

Low-Bitrate Video Compression through Semantic-Conditioned Diffusion

Lingdong Wang^{1,2} Guan-Ming Su² Divya Kothandaraman² Tsung-Wei Huang²
 Mohammad Hajiesmaili¹ Ramesh K. Sitaraman¹

¹University of Massachusetts Amherst

²Dolby Laboratories

{lingdongwang, hajiesmaili, ramesh}@cs.umass.edu

{guanming.su, Divya.Kothandaraman, Tsung-Wei.Huang}@dolby.com

Abstract

Traditional video codecs optimized for pixel fidelity collapse at ultra-low bitrates and produce severe artifacts. This failure arises from a fundamental misalignment between pixel accuracy and human perception. We propose a semantic video compression framework named DiSCo that transmits only the most meaningful information while relying on generative priors for detail synthesis. The source video is decomposed into three compact modalities: a textual description, a spatiotemporally degraded video, and optional sketches or poses that respectively capture semantic, appearance, and motion cues. A conditional video diffusion model then reconstructs high-quality, temporally coherent videos from these compact representations. Temporal forward filling, token interleaving, and modality-specific codecs are proposed to improve multimodal generation and modality compactness. Experiments show that our method outperforms baseline semantic and traditional codecs by 2-10× on perceptual metrics at low bitrates.

1. Introduction

Traditional video compression algorithms, including standard codecs [3, 45] and neural codecs [1, 8, 13, 26, 28, 29], have been optimized to achieve higher pixel-wise accuracy for decades. However, this paradigm collapses when operating in the ultra-low-bitrate regime. With limited available information, traditional codecs produce severe artifacts like blocking, blurring, and temporal flickering that destroy perceptual realism. This failure is not just technical but conceptual. Extreme bandwidth constraints expose the core challenge of compression: deciding *which information matters most* when nearly everything must be discarded. The low-bitrate frontier thus reveals a fundamental misalignment between optimization objectives and human percep-

tion. While conventional codecs focus on minimizing pixel errors, humans interpret videos through a general semantic understanding rather than by analyzing individual pixels.

Semantic compression addresses this misalignment by prioritizing semantic accuracy over pixel fidelity. Instead of directly transmitting pixels, the source video is factorized into a set of semantically expressive and compression-friendly modalities and then reconstructed via a generative decoder. Recent advances in semantic image compression have attempted to represent visual contents by text, sketches, or coarse base images [10, 19, 20, 37, 57]. While semantic video compression remains relatively underexplored, early works like M3-CVC [47] proposed to use key-frames and text description, T-GVC [50] introduced key-frames and motion trajectories. CVG [54], a concurrent work with us, leveraged a combination of text, first-last frames, 3D human poses, segmentation masks, and optical flows [54]. These semantic codecs not only align with human perception but also introduce additional information from the powerful generative prior, enabling realistic reconstruction under stringent bitrate constraints.

However, existing approaches face several limitations. *First*, they address either spatial enhancement [19, 20, 37] or temporal interpolation of video frames [47, 50, 54] separately. In contrast, we propose a joint spatiotemporal generation process that fully leverages the generative potential of diffusion models while harnessing the intrinsic relationship between space and time within media. *Second*, existing works [54] often select semantically overlapping modalities that describe the same scene using multiple perspectives, leading to information redundancy. Instead, we select complementary modalities and introduce a token interleaving strategy to eliminate repetitive descriptions for each frame. *Third*, although semantic modalities are naturally more compact than RGB frames, their potential compactness has not been fully explored. In this work, we develop or finetune specialized codecs tailored to each modality.

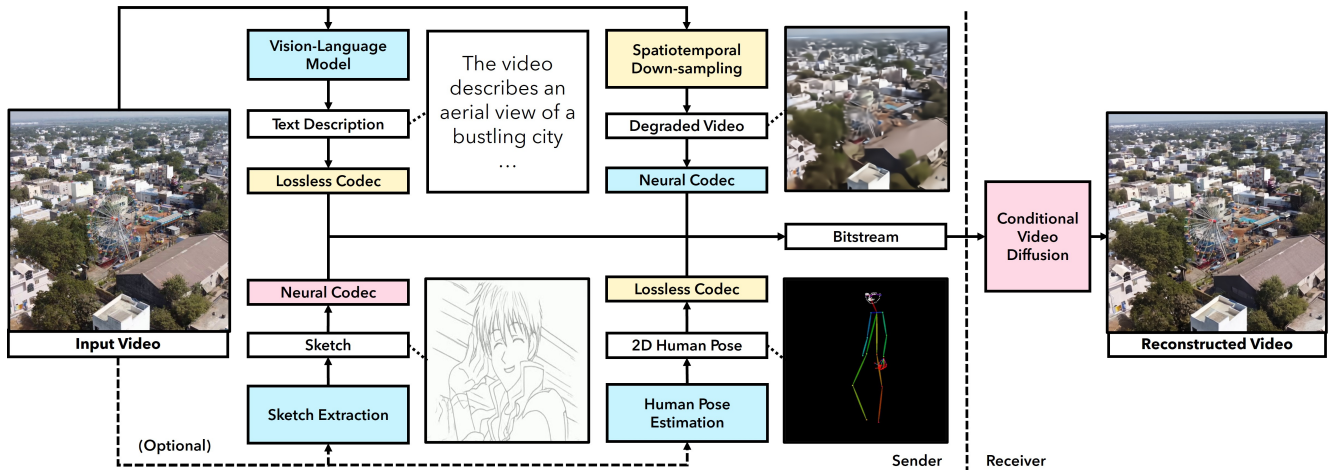


Figure 1. Overview of proposed method. Red means trainable module, blue means frozen module, yellow means non-learning operations.

In this paper, we develop a framework named **DiSCo** (**Diffusion with Semantic Conditioning**) as in Fig. 1. Specifically, the input video is decomposed into complementary components: a **textual description** of content capturing global semantics, generated by pretrained vision-language model [33] and compressed by lossless codec; a **degraded video** preserving coarse appearance, down-sampled via joint spatiotemporal operations and compressed by pretrained neural codec [26]; optional **sketch** or **2D human pose** providing motion cues, extracted by pretrained deep-learning models [6, 53] and compressed via customized lossy or lossless codecs.

At the receiver, we formulate semantic compression as a video-to-video transformation task, converting a degraded multimodal reference video into a high-quality counterpart. To solve this problem, we employ a conditional video diffusion model, fine-tuned from pretrained weights [23] using in-context LoRA [25]. During the spatiotemporal generation process, we introduce a temporal forward-filling scheme to enhance the reconstruction quality. Furthermore, a token-interleaving mechanism is proposed to alter multimodal references, effectively reducing information redundancy during streaming and conditioning. Extensive experiments confirm that our framework significantly surpasses traditional standard or neural codecs by 5 - 10 \times and baseline semantic codecs by 2 - 3 \times in perceptual metrics when operating at low bitrates. Our contributions are as follows:

1. A semantic video compression framework that factorizes a source video into a text description, a spatiotemporally degraded video, and optional sketches or poses.
2. A joint spatiotemporal generation scheme with a forward-filling strategy that fully exploits the generative power of diffusion models.
3. A token interleaving mechanism for efficient, redundancy-free, and streaming-friendly condition-

ing of diffusion Transformers.

4. Specialized codecs for each semantic modality, maximizing compactness and fidelity.
5. Extensive experiments showing that our method outperforms traditional and semantic codecs by 2–10 \times in perceptual quality under low bitrates.

2. Related Works

2.1. Generative Models

The emergence of diffusion denoising models [43] has dramatically advanced the capabilities of generative AI. Their power to synthesize high-fidelity content has been demonstrated in text-to-image, text-to-video, and image-to-video generation tasks. In the context of video processing, these models have been successfully applied to specific restoration sub-tasks. Their aptitude for spatial detail synthesis has been proven in super-resolution [12, 49, 56] and image quality enhancement [9, 60]. Concurrently, their ability to model complex motion has been leveraged for video frame interpolation [7, 16, 61]. However, these applications usually focus solely on either spatial enhancement or temporal interpolation in isolation. In contrast, we introduce a joint spatiotemporal diffusion process, harnessing the full range of diffusion model’s generative potential. This approach enables us to achieve exceptional restoration quality.

2.2. Neural Compression

Neural compression has revolutionized the field by replacing hand-crafted modules with trainable neural networks. While Implicit Neural Representation (INR) [8, 28, 29] proposes to encode an entire video into the weights of a neural network, this approach often suffers from its overfitting nature and limited compression performance. The dominant paradigm, inspired by traditional transform coding, is the

autoencoder architecture. Pioneering works [1, 13] established the core components: a learned transform, a differentiable quantization proxy, and a hyperprior for entropy modeling. State-of-the-art methods like DCVC-RT [26] and DCVC-FM [32] now achieve performance competitive with the latest standard codec, H.266/VVC [3]. Recent efforts have also sought to integrate generative priors into the codec by performing compression in the latent space of a diffusion tokenizer [22], VQ-VAE [27, 41], or GAN [18]. These methods remain unimodal and rely on implicit deep features, whereas our method adopts a completely different paradigm. We factorize video content into explicit multimodal representations, enabling explainable semantic conditioning and complementary cross-modal guidance.

2.3. Semantic Compression

Semantic compression decomposes media into semantically meaningful components, which are then used by a generative model for reconstruction. For semantic image compression, many studies explore representing an image with a text description and a heavily compressed visual base [19, 20, 37]. Other modalities, such as sketches [57] or more complex combinations of text, sketches, and coarse textures [10], have also been investigated. To capture finer details, some methods incorporate object-level information, such as object positions, labels, and individual images [30, 44]. Prior works have extended semantic compression to video by transmitting the first frame alongside sketches and text [55], compressed video with sparse motion trajectories [50], or a combination of first-last frame, text description, 3D human pose, segmentation mask, and optical flow [54]. Other strategies involve sending keyframes with corresponding text descriptions [47], using learned semantic maps [40], or transmitting abstract features [21]. Some of these methods are tailored for specific use cases, such as robustness in noisy wireless channels [21, 40, 55] or portrait video streaming [11]. In comparison, our work introduces an innovative suite of semantic modalities for general contents, equipped with a joint spatiotemporal generation process, a token interleaving scheme, and specialized codecs for each modality to eliminate information redundancy and improve reconstruction quality.

3. Proposed Method

3.1. Multimodal Encoding

Our framework represents a video through a set of compact and expressive semantic modalities. Specifically, the sender factorizes the source video into (1) a textual description for global semantics, (2) a spatiotemporally downsampled video for structural and color guidance, and optionally (3) sketch or pose sequences for motion and geometry cues. Each modality is encoded using a specialized codec

designed for maximal compactness. All encoded modalities are transmitted as separate bitstreams.

Text. We employ a pretrained vision–language model LLaVA [33] to generate concise textual descriptions of the video content. This modality provides global semantics such as objects, scene context, and coarse actions in a highly compact and interpretable form. Since text is compact and error-sensitive, we compress it using lossless codec Zlib which achieves the best empirical performance.

Video. To encode visual structure under extreme bitrate constraints, we aggressively reduce both the spatial and temporal resolutions of the source video. Let D_s and D_t denote the spatial and temporal down-sampling factors, respectively. Only one frame out of every D_t is retained, and each frame is bilinearly down-sampled by D_s times before compression. This produces a spatiotemporally degraded video scaffold that preserves coarse appearance and motion cues. The downsampled sequence is then compressed using a pretrained neural video codec DCVC-RT [26] with low-quality settings. Unlike existing methods that only consider either spatial or temporal dimension, our design simultaneously discards redundant pixels and frames and relies on the diffusion-based decoder to regenerate fine spatial details and temporal continuity.



Figure 2. Conditioning on sketch/pose modality at 0.005 BPP.

Sketch. While text and degraded video are suitable for any content, sketches can be particularly expressive and compact for content featuring object contours, such as animations. Since we found that traditional edge detection algorithms like Canny [5] produce temporally inconsistent results, we adopt a pretrained deep-learning model InformativeDrawings [6] to extract line-drawing sketching. Unlike RGB inputs, sketch videos as shown in Fig. 2 are characterized by a grayscale value range and prominent high-frequency edges. To capitalize on this property, we fine-tune a neural DCVC-RT codec over sketch data to compress it and achieve additional bitrate reduction. More details of

the finetuned codec can be found in Sec. 4.2.

Pose. Human pose serves as a powerful semantic cue in human-centric videos. We utilize a pretrained deep-learning method DWPose [53] to estimate 2D human skeletons in the OpenPose format as shown in Fig. 2. The keypoint coordinates will be quantized into 16-bit unsigned integers. The quantized keypoints, along with metadata such as frame resolution, frame count, and the number of poses per frame, are serialized and compressed into bits using the LZMA algorithm. This custom compression pipeline ensures no precision loss, as the keypoints always fall within the resolution range and are ultimately rendered as pixels.

3.2. Generative Decoding

At the receiver side, the system reconstructs a high-quality video conditioned on the decoded signals.

Text. The textual bitstream is decompressed back to content description, and then concatenated with predefined negative prompts such as “blurry” or “distorted,” encouraging the model to generate visually pleasing results. The text is encoded into embeddings using a pretrained T5 text encoder [42] and fed to the diffusion model.

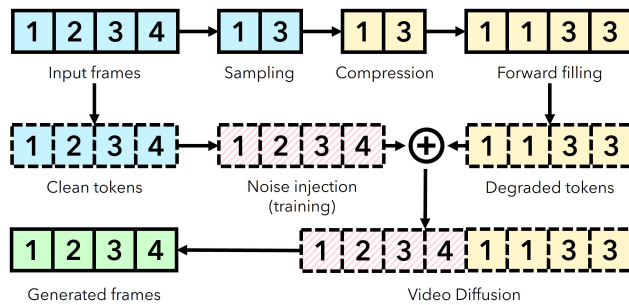


Figure 3. Workflow of the degraded video modality.

Video. The visual bitstream is decompressed into a spatiotemporally degraded video in the pixel domain. Subsequently, we restore its original resolution and frame rate using bilinear spatial upsampling and a temporal **forward filling** technique. Specifically, the forward filling method duplicates the previously sampled frame to fill in the dropped positions, ensuring the continuity of media in the latent space. As illustrated in Fig. 3, for a video with four frames, only the first and third frames are retained if the temporal down-sampling factor is $D_t = 2$. After compression, the forward filling stage pads the video back to four frames on the receiver’s side. Ablation study at Sec. 4.4 proves that this strategy brings significant advantage compared with no filling or zero-value filling. Finally, the reconstructed degraded video is transformed into latent tokens using the pretrained VAE of LTX-Video [23] for further processing.

Sketch and Pose. At the receiver, the sketch bitstream is decompressed and converted into latent representations

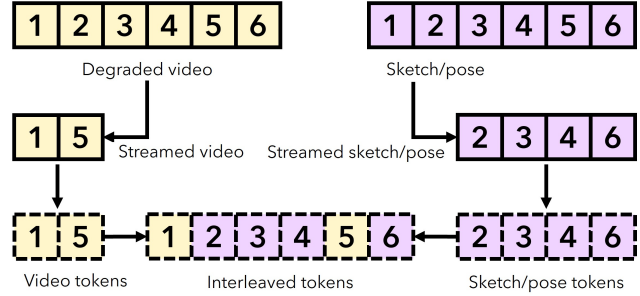


Figure 4. Illustration of token interleaving.

through the same diffusion VAE encoder as the video modality. Alternatively, if the pose modality is enabled, the bitstream will be decoded back to keypoint coordinates, rendered as skeleton images, and then encoded via VAE. We introduce a **token interleaving** strategy to condition the diffusion model on multiple modalities. As shown in Fig. 4, tokens from different modalities are alternated in sequence, with one token from the degraded video followed by several from the auxiliary modality. The interleaved tokens will replace the degraded tokens in Fig. 3 for conditioning. Additionally, only the frames corresponding to the selected tokens will actually be streamed. Instead of concatenating modality tokens [54] or introducing extra modules [58], this interleaving mechanism ensures that each frame is described exclusively by one modality, avoids introducing extra streaming or computational overhead, and prevents semantically overlapping information. Furthermore, we found that token is the minimum granularity for interleaving. A detailed analysis is provided in the appendix.

Diffusion Model. We formulate semantic compression as a conditional video-to-video generation task, where the diffusion model transforms a degraded multimodal reference video into a high-quality counterpart. The core generative decoder is a video diffusion Transformer fine-tuned from the LTX-Video model [23]. Our training strategy is inspired by in-context LoRA [25]. Specifically, only a lightweight low-rank adapter is trained on top of the attention layers, while the large pretrained diffusion Transformer remains frozen. During the training time, as shown in Fig. 3, the input frames will be encoded into clean latent tokens via VAE, and then converted into noisy tokens by injecting Gaussian noises. At inference time, these noisy tokens are purely sampled Gaussian noises. The noisy tokens will then be concatenated with the degraded tokens, or the interleaved tokens if auxiliary modality is enabled, and fed into the video diffusion model. The diffusion Transformer conditions on the text tokens and this concatenated visual tokens. And it is trained to reconstruct the clean latent out of the degraded ones under the supervision of L2 loss following rectified flow matching [35].

4. Experiment

4.1. Experiment Settings

Dataset. We utilize an 8000-video subset of the OpenVid-1M dataset [38] and resize all videos to 512×512 -resolution 57-frame clips as our training data. The testing videos share the same resolution and sequence length unless otherwise specified. We evaluate the general version of our method, which incorporates only text and video modalities, on widely used compression benchmarks, including HEVC-B [39], MCL-JCV [48], and UVG [36] datasets. The auxiliary sketch and pose modalities are most effective for special contents. Therefore, we evaluate our method with an additional sketch modality on a 50-video subset of the animation video dataset Anim400K [4]. Similarly, we test a variant of our method with an additional pose modality on a 50-video subset of the human-centric video dataset Open-HumanVid [31]. Low-quality videos with DOVER scores [51] below 0.4 are filtered out in preprocessing.

Baselines. Our comparison baselines include the standard video codecs H.264/AVC, H.265/HEVC, and H.266/VVC [3]. We use the FFMPEG [45] implementations of these standards, namely x264, x265, and x266. We compare with the state-of-the-art neural compression technique DCVC-RT [26] and DCVC-FM [32]. For semantic compression methods, we compare against CVG [54] and T-GVC [50]. Since recent generative codecs [22, 41, 50, 54] do not provide open-source codes for reproduction, we use the LPIPS and FVD scores reported in the respective papers [50, 54] for comparison and ignore inaccessible results [22, 41]. We also test the generative method PLVC [52], yet it cannot operate in the low-bitrate zone we focus on.

Metrics. We adopt four semantic quality metrics, including LPIPS [59], DISTs [17], FVD [46], and FID [24], that align better with human perception. We also use two traditional metrics, PSNR and SSIM, to measure the pixel fidelity. In addition, we measure the temporal consistency using FloLPIPS [15] and FVMD [34]. We evaluate the overall compression performance using BD-rate [2]. It measures the bitrate savings required to achieve the same quality level as the anchor method, which is set to be H.266.

4.2. Implementation details

Video codec. We utilize a pretrained DCVC-RT [26] neural codec. The Quantization Parameter (QP) for DCVC-RT is set to 0, 8, 16, 24, and 32. Additionally, we apply spatial down-sampling factor D_s of 1, 2, and 4, along with temporal factor D_T of 2, 4, and 8. These settings are specifically designed to achieve low-bitrate video compression.

Sketch codec. We finetune a DCVC-RT codec on sketches extracted from the Anim400K dataset to better align with its unique data distribution, which differs from standard RGB videos. Since the original DCVC-RT training pipeline is

not publicly available, we reproduce it using a two-stage training strategy. In the first stage, we freeze its P-frame network while only training the I-frame network. And vice versa in the second stage to stabilize the training.

Generative decoder. We finetune a video diffusion model from the 13B-parameter LTX-Video [23] pretrained weights. We insert a 256-rank LoRA module containing 625 million parameters and train it using the AdamW optimizer for 8,000 steps. The learning rate is set to 0.0002 with a cosine learning rate scheduler. All models are quantized to BF16 precision. We configure the number of denoising steps to 50. Full details can be found in the appendix.

4.3. Performance Comparison

Table 1. BD-rate (%) comparison on HEVC-B dataset.

Method	LPIPS↓	DISTS↓	FID↓	FVD↓
H.266	0.00	0.00	0.00	0.00
H.265	+58.55	+34.87	+48.36	+36.85
H.264	+108.43	+109.72	+78.91	+79.30
DCVC-FM	+50.83	+112.17	+90.35	+73.48
DCVC-RT	+10.73	+61.10	+51.89	+25.19
DiSCo(ours)	-83.97	-91.69	-90.12	-89.26

Table 2. BD-rate (%) comparison on MCL-JCV dataset.

Method	LPIPS↓	DISTS↓	FID↓	FVD↓
H.266	0.00	0.00	0.00	0.00
H.265	+48.66	+30.16	+68.08	+59.32
H.264	+74.78	+94.01	+102.38	+95.91
DCVC-FM	+22.05	+87.67	+55.76	+17.51
DCVC-RT	-0.34	+52.92	+39.46	+2.42
DiSCo(ours)	-81.74	-86.04	-84.61	-83.55

Table 3. BD-rate (%) comparison on UVG dataset.

Method	LPIPS↓	DISTS↓	FID↓	FVD↓
H.266	0.00	0.00	0.00	0.00
H.265	+60.17	+46.72	+63.14	+34.30
H.264	+84.78	+109.97	+99.56	+75.13
DCVC-FM	+35.70	+92.83	+71.24	+22.45
DCVC-RT	+8.73	+73.66	+74.35	+21.22
DiSCo(ours)	-85.03	-90.88	-90.71	-85.82

Quality Comparison. We compare our proposed method to the baseline approaches visually in Fig. 5. At a low-bitrate video compression rate of 0.005 bits per pixel (BPP), we observe that the state-of-the-art neural codec DCVC-RT exhibits blurry artifacts, while H.266 suffers from blocky and



Figure 5. Visualization of compression results at 0.005 BPP (Bits Per Pixel). Zoom in for details.

motion artifacts. In contrast, our method delivers superior visual quality with rich generated details, preserving both semantic accuracy and pixel fidelity with limited bitrate. We show the rate-distortion curves over HEVC-B, MCL-JCV, and UVG datasets at Fig. 6, Fig. 7, and Fig. 8, respectively. The x-axis represents the BPP, while the y-axis shows the quality score. We also present the corresponding BD-rates for methods that have sufficient data points in Tab. 1, Tab. 2, Tab. 3. Experimental results demonstrate that in the low-bitrate regime, traditional compression algorithms perform the worst. Baseline semantic codecs outperform traditional codecs, as evidenced by their position in the bottom-left region of the rate-distortion curve plot. In comparison, our proposed method significantly outperforms all alternatives in terms of semantic metrics while maintaining comparable pixel fidelity. Our method improves the semantic compression rate by 2 to $3 \times$ over semantic compression baselines and 5 to $10 \times$ over traditional compression baselines.

Temporal Consistency. We present the results of the temporal consistency comparison on HEVC-B dataset in Fig. 9. Quantitative results demonstrate that our method produces significantly more stable videos than baselines. We also provide video visualizations in the appendix.

Generalization Ability. While our method was trained

on 512×512 resolution and 57 frames, it can generalize to higher resolutions of 1440×768 and longer sequences with 233 frames during inference without additional training. Fig. 10 highlights that the zero-shot generalized DiSCo remains a drastic advantage over baselines in terms of perceptual quality and temporal consistency.

4.4. Ablation Study

Table 4. Effect of temporal forward filling on OpenVid-1M.

Method	LPIPS BD-rate (%)
No filling	0.00
Zero filling	+60.56
Forward filling	-17.52

Effect of temporal forward filling. To evaluate our proposed forward filling scheme, we conducted an ablation study using a testing subset of the OpenVid-1M dataset. We compared three approaches: forward filling with a temporal factor of $D_t = 2$, zero-value filling with $D_t = 2$, and no temporal filling with $D_t = 1$. The BD-rate over LPIPS metrics in Tab. 4 demonstrate that filling unsampled positions with zero values leads to a significant drop in performance.

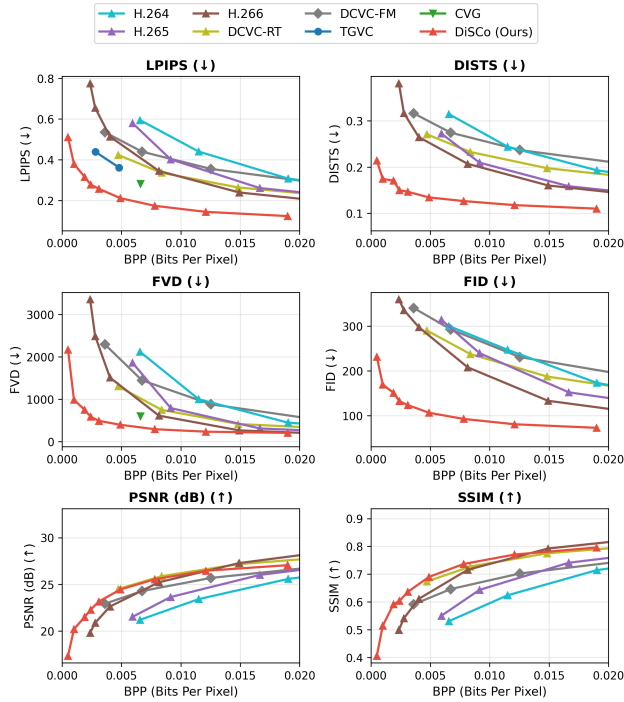


Figure 6. Quality comparison on HEVC-B dataset.

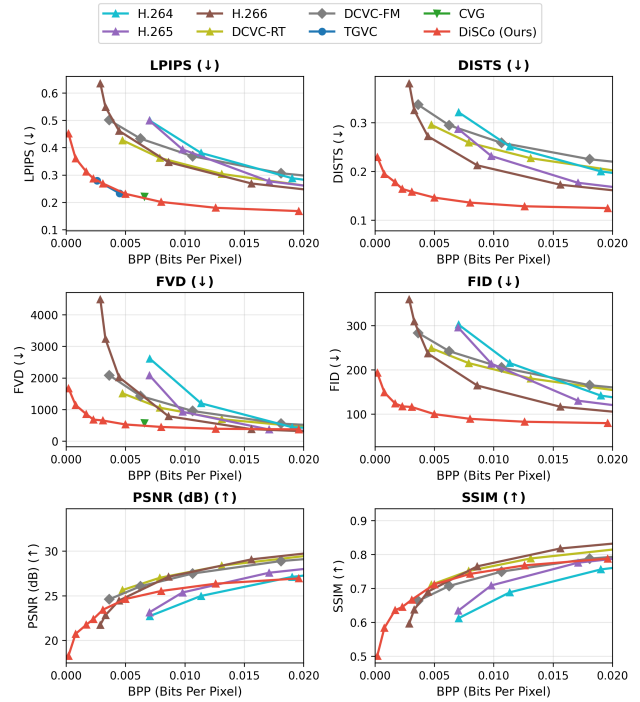


Figure 8. Quality comparison on UVG dataset.

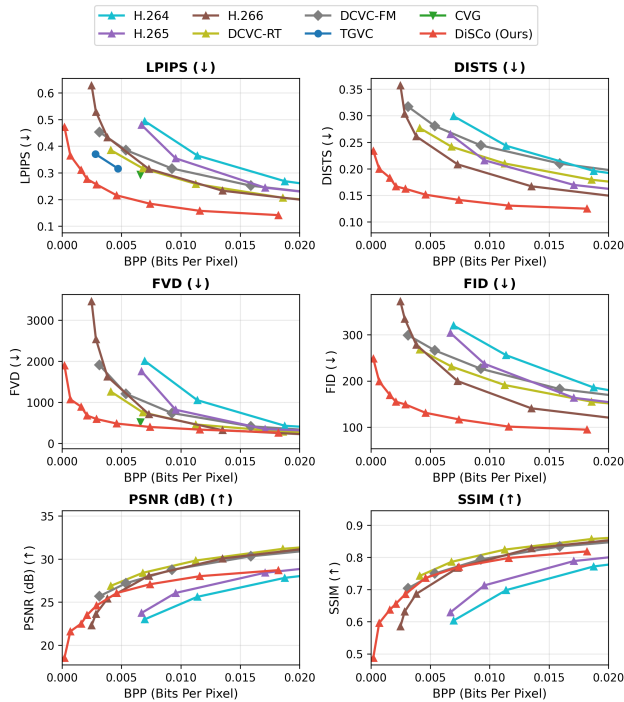


Figure 7. Quality comparison on MCL-JCV dataset.




Figure 9. Temporal consistency comparison on HEVC-B dataset.




Figure 10. Zero-shot generalization on HEVC-B dataset.

In contrast, the proposed forward filling reduces the BD-rate by 17.52% compared with no temporal generation.

Effect of text modality. We perform an ablation study to evaluate the effect of text description on compression performance. Specifically, we compare the performance of conditioning the diffusion decoder on text descriptions generated by LLaVA [33], text generated by Gemini 2.5 [14], and a placeholder text “NA” to simulate the absence

Table 5. Effect of text modality on OpenVid-1M.

Method	LPIPS BD-rate (%)
No text	0.00
With text (Gemini)	-35.96
With text (LLaVA)	-38.34

of textual input. The results presented in Tab. 5 demonstrate that incorporating text modality saves bitrate by over 30%. Among the methods tested, LLaVA provides slightly better empirical results. Consequently, we select LLaVA as our text modality extractor for this study.

Table 6. Effect of sketch modality on Anim400K.

Method	FVD BD-rate (%)
Text + video	0.00
Text + video + sketch	-4.95

Effect of sketch modality. Unlike traditional codecs that rely on a single model for all types of content, semantic compression allows the use of different modalities tailored to specific content types. Since the sketch modality is particularly effective for content that emphasizes object contours, we evaluate its performance using the Anim400K dataset and report the results for positive examples in Tab. 6. Compared to relying solely on text and degraded video, incorporating the sketch modality decreases the bitrate by 4.95%. This validates our proposed token interleaving scheme and the neural codec tailored for sketch data.

Table 7. Effect of pose modality on OpenHumanVid.

Method	FVD BD-rate (%)
Text + video	0.00
Text + video + pose	-3.67

Effect of pose modality. Similarly, we examine the performance of human pose modality over the human-centric video dataset OpenHumanVid. Results in Tab. 7 show that the additional usage of human pose can bring a BD-rate reduction of 3.67% on suitable contents.

4.5. Discussions

Ultra-low bitrate compression. Our method remains robust even at extremely low bitrates, such as 0.00045 BPP or 3.58 kbps bitrate. As illustrated in Fig. 11, this extremely constrained level of information causes the degraded video reference, or the DCVC-RT compression result, to fail in conveying meaningful semantics. For instance, the Ferris wheel and dog are barely recognizable. However, our method reconstructs the object of interest based on the textual description, albeit with differences in appearance due



Figure 11. Ultra-low bitrate compression (0.00045BPP/3.58kbps).

to the absence of color guidance. The results validate the robustness of our method and highlight the effectiveness of complementary semantic modalities. It also reveals that semantic compression consistently achieves acceptable perceptual quality by leveraging additional generative priors. While traditional compression methods typically balance rate and pixel-wise distortion, semantic compression offers a novel trade-off between rate and semantic deviation.

Efficiency. Our method supports fast training and relatively low inference latency. On an H100 GPU, finetuning a 13B-parameter Diffusion decoder requires approximately 4.5 hours. During inference, the total encoding latency is about 97.89 ms per frame, while the decoding latency is 396.20 ms. The decoder can be finetuned over a 2B distilled backbone to further reduce the decoding latency. The comprehensive analysis is provided in the appendix.

5. Conclusion

We introduced DiSCo, a semantic video compression framework that factorizes videos into complementary modalities and reconstructs them through a conditional video diffusion model. By unifying spatiotemporal generation, token interleaving, and in-context LoRA adaptation, the framework effectively leverages generative priors for perceptually rich reconstruction under extreme bitrate constraints. Specialized codecs further enhance compactness and efficiency. Extensive experiments demonstrate significant gains over traditional, neural, and prior semantic codecs, highlighting the promise of our method as a new paradigm for low-bitrate video communication.

Acknowledgments

This research was supported by grants NSF CNS-2106463, CNS-1901137, CNS-2533814, and CAREER-2045641.

References

- [1] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior, 2018. 1, 3
- [2] Gisle Bjontegaard. Calculation of average PSNR differences between RD-curves. Technical Report VCEG-M33, ITU-T SG16/Q6, Austin, TX, USA, April 2001. 5
- [3] Benjamin Bross, Ye-Kui Wang, Yan Ye, Shan Liu, Jianle Chen, Gary J. Sullivan, and Jens-Rainer Ohm. Overview of the versatile video coding (vvc) standard and its applications. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(10):3736–3764, 2021. 1, 3, 5
- [4] Kevin Cai, Chonghua Liu, and David M. Chan. Anim-400k: A large-scale dataset for automated end-to-end dubbing of video, 2024. 5
- [5] John Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6):679–698, 1986. 3
- [6] Caroline Chan, Frédo Durand, and Phillip Isola. Learning to generate line drawings that convey geometry and semantics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7915–7925, June 2022. 2, 3
- [7] Victor Fonte Chavez, Claudia Esteves, and Jean-Bernard Hayet. Time-adaptive video frame interpolation based on residual diffusion, 2025. 2
- [8] Hao Chen, Bo He, Hanyu Wang, Yixuan Ren, Ser Nam Lim, and Abhinav Shrivastava. Nerv: Neural representations for videos. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 21557–21568. Curran Associates, Inc., 2021. 1, 2
- [9] I-Hsiang Chen, Wei-Ting Chen, Yu-Wei Liu, Yuan-Chun Chiang, Sy-Yen Kuo, and Ming-Hsuan Yang. Unirestore: Unified perceptual and task-oriented image restoration model using diffusion prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17969–17979, June 2025. 2
- [10] Ruijie Chen, Qi Mao, and Zhengxue Cheng. Stable diffusion is a natural cross-modal decoder for layered ai-generated image compression. In *2025 Data Compression Conference (DCC)*, pages 361–361, 2025. 1, 3
- [11] Xinyi Chen, Weimin Lei, Wei Zhang, Yanwen Wang, and Mingxin Liu. Ultra-low bitrate predictive portrait video compression with diffusion models. *Symmetry*, 17(6), 2025. 3
- [12] Zheng Chen, Zichen Zou, Kewei Zhang, Xiongfei Su, Xin Yuan, Yong Guo, and Yulun Zhang. Dove: Efficient one-step diffusion model for real-world video super-resolution, 2025. 2
- [13] Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1, 3
- [14] Gheorghe Comanici et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities, 2025. 7
- [15] Duolikun Danier, Fan Zhang, and David Bull. Flolpips: A bespoke video quality metric for frame interpolation. In *2022 Picture Coding Symposium (PCS)*, pages 283–287, 2022. 5
- [16] Duolikun Danier, Fan Zhang, and David Bull. Ldmvfi: Video frame interpolation with latent diffusion models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(2):1472–1480, Mar. 2024. 2
- [17] Keyan Ding, Kede Ma, Shiqi Wang, and Eero P. Simoncelli. Image quality assessment: Unifying structure and texture similarity. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(5):2567–2581, 2022. 5
- [18] Pengli Du, Ying Liu, and Nam Ling. Cgvc-t: Contextual generative video compression with transformers. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 14(2):209–223, 2024. 3
- [19] Haisheng Fu, Jie Liang, Zhenman Fang, and Jingning Han. Selic: Semantic-enhanced learned image compression via high-level textual guidance, 2025. 1, 3
- [20] Yixin Gao, Xin Li, Xiaohan Pan, Runsen Feng, Zongyu Guo, Yiting Lu, Yulin Ren, and Zhibo Chen. Unimic: Towards universal multi-modality perceptual image compression, 2024. 1, 3
- [21] Lei Guo, Wei Chen, Yuxuan Sun, Bo Ai, Nikolaos Pappas, and Tony Q. S. Quek. Diffusion-driven semantic communication for generative models with bandwidth constraints. *IEEE Transactions on Wireless Communications*, 24(8):6490–6503, 2025. 3
- [22] Zongyu Guo, Zhaoyang Jia, Jiahao Li, Xiaoyi Zhang, Bin Li, and Yan Lu. Generative latent video compression, 2025. 3, 5
- [23] Yoav HaCohen, Nisan Chiprut, Benny Brazowski, Daniel Shalem, Dudu Moshe, Eitan Richardson, Eran Levin, Guy Shiran, Nir Zabari, Ori Gordon, Poriya Panet, Sapir Weissbuch, Victor Kulikov, Yaki Bitterman, Zeev Melumian, and Ofir Bibi. Ltx-video: Realtime video latent diffusion, 2024. 2, 4, 5
- [24] 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. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 5
- [25] Lianghua Huang, Wei Wang, Zhi-Fan Wu, Yupeng Shi, Huanzhang Dou, Chen Liang, Yutong Feng, Yu Liu, and Jingren Zhou. In-context lora for diffusion transformers, 2024. 2, 4
- [26] Zhaoyang Jia, Bin Li, Jiahao Li, Wenxuan Xie, Linfeng Qi, Houqiang Li, and Yan Lu. Towards practical real-time neural video compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12543–12552, June 2025. 1, 2, 3, 5
- [27] Zhaoyang Jia, Jiahao Li, Bin Li, Houqiang Li, and Yan Lu. Generative latent coding for ultra-low bitrate image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26088–26098, June 2024. 3
- [28] Hyunjik Kim, Matthias Bauer, Lucas Theis, Jonathan Richard Schwarz, and Emilien Dupont. C3: High-performance and low-complexity neural compression from a single image or video. In *Proceedings of the*

- IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9347–9358, June 2024. 1, 2
- [29] Ho Man Kwan, Ge Gao, Fan Zhang, Andrew Gower, and David Bull. Nvrc: Neural video representation compression. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37, pages 132440–132462. Curran Associates, Inc., 2024. 1, 2
- [30] Chunyi Li, Guo Lu, Donghui Feng, Haoning Wu, Zicheng Zhang, Xiaohong Liu, Guangtao Zhai, Weisi Lin, and Wenjun Zhang. Misc: Ultra-low bitrate image semantic compression driven by large multimodal model. *IEEE Transactions on Image Processing*, 34:335–349, 2025. 3
- [31] Hui Li, Mingwang Xu, Yun Zhan, Shan Mu, Jiaye Li, Kaihui Cheng, Yuxuan Chen, Tan Chen, Mao Ye, Jingdong Wang, and Siyu Zhu. Openhumanvid: A large-scale high-quality dataset for enhancing human-centric video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7752–7762, June 2025. 5
- [32] Jiahao Li, Bin Li, and Yan Lu. Neural video compression with feature modulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26099–26108, June 2024. 3, 5, 1
- [33] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 34892–34916. Curran Associates, Inc., 2023. 2, 3, 7
- [34] Jiahe Liu, Youran Qu, Qi Yan, Xiaohui Zeng, Lele Wang, and Renjie Liao. Fréchet video motion distance: A metric for evaluating motion consistency in videos, 2024. 5
- [35] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow, 2022. 4
- [36] Alexandre Mercat, Marko Viitanen, and Jarno Vanne. Uvg dataset: 50/120fps 4k sequences for video codec analysis and development. In *Proceedings of the 11th ACM Multimedia Systems Conference*, pages 331–336, 2020. 5
- [37] Shimon Murai, Heming Sun, and Jiro Katto. Lmm-driven semantic image-text coding for ultra low-bitrate learned image compression. In *2024 IEEE International Conference on Visual Communications and Image Processing (VCIP)*, pages 1–5, 2024. 1, 3
- [38] Kepan Nan, Rui Xie, Penghao Zhou, Tiehan Fan, Zhenheng Yang, Zhijie Chen, Xiang Li, Jian Yang, and Ying Tai. Openvid-1m: A large-scale high-quality dataset for text-to-video generation, 2025. 5
- [39] Jens-Rainer Ohm, Gary J Sullivan, Heiko Schwarz, Thiow Keng Tan, and Thomas Wiegand. Comparison of the coding efficiency of video coding standards—including high efficiency video coding (hevc). In *IEEE Transactions on Circuits and Systems for Video Technology*, volume 22, pages 1669–1684. IEEE, 2012. 5
- [40] Giovanni Pignata, Eleonora Grassucci, Giordano Cicchetti, and Danilo Comminiello. Lightweight diffusion models for resource-constrained semantic communication, 2024. 3
- [41] Linfeng Qi, Zhaoyang Jia, Jiahao Li, Bin Li, Houqiang Li, and Yan Lu. Generative latent coding for ultra-low bitrate image and video compression. *IEEE Transactions on Circuits and Systems for Video Technology*, 35(10):10500–10515, 2025. 3, 5
- [42] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020. 4
- [43] 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 (CVPR)*, pages 10684–10695, June 2022. 2
- [44] Juan Song, Lijie Yang, and Mingtao Feng. Extremely low-bitrate image compression semantically disentangled by Imms from a human perception perspective, 2025. 3
- [45] Suramya Tomar. Converting video formats with ffmpeg. *Linux J.*, 2006(146):10, June 2006. 1, 5
- [46] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges, 2019. 5
- [47] Rui Wan, Qi Zheng, and Yibo Fan. M3-cvc: Controllable video compression with multimodal generative models. In *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5, 2025. 1, 3
- [48] Heng Wang, Ioannis Katsavounidis, Jiantong Zhou, Jonghoon Park, Shawmin Lei, Xin Zhou, Man-On Pun, Xin Jin, Ronggang Wang, Xin Wang, et al. Mcl-jcv: A jnd-based h.264/avc video quality assessment dataset. *IEEE International Conference on Image Processing (ICIP)*, pages 1509–1513, 2016. 5
- [49] Xijun Wang, Xin Li, Bingchen Li, and Zhibo Chen. Liftvsr: Lifting image diffusion to video super-resolution via hybrid temporal modeling with only 4×rtx 4090s, 2025. 2
- [50] Zhitao Wang, Hengyu Man, Wenrui Li, Xingtao Wang, Xiaopeng Fan, and Debin Zhao. T-gvc: Trajectory-guided generative video coding at ultra-low bitrates, 2025. 1, 3, 5
- [51] Haoning Wu, Erli Zhang, Liang Liao, Chaofeng Chen, Jingwen Hou, Annan Wang, Wenxiu Sun, Qiong Yan, and Weisi Lin. Exploring video quality assessment on user generated contents from aesthetic and technical perspectives. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20144–20154, October 2023. 5
- [52] Ren Yang, Radu Timofte, and Luc Van Gool. Perceptual learned video compression with recurrent conditional gan, 2022. 5
- [53] Zhendong Yang, Ailing Zeng, Chun Yuan, and Yu Li. Effective whole-body pose estimation with two-stages distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 4210–4220, October 2023. 2, 4
- [54] Fangqiu Yi, Jingyu Xu, Jiawei Shao, Chi Zhang, and Xuelong Li. Conditional video generation for high-efficiency video compression, 2025. 1, 3, 4, 5
- [55] Hang Yin, Li Qiao, Yu Ma, Shuo Sun, Kan Li, Zhen Gao, and Dusit Niyato. Generative video semantic communication via

- multimodal semantic fusion with large model, 2025. [3](#)
- [56] Zongsheng Yue, Jianyi Wang, and Chen Change Loy. Resshift: Efficient diffusion model for image super-resolution by residual shifting. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 13294–13307. Curran Associates, Inc., 2023. [2](#)
- [57] Maojun Zhang, Haotian Wu, Guangxu Zhu, Richeng Jin, Xiaoming Chen, and Deniz Gündüz. Semantics-guided diffusion for deep joint source-channel coding in wireless image transmission, 2025. [1](#), [3](#)
- [58] Pingping Zhang, Jinlong Li, Kecheng Chen, Meng Wang, Long Xu, Haoliang Li, Nicu Sebe, Sam Kwong, and Shiqi Wang. When video coding meets multimodal large language models: A unified paradigm for video coding, 2025. [4](#)
- [59] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [5](#)
- [60] Yuhong Zhang, Hengsheng Zhang, Zhengxue Cheng, Rong Xie, Li Song, and Wenjun Zhang. Ssp-ir: Semantic and structure priors for diffusion-based realistic image restoration. *IEEE Transactions on Circuits and Systems for Video Technology*, 35(7):6259–6272, 2025. [2](#)
- [61] Zihao Zhang, Haoran Chen, Haoyu Zhao, Guansong Lu, Yanwei Fu, Hang Xu, and Zuxuan Wu. Eden: Enhanced diffusion for high-quality large-motion video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2105–2115, June 2025. [2](#)