

T-GVC: Trajectory-Guided Generative Video Coding at Ultra-Low Bitrates

Zhitao Wang¹, Hengyu Man^{1*}, Wenrui Li¹, Xingtao Wang¹, Xiaopeng Fan^{1,2,3}, Debin Zhao¹

¹Harbin Institute of Technology

²Peng Cheng Laboratory

³Harbin Institute of Technology Suzhou Research Institute

zhitao.wang.hit@outlook.com, manhengyu@hotmail.com, liwr618@163.com,
{xtwang, fxp, dbzhao}@hit.edu.cn

Abstract

Recent advances in video generation techniques have given rise to an emerging paradigm of generative video coding for Ultra-Low Bitrate (ULB) scenarios by leveraging powerful generative priors. However, most existing methods are limited by domain specificity (e.g., facial or human videos) or excessive dependence on high-level text guidance, which tend to inadequately capture fine-grained motion details, leading to unrealistic or incoherent reconstructions. To address these challenges, we propose **Trajectory-Guided Generative Video Coding** (dubbed T-GVC), a novel framework that bridges low-level motion tracking with high-level semantic understanding. T-GVC features a semantic-aware sparse motion sampling pipeline that extracts pixel-wise motion as sparse trajectory points based on their semantic importance, significantly reducing the bitrate while preserving critical temporal semantic information. In addition, by integrating trajectory-aligned loss constraints into diffusion processes, we introduce a training-free guidance mechanism in latent space to ensure physically plausible motion patterns without sacrificing the inherent capabilities of generative models. Experimental results demonstrate that T-GVC outperforms both traditional and neural video codecs under ULB conditions. Furthermore, additional experiments confirm that our framework achieves more precise motion control than existing text-guided methods, paving the way for a novel direction of generative video coding guided by geometric motion modeling.

Extended version — <https://chigland.github.io/T-GVC/>

1 Introduction

One of the core challenges in video coding lies in effectively modeling inter-frame dependencies to reduce temporal redundancy and enhance coding efficiency. While traditional video coding standards (e.g., H.266/VVC (Bross et al. 2021) and AVS3 (Zhang et al. 2019)) have achieved remarkable progress, their reliance on handcrafted motion compensation and transform-quantization modules exhibits limitations in modeling non-rigid motions and nonlocal spatio-temporal dependencies, particularly under Ultra-Low Bitrate (ULB) scenarios with constrained bandwidth. Benefiting from end-to-end optimization strategy, recently emerging deep learning-based video coding approaches (Li, Li,

and Lu 2024; Jiang et al. 2024; Li et al. 2024; Man et al. 2025; Tang et al. 2025) have demonstrated promising potential to surpass conventional schemes. However, these methods primarily emphasize pixel-level signal fidelity and suffer from semantic information loss in ULB scenarios due to limited temporal context information. Although recent attempts (Yang, Timofte, and Van Gool 2022; Du, Liu, and Ling 2024; Du et al. 2022) have leveraged GAN (Goodfellow et al. 2020) to improve the perceptual quality, they still adopt similar frameworks as (Li, Li, and Lu 2023, 2024; Yang et al. 2020) to model temporal contextual information, which constraints their ability to achieve lower bitrates. Consequently, preserving critical semantic information while enhancing perceptual quality has become a pressing challenge for ULB video compression.

Recent breakthroughs in video generation technology have provided new feasibility for video coding under ULB conditions, catalyzing the emergence of the concept of “generative video coding” (Chen et al. 2024b). Unlike generation tasks in computer vision, generative video coding primarily aims to achieve content-faithful reconstruction under strict bitrate constraints by leveraging the strong priors of generative models as well as compact spatio-temporal guidance. Despite these theoretical advantages, most existing generative video coding schemes are tailored to specific types of videos, such as human face video (Chen et al. 2024a), human body video (Wang et al. 2023; Chen et al. 2025), and small motion video (Yin et al. 2024), due to the limited capabilities of earlier-stage generative models.

The rapid advancement of video diffusion models (VDMs) (Xing et al. 2024b) has significantly intensified research interest in extending generative video coding to natural video content. Specifically, some recent works (Zhang et al. 2024a; Wan, Zheng, and Fan 2025) attempt to model spatio-temporal information based on texts and keyframes, showing the potential to synthesize motion sequences conditioned on sparse and high-level information derived from the original video. Nevertheless, these methods still struggle to ensure temporal consistency and visual stability when applied to complex real-world scenarios. The generated outputs may exhibit unexpected artifacts, such as unnatural motion patterns or missing dynamic objects, compromising the temporal semantic fidelity with respect to the original video. To enable more precise natural video generation, an intuitive

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

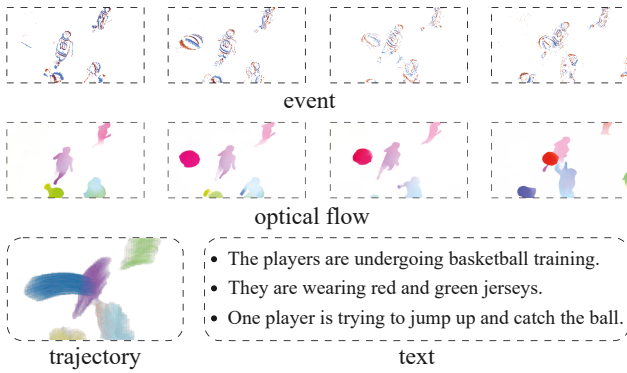


Figure 1: Examples of motion representation for video (event, optical flow, trajectory and text).

idea is to take advantage of low-level motion information, as illustrated in Figure 1. Unfortunately, employing dense motion representation incurs a higher bitrate overhead for VDMs to synthesize high-fidelity videos. **It is crucial to develop an effective strategy to extract compact motion guidance that preserves key temporal semantic information from the original video.**

Another important aspect concerns the conditional guidance mechanisms employed in the generative model. Most existing controllable video generation methods utilize ‘classifier-free guidance’ (Ho and Salimans 2022) to fine-tune pre-trained video diffusion models, such as Stable Video Diffusion (SVD) (Blattmann et al. 2023) and VideoCraft (Chen et al. 2023a, 2024c), or adopt a ControlNet-like adapter (Zhang, Rao, and Agrawala 2023; Niu et al. 2024; Liu et al. 2025) to steer the generation process according to user specifications. However, the loss of critical structural and textural details in compressed keyframes severely impairs the model’s ability to reconstruct motion-aligned videos, particularly under ULB conditions. Moreover, these pre-trained generative models suffer from limited adaptability in the context of video coding tasks, as the compact temporal guidance (e.g., motion trajectory) may deviate substantially from the original conditioning domains. **How to effectively leverage compact guidance information to control the generative model in reconstructing high-fidelity and motion-aligned videos also requires further exploration.**

In response to the critical challenge of generative video coding under ULB conditions, we propose a Trajectory-Guided Generative Video Coding framework. To balance the trade-off between temporal semantic information preservation and bitrate saving, a semantic-aware sparse motion sampling pipeline is first proposed to bridge low-level motion tracking with high-level semantic understanding. Specifically, we track the pixel-wise motion on a pre-defined grid and classify them into distinct motion instances according to motion patterns. To further improve coding efficiency, a sparse sampling of motion instances is performed based on their semantic importance, representing motion as a set of sparse trajectory points. The variations in trajectory points

capture diverse motion patterns (such as translation, occlusion, deformation, etc.), which preserves critical temporal semantic information of the original video while significantly reducing the bitrate. To guide the diffusion-based generative model in reconstructing motion-aligned videos without compromising its inherent generative capabilities, we propose a training-free guidance approach. In contrast to existing methods that constrain on intermediate feature maps, our approach directly imposes constraints on the latent space of diffusion model via a lightweight yet effective guidance function, ensuring that the overall trajectories of target motion instances align with real-world motion paths.

To this end, our principal contributions are summarized as follows:

- We propose a semantic-aware sparse motion sampling pipeline tailored for generative video coding, which bridges low-level motion tracking with high-level semantic understanding by encoding motion as sparse trajectory points, significantly reducing bitrate while preserving temporal semantics.
- A training-free latent space guidance mechanism is designed to enforce trajectory alignment via a lightweight guidance function without compromising the inherent capabilities of the generative model.
- Extensive experiments demonstrate that T-GVC outperforms both traditional codecs and state-of-the-art deep learning-based methods under ULB conditions, establishing a novel paradigm for efficient semantic-aware video coding in resource-constrained scenarios.

2 Related Work

2.1 Generative Video Coding

To ensure satisfactory perceptual quality, generative video coding has been explored in Ultra-Low Bitrate (ULB) scenarios. Early methods in this domain mainly focused on specific content types, such as facial videos (Chen et al. 2024a) or human-centric videos (Wang et al. 2022, 2023). For natural videos, some approaches (Yang, Timofte, and Van Gool 2022; Du et al. 2022; Du, Liu, and Ling 2024) employed GANs to enhance perceptual quality under low bitrate conditions. A more recent framework based on VQ-VAE (Qi et al. 2025) further performed transform coding in the latent space to achieve high realism and high fidelity at ULB. However, constrained by the inherent architectures of conventional frameworks (Yang et al. 2020; Li, Li, and Lu 2023, 2024), these methods are difficult to generalize to even lower bitrate scenarios.

Inspired by Cross-Modal Compression (CMC) in the field of image compression (Li et al. 2021; Zhang et al. 2023; Gao et al. 2023, 2024), Zhang et al. designed Cross-Modal Video Compression (CMVC) to explore the potential of applying multimodal large language models for ULB video compression. In CMVC, a video is segmented into multiple semantically coherent clips, each of which is decomposed into content (a keyframe) and motion (text descriptions). On the decoder side, keyframes are reconstructed first, followed by the generation of intermediate frames guided by the extracted ‘motion’ information. Similarly, Wan, Zheng, and

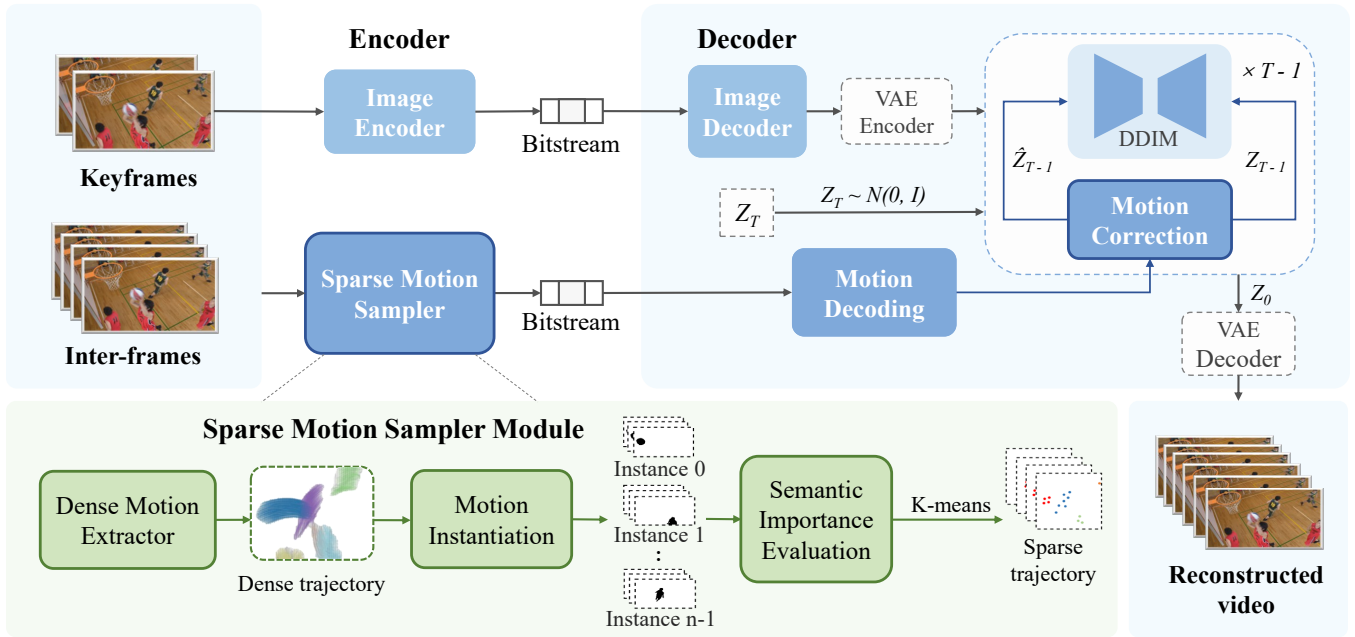


Figure 2: Overview of our T-GVC framework. On the encoder side, each pair of keyframes and corresponding inter-frames are fed into proposed sparse motion sampler to extract motion trajectories. Subsequently, the keyframes and trajectories are encoded into compact bitstreams. On the decoder side, each decoded keyframe pair is encoded into latent features via VAE encoder. These latent features, combined with zero-initialized latent features, form a latent sequence and concatenated with the initial latent noises as input of VDM. The decoded sparse motion trajectories act as guidance conditions during the inference process to correct the motion in latent space. Ultimately, the clean output is decoded by VAE.

Fan utilized a text-guided video diffusion model for video reconstruction, incorporating a more elaborate keyframe selection strategy. While these approaches are able to preserve essential semantic contents at ULB (Li et al. 2025), they still suffer from issues with temporal consistency when handling complex motion patterns. Overall, despite significant progress in generative video coding for specific content types, extending these techniques to natural video, particularly in dynamic and motion-rich scenes, remains an open and challenging topic.

2.2 Motion-Conditioned Diffusion Model

Text-conditioned video diffusion models (VDMs) have demonstrated remarkable capabilities in synthesizing visually striking content. However, they still struggle to achieve precise outcome control, frequently generating unrealistic sequences that violate physical plausibility. To improve temporal coherence and realism, recent efforts introduced motion guidance into VDMs. Early works like MCDiff (Chen et al. 2023b) and DragNUWA (Yin et al. 2023) incorporated sparse motions and multimodal inputs, while later models such as MotionCtrl (Wang et al. 2024) and Tora (Zhang et al. 2024b, 2025b) introduced finer-grained control via explicit trajectory modeling. However, they typically require training new diffusion models or motion control modules from scratch. To mitigate the trade-off between quality and computation cost, more recent studies (Xiao et al. 2024; Zhang et al. 2025a) have explored the use of classifier guidance

(Dhariwal and Nichol 2021) to control the video generation process in a more efficient manner, revealing the presence of intrinsic motion-aware features within the latent representations of VDMs. These approaches have shown the flexibility and effectiveness of the training-free video motion control framework, which provided valuable inspiration for our work.

3 Methodology

The overall structure of the proposed T-GVC framework is shown in Figure 2, which follows the common paradigm of generative video coding. In the remainder of this section, we will describe the entire coding pipeline of T-GVC and elaborate on each module with key symbols summarized in Table 1.

3.1 Keyframe Selection and Compression

Following (Zhang et al. 2024a), we segment the source video into clips and select keyframes based on minimal pairwise semantic similarity, measured via a CLIP-based criteria (Radford et al. 2021). To robustly handle scene transitions, we further incorporate a semantic-aware scene cut detection mechanism based on inter-frame similarity analysis. Specifically, when the semantic similarity between two consecutive frames falls below a pre-defined threshold, both frames are designated as keyframes to ensure seamless scene transitions. Therefore, more than one keyframes may be selected within each video clip. Each pair of consecutive keyframes

Symbol	Description
F_i	The i -th frame in the video clip.
\mathcal{T}	Set of motion trajectories.
\mathcal{T}^i	The i -th motion instance (cluster).
$T_i \in \mathcal{T}$	The i -th trajectory in the trajectory set.
(x_j^i, y_j^i)	Spatial coordinates of T_i in F_j .
$V_j^i \in \{0, 1\}$	Visibility of T_i at F_j .
M_i	Motion mask of \mathcal{T}^i .
z_t^i	The i -th latent feature at time step t .
τ_j^i	Sparse trajectory points of \mathcal{T}^j in F_i .
L	Number of frames in the video clip.

Table 1: Summary of symbols used in T-GVC.

$\{K_i, K_{i+1}\}$, along with the associated intermediate frames, constitutes a new clip $\{F_0, F_1, \dots, F_{L-1}\}$. The clip structure serves as the fundamental processing unit for subsequent motion sampling and semantic-consistent video reconstruction.

Notably, since keyframes encapsulate critical semantic information, their reconstruction on the decoder side must maintain high perceptual quality even under stringent bitrate constraints, thereby preserving essential textural details and structural features to facilitate subsequent intermediate frame reconstruction. As this work primarily focuses on inter-frame coding, we adopt the off-the-shelf MS-ILLM (Muckley et al. 2023) for keyframe compression, owing to its demonstrated ability to preserve superior perceptual fidelity under aggressive bitrate constraints.

3.2 Semantic-Aware Sparse Motion Sampling

To bridge low-level motion tracking with high-level semantic understanding, we propose a semantic-aware sparse motion sampling framework that quantifies the semantic importance of different motion trajectories in the original video. The pipeline comprises three stages:

Dense Trajectory Extraction. Given two keyframes $\{K_i, K_{i+1}\}$ and their intermediate frames, we design a bidirectional tracking scheme to extract dense motion trajectories with improved temporal consistency. For each clip $\{F_0, F_1, \dots, F_{L-1}\}$ with temporal length L , we initialize isometric grids (grid size = $N_{grid} \times N_{grid}$) in the first frame and employ Co-Tracker (Karaev et al. 2024) to trace pixel-wise displacements originating from the grid vertices. The generated forward trajectory sequences over L frames are represented as temporal chains of 3D coordinate points:

$$\mathcal{T}_{fwd} = \left\{ T_i^{fwd} \mid T_i^{fwd} = (x_t^i, y_t^i, V_t^i)_{t=0}^{L-1} \right\} \quad (1)$$

To mitigate trajectory loss caused by occlusions or rotation, we reverse the temporal order of each clip and extract the corresponding backward trajectory sequences \mathcal{T}_{bwd} using the same procedure. Subsequently, the backward trajectories are temporally realigned and concatenated with their forward counterparts to synthesize the final dense trajectories, which holistically preserve the motion characteristics between keyframes:

$$\mathcal{T}_{dense} = \mathcal{T}_{fwd} \cup Flip(\mathcal{T}_{bwd}) \quad (2)$$

$Flip(\cdot)$ denotes the reverse operation in the time dimension.

Motion Instantiation. To differentiate between distinct motion patterns, the raw trajectories are clustered using HDBSCAN (McInnes et al. 2017), a hierarchical density-based algorithm particularly effective in handling variable cluster densities and suppressing noise, which is critical for processing real-world videos with irregular motions. In this stage, T_i is first transformed into 2D trajectories by excluding the points where $V_t^i = 0$. Then, we formulate each T_i as a spatio-temporal feature vector according to their coordinates. By clustering these features, the corresponding trajectories are adaptively categorized into different motion instances $\mathcal{T}^0, \mathcal{T}^1, \dots, \mathcal{T}^{n-1}$ and form the corresponding motion masks, as illustrated in Figure 2.

Semantic Importance Evaluation. The extracted dense motion trajectories inherently encapsulate the spatio-temporal dynamics of the source video, enabling the decoder to reconstruct inter-frame semantic relationships through diffusion-based generation (Section 3.3). However, noisy regions may be mistakenly identified as motion instances due to inaccuracies in motion clustering. Additionally, naively encoding these raw trajectories introduces significant bitrate overhead. To mitigate the influence of the accuracy of motion clustering and reduce inter-frame bitrate while preserving semantically critical motion information, we propose a trajectory sampling strategy driven by semantic loss awareness.

For each motion instance \mathcal{T}^i , we calculate the semantic importance score of motion S_{inter}^i by multiplying the intra-semantic importance score S_{intra}^i with the length of the trajectory. The intra-semantic importance score is defined as:

$$S_{intra}^i = \sum_{j=0}^{L-2} \left\| D(F_j, F_{j+1}) - D[P(F_j, M^i), P(F_{j+1}, M^i)] \right\|_1 \quad (3)$$

where $D(\cdot)$ denotes the CLIP-based similarity, and $P(F_j, M^i)$ represents that the region covered by the motion mask M^i in F_j is replaced by the surrounding pixels to avoid feature mismatch for CLIP.

We select the motion whose inter-frame semantic importance exceeds an empirically chosen threshold and extract sparse trajectories τ^i through k-means clustering, where the keypoint quantity K for each instance is determined by the semantic importance score S_{inter}^i and the quantity of trajectories N^i :

$$K^i = \min\left\{ (\alpha \cdot S_{inter}^i + \beta \cdot \frac{N^i}{N_{total}}) \cdot K_{max}, K_{max} \right\} \quad (4)$$

N_{total} represents the grid quantity while hyperparameters α and β control the weighting ratios of different metrics.

Considering that the video resolution is downsampled by a factor of 8 after projected into the diffusion latent space on the decoder side, the trajectory coordinates are quantized into the same numerical range to ensure alignment with the

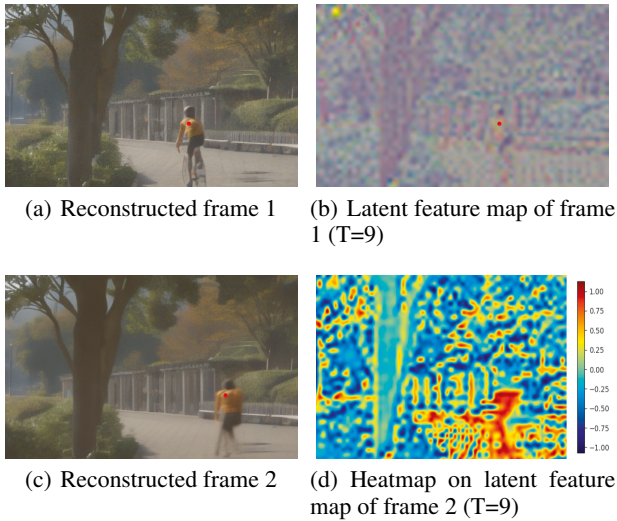


Figure 3: Illustration of the semantic correlation of latent features (upscaled to the same resolution as the original frame) in the same trajectory across two reconstructed frames. Given the red trajectory point in (b), we plot (d) according to the similarity between the feature on the point and the latent features of frame 2.

compressed latent representation. Finally, the initial trajectory data from the first frame, along with the coordinate displacements between adjacent frames are losslessly encoded into a compact bitstream, which would be decoded to guide the reconstruction of the inter-frames on the decoder side.

3.3 Video Reconstruction

Trajectory-Based Motion Guidance. Inspired by (Xiao et al. 2024; Zhang et al. 2025a), we utilize motion trajectories to guide the generation process in a training-free manner by imposing trajectory-aware constraints during the denoising diffusion process. In (Xiao et al. 2024; Zhang et al. 2025a), the trajectories are extracted from the feature maps of U-Net blocks in VDM. However, these approaches incur substantial computational overhead during inference, particularly for generative models with a large number of parameters. Moreover, since the motion of generated results is globally controlled, it is difficult to guide the model’s output with fine-grained precision. In contrast, T-GVC directly imposes constraint on the latent noise, leveraging the insight that the latent space in video diffusion contains crucial structural information and semantic correlation. As shown in Figure 3, the region with the highest similarity in Figure 3(a) corresponds to the position of the red point in Figure 3(c), which belongs to the same trajectory as the red point in Figure 3(a) or Figure 3(b).

For each clip of the target video with length L , the pipeline on the decoder side begins with decoding compressed keyframes and associated motion trajectories. After that, the keyframes $\{\hat{K}_i, \hat{K}_{i+1}\}$ are projected into latent space via VAE (Kingma, Welling et al. 2013) encoder and concatenated with the initial noisy latent $z_T =$

$\{z_T^0, z_T^1, \dots, z_T^{L-1}\}$ as conditions, providing the key textual and structural features of the original video.

Let $F(z_t^i, \tau_i^j)$ denote the feature of the i -th latent z_t^i at time step t within the region covered by the sparse trajectories τ_i^j of the motion instance \mathcal{T}^j . To reconstruct the temporal semantic information in the guidance of trajectories, we define a loss function to measure how well the trajectories of z_t align with the trajectory guidance, thereby constraining the denoising process:

$$\mathcal{L}_m = \sum_{i=1}^{L-2} \sum_{j=0}^{n-1} \left[\alpha_i \cdot S_0(z_t^i, \tau_i^j) + \beta_i \cdot S_{L-1}(z_t^i, \tau_i^j) \right] \quad (5)$$

where $S_k(z_t^i, \tau_i^j) = \|F(z_t^k, \tau_i^j) - F(z_t^i, \tau_i^j)\|_1$, captures the non-local similarity of the latent sequence, while α_i and β_i control the level of the similarity.

We can further update the output ϵ_θ of the model in each time step via the gradient of \mathcal{L}_m :

$$\hat{\epsilon}_\theta(z_t, t) = \epsilon_\theta(z_t, t) + s(t) \cdot \nabla_{z_t} \mathcal{L}_m(f(z_t)) \quad (6)$$

where $f(\cdot)$ is the map from noise latent z_t to the clean latent z_0 .

Video Generation with Variable Length. Most existing VDMs are trained to generate videos of fixed length (e.g., 16 frames). To this end, we further design a variable-length generation scheme to enhance the flexibility of pre-trained VDMs. Specifically, for each video clip with length $L \leq 16$, we first interpolate the sparse trajectories τ to $\hat{\tau}$ with length $\hat{L} = 16$ and assign corresponding positional markers to the original trajectories. The denoised latent frames at these marked positions are then selected and concatenated to form the final video. For each video clip with length $L > 16$, we perform a dual-stage generation strategy where the trajectories are divided into two segments, with each segment guiding one generation stage separately.

4 Experiments

4.1 Experimental Setup

Dataset and Metrics. Following (Wan, Zheng, and Fan 2025), we select HEVC Class B, C (Bossen 2010) as well as UVG (Mercat, Viitanen, and Vanne 2020) and MCL-JCV (Wang et al. 2016) as the test dataset to evaluate the rate-distortion (R-D) performance. All test videos are resized to 512×320 and consist of 96 frames with the original frame rate. For the calculation of R-D metric, the compression rate is quantified in bits per pixel (bpp), while the distortion is measured using LPIPS (Zhang et al. 2018) for perceptual quality and CLIP-SIM (Radford et al. 2021) for semantic similarity. Lower LPIPS and higher CLIP-SIM at the same bitrate represent better coding performance.

Implementation Details. We select the pre-trained DynamiCrafter (Xing et al. 2024a) video diffusion model as our inter-frame decoder due to its superior quality for open-domain image animation. We remove text prompts from the original model and solely utilize trajectories and keyframes

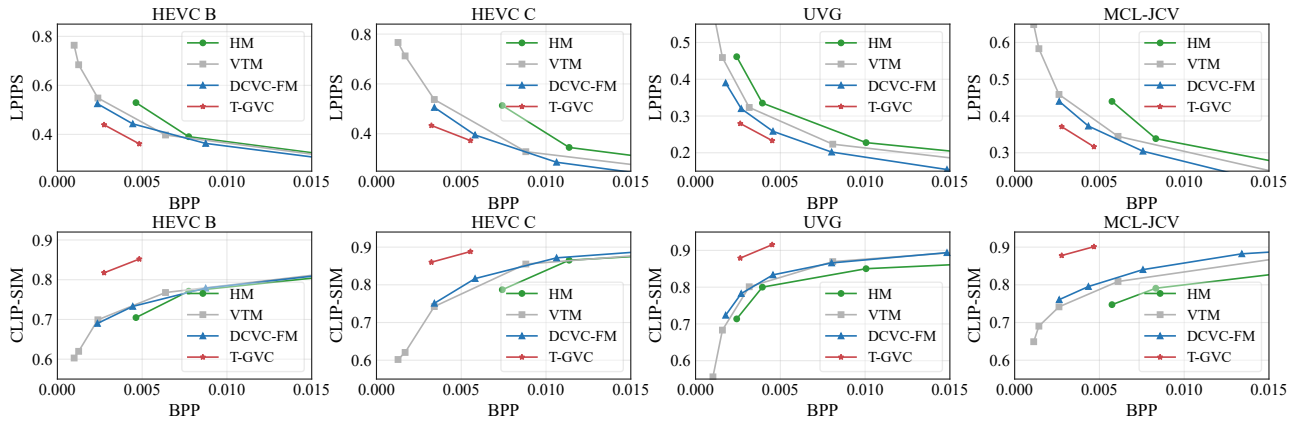


Figure 4: The R-D performance comparison results for HEVC Class B, Class C, UVG and MCL-JCV datasets.




Figure 5: Visual quality comparison: ground truth, DCVC-FM, VTM and proposed T-GVC (top to bottom). The reconstructed frames of our framework demonstrates higher perceptual quality at similar bitrates.

to guide the generation of inter-frames. The number of DDIM denoising steps for each video and the classifier guide scale $s(t)$ are set to 10 and $30 \cdot \sqrt{1 - \alpha_t}$, respectively. For keyframe compression, we adopt MS-ILLM (Muckley et al. 2023) under ‘quality 1’ and ‘quality 2’ settings. In addition, given that the video resolution is downsampled to 64×40 by a factor of 8 after projected into latent space, the grid size for sparse motion sampling is configured as 64 to preserve structural consistency. All experiments are conducted on a single NVIDIA RTX 3090 GPU.

4.2 Quantitative and Qualitative Results

We compare T-GVC with traditional codecs (H.265 (Sullivan et al. 2012) and H.266 (Bross et al. 2021)) and one of the most competitive open-source neural video codec (DCVC-FM (Li, Li, and Lu 2024)) with a wide bitrate range to evaluate its coding performance at ULB. The comparison results with other generative video coding schemes are provided and discussed in Section 4.3. For H.265, we use HM-18.0 with QP = 51, 45 and 39. For H.266, we use VTM-23.7 with QP = 63, 57, 51, 45 and 39. For DCVC-FM, we set the q indexes to 0, 8, 16, 24 and 32, which control the compression level. All codecs above are tested under Low Delay P configuration with intra-period=-1.

The R-D curves for the test video classes are presented in

Models	Settings	BD-rate (%)
Text	W/o motion bitrate	-4.45
	W/ motion bitrate	31.92
Trajectory	W/o motion bitrate	-16.84
	W/ motion bitrate	-3.06
Text+Trajectory	W/o motion bitrate	-13.33
	W/ motion bitrate	34.97

Table 2: Ablation study on the guidance mechanism.

Figure 4, where the bitrate of our T-GVC is controlled by adjusting the quality of keyframes. As observed, T-GVC outperforms both traditional codecs and the neural video codec in terms of perceptual reconstruction quality and semantic fidelity at ULB conditions. Moreover, benefiting from powerful generative priors, T-GVC achieves lower bitrate points (below 0.005 bpp) than GAN- or VQ-VAE-based methods (Yang, Timofte, and Van Gool 2022; Qi et al. 2025). Figure 5 shows the visual results for the sequence ‘BasketballDrill_832x480_50’ from the HEVC Class C dataset. To ensure a fair comparison, reconstructed sequences from different methods are selected at comparable bitrates. It is observed that the reconstructed videos generated by T-GVC demonstrate superior fidelity in the background regions and key foreground objects (e.g., basketball and the players in red). In contrast, VTM exhibits noticeable blocking artifacts while DCVC-FM suffers from significant high-frequency detail loss, leading to degraded texture quality.

It is worth noting that although some texture and structural details in the reconstructed frames might be partially missing or inaccurate (e.g., the direction of the ball and the appearance of players in green), the overall visual quality remains satisfactory for the ULB scenarios. Under ULB conditions, we primarily focus on perceptual quality within a given bitrate range. While key semantic information (e.g., the player jumping to catch the ball) is preserved, certain visual artifacts—such as the deblurring effects or minor spatial deviations—are acceptable.



Figure 6: Visual quality comparison for ablation study on the guidance mechanism.

4.3 Ablation Study

To further verify the rationality of our design, we conduct ablation studies on two core components of the T-GVC framework: the guidance mechanism and the motion sampler module. Video sequences characterized by large motions, including ‘ParkScene’, ‘BasketballDrill’, ‘PartyScene’, ‘videoSRC02’ and ‘videoSRC05’ from Class B, Class C and MCL-JCV, are selected for evaluation. The coding efficiency is assessed using the Bjøntegaard Delta rate (BD-rate) (Bjontegaard 2001), where PSNR is replaced by $1 - \text{LPIPS}$. Negative BD-rate values indicate bitrate saving for equivalent perceptual quality. We apply DynamiCrafter with empty prompts (i.e., only guided by keyframes) as the comparison anchor in our experiments. The reported BD-rate results are calculated by averaging all test sequences.

Ablation Study on the Guidance Mechanism. In prior generative video coding methods (Zhang et al. 2024a; Wan, Zheng, and Fan 2025), text is commonly adopted as motion guidance. To demonstrate the superiority of our guidance mechanism, we implement three variants of generative frameworks based on DynamiCrafter: one guided by text (denoted as Text), one guided by trajectories (denoted as Trajectory) and the other by both text and trajectories (denoted as Text+Trajectory). Trajectories are sampled through our proposed sparse motion sampler, while text descriptions are generated from the source video using Hunyuan-Large (Sun et al. 2024) without limitation on length. The comparison results are presented in Table 2, where “W/o motion bitrate” refer to excluding the bitrate consumed by motion-related information when calculating BPP, while “W/ motion bitrate” indicate that this bitrate is included. It is observed that the trajectory-guided approach consistently achieves better coding performance under both evaluation conditions. In particular, the perceptual quality improvement brought by trajectory guidance makes up for the additional bitrate overhead, offering a distinct advantage over text-based guidance. When overhead is not considered, the trajectory-guided method still outperforms the text-guided method. Combining text and trajectory guidance may enhance texture on simple-motion sequences. Some visual examples in Figure 6 further demonstrate that trajectory guid-

Models	Settings	BD-rate (%)
Dense	W/o motion bitrate	-17.88
	W/ motion bitrate	188.77
Random	W/o motion bitrate	-7.95
	W/ motion bitrate	15.05
Sparse	W/o motion bitrate	-16.84
	W/ motion bitrate	-3.06

Table 3: Ablation study on the sparse motion sampler.

ance can perform more precise motion control, especially in dynamic areas. Text guidance can improve perceptual quality to some extent and combining text and trajectory guidance may enhance texture on simple-motion sequences. However, in more complex cases, text prompts may introduce artifacts that conflict with trajectory-aligned motion, leading to degraded results. For instance, as shown in the first column of Figure 6, the basketball appears behind the net, which violates the physical plausibility.

Ablation Study on the Motion Sampler Module. To validate the effectiveness of the semantic-aware motion sampling pipeline proposed in T-GVC, we further conduct an ablation study on the motion sampler module. The ablation results are presented in Table 3. Specifically, “Sparse” refers to motion trajectories extracted with our semantic-aware motion sampling pipeline. “Dense” denotes applying dense trajectories as guidance without sampling. “Random” corresponds to randomly selecting sparse trajectories with k keypoints ($K < K_{max}$) for each motion instance. We set the maximum keypoint quantity K_{max} as 15 for the optimal R-D performance. Although dense trajectories can provide more precise motion representation for video reconstruction, such excessive bitrate overhead is intolerable for the video coding task. In contrast, our motion sampling pipeline achieves significant bitrate reduction in motion modeling, while preserving perceptual quality comparable to that of dense trajectory guidance. Compared with a random sampling strategy, our approach thoroughly considers the impact of each motion instance on the video semantics, effectively reducing the bitrate while retaining key motion information.

5 Conclusion

In this paper, we propose T-GVC, a novel Trajectory-Guided Generative Video Coding framework designed for ULB scenarios. Unlike existing methods that rely on domain-specific priors or high-level text guidance, T-GVC introduces a semantic-aware sparse motion sampling strategy, extracting trajectories based on semantic importance to retain key temporal information with minimal bitrate overhead. To further enhance motion realism, a training-free guidance mechanism is integrated into the diffusion process, enabling physically plausible reconstructions. Experimental results demonstrate that T-GVC significantly outperforms traditional codecs and state-of-the-art generative methods, especially in terms of motion control precision.

Acknowledgments

This work was supported in part by the National Key R&D Program of China (2023YFA1008500), the National Natural Science Foundation of China (NSFC) under grants U22B2035 and 62502116, and China Post-Doctoral Science Foundation under Grant 2025M774315.

References

- Bjontegaard, G. 2001. Calculation of average PSNR differences between RD-curves. *ITU SG16 Doc. VCEG-M33*.
- Blattmann, A.; Dockhorn, T.; Kulal, S.; Mendelevitch, D.; Kilian, M.; Lorenz, D.; Levi, Y.; English, Z.; Voleti, V.; Letts, A.; et al. 2023. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*.
- Bossen, F. 2010. Common test conditions and software reference configurations. In *3rd. JCT-VC Meeting, Guangzhou, CN, October 2010*.
- Bross, B.; Wang, Y.-K.; Ye, Y.; Liu, S.; Chen, J.; Sullivan, G. J.; and Ohm, J.-R. 2021. Overview of the versatile video coding (VVC) standard and its applications. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(10): 3736–3764.
- Chen, B.; Chen, J.; Wang, S.; and Ye, Y. 2024a. Generative face video coding techniques and standardization efforts: A review. In *2024 Data Compression Conference (DCC)*, 103–112. IEEE.
- Chen, B.; Liao, R.-L.; Chen, J.; and Ye, Y. 2025. Rethinking Generative Human Video Coding with Implicit Motion Transformation. *arXiv preprint arXiv:2506.10453*.
- Chen, B.; Yin, S.; Chen, P.; Wang, S.; and Ye, Y. 2024b. Generative visual compression: A review. In *2024 IEEE International Conference on Image Processing (ICIP)*, 3709–3715. IEEE.
- Chen, H.; Xia, M.; He, Y.; Zhang, Y.; Cun, X.; Yang, S.; Xing, J.; Liu, Y.; Chen, Q.; Wang, X.; et al. 2023a. Videocrafter1: Open diffusion models for high-quality video generation. *arXiv preprint arXiv:2310.19512*.
- Chen, H.; Zhang, Y.; Cun, X.; Xia, M.; Wang, X.; Weng, C.; and Shan, Y. 2024c. Videocrafter2: Overcoming data limitations for high-quality video diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7310–7320.
- Chen, T.-S.; Lin, C. H.; Tseng, H.-Y.; Lin, T.-Y.; and Yang, M.-H. 2023b. Motion-conditioned diffusion model for controllable video synthesis. *arXiv preprint arXiv:2304.14404*.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Du, P.; Liu, Y.; and Ling, N. 2024. Cgvc-t: Contextual generative video compression with transformers. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*.
- Du, P.; Liu, Y.; Ling, N.; Ren, Y.; and Liu, L. 2022. Generative video compression with a transformer-based discriminator. In *2022 Picture Coding Symposium (PCS)*, 349–353. IEEE.
- Gao, J.; Jia, C.; Huang, Z.; Wang, S.; Ma, S.; and Gao, W. 2024. Rate-distortion optimized cross modal compression with multiple domains. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(8): 6978–6992.
- Gao, J.; Li, J.; Jia, C.; Wang, S.; Ma, S.; and Gao, W. 2023. Cross modal compression with variable rate prompt. *IEEE Transactions on Multimedia*, 26: 3444–3456.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Jiang, W.; Li, J.; Zhang, K.; and Zhang, L. 2024. ECVC: Exploiting Non-Local Correlations in Multiple Frames for Contextual Video Compression. *arXiv preprint arXiv:2410.09706*.
- Karaev, N.; Rocco, I.; Graham, B.; Neverova, N.; Vedaldi, A.; and Rupperecht, C. 2024. Cotracker: It is better to track together. In *European Conference on Computer Vision*, 18–35. Springer.
- Kingma, D. P.; Welling, M.; et al. 2013. Auto-encoding variational bayes.
- Li, J.; Jia, C.; Zhang, X.; Ma, S.; and Gao, W. 2021. Cross modal compression: Towards human-comprehensible semantic compression. In *Proceedings of the 29th ACM international conference on multimedia*, 4230–4238.
- Li, J.; Li, B.; and Lu, Y. 2023. Neural video compression with diverse contexts. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 22616–22626.
- Li, J.; Li, B.; and Lu, Y. 2024. Neural video compression with feature modulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26099–26108.
- Li, W.; Han, W.; Man, H.; Zuo, W.; Fan, X.; and Tian, Y. 2025. Language-Guided Graph Representation Learning for Video Summarization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–17.
- Li, Z.; Yuan, Z.; Li, L.; Liu, D.; Tang, X.; and Wu, F. 2024. Object segmentation-assisted inter prediction for versatile video coding. *IEEE Transactions on Broadcasting*.
- Liu, F.-L.; Fu, H.; Wang, X.; Ye, W.; Wan, P.; Zhang, D.; and Gao, L. 2025. SketchVideo: Sketch-based Video Generation and Editing. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 23379–23390.
- Man, H.; Wang, H.; Lu, R.; Wan, Z.; Fan, X.; and Zhao, D. 2025. Content-Aware Dynamic In-Loop Filter With Adjustable Complexity for VVC Intra Coding. *IEEE Transactions on Circuits and Systems for Video Technology*, 35(6): 6114–6128.
- McInnes, L.; Healy, J.; Astels, S.; et al. 2017. hdbscan: Hierarchical density based clustering. *J. Open Source Softw.*, 2(11): 205.

- Mercat, A.; Viitanen, M.; and Vanne, J. 2020. UVG dataset: 50/120fps 4K sequences for video codec analysis and development. In *Proceedings of the 11th ACM multimedia systems conference*, 297–302.
- Muckley, M. J.; El-Nouby, A.; Ullrich, K.; Jégou, H.; and Verbeek, J. 2023. Improving statistical fidelity for neural image compression with implicit local likelihood models. In *International Conference on Machine Learning*, 25426–25443. PMLR.
- Niu, M.; Cun, X.; Wang, X.; Zhang, Y.; Shan, Y.; and Zheng, Y. 2024. Mofa-video: Controllable image animation via generative motion field adaptations in frozen image-to-video diffusion model. In *European Conference on Computer Vision*, 111–128. Springer.
- Qi, L.; Jia, Z.; Li, J.; Li, B.; Li, H.; and Lu, Y. 2025. Generative latent coding for ultra-low bitrate image and video compression. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Sullivan, G. J.; Ohm, J.-R.; Han, W.-J.; and Wiegand, T. 2012. Overview of the high efficiency video coding (HEVC) standard. *IEEE Transactions on circuits and systems for video technology*, 22(12): 1649–1668.
- Sun, X.; Chen, Y.; Huang, Y.; Xie, R.; Zhu, J.; Zhang, K.; Li, S.; Yang, Z.; Han, J.; Shu, X.; et al. 2024. Hunyuan-large: An open-source moe model with 52 billion activated parameters by tencent. *arXiv preprint arXiv:2411.02265*.
- Tang, C.; Li, Z.; Bian, Y.; Li, L.; and Liu, D. 2025. Neural Video Compression with Context Modulation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 12553–12563.
- Wan, R.; Zheng, Q.; and Fan, Y. 2025. M3-CVC: Controllable Video Compression with Multimodal Generative Models. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Wang, H.; Gan, W.; Hu, S.; Lin, J. Y.; Jin, L.; Song, L.; Wang, P.; Katsavounidis, I.; Aaron, A.; and Kuo, C.-C. J. 2016. MCL-JCV: a JND-based H. 264/AVC video quality assessment dataset. In *2016 IEEE international conference on image processing (ICIP)*, 1509–1513. IEEE.
- Wang, R.; Mao, Q.; Jia, C.; Wang, R.; and Ma, S. 2023. Extreme generative human-oriented video coding via motion representation compression. In *2023 IEEE International Symposium on Circuits and Systems (ISCAS)*, 1–5. IEEE.
- Wang, R.; Mao, Q.; Wang, S.; Jia, C.; Wang, R.; and Ma, S. 2022. Disentangled visual representations for extreme human body video compression. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6. IEEE.
- Wang, Z.; Yuan, Z.; Wang, X.; Li, Y.; Chen, T.; Xia, M.; Luo, P.; and Shan, Y. 2024. Motionctrl: A unified and flexible motion controller for video generation. In *ACM SIGGRAPH 2024 Conference Papers*, 1–11.
- Xiao, Z.; Zhou, Y.; Yang, S.; and Pan, X. 2024. Video diffusion models are training-free motion interpreter and controller. *arXiv preprint arXiv:2405.14864*.
- Xing, J.; Xia, M.; Zhang, Y.; Chen, H.; Yu, W.; Liu, H.; Liu, G.; Wang, X.; Shan, Y.; and Wong, T.-T. 2024a. Dynamicafter: Animating open-domain images with video diffusion priors. In *European Conference on Computer Vision*, 399–417. Springer.
- Xing, Z.; Feng, Q.; Chen, H.; Dai, Q.; Hu, H.; Xu, H.; Wu, Z.; and Jiang, Y.-G. 2024b. A survey on video diffusion models. *ACM Computing Surveys*, 57(2): 1–42.
- Yang, R.; Mentzer, F.; Van Gool, L.; and Timofte, R. 2020. Learning for video compression with recurrent auto-encoder and recurrent probability model. *IEEE Journal of Selected Topics in Signal Processing*, 15(2): 388–401.
- Yang, R.; Timofte, R.; and Van Gool, L. 2022. Perceptual Learned Video Compression with Recurrent Conditional GAN. In *IJCAI*, 1537–1544.
- Yin, S.; Wu, C.; Liang, J.; Shi, J.; Li, H.; Ming, G.; and Duan, N. 2023. Dragnuwa: Fine-grained control in video generation by integrating text, image, and trajectory. *arXiv preprint arXiv:2308.08089*.
- Yin, S.; Zhang, Z.; Chen, B.; Wang, S.; and Ye, Y. 2024. Compressing Scene Dynamics: A Generative Approach. *arXiv preprint arXiv:2410.09768*.
- Zhang, J.; Jia, C.; Lei, M.; Wang, S.; Ma, S.; and Gao, W. 2019. Recent development of AVS video coding standard: AVS3. In *2019 picture coding symposium (PCS)*, 1–5. IEEE.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3836–3847.
- Zhang, P.; Li, J.; Wang, M.; Sebe, N.; Kwong, S.; and Wang, S. 2024a. When Video Coding Meets Multimodal Large Language Models: A Unified Paradigm for Video Coding. *CoRR*.
- Zhang, P.; Wang, S.; Wang, M.; Li, J.; Wang, X.; and Kwong, S. 2023. Rethinking semantic image compression: Scalable representation with cross-modality transfer. *IEEE Transactions on circuits and systems for video technology*, 33(8): 4441–4445.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhang, X.; Duan, Z.; Gong, D.; and Liu, L. 2025a. Training-Free Motion-Guided Video Generation with Enhanced Temporal Consistency Using Motion Consistency Loss. *arXiv preprint arXiv:2501.07563*.
- Zhang, Z.; Liao, J.; Li, M.; Dai, Z.; Qiu, B.; Zhu, S.; Qin, L.; and Wang, W. 2024b. Tora: Trajectory-oriented diffusion transformer for video generation. *arXiv preprint arXiv:2407.21705*.
- Zhang, Z.; Liao, J.; Meng, X.; Qin, L.; and Wang, W. 2025b. Tora2: Motion and Appearance Customized Diffusion Transformer for Multi-Entity Video Generation. *arXiv preprint arXiv:2507.05963*.