
Loong: Generating Minute-level Long Videos with Autoregressive Language Models

Anonymous Author(s)

Affiliation

Address

email

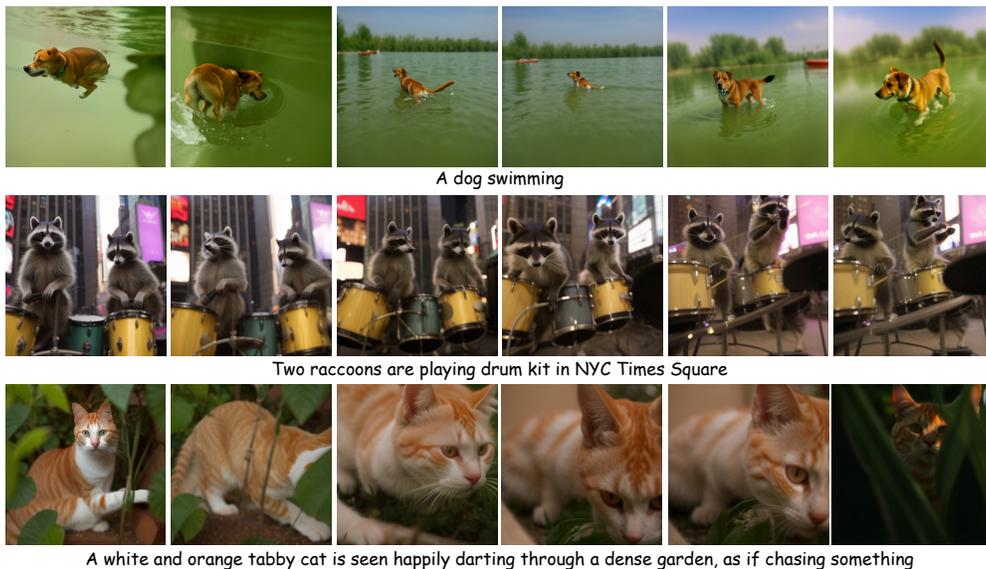


Figure 1: **One-Minute Videos Generated by Loong.** The images are sampled every 10 seconds. Loong is an autoregressive LLM-based model that can generate minute-level long videos with consistent appearance, large motion dynamics, and natural scene transitions.

Abstract

1 It is desirable but challenging to generate content-rich long videos in the scale
2 of minutes. Autoregressive large language models (LLMs) have achieved great
3 success in generating coherent and long sequences of tokens in the domain of
4 natural language processing, while the exploration of autoregressive LLMs for
5 video generation is limited to generating short videos of several seconds. In this
6 work, we conduct a deep analysis of the challenges that prevent autoregressive
7 LLM-based video generators from generating long videos. Based on the obser-
8 vations and analysis, we propose Loong, a pioneering autoregressive LLM-based
9 video generator that can generate minute-long videos. Specifically, we model
10 the text tokens and video tokens as a unified sequence for autoregressive LLMs.
11 We propose progressive short-to-long training with a loss re-weighting scheme to
12 mitigate the loss imbalance problem for long video training. We further investigate
13 inference strategies, including video token re-encoding and sampling strategies,
14 to diminish error accumulation during inference. Our proposed Loong can be
15 trained on 10-second videos and be extended to generate minute-level long videos
16 conditioned on text prompts, as demonstrated by extensive experiments.

17 1 Introduction

18 Over the past few years, video generation models, including diffusion-based ones [1–8] and language
19 model based approaches [9, 10], have shown impressive results in generating short videos of a
20 few seconds. To capture more comprehensive content, it is desirable to generate long videos
21 with consistent appearance, larger motion dynamics, and natural scene transitions. Despite recent
22 works [11–13] to generate long videos with diffusion-based video generators, generating content-rich
23 long videos on the scale of minutes remains largely underexplored and challenging.

24 Autoregressive large language models (LLMs) have shown remarkable success in generating long
25 and coherent text sequences [14–19], demonstrating their ability to capture long-range dependencies
26 and complex temporal patterns. Inspired by the success of autoregressive LLMs in other modalities
27 and their flexibility in unifying various modalities and tasks, recent works [9, 10] have explored
28 autoregressive language models for video generation. Those approaches map videos into discrete
29 tokens and use text tokens as conditioning to generate the video tokens by next-token prediction
30 with decoder-only transformers. State-of-the-art autoregressive LLM-based video generator [10] can
31 generate high-quality 2-second short video clips and iteratively extend to 10-second coherent videos.

32 Despite demonstrating the ability of long sequence generation in NLP and being explored for short
33 video generation, the potential of LLMs to generate minute-level, content-rich, and dynamic videos
34 remains unexplored. In natural language processing, LLMs can be trained on long sequences
35 and extended beyond the training length. However, we empirically observe that either training
36 autoregressive LLMs on long video sequences or extending short video generators to generate long
37 videos leads to unsatisfactory performance for minute-level video generation. A question arises:
38 *What restricts the capability of autoregressive language models for generating long videos?*

39 We hypothesize that the main obstacles are the large redundancy and strong inter-frame dependency
40 among video tokens. The video tokens of the current frame depend heavily on the tokens of the
41 previous frames, leading to two challenges for long video generation: (1) *Imbalanced loss during*
42 *training*. When trained with the next-token prediction objective, predicting early-frame tokens from
43 text prompts is much more difficult than predicting late-frame tokens based on the ground-truth
44 tokens of previous frames. The imbalanced difficulty levels of tokens lead to imbalanced loss during
45 training. The issue becomes more severe as the video length increases, where the accumulated
46 loss of many easy tokens largely surpasses the loss of a few difficult tokens and dominates the
47 gradient direction. (2) *Error accumulation during inference*. While the model predicts the next
48 token conditioned on previous *ground-truth* tokens during training, it has to predict the next token
49 conditioned on previous *predicted* tokens during inference. This training-inference discrepancy leads
50 to error accumulation during inference. Because of the strong inter-frame dependency among video
51 tokens and the large number of video tokens, such error accumulation is non-negligible and causes
52 visual quality degradation for long video inference.

53 In this work, we propose **Loong**, aiming to unleash the power of autoregressive language models to
54 generate long videos in the scale of minutes. Our autoregressive LLM-based video generator consists
55 of two components: a video tokenizer that compresses videos into sequences of discrete video tokens,
56 and an autoregressive LLM that models the unified sequence of text tokens followed by video tokens
57 through next-token prediction. To mitigate the problem of imbalanced loss for long video training,
58 we introduce a progressive short-to-long training strategy that gradually increases the training video
59 length. We further propose loss re-weighting for early frames to prevent the model from being
60 dominated by many easy tokens in the late frames. Moreover, we investigate inference strategies,
61 including the video token re-encoding and sampling strategy, to further extend the video length by
62 iteratively generating the next frames conditioned on previously generated frames. In order to enable
63 training and inference with longer videos, we adopt low-resolution videos for the LLM-based video
64 generator, and leverage a super-resolution and refinement module to further enhance the resolution
65 and fine-grained details of the generated long videos.

66 In summary, we propose Loong, a novel autoregressive LLM-based video generator that can generate
67 content-rich, coherent, and dynamic long videos in the scale of minutes. Based on our observations
68 and analysis of the issues that limit the power of LLMs for long video generation, we propose
69 progressive short-to-long training with a loss weighting scheme to enable model training on 10-
70 second videos. We further investigate inference strategies to extend the 10-second videos to minute-
71 level videos by autoregressive generation strategies designed for long video inference. Our model
72 demonstrates its ability in generating minute-level long videos through extensive experiments.

73 2 Related Work

74 **Video generation.** The mainstream video generation methods can be categorized into GAN-based [20–
75 22], Diffusion-based [23, 7, 24–26, 3, 6, 27, 28, 13] and language-model-based [29, 10, 30, 31].
76 Among them, Diffusion-based methods have recently gained the most popularity. Most Diffusion-
77 based methods encode videos into latent space [32] for efficient training and utilize progressive
78 inference strategies [25, 33, 34] to generate videos with high spatial-temporal resolution. With a new
79 scalable Diffusion Transformer [35] architecture, Sora [13] has further pushed video generation to a
80 new stage. Different from diffusion-based video generation methods, our work aims to explore and
81 unleash the potentiality of language models for long video generation, as their ability for modeling
82 long sequence and scaling up have been proved in NLP.

83 **Image and video generation with language models.** Language models have recently been explored
84 for visual generation, with most works focusing on tokenizing visual data into a form that can be
85 processed by these models. Quantization techniques like VQ-VAE [36, 37] are commonly used, and
86 transformers are employed to model the resulting tokens. For image generation, autoregressive or
87 masked transformers are prevalent [38–42]. In short video generation, image-level or video-level
88 tokenizers are utilized, incorporating spatial-temporal compression and causal structures. Trans-
89 formers model the spatial-temporal relationships, with various techniques proposed, such as sparse
90 attention, spatial-temporal attention, large-scale pre-training, and improved tokenization [9, 43–46].
91 VideoPoet [10] stands out as a multimodal model using bidirectional attention for conditioning, while
92 our method aligns better with the language model paradigm by using unidirectional attention for both
93 text and video. However, these short video generation models focus on producing 1-5 second clips,
94 limiting their ability to capture complex events and maintain consistency over longer durations.

95 **Long video generation.** Previous works have explored long video generation using various ap-
96 proaches. LongVideoGAN [47], NUWA-XL [48], and GAIA-1 [49] utilized GAN-based methods,
97 diffusion-over-diffusion techniques, or world models but were limited to specific domains. More
98 recently, video diffusion models have been extended for longer video generation. FreeNoise [50]
99 and Gen-L [11] focus on sampling noise vectors and aggregating overlapping short video segments,
100 respectively, while StreamingT2V [12] proposes an autoregressive approach with memory blocks
101 for consistency and appearance preservation. In the language model domain, Phenaki [30] generates
102 variable-length videos using a masked video transformer. Despite these advancements, generating
103 long videos with rich motion dynamics, consistent appearance, and high visual quality in the open
104 domain remains a challenge.

105 3 Method

106 We present Loong, an autoregressive LLM-based model for generating long videos in the scale of
107 minutes. We introduce the overall framework, composed of the video tokenizer and the LLM-based
108 video generator, in Sec. 3.1. We analyze the problem with long video training and propose the
109 progressive short-to-long training with loss re-weighting scheme, enabling training on 10-second
110 videos, in Sec. 3.2. We further investigate inference strategies to extend the generated video length to
111 the minute level and post-processing techniques to enhance the spatial resolution of generated videos
112 in Sec. 3.3.

113 3.1 Overall Framework

114 Inspired by previous work in LLM-based image generation and video generation models [38, 41, 46,
115 31, 10], Loong is designed with two components: a video tokenizer that efficiently compresses the
116 videos into discrete tokens, and a decoder-only transformer that autoregressively predicts next video
117 tokens based on text tokens.

118 **Video Tokenizer.** In order to enable spatial-temporal joint compression and joint modeling of images
119 and videos, we leverage causal 3D CNN architecture for the tokenizer, inspired by MAGViT2 [31].
120 The encoded spatial-temporal features are quantized into discrete tokens with Clustering Vector
121 Quantization (CVQ) [51], an improved version of VQGAN [37] designed to enhance codebook
122 utilization. To extend the temporal coverage of videos within a limited number of tokens, we work
123 with low-resolution videos and leave super-resolution for the post-processing in Sec. 3.3. The

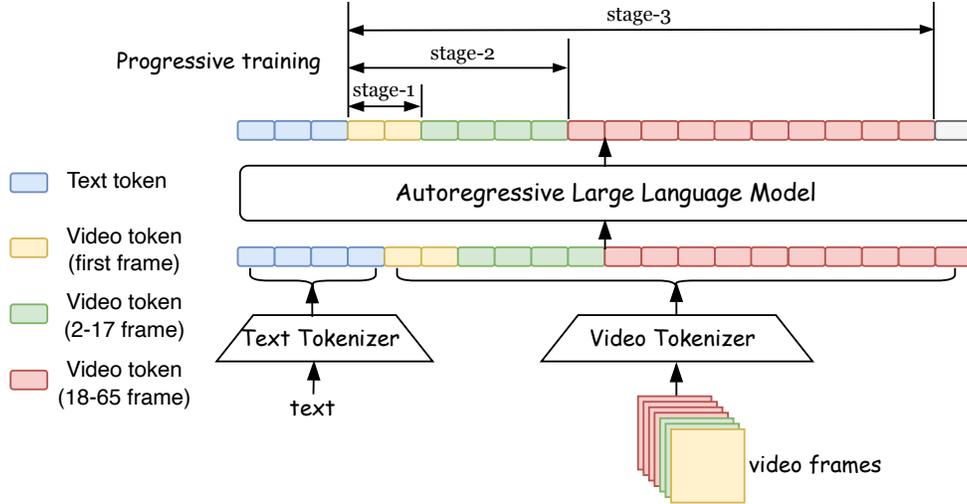


Figure 2: **Overall Framework and the Training process of Loong**. Given the input text tokens, the model predict video tokens autoregressively. All the text and video information is formulated into a unidirectional discrete token sequence, where the model predicts the next token based on the previous tokens. Video Tokenizer is utilized to convert video frames into discrete tokens. We use different color to represent first frame, short clip and long clip separately. We follow a progressive training pipeline to train on long videos.

124 tokenizer can compress a 10-second video (65 frames, 128×128 resolution for each frame) into a
 125 sequence of $17 \times 16 \times 16$ discrete tokens with a vocabulary size of 8192.

126 **Autoregressive LLM-based Video Generation.** With the video frames converted into discrete
 127 tokens, we can now model the text and video tokens as a unified sequence and formulate text-to-video
 128 generation as autoregressively predicting video tokens conditioned on the text tokens with decoder-
 129 only Transformers. The process is illustrated in Fig. 2. For simplicity, we omit the special separate
 130 tokens in the following formulation. Let $\mathbf{t} = \{t_1, t_2, \dots, t_N\}$ represent the sequence of text tokens,
 131 where N is the number of text tokens. Similarly, let $\mathbf{v} = \{v_1, v_2, \dots, v_L\}$ represent the sequence of
 132 video tokens, where L is the number of video tokens. The autoregressive LLM models the unified
 133 token sequence $\mathbf{s} = [\mathbf{t}; \mathbf{v}]$ and is trained with the next-token prediction loss for the video tokens.

$$\mathcal{L} = - \sum_{i=1}^L \log p(v_i | v_{<i}, \mathbf{t}) \quad (1)$$

134 where v_i denotes the i -th token in the video sequence \mathbf{v} , and $v_{<i}$ denotes all the video tokens
 135 preceding v_i .

136 **Discussion.** Different from VideoPoet [10], which encodes text with a pretrained T5 text encoder [52]
 137 and applies bidirectional attention for the input condition tokens and causal attention for the video
 138 tokens, our approach does not rely on a pretrained text encoder. Instead, we formulate the text
 139 tokens and video tokens as a unified token sequence and apply causal attention to all tokens. Our
 140 unified autoregressive modeling of text tokens and video tokens provides a simpler formulation that is
 141 consistent with modern GPT-style LLMs [16]. This design may lead to potential benefits in extending
 142 our model to multimodal LLMs that unify different modalities and different tasks for understanding
 143 and generation.

144 3.2 Progressive Short-to-Long Training with Loss Re-weighting

145 Most video generation models are trained on short video clips, typically no more than 4 seconds,
 146 which limits their ability to capture long-term dependencies and complex dynamics in longer videos.
 147 To address this limitation, it is essential to train these models on videos with longer durations, enabling
 148 them to learn and generate more coherent and contextually rich video content.

149 However, training directly on long videos leads to suboptimal performance, even when the model is
 150 trained for a large number of iterations. We illustrate the loss curve of different frame ranges when

151 training on 65-frame videos (with 4,356 tokens, covering 10 seconds) in Fig. 3. We empirically
 152 observe that tokens from early frames (frames 1-17) have larger losses than those from later frames
 153 (tokens from frames 50-65 have the smallest average loss). During training, the model learns
 154 through next-token prediction, where it is much easier to predict tokens of later frames given the
 155 previous ground-truth video and text tokens. In comparison, predicting early-frame tokens with little
 156 visual cues from previous frames is more challenging. The imbalanced loss is a severe problem for
 157 long-sequence training because the accumulated loss of the many easy-to-predict tokens from later
 158 frames (18-65) surpasses the loss of the few difficult-to-predict tokens from early frames (1-17) and
 159 dominates the gradient direction, leading to suboptimal visual quality in the generated videos.

160 To mitigate the aforementioned challenge of imbal-
 161 anced video token difficulties, we propose a pro-
 162 gressive short-to-long training strategy with loss re-
 163 weighting, demonstrated in the following.

164 **Progressive short-to-long training.** In order to allow
 165 the model to first learn the text-conditioned appear-
 166 ance and motion of short videos, and then smoothly
 167 adjust to longer-range dependencies and more com-
 168 plex motion patterns in longer videos, we factorize
 169 training into three stages which gradually increases
 170 the training video length, as illustrated in the Fig. 2:
 171 (1) In *stage-1*, we pretrain the model with text-to-
 172 image generation on a large dataset of static images,
 173 which helps the model to establish a strong founda-
 174 tion for modeling per-frame appearance and structure.
 175 (2) In *stage-2*, we continue to train the model jointly
 176 on images and short video clips of 17 frames, where
 177 the model learns to capture short-term temporal de-
 178 pendencies and motion patterns while preserving the
 179 per-frame visual quality. (3) In *stage-3*, we increase
 180 the number of video frames to 65, covering a tempo-
 181 ral range of 10 seconds, and continue joint training.

182 **Loss re-weighting for early frames.** To further
 183 strengthen the supervision of early frames and to
 184 prevent the model from forgetting the stage-1 and stage-2 priors, we propose a loss re-weighting
 185 scheme for stage-3. To be specific, we apply larger loss weights for the tokens of early frames, and
 186 the overall weighted loss is formulated as

$$\mathcal{L}_{\text{weighted}} = -(1 + \lambda) \sum_{i=1}^K \log p(v_i | v_{<i}, \mathbf{t}) - \sum_{i=K+1}^L \log p(v_i | v_{<i}, \mathbf{t}), \quad (2)$$

187 where the first term denotes the loss for the K tokens corresponding to the early frames (the first 17
 188 frames), and the second term denotes the loss for the $L - K$ tokens corresponding to the later frames
 189 (frames 18-65). λ is a positive value to strengthen the loss weight of early frames.

190 With the loss weighting and progressive training strategy, our model effectively mitigates the issues
 191 of long video training. As the model is trained on a temporal range of 10 seconds, it can generate
 192 videos of up to 10 seconds with improved temporal coherence and consistency while maintaining the
 193 strong appearance and motion priors learned from the image and short video clips.

194 3.3 Inference Strategies for Extending Video Length and Resolution

195 Large language models are proven to be length-generalizable, so we expect the LLM-based video
 196 generator trained on 10-second videos to be extended to generate longer videos autoregressively.
 197 However, generalizing beyond the training video duration is non-trivial and may lead to error
 198 accumulation and quality degradation. For instance, a one-minute video corresponds to approximately
 199 26, 112 video tokens under our current settings, which is significantly longer than most text sequences
 200 typically encountered in language modeling tasks. The considerable length and the large inter-
 201 frame dependency among video tokens pose challenges for extending the LLM-based generator for

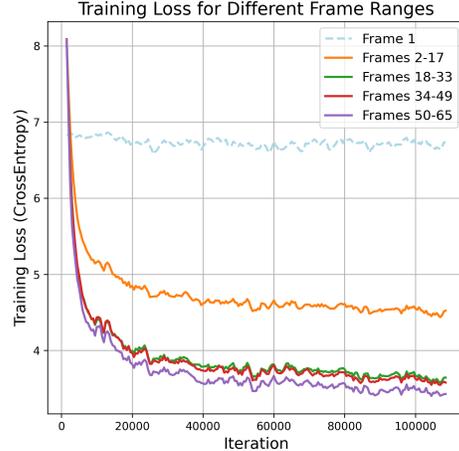


Figure 3: **Imbalanced Training Losses When Training Directly on Long Videos.** The training loss for late frames (18-65) is smaller than that of early frames (1-17), and the loss for the first frame remains relatively high, leading to suboptimal visual quality in the early frames (despite the model being pre-trained on text-to-image).

202 long video generation. In this subsection, we investigate inference strategies to generate minute-
203 level videos and post-processing methods like video super-resolution and refinement to generate
204 higher-quality videos.

205 **Video token re-encoding.** A natural way of extending videos beyond the training duration is
206 to iteratively generate the tokens of the next video clip, conditioned on the text prompts and the
207 previously generated tokens of the current video clip, exploiting the benefit of autoregressive language
208 models. However, this strategy leads to severe video quality degradation for video frames beyond
209 the training range. With further analysis, we find that this issue stems from the token misalignment
210 caused by the causal video tokenizer. To be specific, the tokens from the last n frames in a video clip
211 are derived based on the context of all previous frames, while the tokens from the first n frames in
212 a new video clip are derived without the context of the previous video clip. Therefore, generating
213 tokens for the new clip directly conditioned on previous tokens leads to distribution shift in the
214 input features for LLMs. To address this issue, we decode the LLM-generated video tokens to the
215 pixel-space videos and then re-encode the last n frames with the video tokenizer. The re-encoded
216 video tokens and the text tokens serve as the conditions to generate the tokens of the next video clip.

217 **Sampling strategy.** Decoding video tokens with autoregressive language models is prone to *error*
218 *accumulation* because of the autoregressive nature of the model and the strong inter-frame dependen-
219 cies of video tokens. Errors in predicting one token can propagate and influence the generation of
220 subsequent tokens, leading to a degradation in video quality as the length increases. To mitigate this
221 issue, we draw inspiration from the Top- k sampling strategy commonly used in NLP tasks. During
222 the token sampling process, we only sample from the Top- k most probable tokens, ensuring that the
223 generated tokens are of high quality. By focusing on the most likely tokens, we reduce the influence
224 of potential errors on subsequent token generation, effectively alleviating the error accumulation
225 problem. On the other hand, we also observe that too small values of k ($k = 1$ degrades to greedy
226 decoding) lead to almost static videos with little motion. To balance dynamic motion and error
227 accumulation, we choose $k = 50$ for our model.

228 **Super-resolution and refinement.** As introduced in Sec. 3.1, our video tokenizer and LLM-based
229 video generator operates on the low-resolution 128×128 videos. This design trades off spatial
230 resolution for longer video sequences during training and inference. We apply off-the-shelf super-
231 resolution and refinement models [53–56] on the LLM-generated low-resolution videos. This
232 module serves as a post-processing to enhance the spatial resolution and fine-grained visual details of
233 videos, without affecting the content and motion of the generated videos.

234 4 Experiments

235 4.1 Implementation Details

236 **Model Architecture.** Our video generation model follows the same architecture as LLaMA [18], with
237 a largest size of 7B parameters. We train the models from scratch, without using any text-pretrained
238 weights. The vocabulary consists of 32,000 tokens for text, 8,192 tokens for video, and 10 special
239 tokens, resulting in a total vocabulary size of `vocab_size = 40,202`. For the video tokenizer, we
240 attempt to reproduce the architecture of MAGViT2 [31], utilizing the Clustering Vector Quantization
241 (CVQ) [51] method for quantization. The model compresses the spatial dimensions (width and
242 height) by a factor of 8 and the temporal dimension by a factor of 4.

243 **Training.** Our models are trained on a combination of the CC3M [57] and LAION-2B [58] image
244 datasets, as well as the WebVid-10M [59] video training set and 5.5M self-collected video clips. The
245 training process follows the progressive strategy described in Sec. 3.2. We first pre-train the model on
246 the combined image dataset for 200k iterations, followed by joint training on images and 17-frame
247 video clips from the combined video dataset for another 200k iterations with a batch size of 512. We
248 then jointly train on 65 frames (covering 10 seconds) for 100k iterations with a batch size of 256.
249 The λ is set to 1.0 for the weighted loss of Eq. (2). In each stage, we use AdamW optimizer with a
250 base learning rate of $1.0e-4$. The learning rate is scheduled using a linear warmup for the first 10,000
251 iterations, followed by a cosine annealing decay until reaching the maximum iteration count. For
252 the training of the tokenizer, we also use a progressive approach on the same dataset, increasing the
253 video length from 1 to 17 to 65 frames while maintaining a resolution of 128×128 , with a batch size
254 of 64 and training for 400k iterations.



Figure 4: **Effectiveness of the Progressive Training with Loss Re-weighting.** We sample 4 frames from the 17 earlier frames of the video generation results, to show the performance of models trained with or without our training strategy. The top row shows results of the model trained directly on long video, the appearance of objects degrades largely. The bottom row shows the results model trained with our proposed training approach, the appearance preserves effectively.



Figure 5: **Effectiveness of Token Re-encoding during Video Extension.** For each sample, the left two images show the results before the extension process, and the right two images show the results after extension. Without token re-encoding, the extension fails to generate visually consistent content.

255 4.2 Ablation Study

256 In this section, we conduct ablation studies to evaluate the effectiveness of our main design choices.
 257 Unless otherwise specified, we use the 3B model with an output spatial resolution of 128×128 ,
 258 without any super-resolution and refinement module. To reduce computational cost, we train the
 259 models for half the number of iterations compared to the full setting described in Sec. 4.1. Due to the
 260 lack of a general long video generation benchmark, we build a custom one by selecting the top-1000
 261 longest clips from the WebVid [59] validation set and slicing each to 27 seconds, the duration of
 262 the shortest among them. We employ two performance metrics on this benchmark: Fréchet Video
 263 Distance (FVD)[60] and Video-Text Matching (VTM) score calculated using CLIP (ViT-L/14)[61].
 264 It is worth noting that these metrics serve as references, and human evaluation should be considered a
 265 more accurate quality assessment of the generated videos. We use the text prompt sets from prior
 266 works [4, 6, 62, 3, 13] to generate videos for visualization.

267 **Model Scaling.** Scalability is an important characteristic of LLMs. To study scaling behavior of our
 268 model, we evaluate performance of the models with different sizes. Tab. 1 presents the quantitative
 269 results of our models with 700M, 3B and 7B parameters on the custom benchmark. We observe that
 270 larger models achieve better FVD and VTM scores, demonstrating the scalability of our approach.

271 **Progressive Training with Loss Re-weighting.** To validate
 272 the effectiveness of our proposed training strategy, we compare
 273 our model with models directly trained on long videos.
 274 Both models are pretrained on images and then trained on long
 275 videos using different strategies. Fig. 4 compares the gener-
 276 ated frames from a single generation stage without extension
 277 for each model. It is clear that the videos generated by the
 278 directly-trained models suffer from significant object appear-
 279 ance degradation, losing much of the structural information.
 280 In contrast, our model, trained with the progressive training
 281 approach, effectively preserves the appearance details.

Table 1: **Scalability of Loong.** The performance improves as the model size increases.

	FVD _{I3D} ↓	VTM _c ↑
700M	633	21.5
3B	572	22.8
7B	432	24.1

282 **Video Token Re-encoding.** Fig. 5 illustrates the importance of token re-encoding during the video
 283 extension process. Without proper token re-encoding, the model fails to maintain visual consistency
 284 when extending the video, resulting in abrupt changes in appearance and content. In contrast, by
 285 employing our token re-encoding technique, the extended frames seamlessly continue the video with
 286 coherent visual style and content.

287 **Sampling Strategy for Inference.** We compare three sampling strategies when predicting each token:
 288 greedy decoding ($k = 1$), top- k sampling, and multinomial sampling from the whole vocabulary (k
 289 equals video token vocabulary size). As shown in Fig. 6, greedy decoding generates stable results but
 290 lacks diversity, while multinomial sampling produces more dynamic content at the cost of quality.
 291 Top- k sampling ($k = 50$) balances stability and diversity. A smaller k value prioritizes stability,
 292 resulting in less diverse motion, while a larger k allows for more dynamic and varied content at the
 293 risk of introducing instability. In the process of video extension, selecting an appropriate k value is
 294 crucial for maintaining consistency and mitigating error accumulation over longer sequences.

295 4.3 Comparison to State-of-the-Art Methods

Table 2: Comparison on zero-shot text-to-short-video benchmarks.

Model	CogVideo[45]	MagicVideo[7]	ModelScopeT2V[63]	Show-1[28]	VideoPoet[10]	Loong
CLIPSIM	0.2631	-	0.2930	0.3072	0.3049	0.2903
FVD	1294	998	550	538	213	274

296 **Zero-shot Text to Short Video Generation.** Al-
 297 though our approach is not specifically designed
 298 for short video generation, we compare our per-
 299 formance on the MSR-VTT dataset [64] using
 300 CLIP similarity (CLIPSIM) [44] and FVD [60]
 301 metrics, evaluated on 16 frames. As shown
 302 in Tab. 2, our FVD score is the second-best,
 303 only slightly behind VideoPoet [10] (pretrained).
 304 However, our CLIPSIM score is lower compared
 305 to some other methods, which can be attributed
 306 to the fact that our approach is trained from
 307 scratch without utilizing any pre-trained text
 308 weights. In contrast, methods with higher CLIP-
 309 SIM scores, such as VideoPoet, leverage pre-
 310 trained language models like T5 [52] for text
 311 encoding, while diffusion-based methods often
 312 employ CLIP [61] text embeddings, which are
 313 already trained on the CLIP dataset. Despite not
 314 using pre-trained text models, our method still
 315 achieves competitive performance, demonstrat-
 316 ing its effectiveness in capturing the semantic
 317 relationship between text and video.

318 **User Study on Long Video Generation.** We
 319 conduct a user study to compare our method
 320 with StreamingT2V [12], a state-of-the-art open-
 321 sourced long video generation method built on
 322 Stable Video Diffusion [26]. We use 50 text
 323 prompts from prior works [4, 6, 62, 3] to generate 1-min videos. In the study, users are presented
 324 with 2 videos generated by the two models, conditioned on the same text. They are asked to choose
 325 the preferred video based on visual text matching and content consistency. The videos are presented
 326 randomly, and users are not informed about the models. We collect 440 responses. As shown in
 327 Fig. 8, our model outperforms StreamingT2V in both content consistency (win rate 0.83 vs. 0.125)
 328 and visual text matching (win rate 0.65 vs. 0.19).

329 4.4 Visualization Results

330 In this section, we show the results of our model under different text-to-video generation scenarios.



Figure 6: **Study on Sampling Strategies.** Results of three different inference sampling strategies. Greedy decoding produces stable results but lacks diversity between frames. Multinomial sampling generates more dynamic and diverse content but with lower quality. Top- k sampling achieves a balance between stability and diversity. k is set to 50 in this experiment.

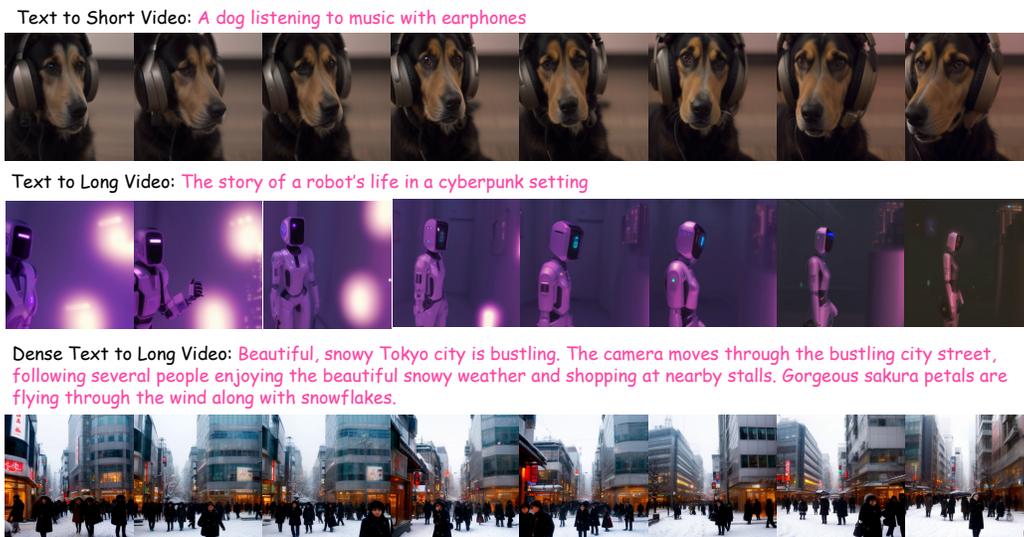


Figure 7: **Generated videos from Loong across various text-to-video scenarios.**

331 **Text to Short Video.** In the top row of the figure, we show results of short video generation. As
 332 shown in the figure, our approach exhibits the capability to generate short videos with rich details
 333 and high fidelity while maintaining strong alignment with the given text descriptions.

334 **Text to Long Video.** The second row shows frames sampled from a long video generated by our
 335 model, conditioned on a concise text description. These examples demonstrate that our approach can
 336 generate long videos containing diverse content and larger dynamic changes compared to short video
 337 generation, while maintaining semantic alignment with the given text.

338 **Dense Text to Long Video.** Although not explicitly trained on dense captions, we found that
 339 our model can effectively adapt to long video generation in a zero-shot manner. Our
 340 model demonstrates the capability to generalize to dense caption conditions without requiring
 341 specialized finetuning. As illustrated in the last row of Fig. 7, The generated long videos
 342 maintain semantic alignment with the provided dense captions, showcasing rich content that cor-
 343 responds to the detailed descriptions, including multiple characters, weather, scenery, and build-
 344 ing information. However, we observe that the generated images appear slightly blurry. We attribute this to the low resolution of our transformer’s
 345 output, which may result in blurriness when generating highly detailed content. We hope that future
 346 advancements in tokenizer compression capabilities and longer LLMs could help address this issue.

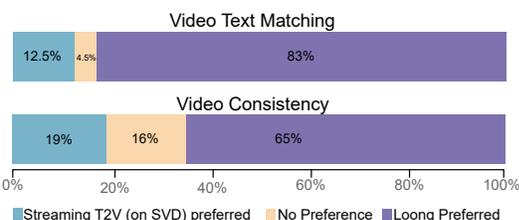


Figure 8: **User Study on 1-min videos.** Comparison with the StreamingT2V on SVD model. Our model is more preferred by human raters in terms of both visual text match and content consistency.

354 5 Conclusion and Discussions

355 We propose Loong, the pioneering autoregressive LLM-based video generation model that can
 356 generate minute-level long videos with consistent appearance, large motion dynamics, and natural
 357 scene transitions. We overcome the challenges of long video training with the progressive short-
 358 to-long training scheme with loss re-weighting. We also investigate inference strategies to extend
 359 generated videos beyond training duration. Our experiments demonstrate the effectiveness of our
 360 approach in generating minute-level long videos. We discuss the limitations in the Appendix.

361 **Border impact.** The model can be deployed to assist visual artists and film producers on video
 362 creation, enhancing their efficiency. It can also be deployed for entertainment purposes. On the
 363 other hand, it may be used for generating fake content and delivering misleading information. The

364 community should be aware of the potential social impacts. It is necessary to develop techniques to
365 detect and watermark the videos generated by machine learning models.

366 **References**

- 367 [1] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J
368 Fleet. Video diffusion models. *arXiv preprint arXiv:2204.03458*, 2022.
- 369 [2] Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis
370 Germanidis. Structure and content-guided video synthesis with diffusion models. In *Proc. IEEE*
371 *Int. Conf. Comp. Vis.*, pages 7346–7356, 2023.
- 372 [3] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiuyan Hu,
373 Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without
374 text-video data. In *Proc. Int. Conf. Learn. Representations*, 2022.
- 375 [4] Omer Bar-Tal, Hila Chefer, Omer Tov, Charles Herrmann, Roni Paiss, Shiran Zada, Ariel Ephrat,
376 Junhwa Hur, Yuanzhen Li, Tomer Michaeli, et al. Lumiere: A space-time diffusion model for
377 video generation. *arXiv preprint arXiv:2401.12945*, 2024.
- 378 [5] Yan Zeng, Guoqiang Wei, Jiani Zheng, Jiaxin Zou, Yang Wei, Yuchen Zhang, and Hang Li.
379 Make pixels dance: High-dynamic video generation. *arXiv:2311.10982*, 2023.
- 380 [6] Rohit Girdhar, Mannat Singh, Andrew Brown, Quentin Duval, Samaneh Azadi, Sai Saketh
381 Rambhatla, Akbar Shah, Xi Yin, Devi Parikh, and Ishan Misra. Emu video: Factorizing
382 text-to-video generation by explicit image conditioning. *arXiv preprint arXiv:2311.10709*,
383 2023.
- 384 [7] Daquan Zhou, Weimin Wang, Hanshu Yan, Weiwei Lv, Yizhe Zhu, and Jiashi Feng. Magicvideo:
385 Efficient video generation with latent diffusion models. *arXiv preprint arXiv:2211.11018*, 2022.
- 386 [8] Weimin Wang, Jiawei Liu, Zhijie Lin, Jiangqiao Yan, Shuo Chen, Chetwin Low, Tuyen Hoang,
387 Jie Wu, Jun Hao Liew, Hanshu Yan, et al. Magicvideo-v2: Multi-stage high-aesthetic video
388 generation. *arXiv preprint arXiv:2401.04468*, 2024.
- 389 [9] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. Videogpt: Video generation
390 using vq-vae and transformers. *arXiv preprint arXiv:2104.10157*, 2021.
- 391 [10] Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Rachel Hornung, Hartwig
392 Adam, Hassan Akbari, Yair Alon, Vighnesh Birodkar, et al. Videopoet: A large language model
393 for zero-shot video generation. *arXiv preprint arXiv:2312.14125*, 2023.
- 394 [11] Fu-Yun Wang, Wenshuo Chen, Guanglu Song, Han-Jia Ye, Yu Liu, and Hongsheng Li.
395 Gen-l-video: Multi-text to long video generation via temporal co-denoising. *arXiv preprint*
396 *arXiv:2305.18264*, 2023.
- 397 [12] Roberto Henschel, Levon Khachatryan, Daniil Hayrapetyan, Hayk Poghosyan, Vahram Tade-
398 vosyan, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Streamingt2v: Consistent,
399 dynamic, and extendable long video generation from text. *arXiv preprint arXiv:2403.14773*,
400 2024.
- 401 [13] OpenAI. Sora: Creating video from text. <https://openai.com/sora>, 2024.
- 402 [14] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language
403 understanding by generative pre-training. 2018.
- 404 [15] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language
405 models are unsupervised multitask learners. 2019.
- 406 [16] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
407 Arvind Neelakantan, Pranav Shyam, Sastry, et al. Language models are few-shot learners. In H.
408 Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Proc. Advances in Neural*
409 *Inf. Process. Syst.*, 2020.
- 410 [17] Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-
411 baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al.
412 Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv*
413 *preprint arXiv:2403.05530*, 2024.
- 414 [18] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timo-
415 thée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open
416 and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- 417 [19] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei,
418 Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open
419 foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- 420 [20] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,
421 Aaron Courville, and Yoshua Bengio. Generative adversarial networks. In *Proc. Advances in*
422 *Neural Inf. Process. Syst.*, volume 27, 2014.

- 423 [21] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics.
424 In *Proc. Advances in Neural Inf. Process. Syst.*, 2016.
- 425 [22] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing
426 motion and content for video generation. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2018.
- 427 [23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proc.*
428 *Advances in Neural Inf. Process. Syst.*, 2020.
- 429 [24] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and
430 David J. Fleet. Video diffusion models. *arXiv preprint arXiv:2204.03458*, 2022.
- 431 [25] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko,
432 Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High
433 definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022.
- 434 [26] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Do-
435 minik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion:
436 Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.
- 437 [27] Xin Li, Wenqing Chu, Ye Wu, Weihang Yuan, Fanglong Liu, Qi Zhang, Fu Li, Haocheng Feng,
438 Errui Ding, and Jingdong Wang. Videogen: A reference-guided latent diffusion approach for
439 high definition text-to-video generation. *arXiv preprint arXiv:2309.00398*, 2023.
- 440 [28] David Junhao Zhang, Jay Zhangjie Wu, Jia-Wei Liu, Rui Zhao, Lingmin Ran, Yuchao Gu,
441 Difei Gao, and Mike Zheng Shou. Show-1: Marrying pixel and latent diffusion models for
442 text-to-video generation. *arXiv preprint arXiv:2309.15818*, 2023.
- 443 [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
444 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proc. Advances in Neural Inf.*
445 *Process. Syst.*, 2017.
- 446 [30] Ruben Villegas, Mohammad Babaeizadeh, Pieter-Jan Kindermans, Hernan Moraldo, Han Zhang,
447 Mohammad Taghi Saffar, Santiago Castro, Julius Kunze, and Dumitru Erhan. Phenaki: Variable
448 length video generation from open domain textual descriptions. In *Proc. Int. Conf. Learn.*
449 *Representations*, 2022.
- 450 [31] Lijun Yu, José Lezama, Nitesh B Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen,
451 Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G Hauptmann, et al. Language model beats
452 diffusion-tokenizer is key to visual generation. In *Proc. Int. Conf. Learn. Representations*,
453 2024.
- 454 [32] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
455 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF*
456 *conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- 457 [33] Yingqing He, Tianyu Yang, Yong Zhang, Ying Shan, and Qifeng Chen. Latent video diffusion
458 models for high-fidelity video generation with arbitrary lengths. *arXiv:2211.13221*, 2022.
- 459 [34] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja
460 Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent
461 diffusion models. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 22563–22575, 2023.
- 462 [35] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proc. IEEE*
463 *Int. Conf. Comp. Vis.*, pages 4195–4205, 2023.
- 464 [36] Aaron van den Oord, Oriol Vinyals, and koray kavukcuoglu. Neural discrete representation
465 learning. In *Proc. Advances in Neural Inf. Process. Syst.*, 2017.
- 466 [37] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution
467 image synthesis. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 12873–12883, 2021.
- 468 [38] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark
469 Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *Proc. Int. Conf. Mach. Learn.*,
470 pages 8821–8831, 2021.
- 471 [39] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay
472 Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive
473 models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2022.
- 474 [40] Huiwen Chang, Han Zhang, Jarred Barber, Aaron Maschinot, Jose Lezama, Lu Jiang, Ming-
475 Hsuan Yang, Kevin Patrick Murphy, William T. Freeman, Michael Rubinstein, Yuanzhen Li,
476 and Dilip Krishnan. Muse: Text-to-image generation via masked generative transformers. In
477 *Proc. Int. Conf. Mach. Learn.*, pages 4055–4075, 2023.
- 478 [41] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T. Freeman. Maskgit: Masked
479 generative image transformer. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, June 2022.

- 480 [42] Lijun Yu, Yong Cheng, Zhiruo Wang, Vivek Kumar, Wolfgang Macherey, Yanping Huang,
481 David Ross, Irfan Essa, Yonatan Bisk, Ming-Hsuan Yang, et al. Spae: Semantic pyramid
482 autoencoder for multimodal generation with frozen llms. In *Proc. Advances in Neural Inf.*
483 *Process. Syst.*, 2024.
- 484 [43] Songwei Ge, Thomas Hayes, Harry Yang, Xi Yin, Guan Pang, David Jacobs, Jia-Bin Huang, and
485 Devi Parikh. Long video generation with time-agnostic vqgan and time-sensitive transformer.
486 In *Proc. Eur. Conf. Comp. Vis.*, pages 102–118, 2022.
- 487 [44] Chenfei Wu, Lun Huang, Qianxi Zhang, Binyang Li, Lei Ji, Fan Yang, Guillermo Sapiro,
488 and Nan Duan. Godiva: Generating open-domain videos from natural descriptions.
489 *arXiv:2104.14806*, 2021.
- 490 [45] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-
491 scale pretraining for text-to-video generation via transformers. In *Proc. Int. Conf. Learn.*
492 *Representations*, 2022.
- 493 [46] Lijun Yu, Yong Cheng, Kihyuk Sohn, José Lezama, Han Zhang, Huiwen Chang, Alexander G
494 Hauptmann, Ming-Hsuan Yang, Yuan Hao, Irfan Essa, et al. Magvit: Masked generative video
495 transformer. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2023.
- 496 [47] Tim Brooks, Janne Hellsten, Miika Aittala, Ting-Chun Wang, Timo Aila, Jaakko Lehtinen,
497 Ming-Yu Liu, Alexei A Efros, and Tero Karras. Generating long videos of dynamic scenes. In
498 *Proc. Advances in Neural Inf. Process. Syst.*, 2022.
- 499 [48] Shengming Yin, Chenfei Wu, Huan Yang, Jianfeng Wang, Xiaodong Wang, Minheng Ni,
500 Zhengyuan Yang, Linjie Li, Shuguang Liu, Fan Yang, et al. Nuwa-xl: Diffusion over diffusion
501 for extremely long video generation. *arXiv preprint arXiv:2303.12346*, 2023.
- 502 [49] Anthony Hu, Lloyd Russell, Hudson Yeo, Zak Murez, George Fedoseev, Alex Kendall, Jamie
503 Shotton, and Gianluca Corrado. Gaia-1: A generative world model for autonomous driving.
504 *arXiv preprint arXiv:2309.17080*, 2023.
- 505 [50] Haonan Qiu, Menghan Xia, Yong Zhang, Yingqing He, Xintao Wang, Ying Shan, and Ziwei
506 Liu. Freenoise: Tuning-free longer video diffusion via noise rescheduling. In *Proc. Int. Conf.*
507 *Learn. Representations*, 2024.
- 508 [51] Chuanxia Zheng and Andrea Vedaldi. Online clustered codebook. In *Proc. IEEE Int. Conf.*
509 *Comp. Vis.*, 2023.
- 510 [52] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena,
511 Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified
512 text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020.
- 513 [53] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
514 resolution image synthesis with latent diffusion models. In *Proc. IEEE Conf. Comp. Vis. Patt.*
515 *Recogn.*, pages 10684–10695, 2022.
- 516 [54] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh
517 Agrawala, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image
518 diffusion models without specific tuning. In *Proc. Int. Conf. Learn. Representations*, 2023.
- 519 [55] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image
520 diffusion models. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 3836–3847, 2023.
- 521 [56] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion
522 for editing real images using guided diffusion models. In *Proceedings of the IEEE/CVF*
523 *Conference on Computer Vision and Pattern Recognition*, pages 6038–6047, 2023.
- 524 [57] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A
525 cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of*
526 *the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long*
527 *Papers)*, 2018.
- 528 [58] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman,
529 Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-
530 5b: An open large-scale dataset for training next generation image-text models. In *Proc.*
531 *Advances in Neural Inf. Process. Syst.*, volume 35, 2022.
- 532 [59] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video
533 and image encoder for end-to-end retrieval. In *Proc. IEEE Int. Conf. Comp. Vis.*, 2021.
- 534 [60] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski,
535 and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges.
536 *arXiv preprint arXiv:1812.01717*, 2018.

- 537 [61] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
538 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
539 models from natural language supervision. In *Proc. Int. Conf. Mach. Learn.*, pages 8748–8763.
540 PMLR, 2021.
- 541 [62] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko,
542 Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High
543 definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022.
- 544 [63] Jiuniu Wang, Hangjie Yuan, Dayou Chen, Yingya Zhang, Xiang Wang, and Shiwei Zhang.
545 Modelscope text-to-video technical report. *arXiv preprint arXiv:2308.06571*, 2023.
- 546 [64] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for
547 bridging video and language. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 5288–5296,
548 2016.