

Generative Video Compression with One-Dimensional Latent Representation

Zihan Zheng^{1*} Zhaoyang Jia^{1*} Naifu Xue^{2*} Jiahao Li³ Bin Li³ Zongyu Guo³
Xiaoyi Zhang³ Zhenghao Chen⁴ Houqiang Li¹ Yan Lu³

¹ University of Science and Technology of China ² Communication University of China

³ Microsoft Research Asia ⁴ University of Newcastle

{zzh2003, jzy_ustc}@mail.ustc.edu.cn, lihq@ustc.edu.cn, zhenghao.chen@newcastle.edu.au
{li.jiahao, libin, zongyuguo, xiaoyizhang, yanlu}@microsoft.com, aaronxuenf@cuc.edu.cn

Abstract

Recent advancements in generative video codec (GVC) typically encode video into a 2D latent grid and employ high-capacity generative decoders for reconstruction. However, this paradigm still leaves two key challenges in fully exploiting spatial-temporal redundancy: Spatially, the 2D latent grid inevitably preserves intra-frame redundancy due to its rigid structure, where adjacent patches remain highly similar, thereby necessitating a higher bitrate. Temporally, the 2D latent grid is less effective for modeling long-term correlations in a compact and semantically coherent manner, as it hinders the aggregation of common contents across frames. To address these limitations, we introduce **Generative Video Compression with One-Dimensional (1D) Latent Representation (GVCID)**. GVCID encodes the video data into extreme compact 1D latent tokens conditioned on both short- and long-term contexts. Without the rigid 2D spatial correspondence, these 1D latent tokens can adaptively attend to semantic regions and naturally facilitate token reduction, thereby reducing spatial redundancy. Furthermore, the proposed 1D memory provides semantically rich long-term context while maintaining low computational cost, thereby further reducing temporal redundancy. Experimental results indicate that GVCID attains superior compression efficiency, where it achieves bitrate reductions of 60.4% under LPIPS and 68.8% under DISTs on the HEVC Class B dataset, surpassing the previous video compression methods. Project: <https://gvclid.github.io/>

1. Introduction

As video data continues to grow, achieving efficient compression while preserving human perceptual quality has be-

*This work was done when Zihan Zheng, Zhaoyang Jia and Naifu Xue were full-time interns at Microsoft Research Asia.

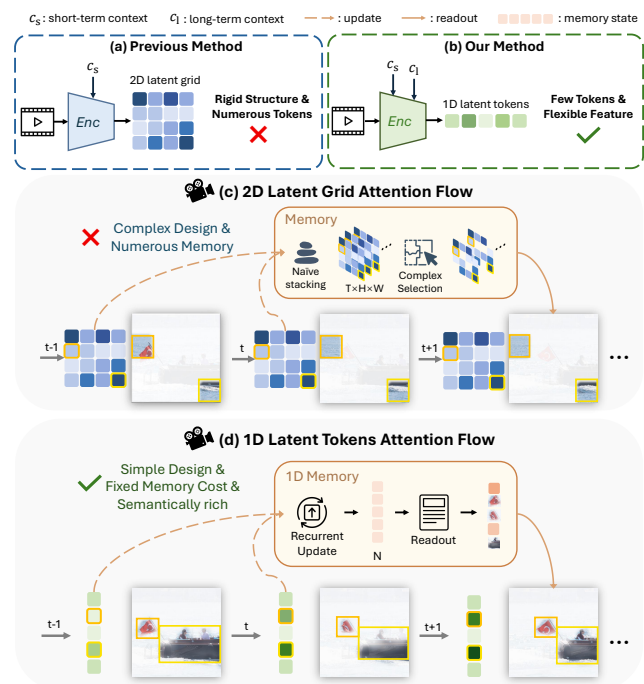


Figure 1. **Method comparison.** (a) Previous generative video codecs [23, 29, 44] encode videos into dense 2D latent grids with rigid spatial structures using short-term context c_s , resulting in numerous inflexible tokens. (b) Our method exploits short-term c_s and 1D-based long-term context c_l to encode videos into a few flexible 1D latent tokens. **Attention map comparison.** (c) 2D latent grids preserve fixed spatial correspondences between tokens and image patches, limiting redundancy exploitation and requiring complex memory designs [3, 30, 36]. (d) Our 1D latent tokens adaptively attend to semantic regions, while the 1D memory, managed by a few 1D tokens, efficiently preserves long-term context in a semantically coherent and computationally efficient manner.

come critical for transmitting high-fidelity content at low bitrates in applications such as video streaming and conferencing. However, both traditional codecs [1–3] and recent neural video codecs (NVCs) [16, 22, 33], which are

typically optimized for generic distortion measures, often yield blurry reconstructions due to limited preserved information [6]. To improve perceptual quality, recent generative video codecs (GVCs) [23, 29, 42, 44] encode raw video into latent representations (typically 2D feature maps) and decode them with powerful generative models [31, 35] to add fine details. Coupled with perceptual loss, these GVCs achieve better perceptual performance. Regardless of their architectural designs or optimization objectives (distortion- or perception-oriented), both traditional and neural approaches ultimately transform the raw signal into 2D spatial structures, such as blocks in traditional codecs or 2D feature maps in neural ones, for subsequent processing such as entropy coding and decoding.

Although 2D representations are intuitive and naturally align with the spatial structure of images, representing frames as 2D latent grids introduces several limitations. Firstly, each token in a 2D latent grid models only its corresponding image patch and focuses more on spatial modeling than object semantics. They capture only variations within its local region, without understanding how the object moves as a whole. Secondly, the 2D latent grid is dense and redundant due to its rigid spatial structure and additional spatial dimension: each image patch is encoded into a token of equal size, resulting in an excessive number of tokens. For instance, both complex and simple regions are assigned the same number of tokens, introducing unnecessary redundancy. These limitations stem from the 2D latent grid’s preservation of spatial structure, which enforces a rigid mapping between 2D latent grids and image patches, thereby restricting the model’s capacity to exploit image redundancy for a more compact latent representation.

To overcome these limitations imposed by 2D structures, we propose compressing videos into a 1D latent representation. Notably, 1D latent representations have been shown to capture high-level semantics efficiently and yield more compact latent spaces in image processing tasks [18, 25, 45, 46]. In particular for image compression [17, 40, 41], prior work [40] has encoded images into 1D latent sequences, achieving extremely low bitrates while preserving competitive perceptual performance. In this work, our 1D representation refers strictly to the compressed latent bottleneck, where a set of unstructured, learnable 1D tokens serving as the 1D latent representation aggregate information from 2D patch embeddings via a ViT, rather than simply flattening them into a sequence. These tokens effectively aggregate essential information from 2D patches without enforcing fixed spatial correspondences (as illustrated in Fig. 1), freeing them from the rigidity and locality constraints of 2D grids. This design decouples the number of tokens from the spatial resolution, contributing to a compact and semantically rich latent space that enables more flexible modeling of semantic variations and improves compression effi-

ciency. Furthermore, by removing rigid spatial correspondences and eliminating the additional spatial dimension, the 1D latent tokens facilitates token reduction, allowing a compact set of latents to efficiently represent video frames and thereby lowering the bitrate.

Beyond effectively removing spatial redundancy, the proposed 1D latent representation enhances long-range dependency modeling, thereby strengthening conditional video compression. Prior studies [27, 28] have demonstrated the benefit of modeling long-term correlations for conditional coding. Yet, their reliance on 2D latent representations limits the ability to provide temporally consistent semantic features. These representations primarily capture spatial variations across frames rather than semantic dynamics, which ultimately constrains long-term context preservation. Hence, leveraging its compactness and enhanced semantic feature modeling, as illustrated in Fig. 1, we further propose a 1D memory that recurrently updates a fixed-size memory state with a small set of 1D latent tokens. Owing to the semantically rich input, the 1D memory can extract temporally consistent semantic features without being constrained by spatial structures, thereby providing a semantically rich and coherent context. Moreover, due to the compactness of the 1D input, the 1D memory can store richer semantic contextual information within the same memory size while keeping the computational cost low. Benefiting from these properties, our memory adopts a simple yet effective design (Transformer) for long-term context modeling. The proposed 1D memory captures abstract long-term semantics through a small number of semantically rich 1D tokens, while the features from the preceding frame preserve fine-grained structural details. Together, they provide a more comprehensive context and further reduce redundancy.

Integrating our design of 1D latent tokens and the 1D memory, we propose **Generative Video Compression with One-Dimensional (1D) Latent Representation (GVC1D)**, demonstrating that the 1D latent representation achieves better performance by removing the rigid spatial structure of 2D representations. Combined with our long-term context constructed from a few 1D tokens, our method achieves state-of-the-art (SOTA) compression performance across multiple datasets and metrics, demonstrating particularly strong performance. Compared to the previous perceptual codec [29], our method reduces the bitrate by 60.4% in LPIPS and 68.8% in DISTS on the HEVC-B dataset. Furthermore, our approach provides visually superior results relative to existing video codecs [1, 22, 29].

In summary, our contributions are as follows:

- We propose a novel video compression method based on 1D latent representations. By removing the rigid 2D spatial structure and reducing the number of tokens, our model encodes videos into a few flexible 1D tokens, en-

- abling more efficient and adaptive video compression.
- We introduce a long-term memory module that leverages 1D latent tokens to construct semantically rich context, enhancing conditional coding through improved contextual modeling.
- Our approach achieves state-of-the-art performance across multiple benchmarks, reducing bitrate while maintaining high perceptual and visual quality.

2. Related Work

Learned Video Compression. Learned video codecs achieve compression by leveraging neural networks to remove redundancy. For neural video codecs (NVCs) primarily optimized for PSNR, the DCVC series [16, 19–22, 32] has demonstrated remarkable potential by replacing traditional residual coding with conditional coding. Through continuous architectural improvements, the latest DCVC-RT [16] even surpasses advanced traditional codecs [1–3] while maintaining high coding speed. However, since these models are mainly optimized for PSNR, they often suffer from severe detail loss at low bitrates, producing blurry frames. Consequently, improving perceptual quality has become a crucial research direction. For generative video codecs (GVCs), the goal is to produce more realistic results using powerful generative models. For example, GLC-Video [15, 29] achieves high perceptual compression quality at low bitrates by performing transform coding in the 2D latent space of a vector-quantized variational autoencoder (VQ-VAE) [35], which is better aligned with human perception but suffers from limited performance due to the restricted capacity of VQ-VAE. Other models, such as DiffVC [23], employ more powerful pretrained diffusion models [31] to enhance perceptual quality; however, their bitrates remain relatively high, failing to fully demonstrate the advantages of GVCs at low bitrates. Our method also focuses on improving perceptual quality. Unlike previous approaches that primarily enhance codec modules while overlooking the critical role of token reduction and token structure in compression efficiency, we address this challenge from a representational perspective by employing a small number of flexible and semantically rich 1D latent representations to improve perceptual quality.

1D Tokenization. In recent years, 1D image tokenization has achieved remarkable progress in image generation [18, 25, 45, 46], enabling the synthesis of high quality images with lots of details. TiTok [46] demonstrated the potential of discrete 1D latent tokens as a replacement for the 2D latent grid, achieving more compact semantic compression that enhances generation quality, while reducing the number of tokens to improve generation speed. TATiTok [18] extended discrete 1D latent tokens to continuous formats and incorporated textual information, thereby

enhancing alignment with text captions. One-D-Piece [25] further advanced this direction by employing the Tail Token Drop regularization technique to develop a variable-length discrete image tokenizer. In the field of image compression, DLF [40] employs discrete 1D tokenization [46] to extract compact semantic representations, achieving superior compression performance through a semantic branch. However, this discrete 1D tokenization is unsuitable for video compression, as the discrete format undermines the video temporal consistency. Moreover, most existing video compression methods still rely on 2D representation. Yet, the dense structure and rigid spatial constraints of 2D representation restrict their ability to achieve efficient compression. Motivated by the continuous 1D tokenizers introduced in TATiTok [18], we enhance video codec efficiency by reducing the number of tokens and exploiting the compact semantic representations offered by 1D latent tokens.

Long-term Context. Long-term context has been extensively leveraged in various tasks, such as video understanding [8, 12, 30], to capture long-range dependencies. Some approaches [36, 38] use fixed-size or bounded dynamic memory to preserve long-term context but suffer from information loss due to dense input features and limited memory capacity. In contrast, other methods [14, 47] reduce computational overhead by adopting suboptimal state-space models (SSMs), thereby sacrificing the strong representational capability of Transformers. In video compression, DCVC-LCG [28] introduces a complex intra-cluster search mechanism based on dense 2D features to capture long-term context. In contrast, our method manages the 1D memory using only a few semantically rich 1D latent tokens. The compact input enables the use of a full Transformer while maintaining low computational cost and mitigating information forgetting. Furthermore, the semantic richness of these 1D tokens allows the memory to naturally focus on compact and meaningful representations without requiring additional complex designs.

3. Methodology

3.1. Framework Overview

Existing generative video codecs [23, 29, 44] typically encode the current frame x_t into a dense 2D latent grid using the short-term context from the preceding frame. However, this approach constrains their ability to effectively exploit the redundancy inherent in videos due to the rigid spatial structure of the 2D latent grid and the inefficient use of temporal information. In contrast, our method encodes x_t into a few 1D latent tokens by leveraging both the short- and long-term context. As illustrated in Fig. 2, we first employ a Vision Transformer-based [10] encoder to encode the current frame x_t into a few 1D latent tokens y_t . Next, we perform entropy coding in these 1D latent tokens to achieve efficient

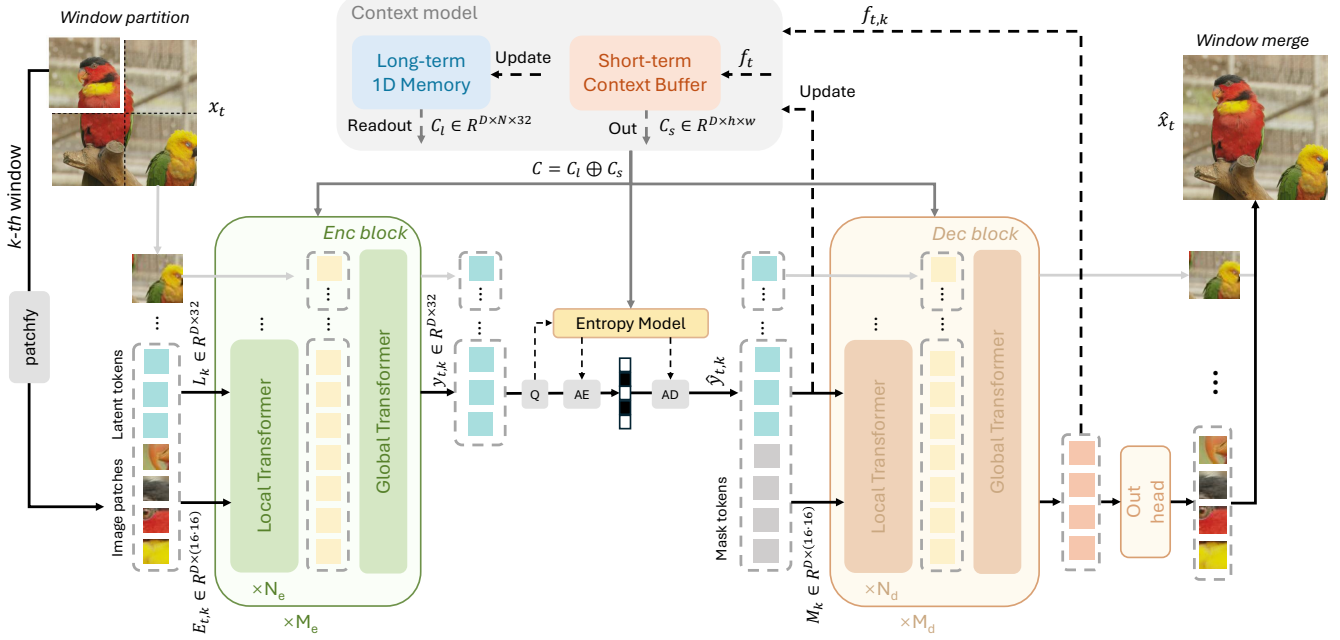


Figure 2. Framework overview. Q, AE and AD represent quantization, arithmetic encoder and decoder, respectively. The input image x_t is first embedded into patches and then fed into the encoder composed of local and global transformers to produce y_t . The local transformer handles each window independently with low cost, while the global transformer captures inter-window correlations. Subsequently, an entropy model then performs autoregressive entropy coding on y_t . The decoder adopts an architecture similar to the encoder to reconstruct the image \hat{x}_t . All are guided by a context model combining long-term 1D memory and short-term context buffer to provide comprehensive temporal context.

compression. Finally, these 1D latent tokens are decoded to reconstruct the frame \hat{x}_t . In addition, all three stages are guided by our context model, which integrates two types of context: long-term context from previous 1D latent tokens and short-term context from the preceding frame.

Encoding. Our method takes a video as input. For the current frame $x_t \in \mathbb{R}^{3 \times H \times W}$, it is first embedded into $Emb(x_t) \in \mathbb{R}^{D \times h \times w}$ through patch embedding, with a patch size of $h = H/16$ and $w = W/16$. The embedded features $Emb(x_t)$ are then flattened into a 1D sequence $E_t \in \mathbb{R}^{D \times (h \cdot w)}$. Next, The embedding E_t is concatenated with learnable 1D latent tokens $L \in \mathbb{R}^{D \times (N \cdot 32)}$, where $N = \frac{h \times w}{16 \times 16}$, and the context C from the context model, which integrates both long-term context $C_l \in \mathbb{R}^{D \times (N \cdot 32)}$ and the flattened short-term context $C_s \in \mathbb{R}^{D \times (h \cdot w)}$. The combined tokens are then fed into the encoder to generate compact 1D tokens y_t .

$$y_t = Enc(E_t \oplus L \oplus C), \quad (1)$$

$$C = C_l \oplus C_s. \quad (2)$$

Specifically, our encoder consists of M_e encoding blocks, and each block contains N_e layers of local transformers followed by one layer of global transformer:

$$Enc = \bigcup_{i=1}^{M_e} \left\{ \bigcup_{j=1}^{N_e} \text{LocalTrans}_{i,j}, \text{GlobalTrans}_i \right\}. \quad (3)$$

For the local transformer, we partition the combined tokens into N windows and process each window in parallel. Specifically, the embedding E_t and short-term context C_s are divided into 16×16 windows, while the 1D latent tokens L and long-term context C_l are divided into 32 tokens per window. Within each window, these four components are concatenated and fed into the local transformer, which leverages the Vision Transformer (ViT) to encode image tokens into 1D latent tokens. Through cascaded local transformer layers, the 1D latent tokens effectively aggregate key semantic and motion information from the image tokens and the context C , thereby producing a compact 1D token representation with rich semantics.

To model the correlations across windows, we introduce a global transformer layer after every N_e local transformer layers. Specifically, the 2D image tokens E_t , short-term context C_s , 1D latent tokens L , and long-term context C_l are first recombined and assigned 2D Rotary Positional Embeddings (RoPE) [13] to independently model relative positional relationships. The resulting tokens are then fed into the global transformer, which employs a ViT to process the entire sequence and facilitate global information exchange. This design allows independent windows to interact through global attention while maintaining low computational cost by employing only a few global transformer layers.

Consequently, we encode these compact 1D tokens to produce $y_t \in \mathbb{R}^{D \times (N \cdot 32)}$. Unlike previous methods [29]

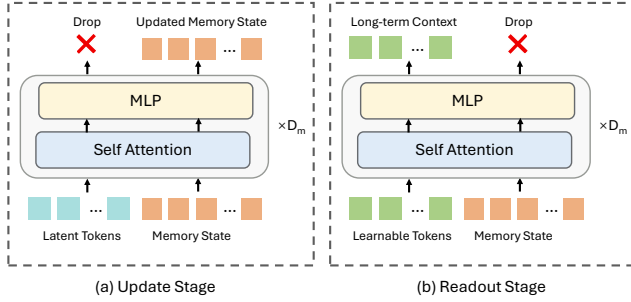


Figure 3. 1D memory. D_m denote the number of Transformer layers. We employ simple yet effective Transformer layers to manage long-term context.

that encode frames into dense 2D latent grids, our approach compresses frames into 1D latent tokens. By removing rigid spatial correspondences, these tokens provide greater flexibility in modeling frames, capture richer semantics, and form a more compact latent space, thereby achieving lower bitrates and improved perceptual reconstruction quality.

Entropy Model. To estimate the probability distribution of 1D latent tokens for arithmetic coding, we employ an autoregressive Transformer (AR Transformer)-based entropy model. Specifically, the 1D latent representation y_t is first quantized into $Q(y_t)$, after which an AR Transformer models its distribution for entropy coding. During this process, the AR Transformer predicts the probability of each quantized token sequentially, conditioned on all previously decoded tokens. Consequently, the autoregressive modeling effectively captures inter-dependencies among tokens, thereby achieving lower bitrates. Notably, the number of 1D latent tokens is small, and the 1D latent tokens from N windows can be processed in parallel, keeping the computational cost of the AR Transformer manageable.

Decoding. Our decoder follows an architecture similar to that of the encoder. Specifically, a sequence of learnable mask tokens $M \in \mathbb{R}^{D \times (h \cdot w)}$ is first integrated with the 1D latent tokens \hat{y}_t and the context C . Through the decoding process, these mask tokens iteratively extract information from \hat{y}_t and C , and the reconstructed frame \hat{x}_t is then generated through a convolution output head.

$$\hat{x}_t = \text{Out}(\text{Dec}(\hat{y}_t \oplus M \oplus C)). \quad (4)$$

3.2. Context Model

In this section, we present our context model, which consists of two components: a long-term 1D memory and a short-term context buffer. Unlike previous methods [16, 22, 29], our model incorporates long-term context derived from 1D tokens to provide richer temporal information. The short-term context captures fine grained features but lacks awareness of long-range dependencies, whereas the long-term context delivers global semantic context that supple-

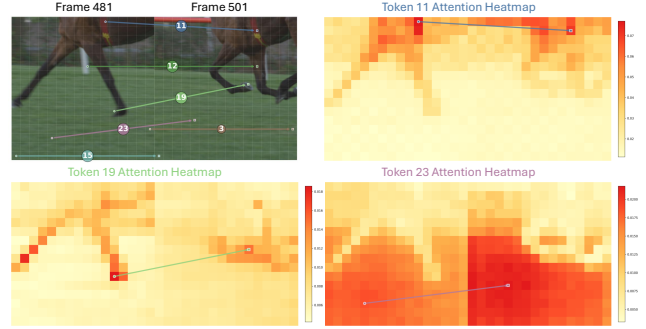


Figure 4. Visualization of 1D latent token outflows across two frames during object motion. In the two figures, the lines connect points corresponding to the maximum attention weights of each token, with the numbers indicating token indices (e.g., token 19 focuses on the horse’s left foreleg in both frames). Detailed attention maps show that each token consistently focuses on the same semantic region across frames, effectively capturing object motion.

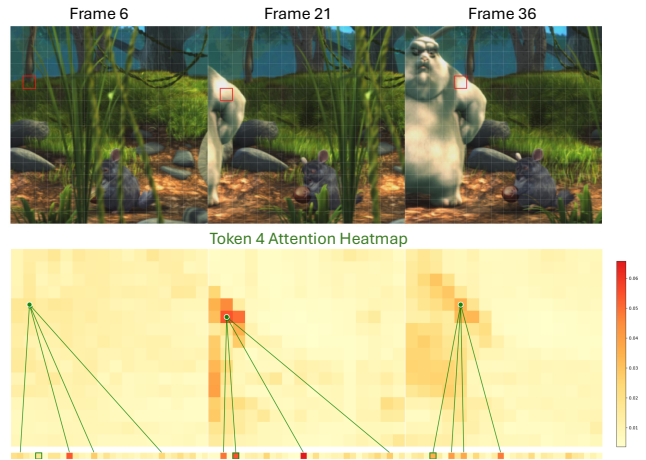


Figure 5. Visualization of the outflow variation of a 1D latent token (index 4) as a new object appears. The red boxes in the first row mark image patches with the highest attention weights, while the green lines in the second row link them to the top four 1D latent tokens with the strongest attention. The bottom row is the 1D latent tokens attention weights corresponding to the maximum weight image patch (red boxes). As new content emerges, attention weights gradually shift from previously active tokens to newly activated ones.

ments the short-term context. Together, these two components form a comprehensive temporal representation, effectively exploiting temporal dependencies and further reducing redundancy in videos.

For the short-term context, similar to previous methods [22], we use the decoder features f_t from the preceding frame to provide detailed short-term context. For the long-term context, we design a simple yet effective 1D memory that provides global semantic features using the Transformer architecture. As shown in Fig. 3, our 1D memory maintains a fixed size memory state and comprises two

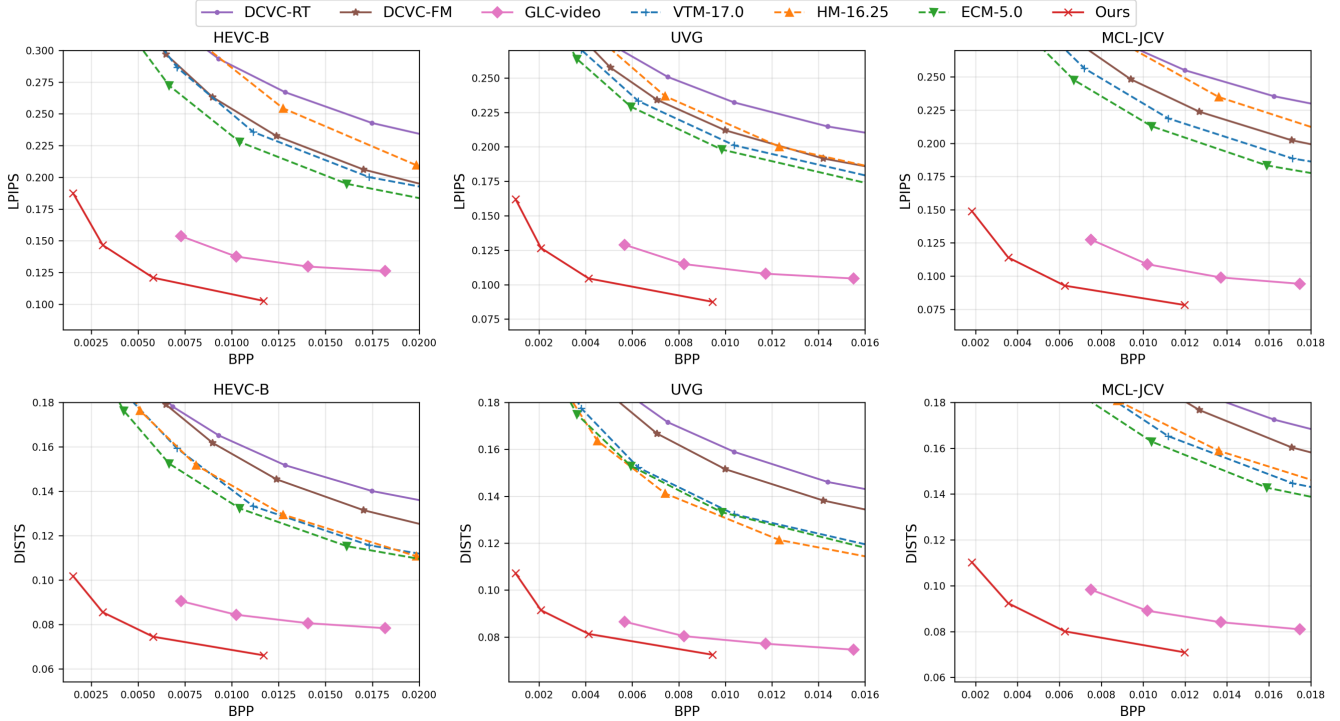


Figure 6. Rate-distortion curves on the HEVC-B, the UVG and the MCL-JCV datasets. 96 frames are tested in RGB colorspace with $\text{intra-period}=-1$.

stages for recurrently updating the memory and reading out long-term context. In the update stage, a few 1D latent tokens \hat{y}_t are used to update the memory state, while in the readout stage, learnable query tokens retrieve the long-term context from the memory. All operations are performed in the 1D latent space, which is free from rigid 2D spatial constraints and thus enables flexible memory management without complex designs. By leveraging compact and semantically rich 1D latent tokens to manage memory, our 1D memory provides more global and semantic context than the fine-grained short-term context. Moreover, since the number of 1D latent tokens is considerably smaller than that of the 2D latent grids used in previous works, the model can store more information within a fixed memory capacity, effectively alleviating information forgetting.

3.3. Analysis on 1D Latent Tokens

To validate that our 1D representation can effectively and flexibly model videos in a semantic manner, thereby achieving better perceptual quality, we follow FlowCut [34] and visualize the attention maps of the 1D latent tokens corresponding to the final reconstructed frames. Specifically, for an attention map $A \in \mathbb{R}^{M \times M}$, we visualize the outflow of each 1D latent token, where the outflow represents the primary destination of its information, defined as $\text{argmax}_A[:, j]$. We illustrate two representative video scenarios: object motion and the appearance of new content, based on the attention map from the last decoder layer.

Object Motion. As shown in Fig. 4, each token consistently focuses on the same semantic region across two frames, even under large motion. For example, token 23 attends to the grassland in both frames. In more challenging cases involving large motion, such as token 19 focusing on the horse’s left foreleg in both frames, more detailed information is required. Consequently, the attention map of token 19 covers a smaller area than that of token 23. These observations demonstrate that our 1D latent token design effectively tracks consistent semantic regions throughout the video while capturing intrinsic motion. Moreover, it can adaptively assign attention weights across the entire image without being constrained by the 2D spatial structure.

New Content. As shown in Fig. 5, we visualize the attention maps of token 4 and the variation in attention weights of the 1D latent tokens corresponding to newly appeared content. Specifically, as the rabbit gradually emerges in the video, the attention map shifts its focus from the blank background to the rabbit. Initially, the token focuses on the grass background. As the rabbit emerges, the attention weights of certain 1D latent tokens progressively increase to capture the new content, indicating that they begin to attend to the newly appeared object. Finally, the attention weights are re-assigned to adapt to the new scene. This observation demonstrates the effectiveness of our 1D latent token design in handling the emergence of new objects, as the 1D tokens can dynamically reallocate attention to newly appearing content.

Table 1. BD-Rate (%) comparison in the RGB colorspace on the HEVC-B, UVG, and MCL-JCV datasets (lower is better). 96 frames are evaluated with an intra-period of -1.

Method	HEVC-B		UVG		MCL-JCV	
	LPIPS	DISTS	LPIPS	DISTS	LPIPS	DISTS
GLC-video [29]	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Ours	-60.4%	-68.8%	-66.0%	-47.9%	-62.1%	-61.5%

Method	HEVC-B		UVG		MCL-JCV	
	PSNR	MS-SSIM	PSNR	MS-SSIM	PSNR	MS-SSIM
GLC-video [29]	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Ours	-53.8%	-45.1%	-62.5%	-45.8%	-54.7%	-33.0%

4. Experiment

4.1. Experimental Protocol

Training Details. Our model is trained on the Vimeo [4, 43] and OpenVid-HD [26] datasets. We train a variable-bitrate model using the loss function $L = R + \lambda D$, where R and D denote the bitrate and distortion, respectively. The trade-off parameter λ is sampled at 8 points using log-space linear interpolation within the range [0.07, 1.5]. The detailed training strategy and loss formulation are provided in the supplementary material.

Test Details. We test our model on the HEVC-B [11], UVG [24], and MCL-JCV [37] datasets. All experiments are conducted under low-delay encoding settings (intra-period=-1) in the RGB colorspace. The BD-Rate [5] metric is employed for performance evaluation. To assess reconstruction quality, we use the reference perceptual metrics LPIPS [48] and DISTS [9]. The bitrate is measured in bits per pixel (bpp), and the PSNR and MS-SSIM [39] results are also reported in the supplementary material.

Baselines. For traditional codecs, we compare against HM-16.25 [2], VTM-17.0 [3], and ECM-5.0 [1], which are primarily optimized for PSNR. For neural methods, we include the PSNR-oriented codecs DCVC-FM [22] and DCVC-RT [16], as well as the perceptual generative codec GLC-Video [29], which targets perceptual quality.

4.2. Experimental Results

Quantitative Evaluation. Fig. 6 shows the rate-distortion curves across datasets. Our method consistently outperforms prior methods on perceptual metrics, achieving lower LPIPS and DISTS at the same bitrate. As reported in Table 1, compared with GLC-Video, our method reduces bitrate on HEVC-B by 60.4% / 68.8% for LPIPS / DISTS and 53.8% / 45.1% for PSNR / MS-SSIM. Detailed rate-distortion curves for PSNR and MS-SSIM are provided in the supplementary material. These results highlight the effectiveness of our 1D latent token design, which reduces token count and leverages long-term context to achieve lower bitrates while maintaining high perceptual quality.

Qualitative Evaluation. In Fig. 7, we present visual comparisons on the UVG and MCL-JCV datasets. The neural

Table 2. BD-Rate comparison for different token sizes. The token size of 32×16 is used as the anchor.

Token Size (Number \times Channel)	HEVC-B	UVG	MCL-JCV
16×32	41.7%	47.6%	37.3%
16×64	20.1%	23.4%	12.2%
32×16 (Ours)	0.0%	0.0%	0.0%
32×32	14.2%	21.5%	10.2%
64×16	24.2%	22.8%	17.6%

Table 3. BD-Rate comparison of different model variants. AR denotes the autoregressive entropy model, and Memory denotes the memory component, where 1D and 2D indicate the use of 1D or 2D features to manage the memory. We use setting (4) as the anchor, which is finally adopted by our method.

Setting	AR	Memory	HEVC-B	UVG	MCL-JCV
(1)	\times	\times	67.8%	67.4%	41.8%
(2)	\checkmark	\times	20.1%	40.6%	24.2%
(3)	\checkmark	2D	11.5%	16.8%	7.3%
(4)	\checkmark	1D	0.0%	0.0%	0.0%

video codec DCVC-FM and the traditional codec ECM-5.0 struggle to reproduce realistic details at low bitrates, leading to blurred regions (e.g., the railing in the first example and the basket pattern in the third example). The perceptual codec GLC-Video enhances visual realism but introduces noticeable artifacts and distortions, such as unrealistic facial features and clothing textures in the second example. In contrast, our method produces reconstructions with fewer artifacts, more accurate colors, and finer details, even at lower bitrates.

4.3. Ablation Study

In this section, we conduct ablation studies on different latent token sizes, the 1D memory, and the autoregressive entropy model. All model variants are evaluated on HEVC-B, UVG, and MCL-JCV using the LPIPS metric.

1D Latent Token Size. We evaluate the impact of different 1D latent token sizes. As shown in Table 2, a smaller token size (e.g., 16×32 , where 16 denotes the number of tokens and 32 denotes the token channel) restricts the model’s capacity, thereby degrading reconstruction quality and resulting in a worse BD-Rate. In contrast, a larger token size (e.g., 64×16) increases the bitrate due to its higher dimensionality, which also leads to a worse BD-Rate. Hence, selecting an appropriate token size is essential to achieve a balanced trade-off between bitrate and reconstruction quality.

1D Memory. In this variant, we first remove the 1D memory to assess the importance of long-term contextual information. As shown in Table 3, comparing settings (2) and (4) reveals that removing the long-term context results in a BD-Rate increase of over 40% on the UVG dataset, demonstrating the necessity of long-term context. Next, we replace the 1D memory with a 2D memory constructed from 2D short-term context to examine the effectiveness of the 1D memory

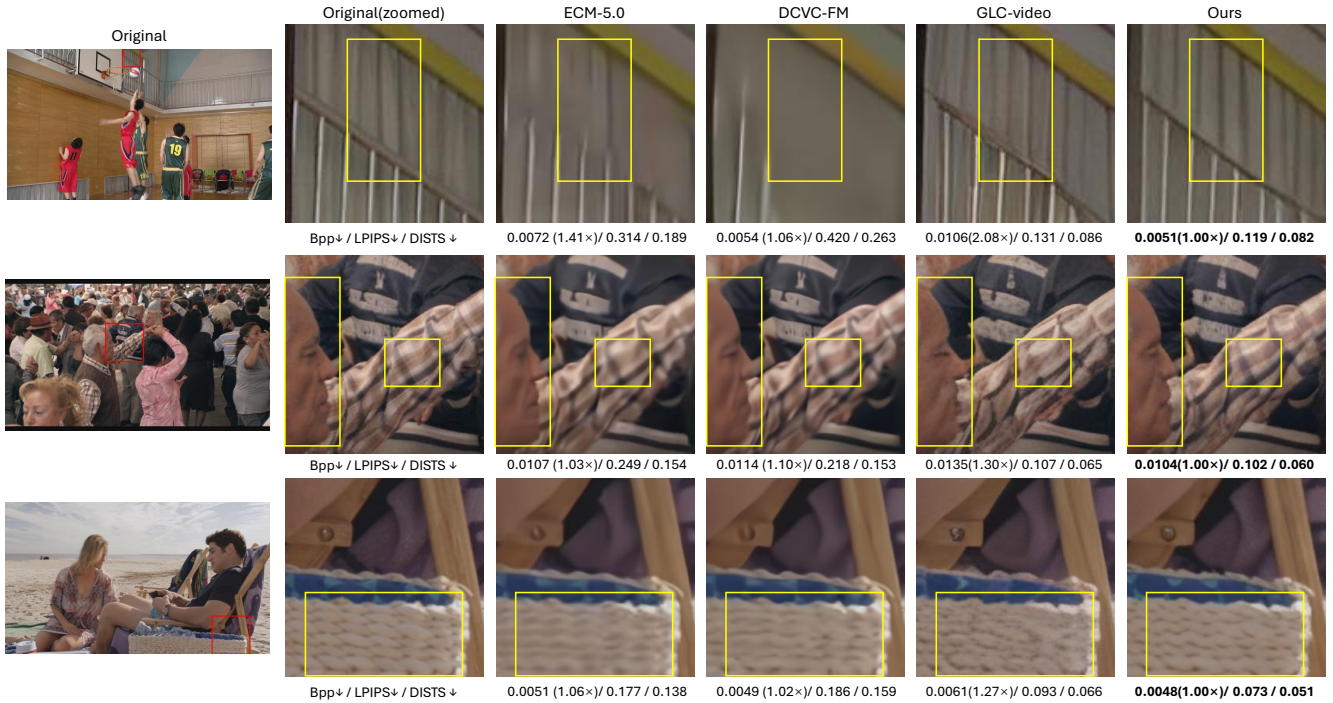


Figure 7. Qualitative examples on the MCL-JCV datasets.

Table 4. Complexity analysis using fp16 precision at 1080P resolution. The tests are conducted on an NVIDIA A100 GPU.

Model	Enc. Time (s)	Dec. Time (s)
DCVC-FM [22]	0.183	0.190
GLC-video [29]	0.290	0.183
Ours	0.262	0.207

design. By comparing settings (3) and (4), we observe that this replacement leads to a BD-Rate degradation of more than 16% on the UVG dataset, confirming that using 1D features for memory management is more effective.

Autoregressive Entropy Model. We further evaluate the effectiveness of our proposed autoregressive entropy model by replacing it with a non-autoregressive version, under the same setting without the 1D memory. As shown in Table 3, comparing setting (1) and (2), replacing the autoregressive entropy model leads to a worse BD-Rate. These results confirm that the autoregressive entropy model effectively reduces redundancy among tokens.

4.4. Complexity Analysis

We compare the average coding time in Table 4. Compared with DCVC-FM, our method is only slightly slower due to its larger model size. However, the larger model provides stronger capacity, enabling our method to achieve significantly better perceptual quality with richer details. Compared with GLC-Video, our method attains a similar overall coding speed, with the encoding even slightly faster, while delivering notably higher visual quality. Overall, our approach improves compression efficiency while maintaining

a relatively fast coding speed. This efficiency stems from performing fast compression in a compact latent space with fewer tokens, utilizing a parallel window processing design, and employing advanced techniques such as FlashAttention [7] to accelerate computation.

5. Conclusion

In this paper, we present a novel method for modeling videos using 1D latent tokens instead of conventional dense 2D latent grids. We first analyze the limitations of 2D latent grids, including the rigid 2D spatial structure and the redundancy of token numbers. To address these issues, we represent videos with 1D latent tokens. Without the constraints of the rigid spatial correspondences, 1D latent tokens can effectively capture intrinsic object motion and semantics while flexibly reducing the number of tokens. Moreover, our 1D memory incorporates semantical long-term context to further exploit temporal information. By leveraging compact and semantically rich 1D latent tokens to manage memory, our design effectively mitigates information forgetting while maintaining low computational cost. Experimental results demonstrate that our model achieves state-of-the-art reconstruction quality while maintaining a relatively fast coding speed.

Limitations. Due to the inherent information loss from representing each frame with only 32 1D tokens and the limited token capacity, our model currently targets lossy compression at low bitrates. Future work will explore extending it towards lossless video compression.

References

- [1] ECM. <https://vcgit.hhi.fraunhofer.de/ecm/ECM>. 1, 2, 3, 7
- [2] HM. <https://vcgit.hhi.fraunhofer.de/jvet/HM>. 7
- [3] VTM. https://vcgit.hhi.fraunhofer.de/jvet/VVCSsoftware_VTM. 1, 3, 7
- [4] Original Vimeo links. https://github.com/anchen1011/toflow/blob/master/data/original_vimeo_links.txt. 7
- [5] Gisle Bjontegaard. Calculation of average PSNR differences between RD-curves. *VCEG-M33*, 2001. 7
- [6] Yochai Blau and Tomer Michaeli. Rethinking lossy compression: The rate-distortion-perception tradeoff. In *International Conference on Machine Learning*, pages 675–685. PMLR, 2019. 2
- [7] Tri Dao. FlashAttention-2: Faster attention with better parallelism and work partitioning. In *International Conference on Learning Representations (ICLR)*, 2024. 8
- [8] Shangzhe Di, Zhelun Yu, Guanghao Zhang, Haoyuan Li, TaoZhong, Hao Cheng, Bolin Li, Wangui He, Fangxun Shu, and Hao Jiang. Streaming video question-answering with in-context video KV-cache retrieval. In *The Thirteenth International Conference on Learning Representations*, 2025. 3
- [9] 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, 2020. 7
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 3
- [11] D Flynn, K Sharman, and C Rosewarne. Common Test Conditions and Software Reference Configurations for HEVC Range Extensions, document JCTVC-N1006. *Joint Collaborative Team Video Coding ITU-T SG*, 16. 7
- [12] Bo He, Hengduo Li, Young Kyun Jang, Menglin Jia, Xuefei Cao, Ashish Shah, Abhinav Shrivastava, and Ser-Nam Lim. Ma-Imm: Memory-augmented large multimodal model for long-term video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13504–13514, 2024. 3
- [13] Byeongho Heo, Song Park, Dongyoon Han, and Sangdoon Yun. Rotary position embedding for vision transformer. In *European Conference on Computer Vision*, pages 289–305. Springer, 2024. 4
- [14] Syed Ariff Syed Hesham, Yun Liu, Guolei Sun, Henghui Ding, Jing Yang, Ender Konukoglu, Xue Geng, and Xudong Jiang. Exploiting temporal state space sharing for video semantic segmentation. In *CVPR*, 2025. 3
- [15] 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*, pages 26088–26098, 2024. 3
- [16] Zhaoyang Jia, Bin Li, Jiahao Li, Wenxuan Xie, Linfeng Qi, Houqiang Li, and Yan Lu. Towards practical real-time neural video compression. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-25, 2024*, 2025. 1, 3, 5, 7
- [17] Zhaoyang Jia, Zihan Zheng, Naifu Xue, Jiahao Li, Bin Li, Zongyu Guo, Xiaoyi Zhang, Houqiang Li, and Yan Lu. Cod: A diffusion foundation model for image compression. *arXiv preprint arXiv:2511.18706*, 2025. 2
- [18] Dongwon Kim, Ju He, Qihang Yu, Chenglin Yang, Xiaohui Shen, Suha Kwak, and Liang-Chieh Chen. Democratizing text-to-image masked generative models with compact text-aware one-dimensional tokens. *arXiv preprint arXiv:2501.07730*, 2025. 2, 3
- [19] Jiahao Li, Bin Li, and Yan Lu. Deep contextual video compression. *Advances in Neural Information Processing Systems*, 34, 2021. 3
- [20] Jiahao Li, Bin Li, and Yan Lu. Hybrid spatial-temporal entropy modelling for neural video compression. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2022.
- [21] Jiahao Li, Bin Li, and Yan Lu. Neural video compression with diverse contexts. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, Canada, June 18-22, 2023*, 2023.
- [22] Jiahao Li, Bin Li, and Yan Lu. Neural video compression with feature modulation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 17-21, 2024*, 2024. 1, 2, 3, 5, 7, 8
- [23] Wenzhuo Ma and Zhenzhong Chen. Diffusion-based perceptual neural video compression with temporal diffusion information reuse. *ACM Trans. Multimedia Comput. Commun. Appl.*, 2025. Just Accepted. 1, 2, 3
- [24] 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 297–302, 2020. 7
- [25] Keita Miwa, Kento Sasaki, Hidehisa Arai, Tsubasa Takahashi, and Yu Yamaguchi. One-d-piece: Image tokenizer meets quality-controllable compression. *arXiv preprint arXiv:2501.10064*, 2025. 2, 3
- [26] 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. In *The Thirteenth International Conference on Learning Representations*, 2025. 7
- [27] Huu-Tai Phung, Zong-Lin Gao, Yi-Chen Yao, Kuan-Wei Ho, Yi-Hsin Chen, Yu-Hsiang Lin, Alessandro Gnutti, and Wen-Hsiao Peng. Mh-lvc: Multi-hypothesis temporal prediction for learned conditional residual video coding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19649–19658, 2025. 2

- [28] Linfeng Qi, Zhaoyang Jia, Jiahao Li, Bin Li, Houqiang Li, and Yan Lu. Long-term temporal context gathering for neural video compression. In *European Conference on Computer Vision*, pages 305–322. Springer, 2024. 2, 3
- [29] 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*, 2025. 1, 2, 3, 4, 5, 7, 8
- [30] Rui Qian, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Shuangrui Ding, Dahua Lin, and Jiaqi Wang. Streaming long video understanding with large language models. *Advances in Neural Information Processing Systems*, 37:119336–119360, 2024. 1, 3
- [31] 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*, pages 10684–10695, 2022. 2, 3
- [32] Xihua Sheng, Jiahao Li, Bin Li, Li Li, Dong Liu, and Yan Lu. Temporal context mining for learned video compression. *IEEE Transactions on Multimedia*, 2022. 3
- [33] Chuanbo Tang, Zhuoyuan Li, Yifan Bian, Li Li, and Dong Liu. Neural video compression with context modulation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 12553–12563, 2025. 1
- [34] Jintao Tong, Wenwei Jin, Pengda Qin, Anqi Li, Yixiong Zou, Yuhong Li, Yuhua Li, and Ruixuan Li. Flowcut: Rethinking redundancy via information flow for efficient vision-language models. *arXiv preprint arXiv:2505.19536*, 2025. 6
- [35] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 2, 3
- [36] Hengyi Wang and Lourdes Agapito. 3d reconstruction with spatial memory. In *2025 International Conference on 3D Vision (3DV)*, pages 78–89. IEEE, 2025. 1, 3
- [37] Haiqiang Wang, Weihao Gan, Sudeng Hu, Joe Yuchieh Lin, Lina Jin, Longguang Song, Ping Wang, Ioannis Katsavounidis, Anne Aaron, and C-C Jay Kuo. MCL-JCV: a JND-based H. 264/AVC video quality assessment dataset. In *2016 IEEE international conference on image processing (ICIP)*, pages 1509–1513. IEEE, 2016. 7
- [38] Qianqian Wang, Yifei Zhang, Aleksander Holynski, Alexei A Efros, and Angjoo Kanazawa. Continuous 3d perception model with persistent state. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 10510–10522, 2025. 3
- [39] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multi-scale structural similarity for image quality assessment. In *The thirty-seventh asilomar conference on signals, systems & computers, 2003*, pages 1398–1402. Ieee, 2003. 7
- [40] Naifu Xue, Zhaoyang Jia, Jiahao Li, Bin Li, Yuan Zhang, and Yan Lu. Dlf: Extreme image compression with dual-generative latent fusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025. 2, 3
- [41] Naifu Xue, Zhaoyang Jia, Jiahao Li, Bin Li, Yuan Zhang, and Yan Lu. One-step diffusion-based image compression with semantic distillation. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. 2
- [42] Naifu Xue, Zhaoyang Jia, Jiahao Li, Bin Li, Zihan Zheng, Yuan Zhang, and Yan Lu. Single-step diffusion-based video coding with semantic-temporal guidance. *arXiv preprint arXiv:2512.07480*, 2025. 2
- [43] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision*, 127(8): 1106–1125, 2019. 7
- [44] Ren Yang, Radu Timofte, and Luc Van Gool. Perceptual video compression with recurrent conditional gan. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2022. 1, 2, 3
- [45] Qihang Yu, Ju He, Xueqing Deng, Xiaohui Shen, and Liang-Chieh Chen. Randomized autoregressive visual generation. *arXiv preprint arXiv:2411.00776*, 2024. 2, 3
- [46] Qihang Yu, Mark Weber, Xueqing Deng, Xiaohui Shen, Daniel Cremers, and Liang-Chieh Chen. An image is worth 32 tokens for reconstruction and generation. *NeurIPS*, 2024. 2, 3
- [47] Olga Zatsarynna, Emad Bahrami, Yazan Abu Farha, Gianpiero Francesca, and Juergen Gall. Manta: Diffusion mamba for efficient and effective stochastic long-term dense action anticipation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3
- [48] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 7