



VideoWorld 2: Learning Transferable Knowledge from Real-world Videos

Zhongwei Ren^{1,2,*}, Yunchao Wei², Xiao Yu², Guixun Luo², Yao Zhao², Bingyi Kang¹,
Jiashi Feng¹, Xiaojie Jin^{1,*,†}

¹ByteDance Seed, ²Beijing Jiaotong University

*Equal Contribution, †Project Lead

Abstract

Learning transferable knowledge from unlabeled video data and applying it in new environments is a fundamental capability of intelligent agents. This work presents **VideoWorld 2**, which extends VideoWorld [51] and provides the first investigation of learning transferable knowledge for complex, long-horizon tasks directly from raw real-world videos. At its core, VideoWorld 2 introduces a dynamics-enhanced Latent Dynamics Model (dLDM) that decouples action dynamics from visual appearance: a pretrained video diffusion model handles visual appearance modeling, enabling the dLDM to learn latent codes that focus on compact and meaningful task-related dynamics. These latent codes are then modeled autoregressively to learn task policies and support long-horizon reasoning. We evaluate VideoWorld 2 on challenging real-world handcraft making tasks, where prior video generation and latent-dynamics models struggle to operate reliably. Remarkably, VideoWorld 2 achieves up to **70% improvement in task success rate** and produces coherent long execution videos. In robotics, we show that VideoWorld 2 can acquire transferable manipulation knowledge from the Open-X dataset, which substantially improves task performance on CALVIN, demonstrating strong cross-domain generalization. This study reveals the potential of learning transferable world knowledge directly from raw videos, with all code, data, and models to be open-sourced for further research.

Date: February 11, 2026

Correspondence: Xiaojiao Jin, Yunchao Wei

Project Page: <https://VideoWorld2.github.io/>

1 Introduction

Current AI models primarily learn knowledge¹ from large-scale text data [2, 5, 10, 16, 18, 33, 57, 60, 70]. However, text alone cannot fully capture the rich information of the real visual world, including world dynamics, spatial relationships, and underlying physical laws. In contrast, animals in nature can acquire knowledge directly from visual signals, and generalize it to solve tasks across diverse scenarios. For instance, a child can reproduce paper-folding skills demonstrated in a video using different paper materials, without any language instruction. Given the vast abundance of video content available on the internet, enabling AI models to learn *generalizable knowledge* [31] from raw video data holds significant promise for scaling their knowledge acquisition and is fundamental to their ability to execute tasks effectively in both real-world and

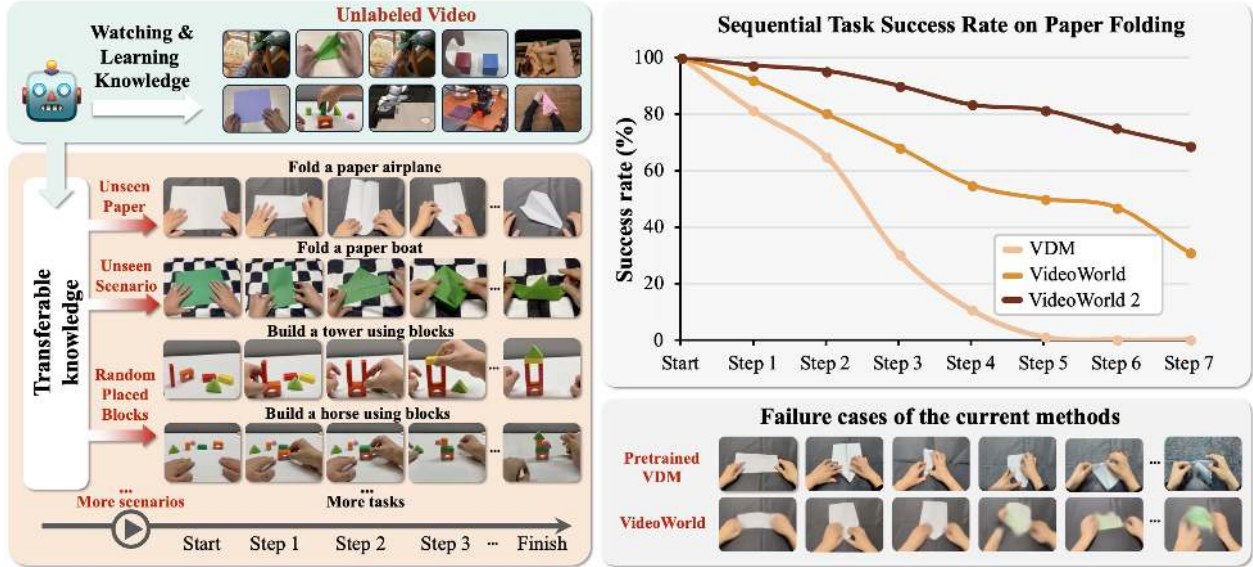


Figure 1 (left) VideoWorld 2 explores how to learn transferable knowledge from unlabeled real-world videos. We construct a handicraft benchmark to evaluate the learned knowledge. (right) Comparison of different frameworks in success rate for long-horizon paper folding tasks in Video-CraftBench. We split the task into seven key steps and evaluate sequential success rates (detailed in Sec. 4). VDM (e.g., Wan2.2 14B [59]) produces high visual fidelity but fails to learn task-relevant dynamics or long-horizon policies. VideoWorld [51] improves policy learning but suffers from poor visual quality in real-world scenarios. The bottom-right figure presents the failure cases of these baseline methods. VideoWorld 2 learns more robust latent dynamics while also achieving significantly better visual quality, enabling generalizable long-horizon knowledge learning from videos.

digital environments.

VideoWorld [51] is among the first works to explore learning knowledge from synthetic videos. It investigates the acquisition of rules, as well as reasoning and planning capabilities, from Go game records and simulated robotics environments. The study demonstrates that models can learn such knowledge solely from visual signals using an autoregressive video generation paradigm. However, extending this paradigm beyond synthetic domains remains an open challenge. Real-world videos exhibit substantial visual diversity, complex action dynamics, and often involve long-horizon, multi-step interactions. These characteristics prevent the training approach and model design of VideoWorld from being directly applied to realistic settings. When presented with minute-long, multi-step real-world task videos, VideoWorld fails to extract the core task-solving knowledge or generalize it to novel scenarios through observation alone—even for tasks such as paper folding that are easily mastered by children (See Sec. 5.3 for details). These limitations naturally lead to the following question: **Can AI models learn transferable knowledge for complex, long-horizon tasks directly from unlabeled real-world videos?**

In this work, we extend VideoWorld to real-world settings. We first establish a Video-Craft benchmark to systematically evaluate the capabilities of existing models across two representative real-world environments. As illustrated in Fig. 1, the first environment focuses on **handicraft making**, which serves as a challenging testbed for learning task knowledge directly from raw videos. These videos involve fine-grained manipulation under diverse environmental appearances, including deformable materials (e.g., paper), viewpoint variation, and frequent occlusions between hands and objects. Moreover, the videos are minute-long and consist of multiple interaction steps, presenting substantially greater complexity and longer temporal horizons than typical entertainment-oriented videos. Second, we study **robotic manipulation** by learning from the Open-X dataset [48], which contains real-world demonstration videos, and evaluating on the CALVIN environment [45]. Since CALVIN features tasks and visual settings that differ from those in Open-X, this setup enables a

¹Following prior work [51], we use “knowledge” to refer broadly to the rules, reasoning, and planning abilities required for task completion.

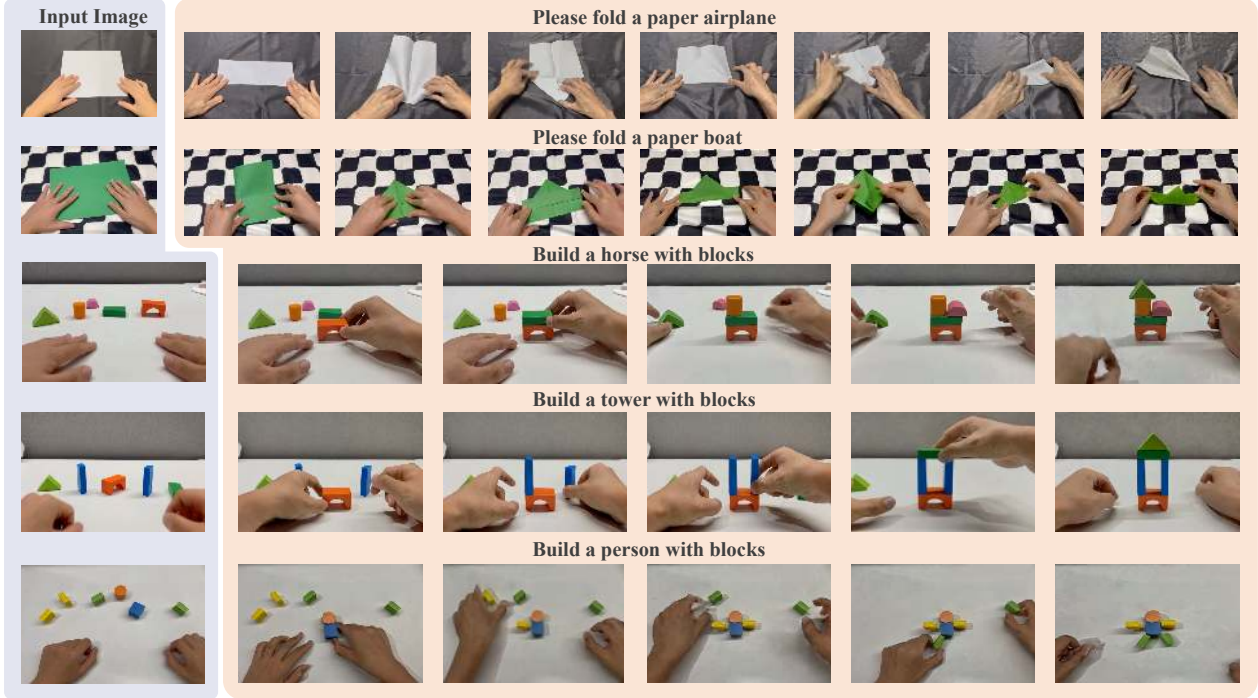


Figure 2 Qualitative Results. VideoWorld 2 learns transferable knowledge and generates long-horizon videos in unseen environments. This figure shows the output on long-horizon handicraft tasks.

rigorous assessment of the generalization capability of the learned knowledge. Crucially, learning from these benchmark videos requires models to extract task-critical actions from agent motion dynamics and to derive robust policies that can reason and plan correct action sequences across diverse environments.

To address this benchmark, we first apply VideoWorld directly to real-world videos. Compared to its original training data, these videos exhibit substantially more complex motion dynamics and visual appearances. We observe that VideoWorld struggles to extract task-critical actions and to learn policies generalizable to novel environments. As shown in Fig. 1 (bottom left), the model’s predictions contain severe errors, including distorted hand poses, incorrect object shapes, and inconsistent environmental appearances, ultimately failing to produce accurate and visually coherent action sequences. Moreover, although state-of-the-art video generation models [1, 36, 59] can generate high-fidelity visual details, they similarly fail to faithfully represent task execution. We conjecture that these models are unable to disentangle task-core actions embedded in visual changes and motion dynamics, and instead overfit to irrelevant visual details. This overfitting leads to degraded task performance and unstable long-horizon generation. In contrast, humans naturally prioritize essential actions while filtering out extraneous variations.

Inspired by these observations, we propose **VideoWorld 2**, which aims to robustly acquire knowledge by explicitly decoupling appearance modeling from action learning. The key innovation over VideoWorld is the incorporation of prior appearance knowledge during action extraction, which prevents irrelevant appearance information from being entangled with the learned action representations. We implement this idea through a novel **dynamics-enhanced Latent Dynamics Model (dLDM)**. The dLDM comprises a causal VQ-VAE and a pretrained Video Diffusion Model (VDM). The VQ-VAE compresses future visual changes into discrete latent codes that capture task-relevant actions, while the VDM is responsible for modeling visual appearance, generating high-fidelity reconstructions of task sequences conditioned on these codes. By offloading appearance modeling to the VDM, VideoWorld 2 encourages the latent action codes to focus on concise, semantically meaningful, and transferable actions rather than superficial appearance details. As a result, the learned action policies are significantly more robust and generalizable than those produced by VideoWorld. As illustrated by the qualitative results in Fig. 2, VideoWorld 2 successfully executes long-horizon handicraft making tasks in unseen environments, producing coherent task execution videos.

Beyond these qualitative demonstrations, we further compare VideoWorld 2 with prior methods on Video-CraftBench. As shown in Fig. 1 (right), VideoWorld 2 excels at long-horizon tasks, outperforming competing methods by learning more generalizable action representations and achieving significantly higher visual quality. In robotic, pre-training on Open-X further enables VideoWorld 2 to acquire effective manipulation knowledge. This leads to substantial improvements in action prediction performance on the CALVIN environment, demonstrating that the learned actions and policies transfer robustly across diverse robotic tasks.

Our contributions are summarized as follows:

- We are the first to investigate learning transferable knowledge for complex, long-horizon tasks directly from raw real-world videos. Our findings show that disentangling action dynamics from visual appearance modeling is critical for effective knowledge acquisition.
- We propose the dynamics-enhanced Latent Dynamics Model (dLDM), which decouples task-relevant dynamics from visual appearance, substantially improving the quality and transferability of learned knowledge for both handicraft and robotic manipulation tasks.
- We introduce Video-CraftBench, a new benchmark designed to address the largely unexplored challenge of fine-grained, long-horizon visual reasoning in real-world handicraft tasks. This benchmark provides a foundation for future research on learning transferable knowledge directly from raw videos.

2 Related Work

2.1 Video Generation

Video generation research is dominated by two main paradigms. First, diffusion-based models set the standard for photorealism, with models like Sora [46], Veo [20], HunyuanVideo [36], Wan [59], and CogVideoX [71] generating high-fidelity videos. Second, autoregressive (AR) models, inspired by LLMs [81], excel at causal sequence modeling. Works like Lumos-1 [77], VideoPoet [35], and NOVA [21] frame video generation as a next-token prediction task similar to text. Some leading platforms like Cosmos [1] explore both diffusion and AR branches for comprehensive generation, supporting downstream tasks such as planning [23] and real-time interactivity [71, 74]. VideoWorld 2 combines the strengths of both paradigms. It leverages the appearance priors of diffusion models to *disentangle* core actions from visual changes, thereby enabling autoregressive modeling of long-horizon policies based on these dynamics.

Furthermore, while some video generation models also explore “disentanglement”, VideoWorld 2 differs substantially in both objective and methodology. First, in prior works [9, 38, 39, 44, 53, 61, 68], disentanglement typically refers to separating motion from appearance for applications like style transfer or visual editing (e.g., camera movement or object-specific changes). In contrast, VideoWorld 2 targets a more demanding objective: reducing task-irrelevant information to learn transferable visual dynamics for complex long-horizon tasks. Second, technically, many works [9, 38, 39] rely on explicit geometric supervision to isolate motion signals. Others [53, 61] capture only coarse, global motion semantics rather than temporal dynamics, while methods such as [44, 68] depend on handcrafted residual encoding or manually separated parameter groups. In contrast, VideoWorld 2 employs a dynamics-enhanced latent dynamics model (dLDM) that suppresses appearance variation and captures task-relevant visual dynamics directly from unlabeled videos, yielding representations suitable for complex task execution, which is beyond the capability of prior approaches.

2.2 World Models

World models, which aim to learn physical dynamics for controllable simulation [22, 84], are approached from several perspectives. The video generation community [19, 30, 32, 49, 54, 55, 75] treats them as controllable synthesizers, emphasizing fidelity and consistency. The reinforcement learning/robotics communities [4, 12, 27–29, 37, 43, 50, 65] view them as learned dynamics simulators for sample-efficient planning. JEPA-style approaches [3, 7, 8, 25, 34, 41, 82, 83] avoid pixel-level reconstruction, instead forecasting in an abstract space to benefit downstream tasks. While these perspectives primarily focus on learning short-term world dynamics for synthesis or planning, VideoWorld 2 addresses the distinct challenge of learning transferable knowledge for complex, long-horizon tasks directly from unlabeled real-world videos.

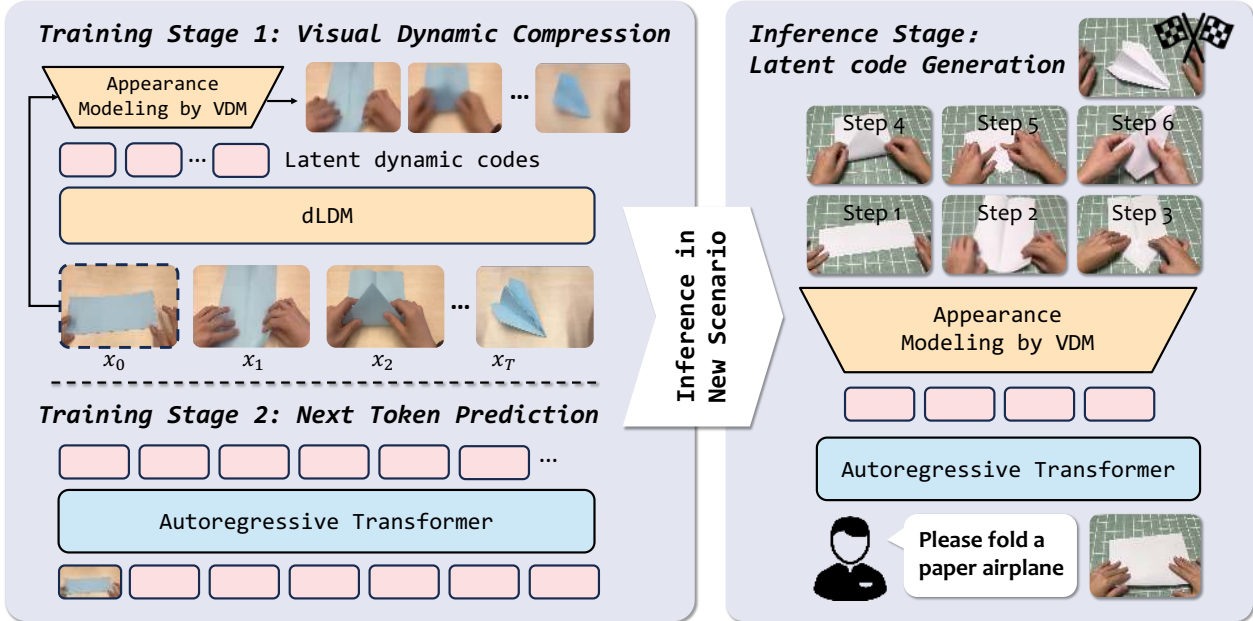


Figure 3 Overview of the VideoWorld 2 model architecture. (Left) First, the dLDM compresses future visual changes into compact and generalizable latent codes. These codes are then modeled by an autoregressive transformer. (Right) In inference, the transformer predicts latent codes for a new, unseen environment from the input image, which are subsequently decoded into task execution videos.

2.3 Learning from unlabeled videos

A key challenge in learning from unlabeled video is extracting meaningful and transferable representations of visual dynamics. While some methods [69] utilize content-similar video pairs to extract general representations, such paired data is scarce. Some recent works explore fully unsupervised learning of implicit latent actions via strategies like forward-inverse cycle consistency [52, 73], VQ-VAE quantization [15, 72], or future frame prediction [11, 13, 14, 24, 56]. VideoWorld 2 distinguishes itself in several ways. **First**, our dLDM models long-horizon visual dynamics, whereas existing latent action models typically focus on short-horizon or pairwise transitions. Moreover, VideoWorld 2 incorporates a pre-trained video generation model to guide the learning of more transferable latent dynamics, while other approaches rely on standard reconstruction decoders. **Second**, VideoWorld 2 targets minute-long tasks like handicraft involving complex visual dynamics, multiple stages, and substantial appearance variation. This setting is far more challenging than the short-horizon tasks addressed by existing LAMs, such as grasping, toggling, or 2D games [17, 40, 42].

Regarding VideoWorld [51], it struggles to decouple task-relevant dynamics from visual appearance, leading to appearance drift and motion errors in unseen environments (see Fig. 8). By delegating appearance modeling to the pre-trained VDM, VideoWorld 2 allows its latent space to focus on concise and transferable dynamics rather than appearance details, substantially improving robustness and transferability

3 Approach

3.1 Learning Knowledge from Unlabeled Videos

Generative knowledge learning. A video can be viewed as a demonstration trajectory that captures world state transitions and the underlying action policy, which constitute the knowledge to be learned. Accordingly, we propose to use a generative model to capture these latent policies and action dynamics directly from video, without reliance on language supervision. Formally, we define this setting as a tuple $\mathcal{G} = \langle \mathcal{X}, \mathcal{A}, \rho \rangle$, where \mathcal{X} denotes the observation space, \mathcal{A} is the action space, and ρ is a video generator. Given a sequence of video frames $x \in \mathcal{X}$, our goal is to train ρ to model the conditional distribution of the next frame x_{t+1} given the observation history $x_{0:t}$. This formulation enables the acquisition of task knowledge through video generation,

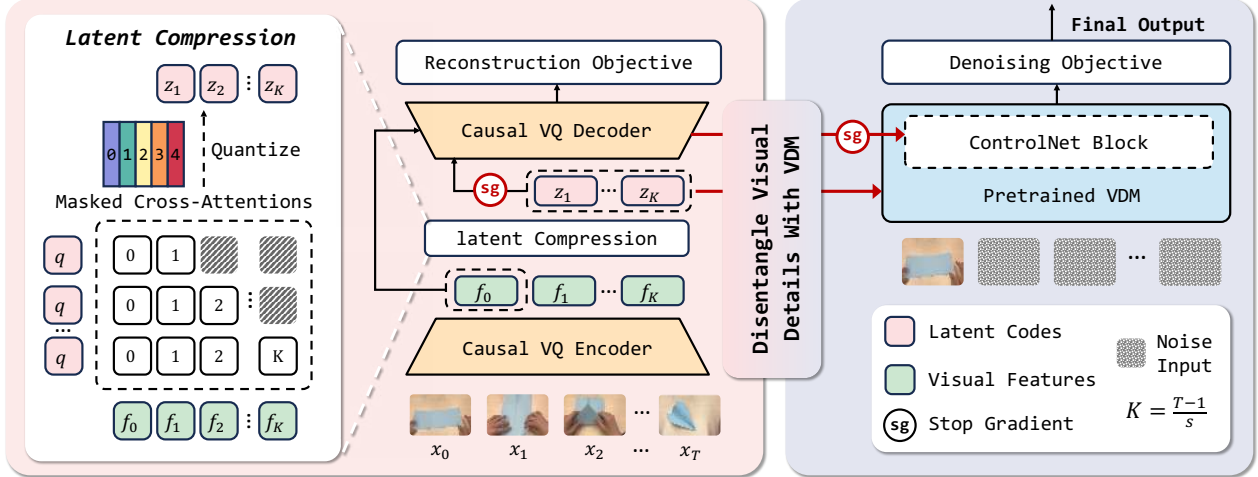


Figure 4 The proposed dynamics-enhanced latent dynamics model (dLDM). (Left) Latent dynamic model in VideoWorld [51]. Visual changes between the first and subsequent frames are compressed into a set of latent codes. (right) The dLDM proposed in VideoWorld 2. It employs a pre-trained VDM as an appearance prior, yielding better latent codes and facilitating high-fidelity video output.

eliminating the need for explicit labels.

To transfer the learned knowledge into executable actions for downstream tasks, ρ maps the visual state transitions into action space and thus serve as a policy model $\pi(\cdot|x_{0:t}) : \mathcal{X} \rightarrow \mathcal{A}$, which predicts actions based on historical observations and executes the tasks accordingly.

Basic generative framework. We adopt mainstream video generation models [1, 36, 59] as the basic framework. They typically utilize a VQ-VAE to encode videos into a compressed representation. The generative process operates within this space to predict future state, which are subsequently decoded back to RGB domain. However, these representations require thousands of discrete tokens or continuous embeddings to capture the full spectrum of visual information, inevitably leading to spatiotemporal redundancy and a sparse distribution of knowledge. As demonstrated in VideoWorld [51], such representations result in an inefficient encoding of the visual changes and motion dynamics tied to critical decisions and actions, thereby hindering the framework from acquiring essential task knowledge from raw videos.

Latent dynamic model. To mitigate the above limitation and facilitate the learning of knowledge, VideoWorld introduces a Latent Dynamic Model (LDM) that compresses future visual changes into a set of compact latent codes, effectively capturing the motion dynamics of multi-step actions.

Specifically, as shown in the left panel of Fig 4, the LDM utilizes a MAGVITv2-style [76] causal codec. The encoder first maps an input clip x of length T to a feature sequence $f_{0:K}$, where $K = 1 + \lfloor \frac{T-1}{s} \rfloor$ and s is the temporal downsampling stride. Next, it defines N learnable query embeddings $q = \{q^n\}_{n=1}^N$. As shown in Fig. 4 (left), these queries use cross-attention to capture change information in $\{f_{0:k}\}_{k=1}^K$, yielding a continuous representation $z = \{z_k\}_{k=1, n=1}^{K, N}$. This representation is then quantized to prevent the LDM from learning shortcuts (e.g., trivially copying f_k to z_k). Finally, the decoder uses f_0 and the quantized z to reconstruct the subsequent frames in a causal manner. The training objective is to minimize the ℓ_2 distance between the original and reconstructed frames. We refer to these embeddings as latent dynamics codes.

3.2 Dynamics-enhanced Latent Dynamic Model

While VideoWorld effectively handles long-term reasoning in synthetic environments by encapsulating temporal dynamics into compact latent codes, it still struggles with real-world application. As shown in the bottom of Fig. 8, the model fails to generate coherent long-horizon sequences or generalize to new environments. Specifically, applying the latent dynamics learned from handcrafting training videos to novel environments with different desktops, paper materials and arms, frequently results in significant scene drift and erroneous

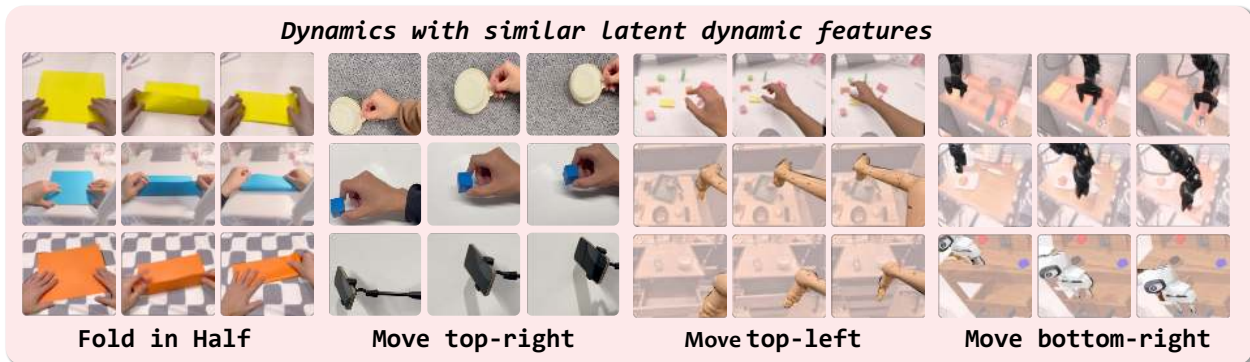


Figure 5 Video clips with similar latent dynamic features. The text below represents the dynamic type.

actions. We also evaluate several state-of-the-art video generation models [1, 36, 59]. As shown in Sec. 5.2 and Fig. 8, although these models can produce high-fidelity frames, they also fail to reproduce fine-grained, long-horizon skills in new contexts, even when detailed language instructions are provided.

Entanglement of dynamics and appearance. We conjecture that these limitations stem from the insufficient disentanglement of action dynamics and visual appearance. In the basic framework, the joint modeling of these components impedes the effective extraction of task-core action knowledge. Regarding VideoWorld, we find its learn latent codes capture irrelevant visual details—such as background motion, lighting changes, textures, and camera displacement—making the model sensitive to environmental changes, as further discussed in Sec. 5.5. Consequently, this incomplete separation ultimately leads to poor consistency in long-horizon tasks and limits the model’s ability to generalize to new environments.

Dynamics-enhanced latent dynamic model. To address this issue, we introduce a dynamics-enhanced Latent Dynamic Model (dLDM). The key mechanism is to replace the original LDM decoder with a pretrained Video Diffusion Model (VDM). By leveraging the capacity of VDM for high-fidelity reconstruction to handle the appearance modeling, we allow the LDM encoder and learnable query embeddings to focus exclusively on capturing task-relevant visual changes. Although the VDM does not contain any knowledge of the target task’s dynamics, it is highly effective at producing realistic visual content once given appropriate dynamics guidance, making it well-suited for this decomposition (As shown in Fig. 2 and Fig. 7).

Concretely, as shown in the right panel of Fig. 4, the dLDM consists of a causal VQ-VAE that encodes future visual changes into discrete latent codes and a pretrained VDM that reconstructs high-fidelity frames. The latent codes are provided to the VDM through a projection layer and causal cross-attention. Since the VDM handles appearance modeling, the latent codes are relieved from encoding fine-grained visual details and can instead focus on capturing task-relevant dynamics. To maintain temporal correctness, we enforce causal attention in the VDM so that features at time t attend only to information up to time t .

Directly training the VDM to generate future frames from noise would be extremely slow and prone to incorrect motion, as it has never been trained on the target tasks like long-horizon handicraft making. Therefore, we reuse the VQ-VAE decoder to reconstruct latent codes into low-fidelity, motion-rich outputs, providing coarse temporal cues such as hand movements and object displacements. This signal is fed into the VDM via a gradient-stopped, ControlNet-like [78] branch. This stabilizes training and allows the VDM to focus on refining appearance rather than inferring motion from scratch. Additionally, we stop the gradient flow of the decoder to the latent codes to prevent the introduction of irrelevant noise. We ablate this design in Sec. 5.5.

By leveraging the VDM, the learned latent codes are much less affected by changes in appearance and stay consistent across environments. As shown in Sec. 7, latent codes corresponding to the similar action exhibit tighter intra-class alignment and reduced cross-environment variance, indicating robust and transferable dynamics.

Auto-regressive transformer. After extracting the latent codes, we use an autoregressive transformer to model the latent dynamics sequence. For each video $x_{0:T}$, the dLDM extracts latent codes $\{z_k^n\}_{k=1, n=1}^{K, N}$. We

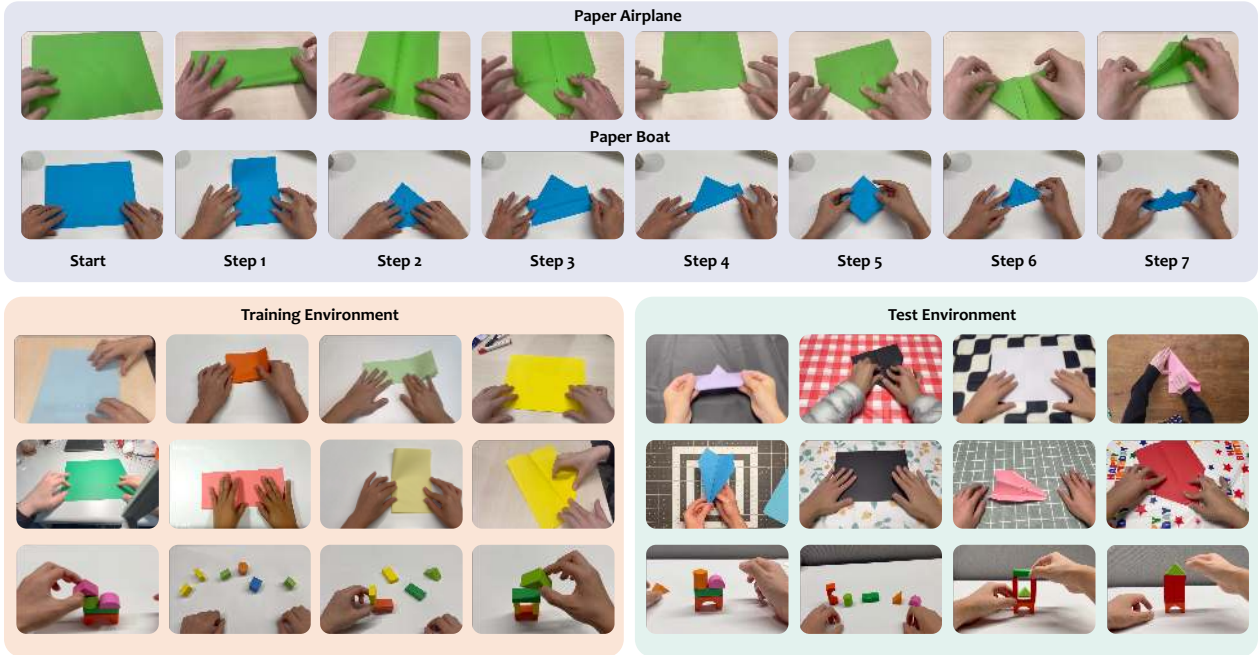


Figure 6 Overview of Video-CraftBench and keys steps of paper folding tasks. Best viewed in color.

flatten these into a sequence and train the transformer to predict them, conditioned on the initial frame x_0 and the task instruction. This allows the model to learn the long-term patterns in complex tasks.

During inference (Fig. 5, right), given a single input frame from a new, unseen environment, the transformer predicts future latent dynamics based on its learned task representation, and the dLDM decodes them into coherent long-horizon execution videos. This allows VideoWorld 2 to transfer its learned dynamics to new environments and execute extended action sequences beyond those observed during training.

4 Video-CraftBench

Our objective is to learn transferable knowledge for complex and long-horizon tasks from real-world videos. We prioritize tasks that demand multi-step planning and feature delicate manipulations. To this end, we introduce the Video-CraftBench, a dataset of first-person video tutorials covering five long-horizon handcraft tasks: folding a paper airplane, folding a paper boat, and building a tower/horse/person using blocks.

4.1 Dataset Generation

These handcraft tasks feature fine-grained, visually diverse manipulations that are difficult to articulate with language. We collected ~ 7 hours (split into $\sim 9.5k$ clips) of tutorial videos via manual recording and internet sources, enabling long-horizon knowledge evaluation. The paper folding tasks typically last 40-80 seconds and block-building tasks last 20-30 seconds. The benchmark is split into training and test sets. The test set (~ 150 separately collected videos) features novel backgrounds, paper textures, and block arrangements not seen in training to assess knowledge transferability. Supp. B provides more details.

4.2 Evaluation

Our evaluation focuses on two main aspects: the accuracy of key actions and the visual quality of the generated videos. Accordingly, we use the following metrics for evaluation:

- **Sequential task success rate:** We decompose paper folding into 7 key steps (Fig. 6) and train a DINOv2-based classifier [47] to detect their completion. The classifier is trained to evaluate only action correctness, disregarding appearance consistency (e.g., appearance drift from the initial frame). To

train it, we sample and annotate frames from the training set, test set, and manually-verified generated videos, labeling each as a specific step or “failed”. For evaluation, we generate 3 video rollouts per test case conditioned on the first frame. A step is only successful if all preceding steps are complete, allowing us to assess long-horizon knowledge acquisition. For the shorter block-building tasks, a similar classifier is trained, but only to verify the final generated state.

- **Visual quality:** We also assess visual quality in the test environment, using standard LPIPS [79] and SSIM [64] metrics to quantify visual fidelity and content consistency with the input.

5 Experiments

5.1 Implementation Details

We employ the NVIDIA Cosmos AR 4B model [1] as our auto-regressive transformer, repurposing its next-token prediction capability to predict our latent codes. For the appearance prior, we utilize the Cosmos DiT 2B model [1], leveraging its high-fidelity I2V capabilities to generate 93-frame (~ 5 s at 16 fps), 480px videos. The dLDM also processes a 93-frame clip at a time by default, using a vocabulary size of 1000 (FSQ levels [8, 5, 5, 5]) and an embedding length $N = 4$. To improve training efficiency, dLDM first applies a short warm-up where the latent codes are optimized solely using the original reconstruction objective. This warm-up is similar to the training strategy of the original VideoWorld LDM, enabling the latent codes to rapidly learn to compress visual changes and motion dynamics, while allowing the decoder to reconstruct low-fidelity video clips containing agent motion trajectories based on the initial frame and codes. Consequently, when we switch to the disentangled scheme, the decoder can provide robust motion conditioning to stabilize training. The trainable parameters in our framework include the AR transformer, the dLDM autoencoder, the DiT, and the projection layer mapping the latent codes to the DiT.

5.2 Benchmarks and Baselines

Benchmarks. We evaluate VideoWorld 2 on two benchmarks: Video-CraftBench assesses learning from long-horizon, real-world tasks with fine-grained actions. For robotics, we train our latent representation on the large-scale OpenX dataset and test its knowledge transfer to CALVIN. CALVIN features 34 tasks and uses a long-horizon, sequential evaluation protocol similar to our paper folding setup: models must complete a 5-task sequence, where success is contingent on all preceding tasks.

Baselines. To demonstrate the effectiveness of VideoWorld 2, we establish three strong baselines for comparative analysis on the aforementioned benchmarks:

- Pre-trained video generation model: We select four state-of-the-art video generation models: NVIDIA Cosmos AR 4B [1], NVIDIA Cosmos DiT 2B [1], Wan2.2 14B [59], and HunyuanVideo 13B [36]. We fine-tune them on Video-CraftBench. To better align with their original training, we also provide detailed language annotations for each key step using Qwen2.5-VL 72B [6] to aid generation. During inference, these models generate task execution videos conditioned on an initial image and text instructions. They then generate subsequent video clips autoregressively, based on their own previous output.
- Latent action models: We compare our approach with five concurrent works [15, 24, 62, 67, 72] that leverage latent action models. These methods compress inter-frame visual changes into latent codes, serving as a pre-training objective for manipulation tasks while allowing for video reconstruction. Consequently, we apply their latent extraction methodologies to our benchmark to assess their ability to extract knowledge from long-horizon real-world videos.
- VideoWorld [51]: This is the most relevant baseline to our work. It uses the original latent dynamic model for video knowledge learning.

5.3 Results on Video-CraftBench

Training with Video-CraftBench only. In Tab. 1, we first evaluate pre-trained video generation models fine-tuned on Video-CraftBench (row 1-4). While these models achieve high success rates ($>68\%$) in the first

Method	Fine-tuning	Sequential Paper folding Success Rate (%)							Block Success Rate (%)			Visual Quality	
		1	2	3	4	5	6	7	Human	Tower	Horse	SSIM↑	LPIPS↓
<i>Pre-trained Video Generation Model</i>													
Cosmos AR 4B [1]	Craft-text	68.4	56.7	11.5	3.3	0.0	0.0	0.0	10.1	18.0	12.0	0.643	0.312
Cosmos DiT 2B [1]	Craft-text	73.4	63.3	20.0	6.7	0.0	0.0	0.0	24.2	21.3	19.7	0.680	0.264
Hunyuan-13B [36]	Craft-text	76.9	68.1	27.5	5.8	0.0	0.0	0.0	30.9	38.4	31.5	0.703	0.255
Wan 2.2 14B [59]	Craft-text	81.2	75.0	30.4	10.6	0.0	0.0	0.0	39.7	42.6	34.1	0.719	0.237
<i>Learning with Latent Action/Dynamic Model</i>													
LAPA [72]	Craft				N.A.					N.A.			N.A.
Moto [15]	Craft	19.1	11.7	3.3	0.0	0.0	0.0	0.0	11.5	10.1	9.8	0.585	0.394
AdaWorld [72]	Craft	43.6	39.8	27.4	10.8	0.0	0.0	0.0	20.7	13.1	15.0	0.611	0.378
VideoWorld [51]	Craft	<u>70.3</u>	<u>66.7</u>	<u>42.5</u>	<u>21.3</u>	<u>6.7</u>	0.0	0.0	<u>23.8</u>	<u>33.9</u>	<u>27.8</u>	<u>0.680</u>	<u>0.351</u>
VideoWorld 2	Craft	97.2	95.3	90.0	83.3	81.4	74.6	68.8	70.0	81.5	80.9	0.770	0.205
iVideoGPT [67]	OpenX & Craft	23.1	18.7	13.3	3.7	0.0	0.0	0.0	15.3	11.0	12.6	0.588	0.390
Moto [15]	OpenX & Craft	43.1	35.3	30.7	25.5	18.3	9.7	0.0	17.4	15.3	16.0	0.596	0.387
AdaWorld [72]	OpenX & Craft	49.5	41.6	34.8	30.7	22.3	19.8	13.0	37.4	29.8	29.1	0.624	0.365
CoLA [62]	OpenX & Craft	83.5	74.4	<u>69.1</u>	<u>64.8</u>	<u>52.3</u>	<u>49.8</u>	<u>40.2</u>	<u>54.1</u>	52.4	<u>49.9</u>	<u>0.668</u>	<u>0.289</u>
VideoWorld [51]	OpenX & Craft	<u>91.7</u>	<u>75.0</u>	68.2	63.1	51.7	48.2	31.9	47.3	<u>52.7</u>	49.8	0.601	0.389
VideoWorld 2	OpenX & Craft	98.2	96.4	90.1	86.7	83.3	81.7	72.3	74.0	83.0	85.8	0.774	0.193

Table 1 Comparison on Video-CraftBench. Craft-text denotes using step-by-step textual description.

step of the paper folding task and up to 38% in the block stacking task, their performance deteriorates rapidly for complete, minute-long sequences. By step 4 of paper folding (Column 7), success rates drop to $\leq 10.6\%$, with a total failure to generate subsequent steps. Crucially, this failure persists even though we have provided detailed textual descriptions for each step during training and inference to leverage their text-conditioning capabilities. This demonstrates that current models struggle to learn complex, long-horizon knowledge from real-world videos, underscoring the need for a more effective paradigm.

Subsequently, we evaluate existing methods that leverage latent action or dynamic models. Following their protocols, we train these baselines on Video-CraftBench and train our AR transformer to predict their latent codes. Note that due to LAPA [72]’s structural constraints, decoding its codes into long-horizon sequences causes severe degradation. Thus, we mark its results as “N.A.”. For the other three baselines, Moto [15] employs a pre-trained vision encoder to extract dynamics, AdaWorld [24] utilizes an auxiliary diffusion head for video synthesis, and VideoWorld [51] features capabilities for long-horizon dynamics compression and generation. Although they generate coherent videos, they fail to generalize to novel environments. Specifically, none successfully completed the full paper folding sequence under new settings (e.g., different desktops or paper styles). As shown in Fig. 8, their predictions exhibit substantial action errors and scene distortions, indicating that their latent codes overfit to irrelevant appearance information.

In contrast, VideoWorld 2 (row 9) generates complete and coherent task sequences in test environments. Remarkably, without requiring large-scale pre-training for these latent codes, training solely on Video-CraftBench achieves a success rate of 68.8% on the paper folding task and up to 81.5% on the block stacking task. This demonstrates that our dLDM efficiently extracts core task-relevant action information while filtering out extraneous details like background clutter, object variations, and camera noise. Consequently, it generalizes effectively to unseen environments. Furthermore, benefiting from the appearance priors of the VDM, VideoWorld 2 not only generates accurate actions but also produces videos with higher SSIM and PSNR metrics. Fig. 2 and Supp. C visualize our generated sequences.

Data scaling with OpenX. Next, we evaluate whether the latent codes of VideoWorld 2 can benefit from larger-scale data. To this end, we incorporate the OpenX dataset [48], which contains a vast collection of manipulation demonstrations. Compared to Video-CraftBench, OpenX features a diverse array of robotic agents and environmental appearances, making it an ideal testbed for assessing VideoWorld 2’s ability to filter out agent-specific and environmental factors while scaling its action extraction capabilities.

Specifically, we first train dLDM and baseline models on a combination of OpenX and Video-CraftBench, followed by training the AR transformer exclusively on Video-CraftBench using the resulting latent codes. As shown in line 10-12 of Tab. 1, the baselines show improved quality, raising success rates across tasks. However, Moto and iVideoGPT still fail to complete the full folding sequence. While VideoWorld can generate complete

Idx	Method	Pretraining Type	Pretraining	Fine-tuning	Sequential Task Success Rate (%)					
					1	2	3	4	5	Avg. Len.
<i>In-Domain Latent Pretraining</i>										
1	Transformer (Oracle)	-	-	ABCD→D	80.9	55.6	44.5	31.3	24.6	2.36
2	Transformer (Oracle)	-	-	10% data	50.5	35.4	20.1	5.20	0	1.11
3	LAPA [72]	Latent	ABCD→D	10% data	74.4	45.8	25.2	15.3	2.30	1.49
4	VideoWorld 2	Latent	ABCD→D	10% data	75.8	47.9	31.8	20.4	9.70	1.87
<i>Cross-Domain Latent Pretraining</i>										
5	Transformer (Oracle)	Video	OpenX	ABCD→D	85.9	<u>60.4</u>	46.0	30.7	23.0	2.46
6	LAPA [72]	Latent	OpenX	ABCD→D	84.0	58.8	<u>46.2</u>	<u>35.4</u>	<u>27.0</u>	<u>2.51</u>
7	VideoWorld 2	Latent	OpenX	ABCD→D	<u>88.5</u>	64.6	55.8	47.5	30.9	2.88

Table 2 Comparison on CALVIN benchmark. “10% data” means 2k trajectories randomly sampled from the ABCD→D training set. “Oracle” means directly training the transformer using ground-truth action label as supervision.

sequences, its final-step success rate remains low at 31.9%. We also examine CoLA [62], a concurrent work that also uses a VDM to optimize latent action codes. However, it is limited to short 2-frame transitions and ignores the structured temporal cues from coarse VAE outputs. Adopting CoLA’s training scheme yields a limited final-step success rate of 40.2%. In contrast, VideoWorld 2 further improves upon its strong baseline performance, achieving a final-step success rate of 72.3% on paper folding and up to 85.8% on block stacking. As shown in the right panel of Fig. 5, similar latent codes in OpenX and Video-CraftBench correspond to videos with similar motion patterns despite differing environments and agents, demonstrating our transferability across diverse agents and settings.

5.4 Results on CALVIN

In-domain latent pre-training on CALVIN. To demonstrate that our latent codes aid in learning manipulation skills, we conduct a latent pre-training experiment, following LAPA. We first pre-train the AR transformer on latent codes from 22k CALVIN trajectories, then fine-tune it on only 2k ground-truth action labels. As shown in Tab. 2, this pre-training strategy significantly boosts VideoWorld 2’s performance (Idx 4) on long-horizon tasks compared to the baseline (Idx 2). Notably, its performance approaches the model trained on the 22k full action label dataset (Idx 1), demonstrating high data efficiency. While LAPA (Idx 3) also benefits from this protocol, it still lags behind VideoWorld 2, highlighting the effectiveness of our latent codes.

Cross-domain latent pre-training on CALVIN. We further evaluate transferability by pre-training on the 1.3M OpenX dataset and fine-tuning on 22k CALVIN trajectories. As shown in Tab. 2, this pre-training significantly boosts success rates for LAPA (Idx 6) and VideoWorld 2 (Idx 7) compared to training only on labeled CALVIN data (Idx 1). Consistent with previous findings, VideoWorld 2 benefits more from this pre-training on long-horizon tasks. We also include a video next-token prediction baseline pre-trained on OpenX (Idx 5). The results show that our latent pre-training is more effective, validating it as a more efficient knowledge transfer paradigm than pre-training directly on videos. Supp. A provides more training details.

5.5 Ablation Study

Code transferability is enhanced by VDM. In Tab. 3a, we first ablate the introduction of the VDM prior (line 1 vs. 2). We find that incorporating the VDM enhances codes transferability and visual quality, yielding a ~30% increase in success rate and an improved LPIPS score. This result validates the effectiveness of our appearance-dynamic decoupling design. To provide a more intuitive illustration, we visualize the latent codes using UMAP in Fig. 7. We randomly sample 4000 trajectories from CALVIN and Bridge (included in OpenX) [58] and label each by its robot arm action (up, down, left, or right). As shown in Fig. 7 (left), with VDM, latent codes for the same action align much more tightly across environments, indicating more consistent and transferable dynamics. In contrast, without VDM (Fig. 7, right), the codes separate strongly by environment, revealing a loss of cross-environment consistency.

Effect of the original decoder. In Sec.3, we first warm up the dLDM with the original decoder during training. In subsequent stages, we discard the reconstruction loss from the VQ-VAE decoder to prevent noise injection.

Pre-trained VDM	Decoder Stop-Grad	Ctrl-Net	Video-CraftBench			Query length N	Video-CraftBench			CALVIN Avg.Len.	Proj Layer	Interplay Attn type	Paper Block	
			Paper	Block	LPIPS \downarrow		Paper	Block	LPIPS \downarrow				Paper	Block
			0.0	28.5	0.312	baseline	0.0	28.5	0.312	1.11	MLP cross	52.0	61.3	
✓			30.3	45.2	0.297	1	41.9	65.0	0.221	1.53	+self cross	52.3	61.8	
✓	✓		47.3	54.7	0.275	2	55.1	69.5	0.210	1.64	MLP causal cross	69.8	78.6	
✓		✓	<u>51.1</u>	52.0	<u>0.213</u>	4	68.8	77.5	<u>0.205</u>	<u>1.87</u>	+self causal cross	72.3	80.9	
✓	✓	✓	68.8	77.5	0.205	8	<u>65.0</u>	77.5	0.195	1.88				

(a) dLDM architecture.

(b) Query embedding lengths.

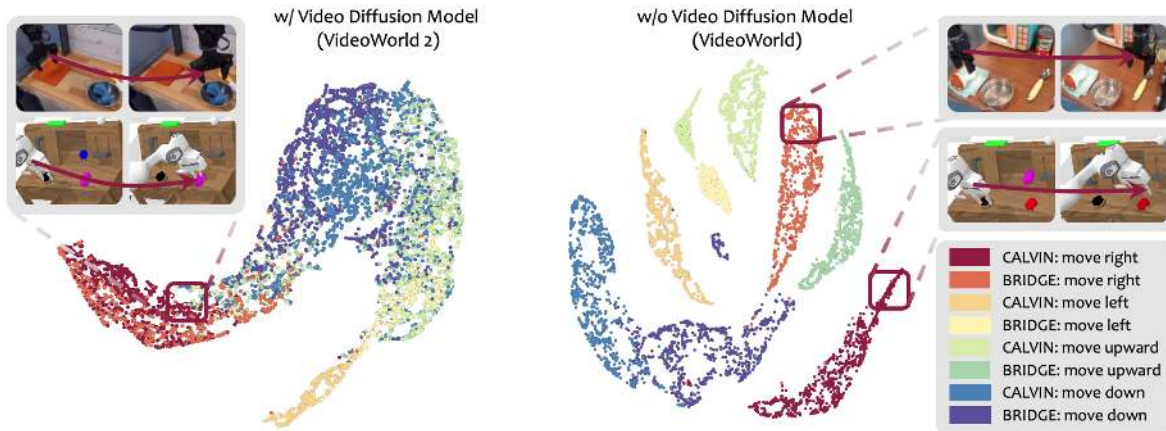
(c) LDM/VDM interplay.

Codebook size	Video-CraftBench			CALVIN Avg.Len.	Length T	Video-CraftBench			CALVIN Avg.Len.	Training Strategy	Video-CraftBench	
	Paper	Block	LPIPS \downarrow			Paper	Block	LPIPS \downarrow			Paper	Block
baseline	0.0	28.5	0.312	1.11	2	19.1	38.7	1.55	baseline	0.0	28.5	
8	20.1	29.9	0.258	1.65	9	55.4	68.7	1.61	random	0.0	0.0	
1000	68.8	77.5	0.205	1.87	49	65.3	76.2	<u>1.80</u>	freeze	31.7	40.2	
4096	<u>50.4</u>	<u>59.6</u>	<u>0.208</u>	1.90	93	<u>68.8</u>	77.5	1.87	lora	50.9	62.3	
64,000	29.4	36.0	0.230	<u>1.89</u>	177	69.0	<u>76.8</u>	1.79	full	68.8	77.5	

(d) Codebook sizes of dLDM.

(e) Compression lengths of dLDM.

(f) Training strategies of VDM.

Table 3 Ablation studies. We train with Video-CraftBench only and evaluate CALVIN in its in-domain setting.**Figure 7 UMAP of latent codes.** With a pretrained VDM, latent codes for the same action align more tightly than those in VideoWorld [51] across environments, indicating more consistent and transferable dynamics. (Left) In distinct environments like Bridge [58] and CALVIN [45], latent representations of the same action (e.g., the robotic arm moving right) are highly similar. (Right) Conversely, in VideoWorld, latent codes for the same action exhibit significant divergence across environments and fail to cluster effectively in the UMAP visualization.

However, we observe that after warm-up, the decoder can reconstruct latent codes into videos that preserve coherent object motion (e.g., hand movements and displacements), although visual details remain blurry. Consequently, we utilize these reconstructions as conditions for the VDM via a ControlNet-like structure to enhance output quality and stabilize training. Tab.3a further investigates this decoder.

First, we stop the gradient from the decoder without using the reconstructed video (row 3) to verify if this decoder introduces noise. Compared to row 2, this yields a $\sim 20\%$ improvement in success rate, suggesting that the original decoder indeed introduces extraneous noise that degrades latent representation performance. Second, we investigate the utility of the reconstructed video itself. We find that utilizing this motion conditioning (rows 4 and 5) stabilizes video output, yielding a substantial improvement of ~ 0.9 in LPIPS and up to $\sim 20\%$ in task success rate. This benefit is more pronounced in the longer paper folding task compared to block stacking, demonstrating the effectiveness of our method for long-horizon video generation.

Query embedding length of dLDM. Tab. 3b ablates the number of dLDM query embeddings N . A larger N captures more information but also risks encoding noise and increases the sequence length for the AR transformer. We find that $N = 1$ already achieves a respectable task success rate, while $N = 4$ provides the best performance balance. Increasing N to 8 offers a slight LPIPS benefit but introduces more noise, which




Figure 8 Visualization of failure cases in other baselines on the Video-CraftBench. Best viewed in color.

reduces the overall success rate.

Interplay between VDM and LDM. Tab. 3c ablates the interplay mechanism between the LDM and VDM, which consists of a projection layer (MLP and causal self-attention) and causal cross-attention. The causal cross-attention ensures generation relies solely on the current time step’s latents, effectively preventing information leakage. Results show that incorporating self-attention into the projection layer and utilizing causal cross-attention significantly enhance LDM-VDM interaction, validating the effectiveness of our design

Horizon length in dLDM. Tab. 3e ablates the dLDM context length T . At $T = 2$ (a LAPA-like setting), codes remain transferable, but performance on the long-range paper folding task is poor due to a lack of temporal perception. This minimal context also provides low-quality motion guidance for the VDM, leading to poor LPIPS scores. Performance improves as T increases, plateauing at $T = 93$, which corresponds to the maximum context length of our Cosmos VDM.

dLDM codebook size. Tab. 3d examines the impact of the dLDM codebook size (adjusted via FSQ levels) on policy learning. We find that CALVIN, with its simpler action space, benefits from a relatively small codebook. In contrast, the more complex Video-CraftBench requires a larger codebook to achieve significant improvements. However, an excessively large codebook risks encoding extraneous noise and can hinder dLDM training convergence, a finding consistent with VideoWorld [51].

Training strategies of VDM. In Tab. 3f, we examine different training configurations for the video diffusion

model. Our default approach fully fine-tunes the pre-trained VDM to exploit its appearance priors. In contrast, training a randomly initialized VDM (row 1) leads to model collapse and fails to generate valid videos. We also experiment with freezing the pre-trained VDM and updating only the VQ-VAE components (encoder, quantizer, and projection layer). Although this setup yields performance gains, it falls significantly short of LoRA or full fine-tuning. We believe this gap exists because the VDM needs further adaptation to capture the fine-grained manipulation details specific to Video-CraftBench.

6 Further Discussion with Other Works

We provide a comprehensive analysis, highlighting the differences between VideoWorld 2 and prior works in both objectives and methodology.

Regard task objectives. Most prior works focus on short-horizon prediction, simulated control, or reconstruction-centric tasks. In contrast, VideoWorld 2 targets complex, minute-long real-world tasks, where appearance variation and long-horizon error accumulation are dominant. This exposes appearance interference as a critical bottleneck, necessitating our dLDM.

Regarding “disentanglement”. Prior works [13, 51, 63, 66, 67, 72] rely on VAE-style reconstruction objectives to separate dynamics from appearance. As demonstrated in VideoWorld 2, this approach is insufficient for complex real-world videos, as reconstruction objectives compel latent codes to encode task-irrelevant details. VideoWorld 2 instead mitigates appearance interference by leveraging a pretrained VDM for appearance modeling, forcing latents to focus on task-relevant dynamics. While CoLA [62] also employs a VDM, it is limited to short 2-frame transitions. In contrast, VideoWorld 2 models multi-step dynamics and reuses coarse VAE decoder outputs to provide structured temporal cues—a mechanism critical for long-horizon tasks (Tab. 3a, rows 3 vs. 5). Replacing our dLDM with prior works’ tokenizers [67] or adopting CoLA’s training scheme significantly degrades long-sequence performance, validating our design.

Regarding the effect of the VDM. The role of the VDM in VideoWorld 2 also differs from prior approaches. Works such as ϵ -VAE [80] and DiVoT [26] utilize VDMs for pixel-level reconstruction, necessitating that latent codes encode both appearance and dynamics. Others, like IGOR [13] and AdaWorld [24], employ VDMs primarily for rendering. This disconnects the VDM from latent optimization, failing to improve task objectives. In contrast, VideoWorld 2 offloads appearance modeling to the VDM, compelling latents to focus exclusively on task-relevant dynamics. Compared to CoLA [62], VideoWorld 2 benefits from multi-frame modeling and a VDM architecture specifically tailored for long-horizon learning.

7 Conclusion

In this work, we explore knowledge learning for complex, long-horizon tasks from raw videos through experiments on Video-CraftBench and in robotic manipulation environments. We find that decoupling visual appearance from core actions is crucial for real-world knowledge learning. Based on this finding, we propose VideoWorld 2, a model featuring a dynamics-enhanced latent dynamics model (dLDM) that leverages a pre-trained VDM to learn generalizable, transferable policies directly from video. While our method shows significant potential, we leave its continued scaling to future work, moving toward the goal of enabling AI to learn the vast knowledge encapsulated in the real world.

References

- [1] Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, Yin Cui, Yifan Ding, et al. Cosmos world foundation model platform for physical ai. [CoRR](#), 2025.
- [2] Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm 2 technical report. [arXiv preprint arXiv:2305.10403](#), 2023.
- [3] Mido Assran, Adrien Bardes, David Fan, Quentin Garrido, Russell Howes, Matthew Muckley, Ammar Rizvi, Claire Roberts, Koustuv Sinha, Artem Zholus, et al. V-jepa 2: Self-supervised video models enable understanding, prediction and planning. [arXiv preprint arXiv:2506.09985](#), 2025.
- [4] Junyeob Baek, Yi-Fu Wu, Gautam Singh, and Sungjin Ahn. Dreamweaver: Learning compositional world models from pixels. [arXiv preprint arXiv:2501.14174](#), 2025.
- [5] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. [arXiv preprint arXiv:2309.16609](#), 2023.
- [6] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibong Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. [arXiv preprint arXiv:2502.13923](#), 2025.
- [7] Federico Baldassarre, Marc Szafraniec, Basile Terver, Vasil Khalidov, Francisco Massa, Yann LeCun, Patrick Labatut, Maximilian Seitzer, and Piotr Bojanowski. Back to the features: Dino as a foundation for video world models. [arXiv preprint arXiv:2507.19468](#), 2025.
- [8] Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mahmoud Assran, and Nicolas Ballas. Revisiting feature prediction for learning visual representations from video. [arXiv preprint arXiv:2404.08471](#), 2024.
- [9] Juan Luis Gonzalez Bello, Xu Yao, Alex Whelan, Kyle Olszewski, Hyeonwoo Kim, and Pablo Garrido. Videospats: Video spatiotemporal splines for disentangled occlusion, appearance and motion modeling and editing. [arXiv preprint arXiv:2504.07146](#), 2025.
- [10] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. [Advances in neural information processing systems](#), 33:1877–1901, 2020.
- [11] Jake Bruce, Michael D Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, et al. Genie: Generative interactive environments. In [Forty-first International Conference on Machine Learning](#), 2024.
- [12] Maxime Burchi and Radu Timofte. Accurate and efficient world modeling with masked latent transformers. [arXiv preprint arXiv:2507.04075](#), 2025.
- [13] Xiaoyu Chen, Junliang Guo, Tianyu He, Chuheng Zhang, Pushi Zhang, Derek Cathera Yang, Li Zhao, and Jiang Bian. Igor: Image-goal representations are the atomic control units for foundation models in embodied ai. [arXiv preprint arXiv:2411.00785](#), 2024.
- [14] Xiaoyu Chen, Hangxing Wei, Pushi Zhang, Chuheng Zhang, Kaixin Wang, Yanjiang Guo, Rushuai Yang, Yucen Wang, Xinquan Xiao, Li Zhao, et al. Villa-x: enhancing latent action modeling in vision-language-action models. [arXiv preprint arXiv:2507.23682](#), 2025.
- [15] Yi Chen, Yuying Ge, Weiliang Tang, Yizhuo Li, Yixiao Ge, Mingyu Ding, Ying Shan, and Xihui Liu. Moto: Latent motion token as the bridging language for learning robot manipulation from videos. In [ICCV](#), pages 19752–19763, 2025.
- [16] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. [Journal of Machine Learning Research](#), 24(240):1–113, 2023.
- [17] Karl Cobbe, Chris Hesse, Jacob Hilton, and John Schulman. Leveraging procedural generation to benchmark reinforcement learning. In [International conference on machine learning](#), pages 2048–2056. PMLR, 2020.

- [18] Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. [arXiv preprint arXiv:2507.06261](#), 2025.
- [19] Yufeng Cui, Honghao Chen, Haoge Deng, Xu Huang, Xinghang Li, Jirong Liu, Yang Liu, Zhuoyan Luo, Jinsheng Wang, Wenxuan Wang, et al. Emu3. 5: Native multimodal models are world learners. [arXiv preprint arXiv:2510.26583](#), 2025.
- [20] DeepMind. Veo – deepmind’s video evolution model. <https://deepmind.google/models/veo/>, 2025. Accessed: November 11, 2025.
- [21] Haoge Deng, Ting Pan, Haiwen Diao, Zhengxiong Luo, Yufeng Cui, Huchuan Lu, Shiguang Shan, Yonggang Qi, and Xinlong Wang. Autoregressive video generation without vector quantization. [arXiv preprint arXiv:2412.14169](#), 2024.
- [22] Jingtao Ding, Yunke Zhang, Yu Shang, Yuheng Zhang, Zefang Zong, Jie Feng, Yuan Yuan, Hongyuan Su, Nian Li, Nicholas Sukiennik, et al. Understanding world or predicting future? a comprehensive survey of world models. [ACM Computing Surveys](#), 58(3):1–38, 2025.
- [23] Yilun Du, Sherry Yang, Bo Dai, Hanjun Dai, Ofir Nachum, Josh Tenenbaum, Dale Schuurmans, and Pieter Abbeel. Learning universal policies via text-guided video generation. [NeurIPS](#), 36:9156–9172, 2023.
- [24] Shenyuan Gao, Siyuan Zhou, Yilun Du, Jun Zhang, and Chuang Gan. Adaworld: Learning adaptable world models with latent actions. [arXiv preprint arXiv:2503.18938](#), 2025.
- [25] Quentin Garrido, Mahmoud Assran, Nicolas Ballas, Adrien Bardes, Laurent Najman, and Yann LeCun. Learning and leveraging world models in visual representation learning. [arXiv preprint arXiv:2403.00504](#), 2024.
- [26] Yuying Ge, Yizhuo Li, Yixiao Ge, and Ying Shan. Divot: Diffusion powers video tokenizer for comprehension and generation. [arXiv preprint arXiv:2412.04432](#), 2024.
- [27] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. [arXiv preprint arXiv:1912.01603](#), 2019.
- [28] Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models. [arXiv preprint arXiv:2010.02193](#), 2020.
- [29] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. [arXiv preprint arXiv:2301.04104](#), 2023.
- [30] Danijar Hafner, Wilson Yan, and Timothy Lillicrap. Training agents inside of scalable world models. [arXiv preprint arXiv:2509.24527](#), 2025.
- [31] William Huitt and John Hummel. Piaget’s theory of cognitive development. [Educational psychology interactive](#), 3(2):1–5, 2003.
- [32] Joel Jang, Seonghyeon Ye, Zongyu Lin, Jiannan Xiang, Johan Bjorck, Yu Fang, Fengyuan Hu, Spencer Huang, Kaushil Kundalia, Yen-Chen Lin, et al. Dreamgen: Unlocking generalization in robot learning through video world models. [arXiv preprint arXiv:2505.12705](#), 2025.
- [33] Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. Mistral 7b. [arXiv preprint arXiv:2310.06825](#), 2023.
- [34] Efstathios Karypidis, Ioannis Kakogeorgiou, Spyros Gidaris, and Nikos Komodakis. Dino-foresight: Looking into the future with dino. [arXiv preprint arXiv:2412.11673](#), 2024.
- [35] Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Grant Schindler, Rachel Hornung, Vighnesh Birodkar, Jimmy Yan, Ming-Chang Chiu, et al. Videopoet: A large language model for zero-shot video generation. [arXiv preprint arXiv:2312.14125](#), 2023.
- [36] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. Hunyuanvideo: A systematic framework for large video generative models. [arXiv preprint arXiv:2412.03603](#), 2024.
- [37] Maria Krinner, Elie Aljalbout, Angel Romero, and Davide Scaramuzza. Accelerating model-based reinforcement learning with state-space world models. [arXiv preprint arXiv:2502.20168](#), 2025.

- [38] Pulkit Kumar, Namitha Padmanabhan, Luke Luo, Sai Saketh Rambhatla, and Abhinav Shrivastava. Trajectory-aligned space-time tokens for few-shot action recognition. In *ECCV*, pages 474–493. Springer, 2024.
- [39] Ruineng Li, Daitao Xing, Huiming Sun, Yuanzhou Ha, Jinglin Shen, and Chiuman Ho. Tokenmotion: Decoupled motion control via token disentanglement for human-centric video generation. In *CVPR*, pages 1951–1961, 2025.
- [40] Xuanlin Li, Kyle Hsu, Jiayuan Gu, Oier Mees, Karl Pertsch, Homer Rich Walke, Chuyuan Fu, Ishikaa Lunawat, Isabel Sieh, Sean Kirmani, et al. Evaluating real-world robot manipulation policies in simulation. In *Conference on Robot Learning*, pages 3705–3728. PMLR, 2025.
- [41] Yingyan Li, Lue Fan, Jiawei He, Yuqi Wang, Yuntao Chen, Zhaoxiang Zhang, and Tieniu Tan. Enhancing end-to-end autonomous driving with latent world model. *arXiv preprint arXiv:2406.08481*, 2024.
- [42] Corey Lynch, Ayzaan Wahid, Jonathan Tompson, Tianli Ding, James Betker, Robert Baruch, Travis Armstrong, and Pete Florence. Interactive language: Talking to robots in real time. *IEEE Robotics and Automation Letters*, 2023.
- [43] Haoyu Ma, Jialong Wu, Ningya Feng, Chenjun Xiao, Dong Li, Jianye Hao, Jianmin Wang, and Mingsheng Long. Harmonydream: Task harmonization inside world models. *arXiv preprint arXiv:2310.00344*, 2023.
- [44] Yue Ma, Yulong Liu, Qiyuan Zhu, Ayden Yang, Kunyu Feng, Xinhua Zhang, Zhifeng Li, Sirui Han, Chenyang Qi, and Qifeng Chen. Follow-your-motion: Video motion transfer via efficient spatial-temporal decoupled finetuning. *arXiv preprint arXiv:2506.05207*, 2025.
- [45] Oier Mees, Lukas Hermann, Erick Rosete-Beas, and Wolfram Burgard. Calvin: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks. *IEEE Robotics and Automation Letters (RA-L)*, 7(3):7327–7334, 2022.
- [46] OpenAI. Sora 2: Advancing video generation models. <https://openai.com/index/sora-2/>, 2025. Accessed: November 11, 2025.
- [47] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *Transactions on Machine Learning Research Journal*, 2024.
- [48] Abby O’Neill, Abdul Rehman, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandekar, Ajinkya Jain, et al. Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6892–6903, 2024.
- [49] J Parker-Holder, P Ball, J Bruce, V Dasagi, K Holsheimer, C Kaplanis, A Moufarek, G Scully, J Shar, J Shi, et al. Genie 2: A large-scale foundation world model. *URL: https://deepmind.google/discover/blog/genie-2-a-large-scale-foundation-world-model*, 2024.
- [50] Ryan Po, Yotam Nitzan, Richard Zhang, Berlin Chen, Tri Dao, Eli Shechtman, Gordon Wetzstein, and Xun Huang. Long-context state-space video world models. *arXiv preprint arXiv:2505.20171*, 2025.
- [51] Zhongwei Ren, Yunchao Wei, Xun Guo, Yao Zhao, Bingyi Kang, Jiashi Feng, and Xiaojie Jin. Videoworld: Exploring knowledge learning from unlabeled videos. In *CVPR*, pages 29029–29039, 2025.
- [52] Dominik Schmidt and Minqi Jiang. Learning to act without actions. *arXiv preprint arXiv:2312.10812*, 2023.
- [53] Yifan Shen, Peiyuan Zhu, Zijian Li, Shaoan Xie, Namrata Deka, Zongfang Liu, Zeyu Tang, Guangyi Chen, and Kun Zhang. Controllable video generation with provable disentanglement. *arXiv preprint arXiv:2502.02690*, 2025.
- [54] Qiao Sun, Liujia Yang, Wei Tang, Wei Huang, Kaixin Xu, Yongchao Chen, Mingyu Liu, Jiange Yang, Haoyi Zhu, Yating Wang, et al. Learning primitive embodied world models: Towards scalable robotic learning. *arXiv preprint arXiv:2508.20840*, 2025.
- [55] Aether Team, Haoyi Zhu, Yifan Wang, Jianjun Zhou, Wenzheng Chang, Yang Zhou, Zizun Li, Junyi Chen, Chunhua Shen, Jiangmiao Pang, et al. Aether: Geometric-aware unified world modeling. *arXiv preprint arXiv:2503.18945*, 2025.
- [56] Bahey Tharwat, Yara Nasser, Ali Abouzeid, and Ian Reid. Latent action pretraining through world modeling. *arXiv preprint arXiv:2509.18428*, 2025.

- [57] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 2023. URL <https://arxiv.org/abs/2302.13971>.
- [58] Homer Rich Walke, Kevin Black, Tony Z Zhao, Quan Vuong, Chongyi Zheng, Philippe Hansen-Estruch, Andre Wang He, Vivek Myers, Moo Jin Kim, Max Du, et al. Bridgedata v2: A dataset for robot learning at scale. In *Conference on Robot Learning*, pages 1723–1736. PMLR, 2023.
- [59] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.
- [60] Jiacong Wang, Bohong Wu, Haiyong Jiang, Zhou Xun, Xin Xiao, Haoyuan Guo, and Jun Xiao. World to code: Multi-modal data generation via self-instructed compositional captioning and filtering. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4608–4623, 2024.
- [61] Yaohui Wang, Piotr Bilinski, Francois Bremond, and Antitza Dantcheva. G3an: Disentangling appearance and motion for video generation. In *CVPR*, pages 5264–5273, 2020.
- [62] Yucen Wang, Fengming Zhang, De-Chuan Zhan, Li Zhao, Kaixin Wang, and Jiang Bian. Co-evolving latent action world models. *arXiv preprint arXiv:2510.26433*, 2025.
- [63] Yuchi Wang, Junliang Guo, Xinyi Xie, Tianyu He, Xu Sun, and Jiang Bian. Vidtwinn: Video vae with decoupled structure and dynamics. *arXiv preprint arXiv:2412.17726*, 2025.
- [64] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [65] Zizhao Wang, Kaixin Wang, Li Zhao, Peter Stone, and Jiang Bian. Dyn-o: Building structured world models with object-centric representations. *arXiv preprint arXiv:2507.03298*, 2025.
- [66] Jialong Wu, Haoyu Ma, Chaoyi Deng, and Mingsheng Long. Pre-training contextualized world models with in-the-wild videos for reinforcement learning. *arXiv preprint arXiv:2305.18499*, 2023.
- [67] Jialong Wu, Shaofeng Yin, Ningya Feng, Xu He, Dong Li, Jianye Hao, and Mingsheng Long. ivideogpt: Interactive videogpts are scalable world models. *NeurIPS*, 37:68082–68119, 2024.
- [68] Jianwen Xie, Ruiqi Gao, Zilong Zheng, Song-Chun Zhu, and Ying Nian Wu. Motion-based generator model: Unsupervised disentanglement of appearance, trackable and intrackable motions in dynamic patterns. In *AAAI*, volume 34, pages 12442–12451, 2020.
- [69] Mengda Xu, Zhenjia Xu, Cheng Chi, Manuela Veloso, and Shuran Song. Xskill: Cross embodiment skill discovery. In *Conference on robot learning*, pages 3536–3555. PMLR, 2023.
- [70] An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.
- [71] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024.
- [72] Seonghyeon Ye, Joel Jang, Byeongguk Jeon, Sejune Joo, Jianwei Yang, Baolin Peng, Ajay Mandlekar, Reuben Tan, Yu-Wei Chao, Bill Yuchen Lin, et al. Latent action pretraining from videos. *arXiv preprint arXiv:2410.11758*, 2024.
- [73] Weirui Ye, Yunsheng Zhang, Pieter Abbeel, and Yang Gao. Become a proficient player with limited data through watching pure videos. In *The Eleventh International Conference on Learning Representations*, 2022.
- [74] Jiwen Yu, Yiran Qin, Haoxuan Che, Quande Liu, Xintao Wang, Pengfei Wan, Di Zhang, Kun Gai, Hao Chen, and Xihui Liu. A survey of interactive generative video. *arXiv preprint arXiv:2504.21853*, 2025.
- [75] Jiwen Yu, Yiran Qin, Xintao Wang, Pengfei Wan, Di Zhang, and Xihui Liu. Gamefactory: Creating new games with generative interactive videos. *arXiv preprint arXiv:2501.08325*, 2025.

- [76] Lijun Yu, Jose Lezama, Nitesh Bharadwaj Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G Hauptmann, et al. Language model beats diffusion-tokenizer is key to visual generation. In ICLR, 2024.
- [77] Hangjie Yuan, Weihua Chen, Jun Cen, Hu Yu, Jingyun Liang, Shuning Chang, Zhihui Lin, Tao Feng, Pengwei Liu, Jiazheng Xing, et al. Lumos-1: On autoregressive video generation from a unified model perspective. arXiv preprint arXiv:2507.08801, 2025.
- [78] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In CVPR, pages 3836–3847, 2023.
- [79] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In CVPR, pages 586–595, 2018.
- [80] Long Zhao, Sanghyun Woo, Ziyu Wan, Yandong Li, Han Zhang, Boqing Gong, Hartwig Adam, Xuhui Jia, and Ting Liu. Epsilon-vae: Denoising as visual decoding. arXiv preprint arXiv:2410.04081, 2025.
- [81] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. arXiv preprint arXiv:2303.18223, 1(2), 2023.
- [82] Ruijie Zheng, Jing Wang, Scott Reed, Johan Bjorck, Yu Fang, Fengyuan Hu, Joel Jang, Kaushil Kundalia, Zongyu Lin, Loic Magne, et al. Flare: Robot learning with implicit world modeling. arXiv preprint arXiv:2505.15659, 2025.
- [83] Gaoyue Zhou, Hengkai Pan, Yann LeCun, and Lerrel Pinto. Dino-wm: World models on pre-trained visual features enable zero-shot planning. arXiv preprint arXiv:2411.04983, 2024.
- [84] Zheng Zhu, Xiaofeng Wang, Wangbo Zhao, Chen Min, Nianchen Deng, Min Dou, Yuqi Wang, Botian Shi, Kai Wang, Chi Zhang, et al. Is sora a world simulator? a comprehensive survey on general world models and beyond. arXiv preprint arXiv:2405.03520, 2024.

Appendix

In this supplementary material, we provide additional details and analyses:

- Sec. A describes the training setup, the structure of the dLDM, and other implementation details.
- Sec. B provides further analysis of Video-CraftBench.
- Sec. C provides additional visualizations of both our method and the baselines.

A Implementation Details

Training details. We present the detailed training configurations of dynamics-enhanced latent dynamics model and auto-regressive transformer in Tab. 4.

Dynamics-enhanced latent dynamics model (dLDM). As shown in Fig. 4, the proposed dLDM consists of four components. (i) A causal encoder extracts visual features. (ii) A set of learnable queries attends to the encoder features to extract visual changes and produce the latent dynamic codes. (iii) A decoder reconstructs frames from the latent dynamic codes and the feature of the first frame, providing coarse motion cues. (iv) A pretrained Video Diffusion Model (VDM) generates high-fidelity future frames by taking three inputs: the first frame, the low-resolution decoder outputs, and the latent dynamic codes. The latent codes are injected into the VDM via causal cross-attention, allowing the codes to concentrate on task-relevant dynamics while the VDM handles visual appearance modeling. We provide PyTorch-style pseudocode for the full dLDM and its components in Alg. ??.

Latent pre-training on CALVIN. In Sec. 5.4 (line 432), we present a latent pre-training experiment where latent dynamic tokens serve as prediction targets. This enables the model to learn manipulation knowledge from unlabeled data, thereby facilitating better adaptation to ground-truth actions. Following LAPA [72]’s protocol, given a task trajectory $x_{0:T}$, we train the AR Transformer to predict the quantized latent dynamic embeddings $\{z_k^n\}_{k=1, n=1}^{K, N}$ (defined in Sec. 3 line 198), conditioned on the language instruction and the first image x_0 . We extend the AR transformer’s vocabulary to include these latent dynamic tokens. Since this pre-training does not rely on ground-truth actions, it can use *any unlabeled video data*.

Following latent pre-training, we further fine-tune the model with ground-truth labels to derive actions executable in the simulation environment. We append a new action head (an MLP layer) to the transformer to map the hidden states, which are originally used for latent token prediction, to real actions via an ℓ_2 loss. To ensure compatibility with cross-environment data, we exclusively use CALVIN’s static camera images during both pre-training and fine-tuning, omitting wrist images and state parameters. This latent pre-training should provide the AR transformer with a manipulation prior, leading to better fine-tuning performance.

B Details in Video-CraftBench

Key steps for evaluation. In Sec. 4.2 (line 314), we define seven key steps in the paper folding task to evaluate the task success rate. Fig. 6 (top) provides a schematic diagram of these key steps.

Train and test environments. Fig. 6 (bottom) illustrates the Video-CraftBench training and testing environments. To evaluate the model’s generalization to new environments, the test set features diverse variations compared to the training set. Specifically, the paper folding test set differs in background, texture, paper appearance, and camera viewpoints, while the block-building varies in initial block arrangement, color combinations, and camera angles.

Task analysis. Fig. 9 shows the task duration distribution of Video-CraftBench. The largest portion (37.3%) falls within 45–60 seconds, followed by 27.1% at 60–90 seconds, while short tasks (20–30s) make up only 10.9%. This distribution highlights our emphasis on long-horizon tasks. Regarding task types (Fig. 10), paper folding accounts for 55.3% (5.2 hours) of the total duration due to its longer execution time, while block building comprises 1.8 hours.

Config	dLDM	AR Transformer
optimizer	AdamW	AdamW
base learning rate	1e-4	3e-4
weight decay	0.1	0.05
optimizer momentum	$\beta_1, \beta_2=0.9, 0.99$	$\beta_1, \beta_2=0.9, 0.98$
batch size	128	256
learning rate schedule	WarmupDecayLR	WarmupDecayLR
warmup iterations	2e+3	5e+3
max iterations	1e+5	5e+4
augmentations	None	None
Training Loss	L2 & Denoising loss	CE loss
Training Target	Recon. & Denoising	Next token pred.

Table 4 Training configurations for the dynamics-enhanced latent dynamics model and auto-regressive transformer.




Figure 9 Video duration distribution.

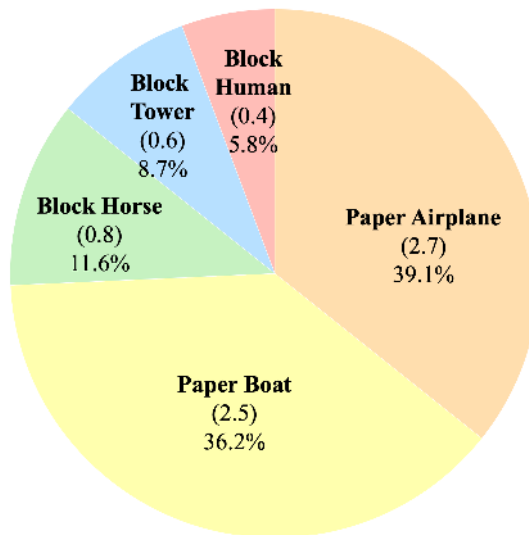


Figure 10 Task type distribution.

Task classifier of Video-CraftBench. We detail the training of the classifier used to evaluate key step completion in Sec. 4.2 (line 314). Its primary objective is to determine whether the model has reached key steps by analyzing paper shapes and block arrangements. Appearance distractions, such as texture, color, and the background drift shown in Fig. 8 (last row), are disregarded. This visual quality is assessed separately via LPIPS and SSIM. Therefore, we initialize the classifier with DINOv2-Base (86M parameters), chosen for its exceptional geometric awareness and ability to extract fine-grained features.

To train this classifier, we construct a dataset of $\sim 25k$ labeled frames. First, we extract $\sim 15k$ frames of both key and non-key steps from the training and testing environments. To further enhance diversity and mitigate bias, we generate task execution videos using all models involved in the evaluation, manually verify successful trajectories, and extract an additional $\sim 10k$ frames. We then attach a classification head to the DINOv2 backbone and fine-tune the entire model. The classifier achieves a test accuracy of 96.1%, ensuring reliable judgment for subsequent evaluations.

Algorithm 1: dLDM Algorithm (No Underline Hack)

Input : Video sequence $V \in \mathbb{R}^{(1+T) \times h \times w \times 3}$ **Variables:** Learnable embeddings $Q_{ldm} \in \mathbb{R}^{N \times C}$

```
1 Function
2   encoder(video):                                     // Encoder Function
3    $f \leftarrow \text{Causal3DCNN}(\text{video})$ 
4   for layer in encoder_layers do
5   |  $f \leftarrow \text{layer}(f)$ 
6   end
7   return  $z, f[0]$ 
8 end
9 Function
10  ldm_qformer(f):
11   $q\_list \leftarrow []$ 
12  for  $k \leftarrow 2$  to  $K$  do
13  |  $f_k \leftarrow f[:k]$ 
14  |  $q_k \leftarrow \text{MLP}(\text{CrossAttention}(Q_{ldm}, f_k, f_k))$ 
15  |  $q\_list.append(q_k)$ 
16  end
17  return  $\text{stack}(q\_list)$ 
18 end
19 Function
20  decoder( $z, \text{first\_f}$ ):
21   $rec \leftarrow \text{repeat}(\text{first\_f}, K, \text{dim} = 0)$ 
22  for _ in decoder_layers do
23  |  $rec \leftarrow \text{Causal3DCNN}(rec)$ 
24  |  $rec \leftarrow \text{up\_scale}(rec)$ 
25  end
26  return  $rec$ 
27 end
28 Function
29  main(video):
30   $z, \text{first\_f} \leftarrow \text{encoder}(\text{video})$ 
31   $z \leftarrow \text{FSQ}(z)$ 
32  if is_train then
33  |  $rec \leftarrow \text{decoder}(z.detach(), \text{first\_f})$ 
34  | return  $\text{MSE}(rec, \text{video}) + \text{VDM}(\text{video}, z, rec)$ 
35  end
36  return  $z$ 
37 end
```

C Visualizations

Fig. 11 presents additional visualizations of VideoWorld 2 on Video-CraftBench. We also include more video demonstrations in the supplementary material. Thanks to the dLDM, VideoWorld 2 generates accurate long-horizon execution sequences in novel environments, producing visually coherent, high-quality videos. Note that the VDM processes only 93 frames at a time, while full task sequences can span thousands of frames. Long videos are therefore generated auto-regressively by extending each segment from the final frame of the previous one. Because the VDM’s inherent reconstruction noise accumulates over time, visual artifacts such as lighting, texture, or color shifts may gradually appear. Nevertheless, the key steps within the generated sequences remain accurate.



Figure 11 More visualization of VideoWorld 2 on Video-CraftBench.