



# USENIX

THE ADVANCED COMPUTING  
SYSTEMS ASSOCIATION

## Morphe: High-Fidelity Generative Video Streaming with Vision Foundation Model

Tianyi Gong, Zijian Cao, and Zixing Zhang, *The Chinese University of Hong Kong, Shenzhen, and Shenzhen Future Network of Intelligence Institute;*  
Jiangkai Wu and Xinggong Zhang, *Peking University;* Shuguang Cui  
and Fangxin Wang, *The Chinese University of Hong Kong, Shenzhen,  
and Shenzhen Future Network of Intelligence Institute*

<https://www.usenix.org/conference/nsdi26/presentation/gong>

This paper is included in the Proceedings of the 23rd USENIX Symposium  
on Networked Systems Design and Implementation.

May 4–6, 2026 • Renton, WA, USA

ISBN 978-1-939133-54-0

Open access to the Proceedings of the 23rd USENIX Symposium  
on Networked Systems Design and Implementation is sponsored by



جامعة الملك عبد الله  
للعلوم والتقنية  
King Abdullah University of  
Science and Technology

# Morphe: High-Fidelity Generative Video Streaming with Vision Foundation Model

Tianyi Gong<sup>1,2</sup>, Zijian Cao<sup>1,2</sup>, Zixing Zhang<sup>1,2</sup>, Jiangkai Wu<sup>3</sup>, Xinggong Zhang<sup>3</sup>,  
Shuguang Cui<sup>1,2</sup>, Fangxin Wang<sup>1,2</sup>

<sup>1</sup>*School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen,*

<sup>2</sup>*Shenzhen Future Network of Intelligence Institute,*

<sup>3</sup>*Wangxuan Institute of Computer Technology, Peking University*

## Abstract

Video streaming is a fundamental Internet service, while the quality still cannot be guaranteed especially in poor network conditions such as bandwidth-constrained and remote areas. Existing works mainly work towards two directions: traditional pixel-codec streaming nearly approaches its limit and is hard to step further in compression; the emerging neural-enhanced or generative streaming usually fall short in latency and visual fidelity, hindering their practical deployment.

Inspired by the recent success of vision foundation model (VFM), we strive to *harness the powerful video understanding and processing capacities of VFM to achieve generalization, high fidelity and loss resilience for real-time video streaming with even higher compression rate*. We present Morphe<sup>1</sup>, the first revolutionized paradigm that enables VFM-based end-to-end generative video streaming towards this goal. Specifically, Morphe employs joint training of visual tokenizers and variable-resolution spatiotemporal optimization under simulated network constraints. Additionally, a robust streaming system is constructed that leverages intelligent packet dropping to resist real-world network perturbations. Extensive evaluation demonstrates that Morphe achieves comparable visual quality while saving 62.5% bandwidth compared to H.265, and accomplishes real-time, loss-resilient video delivery in challenging network environments, representing a milestone in VFM-enabled multimedia streaming solutions.

## 1 Introduction

Nowadays video content dominates the Internet traffic, accounting for more than 85% based on [27]. However, accessing high-quality video content is often hindered by insufficient and unstable network bandwidth between client and server, which is even worse for high-resolution live video streaming. In such network-constrained scenarios as high-speed railway

and remote areas, people's demand for premium media content remains difficult to satisfy.

**The Main Roadmap:** To tackle this problem, multimedia streaming employs codec optimization and adaptive bitrate control, and network-aware protocols to maintain video quality under constrained conditions. *Traditional methods* (e.g., codec with H.264/265 [30, 35] with adaptive bitrate [28]) focus on pixel-level encoding and compression, pursuing pixel-perfect reconstruction. However, when confronted with insufficient or fluctuating bandwidth, they struggle to maintain optimal performance, leading to compression artifacts or playback interruptions. While widely adopted in modern streaming service, pixel-based codec methods nearly approach their physical limits in bandwidth-constrained environments where visual quality compromises become necessary.

*Neural-enhanced streaming* [14, 15, 43] emerges as a promising direction, which leverages learning-based post-processing to achieve substantial bandwidth savings, but often suffers from considerable computational overhead as well as increased latency. *Generative streaming approaches* [2, 33] attempt to reconstruct video content from compact semantic representations, essentially generating frames from high-level descriptions rather than transmitting pixels [48]. These methods offer promising bandwidth efficiency but often suffer from unacceptable visual quality degradation, typically stemming from poor generalizability. Moreover, due to the dense mapping between semantics and pixels and the lack of recovery mechanism, a slight packet loss can lead to serious distortion.

The core tenet of video streaming lies in **streaming high-fidelity, real-time, and loss-resilient video content with highest possible video compression**. However, under complex network conditions, existing streaming systems—both traditional pixel-based codecs and emerging neural approaches—struggle to balance this triangular relationship effectively. Compromises in any dimension continue to severely degrade Quality of Service (QoS) across all current streaming paradigms.

**Opportunities and Challenges:** Vision Foundation Models (VFMs) [4] have attracted considerable attention owing

The first three authors, Tianyi Gong, Zijian Cao, and Zixing Zhang, contributed equally to this paper. Fangxin Wang is the corresponding author.

<sup>1</sup>Our project page can be available at <https://inml-tygong.github.io/Morphe>

to their exceptional capabilities in video understanding [45] and generation [37]. VFMs present a transformative opportunity for content-aware video streaming systems through their unique "understanding before generation" paradigm. Pre-trained on large-scale visual corpora via unsupervised learning, VFMs excel at distilling redundant visual data into compact semantic representations that naturally align with human perceptual priorities—prioritizing faces, salient objects, and textual elements over low-level textural details. This semantic-aware encoding enables extraordinary compression gains (exceeding  $300\times$  in recent studies [24]) while maintaining perceptually-optimized quality beyond conventional pixel-level metrics. Moreover, the zero-shot generalization capability inherent to VFMs ensures robust performance when confronted with real-world data variability. Inspired by the remarkable potential of VFMs, we are motivated to investigate a question: *Can we achieve Vision Foundation Model-based video coding towards high-fidelity, loss-resilient, and real-time streaming service?*

Unfortunately, through our analysis in §2.4, applying VFM for video streaming still faces triple challenges:

**C1: How to maximize VFM-based compression with high-fidelity visual recovery?** VFMs are fundamentally designed for semantic understanding and generation tasks, primarily focusing on low-frequency content, which inherently conflicts with high-fidelity video delivery that demands detailed texture preservation. Beyond severe detail loss, spatiotemporal discontinuities also pose significant challenges when adapting VFMs to video streaming, potentially stemming from high spatiotemporal compression ratios and improper video chunking strategies. The fundamental mismatch between semantic intelligence and high-fidelity transmission requirements explains why simply migrating VFMs to streaming scenarios fails to reconstruct photorealistic video content.

**C2: How to guarantee low-latency streaming with computationally intensive VFMs?** VFMs typically demand substantial computational overhead for tokenization and inference, creating inherent tension with real-time streaming requirements. The computational bottleneck primarily stems from complex encoder-decoder architectures that process visual inputs through multi-stage operations including 3D Haar wavelet transforms, causal residual blocks, and causal spatiotemporal attention mechanisms [1]. Directly applying VFMs to high-resolution video streams exacerbates this issue, as the sequential processing nature and attention complexity create significant processing latency.

**C3: How to ensure loss-resilient and network-adaptive streaming with VFM-based codec?** Traditional streaming systems rely on redundancy-based error correction, but VFM-based systems face unique vulnerabilities due to their semantic compression paradigm. When network conditions fluctuate dramatically, excessive bandwidth leads to underutilization, while insufficient bandwidth may cause packet losses. Packet losses in semantically compressed streams are particularly

Table 1: Comparison of Different Streaming Paradigms.

Technical Paradigm	Fidelity	Efficiency	Robustness
Traditional Codec [35] [30]	Low	High	High
Neural Enhanced [44]	Medium	Medium	High
Generative Neural Codecs [7]	Medium	Medium	High
Diffusion-based [36]	Medium	Low	Low
VFM-based Morphe	High	High	High

devastating, as they can trigger catastrophic quality degradation and temporal inconsistencies that propagate across frames. Current VFMs lack error recovery mechanisms and network-aware adaptation capabilities that can dynamically adjust compression strategies based on channel conditions.

**Design and Contribution:** To address the aforementioned challenges, we propose Morphe<sup>2</sup>, an innovative end-to-end video streaming system empowered by VFMs. To the best of our knowledge, Morphe stands out as the first system to leverage vision foundation models to achieve high-fidelity, low-latency, and loss-resilient advanced video service, as highlighted again in the Table 1. The core of our design lies in three aspects: a generative codec fine-tuned from VFM that simultaneously ensures high compression ratios and high fidelity, a resolution scaling accelerator with neural-enhanced processing that significantly reduces computational complexity, and an adaptive streaming controller that responds to network fluctuations. Specifically, Morphe incorporates the following three innovative design components to efficiently adapt VFMs for multimedia networking.

- **Visual-enhanced Generative Codec:** To address C1, we design an asymmetric spatiotemporal compression strategy that fine-tunes Vision Foundation Models (VFMs) and minimizes visual artifacts during encoding-decoding. This module automatically optimizes compression efficiency while preserving visual quality through advanced generative reconstruction and inter-frame smoothing, achieving unprecedented video fidelity at extreme compression ratios with photorealistic playback quality.
- **Resolution Scaling Accelerator:** To address C2, we propose a comprehensive acceleration framework that can significantly improve system efficiency. This module automatically balances computational complexity and visual quality through intelligent downsampling during preprocessing and deep learning-based super resolution during post-processing, enabling real-time 1080p streaming under resource constraints.
- **Network-Adaptive Streaming Controller:** To address C3, we establish a robust streaming architecture with scalable bitrate control and recovery mechanisms towards network dynamics and bandwidth limitations. This module automatically adapts streaming parameters to network conditions

<sup>2</sup>Morphe derives from the Greek deity of dreams who crafted and shaped human forms. This indicates the core function: to intelligently shape and reconstruct high-fidelity video from highly compressed latent features.

through adaptive bitrate algorithms and self error correction strategies, ensuring seamless video delivery even under severe packet loss conditions.

Morphe reduces bitrate by 62.5% compared to H.265 while maintaining comparable visual quality. The system achieves 65 fps real-time streaming on a server with a single NVIDIA RTX 3090 and 94.2% bandwidth utilization in real network transmission. The contributions of this paper are summarized as follows:

- We identify the key challenges of applying VFMs to video streaming systems and demonstrate that existing compression and streaming approaches are inadequate for achieving high-fidelity, real-time, loss-resilient video delivery under network constraints (§2).
- We then design Morphe, the first VFM-powered end-to-end video streaming system that incorporates temporal/spatial visual enhancement modules for improved codec quality, comprehensive preprocessing and postprocessing optimizations for system efficiency, and an intelligent streaming controller to handle network fluctuations (§4, §5, §6).
- Comprehensive experiments and simulations are provided to demonstrate the system’s applicability and scalability, showing superior quality, real-time HD streaming under bandwidth constraints, and robust performance despite packet loss under challenging environments (§8).

## 2 Background and Motivation

### 2.1 Bandwidth-starved Multimedia Streaming

#### The Bandwidth Concern in Global Video Communication.

Consider business travelers in rural areas joining critical video conferences through degraded 2G networks supporting only 100 kbps—far below the 300 kbps minimum for intelligible video calls. Despite video communication’s ubiquity, billions remain excluded from high-quality visual connectivity due to severe bandwidth limitations.

This constraint extends beyond developing regions. Emergency responders require crystal-clear video feeds through compromised disaster networks. Remote industrial facilities demand real-time equipment monitoring despite limited connectivity. High-speed transportation systems face intermittent connections, yet passenger safety depends on uncompromised video quality. As illustrated in the Figure 1, our investigation and collection of network traces from high-speed rail journeys and vehicle driving in edge areas reveal that many real-world scenarios still suffer from harsh network conditions. These scenarios share a common challenge: delivering high quality video content under extreme bandwidth constraints.

### 2.2 Traditional Video Streaming

**Exhausted Performance Under Low Bandwidth.** Current video coding technologies reveal architectural limitations when pushed to extreme bandwidth constraints. Traditional codecs like H.264 [35] and H.265 [30] employ pixel-based

intra-frame coding and inter-frame prediction paradigms optimized for moderate compression ratios. These codecs achieve compression primarily by eliminating spatial redundancy through correlation between adjacent pixel blocks and temporal redundancy through inter-frame prediction techniques that exploit similarity between consecutive frames. However, this redundancy-removal approach encounters a fundamental ceiling—the information-theoretic physical constraints that bound achievable compression ratios.

Neural-embedded codecs [10, 18, 20, 21] improve upon traditional methods by using neural networks to learn complex motion patterns and correlations, but they remain constrained by the conventional redundancy elimination framework. This architectural limitation bounds their potential, as they still rely on pixel-level reconstruction assumptions. Below 50 kbps, these approaches fail catastrophically, producing videos with severe blocking artifacts, ghosting, and detail loss that make them unusable for practical viewing.

## 2.3 Learning-based Video Streaming

### 2.3.1 Neural-Enhanced Streaming


#### Insufficient Learning and Limited Generalizability.

Neural-enhanced video compression leverages deep learning to achieve superior rate-distortion trade-offs by transmitting highly distorted low-bitrate videos and restoring them using neural network post-processing. Unlike traditional redundancy-based approaches, this method actively discards most information, relying on neural networks’ powerful restoration capabilities to recover lost details while achieving substantial compression ratios.

However, these AI-driven networks face significant vulnerabilities due to domain gaps between training datasets and real-world deployment scenarios [41]. Performance degrades with out-of-distribution content, introducing artifacts like unnecessary high-frequency details and hallucination effects that compromise authenticity. While solutions like NAS [15, 43, 44] propose fine-tuning networks for each video segment, transmitting these adapted models increases bitrate, creating a fundamental trade-off between adaptation capability and bandwidth efficiency.

### 2.3.2 Loss-Resilient Generative Codecs

The evolution of loss-resilient generative codecs has progressed through distinct technological paradigms. Early systems like GRACE [7] pioneered dropout-based training but relied on oversimplified models assuming uniform random packet loss, degrading under real network conditions with temporal clustering. Token-based approaches like Reparo [19] advanced the field using VQ-GAN [11] codebooks to map video content into discrete indices, achieving dramatic bitrate reductions while maintaining perceptual quality. Beyond the aforementioned weak generalization and task-specific limitations, this approach introduces other fundamental problems:



(a) Network trace during train (b) Network trace during countryside self-driving tours

Figure 1: Case Study of Bandwidth-Constrained Multimedia Communication.




Figure 2: Visual perception of different streaming technologies at 400 kbps . Most existing solutions show severe artifacts under poor network conditions.

**Malfunctioning Motion Modeling.** Systems like GRACE model frames independently, resulting in weak temporal consistency that generates severe mosaic artifacts around motion regions under extreme compression. This frame-independent modeling approach, rather than considering the video holistically, leads to catastrophic error propagation - a classic video transmission problem where corrupted frames cascade through temporal dependencies.

**Frame Resolution Inflexibility.** Token-based systems exhibit strict resolution dependencies tied to training configurations. Models trained on 720p content cannot accommodate higher-resolution inputs like 1080p without complete retraining of codebooks and networks. This inflexibility limits deployment scalability, necessitating multiple specialized models for different resolutions and increasing system complexity.

### 2.3.3 Diffusion-based Generative Streaming

Diffusion-based approaches for video streaming have demonstrated significant potential, with DiffusionStream [16] establishing feasibility for real-time transmission. Systems like Promptus [36] advance this paradigm using compact semantic prompts to guide diffusion-based generation, achieving




Figure 3: Video tokenization pipeline in Vision Foundation Models. An input video is encoded into tokens, which are usually much more compact than the input video. The decoder then reconstructs the input video from the tokens.

remarkable generalization across diverse content types and arbitrary resolutions. However, diffusion-based streaming faces two critical limitations:

**Imperfect Efficiency-Quality Trade-offs.** Diffusion models rely on extensive denoising steps, requiring sophisticated optimization and complex hardware acceleration to approach real-time performance. Achieving superior visual quality demands prolonged training periods, restricting deployment to offline playback scenarios. Additionally, diffusion controllability remains weak, frequently producing AI artifacts—temporal inconsistencies or unrealistic textures—that are easily detectable and fall short of photorealistic standards required for professional applications.

**Poor Network Resilience.** Existing diffusion-based designs ignore fluctuating network conditions, lacking robust packet loss resilience. The sequential nature of diffusion generation creates vulnerabilities where prompt corruption or incomplete transmission cascades into complete frame reconstruction failures. Diffusion Model-based systems struggle to maintain semantic coherence when prompt information is partially lost, making them unsuitable for unreliable network environments typical of mobile streaming.

Along the development of video streaming, we continue to witness many limitations. As shown in Figure 2, current streaming systems suffer significant video quality degradation when facing insufficient bandwidth and network jitter.

## 2.4 Advancing Vision Foundation Models Towards High-Quality Video Streaming

Vision Foundation Models (VFM) have emerged as transformative architectures offering unprecedented semantic understanding and generalizability for video compression. These models employ encoder-decoder architectures that compress visual data while preserving semantic content, as illustrated in Fig.3. The encoding function  $f_\phi$  transforms input videos  $x_{0:T} \in \mathbb{R}^{(1+T) \times H \times W \times 3}$  into compact latent representations  $z_{0:T'} \in \mathbb{R}^{(1+T') \times H' \times W' \times C}$  through multi-dimensional downsampling with spatial compression factors  $s_{HW}$  and temporal compression factor  $s_T$ , which are reconstructed via decoder  $g_\theta$ . This architecture enables semantic compression with significant gains based on perceptual importance.

However, directly transplanting VFMs to multimedia

streaming is not feasible. The fundamental tension between VFMs’ understanding-oriented design and streaming’s quality-centric demands reveals deep-rooted incompatibilities. These gaps, as outlined below, represent the core challenges we must address:

**High Fidelity versus Perceptually Lossy VFMs.** Video streaming demands photorealistic playback experiences, yet VFMs primarily focus on video understanding and generation, where low-frequency content representation plays a crucial role, inevitably losing substantial high-frequency details essential for visual fidelity. Current VFM tokenizers fundamentally sacrifice high-quality detail for semantic understanding, creating an unbridgeable gap for high-fidelity streaming applications. Furthermore, the lossy compression inherent in VFM tokenizers leads to perceptual artifacts—spatial blurring, temporal flickering, and phantom details that violate human visual perception expectations. Persistent visual distortions and poor frame-to-frame consistency result in uncomfortable viewing experiences.

**Low Latency versus Computationally Intensive VFMs.** Video streaming requires seamless and smooth playback experiences, yet VFMs are inherently computationally intensive designs that fundamentally contradict real-time streaming requirements. Our investigation reveals that VFM processing times for high-resolution video content are prohibitively high for real-time applications. The inference latency of these AI-native architectures creates insurmountable bottlenecks when processing 1080p streams at standard frame rates, as demonstrated by our investigation shown in Table 2. This establishes an inherent incompatibility between VFM computational demands and the real-time constraints essential for seamless video streaming.

**Loss Resilience versus Network-Agnostic VFMs.** Video streaming systems must adapt to challenging network environments with severe packet loss, demanding sophisticated quality control to maintain QoS. However, current VFMs completely lack network-awareness capabilities, operating as isolated black boxes without network condition integration. Packet loss during transmission can catastrophically corrupt semantic tokens, leading to complete reconstruction failures rather than graceful quality degradation. To enable VFM-based systems to survive real-world networks, loss-resilient strategies and network-adaptive tokenization mechanisms are essential, yet VFMs fundamentally lack these prerequisites for robust streaming systems.

### 3 Morphe Overview

In response to the above challenges and insights, we design three modules for Morphe: Visual-enhanced Generative Codec (§4), Resolution Scaling Accelerator (§5), and Network-Adaptive Streaming Controller (§6). The workflow of Morphe is illustrated in the Figure 4, with specific details listed as follows:

Table 2: Comparative Analysis of Vision Foundation Models for Video Encoding and Decoding.

Model	Precision	Enc.(FPS)	Dec.(FPS)
VideoVAE Plus [38]	fp16	2.12	1.47
Cosmos [1]	fp16	6.21	5.08
CogVideoX-VAE [42]	fp16	5.52	1.95

**Visual-enhanced Generative Codec (VGC)** performs high-fidelity compression on preprocessed videos through residual codec, temporal smoothing, and intelligent token dropping. Unlike traditional VFM models, VGC fine-tuned to focus on visual quality keeps important visual details and allows flexible control of bitrate. The codec generates variable-rate data streams that can be dynamically adjusted based on NASC’s network feedback, achieving extreme compression ratios with minimal degradation.

**Resolution Scaling Accelerator (RSA)** dynamically selects the appropriate resolution for VGC encoding based on bandwidth to balance quality and latency, enabling real-time 1080p streaming under resource constraints. The preprocessed down-sampled frames are sent to VGC for compression, for decoded streams, a neural super-resolution network jointly leverages VGC’s decoder features to restore full resolution.

**Network-Adaptive Streaming Controller (NASC)** orchestrates the streaming of VGC-compressed tokens across fluctuating network conditions. NASC achieves scalable bitrate control by jointly managing VGC’s bitrate adjustment and provides robust streaming protocols with self-recovery mechanisms. When bandwidth varies or packet losses occur, NASC prioritizes critical semantic tokens and reconstructs missing information through temporal-spatial correlations, ensuring stable video delivery even under lossy network conditions.

## 4 Visual-enhanced Generative Codec

As shown in Figure 4 VGC encoder and decoder, this section explains VGC codec in three parts. We begin with the overall architecture of the codec model (§4.1), then describe how we suppress video jitter and improve stability (§4.2), and finally present how to implement scalable video encoding to achieve adjustable compression ratios (§4.3).

### 4.1 Generative Codec Architecture

To balance high compression ratio with visual quality, VGC employs a multi-interface codec framework with asymmetric compression. This section first presents the codec architecture and then describes the multi-interface framework.

**Generative Codec of VGC.** Original video foundation models are primarily designed to extract semantic features of video for subsequent video generation or understanding. To serve diverse downstream tasks, they commonly offer two standard compression settings: (1)  $8 \times$  temporal downsampling with  $16 \times 16$  spatial downsampling, which yields a high compression ratio but introduces noticeable temporal jitter and loss

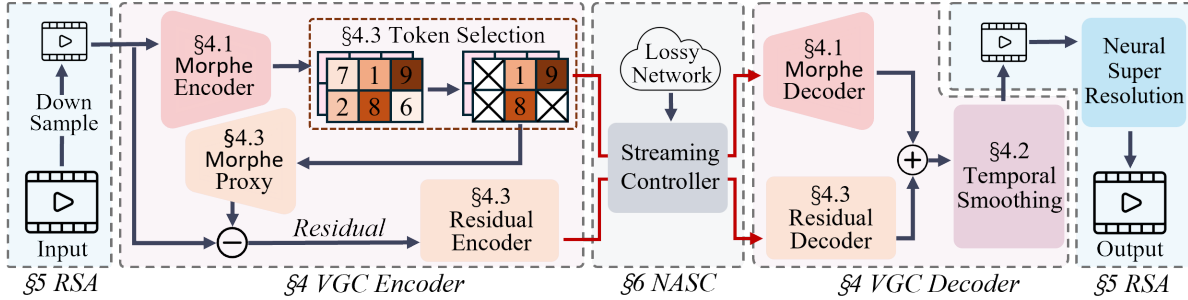


Figure 4: Morphe Overview. It consists of three core modules: Visual-enhanced Generative Codec (§4), Resolution Scaling Accelerator (§5), and Network-Adaptive Streaming Controller (§6).

of spatial detail. And (2)  $4\times$  temporal downsampling with  $8\times 8$  spatial downsampling, which preserves more detail and temporal smoothness but achieves a lower compression ratio that is ill-suited to bandwidth-constrained networks. This fundamental problem between compression and visual quality raises a key design question: *Can we find a balance that preserves critical spatial detail, minimizes temporal jitter, and still fits within narrowband transmission?*

From a systematic analysis of streaming video traffic, we find that the bandwidth required to increase spatial detail far exceed that needed to maintain temporal smoothness. This insight guides our design: prioritize spatial detail and allocate more compression to the temporal dimension. Concretely, we propose an asymmetric configuration that keeps  $8\times 8$  spatial compression while increasing temporal compression to  $8\times$ . This choice maintains a high compression ratio while preserving high-frequency spatial details to meet high-fidelity video demands over narrow bandwidth. Although stronger temporal compression will introduce jitter, the temporal smoothing enhancement techniques in §4.2 effectively mitigate this issue, enabling Morphe to optimize the trade-off between compression efficiency and perceptual quality.

**Multi-objective Optimization Framework.** To enhance video reconstruction quality and system efficiency, we incorporate a multi-interface framework into VGC featuring two critical interfaces: an internal gradient interface for iterative optimization and an external data interface for visual enhancement. The gradient interface enables adaptive training objectives across different stages, while the external interface integrates visual enhancement modules and network control components. VGC’s architecture supports joint scheduling for visual quality optimization and network adaptation, providing unified control for optimization algorithms and ensuring high-quality streaming under challenging network conditions.

## 4.2 Temporal Consistency Enhancement

To mitigate jitter caused by aggressive temporal compression, VGC designs a pixel blending scheme that leverages codec co-optimization. It combines model optimization with inter-frame smoothing to suppress temporal artifacts.

**Pixel Blending and Codec Optimization.** Inspired by con-

ventional video codecs, VGC adopts Group of Pictures (GoP) for segment-based encoding and decoding. However, the combination of per-GoP processing and strong temporal compression inevitably causes temporal jitter at GoP boundaries. A straightforward idea is to blend pixels across adjacent GoPs near their boundary to achieve smooth transitions. However, since VGC encodes each GoP independently, spatial misalignment and brightness inconsistencies arise. Direct blending under these conditions often worsens artifacts and blurring instead of improving visual coherence.

To resolve this, VGC uses the internal gradient interface introduced in §4.1 and proposes an iterative codec co-optimization scheme for boundary-frame blending. The scheme proceeds in two stages.

First, to exploit priors in existing vision foundation model and suppress flicker, we add a dedicated constraint tailored to GoP boundary alignment. Specifically, we compute a pixel-level loss over the boundary region:

$$L_{\text{temporal}} = \frac{1}{n} \sum_{i=1}^n \|\hat{x}_{\text{curr},i} - x_{\text{prev},T-n+i}\| \quad (1)$$

where  $\hat{x}_{\text{curr},i}$  denotes the first  $n$  reconstructed frames of the current GoP and  $x_{\text{prev},T-n+i}$  denotes the last  $n$  frames of the previous GoP. While this pixel-difference supervision does not fully capture temporal coherence, it effectively exploits model priors and reduces training cost. With this constraint, frames near the GoP boundary become closer in pixel space.

Then blending is applied to the optimized boundary region:

$$\hat{x}_{\text{blend},i} = \alpha_i \cdot \hat{x}_{\text{prev},T-n+i} + (1 - \alpha_i) \cdot \hat{x}_{\text{curr},i} \quad (2)$$

where  $\alpha_i$  is the linear interpolation weight, which we set as  $\alpha_i = \frac{n-i}{n}$ . Because training enforces pixel proximity at GoP boundary, this smoothing yields a seamless transition from the previous GoP to the current one. Using this smoothing mechanism, VGC significantly improves temporal smoothness without incurring any additional transmission cost.

## 4.3 Scalable Video Coding

To expand the operating bitrate range and visual quality of codec, Morphe VGC employs scalable video encoding. In this

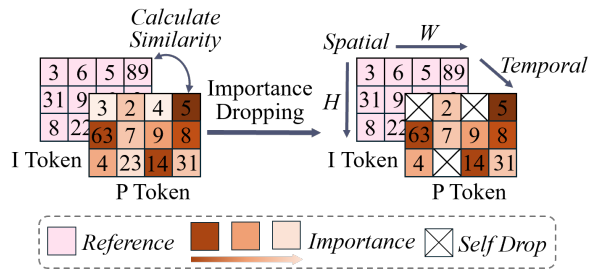


Figure 5: Illustration of Similarity-based Token Selection.

section, we first introduces similarity-based token selection, then describes pixel residual computation and transmission. **Similarity-based Token Selection.** In narrow bandwidth network, even aggressive codec compression may exceed available bandwidth. A straightforward way is to randomly drop semantic tokens, but this ignores semantic importance, risking critical visual content loss and quality degradation. Inspired by I/P frame structures in traditional codecs, we observe significant redundancy within GoP semantic representations. This enables further compression: by identifying and selectively discarding similar tokens, we can preserve the transmission of core visual content even under very limited bandwidth.

VGC adopts an I/P structure like traditional codecs [30, 35]. We set each GoP contains 9 frames. The first frame serves as a reference I frame and is compressed only spatially, yielding a full semantic token matrix. The subsequent 8 frames (P frames) are jointly compressed in spatial and temporal at  $8\times$ , producing another token matrix. This design naturally provides a reference baseline for token importance estimation.

As shown in Figure 5, for the P-frame token matrix, we evaluate the importance of each token by its similarity to the I-frame tokens at the same spatial location. Concretely, for the  $k$ -th GoP, let  $T_k^P$  denote the token matrix from the 8 jointly compressed frames and  $T_k^I$  the token matrix from the first (I) frame. We compute a per-token cosine similarity:

$$S_{i,j} = \frac{T_k^{P(i,j)} \cdot T_k^{I(i,j)}}{\|T_k^{P(i,j)}\| \|T_k^{I(i,j)}\|} \quad (3)$$

where  $T_k^{P(i,j)}$  and  $T_k^{I(i,j)}$  are the token vectors at location  $(i, j)$  in the P and I token. A higher  $S_{i,j}$  indicates greater temporal redundancy relative to the reference; such tokens can be deprioritized and, under bandwidth pressure, dropped first.

Given the computed similarities, the system selects a dynamic threshold  $\tau$  based on the current available bandwidth and marks tokens with  $S_{i,j} > \tau$  as discardable. This cosine-based scoring is efficient, relying only on vector inner products and norms, and thus introduces negligible overhead.

**Pixel Residuals Compression.** To enhance high-frequency details, and provide a solution for scalable bitrate control we introduce a pixel-residual computation mechanism.

On the encoder, as shown in Figure 4, we employ a proxy model to convert encoded feature back to pixels in real time.

Comparing the original frame  $x$  with its reconstruction  $\hat{x}$  yields a per-pixel residual  $r = x - \hat{x}$  that precisely captures information lost during compression. However, directly transmitting raw residuals is impossible: for 1080p at 30 fps, the residual stream is around 1.39 Gbps, which is far beyond narrow-link budgets and at odds with our compression goals.

Therefore, we design a novel residual pipeline. First, we observe that residuals across neighboring frames are often similar, especially in static or slow-motion scenes. Exploiting this, we aggregate residuals over a temporal window of length  $T$  by averaging:

$$r_{\text{compressed}}(h, w, c) = \frac{1}{T} \sum_{t=1}^T r(h, w, c, t) \quad (4)$$

where  $h$  and  $w$  represent the height and width of the video and  $c$  represents the number of channels of the residuals. This temporal compression reduces both volume and noise residuals, since random noise diminishes under averaging. The compressed residual at each pixel is then distributed back to all frames in the window to ensure high visual quality.

Next, we leverage the sparse nature of residuals to achieve additional compression. Because VGC already preserves most of the salient content, the residual number are typically small. We apply thresholding with  $\theta$ , zeroing elements whose absolute value is below  $\theta$ , yielding a highly sparse residual matrix.

Finally, we use arithmetic entropy coding from traditional video codecs to losslessly compress the sparse residuals. Because of its high sparsity, arithmetic coding achieves excellent compression efficiency. Together, these steps compress residuals by nearly  $2000\times$  within acceptable distortion, making transmission feasible even on narrow bandwidth networks.

## 5 Resolution Scaling Accelerator

To address the computational inefficiency of original VFMs, this section presents a resolution scaling-based accelerator that improves system efficiency for our VGC.

**Adaptive Resolution Control.** Reducing input video resolution is the most direct way to lower both encoding compute and transmission bitrate. This offers multiple benefits: reduced bitrates make resolution control the main rate adjustment in Morphe, enabling quick response to bandwidth changes, and lowers encoding and decoding complexity, improving efficiency on resource-limited devices. We develop an adaptive resolution control framework to achieve this.

However, the conventional downsample then super resolution (SR) pipeline faces two key challenges in semantic video coding. First, reconstructions from the VGC decoder often contain noise and artifacts, applying super resolution naively amplifies these degradations and harms perceptual quality. Second, because the codec is built on vision foundation models pretrained on high-quality data, feeding low-resolution inputs after downsampling degrades reconstruction performance and fails to deliver the desired visual quality.

To address these issues, we propose staged optimization where the codec is trained to produce outputs well-suited for super resolution, rather than forcing a lightweight SR model to adapt to codec artifacts. This design is driven by real-time constraints requiring lightweight SR networks with limited depth and parameters, which struggle to learn complex denoising mappings. In contrast, the codec’s richer structure and capacity make it better positioned to learn how to emit outputs matching a target distribution.

The procedure employs a two-stage approach: Stage 1 independently trains a lightweight super resolution model on large-scale data augmented with random noise, compression artifacts, and quality degradations to establish robust priors and define the expected input distribution. Stage 2 freezes the SR weights and jointly fine-tunes the VGC on high-definition video to align its reconstructions with the SR model’s training distribution, reframing the objective as distribution alignment rather than learning restoration mappings from scratch.

We perform joint fine-tuning end-to-end via the internal gradient interface in §4.1. Because the super resolution model is already pretrained in simulated degradations, the codec’s gradient space is substantially reduced, making training more stable and efficient. This reverse-adaptation strategy improves the quality and compute trade-off.

## 6 Network-Adaptive Streaming Controller

In this section, we propose networking integration techniques for Morphe, where NASC enables scalable bitrate control and robust streaming while achieving loss resilience.

### 6.1 Scalable Bitrate Control

To accommodate dynamic network bandwidth, Morphe achieves continuous and scalable bitrate adaptation. In this section, we first introduce how Morphe controls the bitrate by coordinating three rate-control techniques: adaptive resolution control (§5), pixel-residual transmission and intelligent token dropping (§4.3). Then, we describe how to adaptively select the target bitrate based on network bandwidth.

**Rate-control Framework.** Morphe employs an anchor-based linear control scheme. We define two key anchors—the 3× downsampling anchor ( $R_{3x}$ ) and the 2× downsampling anchor ( $R_{2x}$ )—which serve as thresholds for enabling different rate-control strategies. Given the measured available bandwidth  $B_{avail}$ , the controller selects a strategy bundle in real time, with relevant details available in the Appendix A.1.

In the extremely low bandwidth mode, the system encodes at 3× downsampling to encode tokens and then applies similarity-aware token selection from §4.3, adapting the drop rate to the current bandwidth. In the low bandwidth mode, the system preserves the full 3× downsampled semantic representation and allocates the remaining bandwidth to pixel residuals from §4.3, ensuring that salient regions receive higher-fidelity reconstruction. When bandwidth is sufficient, the controller switches to 2× downsampling for higher base

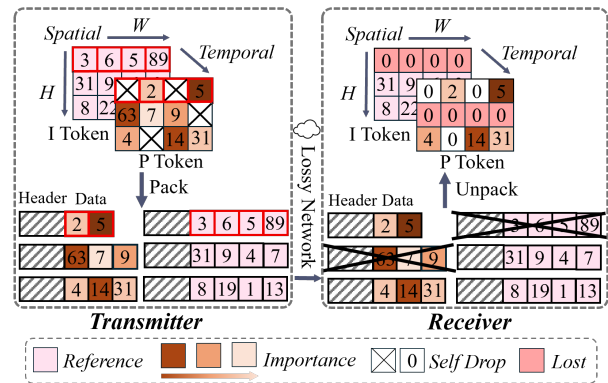


Figure 6: Token Packetization and Hybrid Loss Handling.

quality and continues to devote surplus bandwidth to residual transmission to further enhance detail. To ensure smooth viewing, mode transitions use hysteresis to avoid oscillations due to bandwidth jitter. When the resolution changes, the system applies the temporal smoothing techniques from §4.2 to preserve coherence across frames.

This hierarchical control strategy allows Morphe to combine multiple mechanisms intelligently and to deliver better visual quality and bandwidth efficiency than conventional codecs across diverse network conditions (§8.3).

**Adaptive Bitrate Selection.** NASC adopts a receiver-driven control architecture and uses BBR (Bottleneck Bandwidth and Round-trip propagation time) [6] for bandwidth estimation. The receiver reports the estimated available bandwidth every 100 ms, the sender then invokes the above bitrate controller to adjust encoding parameters on the fly.

### 6.2 Robust Streaming Protocol

Conventional video streaming systems adopt multiple packet loss resilience mechanisms, typically combining forward error correction (FEC), automatic repeat request (ARQ), and decoder-side error concealment [31]. However, under adverse conditions these mechanisms often induce high latency or sharp swings in visual quality [17].

Conventional codecs handle missing pixels through simple interpolation. In contrast, semantic codecs leverage high-level representations and can reconstruct missing content by exploiting surrounding semantic and spatiotemporal context. We therefore design a streaming protocol tailored to a semantic codec, tightly coupling transport with semantic reconstruction. The protocol comprises two core components: an intelligent packetization strategy, and a hybrid loss handling strategy. These components work in concert to deliver resilient video without injecting extra redundancy.

**Intelligent Packetization.** NASC introduces a token-oriented packetization scheme that reframes packet loss as a denoising problem the network has already learned to handle. As shown in Figure 6, for an  $H \times W$  token matrix emitted by the encoder, we package row by row into  $H$  packets of length  $W$ . Each



Figure 7: Showing Devices Used in System Prototype. Two NVIDIA Jetson serve as transmitter and receiver connected via network cable for end-to-end transmission.

packet carries a header and a payload of valid tokens. The header records the row index in the original matrix and a position mask for tokens in the packet, represented as a binary vector of length  $W$  where one denotes a valid token and zero denotes a proactively dropped token.

A key innovation is the unified treatment of missing information. At the receiver, the decoder reassembles each received row using the header’s row index and mask, placing valid tokens at their locations and filling zeros at masked positions corresponding to proactively dropped tokens. If a packet is lost, the entire row is zero filled. From the decoder’s perspective, similarity-based token dropping and network loss appear identically as zero-valued noise.

Because Morphe has been jointly trained to reconstruct from incomplete token matrices (§7), packet loss merely increases the amount of noise to process rather than introducing a new failure mode. The system leverages the codec’s intrinsic capability and does not require additional error-correction layers to remain robust.

**Hybrid Loss Design.** Morphe differentiates loss policies for semantic tokens and residuals. For semantic tokens, the system decodes partial data directly and triggers retransmission only when the loss rate exceeds a preset threshold, typically 50%. For residuals, the policy is looser. If a residual packet is lost, the corresponding frame simply skips residual enhancement rather than waiting for a resend. The basic insight is that semantic tokens carry the core content, whereas residuals primarily add detail. Under constrained or unstable networks, sustaining continuous delivery of core content is more important than perfectly reconstructing detail.

This content-aware streaming protocol yields graceful quality degradation instead of hard stalls when conditions worsen, substantially improving user experience. At the same time, because the system can operate on incomplete data, retransmissions are reduced, which lowers end-to-end latency and improves bandwidth utilization.

## 7 Implementation of Morphe

Morphe is implemented with approximately 35K lines of Python and C++ code, including its codec model, scalable

bitrate adaptation, streaming protocol control and a WebRTC-based prototype.

**Codec Training.** The Morphe codec architecture builds on the recent Cosmos vision foundation model [1]. We fine-tune Morphe from pretraining this VFM on the UltraVideo-Long dataset [39] using a two-stage training strategy. Stage one trains the base model without packet loss to optimize inter-frame temporal smoothness and adaptive resolution support. Stage two introduces random token drop training (§4.3) with drop rates sampled uniformly from [0%, 25%]. Unlike decoder-only error concealment, Morphe employs joint training to align the encoder and decoder for handling token loss. The training objective combines multiple loss functions.<sup>3</sup> We train VGC on 16 NVIDIA A100 GPUs consuming approximately 195 GPU hours total. Due to space limitations, we provide more details about Morphe training in Appendix A.2.

**Codec Wrapper.** To enable end-to-end video streaming, we implement a wrapper that integrates the VGC into the transmission pipeline. We compile Morphe with PyTorch JIT version 2.6 and use BFP16 for encoding and decoding computations. The wrapper maintains model parameter states at both sender and receiver, eliminating initialization overhead. For instance, when the encoder adjusts its downsampling factor, the decoder’s super-resolution module can immediately switch to the corresponding model parameters without reloading network weights over PCIe.

**Morphe Prototype.** We build a Morphe prototype based on WebRTC [3], running on two NVIDIA Jetson AGX Orin (32 GB) devices respectively. As shown in Figure 7, one Jetson serves as the sender and the other as the receiver, each running the wrapped Morphe codec model. During operation, the codec interacts with the scalable bitrate control (§6.1) module through the external data interface to enable dynamic rate adaptation. The two devices connect via network cable. To test our robust streaming protocol (§6.2) in lossy network conditions, we also build a trace-based network emulator. This network emulator acts as a relay, receiving and forwarding packets between the sender and the receiver. In the emulator, packets are randomly dropped and mahimahi [23] is used to replay the real network traces from Puffer [40]. Both sender and receiver use QT to render the playback interface, which displays the real-time coding FPS in the top left corner.

## 8 Evaluation

### 8.1 Setup

**Test Videos:** Our test dataset consists of 100 unique 10-second videos at 1080p resolution, including both 30 and 60 FPS (total 1000 seconds), carefully chosen from four public datasets: UVG [22], UHD [39], UGC [32], and Inter4K [29]. The comprehensive and diverse dataset allows us to test Morphe’s

<sup>3</sup>Please note that the  $\ell_1$  norm between  $\hat{x}$  and  $x$  is strongly correlated with the PSNR of  $\hat{x}$  (with respect to  $x$ ). To avoid this bias, in our evaluation in §8 we report quality improvements using alternative metrics such as VMAF, etc.




Figure 8: Rate-Distortion Performance of Our Generative Codec (UGC Dataset).

performance in real-life video streaming tasks.

**Baselines:** To comprehensively evaluate Morphe’s performance, we selected a set of representative video codecs as baselines. Among these, we first included H.264 [35], H.265 [30] and H.266 [5], which are the most widely adopted codecs (implemented via FFmpeg v7.1.3 [12]). Furthermore, we experimented with recent advanced methods such as Grace [7] and NAS [44], the representative of neural codecs, and Diffusion Model based generative codec Promptus [36].

**Metrics:** Following the latest practices in video streaming performance evaluation [8], we select 3 aspects:

- **Visual Fidelity** is measured using four main metrics: SSIM [34], VMAF [25], LPIPS [47], and DISTS [9]. To assess the temporal stability of neural streaming, we introduce inter-frame temporal consistency metrics. Our primary goal is to maintain a high-quality user viewing experience at low bitrates. These metrics collectively reflect human subjective perception quality for video playback performance.
- **Real-time performance** for Morphe is evaluated by measuring the per-frame processing time for both the encoder and decoder, as well as the per-frame transmission delay under complex network conditions.
- **Video Loss-resilience** refers to the system’s robustness and stability when facing packet loss conditions. We report both the visual quality under packet loss and the client-side rendered frame rate under poor network conditions.

## 8.2 Evaluation of Visual Performance

**Compression versus Visual Quality.** We conducted extensive experiments to validate that our system effectively balances the trade-off between aggressive compression and high fidelity reconstruction. Figure 8 demonstrates that our generative codec significantly outperforms traditional codecs and generative baselines across most metrics on the UGC dataset within the selected bandwidth range. Our method achieves a VMAF score of 85.17 compared to H.266’s 57.61 and H.265’s 55.85 at 400kbps, indicating superior perceptual quality under low bandwidth constraints. Although H.266 exhibits excellent performance in terms of SSIM and DISTS metrics, its computational complexity and the current limitations in hardware acceleration development constrain its practical efficiency. Notably, when transitioning from 350kbps to 400kbps through




Figure 9: Visual Metrics in Four Datasets.

quality control mechanisms, the video downsampling ratio has been adjusted from  $3\times$  to  $2\times$ . Benefiting from efficient spatiotemporal compression and downsampling preprocessing, achieving over thousand-fold compression demonstrates promising prospects for applications.

**Morphe Generalizability.** We validated our model’s performance across four diverse datasets. Figure 9 demonstrates our method’s superior generalization capability, consistently achieving higher average VMAF scores compared to other methods while maintaining competitive performance with generative baselines across all perceptual metrics (SSIM, LPIPS, DISTS), thereby validating robust cross-domain applicability. We provide more details about the visual comparison of Morphe and other baselines for video frames transmitted at 400kbps bitrate in Appendix A.3.

**Video Consistency.** As illustrated in Figure 10, we present the distribution curves of visual temporal consistency across two metrics. We compute the inter-frame residuals between consecutive frames of each videos and assess the quality by comparing them to the residuals of the original video. Analysis of pixel-level variations between consecutive frames reveals that all neural-based systems, including GRACE and Promptus, demonstrate considerably more temporal flickering than traditional codecs. To enhance the temporal coherence of vision foundation models, we introduce a temporal smoothing technique (§4.2). Our system reduces video flickering and improves visual quality compared to the version without this technique. The proposed Morphe system achieves




Figure 10: Comparison of Video Temporal Consistency.

Table 3: Computational Overhead for Different Devices

Device	Scale	GPU Memory (GB)	Encoder (FPS)	Decoder (FPS)
RTX3090	3x	8.86	98.51	65.74
	2x	17.09	47.14	32.03
A100	3x	7.96	101.23	83.33
	2x	16.24	52.54	40.19
Jetson	3x	15.21	61.17	43.45
	2x	23.87	31.87	24.93

superior pixel structure stability and consistency, substantially outperforming existing neural-based streaming methods.

### 8.3 Evaluation of Networking System

**Real-time Performance.** We conducted comprehensive experiments across different platforms, including Nvidia RTX 3090, A100, and Jetson Orin devices, with results shown in Table 3. Our evaluation focuses on key streaming metrics including encoding/decoding frame rates, memory consumption, and scalability under different resolution scaling factors (2× and 3×). The results demonstrate that our system achieves real-time performance across all platforms: the A100 delivers 101 FPS encoding and 83 FPS decoding at 3× scaling, while the RTX 3090 provides balanced performance suitable for consumer-grade streaming. Notably, even the resource-constrained Jetson Orin maintains practical frame rates (61/43 FPS at 3× scaling), validating our system’s deployability on edge devices. These measurements demonstrate the system’s capability for live video streaming and provide optimal deployment strategies for various streaming applications ranging from cloud servers to edge devices.

**Loss Resilience.** To evaluate the loss resilience of the systems in complex network environment, we tested the frame delay and video stuttering (measured by average rendering frame rate [26]) under packet loss rates ranging from 5% to 25% at 400kbps. (Morphe can tolerate packet loss up to 30%, but we cap it at 25% to ensure visual quality.) Figure 11 demonstrates that our system maintains low latency even under severe packet loss conditions. At 25% packet loss, Morphe and Grace achieve sub-150ms delay for over 90% of frames, while H.266 suffers substantial degradation with only 40% of frames meeting this threshold. This validates the effectiveness of our loss-resilient streaming design. Figure 12 demonstrates the rendered frame rate performance under packet loss, measured by average FPS at two target frame rates. Our system and Grace maintain stable performance across all loss rates, sustaining near-target FPS even at 25%




Figure 11: Frame Latency Distribution for Video Streaming at Different Packet Loss Rates. The dashed lines indicate the median values.




Figure 12: Comparison of Decoded and Rendered Frame Rates Between Videos Encoded at 30 fps and 60 fps. The decline in rendered frames suggests that packet loss has either prevented frame rendering entirely or caused substantial visual quality degradation.

loss. In contrast, H.266 suffers severe degradation, dropping to nearly 1 FPS at 25% loss, rendering the video essentially unwatchable.

Beyond system transmission performance in complex network environments, this study evaluates visual quality under packet loss. We tested multiple visual metrics at 400kbps with packet loss rates from 5% to 25%, as shown in Figure 13. For key metrics measuring video semantic information and objective quality—VMAF, LPIPS, and DIST5—Morphe not only achieved optimal performance but also exhibited minimal degradation, demonstrating strong stability. In contrast, traditional video codecs showed significant performance decline as packet loss increased. While the GRACE model degraded more gradually, its pixel-based nature resulted in inherently lower visual quality under poor networks.

By integrating semantic streaming with VFMs’ generative capabilities, Morphe achieves superior high-fidelity video restoration in low-bandwidth, lossy networks while maintaining low latency and smooth playback, ensuring excellent visual quality and user experience.

**Bitrate Control.** To evaluate system adaptation to dynamic network conditions, we used MahiMahi network emulator to simulate bitrate fluctuations between 200-500 kbps with 30-second periods. GRACE was excluded due to unavailable open-source bitrate control. As shown in Figure 14, Morphe demonstrated excellent bitrate tracking, with curve closely aligning with targets. This indicates effective scalable bitrate control that accurately adjusts to bandwidth changes, ensuring stable streaming under complex network conditions. In contrast, traditional encoders showed significant perfor-

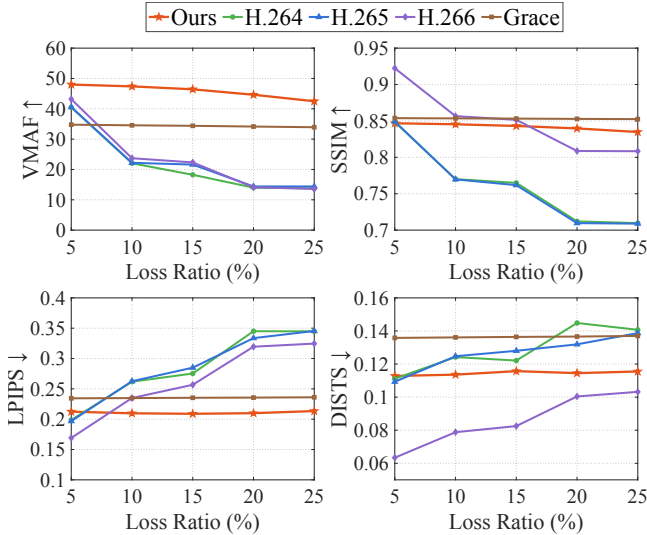


Figure 13: Visual Performance Benchmark under Different Packet Loss Rates.




Figure 14: Network Bandwidth Trace Tracking Experiment.

mance variations. H.264 and H.266 tracked the target relatively well but exhibited notable fluctuations and deviations. H.265 showed severe oscillations with large discrepancies from target bitrates, reaching up to 859.5 kbps during rapid changes, causing network congestion and packet loss.

## 8.4 Ablation Study

Table 4 presents an ablation study evaluating each component’s contribution to system performance, with latency in table referring to encoder/decoder delays. Removing Intelligent Self Drop (§4.3) causes the most severe degradation, with VMAF dropping from 60.76 to 20.31 and LPIPS/DISTS deteriorating significantly, demonstrating its critical role in maintaining visual quality. The RSA (§5) module improves system efficiency by reducing latency, making computationally intensive AI models more suitable for real-time streaming. Although the Residual (§4.3) incurs extra computational cost from the proxy model, it significantly enhances bandwidth utilization and enables finer-grained quality scheduling. Together, they validate the necessity of each component in balancing quality and latency trade-offs.

We tested Morphe under 50% packet loss using both Intelligent Self Drop and Random Drop strategies. As shown in Figure 16 in the Appendix A.4, Intelligent Drop outperformed Random Drop in both visual quality and quantitative

Table 4: Ablation Study of Individual Module Contributions

Method	VMAF↑	SSIM↑	LPIPS↓	DISTS↓	Latency (ms)
w/o RSA	59.72	0.84	0.22	0.14	644.61 / 874.76
w/o Residual	60.54	0.85	0.20	0.13	<b>78.21 / 98.36</b>
w/o Self Drop	20.31	0.73	0.41	0.23	89.82 / 136.87
Morphe	<b>60.76</b>	<b>0.86</b>	<b>0.18</b>	<b>0.11</b>	91.36 / 136.91

Note: Latency values represent encoding/decoding latency (in milliseconds) for a chunk of 9 frames, respectively.

metrics. The semantic importance-based self-drop mechanism effectively minimized visual degradation in critical regions.

To evaluate the effectiveness of temporal smoothing, we conduct temporal video segment analysis and visualization. As shown in Figure 17 in Appendix A.4 and the quantitative results in Figure 10, the system without this mechanism exhibits significantly greater temporal flickering effects, severely degrading the viewing experience.

## 9 Conclusion and Discussion

We present Morphe, the first VFM-powered video streaming system delivering high-fidelity, low-latency transmission under challenging network conditions. Our system achieves a 62.5% bitrate reduction compared to H.265 while maintaining comparable visual quality, with real-time streaming and 94.2% bandwidth utilization in live deployments.

Our VFM-driven approach has trade-offs: fine-grained details like small signage may be lost, and text-heavy content remains challenging due to limited textual training in current VFMs—a common limitation across generative streaming methods. Despite these constraints, Morphe represents a significant advance in video streaming, demonstrating that VFM-based video delivery is ready for real-world applications.

## Acknowledgments

We sincerely thank our shepherd Junchen Jiang, anonymous reviewers for their valuable feedback. The work was supported in part by the National Natural Science Foundation of China (Grant No. 62293482), the Basic Research Project No. HZQB-KCZYZ-2021067 of Hetao Shenzhen-HK S&T Cooperation Zone, the National Key Research and Development Program of China (Grant No. 2024YFB2907000), the National Natural Science Foundation of China (Grant No. 62471423), the Shenzhen Science and Technology Program (Grant No. JCYJ20241202124021028 and Grant No. JCYJ20230807114204010), the Guangdong Talents Program (Grant No. 2024TQ08X346), the Shenzhen Outstanding Talents Training Fund 202002, the Young Elite Scientists Sponsorship Program of CAST (Grant No. 2022QNRC001), the Guangdong Provincial Key Laboratory of Future Networks of Intelligence (Grant No. 2022B1212010001) and the Shenzhen Key Laboratory of Big Data and Artificial Intelligence (Grant No. SYSPG20241211173853027). The work of Xing-gong Zhang and Jiangkai Wu was supported by the National Natural Science Foundation of China (Grant No. 62431017).

## References

- [1] Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, Yin Cui, Yifan Ding, et al. Cosmos world foundation model platform for physical ai. arXiv preprint arXiv:2501.03575, 2025.
- [2] Emanuele Artioli. Generative ai for http adaptive streaming. In Proceedings of the 15th ACM Multimedia Systems Conference, pages 516–519, 2024.
- [3] Adam Bergkvist, Daniel C. Burnett, Cullen Jennings, Anant Narayanan, and Bernard Aboba. WebRTC 1.0: Real-time communication between browsers. W3C recommendation, W3C, 2021. Accessed: 2025-09-19.
- [4] Rishi Bommasani. On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258, 2021.
- [5] Benjamin Bross, Ye-Kui Wang, Yan Ye, Shan Liu, Jianle Chen, Gary J Sullivan, and Jens-Rainer Ohm. Overview of the versatile video coding (vvc) standard and its applications. IEEE Transactions on Circuits and Systems for Video Technology, 31(10):3736–3764, 2021.
- [6] Neal Cardwell, Yuchung Cheng, C Stephen Gunn, Soheil Hassas Yeganeh, and Van Jacobson. Bbr: Congestion-based congestion control: Measuring bottleneck bandwidth and round-trip propagation time. Queue, 14(5):20–53, 2016.
- [7] Yihua Cheng, Ziyi Zhang, Hanchen Li, Anton Arapin, Yue Zhang, Qizheng Zhang, Yuhang Liu, Kuntai Du, Xu Zhang, Francis Y Yan, et al. {GRACE}:{Loss-Resilient}{Real-Time} video through neural codecs. In 21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24), pages 509–531, 2024.
- [8] Luca De Cicco and Saverio Mascolo. An adaptive video streaming control system: Modeling, validation, and performance evaluation. IEEE/ACM Transactions On Networking, 22(2):526–539, 2013.
- [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.
- [10] Abdelaziz Djelouah, Joaquim Campos, Simone Schaub-Meyer, and Christopher Schroers. Neural inter-frame compression for video coding. In Proceedings of the IEEE/CVF international conference on computer vision, pages 6421–6429, 2019.
- [11] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 12873–12883, 2021.
- [12] FFmpeg Developers. Ffmpeg (version 7.1.3). Software available at <https://ffmpeg.org/>, 2025.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In European conference on computer vision, pages 630–645. Springer, 2016.
- [14] Seunghwa Jeong, Bumki Kim, Seunghoon Cha, Kwanggyoon Seo, Hayoung Chang, Jungjin Lee, Younghui Kim, and Junyong Noh. Real-time cnn training and compression for neural-enhanced adaptive live streaming. IEEE Transactions on Pattern Analysis and Machine Intelligence, 46(9):6023–6039, 2024.
- [15] Jaehong Kim, Youngmok Jung, Hyunho Yeo, Juncheol Ye, and Dongsu Han. Neural-enhanced live streaming: Improving live video ingest via online learning. In Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication, pages 107–125, 2020.
- [16] Akio Kodaira, Chenfeng Xu, Toshiki Hazama, Takanori Yoshimoto, Kohei Ohno, Shogo Mitsuohori, Soichi Sugano, Hanying Cho, Zhijian Liu, and Kurt Keutzer. Streamdiffusion: A pipeline-level solution for real-time interactive generation. arXiv preprint arXiv:2312.12491, 2023.
- [17] Marc Lambooj, Wijnand IJsselsteijn, Marten Fortuin, Ingrid Heynderickx, et al. Visual discomfort and visual fatigue of stereoscopic displays: a review. Journal of imaging science and technology, 53(3):30201–1, 2009.
- [18] Jiahao Li, Bin Li, and Yan Lu. Neural video compression with feature modulation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 26099–26108, 2024.
- [19] Tianhong Li, Vibhaalakshmi Sivaraman, Pantea Karimi, Lijie Fan, Mohammad Alizadeh, and Dina Katabi. Reparo: Loss-resilient generative codec for video conferencing. arXiv preprint arXiv:2305.14135, 2023.
- [20] Jianping Lin, Dong Liu, Houqiang Li, and Feng Wu. M-lvc: Multiple frames prediction for learned video compression. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 3546–3554, 2020.

- [21] Guo Lu, Wanli Ouyang, Dong Xu, Xiaoyun Zhang, Chunlei Cai, and Zhiyong Gao. Dvc: An end-to-end deep video compression framework. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 11006–11015, 2019.
- [22] 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.
- [23] Ravi Netravali, Anirudh Sivaraman, Somak Das, Ameesh Goyal, Keith Winstein, James Mickens, and Hari Balakrishnan. Mahimahi: Accurate Record-and-Replay for HTTP. In 2015 USENIX Annual Technical Conference (USENIX ATC 15), pages 417–429, Santa Clara, CA, July 2015. USENIX Association.
- [24] Huu-Tai Phung, Yu-Hsiang Lin, Yen-Kuan Ho, and Wen-Hsiao Peng. Exploring autoregressive vision foundation models for image compression. arXiv preprint arXiv:2509.05169, 2025.
- [25] Reza Rassool. Vmaf reproducibility: Validating a perceptual practical video quality metric. In 2017 IEEE international symposium on broadband multimedia systems and broadcasting (BMSB), pages 1–2. IEEE, 2017.
- [26] Michael Rudow, Francis Y Yan, Abhishek Kumar, Ganesh Ananthanarayanan, Martin Ellis, and KV Rashmi. Tambur: Efficient loss recovery for videoconferencing via streaming codes. In 20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23), pages 953–971, 2023.
- [27] Sandvine. 2023 global internet phenomena report. Technical report, 2023.
- [28] Kevin Spiteri, Rahul Uргаonkar, and Ramesh K Sitaraman. Bola: Near-optimal bitrate adaptation for online videos. IEEE/ACM transactions on networking, 28(4):1698–1711, 2020.
- [29] Alexandros Stergiou and Ronald Poppe. Adapool: Exponential adaptive pooling for information-retaining downsampling. IEEE Transactions on Image Processing, 32:251–266, 2022.
- [30] Vivienne Sze, Madhukar Budagavi, and Gary J Sullivan. High efficiency video coding (hevc). Integrated circuit and systems, algorithms and architectures, 39:40, 2014.
- [31] Yao Wang and Qin-Fan Zhu. Error control and concealment for video communication: A review. Proceedings of the IEEE, 86(5):974–997, 1998.
- [32] Yilin Wang, Sasi Inguva, and Balu Adsumilli. Youtube ugc dataset for video compression research. In 2019 IEEE 21st international workshop on multimedia signal processing (MMSP), pages 1–5. IEEE, 2019.
- [33] Yuancheng Wang, Haoyue Zhan, Liwei Liu, Ruihong Zeng, Haotian Guo, Jiachen Zheng, Qiang Zhang, Xueyao Zhang, Shunsi Zhang, and Zhizheng Wu. Maskgct: Zero-shot text-to-speech with masked generative codec transformer. arXiv preprint arXiv:2409.00750, 2024.
- [34] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing, 13(4):600–612, 2004.
- [35] Thomas Wiegand, Gary J Sullivan, Gisle Bjontegaard, and Ajay Luthra. Overview of the h. 264/avc video coding standard. IEEE Transactions on circuits and systems for video technology, 13(7):560–576, 2003.
- [36] Jiangkai Wu, Liming Liu, Yunpeng Tan, Junlin Hao, and Xinggong Zhang. Promptus: Can prompts streaming replace video streaming with stable diffusion. arXiv preprint arXiv:2405.20032, 2024.
- [37] Yecheng Wu, Zhuoyang Zhang, Junyu Chen, Haotian Tang, Dacheng Li, Yunhao Fang, Ligeng Zhu, Enze Xie, Hongxu Yin, Li Yi, et al. Vila-u: a unified foundation model integrating visual understanding and generation. arXiv preprint arXiv:2409.04429, 2024.
- [38] Yazhou Xing, Yang Fei, Yingqing He, Jingye Chen, Jiaxin Xie, Xiaowei Chi, and Qifeng Chen. Large motion video autoencoding with cross-modal video vae. arXiv preprint arXiv:2412.17805, 2024.
- [39] Zhucun Xue, Jiangning Zhang, Teng Hu, Haoyang He, Yinan Chen, Yuxuan Cai, Yabiao Wang, Chengjie Wang, Yong Liu, Xiangtai Li, et al. Ultravideo: High-quality uhd video dataset with comprehensive captions. arXiv preprint arXiv:2506.13691, 2025.
- [40] Francis Y Yan, Hudson Ayers, Chenzhi Zhu, Sadjad Fouladi, James Hong, Keyi Zhang, Philip Levis, and Keith Winstein. Learning in situ: a randomized experiment in video streaming. In 17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20), pages 495–511, 2020.
- [41] Xi Yang, Wangmeng Xiang, Hui Zeng, and Lei Zhang. Real-world video super-resolution: A benchmark dataset and a decomposition based learning scheme. In Proceedings of the IEEE/CVF international conference on computer vision, pages 4781–4790, 2021.

- [42] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024.
- [43] Hyunho Yeo, Chan Ju Chong, Youngmok Jung, Juncheol Ye, and Dongsu Han. Nemo: enabling neural-enhanced video streaming on commodity mobile devices. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, pages 1–14, 2020.
- [44] Hyunho Yeo, Youngmok Jung, Jaehong Kim, Jinwoo Shin, and Dongsu Han. Neural adaptive content-aware internet video delivery. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, pages 645–661, 2018.
- [45] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021.
- [46] Aiping Zhang, Zongsheng Yue, Renjing Pei, Wenqi Ren, and Xiaochun Cao. Degradation-guided one-step image super-resolution with diffusion priors. *arXiv preprint arXiv:2409.17058*, 2024.
- [47] 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.
- [48] Pengyuan Zhou, Lin Wang, Zhi Liu, Yanbin Hao, Pan Hui, Sasu Tarkoma, and Jussi Kangasharju. A survey on generative ai and llm for video generation, understanding, and streaming. *arXiv preprint arXiv:2404.16038*, 2024.

## A Appendix

### A.1 Scalable Bitrate Control Algorithm

Algorithm 1 presents Morphe’s bitrate control procedure. The COMPUTE\_RESIDUAL and TOKEN\_SELECTION modules are from the scalable video coding techniques described in §4.3.

### A.2 Morphe Training Flow

**Base Codec Training Framework.** Morphe fine-tunes a pre-trained visual foundation model (VFM) backbone via the internal gradient interface described in §4.1. The first training stage focuses on two core capabilities: inter-chunk temporal

---

#### Algorithm 1 Scalable Bitrate Control

---

**Require:** Available bandwidth  $B_{\text{avail}}$ , video frame  $F$

**Ensure:** Encoded bitstream

```

1: if  $B_{\text{avail}} < R_{3x}$  then
2:    $y \leftarrow \text{ENCODE\_3X}(F)$ 
3:    $y_{\text{filtered}} \leftarrow \text{TOKEN\_SELECTION}(y, B_{\text{avail}})$ 
4:   return  $y_{\text{filtered}}$ 
5: else if  $R_{3x} \leq B_{\text{avail}} < R_{2x}$  then
6:    $y \leftarrow \text{ENCODE\_3X}(F)$ 
7:    $\text{residual} \leftarrow \text{COMPUTE\_RESIDUAL}(F, y, B_{\text{avail}} - R_{3x})$ 
8:   return  $\{y, \text{residual}\}$ 
9: else
10:   $y \leftarrow \text{ENCODE\_2X}(F)$ 
11:   $\text{residual} \leftarrow \text{COMPUTE\_RESIDUAL}(F, y, B_{\text{avail}} - R_{2x})$ 
12:  return  $\{y, \text{residual}\}$ 
13: end if

```

---

smoothing (§4.2) and adaptive resolution control (§5). The optimization objective is defined as:

$$L = \alpha L_{\text{pixel}}(\hat{x}, x) + (1 - \alpha) L_{\text{flow}}(\hat{x}, x) \quad (5)$$

where  $L_{\text{pixel}}$  denotes the  $\ell_1$  pixel reconstruction loss,  $L_{\text{flow}}$  denotes the optical flow consistency loss, and the weighting parameter  $\alpha$  is set to 0.8. The flow loss encourages temporal coherence across Group of Pictures (GoP) boundaries, while the pixel loss ensures spatial detail preservation.

To accommodate adaptive video encoding, the training data is augmented by resizing the original 1080P videos to 540P and 720P resolutions. During training, video frames of varying resolutions are randomly sampled to enhance the model’s adaptability to multi-resolution inputs.

Furthermore, an adversarial GAN loss [46] term  $L_{\text{adv}}$  is incorporated to enrich video detail fidelity. A hyperparameter  $\gamma = 0.1$  balances the contribution of the GAN loss. This adversarial objective promotes the generation of visually plausible details, enabling the model to capture finer and more natural textures and dynamic patterns.

Accordingly, the overall training loss is formulated as:

$$L_{\text{total}} = L + \gamma \cdot L_{\text{adv}} \quad (6)$$

This training flow allows Morphe to maintain spatiotemporal consistency while flexibly supporting multi-resolution video inference, significantly enhancing the richness and realism of generated video details.

**Video Super Resolution Model Training.** We first use the Codec model trained in the previous stage to encode 540P videos at the lowest bitrate. This encoding and decoding process introduces some loss and compression artifacts. The resulting degraded videos form the training dataset for the super-resolution model. The super resolution (SR) model is built using a residual convolutional network [13]. It is initially trained on this dataset to learn how to restore the original

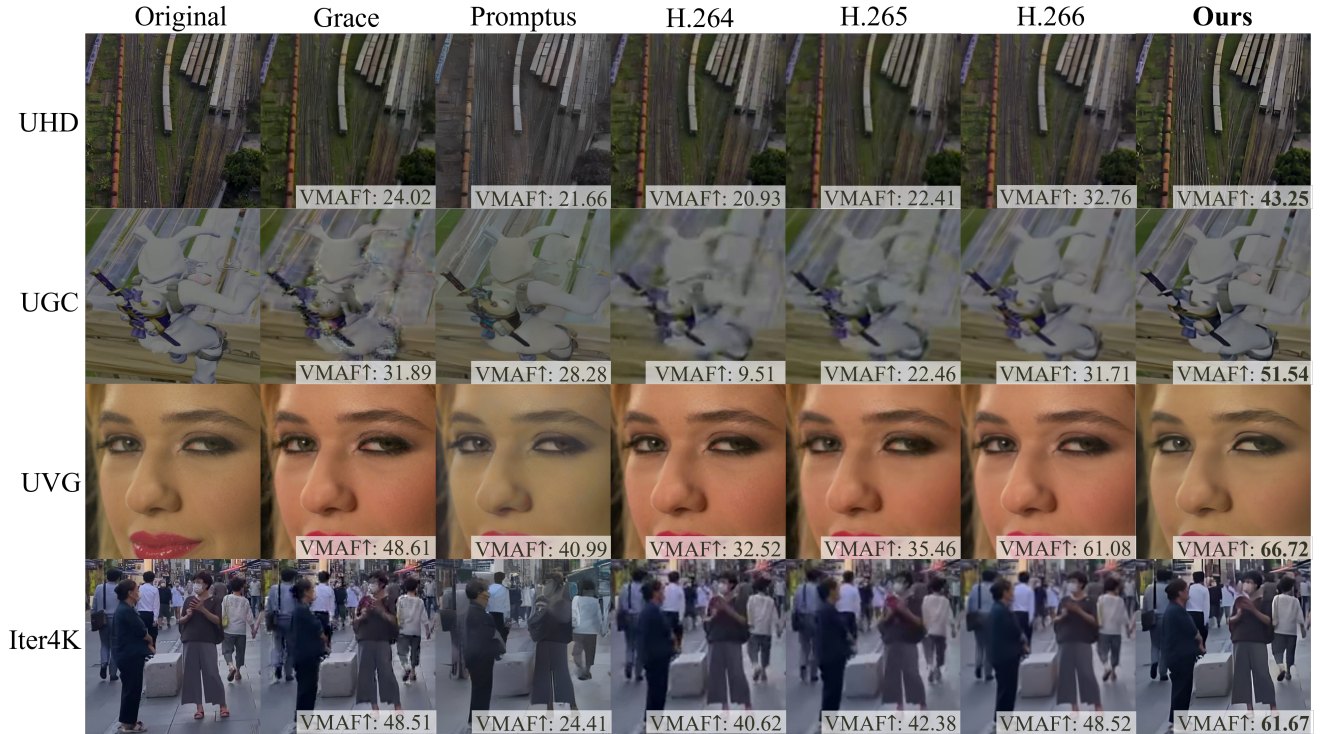


Figure 15: Visual quality comparison on four datasets.

videos from blurred, jittered, and compressed inputs. The loss during this stage follows the same formulation as in Equation (5), which combines pixel reconstruction and optical flow consistency losses to preserve spatial details and temporal coherence. After pretraining, the weights of the SR model are fused with the decoder. We freeze the SR model weights and only fine-tune the codec weights. During this joint training, the original 1080P videos are linearly downsampled by random scales of  $2\times$  or  $3\times$ . These low-resolution frames are fed into the codec model for inference, producing super-resolved videos. To further improve SR performance, the output videos are supervised with the corresponding original high-resolution videos using the GAN loss defined in Equation (6). The training process leverages the internal gradient interface introduced in §4.1 for gradient backpropagation and iterative optimization. By freezing the SR model and only updating the codec weights, the encoder learns to generate pixel features that match the expected distribution of the SR model. This helps enhance the SR restoration quality and final video fidelity.

**Robustness Training via Random Drop.** After establishing basic codec capability, stage two introduces token-drop training to enhance loss resilience.

For a P-frame semantic tensor  $y \in \mathbb{R}^{H \times W \times T}$ , given similarity scores  $S_{i,j}$  and a target drop rate  $\gamma$ , we selectively discard high-similarity tokens. Dropped positions are zero filled to form a partial representation  $\tilde{y}$ . The decoder  $g_\theta(\cdot)$  consumes

$\tilde{y}$  and learns to reconstruct the original video, using the same objective as stage one:

$$L_{\text{robust}} = \alpha L_{\text{pixel}}(g_\theta(\tilde{y}), x) + (1 - \alpha) L_{\text{flow}}(g_\theta(\tilde{y}), x) \quad (7)$$

Crucially, via the internal gradient interface, gradients back-propagate into the encoder so it learns to emit semantic representations that are robust to token loss. During training, we first perform inference at the lowest bitrate. In this process, we apply a similarity-based token dropping method §4.3 to simulate autonomous packet loss. To adapt to varying network packet loss conditions, we uniformly sample the drop rate from the range  $[0, 25\%]$ . This equips the model with the ability to handle different levels of missing information.

Why is Morphe more resilient to loss? Unlike decoder-only error concealment, joint training aligns the encoder and decoder to cope with token drops. First, the decoder learns to exploit reference information in the I-frame semantic matrix to infer and complete missing tokens in P frames. Second, the encoder learns to organize the semantic space so that redundant content shared by I and P frames lies closer, which improves the accuracy of similarity estimation.

This co-optimization has a direct benefit under bandwidth constraints. The system can safely discard high-similarity tokens that are redundant and prioritize lower-similarity tokens that carry novel information. Because the encoder imposes a more structured semantic organization, the decoder can more



Figure 16: Our Token Dropping Solution vs Random Drop.

accurately recover dropped P-frame content from retained I-frame information. This adaptive allocation, coupled with cooperative training, allows Morphe to maintain acceptable visual quality even with token loss rates up to 30%, achieving robustness at the semantic and pixel layer.

Throughout the entire training process, we use learning rates that decay gradually from  $1 \times 10^{-5}$  to  $2 \times 10^{-8}$ . This learning rate schedule helps stabilize training by allowing larger parameter updates in the early stages for faster convergence, while progressively reducing the update magnitude towards the end to fine-tune the model and avoid overshooting. As a result, the model can achieve better generalization and improved performance.

### A.3 Details of Visual Comparison

**Visual quality comparison.** Figure 15 compares the visual quality of Morphe against four baseline methods across different datasets. We evaluate each method on four representative datasets. All methods are compared at the same bitrate (400Kbps) using VMAF as the perceptual quality metric. Morphe achieves the highest VMAF scores across all test scenarios. Note that while the Morphe codec can encode and decode videos of arbitrary sizes and dimensions, we crop 1080p videos to square format for presentation clarity. The visual comparisons show that Morphe maintains significant advantages even on challenging user generated content like UGC, demonstrating its robustness to diverse video types. Traditional codecs H.264, H.265 and H.266 perform worst across all tests. The neural codec Grace outperforms traditional methods but still lags considerably behind ours. For YUV videos in the UVG dataset, all methods except ours introduce severe color shifts that degrade viewing experience.

Traditional encoders exhibit uniform quality degradation with perceptible compression artifacts and blocking effects. GRACE shows significant mosaic distortions between moving objects and backgrounds, while Promptus, though preserving fine details, suffers from semantic inconsistencies due to diffusion model limitations.

### A.4 Ablation Study

To evaluate the efficacy of our intelligent token dropping mechanism, we compare it against naive random dropping at




Figure 17: Ablation Study of Temporal Smoothing

50% reduction rate in Figure 16. The results reveal substantial performance differences: our method achieves VMAF of 50.17 and LPIPS of 0.18, whereas random dropping yields VMAF of only 20.31 and LPIPS of 0.40. The visual comparison clearly demonstrates that random dropping introduces severe corruption artifacts, particularly affecting texture-rich regions such as clothing and background details. Conversely, our intelligent approach maintains perceptual coherence by prioritizing semantically salient tokens, achieving approximately 2.5× higher VMAF scores and 55% lower perceptual distortion. These results underscore the importance of content-aware token selection for quality-constrained video transmission.

Figure 17 presents a visual comparison of temporal consistency with and without our smoothing mechanism. The temporal profile visualization reveals that the system without temporal smoothing (Ours w/o) exhibits significantly more color variation and discontinuities across consecutive frames, as evidenced by the pronounced fluctuations in the profile bands. In contrast, our complete system demonstrates substantially smoother transitions between frames, with more consistent color patterns throughout the sequence. This reduction in frame-to-frame variation directly translates to reduced flickering artifacts and improved viewing experience.