diffusion Based Models. As claimed, our method achieves $20\times$ fewer model parameters and requires $9000\times$ less training data than video-based diffusion models [17, 37, 48]. We include quantitative results of the number of parameters and training data size of both image-diffusion-based methods and video-based diffusion methods in Tab. 2. Since MCVD [42] gets much worse performance than recent SOTAs (see Tab. 1), we exclude it for comparison. VIDIM [17] is trained with WebVid-10M [2] and private videos, where WebVid-10M is an extension of WebVid-2M, which has 13K hours of videos, corresponding to 1.4T frames. Vimeo Triplet 90K contains 51K triplets, corresponding to 153K frames, and therefore, our method gets at least $9000\times$ less training data than VIDIM. The

Table 2. Training and inference cost of recent diffusion-based methods. * means VIDIM [17] is trained additionally with private datasets. † means the performance of MCVD is much worse than SOTAs of recent years. Runtime is measured in seconds per frame with 480×720 resolution. N/A indicates that VIDIM [17] is not open-sourced, and NF indicates that Dreammover [37] needs fine-tuning for every single inference example, which takes minutes.

Method	Input type	# Parameters	# Training Data	Pretrained	Runtime
LDMVFI'24 [9]	Images	439.0M	51K Triplets	✗	2.48
Consec. BB'24 [26]	Images	146.4M	51K Triplets	✗	1.62
MCVD'22 [42]	Videos	27.3M†	51K Triplets†	✗	52.55
VIDIM'24 [17]	Videos	>1B	>10M videos*	✗	N/A
Dreammover'24 [37]	Videos	943.2M	244.7M Videos	✓	NF
ViBiDSampler'24 [48]	Videos	943.2M	244.7M Videos	✓	8.48
Ours	Videos	46.7M	51K Triplets	✗	0.69




Figure 5. **Visualization of 3D Wavelet Transform.** The 3D wavelet transform can extract frequency information along different directions (time, height, width). We visualize results by applying a high pass filter in the time dimension with a combination of (low, low), (low, high), and (high, low) filters in spatial dimensions.

number of parameters, based on Tab. 2, is $20\times$ fewer. Even comparing with image-based diffusion methods [9, 26], our number of parameters is over $3\times$ fewer.

4.3. Ablation Studies

Temporal Aware Design. We include ablation studies in Tab. 3 to indicate the effectiveness of each component. We start with our method and gradually remove 3D wavelet, replace cross-attention aggregation by average pooling, remove temporal attention, and replace 3D convolution by 2D. We also include Ours[†], where we replace Image Encoder by a temporal aware design in our full model. The performance gets worse since multi-level features are changed due to zero replacement in this setup, indicating the necessity of the Image Encoder. We additionally show three example visualizations of 3D wavelet transform in Fig. 5. We apply high-pass and low-pass filters as convolution kernels to the input videos. By switching between high-pass and low-pass filters in temporal, height, and width dimensions, we can effectively extract 8 different frequency maps. Our temporal autoencoder plays the primary role, with the 3D wavelet feature gating contributing pixel-level information that further enhances the results.

Distribution Shift in Brownian Bridge. The Brownian Bridge in Consec. BB [26] connects the encoder features of adjacent frames. Since adjacent frames are similar, features connected with Brownian Bridge Diffusion are almost identical, and therefore the diffusion is approximately

Table 3. Ablation studies. The result is in FID. The - signs mean that the component is removed. The removal is one at a time. † means that we use the full version of our method, but the Image Encoder is temporal-aware. ‡ means we disable feature sharing from †.

Dataset	Xiph-4K	Xiph-2K	SNU-FILM-extreme
Ours	19.114	9.901	29.868
- 3D Wavelet	19.247	10.092	30.717
- Cross-attention aggregation	19.663	10.499	30.903
- Temporal attention	19.944	10.911	32.061
- 3D Convolution	23.481	12.679	33.155
Ours [†]	22.731	13.410	34.982
Ours [‡]	20.004	10.769	32.018




Figure 6. $MAPE(\mathcal{E}(I_n), \mathcal{E}(I_0))$ in the setup of Consec. BB [26] and $MAPE(\mathcal{E}(V), \mathcal{E}(\tilde{V}))$ in our method. \tilde{V} is the video clip with the intermediate frame replaced with 0s. The MAPE in Consec. BB is less than 1%, resulting in a rough identity transformation. In our method, MAPE is 40-50%, and therefore the Brownian Bridge learns to reduce this gap.

an identity mapping. However, our design mitigates this issue by replacing I_n in the video clip V (results denoted as \tilde{V}). Following proposition 1, we do a t-test for our method and Consec. BB [26]. The dataset is selected to be SNU-FILM extreme [7] for the t-test. The t-statistic computed in our method is more than 21, which is much larger than the threshold of significance level 0.001: 3.291. This means that we can reject H_0 . However, the t-statistic computed in the setup of Consec. BB is about 0.0001, where we cannot reject H_0 . We show some examples in Fig. 6. The Mean Absolute Percentage Error (MAPE) is used to measure the distribution shift of encoded features, where $MAPE(x, y) = \mathbb{E}|\frac{x-y}{x}|$, indicating the average percentage change. In these examples, we can see that the MAPE in Consec. BB is close to 0, but our method gives a large MAPE, experimentally showing the distributional shift.

5. Conclusion

In this paper, we introduce our TLB-VFI to extract temporal information in both latent space (temporal blocks) and pixel space (3D wavelet feature gating). With such a design, our method achieves state-of-the-art results and solves the limitations of recent diffusion-based methods. Our method is highly flexible as video diffusion-based models that we can handle more than 3 frames as input even though our method is trained with triplets. Meanwhile, we do not require large-scale training like video diffusion-based models because we take advantage of flow estimation to guide the generation.

References

- [1] Dawit Mureja Argaw and In So Kweon. Long-term video frame interpolation via feature propagation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2, 3
- [2] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *IEEE International Conference on Computer Vision*, 2021. 7
- [3] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023. 3
- [4] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *European Conf. on Computer Vision (ECCV)*, 2012. 6
- [5] Xianhang Cheng and Zhenzhong Chen. Video frame interpolation via deformable separable convolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020. 2, 3
- [6] Jinsoo Choi, Jaesik Park, and In So Kweon. High-quality frame interpolation via tridirectional inference. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021. 2, 3
- [7] Myungsub Choi, Heewon Kim, Bohyung Han, Ning Xu, and Kyoung Mu Lee. Channel attention is all you need for video frame interpolation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020. 2, 5, 6, 8
- [8] Duolikun Danier, Fan Zhang, and David Bull. Flolpips: A bespoke video quality metric for frame interpolation. In *2022 Picture Coding Symposium (PCS)*. IEEE, 2022. 5
- [9] Duolikun Danier, Fan Zhang, and David Bull. Ldmvfi: Video frame interpolation with latent diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1472–1480, 2024. 2, 3, 4, 5, 6, 8
- [10] A. Dosovitskiy, P. Fischer, E. Ilg, P. Häusser, C. Hazırbaş, V. Golkov, P. v.d. Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2015. 6
- [11] Saikat Dutta, Arulkumar Subramaniam, and Anurag Mittal. Non-linear motion estimation for video frame interpolation using space-time convolutions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022. 2, 3
- [12] John Flynn, Ivan Neulander, James Philbin, and Noah Snavely. Deepstereo: Learning to predict new views from the world’s imagery. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016. 1
- [13] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 2017. 5
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2, 3
- [15] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646, 2022. 3
- [16] Zhewei Huang, Tianyuan Zhang, Wen Heng, Boxin Shi, and Shuchang Zhou. Real-time intermediate flow estimation for video frame interpolation. In *European Conference on Computer Vision*, 2022. 2, 3
- [17] Siddhant Jain, Daniel Watson, Eric Tabellion, Ben Poole, Janne Kontkanen, et al. Video interpolation with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7341–7351, 2024. 2, 3, 6, 7, 8
- [18] 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/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 3, 6
- [19] 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 (CVPR)*, 2022. 6
- [20] 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*, 2020. 2, 3
- [21] Bo Li, Kaitao Xue, Bin Liu, and Yu-Kun Lai. BbDM: Image-to-image translation with brownian bridge diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 2, 3, 5
- [22] Zhen Li, Zuo-Liang Zhu, Ling-Hao Han, Qibin Hou, Chun-Le Guo, and Ming-Ming Cheng. AMT: All-pairs multi-field transforms for efficient frame interpolation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 3, 6
- [23] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. DPM-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 2022. 3
- [24] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. DPM-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*, 2022. 3
- [25] Liying Lu, Ruizheng Wu, Huaijia Lin, Jiangbo Lu, and Jiaya Jia. Video frame interpolation with transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2, 3, 6
- [26] Zonglin Lyu, Ming Li, Jianbo Jiao, and Chen Chen. Frame interpolation with consecutive brownian bridge diffusion. *arXiv preprint arXiv:2405.05953*, 2024. 2, 3, 4, 5, 6, 8
- [27] S.G. Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1989. 5

- [28] N. Mayer, E. Ilg, P. Häusser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [29] Simon Niklaus and Feng Liu. Softmax splatting for video frame interpolation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 2, 5, 6
- [30] 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*, 2017. 2, 3
- [31] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive separable convolution. In *Proceedings of the IEEE international conference on computer vision*, 2017. 2, 3
- [32] Junheum Park, Jintae Kim, and Chang-Su Kim. Biformer: Learning bilateral motion estimation via bilateral transformer for 4k video frame interpolation. In *Computer Vision and Pattern Recognition*, 2023. 2, 3
- [33] 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 *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5
- [34] Markus Plack, Karlis Martins Briedis, Abdelaziz Djelouah, Matthias B Hullin, Markus Gross, and Christopher Schroers. Frame interpolation transformer and uncertainty guidance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 2, 3
- [35] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022. 3
- [36] Sheldon M Ross. *Stochastic processes*. 1995. 3
- [37] Liao Shen, Tianqi Liu, Huiqiang Sun, Xinyi Ye, Baopu Li, Jianming Zhang, and Zhiguo Cao. Dreammover: Leveraging the prior of diffusion models for image interpolation with large motion. *arXiv preprint arXiv:2409.09605*, 2024. 2, 3, 6, 7, 8
- [38] Zhihao Shi, Xiaohong Liu, Kangdi Shi, Linhui Dai, and Jun Chen. Video frame interpolation via generalized deformable convolution. *IEEE transactions on multimedia*, 2021. 2, 3
- [39] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. 3
- [40] 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*, 2020. 6
- [41] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 3, 4
- [42] Vikram Voleti, Alexia Jolicoeur-Martineau, and Chris Pal. Mcvd-masked conditional video diffusion for prediction, generation, and interpolation. *Advances in neural information processing systems*, 2022. 2, 3, 6, 7, 8
- [43] 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*, 2004. 5
- [44] Chao-Yuan Wu, Nayan Singhal, and Philipp Krahenbuhl. Video compression through image interpolation. In *Proceedings of the European conference on computer vision (ECCV)*, 2018. 1
- [45] Guangyang Wu, Xin Tao, Changlin Li, Wenyi Wang, Xiaohong Liu, and Qingqing Zheng. Perception-oriented video frame interpolation via asymmetric blending. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2753–2762, 2024. 2, 3, 4, 5, 6
- [46] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezaatofghi, and Dacheng Tao. Gmflow: Learning optical flow via global matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022. 6
- [47] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision (IJCV)*, 2019. 3, 5, 6
- [48] Serin Yang, Taesung Kwon, and Jong Chul Ye. Vibidsampler: Enhancing video interpolation using bidirectional diffusion sampler. *arXiv preprint arXiv:2410.05651*, 2024. 2, 3, 6, 7, 8
- [49] 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*, 2023. 2, 3, 5, 6
- [50] Guozhen Zhang, Chunxu Liu, Yutao Cui, Xiaotong Zhao, Kai Ma, and Limin Wang. Vfimamba: Video frame interpolation with state space models. *arXiv preprint arXiv:2407.02315*, 2024. 2
- [51] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 5