

000 IMPROVING AUTOREGRESSIVE VIDEO MODELING 001 002 WITH HISTORY UNDERSTANDING 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009

011 Video autoregressive generation (VideoAR) sequentially predicts future frames
012 conditioned on history frames. Despite the advance of recent diffusion-based
013 VideoAR, the role of conditioning signal—internal representations of history
014 frames—remains underexplored. Inspired by the success of strong condition
015 representations in text-conditioned generation, we investigate: *Can better inter-*
016 *internal representations of history frames improve VideoAR performance?* Through
017 systematic analysis, we show that history representation quality positively cor-
018 relates with VideoAR, and that enhancing these representations provides gains
019 that cannot be achieved by refining future frames representations alone. Based on
020 these insights, we propose **MiMo** (Masked History Modeling), a novel framework
021 that seamlessly integrates representation learning into diffusion-based VideoAR.
022 MiMo applies masks to history frame tokens and trains the model to predict
023 masked tokens of current and future frames alongside the diffusion objective,
024 yielding predictive and robust history representations without relying on vision
025 foundation models (VFM) or heavy architectural changes. Extensive experiments
026 demonstrate that MiMo achieves competitive performance in video prediction and
027 generation tasks while substantially improving training efficiency. Our work un-
028 derscores the importance of history representations in VideoAR.

029 1 INTRODUCTION 030

031 Video autoregressive generation (VideoAR) predicts future frames conditioned on previously ob-
032 served or generated frames (the history). The *history-to-future* generation process naturally aligns
033 with the causal structure of video dynamics and enables variable-length generation (Villegas et al.,
034 2022; Yin et al., 2025; Teng et al., 2025). However, early AR approaches (Yan et al., 2021; Hong
035 et al., 2022; Ge et al., 2022; Villegas et al., 2022) significantly underperformed non-AR methods
036 (Brooks et al., 2024; Ho et al., 2022; He et al., 2022b; Guo et al., 2024), primarily due to the dif-
037 ficulty of modeling the complex conditional distribution of future frames given history. Recently,
038 diffusion-based VideoAR (Kondratyuk et al., 2023; Chen et al., 2024a; Song et al., 2025; Gu et al.,
039 2025) has emerged as a promising solution, as it can approximate complex conditionals via iterative
040 denoising of future frames from random noise, conditioned on the history frames.

041 Despite this progress, the conditioning signal—the representation of the history frames—remains
042 underexplored. In text-to-image/video (T2I/T2V) and class-conditioned generation, stronger con-
043 dition representations consistently improve generation quality (Esser et al., 2024; Gao et al., 2024;
044 Kong et al., 2024; Hu et al., 2024b; Wu et al., 2025), which raises a natural question: *Can better*
045 *internal representations of history enhance VideoAR performance?*¹ Intuitively, if the model’s
046 internal representations of history effectively capture the semantics and dynamics of the history
047 frames, predicting coherent future frames should become easier. However, in current diffusion-
048 based VideoAR, history representations are mainly learned via the diffusion objective, which may
049 not be optimal for learning semantically aligned, predictive condition representations. Moreover,
050 good representations do not naturally emerge from VideoAR training, because predicting future
051 frames requires modeling the low-level details of the future, which can hinder representation learning
052 (Yu et al., 2024). This limitation motivates us to design dedicated learning objectives to enhance

053 ¹We focus on internal representations of *clean* history frames as conditions, distinct from methods that
improve representations of *noisy* data within the diffusion process (Yu et al., 2024; Zhang et al., 2025).

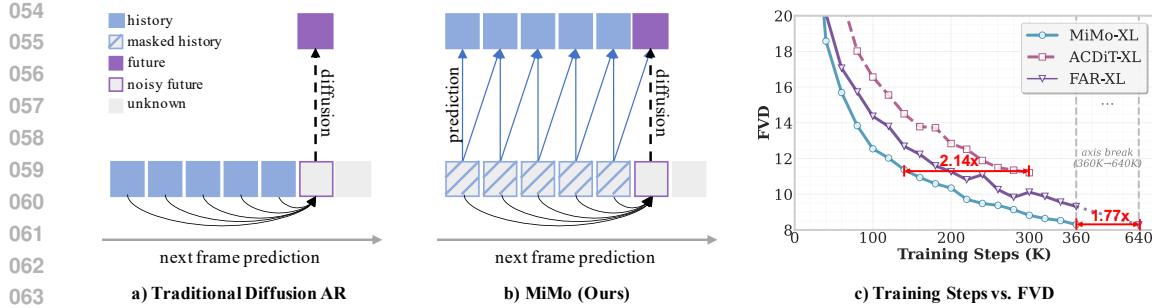


Figure 1: Good representations of history frames improve VideoAR. Our framework, MiMo, incorporates masked modeling into the history frames during training. MiMo achieves significantly faster convergence than baselines *without* using VFM.

history representations and improve VideoAR performance. Importantly, we aim to achieve this without relying on external vision foundation models (VFM) that incur substantial training costs and may suffer from out-of-distribution issues when applied to new video domains.

In this work, we demonstrate that improving history frame representations can indeed enhance VideoAR performance—an improvement that cannot be achieved by refining noisy future frame representations alone. Based on this insight, we propose Masked History Modeling (MiMo), a novel diffusion-based VideoAR framework *without* vision foundation model (VFM), illustrated in Figure 1. MiMo naturally integrates masked modeling (Devlin et al., 2019; He et al., 2022a; Tong et al., 2022), a simple yet effective representation learning technique, into history frame modeling for VideoAR. Our approach works as follows: We first address the train-test discrepancy in recent methods (Chen et al., 2024a; Song et al., 2025) by incorporating clean (noise-free) history frames alongside the noisy future frames being denoised during training, similar to the approaches of Zhou et al. (2025); Hu et al. (2024a). Then, we mask (drop) portions of the history frame tokens and train the model to reconstruct the masked tokens of current and future frames in parallel with the diffusion loss. This dual objective encourages the model to learn robust history representations that help the model to predict future frames, while also improving its robustness to perturbations in history.

Unlike previous work that applies masked modeling to noisy inputs of diffusion models (Gao et al., 2023; Wei et al., 2023), which harms diffusion and requires complicated techniques to mitigate the negative effects, our approach operates on clean history frames. Our approach greatly alleviates interference with future prediction and requires minimal architectural modifications. MiMo substantially improves training efficiency and generation quality through self-supervised visual representation learning and achieves strong VideoAR performance, all without external pretrained VFM (Yu et al., 2024; Zhang et al., 2025). In MiMo, history frames serve dual purposes: as conditions for the denoising of the future frame, and as input for self-supervised representation learning. By unifying history representation learning with future frame diffusion modeling, our framework enables high-quality representations that boost video prediction and generation.

Our main contributions are:

1. We investigate how history frame representations impact VideoAR performance and demonstrate that better representations lead to improved generation quality.
2. We propose MiMo, a simple yet effective VFM-free framework that seamlessly unifies diffusion-based VideoAR with self-supervised history representation learning.
3. Our framework demonstrates competitive video prediction and generation performance in VideoAR, achieving state-of-the-art (SOTA) results on several benchmarks.

2 PRELIMINARIES

VideoAR Given a video $\mathbf{x} = \{x_i \in \mathbb{R}^{H \times W \times 3} | i = 1, \dots, T\}$ with T frames of height H and width W , AR approaches model the temporal sequence by generating future frames *sequentially*

108 conditioned on historical frames, following the natural causal structure of video dynamics. VideoAR
 109 can be formulated by conditional probabilities:
 110

$$p(x_{t+1}|x_{1:t}) = p(\text{future frame}|\text{history frames}), \quad (1)$$

112 where $x_{1:t} = \{x_1, x_2, \dots, x_t\}$ represents the history frames and x_{t+1} is the next frame to generate.
 113

114 **Diffusion-based VideoAR** The conditional probabilities defined by Equation (1) are usually
 115 complex, which can be modeled by diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020).
 116 Diffusion-based AR approaches models Equation (1) by learning to denoise the Gaussian-noised
 117 future frame: $x_{t+1}^{(\tau)} = \alpha_\tau x_{t+1} + \sigma_\tau \epsilon$, conditioned on history frames $x_{1:t}$, where $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 118 and $\tau \in [0, 1]$ is noise level, and $\{\alpha_\tau, \sigma_\tau\}_\tau$ is noise schedule. This is typically done by es-
 119 timating the score function $s_\theta(x_{t+1}^{(\tau)}; \tau, x_{1:t}) \approx \nabla \log p_\tau(x_{t+1}^{(\tau)}|x_{1:t})$ (Vincent, 2011). In prac-
 120 tice, s_θ is often parameterized in alternative forms, such as v-prediction (Salimans & Ho, 2022):
 121 $v_\theta(x_{t+1}^{(\tau)}; \tau, x_{1:t}) \approx \alpha_\tau \epsilon - \sigma_\tau x_{t+1}$.
 122

123 During training, diffusion forcing (Song et al., 2025; Gu et al., 2025) learns v_θ that conditions on
 124 noisy history frames $x_{1:t}^{(\tau_{1:t})} = \{x_1^{(\tau_1)}, x_2^{(\tau_2)}, \dots, x_t^{(\tau_2)}\}$ with independent noise levels $\tau_{1:t}$,
 125

$$\mathcal{L} = \mathbb{E}_{t, \tau_{1:t+1}, \epsilon_{t+1}, \mathbf{x}} \left[\|\alpha_{\tau_{t+1}} \epsilon_{t+1} - \sigma_{\tau_{t+1}} x_{t+1} - v_\theta(x_{t+1}^{(\tau_{t+1})}; \tau_{t+1}, x_{1:t}^{(\tau_{1:t})})\|_2^2 \right]. \quad (2)$$

127 In contrast, complete teacher forcing (CTF) (Hu et al., 2024a; Zhou et al., 2025) eliminates train-test
 128 discrepancy of diffusion forcing by conditioning on clean history frames $x_{1:t}$:
 129

$$\mathcal{L} = \mathbb{E}_{t, \tau_{t+1}, \epsilon_{t+1}, \mathbf{x}} \left[\|\alpha_{\tau_{t+1}} \epsilon_{t+1} - \sigma_{\tau_{t+1}} x_{t+1} - v_\theta(x_{t+1}^{(\tau_{t+1})}; \tau_{t+1}, x_{1:t})\|_2^2 \right] \quad (3)$$

132 During generation, the model iteratively denoises $x_{t+1}^{(\tau)}$ using the learned denoising network, starting
 133 from pure noise and gradually recovering the clean future frame. Once x_{t+1} is fully denoised, it is
 134 appended to history frames for generating the subsequent frame x_{t+2} .
 135

136 3 METHOD

138 3.1 OVERVIEW

140 We hypothesize that good history frame representations benefit VideoAR. To investigate this hypoth-
 141 esis, we first analyze the relationship between history frame representation quality and VideoAR
 142 performance (Section 3.2). Our findings reveal that improving history representations improves
 143 performance, and such improvement cannot be achieved by solely refining noisy future frames.
 144

145 These findings motivate our approach. Additionally, we aim to avoid using VFM (Yu et al., 2024),
 146 as they may perform poorly for out-of-distribution (OOD) data, and adapting or pretraining VFMs
 147 on OOD data requires additional effort and increases complexity. Specifically, we propose **Masked**
 148 **History Modeling** (MiMo), a unified framework that jointly optimizes history frame representation
 149 learning and VideoAR within a single training process (Section 3.3) *without* using VFM. The insight
 150 of MiMo lies in treating history frames as noise-free conditioning signals during both training and
 151 inference, while introducing auxiliary masked video modeling objectives specifically targeting his-
 152 tory frames. The dual objectives ensure that the model develops robust history frame representations
 153 while maintaining strong generative capabilities. MiMo can also be extended to other pretraining
 154 objectives (Oquab et al., 2023; Assran et al., 2023; Jiang et al., 2025; Wang & He, 2025), which we
 155 leave for future work.
 156

157 3.2 EXPLORING REPRESENTATIONS OF HISTORY

158 In this section, we analyze the impact of history frame representations on VideoAR performance.
 159 We aim to understand whether good representations of history frames correlate with better genera-
 160 tion quality, and whether this is necessary—in other words, whether we can achieve all benefits by
 161 solely improving the representations of the noisy future frames being denoised (Yu et al., 2024). To
 162 exclude the influence of first-frame generation quality and focus on understanding the role of his-
 163 tory frames, we conduct experiments on the K600 video prediction task, where the video prediction

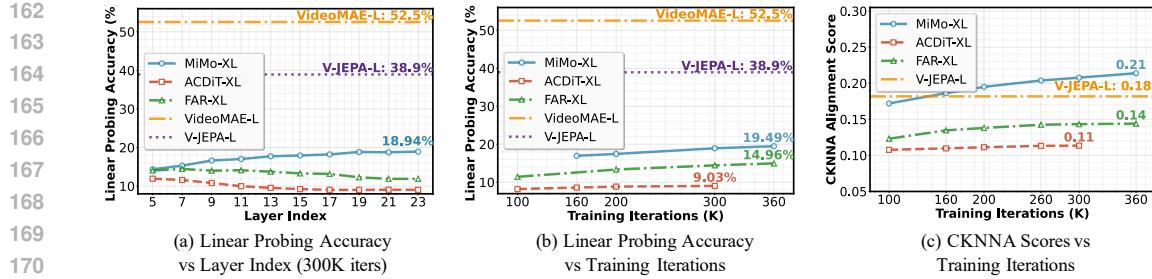


Figure 2: Exploring representations of history frames.

model predicts eleven future frames based on five given context frames. To study representation quality, we perform linear probing on K600 and measure CKNNA (Huh et al., 2024) to assess the similarity between model internal representations and pretrained representations (Yu et al., 2024). We select VideoMAE (Tong et al., 2022) and V-JEPA (Bardes et al., 2024) as VFM. All models use DFoT VAE (Song et al., 2025), with identical hyperparameters across all experiments and history guidance (see Appendix D.5) with scale 1.05 (Song et al., 2025) during inference. Details about evaluations are provided in Appendix F.

History frame representation quality correlates with video prediction performance We empirically investigate the relationship between history frame representation quality and video prediction performance using the models trained as shown in Figure 1, with results summarized in Figure 2. Our main findings include: (a) History frame quality positively correlates with video prediction performance—better models exhibit higher linear probing accuracy and better alignment with VFM (measured by CKNNA). (b) During training, history frame representation quality gradually improves but consistently maintains a significant gap with pretrained models. (c) Our proposed MiMo method effectively improves history frame representation quality. Notably, MiMo changes the layer where linear probing accuracy peaks, as our method introduces decoders in later layers to execute the masked history modeling objective (see Section 3.3).

Improving history frame representations is a feasible way to improve video prediction performance. We investigate whether improving history frame representation quality can enhance video prediction performance by training ACDiT-B models (Hu et al., 2024a), which take clean history frames and noisy future frames as input during training, where both history and future frames can only attend to themselves and history frames. This architecture allows us to explicitly separate the representations of history frames. We compare two approaches: one similar to REPA (Yu et al., 2024), which distills features from VFM into history frame representations; another introduces the MAE objective (He et al., 2022a) in history frames. Table 1 shows that both REPA and self-supervised methods can improve representation quality and subsequently enhance video prediction performance, demonstrating that improving history frame representations is feasible.

Improving noisy future frame representations cannot replace the role of improving history frame representations. Besides history frame representations, the representation quality of noisy future frames also affects diffusion model generation performance (Yu et al., 2024; Zhang et al., 2025). A meaningful question is: Is it sufficient to only improve the representation quality of noisy future frames? Our answer is *no*. We train ACDiT-B models and compare introducing REPA objectives in clean history frames, noisy future frames, or both. Table 1 demonstrates that merely improving noisy future frame representations is insufficient. Simultaneously improving both history and future frames yields benefits beyond just improving noisy future frame representations, indicat-

Table 1: Improving representations of history frames.

| Method | FVD \downarrow | Acc.(%) \uparrow |
|---------|------------------|--------------------|
| ACDiT-B | 54.814 | 6.21 |
| REPA | 40.022 | 16.96 |
| | 40.253 | 17.04 |
| | 36.542 | 19.23 |

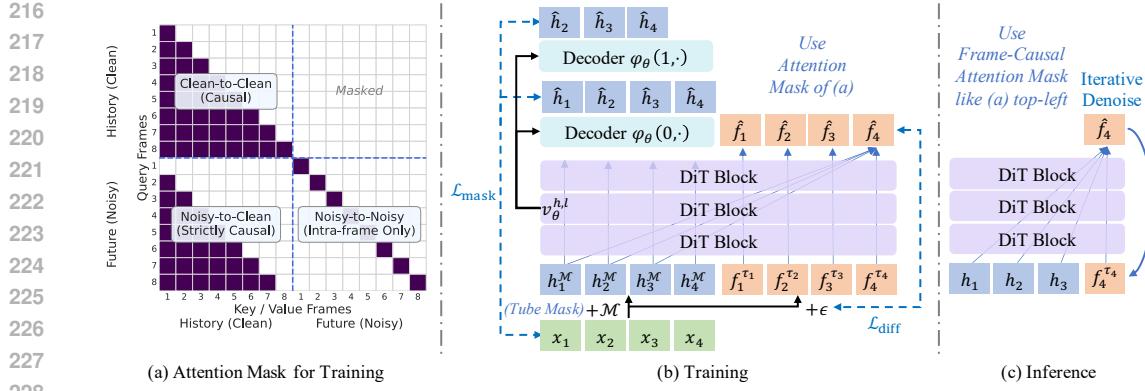


Figure 3: Framework of MiMo. (a) Attention mask used for training. Eight frames are shown. The clean history frames and noisy future frames are allowed to attend to themselves and previous history frames. (b) Training. Four frames are shown. The video clip $\mathbf{x} = \{x_1, \dots, x_3\}$ is both used as history frames \mathbf{h} and masked with a random tube mask \mathcal{M} , and as future frames \mathbf{f} and noised with Gaussian noise ϵ . The prediction targets of masked history modeling are the current and next frames. (c) AR Inference. Three history frames are already generated or provided by the user; the fourth frame is being denoised. After denoising, the fourth frame is appended to the history frames.

ing that history frames contain unique semantics. Note that our attempt to introduce MAE objective in noisy future frames (similar to Wei et al. (2023); Gao et al. (2023)) fails to surpass the performance of the ACDiT baseline without modifying the model’s macro-architecture (also reported by Gao et al. (2023)). We leave such exploration for future work.

3.3 MiMo: MASKED HISTORY MODELING

Motivated by our findings in Section 3.2, we propose MiMo to improve history representations in diffusion-based VideoAR.

Framework Design The core design principle of MiMo is to leverage history frames for dual purposes: (1) as conditions for diffusion-based future frame generation, and (2) as input for self-supervised representation learning through masked modeling. This dual utilization enables the model to develop robust history frame representations that are specifically tailored for video modeling tasks. The design is shown in Figure 3.

During training, MiMo follows CTF (Hu et al., 2024a; Zhou et al., 2025), which exposes clean history frames for representation learning. Given a video clip $\mathbf{x} = \{x_1, x_2, \dots, x_T\}$, we duplicate it as both history frames $\mathbf{h} = \{h_1, h_2, \dots, h_T\}$ and target future frames $\mathbf{f} = \{f_1, f_2, \dots, f_T\}$, where $\mathbf{h} = \mathbf{f} = \mathbf{x}$. The history frame h_t is input without noise; it can attend to itself and all its previous history frames $h_{t' \leq t}$. Future frame f_t is independently noised with Gaussian noise ϵ_t as in DFoT (Chen et al., 2024a; Song et al., 2025); it can attend to itself and all the previous future frames $h_{t' < t}$. This can be implemented by an attention mask as depicted in Figure 3(a).

The diffusion objective for future frame generation is:

$$\mathcal{L}_{diff} = \mathbb{E}_{t, \tau, \epsilon_t, \mathbf{x}, \mathcal{M}} \left[\|\alpha_\tau \epsilon_t - \sigma_\tau f_t - v_\theta(f_t^{(\tau)}; \tau, h_{1:t}^M)\|_2^2 \right], \quad (4)$$

where $x_t^{(\tau)}$ is the noisy version of the future frame x_t at diffusion timestep τ , v_θ is the denoising network (v-prediction (Salimans & Ho, 2022)) conditioned on masked history frames $h_{1:t}^M$, and \mathcal{M} is a random tube mask (Tong et al., 2022) applied on the history frames with a ratio r for masked history modeling (introduced below).

Masked History Modeling To enhance the model’s understanding of history frames, we introduce a masked modeling objective on the clean history frames. We randomly mask a subset of tokens in the history frame h_t and train the model to reconstruct the masked content. *Crucially*, the recon-

270 struction target can be either the tokens of the current history frame or any of the clean future frames
 271 $h_{t' \geq t}$. This distinguishes it from the normal diffusion objective as it allows greater flexibility.
 272

273 Formally, we introduce the reconstruction loss on the masked history frames $h_{1:t}^M$, the reconstruction
 274 target is a set of frames $\mathcal{T}_t = \{t, t+1\}$ which contain both t and its next frame $t+1$:

$$275 \quad \mathcal{L}_{\text{mask}} = \mathbb{E}_{t, \tau, \epsilon_t, \mathbf{x}, \mathcal{M}} \left[\frac{1}{|\mathcal{T}_t|} \sum_{t' \in \mathcal{T}_t} \|h_{t'} - \varphi_\theta(t' - t, v_\theta^{h,l}(f_t^{(\tau)}; \tau, h_{1:t}^M))\|_2^2 \right], \quad (5)$$

278 where φ_θ is a lightweight decoder that predicts the masked tokens of frame $t' \in \mathcal{T}$, and $v_\theta^{h,l}$ is the
 279 denoising network's output features of the l -th layer for the history frames $h_{1:t}$.
 280

281 The unified training objective combines both losses:

$$282 \quad \mathcal{L}_{\text{total}} = \mathcal{L}_{\text{diff}} + \lambda \mathcal{L}_{\text{mask}}, \quad (6)$$

284 where the hyperparameter λ balances the masked modeling objective.
 285

286 **Inference** During inference, as shown in Figure 3, MiMo discards decoder φ_θ and operates in
 287 standard AR fashion with KV cache (Zhou et al., 2025; Hu et al., 2024a; Gu et al., 2025): given
 288 clean history frames $h_{1:t-1}$, the model generates the next future frame f_t through iterative denoising.
 289 The learned history representations enhance the model's ability to maintain temporal coherence and
 290 generate high-quality future content. The framework naturally supports variable-length generation
 291 by iteratively updating the history context with newly generated frames.
 292

293 **Discussion** Compared with masked diffusion that applies a masked modeling objective to denoising
 294 input (Gao et al., 2023; Wei et al., 2023), our approach operates on clean history frames and
 295 mitigates the interference with the diffusion denoising process. Thus, MiMo requires no special
 296 architectural designs that masked diffusion approaches require. Zhou et al. (2025) also corrupt history
 297 frames, but their motivation is to improve robustness to noise in history, and they apply no recon-
 298 struction target to the history frames. Thus, they are still limited in history representations. Our
 299 approach also reduces the computational costs compared with those of Zhou et al. (2025); Hu et al.
 300 (2024a) due to masking.
 301

4 EXPERIMENTS

4.1 SETUP

305 **Tasks and Datasets** We evaluate MiMo on three video modeling tasks: video prediction, un-
 306 conditional video generation, and class-conditional video generation. For video prediction, we use the
 307 Kinetics-600 dataset (Carreira et al., 2018), which consists of 480,000 videos with 600 categories
 308 (class labels are *not* used). Five frames are provided as initial conditions to predict the next eleven
 309 frames. For video generation, we use the UCF-101 dataset (Soomro et al., 2012) with 13,320 videos
 310 across 101 categories. No initial frame is provided, and the model generates 16 frames.
 311

312 **Implementation Details** The architecture is based on DiT (Peebles & Xie, 2023). Our mod-
 313 ifications are: 1) using QK normalization (Henry et al., 2020) to stabilize training, 2) incorporating
 314 RoPE (Su et al., 2024), and 3) using separate LayerNorm (Ba et al., 2016) for clean history frames².
 315 The decoder is a stack of four DiT blocks with the same configuration as the model. Hyperparameter
 316 $\lambda = 0.5$. The learning rate is 8×10^{-4} for Kinetics and 4×10^{-4} for UCF-101, both decayed to 10^{-5}
 317 with cosine schedule. The global batch size is 256 for Kinetics and 128 for UCF-101. The weight
 318 decay is 0.001, and the betas for AdamW (Loshchilov & Hutter, 2017) are (0.9, 0.99). The model
 319 is trained for 360K steps on Kinetics and 180K on UCF-101 with 32 H100 GPUs. For Kinetics, we
 320 use DFoT's VAE (Song et al., 2025) with a compression ratio of $4 \times 8 \times 8$ and sample 17 frames
 321 per clip with resolution 128. For UCF-101, we use FAR's (Gu et al., 2025) per-frame DC-AE (Chen
 322 et al., 2024b) with a compression ratio of 32×32 and sample 16 frames per clip with resolution
 323 256. See Appendix D for more details.

²These modifications moderately affect performance, as shown in Section 4.3.

324
 325 Table 2: System comparison on Kinetics and UCF-101 with video prediction, unconditional video
 326 generation, and conditional video generation tasks. \dagger : Different from the original work, we reimplemented
 327 DFoT using a causal architecture to align with the standard AR practice.

| 328 Method | 329 Type | 330 Kinetics (Pred.) FVD \downarrow | 331 UCF-101 (Uncond.) FVD \downarrow | 332 UCF-101 (Cond.) FVD \downarrow |
|---|-----------------|---|--|--|
| 333 LVDM (He et al., 2022b) | Non-AR | — | 372 | — |
| 334 MAGVIT (Yu et al., 2023a) | | 9.9 | — | 76 |
| 335 MAGVITv2 (Yu et al., 2023b) | | 4.3 | — | 58 |
| 336 Latte (Ma et al., 2024b) | | — | 478 | — |
| 337 TATS (Ge et al., 2022) | AR | — | 420 | 332 |
| 338 Phenaki (Villegas et al., 2022) | | 36.4 | — | — |
| 339 Omni (Wang et al., 2024) | | 32.9 | — | 191 |
| 340 DFoT-XL \dagger (Song et al., 2025) | | 11.1 | — | — |
| 341 ACDiT-XL (Hu et al., 2024a) | | — | — | 111 |
| 342 MAGI-XL (Zhou et al., 2025) | | 11.5 | 298 | — |
| 343 FAR-XL (Gu et al., 2025) | | — | 279 | 108 |
| 344 VAE Reconstruction | AR | 3.7 | 15 | 15 |
| 345 MiMo-XL | | 8.3 | 240 | 98 |

346 **Inference and Evaluation** We follow Song et al. (2025) for evaluation on Kinetics, generating
 347 50,000 random videos and computing the Fréchet Video Distance (FVD) (Karras et al., 2019) on all
 348 frames (including conditioning and generated frames) with the groundtruth videos, both resized to
 349 64×64 . On UCF-101, following FAR, we randomly sample 2,048 videos and compute the FVD
 350 against groundtruth videos, resized to 256×256 .
 351

352 4.2 MAIN RESULTS

353 Table 2 presents a comprehensive comparison of MiMo against state-of-the-art non-AR and AR
 354 methods across three video modeling tasks. For reference, we also report the reconstruction FVD of
 355 the VAE, which represents the upper bound of performance achievable given the groundtruth.
 356

357 **Video Prediction (Pred.)** On the challenging Kinetics-600 video prediction benchmark, MiMo
 358 demonstrates exceptional performance with an FVD score of 8.3, establishing a new state-of-the-art
 359 among AR models. This represents a substantial improvement over previous AR methods, with
 360 our approach significantly outperforming DFoT (FVD: 11.1) despite using the same VAE. The
 361 performance gain directly demonstrates the superiority of MiMo, as both methods share the same
 362 underlying video tokenization and differ primarily in their treatment of historical context. Qualitative
 363 examples are presented in Figure 4(a), where our method generates smooth, realistic continuations.
 364

365 **Unconditional Video Generation (Uncond.)** For unconditional video generation on UCF-101,
 366 MiMo achieves remarkable results, establishing new state-of-the-art performance among AR
 367 approaches. Our method substantially outperforms the previous AR leader FAR by nearly 40 FVD
 368 points (240 vs 279) while utilizing the same DC-AE tokenizer, demonstrating the significant impact
 369 of our masked history modeling approach. Also noteworthy is the comparison with MAGI, which
 370 similarly employs Complete Teacher Forcing (CTF) during training—our method achieves a
 371 considerable performance improvement (FVD: 240 vs 298), validating the effectiveness of our masked
 372 history modeling objective. The generated videos exhibit diverse motions, realistic textures, and
 373 coherent temporal dynamics, as illustrated in the qualitative examples in Figure 4(b).
 374

375 **Class-Conditional Video Generation (Cond.)** In class-conditional video generation on UCF-
 376 101, MiMo again demonstrates superior performance, achieving state-of-the-art results across AR
 377 methods. Our approach surpasses FAR by 10 FVD points (98 vs 108), confirming the consistent
 378 benefits of our approach across different conditioning modalities. The comparison with ACDiT is
 379 also interesting—both methods utilize CTF and share similar architectural foundations, yet MiMo
 380 achieves notable improvements (FVD: 98 vs 111), consistent with our findings in unconditional

378 generation when compared against MAGI. This consistency across tasks reinforces that our performance gains stem from improvements in history representation learning rather than task-specific
 379 optimizations. Representative generated videos are shown in Figure 4(c).
 380
 381



Figure 4: Visualization of generated videos.

4.3 ABLATION STUDY

403 This section ablates the designs of MiMo. All models are based on DiT-B trained on Kinetics for
 404 100K steps with modifications and hyperparameters described in Section 4.1. ACDiT baseline is
 405 MiMo without masked history modeling, similar to MAGI and ACDiT.

Table 3: Comparison with variants of REPA.

| Method | FVD \downarrow |
|-----------------|------------------|
| ACDiT Baseline | 54.814 |
| History | 40.022 |
| REPA Future | 40.253 |
| Both | 36.542 |
| MiMo | 36.601 |
| MiMo +REPA-Both | 34.133 |

Table 4: Comparison of different prediction targets for masked history modeling.

| Target(s) | FVD \downarrow |
|-------------------------|------------------|
| ACDiT Baseline | 54.814 |
| Current Frame | 41.832 |
| Next Frame | 37.782 |
| Current + Next (MiMo) | 36.601 |
| Current + Next+NextNext | 36.263 |

417 **Comparison With REPA** An alternative way to inject good representations into the model is
 418 distilling the features from a VFM, known as REPA (Yu et al., 2024). Table 3 compares MiMo with
 419 several variants of REPA, supervising history frames, future frames, or both. Both REPA and MiMo
 420 can significantly improve convergence, while MiMo performs on par or better than all variants. In
 421 practice, however, VFMs for the user’s domain of interest are not always available, in which cases
 422 MiMo is a viable substitute.

423 MiMo is complementary to VFM-based methods like REPA. MiMo excels at learning task-specific
 424 dynamics from the data, while VFM provides strong semantic priors. Combining MiMo with REPA
 425 in Table 3 yielded further improvements over either method alone. This suggests that MiMo and
 426 VFM capture different aspects of the data.

427 **Prediction Targets of Masked Modeling** One of the merits of MiMo is its flexibility: while
 428 diffusion always predicts the noise-free version of the noised current frame effectively, MiMo can
 429 predict both the current and next frames for masked history modeling. Table 4 compares different
 430 prediction targets for masked modeling, and it is clear that predicting both current and history frames
 431 outperforms predicting either.

We also include a target of predicting the current and next two frames (Current + Next+NextNext) in Table 4. It is beneficial but yields diminishing returns. We hypothesize that predicting a more distant future frame is a significantly harder task, and the increased difficulty does not naively translate into proportional performance gains. Our proposed target (Current + Next) strikes an effective balance without the added complexity of longer-range prediction.

Table 5: Decoder position (placed after the l -th layer) moderately affects performance. DiT-B has 12 layers in total. None means CFT baseline.

| l | None | 12 | 11 | 10 | 9 |
|------------------|--------|--------|--------|--------|--------|
| FVD \downarrow | 54.814 | 36.601 | 35.815 | 35.838 | 37.593 |

Table 6: Architecture modifications.

| Modification | FVD \downarrow |
|---------------------|------------------|
| Vanilla DiT | 37.763 |
| +RoPE | 37.313 |
| +Separate LayerNorm | 36.601 |

Decoder Position The decoder for masked modeling is usually placed after the encoder (He et al., 2022a; Tong et al., 2022). We treat the first l layers of the DiT model as the encoder, and shortcut the output of the l -the layer corresponding to the history frames into the decoder. Table 5 shows the effect of varying l . The performance is robust to l when l is close to the last layer.

Table 6 shows the impact of architecture modifications on performance. Incorporating RoPE and separating LayerNorm layers for history frames both bring moderate gains.

Table 7: Ablations of hyperparameters λ and mask ratios.

| λ | 0.1 | 0.5 | 1.0 | 2.0 | Mask Ratios | [0.25, 0.25] | [0.25, 0.5] | [0.5, 0.75] |
|------------------|--------|--------|--------|--------|------------------|--------------|-------------|-------------|
| FVD \downarrow | 40.213 | 36.601 | 37.443 | 38.910 | FVD \downarrow | 37.539 | 36.601 | 39.121 |

Hyperparamter Ablations Table 7 shows the impact of hyperparameters on MiMo-B models with different λ from 0.25 to 2.0 and mask ratios from 0.25 to 0.75. For λ , a weight of 0.5 provides the best balance, but performance does not degrade sharply for nearby values. For mask ratios, performance remains relatively stable between 0.25 and 0.50; however, a higher mask ratio requires fine-tuning without masking to achieve better performance.

Table 8: Computational costs.

| Method | MiMo-XL | ACDiT-XL | FAR-XL |
|-----------------|---------|----------|--------|
| Wall-Clock Time | 0.750s | 0.788s | 0.704s |
| GFLOPs | 8.22 | 8.81 | 5.94 |

Table 9: Comparison of optimization strategies.

| Strategy | FVD \downarrow |
|----------------|------------------|
| ACDiT Baseline | 54.814 |
| Interleaving | 38.543 |
| MiMo | 36.601 |

Computational Costs Table 8 shows the training computational costs of MiMo, ACDiT (our baseline), and FAR. Compared with our baseline (ACDiT), MiMo reduces training wall-clock time by 5%; compared with FAR, *MiMo increases training wall-clock time by a modest 10%, which is a small price for the significant performance boost* (25% from 279 to 240 on Kinetics, and 14% from 279 to 240 on UCF-101). Compared with ACDiT, MiMo also reduces the FLOPs per training step due to masking (the decoders increase FLOPs). MiMo has higher training FLOPs compared with FAR, but the increase in training wall-clock time is moderate (10%)

486 due to hardware acceleration. Note that MiMo has *no* additional inference cost once the training is
 487 complete.
 488

489
 490 **Alternative Optimization Strategy** While our approach of simply using a weighted sum of the
 491 two losses is standard practice for auxiliary loss training, an alternative approach is optimizing the
 492 diffusion loss and the auxiliary loss interleavingly (a diffusion-only training step is followed by a
 493 mask-only training step). The results are summarized in Table 9. While the interleaving approach
 494 has lower computational costs per step, it leads to slower convergence, which diminishes its speed
 495 gains.
 496

497 5 RELATED WORKS

498
 499
 500 **Autoregressive Visual Generation.** Autoregressive language modeling (Radford, 2018; Radford
 501 et al., 2019) has facilitated the development of visual content generation using discrete visual tokens
 502 (Van Den Oord et al., 2017). In this framework, pre-trained visual tokenizers like VQ-VAE (Van
 503 Den Oord et al., 2017) map visual patches into a discrete latent space, allowing visual generation to
 504 be approached similarly to language modeling. Early works such as DALL-E (Ramesh et al., 2021)
 505 focus on text-to-image generation by learning a joint distribution between text and discrete image
 506 representations using an autoregressive cross-entropy loss. VideoGPT (Yan et al., 2021) extends this
 507 idea to video generation, employing discrete tokens for autoregressive video prediction. VideoPoet
 508 (Kondratyuk et al., 2023) further advances this approach by integrating a causal video tokenizer (Yu
 509 et al., 2023b). OmniTokenizer (Wang et al., 2024) proposes a unified tokenizer for both discrete
 510 and continuous representations. In contrast, our work focuses on frame-level causality rather than
 patch-level, avoiding the limitations of raster-scan order.

511
 512 **Representations and Generative Modeling.** Recent advances in diffusion models highlight the
 513 importance of high-quality representations for generative modeling, as diffusion models inherently
 514 struggle to learn good representations (Yu et al., 2024; Zhang et al., 2025; Jiang et al., 2025; Wang &
 515 He, 2025). In practice, diffusion models are predominantly conditional generative models, where the
 516 conditions can be text prompts in T2I/T2v T2V generation, or history frames in VideoAR. Despite
 517 this prevalence, few studies have investigated how the quality of condition representations affects
 518 generative performance. Existing evidence from text-conditional generation provides compelling
 519 support for exploring this relationship. Replacing CLIP text encoders with large language
 520 models such as T5 and Llama has consistently improved generation quality, particularly for attributes
 521 strongly correlated with text conditions (counting, object reference, text rendering, etc.) (Esser
 522 et al., 2024; Gao et al., 2024; Kong et al., 2024; Hu et al., 2024b). Another evidence is that dis-
 523 tilling class representations improves the performance of class-conditioned image generation (Wu
 524 et al., 2025). These observations naturally extend to VideoAR, where future frames depend on his-
 525 tory frames as conditions, suggesting that enhanced history frame representations may potentially
 526 benefit video generation performance, which is the focus of our work.
 527

528 6 CONCLUTION

529
 530 In this work, we explored the fundamental question of whether good representations of history
 531 frames can improve VideoAR performance. Through systematic analysis, we demonstrated that
 532 enhancing history frame representations significantly benefits VideoAR, a finding that cannot be
 533 achieved by solely refining noisy future frames. Motivated by these insights, we proposed MiMo
 534 (Masked History Modeling), a novel framework that naturally integrates masked modeling into
 535 diffusion-based VideoAR. By applying masks to history frame tokens and training the model to
 536 predict masked tokens of current and future frames alongside denoising tasks, MiMo learns robust
 537 representations that improve VideoAR performance. Our approach requires no VFM or special
 538 architectural modifications. Extensive experiments across multiple benchmarks demonstrate that
 539 MiMo achieves competitive performance in video prediction and generation tasks, establishing new
 state-of-the-art results. Notably, our framework substantially improves training efficiency and gener-
 ation quality, showcasing the effectiveness of unified representation learning and diffusion modeling.

540 REFERENCES

542 Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic inter-
543 polants. *arXiv preprint arXiv:2209.15571*, 2022.

544 Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat,
545 Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding
546 predictive architecture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and
547 Pattern Recognition*, pp. 15619–15629, 2023.

548 Mido Assran, Adrien Bardes, David Fan, Quentin Garrido, Russell Howes, Matthew Muckley, Am-
549 mar Rizvi, Claire Roberts, Koustuv Sinha, Artem Zholus, et al. V-jepa 2: Self-supervised video
550 models enable understanding, prediction and planning. *arXiv preprint arXiv:2506.09985*, 2025.

552 Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint
553 arXiv:1607.06450*, 2016.

555 Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mahmoud
556 Assran, and Nicolas Ballas. Revisiting feature prediction for learning visual representations from
557 video. *arXiv preprint arXiv:2404.08471*, 2024.

558 Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe
559 Taylor, Troy Luhman, Eric Luhman, et al. Video generation models as world simulators. *OpenAI
560 Blog*, 1(8):1, 2024.

562 Jake Bruce, Michael D Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes,
563 Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, et al. Genie: Generative inter-
564 active environments. In *Forty-first International Conference on Machine Learning*, 2024.

565 Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics
566 dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp.
567 6299–6308, 2017.

569 Joao Carreira, Eric Noland, Andras Banki-Horvath, Chloe Hillier, and Andrew Zisserman. A short
570 note about kinetics-600. *arXiv preprint arXiv:1808.01340*, 2018.

572 Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative
573 image transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern
574 Recognition*, pp. 11315–11325, 2022.

575 Boyuan Chen, Diego Martí Monsó, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitz-
576 mann. Diffusion forcing: Next-token prediction meets full-sequence diffusion. *Advances in
577 Neural Information Processing Systems*, 37:24081–24125, 2024a.

578 Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Yue Wu, Zhongdao Wang, James
579 Kwok, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart- α : Fast training of diffusion transformer
580 for photorealistic text-to-image synthesis, 2023. URL <https://arxiv.org/abs/2310.00426>.

583 Junyu Chen, Han Cai, Junsong Chen, Enze Xie, Shang Yang, Haotian Tang, Muyang Li, Yao Lu, and
584 Song Han. Deep compression autoencoder for efficient high-resolution diffusion models. *arXiv
585 preprint arXiv:2410.10733*, 2024b.

586 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep
587 bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of
588 the North American chapter of the association for computational linguistics: human language
589 technologies, volume 1 (long and short papers)*, pp. 4171–4186, 2019.

591 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
592 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An
593 image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint
594 arXiv:2010.11929*, 2020.

594 DC Dowson and BV666017 Landau. The fréchet distance between multivariate normal distributions.
 595 *Journal of multivariate analysis*, 12(3):450–455, 1982.
 596

597 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam
 598 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers
 599 for high-resolution image synthesis. In *Forty-first international conference on machine learning*,
 600 2024.

601 Peng Gao, Le Zhuo, Dongyang Liu, Ruoyi Du, Xu Luo, Longtian Qiu, Yuhang Zhang, Chen Lin,
 602 Rongjie Huang, Shijie Geng, et al. Lumina-t2x: Transforming text into any modality, resolution,
 603 and duration via flow-based large diffusion transformers. *arXiv preprint arXiv:2405.05945*, 2024.

604 Shanghua Gao, Pan Zhou, Ming-Ming Cheng, and Shuicheng Yan. Masked diffusion transformer
 605 is a strong image synthesizer. In *Proceedings of the IEEE/CVF international conference on com-*
 606 *puter vision*, pp. 23164–23173, 2023.

607 Songwei Ge, Thomas Hayes, Harry Yang, Xi Yin, Guan Pang, David Jacobs, Jia-Bin Huang, and
 608 Devi Parikh. Long video generation with time-agnostic vqgan and time-sensitive transformer. In
 609 *European Conference on Computer Vision*, pp. 102–118. Springer, 2022.

610 Yuchao Gu, Weijia Mao, and Mike Zheng Shou. Long-context autoregressive video modeling with
 611 next-frame prediction. *arXiv preprint arXiv:2503.19325*, 2025.

612 Xun Guo, Mingwu Zheng, Liang Hou, Yuan Gao, Yufan Deng, Pengfei Wan, Di Zhang, Yufan Liu,
 613 Weiming Hu, Zhengjun Zha, et al. I2v-adapter: A general image-to-video adapter for diffusion
 614 models. In *ACM SIGGRAPH 2024 Conference Papers*, pp. 1–12, 2024.

615 Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh
 616 Agrawala, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffu-
 617 sion models without specific tuning. *arXiv preprint arXiv:2307.04725*, 2023.

618 Tiankai Hang, Shuyang Gu, Chen Li, Jianmin Bao, Dong Chen, Han Hu, Xin Geng, and Baining
 619 Guo. Efficient diffusion training via min-snr weighting strategy. In *Proceedings of the IEEE/CVF*
 620 *international conference on computer vision*, pp. 7441–7451, 2023.

621 Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked au-
 622 toencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer*
 623 *vision and pattern recognition*, pp. 16000–16009, 2022a.

624 Yingqing He, Tianyu Yang, Yong Zhang, Ying Shan, and Qifeng Chen. Latent video diffusion
 625 models for high-fidelity long video generation. *arXiv preprint arXiv:2211.13221*, 2022b.

626 Alex Henry, Prudhvi Raj Dachapally, Shubham Pawar, and Yuxuan Chen. Query-key normalization
 627 for transformers. *arXiv preprint arXiv:2010.04245*, 2020.

628 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter.
 629 Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in*
 630 *neural information processing systems*, 30, 2017.

631 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint*
 632 *arXiv:2207.12598*, 2022.

633 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*
 634 *neural information processing systems*, 33:6840–6851, 2020.

635 Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J
 636 Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–
 637 8646, 2022.

638 Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pre-
 639 training for text-to-video generation via transformers. *arXiv preprint arXiv:2205.15868*, 2022.

640 Jinyi Hu, Shengding Hu, Yuxuan Song, Yufei Huang, Mingxuan Wang, Hao Zhou, Zhiyuan Liu,
 641 Wei-Ying Ma, and Maosong Sun. Acdit: Interpolating autoregressive conditional modeling and
 642 diffusion transformer. *arXiv preprint arXiv:2412.07720*, 2024a.

648 Xiwei Hu, Rui Wang, Yixiao Fang, Bin Fu, Pei Cheng, and Gang Yu. Ella: Equip diffusion models
 649 with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*, 2024b.
 650

651 Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. The platonic representation
 652 hypothesis. *arXiv preprint arXiv:2405.07987*, 2024.

653 Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by
 654 reducing internal covariate shift. In *International conference on machine learning*, pp. 448–456.
 655 pmlr, 2015.

656

657 Sadeep Jayasumana, Srikumar Ramalingam, Andreas Veit, Daniel Glasner, Ayan Chakrabarti, and
 658 Sanjiv Kumar. Rethinking fid: Towards a better evaluation metric for image generation. In *Pro-
 659 ceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9307–
 660 9315, 2024.

661 Dengyang Jiang, Mengmeng Wang, Liuzhuozheng Li, Lei Zhang, Haoyu Wang, Wei Wei, Guang
 662 Dai, Yanning Zhang, and Jingdong Wang. No other representation component is needed:
 663 Diffusion transformers can provide representation guidance by themselves. *arXiv preprint
 664 arXiv:2505.02831*, 2025.

665

666 Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative
 667 adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern
 668 recognition*, pp. 4401–4410, 2019.

669 Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijaya-
 670 narasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action
 671 video dataset. *arXiv preprint arXiv:1705.06950*, 2017.

672 Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Ad-
 673 vances in neural information processing systems*, 34:21696–21707, 2021.

674

675 Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Grant Schindler, Rachel
 676 Hornung, Vighnesh Birodkar, Jimmy Yan, Ming-Chang Chiu, et al. Videopoet: A large language
 677 model for zero-shot video generation. *arXiv preprint arXiv:2312.14125*, 2023.

678

679 Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li,
 680 Bo Wu, Jianwei Zhang, et al. Hunyuandvideo: A systematic framework for large video generative
 681 models. *arXiv preprint arXiv:2412.03603*, 2024.

682

683 Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural
 684 network representations revisited. In *International conference on machine learning*, pp. 3519–
 3529. PMLR, 2019.

685

686 Shanchuan Lin, Bingchen Liu, Jiashi Li, and Xiao Yang. Common diffusion noise schedules and
 687 sample steps are flawed. In *Proceedings of the IEEE/CVF winter conference on applications of
 688 computer vision*, pp. 5404–5411, 2024.

689

690 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
 for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

691

692 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and
 693 transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

694

695 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint
 696 arXiv:1711.05101*, 2017.

697

698 Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast
 699 ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in neural
 700 information processing systems*, 35:5775–5787, 2022.

701

702 Ge Ya Luo, Gian Mario Favero, Zhi Hao Luo, Alexia Jolicoeur-Martineau, and Christopher
 703 Pal. Beyond fvd: Enhanced evaluation metrics for video generation quality. *arXiv preprint
 704 arXiv:2410.05203*, 2024.

702 Nanye Ma, Mark Goldstein, Michael S Albergo, Nicholas M Boffi, Eric Vanden-Eijnden, and Sain-
 703 ing Xie. Sit: Exploring flow and diffusion-based generative models with scalable interpolant
 704 transformers. In *European Conference on Computer Vision*, pp. 23–40. Springer, 2024a.

705

706 Xin Ma, Yaohui Wang, Gengyun Jia, Xinyuan Chen, Ziwei Liu, Yuan-Fang Li, Cunjian Chen,
 707 and Yu Qiao. Latte: Latent diffusion transformer for video generation. *arXiv preprint*
 708 *arXiv:2401.03048*, 2024b.

709 Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and
 710 projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

711

712 Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models.
 713 In *International conference on machine learning*, pp. 8162–8171. PMLR, 2021.

714 Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov,
 715 Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning
 716 robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.

717

718 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of*
 719 *the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.

720 Alec Radford. Improving language understanding by generative pre-training. *OpenAI blog*, 2018.

721

722 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language
 723 models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

724

725 A. Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and
 726 Ilya Sutskever. Zero-shot text-to-image generation. *ArXiv*, abs/2102.12092, 2021.

727

728 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 729 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-
 730 ence on computer vision and pattern recognition*, pp. 10684–10695, 2022.

731

732 Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv*
 733 *preprint arXiv:2202.00512*, 2022.

734

735 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised
 736 learning using nonequilibrium thermodynamics. In *International conference on machine learn-
 737 ing*, pp. 2256–2265. pmlr, 2015.

738

739 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv*
 740 *preprint arXiv:2010.02502*, 2020a.

741

742 Kiwhan Song, Boyuan Chen, Max Simchowitz, Yilun Du, Russ Tedrake, and Vincent Sitzmann.
 743 History-guided video diffusion. *arXiv preprint arXiv:2502.06764*, 2025.

744

745 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 746 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint*
 747 *arXiv:2011.13456*, 2020b.

748

749 Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions
 750 classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.

751

752 Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: En-
 753 hanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.

754

755 Hansi Teng, Hongyu Jia, Lei Sun, Lingzhi Li, Maolin Li, Mingqiu Tang, Shuai Han, Tianning
 756 Zhang, WQ Zhang, Weifeng Luo, et al. Magi-1: Autoregressive video generation at scale. *arXiv*
 757 *preprint arXiv:2505.13211*, 2025.

758

Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-
 759 efficient learners for self-supervised video pre-training. *Advances in neural information process-
 760 ing systems*, 35:10078–10093, 2022.

756 Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski,
 757 and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges.
 758 *arXiv preprint arXiv:1812.01717*, 2018.

759

760 Dani Valevski, Yaniv Leviathan, Moab Arar, and Shlomi Fruchter. Diffusion models are real-time
 761 game engines. *arXiv preprint arXiv:2408.14837*, 2024.

762

763 Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in
 764 neural information processing systems*, 30, 2017.

765

766 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
 767 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural informa-
 768 tion processing systems*, 30, 2017.

769

770 Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2008.

771

772 Ruben Villegas, Mohammad Babaeizadeh, Pieter-Jan Kindermans, Hernan Moraldo, Han Zhang,
 773 Mohammad Taghi Saffar, Santiago Castro, Julius Kunze, and Dumitru Erhan. Phenaki: Variable
 774 length video generation from open domain textual description. *arXiv preprint arXiv:2210.02399*,
 775 2022.

776

777 Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural compu-
 778 tation*, 23(7):1661–1674, 2011.

779

780 Junke Wang, Yi Jiang, Zehuan Yuan, Binyue Peng, Zuxuan Wu, and Yu-Gang Jiang. Ommitokenizer:
 781 A joint image-video tokenizer for visual generation. *arXiv preprint arXiv:2406.09399*, 2024.

782

783 Runqian Wang and Kaiming He. Diffuse and disperse: Image generation with representation regu-
 784 larization. *arXiv preprint arXiv:2506.09027*, 2025.

785

786 Chen Wei, Karttikeya Mangalam, Po-Yao Huang, Yanghao Li, Haoqi Fan, Hu Xu, Huiyu Wang, Ci-
 787 hang Xie, Alan Yuille, and Christoph Feichtenhofer. Diffusion models as masked autoencoders. In
 788 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 16284–16294,
 789 2023.

790

791 Ge Wu, Shen Zhang, Ruijing Shi, Shanghua Gao, Zhenyuan Chen, Lei Wang, Zhaowei Chen,
 792 Hongcheng Gao, Yao Tang, Jian Yang, et al. Representation entanglement for generation: Train-
 793 ing diffusion transformers is much easier than you think. *arXiv preprint arXiv:2507.01467*, 2025.

794

795 Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. Videogpt: Video generation using
 796 vq-vae and transformers. *arXiv preprint arXiv:2104.10157*, 2021.

797

798 Wilson Yan, Danijar Hafner, Stephen James, and Pieter Abbeel. Temporally consistent transform-
 799 ers for video generation. In *International Conference on Machine Learning*, pp. 39062–39098.
 800 PMLR, 2023.

801

802 Tianwei Yin, Qiang Zhang, Richard Zhang, William T Freeman, Fredo Durand, Eli Shechtman, and
 803 Xun Huang. From slow bidirectional to fast autoregressive video diffusion models. In *Proceed-
 804 ings of the Computer Vision and Pattern Recognition Conference*, pp. 22963–22974, 2025.

805

806 Lijun Yu, Yong Cheng, Kihyuk Sohn, José Lezama, Han Zhang, Huiwen Chang, Alexander G
 807 Hauptmann, Ming-Hsuan Yang, Yuan Hao, Irfan Essa, et al. Magvit: Masked generative video
 808 transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recog-
 809 nition*, pp. 10459–10469, 2023a.

810

811 Lijun Yu, José Lezama, Nitesh B Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong
 812 Cheng, Vighnesh Birodkar, Agrim Gupta, Xiuye Gu, et al. Language model beats diffusion-
 813 tokenizer is key to visual generation. *arXiv preprint arXiv:2310.05737*, 2023b.

814

815 Sihyun Yu, Sangkyung Kwak, Huiwon Jang, Jongheon Jeong, Jonathan Huang, Jinwoo Shin, and
 816 Saining Xie. Representation alignment for generation: Training diffusion transformers is easier
 817 than you think. *arXiv preprint arXiv:2410.06940*, 2024.

810 Xiangdong Zhang, Jiaqi Liao, Shaofeng Zhang, Fanqing Meng, Xiangpeng Wan, Junchi Yan, and
811 Yu Cheng. Videorepa: Learning physics for video generation through relational alignment with
812 foundation models. *arXiv preprint arXiv:2505.23656*, 2025.

813
814 Deyu Zhou, Quan Sun, Yuang Peng, Kun Yan, Runpei Dong, Duomin Wang, Zheng Ge, Nan Duan,
815 Xiangyu Zhang, Lionel M Ni, et al. Taming teacher forcing for masked autoregressive video
816 generation. *arXiv preprint arXiv:2501.12389*, 2025.

817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

864 A DIFFUSION MODELING

866 In this section, we present a brief overview of diffusion-based generative models. These models
 867 learn to approximate target distributions through training denoising neural networks. There are two
 868 correlated approaches: “conventional” diffusion models based on score matching (Appendix A.1),
 869 and flow matching (Appendix A.2), introduced below.

871 A.1 SCORE MATCHING

873 Diffusion models based on score matching (Ho et al., 2020; Kingma et al., 2021; Song et al., 2020b)
 874 generate samples $x \sim p_0(\cdot)$ by learning to invert a noise corruption process (i.e., the diffusion
 875 process) that transforms the data distribution into standard Gaussian noise $\epsilon \sim \mathcal{N}(0, I)$. The forward
 876 diffusion process is defined as:

$$877 p_\tau(x^{(\tau)}|x) = \mathcal{N}(\alpha_\tau x, \sigma_\tau^2 I); \quad \tau \in [0, 1], \quad (7)$$

878 where the coefficients α_τ and σ_τ specify the “noise schedule” that interpolates between data and
 879 noise. Usually $\alpha_0 = 1, \sigma_0 = 0$ and $\alpha_1 = 0, \sigma_1 = 1$, so that $x^{(0)} = x$ and $x^{(1)} = \epsilon$.

$$880 x^{(0)} = x, \quad x^{(1)} = \epsilon \quad (8)$$

881 The generative process is realized by integrating the reverse-time stochastic differential equation
 882 (SDE) (Song et al., 2020b; Lu et al., 2022) that describes the backward diffusion process:

$$884 dx^{(\tau)} = \left[f(\tau)x^{(\tau)} - g^2(\tau)\nabla_{x^{(\tau)}} \log p_\tau(x^{(\tau)}) \right] d\tau + g(\tau) d\bar{w}_\tau, \quad (9)$$

885 where \bar{w}_τ denotes the reverse-time Wiener process, and the drift and diffusion coefficients f and g
 886 are given by:

$$887 f(\tau) = \frac{d \log \alpha_\tau}{d\tau}, \quad g^2(\tau) = -\sigma_\tau^2 \frac{d \log(\alpha_\tau / \sigma_\tau)}{d\tau}. \quad (10)$$

889 A score network $s_\theta(x^{(\tau)}; \tau)$ is trained to approximate $\nabla_{x^{(\tau)}} \log p_\tau(x^{(\tau)})$ via denoising score match-
 890 ing (Vincent, 2011):

$$892 \min_{\theta} \mathbb{E}_{\tau, \epsilon, x^{(0)}, x^{(\tau)}} [\|\epsilon + \sigma_\tau s_\theta(x^{(\tau)}; \tau)\|_2^2]. \quad (11)$$

893 Beyond directly modeling the score $s_\theta(x^{(\tau)}; \tau)$, diffusion models commonly use equivalent param-
 894 eterizations tied to the forward relation $x^{(\tau)} = \alpha_\tau x^{(0)} + \sigma_\tau \epsilon$.

895 **Noise prediction (Ho et al., 2020).** The model predicts the noise $\epsilon_\theta(x^{(\tau)}; \tau) \approx \epsilon$, yielding the score
 896 via

$$897 s_\theta(x^{(\tau)}; \tau) = -\frac{1}{\sigma_\tau} \epsilon_\theta(x^{(\tau)}; \tau), \quad (12)$$

900 and is trained with the MSE objective $\mathbb{E}[\|\epsilon - \epsilon_\theta(x^{(\tau)}; \tau)\|_2^2]$.

901 **Data (clean sample) prediction (Sohl-Dickstein et al., 2015).** The model outputs a denoised esti-
 902 mate $x_\theta(x^{(\tau)}; \tau) \approx x^{(0)}$. Converting to a score gives

$$904 s_\theta(x^{(\tau)}; \tau) = -\frac{x^{(\tau)} - \alpha_\tau x_\theta(x^{(\tau)}; \tau)}{\sigma_\tau^2}, \quad (13)$$

906 which is equivalent to first forming $\hat{\epsilon} = (x^{(\tau)} - \alpha_\tau x_\theta)/\sigma_\tau$ and then using $s_\theta = -\hat{\epsilon}/\sigma_\tau$. Training
 907 objectives is minimizing $\mathbb{E}[\|x^{(0)} - x_\theta(x^{(\tau)}; \tau)\|_2^2]$.

908 **v-prediction (Salimans & Ho, 2022).** A time-dependent linear combination is predicted:

$$909 v_\theta(x^{(\tau)}; \tau) \approx \alpha_\tau \epsilon - \sigma_\tau x^{(0)}. \quad (14)$$

910 From v_θ one can recover all other targets:

$$912 \hat{\epsilon}(x^{(\tau)}; \tau) = \frac{\sigma_\tau x^{(\tau)} + \alpha_\tau v_\theta(x^{(\tau)}; \tau)}{\alpha_\tau^2 + \sigma_\tau^2}, \quad (15)$$

$$914 x_\theta(x^{(\tau)}; \tau) = \frac{\alpha_\tau x^{(\tau)} - \sigma_\tau v_\theta(x^{(\tau)}; \tau)}{\alpha_\tau^2 + \sigma_\tau^2}, \quad (16)$$

$$917 s_\theta(x^{(\tau)}; \tau) = -\frac{1}{\sigma_\tau} \hat{\epsilon}(x^{(\tau)}; \tau). \quad (17)$$

918 The training objective becomes $\mathbb{E}[\|\alpha_\tau \epsilon - \sigma_\tau x^{(0)} - v_\theta(x^{(\tau)}; \tau)\|_2^2]$.
 919

920 All these parameterizations are connected by τ -dependent linear transforms, and thus represent the
 921 same model class. Choosing among them mainly affects optimization stability and the weighting of
 922 errors across noise levels.

923 A.2 FLOW MATCHING

925 Flow matching (Lipman et al., 2022; Liu et al., 2022; Esser et al., 2024) simplifies score matching
 926 by defining the generative process via ordinary differential equations (ODEs). Specifically, given
 927 the same definitions of $x, x^{(\tau)}, \alpha_\tau, \sigma_\tau$ as in Appendix A.1, the generative process is governed by a
 928 probability flow ODE:

$$929 \quad \frac{dx^{(\tau)}}{d\tau} = v(x^{(\tau)}; \tau), \quad (18)$$

931 where the velocity field ³ $v(x^{(\tau)}; \tau)$ satisfies:

$$933 \quad v(x^{(\tau)}, \tau) = \mathbb{E} \left[\frac{dx^{(\tau)}}{d\tau} \middle| x^{(\tau)} \right] = \dot{\alpha}_\tau \mathbb{E}[x^{(0)}|x^{(\tau)}] + \dot{\sigma}_\tau \mathbb{E}[\epsilon|x^{(\tau)}], \quad (19)$$

935 where $\dot{\alpha}_\tau = \frac{d\alpha_\tau}{d\tau}$ and $\dot{\sigma}_\tau = \frac{d\sigma_\tau}{d\tau}$.

936 The flow matching objective trains a neural network $v_\theta(x^{(\tau)}; \tau)$ to minimize:

$$938 \quad \min_{\theta} \mathbb{E}_{\tau, \epsilon, x^{(0)}, x^{(\tau)}} [\|v_\theta(x^{(\tau)}; \tau) - (\dot{\alpha}_\tau x^{(0)} + \dot{\sigma}_\tau \epsilon)\|_2^2]. \quad (20)$$

940 Flow matching and score matching are connected by the score function:

$$941 \quad s(x^{(\tau)}; \tau) = -\frac{1}{\sigma_t} \mathbb{E}[\epsilon|x^{(\tau)}], \quad (21)$$

943 which corresponds to an equivalent reverse-time SDE (Ma et al., 2024a):

$$944 \quad dx^{(\tau)} = v(x^{(\tau)}; \tau) d\tau - \frac{1}{2} \eta_\tau s(x^{(\tau)}; \tau) d\tau + \sqrt{\eta_\tau} d\bar{w}_t, \quad (22)$$

946 where η_τ controls the amount of stochasticity and \bar{w}_t is a reverse-time Wiener process as in Ap-
 947 pendix A.1. Solving Equations (19) and (21), we obtain:

$$948 \quad s(x^{(\tau)}, \tau) = \frac{1}{\sigma_\tau} \cdot \frac{\alpha_\tau v(x^{(\tau)}, \tau) - \dot{\alpha}_\tau x^{(\tau)}}{\dot{\alpha}_\tau \sigma_\tau - \alpha_\tau \dot{\sigma}_\tau}. \quad (23)$$

950 Thus, flow matching and score matching represent the same model class.

951 Flow matching is easy to implement and usually converges faster than score matching in practice
 952 (Liu et al., 2022; Esser et al., 2024). Another advantage of flow matching is the flexibility to choose
 953 the diffusion coefficient η_τ independently of the training process, allowing for post-hoc optimization
 954 of the sampling procedure.

956 B EXTENDED RELATED WORKS

958 **Masked and Diffusion Video Generation** Diffusion models have recently gained prominence in
 959 visual generation tasks (Ho et al., 2020; Rombach et al., 2022; He et al., 2022b; Guo et al., 2023;
 960 Chen et al., 2023; Guo et al., 2024), effectively extending to video generation. Video diffusion mod-
 961 els (Brooks et al., 2024; Ho et al., 2022) utilize bidirectional attention and binary mask embeddings
 962 to facilitate frame-level autoregressive prediction. Notable works such as GameNGen (Valevski
 963 et al., 2024) use bidirectional diffusion models for real-time game generation. However, due to
 964 their bidirectional nature, these models cannot leverage KV Cache for extended video generation,
 965 limiting their scalability. Several masked video generators, such as Genie (Bruce et al., 2024), ex-
 966 tend MaskGIT (Chang et al., 2022) into a causal-attention-based architecture for video generation.
 967 Despite their advantages, these methods suffer from the training-inference gap inherent in masked
 968 autoregressive modeling, which negatively impacts generation quality. In contrast, our approach
 969 fully leverages KV Cache during inference, facilitated by our training paradigm that bridges the
 970 training-inference gap through a novel complete teacher forcing paradigm.

971 ³The velocity field in flow matching is *different* from the v-prediction parameterization in score matching,
 972 though they are correlated: the two parameterizations are connected by Equation (23).

Table 10: Hyperparameters.

| Name | MiMo-B | MiMo-XL | | |
|----------------------|--|--|--------------------|-----------------------|
| Input | | | | |
| Dataset | Kinetics-600 | Kinetics-600 | UCF-101 | UCF-101 |
| Task | prediction | prediction | class cond. | uncond. |
| Input shape | 17 × 128 × 128 | 17 × 128 × 128 | 16 × 256 × 256 | 16 × 256 × 256 |
| VAE | | | | |
| Compression ratio | 4 × 8 × 8 | 4 × 8 × 8 | 32 × 32 | 32 × 32 |
| Latent shape | 5 × 16 × 16 | 5 × 16 × 16 | 16 × 8 × 8 | 16 × 8 × 8 |
| Architecture | | | | |
| Patch size | 1 × 1 × 1 | 1 × 1 × 1 | 1 × 1 × 1 | 1 × 1 × 1 |
| Depth | 12 | 28 | 28 | 28 |
| Embed dim | 768 | 1152 | 1152 | 1152 |
| Num heads | 12 | 16 | 16 | 16 |
| RoPE theta | 100 | 100 | 100 | 100 |
| Decoder | | | | |
| Depth | 4 | 4 | 4 | 4 |
| l | 9 | 23 | 23 | 23 |
| Diffusion | | | | |
| Parameterization | v-prediction | v-prediction | velocity | velocity |
| Noise scheduler | linear | linear | rectified flow | rectified flow |
| Weighting | fused min-SNR $\gamma = 5.0, \rho = 0.96$ | fused min-SNR $\gamma = 5.0, \rho = 0.96$ | logit-normal | logit-normal |
| Sampler | DDIM | DDIM | Euler | Euler |
| Sampling steps | 50 | 50 | 50 | 50 |
| Guidance | history guidance 1.05 | history guidance 1.05 | — | class guidance 2.0 |
| Optimization | | | | |
| Training steps | 100K | 360K | 180K | 180K |
| Batch size | 256 | 256 | 128 | 128 |
| Optimizer | AdamW | AdamW | AdamW | AdamW |
| Learning rate (LR) | 8×10^{-4} | 8×10^{-4} | 4×10^{-4} | 4×10^{-4} |
| Warmup steps | 10K | 10K | 10K | 10K |
| LR schedule | cosine | cosine | cosine | cosine |
| End LR | 10^{-5} | 10^{-5} | 10^{-5} | 10^{-5} |
| Weight decay | 0.001 | 0.001 | 0.001 | 0.001 |
| (β_1, β_2) | (0.9, 0.99) | (0.9, 0.99) | (0.9, 0.99) | (0.9, 0.99) |
| Gradient clipping | 1.0 | 1.0 | 1.0 | 1.0 |
| λ | 0.5 | 0.5 | 0.5 | 0.5 |
| Mask ratios | [0.25, 0.5] | [0.25, 0.5] | [0.25, 0.5] | [0.25, 0.5] |
| EMA decay | 0.999 | 0.999 | 0.9999 | 0.9999 |

C DATASETS

Kinetics-600 Kinetics-600 (Carreira et al., 2018) is a large-scale video action recognition dataset that extends the original Kinetics-400 dataset (Kay et al., 2017), containing approximately 500,000 video clips across 600 human action categories, sourced from YouTube and covering diverse human actions ranging from sports and cooking to dancing and musical performances. The dataset is split into training, validation, and test sets, with each action class containing at least 600 video clips in the training set and 50 clips in both validation and test sets. Videos in Kinetics-600 are characterized by

1026 their temporal dynamics and complex motion patterns, making it a challenging benchmark for video
 1027 understanding tasks. The dataset provides rich temporal information and diverse visual content,
 1028 which makes it particularly suitable for evaluating autoregressive video modeling approaches that
 1029 need to capture long-term temporal dependencies and generate coherent future frames based on
 1030 historical context. Following existing works (Song et al., 2025), we use a resolution of 128×128
 1031 pixels and train on the training set while evaluating on the test set. The model is conditioned on the
 1032 first 5 frames and predicts the next 11 frames, totaling 16 frames.
 1033

1034 **UCF-101** UCF-101 (Soomro et al., 2012) is a widely used action recognition dataset consisting of
 1035 13,320 video clips distributed across 101 action categories. The dataset was collected from YouTube
 1036 and contains realistic videos with significant variations in camera motion, object appearance, pose,
 1037 scale, viewpoint, cluttered background, and illumination conditions. Each action class contains 25
 1038 groups of videos, with each group sharing common features such as similar backgrounds, simi-
 1039 lar viewpoints, etc. UCF-101 covers a diverse range of human actions, including sports activities
 1040 (e.g., basketball, tennis, surfing), musical instrument playing, and daily life activities. Despite be-
 1041 ing smaller in scale compared to Kinetics datasets, UCF-101 remains a fundamental benchmark for
 1042 assessing the generalization capability of video models across different domains and action com-
 1043 plexities, due to its well-curated action categories. We follow the protocol of Gu et al. (2025) and
 1044 use a resolution of 256×256 pixels. The models are trained on the full UCF-101 dataset and
 1045 evaluated with class labels as the only initial condition, generating a total of 16 frames.
 1046

D IMPLEMENTATION DETAILS

1049 Table 10 summarizes the hyperparameters we use in our implementations. The details are discussed
 1050 in the following sections.

D.1 MODEL ARCHITECTURES

1054 **Diffusion Models** We employ the Diffusion Transformer (DiT) (Peebles & Xie, 2023) with full
 1055 3D attention as our backbone. The DiT block is analogous to a vision transformer (ViT) (Dosovit-
 1056 skiy et al., 2020) block and replaces the LayerNorm (Ba et al., 2016) layers with adaptive Layer-
 1057 Norm (AdaLN) (Peebles & Xie, 2023) layers to inject diffusion timestep condition into the features.
 1058 AdaLN works by embedding the timesteps using sinusoidal positional encoding (Vaswani et al.,
 1059 2017) and feeding them to an MLP to predict the shift and bias factors for LayerNorm layers. For
 1060 class-conditioned generation, the class labels are also embedded and added into the timestep em-
 1061 beddings as additional conditions. Following existing works (Song et al., 2025; Gu et al., 2025; Hu
 1062 et al., 2024a), AdaLN is applied separately to each noisy future frame because different frames can
 1063 have different diffusion timesteps during training (Section 3.3). We use QK normalization (Henry
 1064 et al., 2020) to stabilize training. Below, we introduce two other modifications we apply to vanilla
 1065 DiT: separate LayerNorm and 3D RoPE. The vanilla DiT block and our modified DiT block are
 1066 illustrated in Figure 5.

1066 Separate LayerNorm layers instead of AdaLN are applied to the clean frames. Note that the other
 1067 parameters are shared among all frames regardless of whether they are history or future frames. The
 1068 attention mask introduced in Section 3.3 is applied to the attention operation to ensure causality
 1069 between history and future frames.

1070 Additionally, we incorporate axial 3D RoPE (Su et al., 2024) and assign an equal number of channels
 1071 to encode the positions along the T, H, W dimensions.

1073 **VAE and Patch Size of DiT** Video generation models usually work in some compressed latent
 1074 space with reduced space-time dimensions to save computations, due to the sheer volume of pixels
 1075 in video. In all of our experiments, the patch size of DiT is $1 \times 1 \times 1$ (T, H, W), meaning that
 1076 compression is done solely in VAE.

1078 For fair comparison with existing methods, we adopt the pretrained 3D video VAE of DFoT (Song
 1079 et al., 2025) for Kinetics-600 experiments. DFoT’s VAE has a compression ratio of $4 \times 8 \times 8$
 (T, H, W) , where the first frames are separately encoded and the following frames are temporally

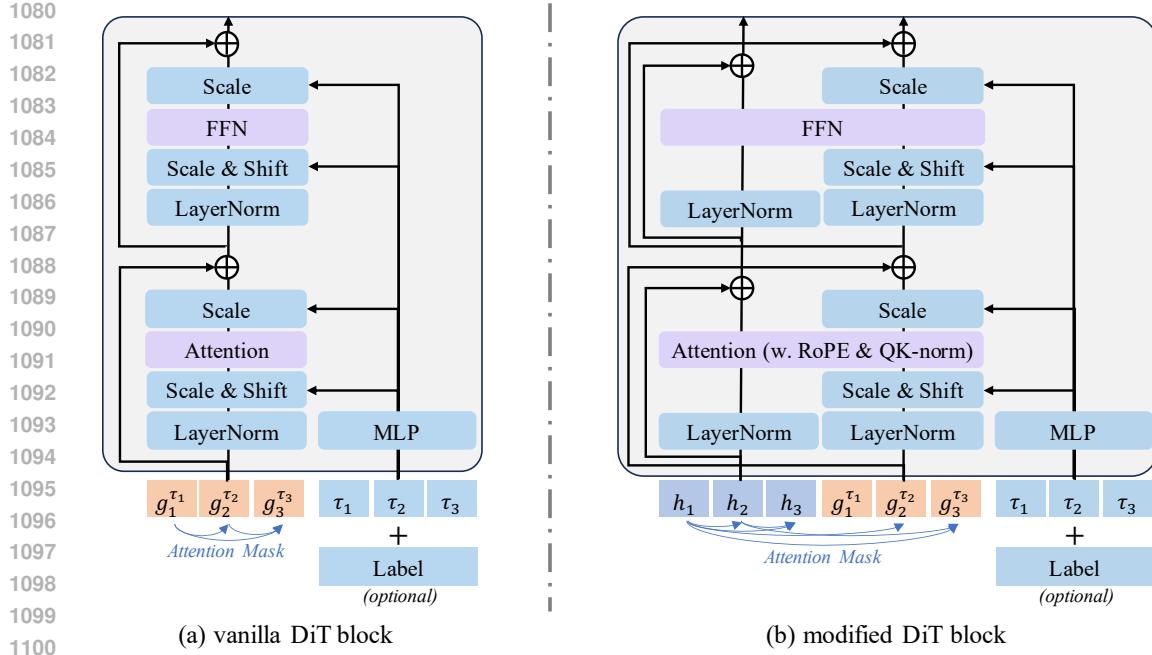


Figure 5: Illustration of vanilla and our modified DiT blocks.

downsampled by a factor of 4, following Yu et al. (2023a). The input resolution is 128×128 pixels with 17 frames, leading to a latent shape of $5 \times 16 \times 16$ per video clip.

We utilize the 2D image DC-AE of FAR (Gu et al., 2025) for UCF-101 experiments. FAR’s DC-AE has a compression ratio of 32×32 with no temporal compression, and it encodes each frame independently. Given input of 16 frames with a resolution of 256×256 , the latent shape is $16 \times 8 \times 8$.

Decoder The decoders take masked history frame features from intermediate DiT layers as the only input, and fill the masked positions with learnable query tokens. Then, the input is fed into a stack of several decoder blocks and reshaped to the same dimensions as the latents of the history (or future) frames as the output. The decoder block is the vanilla ViT block with axial 3D RoPE.

D.2 DIFFUSION

Kinetics-600 Experiments For Kinetics-600, we use a linear noise schedule (Nichol & Dhariwal, 2021) with the v-prediction parameterization (Salimans & Ho, 2022) and zero terminal SNR (Lin et al., 2024). We use the DDIM sampler (Song et al., 2020a) with 50 sampling steps during inference. We also incorporate the fused min-SNR loss weighting (Chen et al., 2024a), a variant of min-SNR loss weighting (Hang et al., 2023) for video diffusion, to accelerate convergence.

Fused min-SNR extends the standard min-SNR loss weighting by accounting for the “signals” from previous frames. The difference between the two methods is the way to compute the signal-to-noise ratio (SNR) used for loss weighting. Using the notations in Section 2, SNR is defined as $\text{SNR}_\tau = \sigma_\tau^2 / \sigma_\tau^2$. Fused min-SNR first normalizes the SNR to $[0, 1]$ by dividing by the maximal value of SNR. Since min-SNR weighting clips the SNR value with the hyperparameter $\gamma > 0$, we normalize by γ . Then, it computes fused SNR S'_t with decaying factor $\rho > 0$:

$$S_t = \text{normalized SNR factor for the } t\text{-th noisy future frame} \in [0, 1], \quad (24)$$

$$\bar{S}_t = \rho \bar{S}_{t-1} + (1 - \rho) S_t \quad (\text{exponentially decayed cumulative SNR}), \quad (25)$$

$$S'_t = 1 - (1 - S_t)(1 - \bar{S}_{t-1}) \quad (\text{fused reweighting factor}), \quad (26)$$

Fused SNR S'_t combines the current frame signal with accumulated history signals, treating them as independent probabilistic events. This accounts for the additional information available from history

```

1134
1135 1 def compute_loss_weight(snr, gamma, prediction_type, decay=None,
1136  causal=True):
1137 2     """Compute SNR weighting.
1138
1139 3     Args:
1140 4         snr (torch.Tensor): per-frame SNR of shape [B, T]
1141 5         gamma (float): clip threshold of min-SNR
1142 6         prediction_type (str): "epsilon", "v_prediction", or "sample"
1143 7         decay (float, optional): if not None, enable fused min-SNR with
1144 8             the specified decay factor
1145 9         causal (bool, optional): whether we are training a causal model
1146 10
1147 11     Returns:
1148 12         weight (torch.Tensor): per-frame loss weight of shape [B, T]
1149 13     """
1150 14     # Compute fused SNR
1151 15     clipped_snr = snr.clamp(max=gamma)
1152 16     if decay is not None:
1153 17         normalized_clipped_snr = clipped_snr / gamma
1154 18         normalized_snr = snr / gamma
1155 19
1156 20         def compute_cum_snr(reverse: bool = False):
1157 21             new_normalized_clipped_snr = (
1158 22                 normalized_clipped_snr.flip(1)
1159 23                 if reverse
1160 24                 else normalized_clipped_snr
1161 25             )
1162 26             cum_snr = torch.zeros_like(new_normalized_clipped_snr)
1163 27             for t in range(0, snr.shape[1]):
1164 28                 if t == 0:
1165 29                     cum_snr[:, t] = new_normalized_clipped_snr[:, t]
1166 30                 else:
1167 31                     cum_snr[:, t] = (
1168 32                         decay * cum_snr[:, t - 1]
1169 33                         + (1 - decay) * new_normalized_clipped_snr[:, t]
1170 34                     )
1171 35             cum_snr = torch.nn.functional.pad(cum_snr[:, :-1], (1, 0, 0,
1172 36                 0), value=0.0)
1173 37             return cum_snr.flip(1) if reverse else cum_snr
1174
1175 38     if causal:
1176 39         cum_snr = compute_cum_snr()
1177 40     else:
1178 41         cum_snr = compute_cum_snr(reverse=True) + compute_cum_snr()
1179 42         cum_snr *= 0.5
1180 43         clipped_fused_snr = 1 - (1 - cum_snr * decay) * (1 -
1181 44             normalized_clipped_snr)
1182 45         fused_snr = 1 - (1 - cum_snr * decay) * (1 - normalized_snr)
1183 46         clipped_snr = clipped_fused_snr * gamma
1184 47         snr = fused_snr * gamma
1185
1186 48     # Compute loss weight
1187 49     if prediction_type == "epsilon": # noise-prediction
1188 50         weight = clipped_snr / snr
1189 51     elif prediction_type == "v_prediction": # v-prediction
1190 52         weight = clipped_snr / (snr + 1)
1191 53     else: # data-prediction
1192 54         weight = clipped_snr
1193
1194 55     return weight
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2099

```

Listing 1: Fused min-SNR (PyTorch psuedo-code)

1188 context in video generation, beyond what standard SNR weighting captures. The S'_t is denormalized
 1189 by multiplying γ and used to compute the loss weighting as normal min-SNR weighting does.
 1190

1191 Listing 1 summarizes the algorithm to compute fused min-SNR weighting.
 1192

1193 **UCF-101 Experiments** For UCF-101, we follow Gu et al. (2025) and use flow matching (Liu
 1194 et al., 2022; Lipman et al., 2022; Albergo & Vanden-Eijnden, 2022) with a “straigh” flow path, i.e.,
 1195 $\alpha_\tau = 1 - \tau$, $\sigma_\tau = \tau$. We also adopt the logit-normal timestep sampling strategy (Esser et al., 2024),
 1196 where the timesteps τ are sampled from a logit-normal distribution (instead of uniformly):
 1197

$$\pi_{\text{ln}}(\tau) = \frac{1}{\sqrt{2\pi}} \frac{1}{t(1-t)} \exp\left(-\frac{\log^2(t/1-t)}{2}\right). \quad (27)$$

1200 We use the Euler integrator sampler (Esser et al., 2024) with 50 sampling steps during inference.
 1201

1203 D.3 TRAINING

1206 Algorithm 1 Training (v-prediction or flow matching)

1207 **Input:** Dataset \mathcal{D} , noise schedule $\{(\alpha_\tau, \sigma_\tau)\}_\tau$, velocity network v_θ , decoder φ_θ , loss weight λ
 1208 **Output:** Trained velocity network v_θ
 1209 1: **while** not converged **do**
 1210 2: Sample video clip $\mathbf{x} = \{x_t\}_{t=1}^T$ from \mathcal{D}
 1211 3: $\mathbf{h} \leftarrow \mathbf{x}$, $\mathbf{f} \leftarrow \mathbf{x}$ // Duplicate \mathbf{x} as history \mathbf{h} and future \mathbf{f}
 1212 4: Sample $\{\tau_t \sim \text{Uniform}[0, 1]\}_{t=1}^T$ and $\{\epsilon_t \sim \mathcal{N}(0, I)\}_{t=1}^T$
 1213 5: Sample random tube mask \mathcal{M}
 1214 6: $h_t^{\mathcal{M}} \leftarrow \text{applyMask}(h_t, \mathcal{M})$ // Apply \mathcal{M} to history frame
 1215 7: $\mathcal{L} \leftarrow 0$
 1216 8: **for** $t = 1$ to T **do**
 1217 9: $f_t^{(\tau_t)} \leftarrow \alpha_{\tau_t} f_t + \sigma_{\tau_t} \epsilon_t$ // Add noise to future frame
 1218 10: $v_{\text{target}} \leftarrow \alpha_{\tau_t} \epsilon_t - \sigma_{\tau_t} f_t$ or $v_{\text{target}} \leftarrow \dot{\alpha}_{\tau_t} f_t + \dot{\sigma}_{\tau_t} \epsilon_t$ // v-prediction or flow matching
 1219 11: $v_{\text{pred}} \leftarrow v_\theta(f_t^{(\tau)}; \tau, h_{1:t}^{\mathcal{M}})$ // Internally apply attention mask (Figure 3(a))
 1220 12: $\mathcal{L}_{\text{diff}} \leftarrow \|v_{\text{pred}} - v_{\text{target}}\|_2^2$ // Diffusion loss (Equation (4))
 1221 13: $v_{\text{feat}}^{h,l} \leftarrow v_\theta^{h,l}(f_t^{(\tau)}; \tau, h_{1:t}^{\mathcal{M}})$ // Output features of the l -th layer for history frames $h_{1:t}$
 1222 14: $\mathcal{T}_t \leftarrow \{t, t+1\}$ // Frame indexes
 1223 15: $\mathcal{L}_{\text{mask}} \leftarrow \frac{1}{|\mathcal{T}_t|} \sum_{t' \in \mathcal{T}_t} \|h_{t'} - \varphi_\theta(t' - t, v_{\text{feat}}^{h,l})\|_2^2$ // Masked history modeling loss
 (Equation (5))
 1224 16: $\mathcal{L} \leftarrow \mathcal{L} + \mathcal{L}_{\text{diff}} + \lambda \mathcal{L}_{\text{mask}}$
 1225 17: **end for**
 1226 18: Update θ using $\nabla_\theta \mathcal{L}$
 1227 19: **end while**

1229
 1230 The training algorithm is summarized in Algorithm 1. Training hyperparameters are summarized in
 1231 Table 10. Note that the masked modeling loss (Equation (5)) is computed in the latent space between
 1232 the latents of the history (or future) frames and the predictions of the decoders.
 1233

1234 D.4 INFERENCE

1235
 1236 The sampling algorithm is summarized in Algorithm 2. Inference hyperparameters are summarized
 1237 in Table 10. The inference process is the same as in other diffusion-based video generation models
 1238 (Song et al., 2025; Gu et al., 2025; Hu et al., 2024a): Given initial conditions (initial frames for
 1239 video prediction, class labels for class-conditioned generation, or no condition for unconditioned
 1240 generation), the model iteratively denoises the next frame starting from pure noise, and appends the
 1241 generated frame after the known (provided as initial conditions or generated) frames until all frames
 are known.

D.5 HISTORY GUIDANCE

We incorporate a simplified version of history guidance (Song et al., 2025) into diffusion-based VideoAR. History guidance takes advantage of the insight that history frames are the conditions for generating the future frames, much like class labels as conditions for class-conditioned generation, and applies classifier-free guidance (CFG) (Ho & Salimans, 2022) with history frames as conditions. Adopting the notations in Section 2, history guidance modifies the score function as

$$s_\theta^w(x_{t+1}^{(\tau)}; \tau, x_{1:t}) = (1-w) \cdot s_\theta(x_{t+1}^{(\tau)}; \tau, \emptyset) + w \cdot s_\theta(x_{t+1}^{(\tau)}; \tau, x_{1:t}), \quad (28)$$

where \emptyset means no history frame and $w > 1$ is the guidance scale. We compute $s_\theta(x_{t+1}^{(\tau)}; \tau, \emptyset)$ by forbidding $(x_{t+1}^{(\tau)})$ to attend to $x_{1:t}$ via attention masking, i.e., by setting the corresponding noisy-to-clean rows in the attention mask (Figure 3(a)) to $-\infty$.

During training, we randomly select $r = 10\%$ future frames and forbid them from attending to the history frames. The training algorithm with history guidance is summarized in Algorithm 3.

During inference, $s_\theta^w(x_{t+1}^{(\tau)}; \tau, \emptyset)$ is computed by Equation (28) and the other process is the same as in normal CFG. The sampling algorithm with history guidance is summarized in Algorithm 4.

E BASELINES

In our work, we primarily consider three baseline methods in Figure 1 and Section 3.2. All the considered baselines are trained with the *same* model architecture and hyperparameters as shown in Table 10 unless otherwise specified, with the only difference being the training strategies.

ACDiT ACDiT (Hu et al., 2024a) also adopts complete teacher forcing as in MiMo. The primary difference between MiMo and ACDiT is that we apply the masked history modeling target on the history frames. Thus, the direct comparison between ACDiT and MiMo clearly demonstrates the advantage of our approach and the benefit of good history representations.

FAR FAR (Gu et al., 2025) adopts diffusion forcing (Chen et al., 2024a; Song et al., 2025), it randomly replaces some noisy frames with their clean version to simulate clean history frames. The better performance of MiMo over FAR demonstrates that MiMo can achieve competitive performance even against the best performing models in a broader context.

REPA REPA (Yu et al., 2024) was originally proposed to improve the representation quality of the noisy images being denoised. We adapt it to diffusion-based VideoAR following the approach of Zhang et al. (2025). Compared with REPA, we focus on the representations of the history frames that serve as conditions in AR modeling, while REPA does not consider the AR context. Also, REPA requires a VFM, but sVFM may not always be available and may misbehave for out-of-distribution (OOD) data, while MiMo does not rely on VFM.

For the analysis in Section 3.2, we align the features of the 4-th layer with a pretrained VideoMAE-L (Tong et al., 2022) using a loss weight of 0.5. The feature dimensions of DiT models and Video-

1296 **Algorithm 3** Training with **history guidance** (v-prediction or flow matching)

1297 **Input:** Dataset \mathcal{D} , noise schedule $\{(\alpha_\tau, \sigma_\tau)\}_\tau$, velocity network v_θ , decoder φ_θ , loss weight λ ,
1298 **drop rate** r

1299 **Output:** Trained velocity network v_θ

1300 1: **while** not converged **do**

1301 2: Sample video clip $\mathbf{x} = \{x_t\}_{t=1}^T$ from \mathcal{D}

1302 3: $\mathbf{h} \leftarrow \mathbf{x}$, $\mathbf{f} \leftarrow \mathbf{x}$ // Duplicate \mathbf{x} as history \mathbf{h} and future \mathbf{f}

1303 4: Sample $\{\tau_t \sim \text{Uniform}[0, 1]\}_{t=1}^T$ and $\{\epsilon_t \sim \mathcal{N}(0, I)\}_{t=1}^T$

1304 5: Sample random tube mask \mathcal{M}

1305 6: $h_t^{\mathcal{M}} \leftarrow \text{applyMask}(h_t, \mathcal{M})$ // Apply \mathcal{M} to history frame

1306 7: $\mathcal{L} \leftarrow 0$

1307 8: **for** $t = 1$ to T **do**

1308 9: $f_t^{(\tau_t)} \leftarrow \alpha_{\tau_t} f_t + \sigma_{\tau_t} \epsilon_t$ // Add noise to future frame

1309 10: $v_{\text{target}} \leftarrow \alpha_{\tau_t} \epsilon_t - \sigma_{\tau_t} f_t$ or $v_{\text{target}} \leftarrow \dot{\alpha}_{\tau_t} f_t + \dot{\sigma}_{\tau_t} \epsilon_t$ // v-prediction or flow matching

1310 11: **if** $\text{Uniform}[0, 1] < r$ **then**

1311 12: $v_{\text{pred}} \leftarrow v_\theta(f_t^{(\tau)}; \tau, \emptyset)$ // Randomly drop history frames

1312 13: **else**

1313 14: $v_{\text{pred}} \leftarrow v_\theta(f_t^{(\tau)}; \tau, h_{1:t}^{\mathcal{M}})$ // Internally apply attention mask (Figure 3(a))

1314 15: **end if**

1315 16: $\mathcal{L}_{\text{diff}} \leftarrow \|v_{\text{pred}} - v_{\text{target}}\|_2^2$ // Diffusion loss (Equation (4))

1316 17: $v_{\text{feat}}^{h,l} \leftarrow v_\theta(f_t^{(\tau)}; \tau, h_{1:t}^{\mathcal{M}})$ // Output features of the l -th layer for history frames $h_{1:t}$

1317 18: $\mathcal{T}_t \leftarrow \{t, t+1\}$ // Frame indexes

1318 19: $\mathcal{L}_{\text{mask}} \leftarrow \frac{1}{|\mathcal{T}_t|} \sum_{t' \in \mathcal{T}_t} \|h_{t'} - \varphi_\theta(t' - t, v_{\text{feat}}^{h,l})\|_2^2$ // Masked history modeling loss

1319 (Equation (5))

1320 20: $\mathcal{L} \leftarrow \mathcal{L} + \mathcal{L}_{\text{diff}} + \lambda \mathcal{L}_{\text{mask}}$

1321 21: **end for**

1322 22: Update θ using $\nabla_\theta \mathcal{L}$

1323 23: **end while**

1324 **Algorithm 4** Sampling with **history guidance** (v-prediction or flow matching)

1325 **Input:** Noise schedule $\{(\alpha_\tau, \sigma_\tau)\}_\tau$, sampling steps N , **guidance scale** w , velocity network v_θ ,
1326 initial frames $x_{1:t_0}$ (\emptyset if $t_0 = 0$)

1327 **Output:** Clean frames $x_{1:T}$

1328 1: **for** $t = t_0 + 1$ to T **do**

1329 2: $x_t \sim \mathcal{N}(0, I)$ // Initialize with noise at $\tau = \tau_N = 1$

1330 3: **for** $i = N$ to 1 **do**

1331 4: $v_{\text{pred}} \leftarrow v_\theta(x_t; \tau_i, x_{1:t-1})$ // Internally apply block-causal attention mask

1332 5: $v_\emptyset \leftarrow v_\theta(x_t; \tau_i, \emptyset)$ // Negative condition

1333 6: $v_{\text{pred}} \leftarrow (1 - w) \cdot v_{\text{pred}} + w \cdot v_\emptyset$ // Apply guidance

1334 7: $x_t \leftarrow \text{Sampler}(x_t; \tau_i, \tau_{i-1}, v_{\text{pred}})$ // Sampler step, $\tau_0 = 0$

1335 8: **end for**

1336 9: $x_{1:t} \leftarrow x_{1:t-1} + \{x_t\}$ // Append generated frame after known frames

1337 10: **end for**

1338
1339 MAE are aligned following the strategy of Zhang et al. (2025), which interpolates the DiT’s representations
1340 to match the feature dimensions of the pre-trained VideoMAE.
1341

F EVLAUTION DETAILS

1342
1343 **Fréchet Video Distance (FVD)** FVD (Unterthiner et al., 2018) is a perceptual metric designed
1344 to evaluate the quality of generated videos by measuring the distributional distance between real
1345 and generated video sequences. Similar to the Fréchet Inception Distance (FID) (Heusel et al.,
1346 2017) used for images, FVD employs a pre-trained 3D convolutional neural network (specifically,
1347 an Inflated 3D ConvNet or I3D model trained on Kinetics-400) (Carreira & Zisserman, 2017) to
1348 extract spatio-temporal features from video clips. Then Fréchet distance (Dowson & Landau, 1982)
1349

1350 is computed between the feature distributions of real and generated videos by fitting multivariate
 1351 Gaussian distributions to the extracted features and calculating the Wasserstein-2 distance (Villani
 1352 et al., 2008) between them. Lower FVD scores indicate higher similarity to real video distributions.
 1353 FVD captures both spatial and temporal aspects of video content, making it a standard evaluation
 1354 tool in video synthesis research. Following prior works (Song et al., 2025), we compute FVD for
 1355 the entire video, including both initial conditioning frames (for Kinetics-600) and generated frames,
 1356 to assess the overall consistency.

1357 **Centered Kernel Nearest-Neighbor Alignment (CKNNA)** CKNNA (Huh et al., 2024) is a non-
 1358 parametric evaluation metric that measures the alignment between two sets of features by analyzing
 1359 their local neighborhood structures. CKNNA relaxes the overly rigid Centered Kernel Alignment
 1360 (CKA) (Kornblith et al., 2019) metric by computing similarity only for the mutual nearest
 1361 neighbors of each feature vector. Given two sets of vectorized features $\{\phi_i \in \mathbb{R}^n\}$ and $\{\psi_i \in \mathbb{R}^m\}$ from
 1362 two models and inner product operator $\langle \cdot, \cdot \rangle$, CKNNA first computes centered kernel matrices:
 1363

$$\bar{\mathbf{K}}_{ij} = \langle \phi_i, \phi_j \rangle - \mathbb{E}_l[\langle \phi_i, \phi_l \rangle], \quad \bar{\mathbf{L}}_{ij} = \langle \psi_i, \psi_j \rangle - \mathbb{E}_l[\langle \psi_i, \psi_l \rangle] \quad (29)$$

1364 The centering operation removes the mean similarity, focusing on relative relationships rather than
 1365 absolute magnitudes. CKNNA restricts the alignment computation to mutual nearest neighbors:
 1366

$$\text{Align}_{\text{knn}}(\mathbf{K}, \mathbf{L}) = \sum_i \sum_j \alpha(i, j) \cdot \bar{\mathbf{K}}_{ij} \bar{\mathbf{L}}_{ij} \quad (30)$$

$$\text{where } \alpha(i, j) = \mathbf{1}[\phi_j \in \text{knn}(\phi_i) \wedge \psi_j \in \text{knn}(\psi_i) \wedge i \neq j] \quad (31)$$

1367 The indicator function $\alpha(i, j)$ ensures we only consider sample pairs whose members are nearest
 1368 neighbors to each other, emphasizing local structural consistency over global alignment. The final
 1369 CKNNA metric is the normalized version:
 1370

$$\text{CKNNA}(\mathbf{K}, \mathbf{L}) = \frac{\text{Align}_{\text{knn}}(\mathbf{K}, \mathbf{L})}{\sqrt{\text{Align}_{\text{knn}}(\mathbf{K}, \mathbf{K}) \cdot \text{Align}_{\text{knn}}(\mathbf{L}, \mathbf{L})}} \quad (32)$$

1371 This normalization bounds the metric to $[0, 1]$, where higher values indicate better preservation of
 1372 local neighborhood structure between the two representation spaces. Intuitively, CKNNA measures
 1373 whether two feature representations maintain similar local similarity structures within their respec-
 1374 tive neighborhoods. Following prior works (Huh et al., 2024; Yu et al., 2024), we evaluate represen-
 1375 tation alignment using CKNNA with $k = 10$ nearest neighbors. We randomly sample 10,000 videos
 1376 from the Kinetics-600 test set and extract globally average pooled features using both a pretrained
 1377 VideoMAE-L (Tong et al., 2022) (as reference) and our models, treating all frames as clean history
 1378 frames. Then, we compute CKNNA between the features of VideoMAE-L and features from each
 1379 layer of the query models, reporting the highest alignment score across all layers.
 1380

1381 **Linear Probing** We follow the linear probing protocol of MAE (He et al., 2022a). Specifically, we
 1382 use the model representations of the clean history frames for linear probing training and evaluation.
 1383 Global average pooling is applied to the output feature map to obtain a single feature vector for each
 1384 video. The feature vector is then fed to a parameter-free BatchNorm (Ioffe & Szegedy, 2015) layer
 1385 and a linear classifier layer. The training batch size is 128, the learning rate is 10^{-3} and decayed to
 1386 0 with a cosine schedule, the weight decay is 0.01, and the training length is 10 epochs. Random
 1387 flipping is used during training. Top-1 accuracy is reported.
 1388

1389 G ADDITIONAL VISUALIZATIONS

1390 G.1 SAMPLES

1391 This section shows the samples generated by MiMo on Kinetics-600 (Figure 6), UCF-101 class-
 1392 conditioned generation (Figure 7), and unconditional generation (Figure 8). Each row is a generated
 1393 video containing 16 frames.
 1394

1395 G.2 ATTENTION HEATMAPS

1396 In Figures 9 and 10 we show the attention heatmaps (marked by red) of two videos, each without
 1397 and with MiMo. The center position of the last frame (marked by a blue dot) serves as the query,
 1398

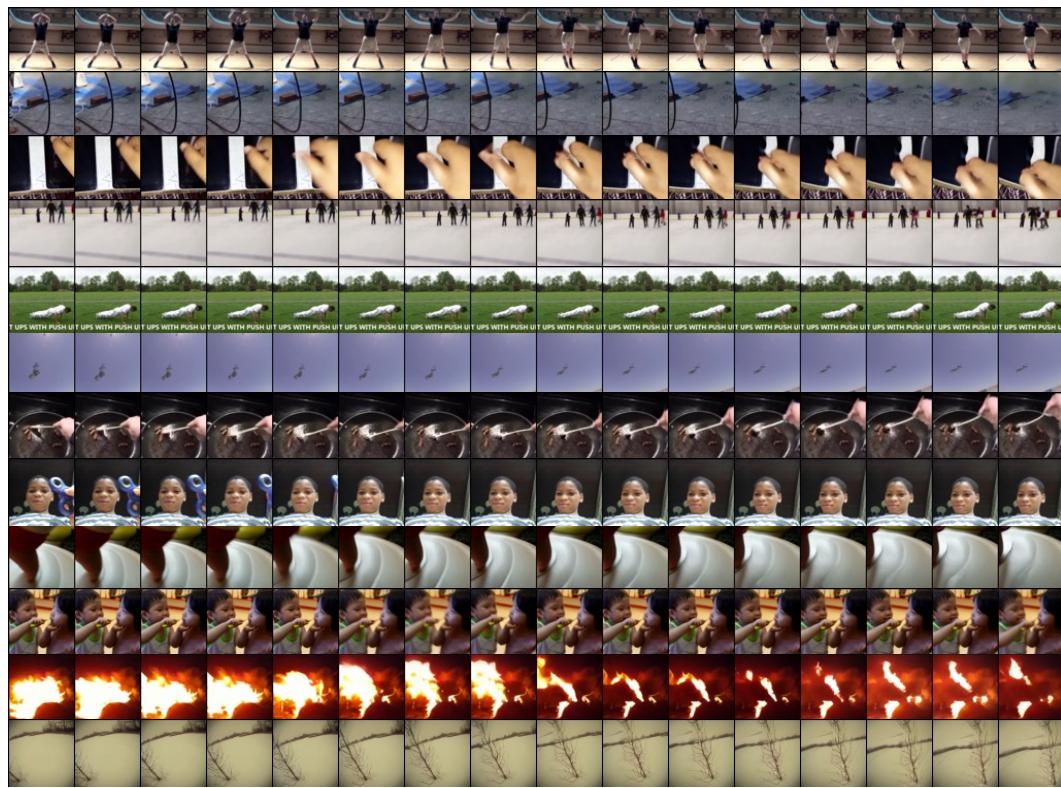


Figure 6: Uncurated samples of Kinetics history frames conditional generation.

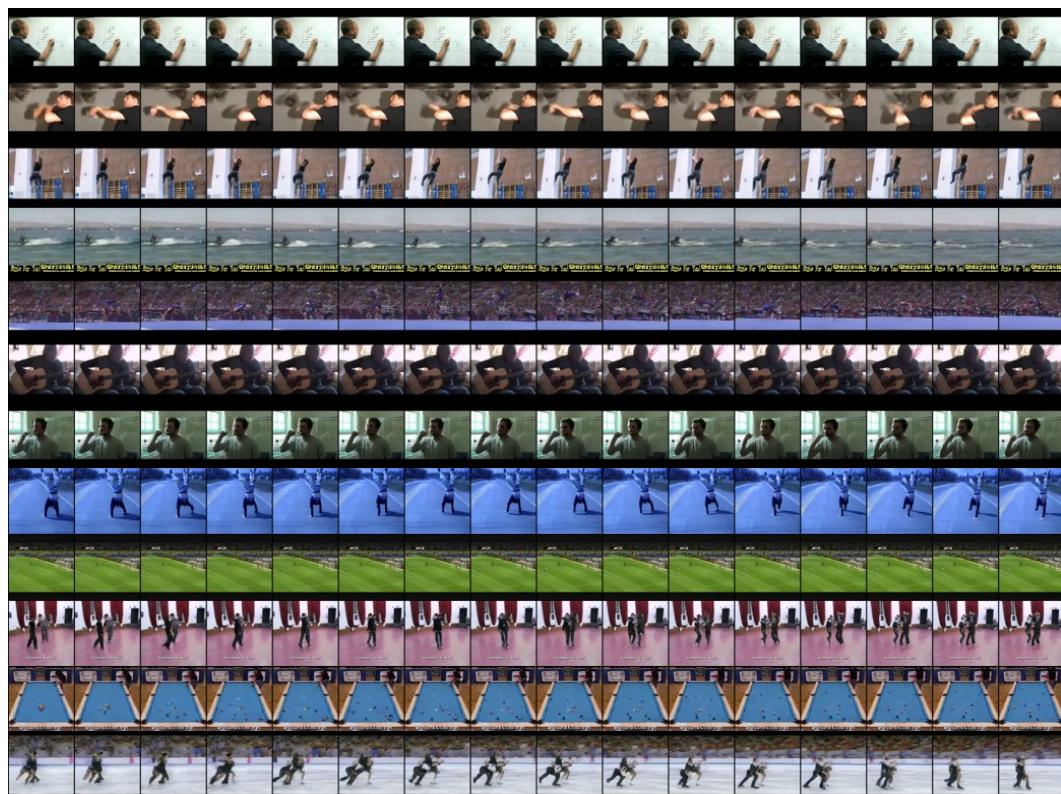


Figure 7: Uncurated samples of UCF-101 class-conditioned generation.

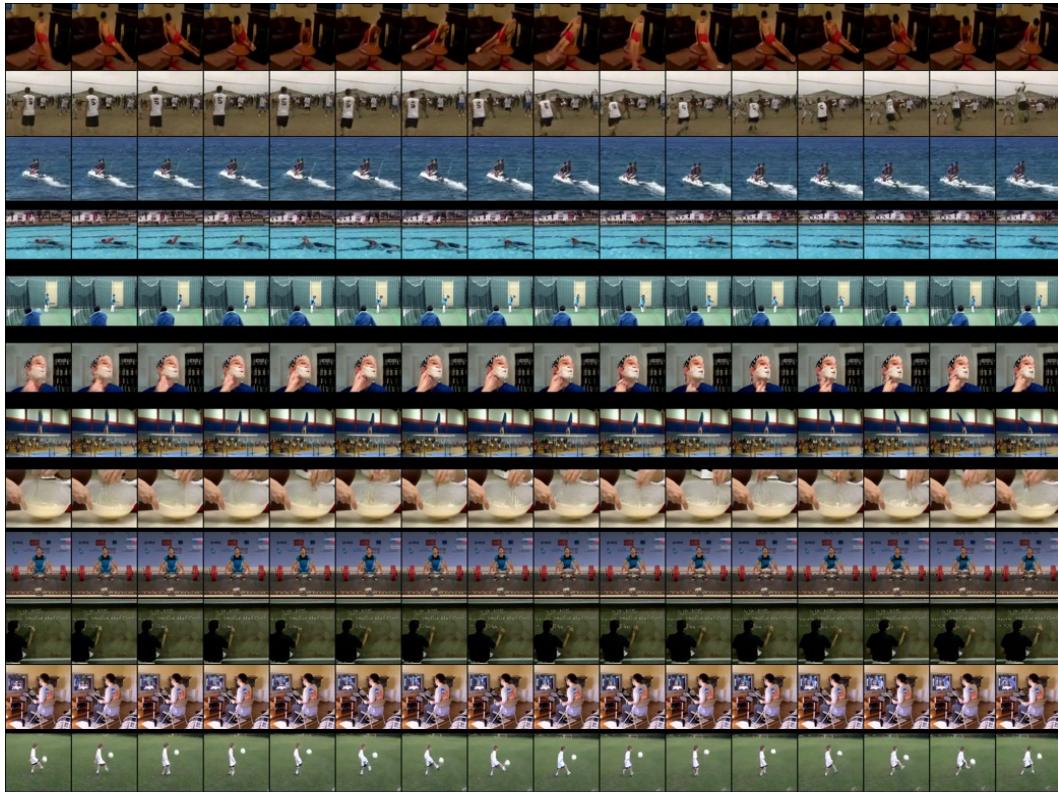


Figure 8: Uncurated samples of UCF-101 unconditional generation.

while the other positions are the keys. We take the attention weights from four random heads and layers 4, 8, 12, 16, 20.

As shown in Figures 9 and 10, without MiMo, attention patterns are more dispersed and less focused, whereas with MiMo, attention heatmaps show more concentrated patterns that exhibit stronger semantic correlations with the query content. Additionally, Figure 10 demonstrates how different transformer layers specialize in matching distinct body parts (e.g., arms, torso, legs), revealing the hierarchical nature of the learned representations.

G.3 EMBEDDING VISUALIZATION

Figure 11 shows the UMAP (McInnes et al., 2018) visualization of video embeddings without and with MiMo. The model is a DiT-XL and is trained for 360K steps on Kinetics-600. As shown in Figure 11, without MiMo, the distribution of video embeddings is mostly uniform, while MiMo introduces some structures in the embedding distribution by learning a more structured representation space.

H ADDITIONAL EXPERIMENTS

H.1 COMPLEMENTARY EVALUATION METRICS

Our evaluation mainly relies on the standard FVD metric. However, FVD is known to have several issues and may not fully capture real-world video dynamics (Luo et al., 2024). To provide a more comprehensive evaluation of MiMo’s performance, we adopt two additional metrics as complements.

VMMD Our VMMD (V-JEPA 2 Maximum Mean Discrepancy) metric is based on the CMMD metric (Javasumana et al., 2024). The VMMD metric benchmarks the perceptual similarity between

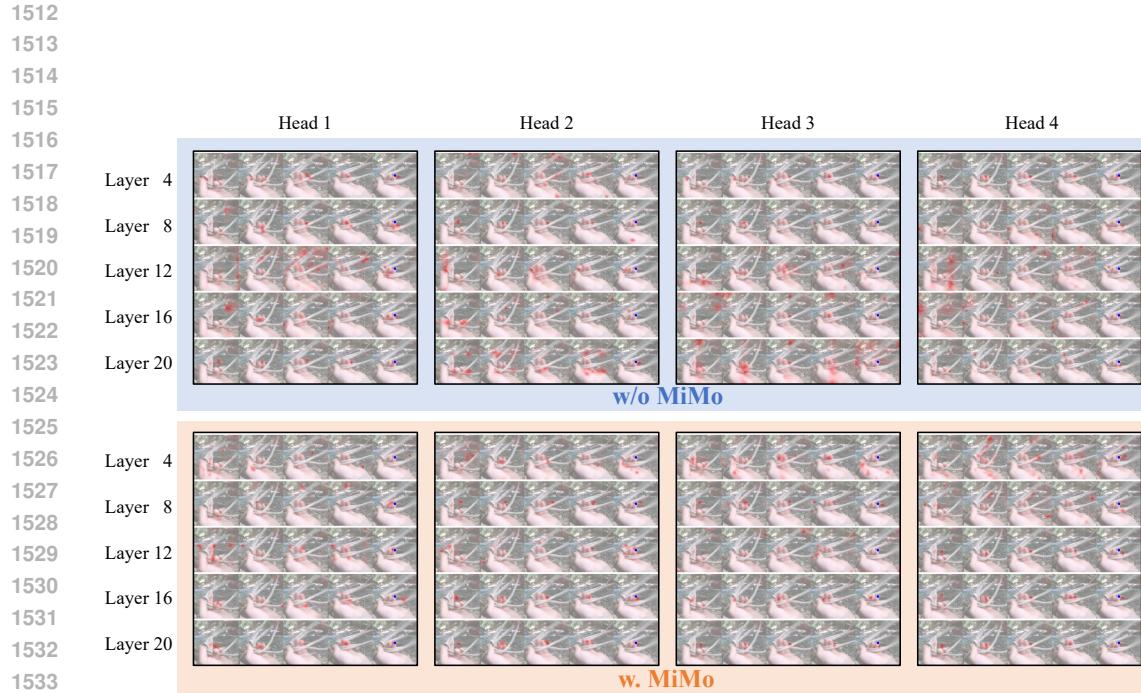


Figure 9: Attention heatmaps without and with MiMo.

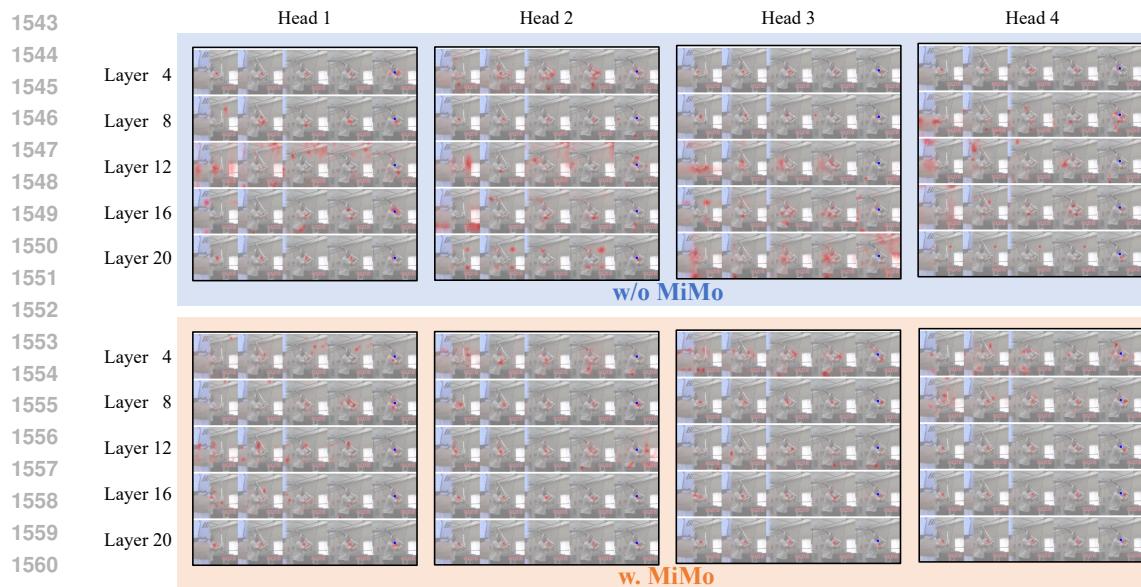


Figure 10: Attention heatmaps without and with MiMo.

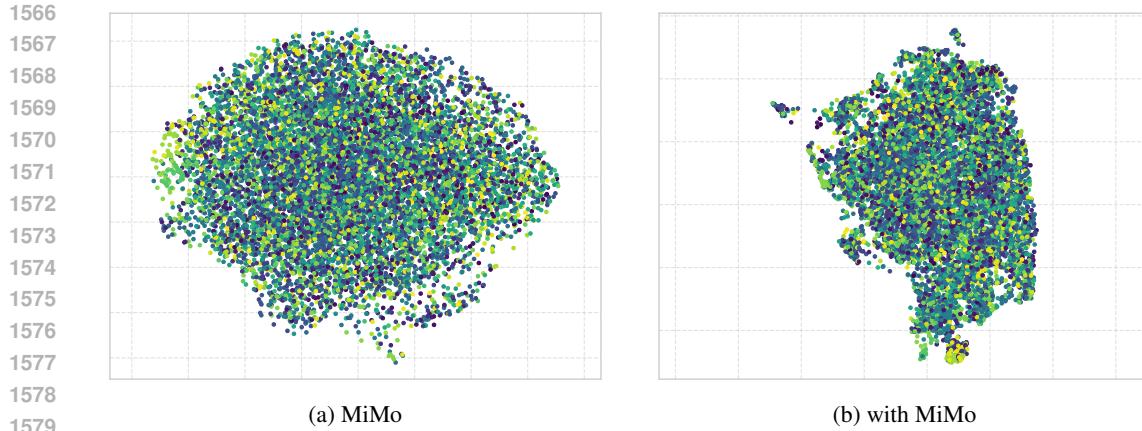


Figure 11: UMAP visualization of video embeddings without and with MiMo.

the generated videos and the reference videos, using the strong V-JEPA 2 (Assran et al., 2025) pretrained model as the judge. It does not rely on the Gaussian assumption of the FVD metric, and gives more faithful evaluation (Jayasumana et al., 2024). Specifically, VMMD replaces the CLIP model in CMMD with V-JEPA 2 Large; other implementations are the same as in CMMD⁴.

Table 11: Comparison of different methods with VMMD and FVD metrics.

| Method | ACDiT-XL | FAR-XL | MiMo-XL |
|--------|----------|--------|--------------|
| FVD↓ | 10.264 | 9.311 | 8.257 |
| VMMD↓ | 1.075 | 1.036 | 0.977 |

Table 11 compares ACDiT, FAR, and MiMo with both VMMD and FVD metrics. The VMMD measurement results are consistent with the FVD, indicating that in our cases, FVD and VMMD can relatively well characterize the generation quality.

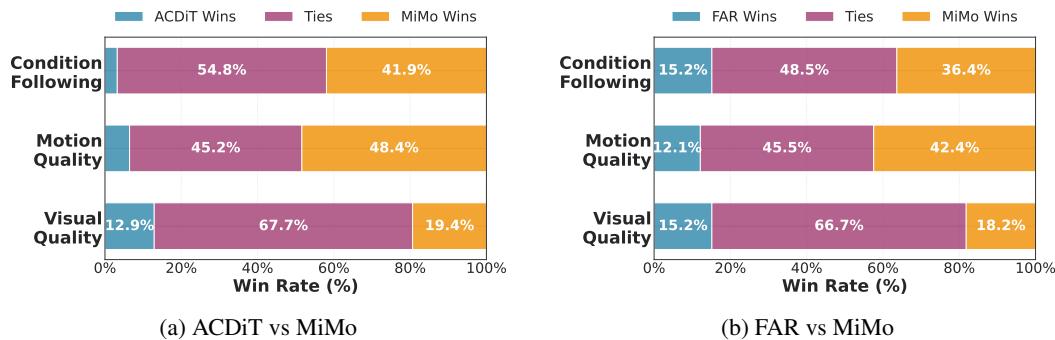


Figure 12: User studies (win rates) of ACDiT, FAR, and MiMo.

User Studies We conduct user studies to better understand what aspects MiMo improves. Five experts are instructed to evaluate 100 tasks, assessing three key dimensions: condition following, motion quality, and visual quality.

- Condition following: the visual and semantic consistency between conditioning frames and generated frames.
- Motion quality: whether there is motion distortion or motion that is semantically inconsistent with the context.

⁴<https://github.com/google-research/google-research/tree/master/cmmd>

1620
 1621 • Visual quality: whether there is frame-level visual distortion or visual components that are
 1622 semantically inconsistent with the context.

1623 Figure 12 summarizes the results for ACDiT-XL, FAR-XL, and MiMo-XL, all trained for 360K steps
 1624 on Kinetics-600. MiMo excels in condition following and motion quality, while the visual quality
 1625 is marginally improved. Additionally, MiMo has higher win rates against ACDiT than against FAR,
 1626 which indicates that lower VMMD and FVD metric values correlate with better perceptual quality
 1627 in our cases.

1628
 1629 **H.2 LONG-HORIZON VIDEO GENERATION**
 1630

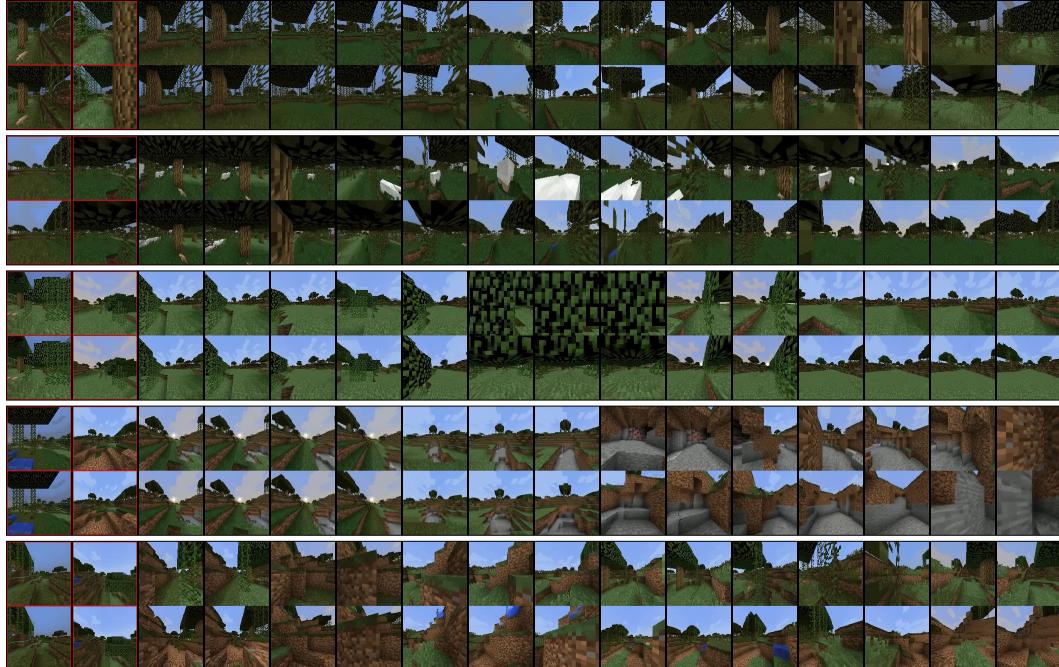
1631 MiMo is effective for long-horizon video generation, as its robust history representation helps mit-
 1632 ate the error accumulation common in autoregressive models. We validate this on the action-
 1633 conditioned Minecraft dataset (Yan et al., 2023), predicting 156 frames from 144, following the
 1634 FAR (Gu et al., 2025) setup.

1635
 1636 **Table 12: Long-horizon video generation on the**
 1637 **Minecraft dataset.**

| Model | Steps | FVD \downarrow |
|--------|-------|------------------|
| FAR-B | 100K | 42.710 |
| | 150K | 33.873 |
| MiMo-B | 100K | 33.829 |

1644
 1645 results confirm that a superior understanding of the past, enforced by MiMo, leads to more plausible
 1646 and coherent long-term video generation.

1647
 1648 The results in Table 12 show that MiMo signifi-
 1649 cantly outperforms the baseline. At 100K train-
 1650 ing steps, MiMo achieves an rFVD of 33.829,
 1651 a 27% improvement over the baseline’s 42.710.
 1652 This performance gap highlights MiMo’s abil-
 1653 ity to maintain long-term coherence. Fur-
 1654 thermore, MiMo accelerates training, reaching
 1655 this performance 1.5x faster than the base-
 1656 line, which requires an additional 50K steps to
 1657 achieve a comparable rFVD. Figure 13 shows
 1658 uncurated samples of generated results. These
 1659 results confirm that a superior understanding of the past, enforced by MiMo, leads to more plausible
 1660 and coherent long-term video generation.



1661
 1662 **Figure 13: Uncurated samples of Minecraft long-horizon video generation. The upper row of each**
 1663 **video is GT, the lower row is the generated sample. Red bounding boxes indicate conditioning**
 1664 **frames.**

1674 I LIMITATIONS AND FUTURE WORK 1675

1676 In this work, we analyze the impact of the DiT’s internal representations of history frames on
1677 VideoAR. Based on our findings, we propose MiMo to improve history representations *without*
1678 utilizing VFM.

1679 However, it remains an open question to improve future frame representations with VFM. Masked
1680 DiT (Gao et al., 2023; Wei et al., 2023) achieves success to some extent but requires elaborate
1681 architecture modifications. Some recent approaches (Jiang et al., 2025; Wang & He, 2025) incor-
1682 porate methodologies from self-supervised learning literature, but it is still unclear whether they
1683 (and MiMo) can beat representation alignment approaches (e.g., REPA (Yu et al., 2024)) that utilize
1684 pretrained VFM on in-distribution data of VFM. It is also unclear whether it is possible to pre-
1685 train a generative model that beats VFM in downstream tasks such as video segmentation, video
1686 grounding, etc.

1687 Furthermore, it is an interesting future direction to explore other training objectives to improve
1688 history representations (Oquab et al., 2023; Assran et al., 2023; Jiang et al., 2025; Wang & He,
1689 2025).

1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727