

# AUTOREGRESSIVE TRANSFORMERS ARE ZERO-SHOT VIDEO IMITATORS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

People interact with the real-world largely dependent on visual signal, which are ubiquitous and illustrate detailed demonstrations. In this paper, we explore utilizing visual signals as a new interface for models to interact with the environment. Specifically, we choose videos as a representative visual signal. And by training autoregressive Transformers on video datasets in a self-supervised objective, we find that the model emerges a zero-shot capability to infer the semantics from a demonstration video, and imitate the semantics to an unseen scenario. This allows the models to perform unseen tasks by watching the demonstration video in an in-context manner, without further fine-tuning. To validate the imitation capacity, we design various evaluation metrics including both objective and subjective measures. The results show that our models can generate high-quality video clips that accurately align with the semantic guidance provided by the demonstration videos, and we also show that the imitation capacity follows the scaling law.<sup>1</sup>



Figure 1: Illustrations of VidIT. Given a query video clip, the model generates different results based on the provided demonstration video clips. The generated results are coherent with the query in content and semantically consistent with the demonstrations. Text descriptions are annotated for clarity and not used as model input.

<sup>1</sup>Anonymous Page: <https://anonymous.4open.science/w/VidImit-page-656D/>

# 1 INTRODUCTION

Humans acquire skills through imitation, a fundamental aspect of learning demonstrated even in early childhood. For instance, when children observe how to hold a spoon, they not only learn this specific action but can also generalize it to grasp different spoons in varied contexts. These demonstrations are primarily conveyed through visual signals, which are abundant in our environment and provide intricate guidance. Similarly, for the development of general AI agents, the ability to learn from visual demonstrations and generalize to new situations is crucial and desirable. This capability to understand and imitate visual signals is not only appealing but also essential for advancing AI towards more human-like learning and interaction.

Previous works (Bar et al., 2022; Wang et al., 2023a;b; Bai et al., 2023) have studied *image in-context learning* in image-based visual tasks. By formulating image perception tasks such as segmentation and detection to supervised image pair demonstrations in delicately designed structures, e.g., a joint grid image (Bar et al., 2022) for MAE models (He et al., 2022) or a sequence (Bai et al., 2023) for Transformer models (Vaswani et al., 2017), they encourage the model to mimic the task and generate the prediction of a given image query. However, these works require explicitly formulating the training samples in source-target pairs, which restricts in-context inference to a supervised capacity. Furthermore, image-level generation struggles to produce coherent and contextually appropriate sequences.

In this paper, we introduce the **Video ImiTator (VidIT)** model, an autoregressive Transformer (Vaswani et al., 2017; Touvron et al., 2023) trained with the next token prediction objective on video datasets. VidIT uses videos as the primary unit for both inputs and outputs, with each training sample consisting of a sequence of frames from a single video, encoded as discrete tokens using a vector-quantized encoder (Van Den Oord et al., 2017; Patil et al., 2024). The model learns to predict future frames based on preceding ones through self-supervised learning. Though demonstration-generation pairs are not provided during training, the autoregressive paradigm enables the model to discern the underlying structures and patterns within the videos, leading to the emergence of a zero-shot imitation capability. Consequently, when the model encounters demonstration videos during inference, it can generalize and apply these learned patterns to generate coherent and contextually appropriate video sequences. This zero-shot capability echoes findings observed in large language models (Brown et al., 2020).

Specifically, demonstration videos are highly versatile and capable of conveying a wide range of information, such as examples for tasks including object manipulation, or movements of the camera in an ego-centric video. Guided by these demonstrations, the model receives a query video clip set in a new scene and generates a subsequent video sequence that mimics the semantic actions from the demonstrations (see Figure 1 for examples). This allows VidIT to address multiple downstream tasks, such as embodied planning and simulating, by letting a query robot imitate the actions demonstrated by humans or other agents. Given that videos excel in describing low-level details (where language may fall short) and temporal dynamics (where images are insufficient), VidIT acts as a crucial interface for models to interact with the real world.

To comprehensively and accurately evaluate the model performance in VidIT, we develop both objective and subjective metrics to assess the generated videos in terms of visual quality, semantic accuracy, and consistency with the prompted demonstrations. Our extensive experiments demonstrate that the model not only produces high-quality video clips but also successfully adheres to the semantic guidance provided by the demonstration examples. In addition, we show that the zero-shot imitation capacity also follows the scaling law (Kaplan et al., 2020) of large models, illustrating the potential of future works. The main contributions of this work are summarized as follows:

- We propose and study the task of video imitation, which enables the model to interact with real-world demonstrations through video generation.
- We train a large Transformer model VidIT that exhibits powerful video imitation learning capacity, which also follows the scaling law of large models.
- We propose various evaluation metrics to evaluate the visual quality and semantic accuracy of generated videos, providing a solid benchmark for the evaluation of video imitation learning.

## 2 RELATED WORK

**Imitation Learning and In-context Learning** Imitation Learning is a crucial paradigm in reinforcement learning, wherein the expected behavior is acquired by mimicking demonstration’s actions (Zare et al., 2024). By learning on state-action pairs from an expert, the model is expected to map the current state to the corresponding actions. Such learning paradigm is close to human behavior and is appealing in interactive circumstances.

This ability to imitate has also been identified as a key feature in large language models (Radford et al., 2019; Brown et al., 2020; Touvron et al., 2023; Chowdhery et al., 2023), which is referred to as *in-context learning* in LLM area. By conditioning on a sequence of text templates, an LLM generates coherent content based on these templates. This paradigm obviates the need for parameter updates and thereby benefiting the downstream usage of LLMs (Brown et al., 2020; Touvron et al., 2023). Different prompting designs (Wei et al., 2022; Guo et al., 2023; Li et al., 2023) empower LLMs to deal with a wide range of natural language understanding and generation tasks.

As for the in-context paradigm in vision area, current research mainly focus on teaching the model to mimic the vision task provided by demonstrations. Pioneering works construct vision imitation as an image inpainting (Bar et al., 2022; Wang et al., 2023a;b) task. Given multiple query-answer image pairs arranged in a grid image, models are optimized to reconstruct the masked answers under the MAE projective (He et al., 2022). Recently, LVM (Bai et al., 2023) flatten image pairs into a sequence and train an auto-regressive Transformer with next token prediction. These models show powerful capacity to imitating vision tasks such as semantics segmentation and detection from the demonstrations, but rely on training on supervised datasets. In addition, when extending the basic elements (i.e., demonstrations, queries and predictions) from images to videos, unfortunately, existing models struggle to handle the increased complexity as they are not designed to capture the spatial-temporal relationships of inputs. Specifically, though LVM (Bai et al., 2023) is able to generate consecutive frames of a query video clip, the video imitation capacity (e.g., generating different consequences with different demonstrations) is not demonstrated.

**Video Generation** The studied video imitation learning can be considered as a conditional video generation problem, which has recently become a prominent focus in research. Text is the most commonly utilized condition, and various text-to-video models have shown promising results in generating high-fidelity videos (Ho et al., 2022b;a; Brooks et al., 2024; Hong et al., 2023; Villegas et al., 2023; Kondratyuk et al., 2023; Yu et al., 2023b). Recently, video generation models are spread to other related domains such as embodied AI (Yang et al., 2024), where they are utilized in visual planning (Du et al., 2023) and simulation (Yang et al., 2023) with actions and observations as conditions, illustrating the potential of this area. As for model architectures, Diffusion models (Ho et al., 2020; Nichol & Dhariwal, 2021) and Transformers (Vaswani et al., 2017) are both widely adopted in different scenarios, and we choose the autoregressive Transformer as the backbone model in this paper, as its abilities to model long-range dependencies and contextual relationships are crucial for video imitators.

## 3 VIDEO IMITATOR

We introduce the proposed Video Imitator in this section. We begin with the problem definition of both image and video imitation in Section 3.1, highlighting the differences between them. We then elaborate on the training and inference pipeline of VidIT in Section 3.2. Dataset selection is discussed in Section 3.3.

### 3.1 PROBLEM DEFINITION

Typically, the image task imitator models follow the setting akin to *question answering*: the demonstration is composed of  $k$  image pairs  $D^I = \{s_i^q, s_i^a\}_{i=1}^k$ , where  $s_i^q \in \mathbb{R}^{(3, n_r, n_r)}$  denotes the input image with resolution  $n_r$  and  $s_i^a$  denotes its target (usually an image with task-specific annotations (Bai et al., 2023)). Given a query image  $x^I$ , the model predicts its target  $y^I$  conditioned on the demonstration:

$$f_{\theta}(x^I) = P(y^I | x^I, D^I), \quad (1)$$

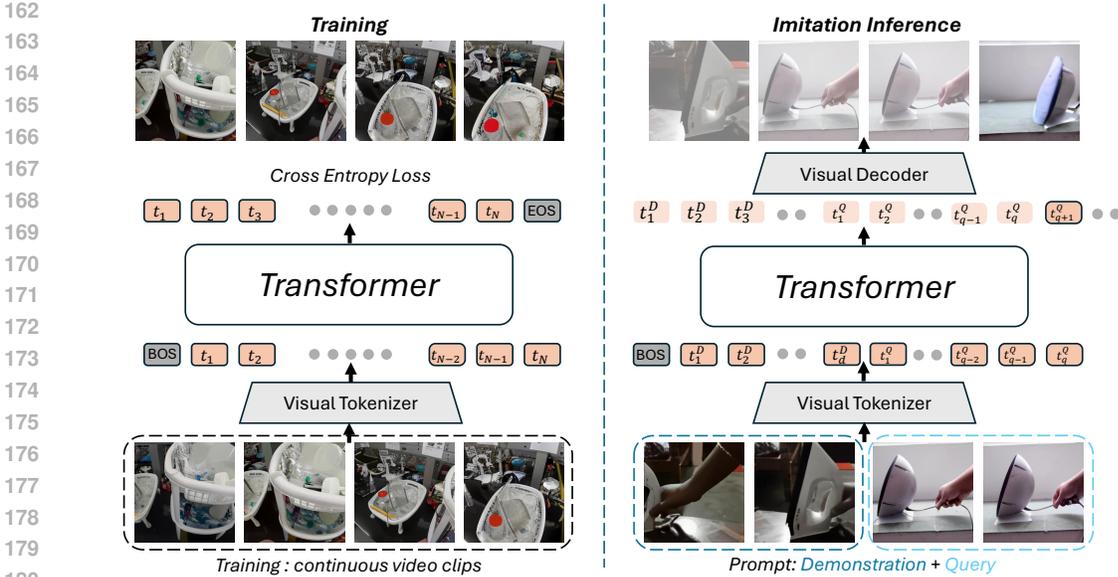


Figure 2: The framework of VidIT. *Left:* Training the VidIT. The data used for training are continuous video clips and the Transformer is trained by next token prediction objective. *Right:* Video Imitation Inference. The model is conditioned on demonstration videos and generates the subsequent frames of a given query video.

where  $f_\theta(\cdot)$  represents the utilized vision model such as an MAE-pretrained ViT (Wang et al., 2023a; Bar et al., 2022) or an auto-regressive Transformer (Bai et al., 2023), which is expected to discern the inherent task from image pairs  $(s_i^d, s_i^q)$  and process  $x$  accordingly.

Unlike image task imitators, VidIT takes videos as basic element, focusing on inheriting semantics from demonstrations rather than discerning specific tasks. Specifically, the demonstration  $D^V = \{(s_i^1, \dots, s_i^{n_i})\}_{i=1}^k$  comprises  $k$  video clips, and each clip consists of  $n_i$  frames. Given a query video clip  $x^V = (s_q^1, \dots, s_q^{n_q})$ , the objective of VidIT can be formulated as:

$$f_\theta(x^V) = P(y^V | x^V, D^V), \quad (2)$$

where  $y^V = (s_y^1, \dots, s_y^{n_y})$  denotes the generated video clip, which should be perceptually coherent with the query  $x^V$  while semantically consistent with the demonstration  $D^V$  at the same time.

### 3.2 APPROACH

Training a Transformer for vision tasks typically consists of two stages, 1) training a visual tokenizer such as VQ-VAE (Van Den Oord et al., 2017; Esser et al., 2021) to convert each image to discrete tokens; 2) each training sample is constructed as a sequence of tokens to train the Transformer decoder. For the first stage, we utilize a public pretrained checkpoint (Patil et al., 2024) with  $16 \times$  spatial compression rate as the VQ tokenizer. Our framework also fits for recent tokenizers with both spatial and temporal compressions (Yu et al., 2023a;b), and we leave it for future work. Formally, for a training video clip  $x^V = (s^1, \dots, s^n)$  with  $n$  frames, the VQ tokenizer converts it to a flat sequence  $x^t = (t_1, \dots, t_{n_t}, t_{n_t+1}, \dots, t_N)$ , where  $n_t$  denotes the number of quantized ids that represents one frame, and  $N = n \cdot n_t$  denotes the total length of the sequence. We append special tokens [bos] and [eos] to the front and end of the sequence, and no special tokens are inserted between sequences of different frames.

Then, in the second stage, we follow the architecture of LLaMA (Touvron et al., 2023) utilizing RMSNorm normalizing (Zhang & Sennrich, 2019) and Rotary Embeddings (Su et al., 2024), and train a Transformer decoder in an autoregressive way:

$$f_\theta(x^t) = \prod_{i=1}^N P(t_i | t_{<i}), \quad (3)$$

where the model predicts each token conditioned on previous ones. Note that in the training stage, each sequence is sampled from one original video, and we do not concatenate video clips from different videos.

**Imitation Inference** The imitation inference format differs from that in the training. Following Equation (2), a number of video clips are selected as the demonstration  $D^V$ , which is appended in front of the query clip  $x_{<j}^t = (t_1, \dots, t_{j-1})$  to let the model generate corresponding responses:

$$f_{\theta}(x_{>j}^t) = \prod_{i=j}^N P(t_i | x_{<i}^t, D^V). \quad (4)$$

The generation results  $x_{\geq j}^t = (t_j, \dots, t_N)$  are vector-quantized IDs, which are then fed to the pretrained VQ decoder to reconstruct as images. We provide an illustration of training and inference pipelines in Figure 2.

**Zero-shot Capacity** The training and inference processes of the model differ in that the demonstration  $D^V$  is not provided during training. Despite this, the model exhibits zero-shot video imitation capabilities. This can be attributed to two key reasons. Firstly, no special separation token is inserted between frames, allowing the preceding sequences  $t_{<i}$  in Equation (3) to be implicitly viewed in a demonstration-query format, i.e.,  $P(t_i | t_{<i, >j}, t_{\leq j})$  where  $j < i$ . This implies that the model has inherently learned to handle sequences resembling the demonstration-query structure during training. Secondly, the autoregressive nature of the Transformer enables it to seamlessly extend its sequence prediction capabilities to scenarios where the demonstration and query come from different videos, thus facilitating smooth generalization to the imitation paradigm. Our experiments further show that training the model with explicit imitation examples does not yield significant improvements over the zero-shot approach. Refer to Sec 5.1 for more discussions.

**Incorporate Other Modalities** We mainly focus on videos as the demonstration in this paper, but our approach can be easily extended to acquiring other modalities such as text. To do so, we just need to transfer original text descriptions into latent representations denoted as  $c$  by a pretrained language model (Raffel et al., 2020), then take  $c$  as an additional condition while training the Transformer (Equation (3)) as well as imitation inference (Equation (4)). We show in experiments Sec 5.5 that our model understands and reacts to both text and video demonstrations.

### 3.3 DATA

While large language models are trained on vast amounts of data (usually trillions of tokens (Hoffmann et al., 2022; Touvron et al., 2023)), high-quality video data is constrained. In addition, VidIT prefers videos that exhibit not only rich content but also clear causal relationships and interactivity. As a result, among various public video datasets, we focus on these accomplish embodied tasks and select two primary datasets as our main training data sources: 1) **Ego4d** (Grauman et al., 2022), an egocentric video dataset featuring abundant first-person activities; and 2) **Kinetics-600** (Carreira et al., 2018), a comprehensive video dataset comprising diverse human activities. Additionally, we incorporate **WebVid** (Bain et al., 2021) that contains a large amount of internet videos, to augment the variety of video content. Notably, due to the typically ambiguous and inconsistent semantic information conveyed by internet videos, we only utilize a small fraction of WebVid, and show that simply increasing the data scale by adding more internet videos does not help improve the imitation capacity of the model (refer to Sec 5.5).

To validate the VidIT’s imitation capability, we choose the **Something-Something v2** (SSv2) as the main evaluation dataset (Goyal et al., 2017) where each video depicts a basic action that occurs in the physical world, such as moving objects in a direction. These videos convey strong semantic information, which can be utilized to construct the demonstrations. We utilize the evaluation split of SSv2 as the evaluation set for all experiments. In addition, we include the **Robotics Transformer-1** (RT-1) (Brohan et al., 2022) dataset and **MineRL**<sup>2</sup> to demonstrates VidIT’s effectiveness on embodied AI and interactive tasks.

<sup>2</sup><https://github.com/minerllabs/minerl>

## 270 4 EXPERIMENTAL SETUPS

271  
272 We design tailored evaluation pipelines to assess the VidIT’s capacity to perform video imitation.  
273 We introduce the proposed evaluation metrics in Section 4.1, and summarize the details of our  
274 implementations in Section 4.2.

### 276 4.1 EVALUATION

277  
278 The evaluation of VidIT should include two aspects, the first is the visual quality as a normal video  
279 generation task, and the second is semantic accuracy which is more critical as it reflects whether  
280 the model understands and follows the semantic guidance provided by demonstrations. To validate  
281 semantic accuracy, we design four kinds of demonstrations.

- 282 • **No Demonstration.** This setting is akin to unconditional video prediction, where the model  
283 is asked to complete the query  $x^V$  without any demonstration.
- 284 • **Random Demonstration.** In this type of demonstration, video clips are chosen from arbi-  
285 trary videos within the entire SSv2 dataset.
- 286 • **In-class Demonstration.** Given a video clip query, the demonstrations are sampled from  
287 videos with the same action label as the query.
- 288 • **Contrastive Demonstration.** In contrast to the previous one, the sampled demonstrations  
289 have the contrast action label to the query.  
290  
291

292 We then evaluate the model in two settings, 1) using ground truth query labels to evaluate results,  
293 to test whether the demonstrations enhance or deviate from the generation results; 2) using demon-  
294 stration labels, to assess whether the model accurately acquires the semantic guidance and generates  
295 corresponding results. We define various metrics to deal with these two settings.

296  
297 **Automatic Metrics** For visual quality, we adopt several widely utilized metrics including **PSNR**  
298 and **LPIPS** (Zhang et al., 2018) to perform a frame-by-frame comparison between the generation  
299 results and the ground truth video clip, i.e., the consequences of the query. Additionally, we include  
300 the **FID** (Heusel et al., 2017) and **FVD** (Unterthiner et al., 2018) score to quantify the distribution  
301 difference. These metrics are suitable for the first setting where ground truth is available.

302 For semantic accuracy, we propose two classification-based metrics:

- 303 • **Video Accuracy (V-Acc).** We utilize an off-the-shelf video classifier (Tong et al., 2022)  
304 trained on the SSv2 dataset to calculate the classification accuracy of the generated video  
305 clips. Specifically, we predict the action label of each generated video clip and com-  
306 pare it with the ground truth label of the query. This metric provides a perceptual as-  
307 sessment of the semantic information in generation results. In particular, we use the  
308 `videomae-base-finetuned-kinetics3` checkpoint.
- 309 • **Probing Accuracy (P-Acc).** While V-Acc relies on a pretrained classifier and judges on vi-  
310 sual signals, we propose another metric named P-Acc to operate on the latent representation  
311 of VidIT. To directly validate the semantic information contained in latent representations,  
312 inspired by the widespread use of probing in vision feature extractors (He et al., 2022;  
313 Huang et al., 2023; Bardes et al., 2023), we train a probing classifier which takes hiddens  
314 in the last Transformer layer of the generation results as input and predicts the correspond-  
315 ing action label.

316 Most automatic metrics are only suitable for the first setting where ground truth video is available.  
317 For the second setting without ground truth, only the probing accuracy can be utilized. As a supple-  
318 mentary, we also introduce human evaluation.

319  
320 **Human Evaluation** We manually select a subset from the SSv2 validation set, by picking out  
321 action labels with contrastive semantics, such as *Pulling [something] from left to right* and *Pulling*  
322 *[something] from right to left*. For a query video clip, we sample two demonstrations from the  
323

<sup>3</sup><https://huggingface.co/MCG-NJU/videomae-base-finetuned-kinetics>

324 same and contrastive class respectively, constructing a pair of evaluation samples (just as shown  
 325 in Figure 1). We involve 10 experienced users to score 20 pairs of samples from several aspects,  
 326 i.e., visual quality, semantic alignment and control. Alignment indicates whether each generation  
 327 is semantically consistent with the demonstration, while control justifies whether the pair of results  
 328 perform contrastive behavior following the demonstrations. Participants were presented with a pair  
 329 of videos at a time and asked to rate each video for each score on a scale of 1 to 5. We calculated  
 330 the average score as the final result.

## 332 4.2 IMPLEMENTATION DETAILS

334 **Preprocessing** For each dataset, we assign a stride to sample frames at regular intervals in order to  
 335 capture human-recognizable video clips. The stride varies among datasets depending on the average  
 336 FPS but remains consistent within each dataset. Subsequently, we resize and center-crop the images  
 337 to a fixed size of  $256 \times 256$  resolution. The pretrained VQ-GAN tokenizer (Patil et al., 2024) takes in  
 338  $256 \times 256$  sized images and produces  $16 \times 16 = 256$  discrete tokens with a compression coefficient  
 339  $f = 16$ . Tokens of different frames are concatenated in the order of original frames, and each  
 340 training sequence contains 16 images and 4096 discrete tokens.

341 **Transformer Architecture** we adopt the LLaMA (Touvron et al., 2023) architecture, a typical  
 342 decoder-only transformer model for auto-regressive modeling with designs advantageous to large-  
 343 scale modeling. we experiment with different sets of hyper-parameters which result in models with  
 344 300M, 700M and 1.1B parameters separately. The details of hyper-parameter settings are shown in  
 345 Table 8.

347 **Inference** In our main experiments, we set  $k = 1$  to restrict the demonstration to one video. In this  
 348 way, the demonstration clip contains 8 frames, while the query and generation results both consist  
 349 of 4 frames. We also study different compositions such as  $k = 2$  demonstrations sampled from  
 350 different videos and each with 4 frames. Please refer to Section 5.5 for more discussions.

352 **Baselines** We compare the proposed VidIT model with recent visual generative models based on  
 353 auto-regressive Transformer. **LVM** (Bai et al., 2023) is a natural baseline that shows strong image  
 354 imitation capacity. **DeLVM** (Guo et al., 2024) is a data efficient version of LVM, which involves  
 355 extra knowledge distilling step and have fewer parameters. **LWM** (Liu et al., 2024) is designed to  
 356 understand and generate long-context video contents, which is trained on aligned text-video pairs.  
 357 To evaluate the generalizability of our model, we train several variants: *Pretrain*, which is trained  
 358 without SSv2 in the training set; *Pretrain w/ In-domain finetune*, which is fine-tuned on SSv2 from  
 359 the *Pretrain* model where each sample consists of continuous 16 frames, similar to the pretraining  
 360 stage; and *Pretrain w/ Imitation finetune*, which is fine-tuned using the same data format as imitation  
 361 inference, consisting of two 8-frame videos from the same class.

## 363 5 RESULTS AND ANALYSIS

### 365 5.1 MAIN RESULTS

367 If not specified, we use our largest 1.1B model to generate results. We first evaluate the VidIT  
 368 results using the ground truth query label. In this setting, all automatic metrics can be utilized. The  
 369 quantitative results are presented in Table 1, where each row corresponds to a different demonstration  
 370 type. We draw the following conclusions from the results.

371 **In-class Demonstrations Contribute to Stronger Semantic Accuracy** Our analysis begins by  
 372 focusing on the result of the *Pretrain* model in the first block. We observe a significant enhancement  
 373 in the semantic accuracy when using in-class video clips as demonstrations, compared to without  
 374 demonstration or using randomly selected ones. Specifically, there is a notable 7.1% / 8.4% improve-  
 375 ment in P-Acc and a 1.8% / 2.0% gain in V-Acc over no or prompting with random demonstrations  
 376 respectively. These findings validate our proposition that when prompted with semantically related  
 377 demonstrations, the model more accurately generates video clips that adhere to the original trace, in  
 a zero-shot way on both the domain and the imitation ability.

Table 1: The Semantic Accuracy metrics and Visual Quality metric of different demonstration type and training strategies. The best result in each block is **bolded** and second best result is underlined.

Demonstration	Semantic Accuracy		Visual Quality			
	V-Acc $\uparrow$	P-Acc $\uparrow$	PSNR $\uparrow$	LPIPS $\downarrow$	FID $\downarrow$	FVD $\downarrow$
<i>Pretrain</i>						
No Demonstration	22.9	29.6	<b>13.20</b>	<b>0.442</b>	<b>20.94</b>	<b>119.51</b>
Random Demonstration	22.7	28.3	13.01	0.450	22.53	131.34
In-class Demonstration	<b>24.7</b>	<b>36.7</b>	<u>13.07</u>	<u>0.447</u>	<u>21.95</u>	<u>125.77</u>
<i>Pretrain w/ Imitation Finetune</i>						
No Demonstration	<u>24.2</u>	<u>26.7</u>	<u>13.02</u>	0.460	22.21	129.68
Random Demonstration	23.1	25.6	13.01	<u>0.454</u>	<u>20.66</u>	<u>115.76</u>
In-class Demonstration	<b>25.7</b>	<b>40.7</b>	<b>13.08</b>	<b>0.450</b>	<b>20.49</b>	<b>113.52</b>
<i>Pretrain w/ In-domain Finetune</i>						
No Demonstration	<u>25.0</u>	30.8	13.09	0.446	20.87	108.12
Random Demonstration	24.9	<u>34.1</u>	<u>13.16</u>	0.440	<u>19.17</u>	<u>95.72</u>
In-class Demonstration	<b>25.9</b>	<b>48.8</b>	<b>13.21</b>	<b>0.438</b>	<b>18.92</b>	<b>88.63</b>

Table 2: Results on automatic and human evaluation of VidIT and baseline models.

Baselines	Automatic Metrics		Human Evaluation		
	P-Acc	PSNR	Quality	Alignment	Control
LVM (Bai et al., 2023)	27.1	12.45	<u>3.50</u>	<u>2.71</u>	1.81
DeLVM (Guo et al., 2024)	<u>32.8</u>	<u>12.69</u>	3.21	2.28	<u>2.11</u>
LWM (Liu et al., 2024)	24.9	11.89	3.39	2.33	2.05
<b>VidIT (ours)</b>	<b>38.5</b>	<b>13.07</b>	<b>4.12</b>	<b>3.73</b>	<b>3.01</b>

**Random Demonstrations Mislead the Semantics** Beyond the inferior performance compared to in-class demonstrations, random ones also reveal a gap of 0.2% in V-Acc and 0.7% in P-Acc to the no demonstration setting. This indicates that prompting with unrelated demonstrations negatively impacts the semantics of the generated results.

**In-domain Finetuning Leads to General Improvement** We further analyze the results presented in the second and third blocks of Table 1, focusing on different fine-tuning strategies. The superiority of in-domain finetuning over the initial *Pretrain* is evident, indicating a straightforward performance boost when in-domain training data is available. In addition, when comparing in-domain finetuning with imitation finetuning, the in-class P-Acc of the former rises to 48.8%, showcasing an 8.1% improvement over the latter. These findings underscore the remarkable *zero-shot capacity* of VidIT, showcasing the model’s capability at adapting to the imitation formatting through exclusive training on continuous video clips, while explicit finetuning with the target format fails to yield significant enhancements.

**Evaluation with Demonstration Label** We evaluate the generation results with the demonstration label to assess the controllability of the model. We use results with both in-class and contrastive demonstrations as described in Section 4.1 to train the probing model. As shown in Table 3, both cases have well above average probing accuracy, indicating that the representations of generation results successfully obtain the semantic information from demonstrations.

Table 3: The Probing accuracy evaluated by demonstration labels.

Demonstration	P-Acc
Contrastive	35.5
In-class	33.8
Random	16.7

## 5.2 COMPARISON WITH BASELINE MODELS

We compare our VidIT with baseline methods mentioned in Section 4.2 on the video imitation abilities. We report both automatic metrics and human evaluation results in Table 2. As shown in the results, our VidIT model significantly outperforms the baseline models in both objective and subjective evaluations. We attribute this difference to the training data formats of VidIT and other models.

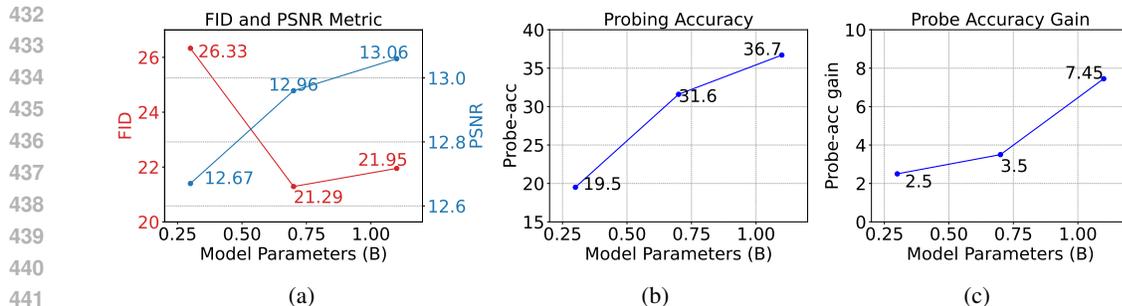


Figure 3: The performance when scaling model parameters.

LVM and DeLVM mainly relies on pairs of images and their annotations, which limits their ability to extract semantic information across consecutive frames and generate coherent sequences correspondingly. Moreover, the training process of LVM are not designed to capture the demonstration’s semantics, leading to degradation on imitation performance.

### 5.3 SCALING BEHAVIOR

As describe in Section 4.2, we set up different sizes of model, respectively 300M, 700M and 1.1B. In this section, we assess the models’ scaling behavior from the following aspects.

**Scalability on Visual Quality** We present the PSNR and FID scores of models in various sizes in Figure 3a. The results indicate that larger models tend to produce samples of higher quality.

**Scalability on Semantic Accuracy** We present both the absolute probing accuracy scores of in-class demonstrations and the performance gain over random demonstrations in Figure 3b and 3c. From both results, we observe a clear trend that larger models yield more informative hidden representations, and offer more precise control over the generation results when conditioned on different demonstrations.

### 5.4 GENERALIZATION ABILITY

In this section, we elaborate on the generalization ability of the VidIT model. Thanks to its self-supervised training paradigm, VidIT can be smoothly adapted to multiple scenarios and tasks with minimal modifications.

Firstly, we apply VidIT to Robotics Transformer dataset (Brohan et al., 2022) to imitate actions of a robotic arm, such as opening and closing drawers and placing objects. The samples on Figure 4 demonstrates that VidIT can successfully mimic various robotic arm actions from demonstrations, completing the frame sequences with the same movements.

Additionally, we highlight that VidIT can serve as an underlying video generation engine within an interactive agent. In this setup, VidIT receives previous observations and generates subsequent frames, which are then converted to action signals via the inverse dynamics mechanism. This process ensures that the action signals reflect VidIT’s imitation intent. As depicted in Figure 7 in Appendix A.1, agents perform specific actions guided by VidIT, showcasing the zero-shot imitation ability of VidIT and revealing its potential to interact with the environment.

We also validate VidIT’s capability as a video task imitator. Following the approach of image task imitators like LVM, we create video segmentation exemplars to form QA pairs on the VIPSeg (Miao et al., 2022) dataset. The results, presented in Figure 8 in Appendix A.2, indicate that VidIT successfully learns the task format from the demonstrations.

### 5.5 ABLATION AND ANALYSIS

**Different Number of Demonstrations** We investigate the impact of different numbers of demonstrations on VidIT performance. Within a 16-frame context window, we fix the query length as 4 frames and the generation as 4 frames, and construct four demonstration formats: 1) no demonstration; 2) one 4-frame demonstration; 3) two 4-frame demonstrations sampled from different videos;

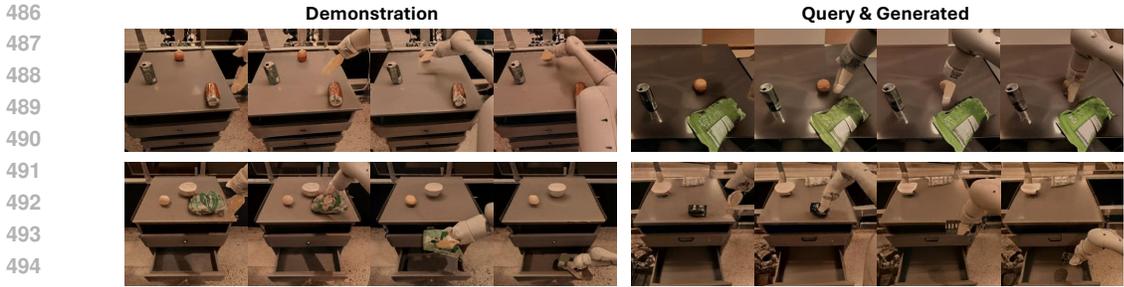


Figure 4: Generation Samples of VidIT on RT-1 Dataset.

4) one 8-frame demonstration. The results in Table 4 demonstrate that providing demonstrations generally improves both P-Acc and V-acc scores. Additionally, providing more ( $2 \times 4$  versus  $1 \times 4$ ) or longer demonstrations ( $1 \times 8$  versus  $1 \times 4$ ) both result in better performance.

Table 4: Semantic accuracy of different demonstration formats.

Demonstration	P-Acc	V-Acc
No	29.6	22.9
$1 \times 4$ frames	35.6	22.9
$2 \times 4$ frames	<b>37.2</b>	24.2
$1 \times 8$ frames	36.7	<b>24.7</b>

Table 5: Results with text as demonstrations.

	V-Acc	FID
Tune w/ text	<b>28.8</b>	<b>19.46</b>
+ Infer w/o text	24.6	19.72
Tune w/o text	24.7	20.02

**Text as Demonstrations** To demonstrate the versatility of our model, we explore incorporating text into the demonstrations. We fine-tune the 300M *Pretrain* model on the SSv2 dataset by appending corresponding text annotations to the front of each video. The text annotations are encoded into latent representations using a pretrained T5-large model (Raffel et al., 2020). The results, presented in Table 5, indicate that our model can seamlessly adopt demonstrations in various modalities. Additionally, adding textual descriptions enhances the model’s imitation capacity.

**Effect of Web Videos** In previous studies, incorporating miscellaneous web videos into the training dataset has been shown to enhance generation quality (Yang et al., 2023; Liu et al., 2024). Here we investigate their effect on VidIT. We sampled varying quantities of video clips from the Webvid dataset during pretraining and present the SSv2 validation perplexity results in Figure 5. We observe that the perplexity does not linearly decrease with the number of web videos, suggesting that VidIT prefers videos conveying clear and consistent semantic information.

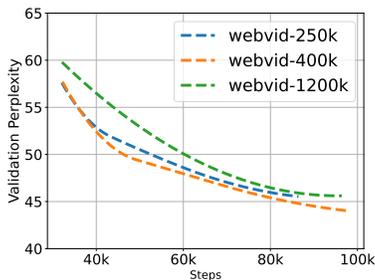


Figure 5: The validation perplexity in various amount of web videos.

## 6 CONCLUSION

In this paper, we introduce Video Imitator, a novel model that extends imitation learning to video data. By converting video frames into sequences of discrete tokens and training the autoregressive Transformer with next token prediction, our VidIT model exhibits zero-shot imitation capabilities, which can generate semantically coherent and contextually appropriate video sequences based on provided demonstrations. Extensive experiments confirm that VidIT effectively captures and conveys the semantic information embedded in the demonstrations, demonstrating its potential to enhance various downstream tasks such as visual planning. Additionally, the model’s versatility is verified by its ability to understand and integrate multi-modal demonstrations simultaneously.

## REFERENCES

- 540  
541  
542 Yutong Bai, Xinyang Geng, Karttikeya Mangalam, Amir Bar, Alan Yuille, Trevor Darrell, Jitendra  
543 Malik, and Alexei A Efros. Sequential modeling enables scalable learning for large vision models.  
544 *arXiv preprint arXiv:2312.00785*, 2023.
- 545 Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and  
546 image encoder for end-to-end retrieval. In *Proceedings of the IEEE/CVF International Conference*  
547 *on Computer Vision*, pp. 1728–1738, 2021.
- 548 Amir Bar, Yossi Gandelsman, Trevor Darrell, Amir Globerson, and Alexei Efros. Visual prompting  
549 via image inpainting. *Advances in Neural Information Processing Systems*, 35:25005–25017,  
550 2022.
- 551  
552 Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mido  
553 Assran, and Nicolas Ballas. V-jepa: Latent video prediction for visual representation learning.  
554 2023.
- 555 Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn,  
556 Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, Julian Ibarz, Brian  
557 Ichter, Alex Irpan, Tomas Jackson, Sally Jesmonth, Nikhil Joshi, Ryan Julian, Dmitry Kalash-  
558 nikov, Yuheng Kuang, Isabel Leal, Kuang-Huei Lee, Sergey Levine, Yao Lu, Utsav Malla, Deek-  
559 sha Manjunath, Igor Mordatch, Ofir Nachum, Carolina Parada, Jodilyn Peralta, Emily Perez,  
560 Karl Pertsch, Jornell Quiambao, Kanishka Rao, Michael Ryoo, Grecia Salazar, Pannag Sanketi,  
561 Kevin Sayed, Jaspriar Singh, Sumedh Sontakke, Austin Stone, Clayton Tan, Huong Tran, Vincent  
562 Vanhoucke, Steve Vega, Quan Vuong, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and  
563 Brianna Zitkovich. Rt-1: Robotics transformer for real-world control at scale. In *arXiv preprint*  
564 *arXiv:2212.06817*, 2022.
- 565 Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe  
566 Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. Video  
567 generation models as world simulators. 2024. URL [https://openai.com/research/  
568 video-generation-models-as-world-simulators](https://openai.com/research/video-generation-models-as-world-simulators).
- 569 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,  
570 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are  
571 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- 572 Joao Carreira, Eric Noland, Andras Banki-Horvath, Chloe Hillier, and Andrew Zisserman. A short  
573 note about kinetics-600. *arXiv preprint arXiv:1808.01340*, 2018.
- 574  
575 Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam  
576 Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm:  
577 Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):  
578 1–113, 2023.
- 579  
580 Yilun Du, Sherry Yang, Bo Dai, Hanjun Dai, Ofir Nachum, Josh Tenenbaum, Dale Schuurmans, and  
581 Pieter Abbeel. Learning universal policies via text-guided video generation. *Advances in Neural*  
582 *Information Processing Systems*, 36, 2023.
- 583 Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image  
584 synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recogni-  
585 tion*, pp. 12873–12883, 2021.
- 586  
587 Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne West-  
588 phal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al.  
589 The” something something” video database for learning and evaluating visual common sense. In  
590 *Proceedings of the IEEE international conference on computer vision*, pp. 5842–5850, 2017.
- 591 Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Gird-  
592 har, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in  
593 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision*  
*and Pattern Recognition*, pp. 18995–19012, 2022.

- 594 Jianyuan Guo, Zhiwei Hao, Chengcheng Wang, Yehui Tang, Han Wu, Han Hu, Kai Han, and  
595 Chang Xu. Data-efficient large vision models through sequential autoregression. *arXiv preprint*  
596 *arXiv:2402.04841*, 2024.
- 597  
598 Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian,  
599 and Yujiu Yang. Connecting large language models with evolutionary algorithms yields powerful  
600 prompt optimizers. *arXiv preprint arXiv:2309.08532*, 2023.
- 601 Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked au-  
602 toencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer*  
603 *vision and pattern recognition*, pp. 16000–16009, 2022.
- 604  
605 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter.  
606 Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in*  
607 *neural information processing systems*, 30, 2017.
- 608 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*  
609 *neural information processing systems*, 33:6840–6851, 2020.
- 610  
611 Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P  
612 Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition  
613 video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022a.
- 614  
615 Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J  
616 Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–  
8646, 2022b.
- 617  
618 Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza  
619 Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Train-  
620 ing compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- 621  
622 Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pre-  
623 training for text-to-video generation via transformers. In *International Conference on Learning*  
*Representations*, 2023.
- 624  
625 Zhicheng Huang, Xiaojie Jin, Chengze Lu, Qibin Hou, Ming-Ming Cheng, Dongmei Fu, Xiaohui  
626 Shen, and Jiashi Feng. Contrastive masked autoencoders are stronger vision learners. *IEEE*  
627 *Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- 628  
629 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child,  
630 Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language  
631 models. *arXiv preprint arXiv:2001.08361*, 2020.
- 632  
633 Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Rachel Hornung, Hartwig  
634 Adam, Hassan Akbari, Yair Alon, Vighnesh Birodkar, et al. Videopoet: A large language model  
635 for zero-shot video generation. *arXiv preprint arXiv:2312.14125*, 2023.
- 636  
637 Bei Li, Rui Wang, Junliang Guo, Kaitao Song, Xu Tan, Hany Hassan, Arul Menezes, Tong Xiao,  
638 Jiang Bian, and JingBo Zhu. Deliberate then generate: Enhanced prompting framework for text  
639 generation. *arXiv preprint arXiv:2305.19835*, 2023.
- 640  
641 Hao Liu, Wilson Yan, Matei Zaharia, and Pieter Abbeel. World model on million-length video and  
642 language with ringattention. *arXiv preprint arXiv:2402.08268*, 2024.
- 643  
644 Jiayu Miao, Xiaohan Wang, Yu Wu, Wei Li, Xu Zhang, Yunchao Wei, and Yi Yang. Large-scale  
645 video panoptic segmentation in the wild: A benchmark. In *Proceedings of the IEEE Conference*  
*on Computer Vision and Pattern Recognition*, 2022.
- 646  
647 Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models.  
In *International conference on machine learning*, pp. 8162–8171. PMLR, 2021.
- Suraj Patil, William Berman, Robin Rombach, and Patrick von Platen. amused: An open muse  
reproduction. *arXiv preprint arXiv:2401.01808*, 2024.

- 648 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language  
649 models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- 650
- 651 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi  
652 Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text  
653 transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
- 654 Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: En-  
655 hanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.
- 656
- 657 Stephen Tian, Chelsea Finn, and Jiajun Wu. A control-centric benchmark for video prediction. In  
658 *International Conference on Learning Representations*, 2023.
- 659 Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. VideoMAE: Masked autoencoders are data-  
660 efficient learners for self-supervised video pre-training. In *Advances in Neural Information Pro-  
661 cessing Systems*, 2022.
- 662
- 663 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée  
664 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and  
665 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- 666
- 667 Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski,  
668 and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges.  
669 *arXiv preprint arXiv:1812.01717*, 2018.
- 670
- 671 Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in  
672 neural information processing systems*, 30, 2017.
- 673
- 674 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,  
675 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural informa-  
676 tion processing systems*, 30, 2017.
- 677
- 678 Ruben Villegas, Arkanath Pathak, Harini Kannan, Dumitru Erhan, Quoc V Le, and Honglak Lee.  
679 High fidelity video prediction with large stochastic recurrent neural networks. *Advances in Neural  
680 Information Processing Systems*, 32, 2019.
- 681
- 682 Ruben Villegas, Mohammad Babaeizadeh, Pieter-Jan Kindermans, Hernan Moraldo, Han Zhang,  
683 Mohammad Taghi Saffar, Santiago Castro, Julius Kunze, and Dumitru Erhan. Phenaki: Variable  
684 length video generation from open domain textual descriptions. In *International Conference on  
685 Learning Representations*, 2023.
- 686
- 687 Xinlong Wang, Wen Wang, Yue Cao, Chunhua Shen, and Tiejun Huang. Images speak in images:  
688 A generalist painter for in-context visual learning. In *Proceedings of the IEEE/CVF Conference  
689 on Computer Vision and Pattern Recognition*, pp. 6830–6839, 2023a.
- 690
- 691 Xinlong Wang, Xiaosong Zhang, Yue Cao, Wen Wang, Chunhua Shen, and Tiejun Huang. Seggpt:  
692 Segmenting everything in context. *arXiv preprint arXiv:2304.03284*, 2023b.
- 693
- 694 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny  
695 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in  
696 neural information processing systems*, 35:24824–24837, 2022.
- 697
- 698 Mengjiao Yang, Yilun Du, Kamyar Ghasemipour, Jonathan Tompson, Dale Schuurmans, and Pieter  
699 Abbeel. Learning interactive real-world simulators. *arXiv preprint arXiv:2310.06114*, 2023.
- 700
- 701 Sherry Yang, Jacob Walker, Jack Parker-Holder, Yilun Du, Jake Bruce, Andre Barreto, Pieter  
Abbeel, and Dale Schuurmans. Video as the new language for real-world decision making. *arXiv  
preprint arXiv:2402.17139*, 2024.
- Lijun Yu, Yong Cheng, Kihyuk Sohn, José Lezama, Han Zhang, Huiwen Chang, Alexander G  
Hauptmann, Ming-Hsuan Yang, Yuan Hao, Irfan Essa, et al. Magvit: Masked generative video  
transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recog-  
nition*, pp. 10459–10469, 2023a.

702 Lijun Yu, Jose Lezama, Nitesh Bharadwaj Gundavarapu, Luca Versari, Kihyuk Sohn, David Min-  
703 nen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G Hauptmann, et al. Language model  
704 beats diffusion-tokenizer is key to visual generation. In *The Twelfth International Conference on*  
705 *Learning Representations*, 2023b.

706 Maryam Zare, Parham M Kebria, Abbas Khosravi, and Saeid Nahavandi. A survey of imitation  
707 learning: Algorithms, recent developments, and challenges. *IEEE Transactions on Cybernetics*,  
708 2024.

709 Biao Zhang and Rico Sennrich. Root mean square layer normalization. *Advances in Neural Infor-*  
710 *mation Processing Systems*, 32, 2019.

711 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable  
712 effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on*  
713 *computer vision and pattern recognition*, pp. 586–595, 2018.

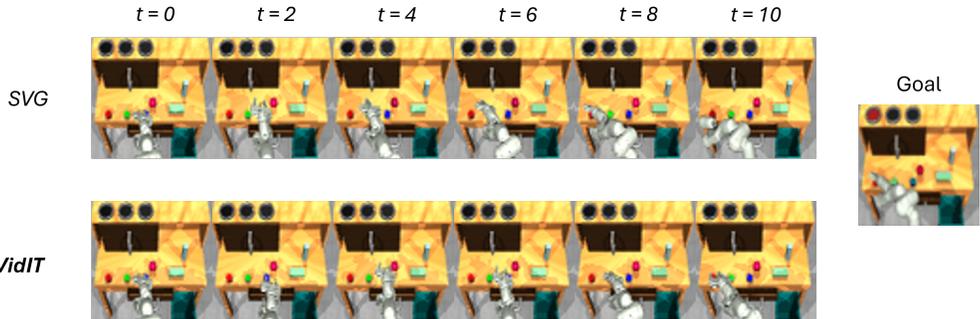
714  
715  
716  
717  
718  
719  
720  
721  
722  
723  
724  
725  
726  
727  
728  
729  
730  
731  
732  
733  
734  
735  
736  
737  
738  
739  
740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755

756 A ADDITIONAL RESULTS

757  
758 In this section we provide more samples on different video datasets and benchmarks that  
759 VidIT model produces. Synthesis videos on Something-something v2 dataset are presented  
760 in Table 10. One can also navigate to [https://anonymous.4open.science/w/  
761 VidImit-page-656D/](https://anonymous.4open.science/w/VidImit-page-656D/) for clearer presented samples. Table 11 includes samples on Robotics  
762 Transformer-1 dataset, showcasing the ability of VidIT to fit on various form of video datasets. Fur-  
763 thermore, VidIT as a versatile generalist, can act as a simulator in reinforcement learning tasks, for  
764 which we provide detailed explanations.

765  
766 A.1 VIDIT AS SIMULATOR

767  
768 We demonstrate that VidIT can also function as a simulator in reinforcement learning tasks by eval-  
769 uating it on the VP<sup>2</sup> (Tian et al., 2023) benchmark. VP<sup>2</sup> is a benchmark for video prediction mod-  
770 els used in robotic manipulation via model-predictive control. VidIT generates future frames with  
771 videos that accomplishing the same task as the demonstration, where the generated frames inversely  
772 reflect the corresponding actions that correctly interact with the environment (i.e. inverse dynam-  
773 ics). We compare the trajectories produced by the SVG (Villegas et al., 2019) baseline and VidIT in  
774 Figure 6. Evaluating on the *Push-red* task, we find that VidIT provides more precise control over  
775 the environment interaction.



788 Figure 6: The trajectory produced by SVG and VidIT respectively.

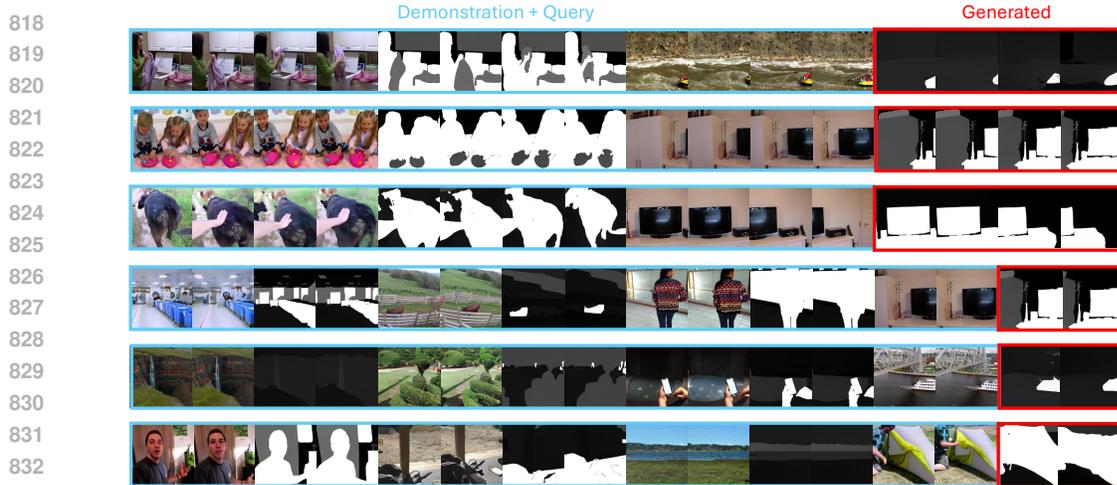
789  
790 Additionally, we evaluate VidIT on the MineRL environment as the underlying video generation  
791 model, where VidIT take in previous game scenes as conditions to imitate and generate subsequent  
792 scenes. Generated game scenes are then fed into inverse dynamics models to generate the next action  
793 signal. We showcase two types of actions: *Chop wood* and *Attack sheep*. In Figure 7, the VidIT  
794 correctly generate game scenes and produce corresponding action signal to finish the action.



Figure 7: The VidIT model functions as video generator for *chop-wood* and *attack-sheep* in MineRL.

## 810 A.2 VIDIT AS TASK IMITATOR

811 We expand the role of VidIT model to the task imitator on video, which mirrors the cases in image  
 812 imitation learning. In this section, we finetune VidIT in VIPSeg dataset (Miao et al., 2022) for a few  
 813 steps and show that VidIT can also smoothly generalize to video perceptual tasks like video seg-  
 814 mentation. As depicted in Figure 8, VidIT successfully learns the task format from demonstration,  
 815 and generates the segmentation outputs of the query video.  
 816



834 Figure 8: The VidIT can also imitate video perceptual tasks such as video segmentation.

## 835 A.3 CASE STUDY: BASELINE METHODS

836 In this section we compare the generated cases of VidIT with the baseline methods. As in Figure 9,  
 837 we see that conditioned on the same frames, VidIT is capable of generating more semantically ac-  
 838 curate and high-quality video than baseline models. Moreover, VidIT is lower in parameter budget,  
 839 highlighting the priority of VidIT as a video imitator model.  
 840

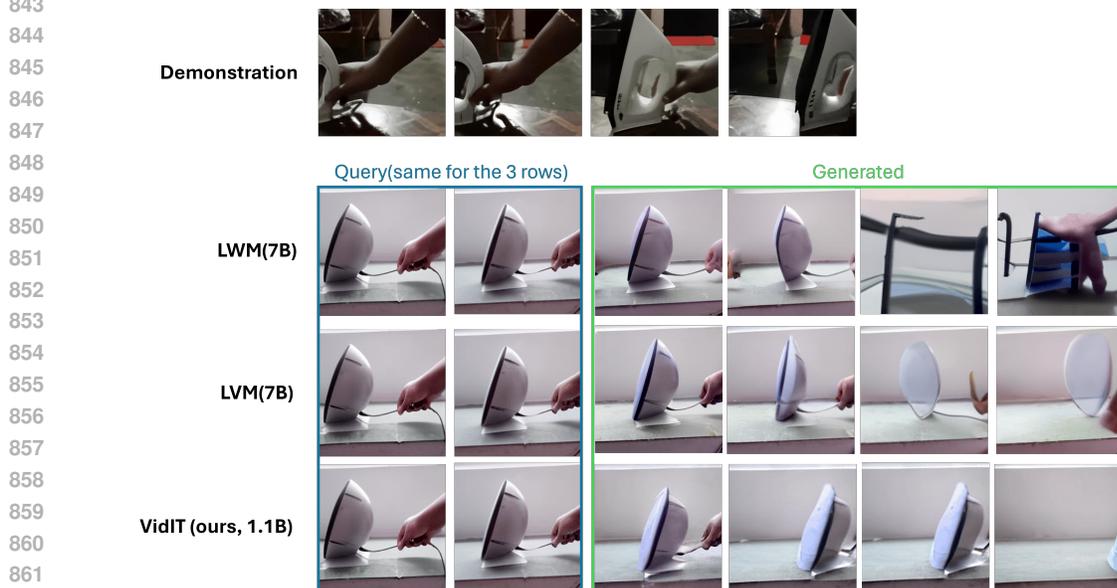


Figure 9: Generated cases of VidIT and baseline methods.

#### A.4 FAILURE CASES

When prompted with semantically unrelated videos, VidIT may struggle to generate coherent content. In this section, we present failure cases where VidIT fails to align the query video with the semantic context of the demonstration video (see Figure 10).

For instance, in the first row, the demonstration video conveys the semantic information of moving a ball downward. However, when starting with a query video depicting an unrelated action, such as folding a blanket, the model struggles to reconcile the two contexts. As a result, it tends to ignore the demonstrated actions and generates content that lacks coherence with the demonstration.

Similar patterns are observed in other cases. These examples reveals VidIT’s potential difficulty in maintaining semantic coherence when prompted with unrelated or ambiguous demonstrations. Addressing this limitation presents a promising direction for future research.

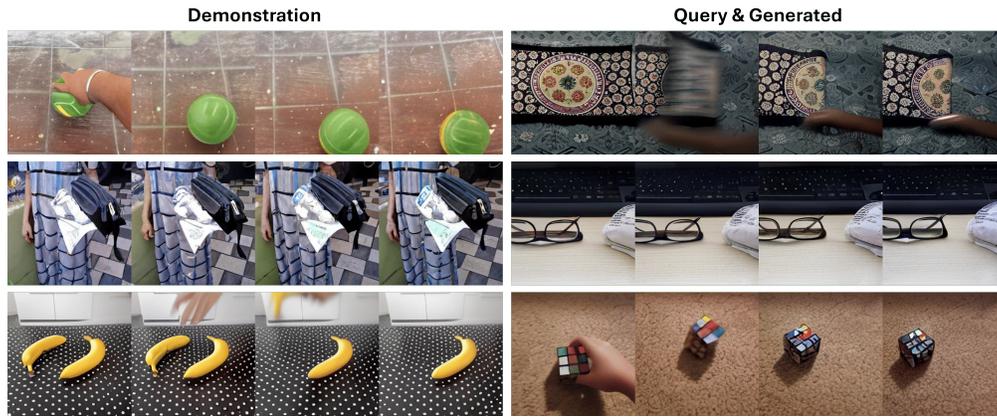


Figure 10: Failure cases of VidIT.

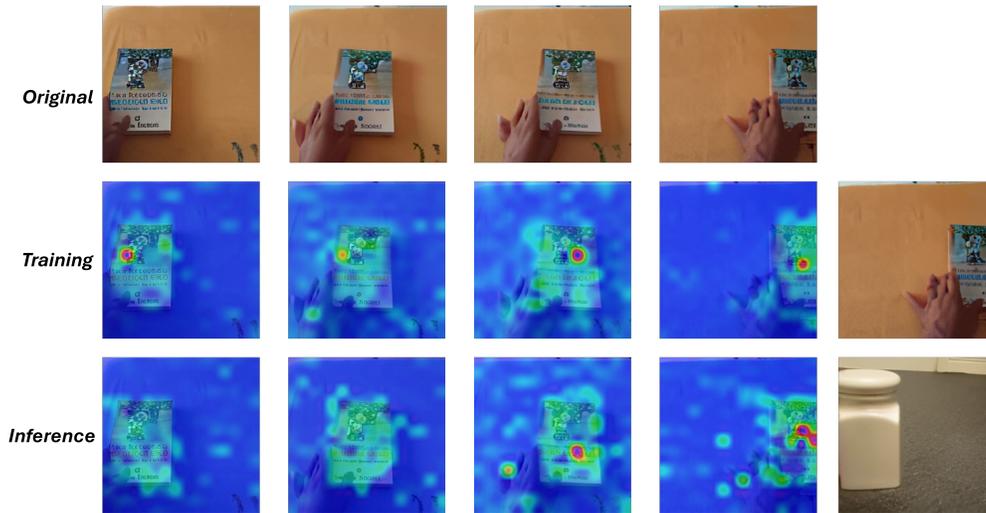


Figure 11: The attention weight visualization during training and inference respectively.

#### A.5 INTERPRETABILITY OF GENERALIZATION

In this section, we highlight the interpretability of VidIT and its ability to generalize to imitation tasks in a zero-shot manner. To demonstrate this, we analyze the attention weights from the language model backbone to compare the similarity of attention patterns when provided with the same demonstration during training and inference.

Specifically, we extract the attention weights from the last layer of the model and average them across all attention heads. Since the tokens are flattened from the original frames in raster scan order, we rebind the tokens to their corresponding regions in the frames for better visualization. The results are shown in Figure 11.

The first row depicts the original frame sequence. The second row represents the training case, where frames are drawn from a single video clip. The third row illustrates the inference case, where the sequence begins from a different query frame. When generating a new token, we observe similar attention patterns across the demonstration frames. Notably, the attention is concentrated around the moving object within the frames.

This phenomenon indicates that the model effectively learns and leverages the semantic information from previous frames, enabling it to perform imitation tasks even when starting from a different scene.

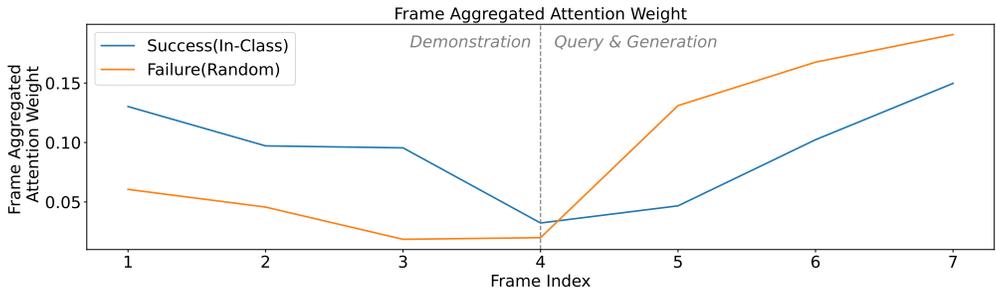


Figure 12: The aggregated attention score distribution in demonstration and query frames.

Additionally, we analyze the failure cases in detail from the perspective of attention weights. Specifically, we examine how new frames attend to previous frames during generation. Intuitively, when a new frame is effectively guided by the demonstration, it tends to allocate more attention to the frames in the demonstration video, rather than only attending to the query frames.

To illustrate this, we visualize the attention distribution of generated frames on previous frames, as shown in Figure 12. Clearly, successful generation samples exhibit a tendency to allocate more attention to the demonstration frames, indicating an ability to comprehend and utilize its semantic information. In contrast, the attention curve for failure cases reveals a higher focus on the query frames, effectively ignoring the guidance provided by the demonstration part.

#### A.6 FULL EVALUATION METRICS ON BASELINES

In this section, we complete Table 2’s results with the all the evaluation metrics in Section 4.1. From the results in Table 6, it is evident that the VidIT model consistently outperforms the baseline models in video imitation tasks, reinforcing its superior capability in capturing both semantic alignment and visual quality.

Table 6: Full automatic metrics for baseline models. The best result is **bolded** .

Demonstration	Semantic Accuracy		Visual Quality			
	V-Acc ↑	P-Acc ↑	PSNR ↑	LPIPS ↓	FID ↓	FVD ↓
LVM (Bai et al., 2023)	21.2	27.1	12.45	0.489	27.58	163.13
DeLVM (Guo et al., 2024)	24.1	32.8	12.69	0.462	23.65	132.09
LWM (Liu et al., 2024)	21.8	24.9	11.89	0.474	29.91	191.56
<b>VidIT (ours)</b>	<b>25.9</b>	<b>38.5</b>	<b>13.07</b>	<b>0.450</b>	<b>22.30</b>	<b>126.32</b>

Table 7: The configuration of different size of models.

	Hidden dim	MLP dim	Num. Heads	Num. Layers
300M	1024	2688	8	22
700M	1536	4096	16	24
1.1B	2048	4096	16	26

Table 8: Model hyperparameters.

Hyperparameter	Value
Learning rate scheduler	inverse sqrt
Learning rate	$5e^{-4}$
Warm up steps	10000
Weight decay	0.01
Optimizer	AdamW
AdamW betas	(0.9, 0.95)
Context length	4096

## B IMPLEMENTATION DETAILS

### B.1 MODEL ARCHITECTURE

We utilize LLaMA as the Transformer architecture in our work. To validate the scaling behavior, we train three different sizes of the VidIT model within the LLaMA architecture: 300M, 700M, and 1.1B. We carefully tune the hidden dimension, MLP intermediate dimension, and the number of Transformer decoder layers to achieve different model sizes. The configuration details of these models are presented in Table 7.

### B.2 HYPERPARAMETERS

The hyperparameters used to train the VidIT model are presented in Table 8. We utilize inverse square root scheduler and start model training with 10,000 warmup steps.

### B.3 DATASET STATISTICS

The summary of dataset used in our paper is presented in Table 9. We provide detailed statistics including number of videos, total of tokens after tokenization and tokens used to train our VidIT model.

### B.4 TRAINING DETAILS

Our largest variant, VidIT 1.1B, is trained on 2x8H100 nodes, with Pytorch DDP parallel strategy integrated in the Pytorch-Lightening trainer. The training time is about 3 hours per epoch and thus the total training GPU hours is 720 (15 epochs).

Table 9: Statistics of training datasets.

Dataset	Sample Stride	Num. Videos	Total tokens(M)	Training tokens(M)
Ego4d	6	14,192	6890.8	2048.0
Kinetics-600	5	426,253	4711.35	1024.0
WebVid	6	2,494,801	50825.2	1024.0
<b>Total</b>	-	-	-	4096.0

## 1026 C LIMITATIONS

1027

1028 Despite the promising results, several limitations exist in this study. Firstly, while VidIT excels at  
1029 generating short video sequences, its performance on longer sequences has not been thoroughly eval-  
1030 uated. Secondly, the model’s dependence on the quality and relevance of provided demonstrations  
1031 can lead to inconsistencies, particularly when demonstrations are noisy or semantically misaligned.  
1032 Future work should address these limitations by utilizing visual tokenizers with temporal compres-  
1033 sion, and scaling or developing robust models to handle imperfect demonstrations.

1034

## 1035 D ETHICS STATEMENT

1036

1037 Positive societal impacts of VidIT include its potential to enhance various downstream applications  
1038 such as embodied planning and robotic simulation by enabling models to imitate actions demon-  
1039 strated in videos. As for negative societal impacts, the widespread adoption of this technique may  
1040 raise concerns about privacy when powerful models gain the ability to analyze and generate video  
1041 content at scale. The datasets utilized in this study are publicly available and have been thoroughly  
1042 reviewed to ensure they do not include personally identifiable information or offensive content.  
1043 Nonetheless, as these datasets are sourced from the Internet, there may still be inherent biases. To  
1044 address this, we have implemented a rigorous filtering process on the training data to minimize the  
1045 potential for the model to generate inappropriate content.

1046

1047

1048

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

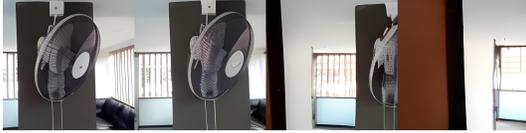
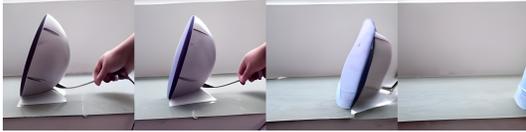
1077

1078

1079

1080  
 1081  
 1082  
 1083  
 1084  
 1085  
 1086  
 1087  
 1088  
 1089  
 1090  
 1091  
 1092  
 1093  
 1094  
 1095  
 1096  
 1097  
 1098  
 1099  
 1100  
 1101  
 1102  
 1103  
 1104  
 1105  
 1106  
 1107  
 1108  
 1109  
 1110  
 1111  
 1112  
 1113  
 1114  
 1115  
 1116  
 1117  
 1118  
 1119  
 1120  
 1121  
 1122  
 1123  
 1124  
 1125  
 1126  
 1127  
 1128  
 1129  
 1130  
 1131  
 1132  
 1133

Table 10: More samples of VidIT generated samples on Something-something v2 dataset.

Demonstration	Query & Generation
	
	
	
	
	
	
	
	

1134  
1135  
1136  
1137  
1138  
1139  
1140  
1141  
1142  
1143  
1144  
1145  
1146  
1147  
1148  
1149  
1150  
1151  
1152  
1153  
1154  
1155  
1156  
1157  
1158  
1159  
1160  
1161  
1162  
1163  
1164  
1165  
1166  
1167  
1168  
1169  
1170  
1171  
1172  
1173  
1174  
1175  
1176  
1177  
1178  
1179  
1180  
1181  
1182  
1183  
1184  
1185  
1186  
1187

Table 11: More examples of VidIT generated samples on RT-1 dataset.

Demonstration	Query & Generation
	
	
	
	
	