

WorldWander: Bridging Egocentric and Exocentric Worlds in Video Generation

Quanjian Song^{2*}, Yiren Song^{1*}, Kelly Peng², Yuan Gao², and Mike Zheng Shou^{1†}

¹Show Lab, National University of Singapore, Singapore ²First Intelligence
<https://github.com/showlab/WorldWander>



Fig. 1: Gallery of our WorldWander. It bridges the egocentric and exocentric worlds in video generation, enabling character-centric and immersive navigation.

Abstract. Recent advances in video world models enable interactive environments with free navigation, making translation between first-person (egocentric) and third-person (exocentric) perspectives increasingly important. However, existing studies focus on unidirectional exocentric-to-egocentric translation, overlooking reference-guided exocentric perspective synthesis. This capability is crucial for gaming and embodied AI applications. Motivated by this, we present WorldWander, an in-context learning framework tailored for translating between egocentric and exocentric worlds in video generation. Building upon advanced video diffusion transformers, WorldWander integrates (i) In-Context Perspective Alignment and (ii) Collaborative Position Encoding to model cross-view synchronization and character consistency. To support our task, we curate EgoExo-8K, a dynamic and scene-rich dataset containing synchronized egocentric-exocentric triplets from both synthetic and real-world scenarios. Experiments demonstrate that WorldWander achieves superior perspective synchronization, character consistency, and generalization, setting a new benchmark for egocentric-exocentric video translation.

Keywords: Egocentric-Exocentric Video Translation

*Equal contribution. †Corresponding author.

1 Introduction

Driven by the advancements in diffusion models [21, 35, 40], video generation has attracted increasing attention. Early research primarily focus on producing temporally coherent videos with controllable attributes such as camera motion [8, 50], scene [39, 53], and style [36, 42, 44]. Building upon these advances, subsequent works [1, 3, 15, 58, 59] explore perspective re-orientation, aiming to generate videos from source perspectives to target perspectives. More recently, emerging world models [2, 9, 20] further extend this paradigm by constructing interactive virtual environments that enable free navigation. As users naturally navigate these environments, the ability to seamlessly switch between first-person (egocentric) and third-person (exocentric) perspectives becomes increasingly important. However, existing approaches [4, 7, 54] are largely limited to one-directional exocentric-to-egocentric translation and overlook the inverse problem of reference-guided exocentric perspective synthesis. Such capability is essential for applications including gaming production, embodied AI, and VR, where flexible perspective switching enables immersive and character-centric exploration in the worlds.

In this work, we first formulate the task of **Egocentric–Exocentric Perspective Translation**, which aims to seamlessly translate video across egocentric and exocentric perspectives during dynamic scene navigation. Accordingly, we identify three primary challenges inherent to this setting: (i) *Perspective Synchronization*. Existing unidirectional egocentric-to-exocentric methods [4, 7, 54] are largely confined to static cameras. In our dynamic scene navigation, they struggle to maintain synchronization of key elements (*e.g.*, appearance, motion rhythm, and environmental dynamics) between these two perspectives. (ii) *Character Consistency*. The novel exo-to-ego translation task requires generating customized exocentric views conditioned on an egocentric input and a reference image. Ensuring character consistency is non-trivial, as the model must associate the first-person observations with the corresponding third-person identity. (iii) *Data Scarcity*. While prior works [7, 19, 25] provide paired egocentric–exocentric datasets, they are heavily restricted to static cameras. Such static setups fail to support our “WorldWander” vision, where a personalized character can freely navigate via dynamic environments with seamless video perspective translation.

To address these issues above, we propose **WorldWander**, an in-context learning framework tailored for translating between egocentric and exocentric worlds in video generation (see Fig. 1). Unlike traditional novel-view synthesis methods relying on 3D geometry or explicit camera pose, WorldWander performs in-context learning directly from triplet datasets. Specifically, it is built upon the state-of-the-art video diffusion transformer Wan2.2 [48] and integrates two core components for efficient fine-tuning: (i) **In-Context Perspective Alignment**, which transforms the conditional and target videos into a shared latent space, followed by collaborative attention to model cross-view alignment and character consistency; and (ii) **Collaborative Position Encoding**, which enables the conditional and target latents to share the unified positional embedding along the temporal dimension, thereby further enhancing cross-view consistency.

To tackle data scarcity, we carefully design a two-player protocol for recording synthetic scenes in GTA-5, alongside a dual-camera setup for capturing real-world scenarios. Ultimately, we curate **EgoExo-8K**, a dynamic and scene-rich dataset featuring synchronized egocentric and exocentric perspectives. It comprises approximately 4,000 synthetic and 4,000 real-world triplets, each containing an egocentric video, its exocentric counterpart, and a reference image. Notably, unlike previous datasets [7, 19, 25] confined to static camera viewpoints, EgoExo-8K features dynamic scene navigation in open-world environments. Comprehensive experiments demonstrate that WorldWander achieves high cross-view consistency and robust generalization in both transformations (ego $\rightarrow$ exo and exo $\rightarrow$ ego), significantly outperforming all existing baselines.

In summary, the main contributions of this paper are threefold:

- We formulate the *Egocentric-Exocentric Video Translation* tasks, where the novel exo-to-ego translation enables customized egocentric synthesis conditioned on a reference image, which is unexplored in prior works.
- We propose an in-context learning framework that is independent of camera pose and depth map, featuring *In-Context Perspective Alignment* and *Collaborative Position Encoding* for geometry-free cross-view translation.
- We carefully curate *EgoExo-8K*, the first dynamic and open-world scene wandering benchmark for the egocentric-exocentric video translation, on which our WorldWander achieves state-of-the-art performance across baselines.

2 Related Work

2.1 Exo-to-Ego Perspective Translation

Paired exocentric-egocentric data like EgoExo4D [7], H2O [19], and TACO [25] have driven recent progress in exocentric-to-egocentric view translation. Essentially, it involves synthesizing first-person (egocentric) perspective from third-person (exocentric) perspective while maintaining spatio-temporal consistency. Early works [27, 28, 31] fuse exocentric and egocentric latents via channel-wise concatenation. Given that this approach suffers from information loss and sluggish convergence, later methods [4, 23] shift towards cross-attention mechanisms to inject exocentric features. Subsequently, several works propose task variants that rely on auxiliary inputs for more accurate generation. For instance, Exo2Ego-V [24] leverages four exocentric views, while EgoExo-Gen [54] requires the first egocentric frame as an additional condition. However, these works typically focus on unidirectional ego-to-exo translation restricted to static viewpoints. In contrast, our work tackle bidirectional ego-exo translation in dynamic wandering scenarios with freely moving cameras. Furthermore, for the novel exo-to-ego translation, our method enables customized egocentric perspective synthesis using a reference character, which remains unexplored in prior works.

2.2 Viewpoint-Control in Video Generation

Driven by industrial demand, research on controllability in diffusion models has grown rapidly [6, 16, 26, 33, 34, 38, 41, 45, 51, 61], providing a solid foundation for

viewpoint-controllable video generation. Early on, some methods [8, 29, 50, 62] typically train an additional encoder to accept camera parameters and produce videos with camera movements; while other approaches [10, 37, 52] explore training-free techniques to simulate camera motion. Over time, some studies such as TrajectoryCrafter [58], Recammaster [1], and Reangle-a-video [15] have explored the task of perspective re-orientation. Given a reference video, these methods aim to transform it to the target perspective while maintaining spatial-temporal consistency. The emergence of world models such as Genie3, Matrix-Game 2.0 [9], and Hunyuan-GameCraft [20] has enabled video diffusion models to construct virtual worlds that users can navigate through camera movements. Despite these advances in video generation, existing approaches still struggle to model egocentric-exocentric transformations, which are the focus of this work.

2.3 Video-to-Video Translation

As a sub-task of controllable video generation, video-to-video translation aims to take an additional video as the condition and transform it into another video. On one hand, different studies explore this task from various dimensions: Uni-VST [36], StyleCrafter [22], and StyleMaster [57] focus on stylizing conditional videos into specific artistic forms; FlowV2V [49] and Render-A-Video [55] explore video editing for controllable outcomes. On the other hand, several works such as AnyV2V [18], I2VEdit [30], and VACE [17] have sought to unify different forms of Video-to-Video Generation within a single framework. These frameworks are capable of performing diverse tasks, including style transfer [13, 46], inpainting [5], outpainting, and others [12, 43, 60]. However, none of these approaches explores the transformations between egocentric and exocentric perspectives. Our work aims to bridge this gap in the field.

3 Preliminary

Video Diffusion Transformer. The state-of-the-art video diffusion transformer [32] is typically composed of a pair of variational encoder $\mathcal{E}$ and decoder $\mathcal{D}$, as well as a transformer-based prediction network v_θ . During training, the encoder $\mathcal{E}$ will transform F frames of the target video into f frames of the latent representation $z_0^{1:f}$, where $f = \frac{F-1}{4} + 1$. According to the rectified-flow [21], the forward process is defined as a straight path from the data distribution to a standard normal distribution. This process is formulated below:

$$z_t^{1:f} = (1 - t) \cdot z_0^{1:f} + t \cdot \epsilon^{1:f}, \quad (1)$$

where $\epsilon^{1:f} \in \mathcal{N}(0, I)$ and t is a random timestep. Given the noisy latent $z_t^{1:f}$, we utilize the network v_θ to regress the vector field via flow matching [21] loss:

$$\min_{\theta} \mathbb{E}_{t \sim \mathcal{U}(0,1)} \|v_\theta(z_t^{1:f}, t, c) - v\|_2^2, \quad (2)$$




Fig. 2: Overall pipeline of WorldWander. The backbone Wan2.2-5B [48] is fine-tuned with proposed *In-Context Perspective Alignment* (Shared Latent Space, Different Noise Levels, and Collaborative Attention) as well as *Collaborative Position Encoding*.

where $v = \epsilon^{1:f} - z_0^{1:f}$ denotes the target vector field, and c represents the additional conditioning signal (e.g., text, image, video).

During inference, the trained v_θ will iteratively map samples from Gaussian distribution back to specific data distribution, by solving the ODE below:

$$dz_t^{1:f} = v_\theta(z_t^{1:f}, t, c)dt. \quad (3)$$

Finally, the decoder $\mathcal{D}$ will reconstruct final video from the denoised latent.

Rotary Position Encoding. In advanced video generation models [20, 48, 56], rotary position embeddings are typically applied within the attention process, thereby distinguishing tokens across spatial and temporal dimensions.

In detail, customized positional embeddings $\mathcal{P}$ will be injected into the query and key representations (Q and K) before the attention computation:

$$\begin{aligned} Q[i, j, k] &\leftarrow Q[i, j, k] \cdot \mathcal{P}[i, j, k], \\ K[i, j, k] &\leftarrow K[i, j, k] \cdot \mathcal{P}[i, j, k], \end{aligned} \quad (4)$$

where i , j , and k indicate the indices along the height, width, and frame. We elaborate in the methodology section on how the positional embeddings $\mathcal{P}$ are constructed to adapt for our egocentric–exocentric video translation.

4 Methodology

In this work, we propose **WorldWander**, an in-context learning framework that bridges first-person (egocentric) and third-person (exocentric) perspectives in video generation, enabling immersive and character-centric world exploration. The overall pipeline is shown in Fig. 2. First, we describe our *Task Formulation* in Sec. 4.1. Then, we present the *In-Context Perspective Alignment* paradigm in Sec. 4.2 and the *Collaborative Position Encoding* strategy in Sec. 4.3, which

model the perspective correspondence and character consistency without auxiliary networks. Subsequently, we provide details of *Efficient Fine-Tuning via LoRA* in Sec. 4.4. Finally, we present curated *EgoExo-8K* dataset in Sec. 4.5.

4.1 Task Formulation

This paper tackles a challenging yet applicable task of egocentric-exocentric video translation. In detail, for the triplet sample $\langle \mathcal{V}^{\text{ego}}, \mathcal{V}^{\text{exo}}, \mathcal{I}^{\text{ref}} \rangle$, egocentric-exocentric translation consists of two sub-tasks: (i) *Exocentric-to-egocentric translation*, where the exocentric video $\mathcal{V}^{\text{exo}}$ is used as condition and transformed into corresponding egocentric video $\mathcal{V}^{\text{ego}}$; and (ii) *Egocentric-to-exocentric translation*, where the egocentric video $\mathcal{V}^{\text{ego}}$ and reference image $\mathcal{I}^{\text{ref}}$ are used as the condition and transformed back into the exocentric video $\mathcal{V}^{\text{exo}}$. Note that neither camera poses nor depth information is provided, which deliberately complicates the task to enforce robust generalization across different cameras and scenarios. Importantly, our work addresses bidirectional ego-exo translation in dynamic wandering scenarios with freely moving cameras, which distinguishes it from previous unidirectional exo-to-ego methods [4, 24, 54] with static viewpoints.

4.2 In-Context Perspective Alignment

Previous controllable video generation methods [8, 50] typically train an auxiliary network to inject conditional information. Although this strategy is effective, it inevitably introduces extra parameters. To address this limitation, we introduce an *In-Context Perspective Alignment* paradigm, which efficiently performs in-context learning within a single backbone network, avoiding auxiliary networks. **Shared Latent Space.** Recall the mentioned triplet sample $\langle \mathcal{V}^{\text{ego}}, \mathcal{V}^{\text{exo}}, \mathcal{I}^{\text{ref}} \rangle$, where $\mathcal{V}^{\text{ego}}$ represents the egocentric video, $\mathcal{V}^{\text{exo}}$ denotes the corresponding exocentric video, and $\mathcal{I}^{\text{ref}}$ depicts the back-view character image. During fine-tuning, instead of training an additional condition encoder, we utilize a unified VAE encoder $\mathcal{E}$ to encode them into a shared latent space $\langle z_0^{\text{ego}}, z_0^{\text{exo}}, z_0^{\text{ref}} \rangle$, as below:

$$z_0^{\text{ego}} = \mathcal{E}(\mathcal{V}^{\text{ego}}); \quad z_0^{\text{exo}} = \mathcal{E}(\mathcal{V}^{\text{exo}}); \quad z_0^{\text{ref}} = \mathcal{E}(\mathcal{I}^{\text{ref}}), \quad (5)$$

where all latent can share semantic space while eliminating additional encoders. **Different Noise Levels.** Building upon the shared latent space, the target latent is noised in the same manner as the rectified flow [21], while the conditional latent remains noise-free. This strategy preserves the integrity of the conditional information while explicitly distinguishing the conditional latents from the target latent, enabling the model to learn their correspondence more effectively.

In detail, for exocentric-to-egocentric translation: the egocentric latent z_0^{ego} is noised into z_t^{ego} according to Eq. (1), while the exocentric latent z_0^{exo} remains noise-free; for egocentric-to-exocentric translation: the exocentric latent z_0^{exo} is noised into z_t^{exo} according to Eq. (1), while the egocentric latent z_0^{ego} and reference character latent z_0^{ref} remains noise-free. Note that the timestep of the noised latent is t , while that of the noise-free latent is 0, accordingly. In the following,




Fig. 3: Loss comparison across different approaches on synthetic scenarios. *Collaborative Attention* demonstrates faster convergence than *Channel-Wise Concatenation*.

we discuss how to enable efficient interaction between the conditional latent and the target latent in the egocentric-exocentric translation task.

Collaborative Attention. To integrate the conditional latent into the target latent within a single backbone network, a simple approach is to concatenate them along the channel dimension for information fusion. This only requires modifying input channels during the patchify stage, while keeping the rest of the backbone unchanged. However, as shown in Fig. 3, direct channel-wise concatenation leads to slow convergence during fine-tuning. We attribute this to the channel fusion, which hinders explicit interaction between the two latent types. To address this, we concatenate the conditional and target latents along the token dimension, integrating collaborative attention to enable explicit interaction.

Specifically, for exocentric-to-egocentric translation: we concatenate the noise-free exocentric latent z_0^{exo} and the noisy egocentric latent z_t^{ego} along the token dimension as the unified input z_t^{uni} ; for egocentric-to-exocentric translation: we concatenate the noise-free reference character latent z_0^{ref} , the noise-free exocentric latent z_0^{ego} , and the noised exocentric latent z_t^{exo} along the token dimension as the unified input z_t^{uni} . Formally, this process can be formulated as follows:

$$z_t^{\text{uni}} = \begin{cases} \text{TokenConcat}([z_0^{\text{exo}}, z_t^{\text{ego}}]), & \text{exo} \rightarrow \text{ego} \\ \text{TokenConcat}([z_0^{\text{ref}}, z_0^{\text{ego}}, z_t^{\text{exo}}]), & \text{ego} \rightarrow \text{exo} \end{cases}. \quad (6)$$

During the self-attention process, the unified input z_t^{uni} is projected through different learnable matrices W_q , W_k , and W_v , and then participates in multi-modal (conditional latent and target latent) attention interactions. Formally, the attention output $\mathcal{O}$ can be formulated as follows:

$$\mathcal{O} = \text{Softmax}\left(\frac{(z_t^{\text{uni}}W_q)(z_t^{\text{uni}}W_k)^\top}{\sqrt{d_k}}\right)(z_t^{\text{uni}}W_v), \quad (7)$$




Fig. 4: Toy examples between naive 3D RoPE and our collaborative position encoding.

where d_k denotes the feature dimension. In this way, the conditional and target latents share the same projection matrices for global interaction, thereby avoiding additional parameters from auxiliary networks.

As illustrated in Fig. 3, the collaborative attention strategy enables the model to achieve faster convergence compared to simple channel-wise concatenation.

4.3 Collaborative Position Encoding

To distinguish input tokens across spatial and temporal dimensions, rotary position embeddings are typically integrated into attention process according to Eq. (4). Recall that our input z_t^{uni} comprises both conditional and target latent. A straightforward approach is to treat them as a whole, uniformly encoding them along the temporal dimension. However, as analyzed in our ablation study, this strategy often leads to incorrect cross-view correspondences between egocentric and exocentric videos. In practice, the conditional and target latents exhibit a frame-wise alignment; uniformly encoding them disrupts this alignment, thereby hindering the learning such correspondences. To overcome this limitation, we design a *Collaborative Position Encoding* strategy to effectively learn the positional mapping between the two latent types, with the toy examples show in Fig. 4.

In detail, the unified latent z^{uni} is first decomposed into corresponding conditional and target latents, which are then independently encoded to share the same position embeddings along the temporal dimension. This strategy enforces a frame-wise positional correspondence. Formally, the position embeddings $\mathcal{P}^{\text{uni}}$ for two sub-tasks are formulated as follows:

$$\mathcal{P}^{\text{uni}} = \begin{cases} \text{Concat}([R(z_0^{\text{exo}}), R(z_t^{\text{ego}})]), & \text{exo} \rightarrow \text{ego} \\ \text{Concat}([R(z_0^{\text{ref}}), R(z_0^{\text{ego}}), R(z_t^{\text{exo}})]), & \text{ego} \rightarrow \text{exo} \end{cases}, \quad (8)$$

where R denotes the positional encoding operation. Subsequently, the collaborative position embeddings $\mathcal{P}^{\text{uni}}$ are integrated into the attention process according to Eq. (4), thus facilitating the interaction between different latents.



Fig. 5: Representative samples of our carefully curated EgoExo-8K dataset. Compared to prior egocentric-exocentric datasets [7, 19, 25], EgoExo-8K uniquely supports dynamic scene navigation and flexible camera tracking, encompassing both synthetic and real-world scenarios across diverse indoor and outdoor environments.

4.4 Efficient Fine-Tuning via LoRA

In this work, we adopt the video generation model Wan2.2-5B [48] as our backbone. Given the pretrained model encodes strong priors from large-scale video data, we fine-tune it with LoRA [11] to efficiently preserve these priors under limited resources. In detail, LoRA injects trainable low-rank matrices $\mathcal{A} \in \mathbb{R}^{r \times k}$ and $\mathcal{B} \in \mathbb{R}^{d \times r}$ into specific layers of the model, while keeping the corresponding original weights $\mathcal{W} \in \mathbb{R}^{d \times k}$ frozen to preserve the priors. Here, r is the rank, k is the input dimension, and d is the output dimension, with $r \ll \min(d, k)$. Once integrated with LoRA, the forward computation is modified below:

$$y' = y + \Delta y = \mathcal{W} \cdot x + \mathcal{B} \cdot \mathcal{A} \cdot x, \quad (9)$$

where x is the input to specific layers and y' is the corresponding output. By default, matrix $\mathcal{B}$ is initialized to zeros to stabilize the fine-tuning process.

During fine-tuning, the backbone is integrated with the proposed *In-Context Perspective Alignment* paradigm and *Collaborative Position Encoding* strategy, and optimized using the flow-matching loss defined in Eq. (2) as supervision.

4.5 EgoExo-8K

High-quality egocentric-exocentric data is a key prerequisite for models to learn cross-view correspondences. Although prior works [7, 19, 25] have introduced paired egocentric-exocentric datasets, these datasets are collected from static viewpoints in confined scenarios, involving limited camera dynamics. Such static setups contradict our core vision: modeling a “WorldWander” paradigm, where

a character freely wanders and navigate across dynamic environments and the viewpoint can seamlessly switch between egocentric and exocentric perspectives. To bridge this gap, we curate the *EgoExo-8K* dataset (see Fig. 5), a dynamic and scene-rich collection of synchronized egocentric-exocentric triplets for perspective transformation during open-world wandering. Each triplet sample is organized as $\langle \mathcal{V}^{\text{ego}}, \mathcal{V}^{\text{exo}}, \mathcal{I}^{\text{ref}} \rangle$, where $\mathcal{V}^{\text{ego}}$ and $\mathcal{V}^{\text{exo}}$ denote the egocentric and exocentric videos, while $\mathcal{I}^{\text{ref}}$ is the corresponding reference image. In summary, EgoExo-8K contains approximately 4,000 triplets derived from both synthetic and real-world scenarios. More samples and analyses can be found in the **Appendix**.

Synthetic Scenarios. The synthetic subset is recorded within the GTA-5 environment, offering a photorealistic and controllable virtual world. In detail, we adopt a two-player collection protocol: one player controls the selected character and wanders across scenes from an egocentric perspective, while the other captures the exocentric perspective via *spectator mode*. Both recordings are strictly synchronized to ensure temporal consistency, with the selected character saved as the reference image. In total, we construct 400+ paired video sequences along with their corresponding reference images, where each video lasts more than one minute. These samples are subsequently filtered and segmented into 5-second clips (300 frames), yielding exactly 4,000 triplet samples. These triplet samples cover diverse indoor and outdoor scenes, featuring navigation, camera tracking, and dynamic interactions such as manipulation and task execution.

Real-World Scenarios. The real-world subset is captured across everyday environments, providing a natural and complex physical counterpart. Specifically, we adopt a dual-camera collection protocol: a participant equipped with a dual-camera setup wanders across diverse scenes, where a head-mounted action camera (Ac-tion5Pro) captures the egocentric perspective, while a rear-mounted panoramic camera (Osmo360) simultaneously records the exocentric view. Similarly, both recordings are strictly synchronized to ensure temporal consistency, while the reference image is randomly sampled from the exocentric video. In total, we construct 10h+ paired video sequences along with their corresponding reference images. These samples also subsequently filtered and segmented into 5-second clips (300 frames), yielding 4,000 triplet samples. These triplet samples also span indoor and outdoor scenes that capture human activities (*e.g.*, walking, cleaning, and carrying objects), providing realistic cross-view correspondences.

5 Experiments

5.1 Experimental details.

Implementation Details. Our WorldWander is built upon the state-of-the-art video diffusion transformer Wan2.2-5B [48], into which LoRA [11] modules are integrated within the attention layers, with the rank $r = 80$. Each video clip is resized and cropped to 704×1280 to match the original configuration of Wan2.2-5B, with the number of frames $F = 49$. Considering the inherent differences between the egocentric-to-exocentric and exocentric-to-egocentric tasks, we fine-tune the model for each task using the same configuration. In detail, we use a

Table 1: Automatic evaluation across methods on both the exocentric-egocentric translation tasks. We report results for both synthetic and real-world scenarios, with the best result in **bold** and the second-best result underlined. Note that in the zero-shot setting, TrajectoryCrafter is unable to generate customized characters; therefore, we do not report its CLIP-I score. * indicates fine-tuning using our egoexo-8K datasets.

Methods	Exocentric-to-Egocentric				Egocentric-to-Exocentric				
	LPIS↓	SSIM↑	FVD↓	VBench↑	LPIS↓	SSIM↑	FVD↓	VBench↑	CLIP-I↑
<i>Synthetic Scenarios</i>									
AnyV2V [18]	0.6627	0.4870	990.3	0.7595	0.6809	0.4516	1285.9	0.7311	0.6240
I2VEdit* [30]	0.6272	0.5154	477.2	0.8037	0.6385	0.4940	737.4	0.7880	<u>0.7470</u>
TrajectoryCrafter [58]	0.6109	0.5068	743.2	0.7967	0.6112	0.5031	745.4	<u>0.8171</u>	–
ReCamMaster [1]	0.6516	0.4039	738.7	0.7634	0.6178	0.4629	915.7	0.7780	0.6537
ReCamMaster* [1]	<u>0.6039</u>	<u>0.5365</u>	<u>298.9</u>	<u>0.8098</u>	<u>0.5732</u>	<u>0.5125</u>	<u>531.5</u>	0.8116	0.7129
WorldWander (Ours)	0.5811	0.5416	230.3	0.8110	0.5550	0.5215	353.7	0.8395	0.7672
<i>Real-World Scenarios</i>									
AnyV2V [18]	0.7227	0.3394	1789.3	0.6813	0.7182	0.2885	1451.2	0.6705	0.5836
I2VEdit* [30]	0.6945	0.3428	1226.3	<u>0.7996</u>	0.6892	<u>0.3253</u>	<u>904.4</u>	<u>0.7815</u>	0.6601
TrajectoryCrafter [58]	0.7058	0.3127	1238.4	0.7977	0.6928	0.3215	1126.5	0.7583	–
ReCamMaster* [1]	<u>0.6752</u>	<u>0.3476</u>	<u>716.1</u>	0.7933	<u>0.6699</u>	0.3048	968.1	0.7787	<u>0.7020</u>
WorldWander (Ours)	0.6054	0.3567	416.5	0.8333	0.5357	0.3469	374.7	0.8142	0.7214

batch size of 4 per GPU and the AdamW optimizer (learning rate 1×10^{-4} , weight decay 1×10^{-2}). All experiments are conducted on four NVIDIA H200 GPUs. Training takes about two days, while inference requires about three minutes.

Evaluation Settings. Our task aims to tackle mutual egocentric-exocentric perspective translation. However, direct comparisons with existing exo-to-ego methods are infeasible: the implementations of 4Diff [4] and EgoExoGen [54] are not publicly available, while EgoExo-V [24] requires four exocentric video inputs, creating a fundamental mismatch with our setting. In light of these, we alternatively select two categories of closely related methods to serve as baselines: (i) *video-to-video translation*, AnyV2V [18] and I2VEdit [30]; (ii) *perspective re-orientation*, TrajectoryCrafter [58] and RecamMaster [1]. Implementation details for all baselines are provided in the **Appendix**. We evaluate the effectiveness of these approaches using our curated EgoExo-8K dataset. Considering the domain gap between synthetic and real-world scenes, we fine-tune the model individually for each scene and reserve 50 corresponding triplets as the test set for evaluation.

5.2 Quantitative Results

In this section, we conduct quantitative comparisons across baselines on both the exocentric-to-egocentric and egocentric-to-exocentric video translation tasks. The evaluation focuses on both synthetic and real-world scenarios, analyzed from two complementary perspectives: *Automatic Evaluation* that provides objective comparison, and *User Study* that captures subjective comparison.

Automatic Evaluation. We evaluate the generated videos across different tasks and from multiple dimensions. For the *exocentric-to-egocentric* translation:




Fig. 6: User study across baselines on the exocentric-egocentric video translation tasks. We report average results of the synthetic and real-world scenarios.

LPIPS and SSIM between generated and ground-truth videos are used to evaluate perspective alignment, FVD [47] with respect to the corresponding ground-truth is used to assess distribution alignment, and the average VBench [14] score is used to measure overall generation quality. For the *egocentric-to-exocentric* translation: we report the same set of metrics and additionally include CLIP-I to evaluate the semantic consistency between the characters in the generated videos and the reference images. We report results both on synthetic and real-world scenarios, with the details illustrated in Tab. 1. Note that in the zero-shot setting, TrajectoryCrafter [58] is unable to generate customized characters; therefore, we do not report its CLIP-I score. In synthetic scenarios, our WorldWander outperforms all baselines in terms of distribution alignment, perspective alignment, overall quality, and character consistency, across both directions of egocentric-exocentric video translation. With respect to real-world scenarios, WorldWander also outperforms all baselines across all evaluation dimensions, showing superior performance in the egocentric-exocentric video translation tasks. An additional point worth noting is that we compare the performance of RecamMaster with and without fine-tuning on our EgoExo-8K; the fine-tuned version shows superior performance, further validating the effectiveness of our datasets.

User Study. To further evaluate different methods on bidirectional *egocentric-exocentric* video translations, we conduct a user study to assess human perceptual preferences. In detail, we design a questionnaire-based evaluation covering both real and synthetic scenes. For exocentric-to-egocentric translation, participants are asked to select the best output among all methods based on (i) generation quality and (ii) perspective consistency. For egocentric-to-exocentric translation, participants are asked to select the best output using the same criteria and additionally evaluate (iii) character consistency. In total, we collect 630 valid responses and report the averaged results for both synthetic and real-world scenarios, with the details illustrated in Fig. 6. Clearly, our WorldWander out-




Fig. 7: Qualitative comparison of different methods on both exocentric-to-egocentric and egocentric-to-exocentric video translation tasks. Best viewed with zooming in.

performs all baselines across different subjective metrics, further demonstrating its human preferences in the egocentric-exocentric video translation task.

5.3 Qualitative Results

We conduct a detailed comparison across baselines for both exocentric-to-egocentric and egocentric-to-exocentric translations, and visualize the results for synthetic and real-world scenarios in Fig. 7. In synthetic scenarios, traditional video-to-video translation methods like AnyV2V [18] and I2VEdit [30] fail to model perspective correspondences. Recent perspective re-orientation approaches, such as Trajectory [58], which only support perspective changes, struggle to preserve identity consistency and handle identity removal in exocentric-to-egocentric transformations. In contrast, our WorldWander exhibits higher perspective consistency and character consistency than other methods, while maintaining superior generation quality. In real-world scenarios, WorldWander also exhibits superior synchronization and quality, further highlighting its generalization. In




Fig. 8: Representative examples for specific human-object interaction.

Table 2: Quantitative ablation about two components: *In-Context Perspective Alignment* and *Collaborative Position Encoding*. We compare two variants: *channel concatenation* and *uniform position encoding*, on both the exocentric-to-egocentric and egocentric-to-exocentric translation task. Due to limited resources, we only report the ablation results for synthetic scenarios, with the best result in **bold**.

Variants	Exocentric-to-Egocentric				Egocentric-to-Exocentric				
	LPIPS↓	SSIM↑	FVD↓	VBench↑	LPIPS↓	SSIM↑	FVD↓	VBench↑	CLIP-I↑
Channel Concat.	0.6008	0.5221	280.2	0.8072	0.6205	0.4708	524.3	0.8377	0.7310
Uni. Timestep	0.5963	0.5209	287.7	0.8063	0.5551	0.5211	365.9	0.8389	0.7429
Uni. Pos. Enc.	0.5907	0.5325	280.3	0.8044	0.5771	0.5083	356.5	0.8357	0.7562
Full Comp.($r=64$)	0.5931	0.5260	282.6	0.8083	0.5471	0.5165	362.9	0.8306	0.7461
Full Comp.($r=80$)	0.5811	0.5416	230.3	0.8110	0.5550	0.5215	353.7	0.8395	0.7672
Full Comp.($r=128$)	0.5819	0.5351	273.0	0.8122	0.5587	0.5202	344.9	0.8392	0.7556

addition to viewpoint translation in scene wandering, we also support translation for specific human-object interactions, and we visualize representative examples in Fig. 8. More comparisons and visualization are provided in **Appendix**.

5.4 Ablation Studies

In this section, we study the effectiveness of the core components: *In-Context Perspective Alignment* and *Collaborative Position Encoding*, and corresponding LoRA rank. Due to limited resources, we study only on synthetic scenarios.

In-Context Perspective Alignment. The core of in-context perspective alignment is the collaborative attention mechanism, which enables the clean conditional latent and the noisy target latent to perform in-context learning. To validate the effectiveness of this, we compare it with a variant that concatenates the two latents along the channel dimension. In Tab. 2, collaborative attention consistently outperforms channel-wise fusion across all metrics in egocentric-exocentric video translation. The qualitative results in Fig. 9 also support this conclusion. Additionally, we compare a uniform timestep with varying timestep; the results in Tab. 2 also demonstrate the superiority of our design.

Collaborative Position Encoding. Recall that our collaborative positional encoding independently encodes the conditional and target latents, thereby sharing the same embeddings along the temporal dimension and achieving better




Fig. 9: Qualitative ablation about two components: *In-Context Perspective Alignment* and *Collaborative Position Encoding*. Best viewed with zooming in.

frame-level perspective correspondence. To further evaluate its superiority, we compare it with a variant that uniformly encodes the conditional and target latents, where the positional embeddings differ across all frames. The quantitative results are reported in Tab. 2, while the qualitative results are illustrated in Fig. 9. Our collaborative position encoding strategy consistently outperforms the uniform position encoding in egocentric-exocentric video translation.

LoRA Rank. We compare the performance of LoRA with ranks of 64, 80, and 128, with the results shown in Fig. 9. Under the limited datas, a rank of 80 shows superior overall performance, so we ultimately selecte it as final setting.

6 Conclusion

This paper formulates the novel *egocentric-exocentric video perspective translation* task. To achieve this, we propose WorldWander, an in-context learning framework tailored for translating between egocentric and exocentric worlds in video generation. Leveraging advanced video diffusion transformers, WorldWander performs LoRA fine-tuning combined with the proposed *In-Context Perspective Alignment* and *Collaborative Position Encoding*. These strategies enable effective modeling of cross-view correspondences and character consistency. To facilitate this research, we further curate *EgoExo-8K*, a dynamic and scene-rich dataset featuring synchronized egocentric-exocentric triplet samples collected from both synthetic and real-world environments. Extensive experiments validate the effectiveness of WorldWander, demonstrating superior perspective synchronization, character consistency, and generalization across diverse scenarios.

Acknowledgements

This project is supported by the Ministry of Education, Singapore, under the Academic Research Fund Tier 1 (FY2023).

References

1. Bai, J., Xia, M., Fu, X., Wang, X., Mu, L., Cao, J., Liu, Z., Hu, H., Bai, X., Wan, P., et al.: Recammaster: Camera-controlled generative rendering from a single video. In: ICCV (2025)
2. Bar, A., Zhou, G., Tran, D., Darrell, T., LeCun, Y.: Navigation world models. In: CVPR (2025)
3. Chen, X., Chen, Z., Song, Y.: Transanimate: Taming layer diffusion to generate rgba video. In: CVPR (2026)
4. Cheng, F., Luo, M., Wang, H., Dimakis, A., Torresani, L., Bertasius, G., Grauman, K.: 4diff: 3d-aware diffusion model for third-to-first viewpoint translation. In: ECCV (2024)
5. Gao, H., Tang, B., Song, Y., Fang, G., He, Z., Yang, J., Shou, M.Z.: Pai-studio: Cinematic video background replacement with camera-aware motion. arXiv preprint arXiv:2606.01399 (2026)
6. Gong, Y., Song, Y., Li, Y., Li, C., Zhang, Y.: Relationadapter: Learning and transferring visual relation with diffusion transformers. In: NeurIPS (2026)
7. Grauman, K., Westbury, A., Torresani, L., Kitani, K., Malik, J., Afouras, T., Ashutosh, K., Baiyya, V., Bansal, S., Boote, B., et al.: Ego-exo4d: Understanding skilled human activity from first-and third-person perspectives. In: CVPR (2024)
8. He, H., Xu, Y., Guo, Y., Wetzstein, G., Dai, B., Li, H., Yang, C.: Cameractrl: Enabling camera control for video diffusion models. In: ICLR (2025)
9. He, X., Peng, C., Liu, Z., Wang, B., Zhang, Y., Cui, Q., Kang, F., Jiang, B., An, M., Ren, Y., et al.: Matrix-game 2.0: An open-source, real-time, and streaming interactive world model. arXiv preprint arXiv:2508.13009 (2025)
10. Hou, C., Wei, G., Zeng, Y., Chen, Z.: Training-free camera control for video generation. In: ICLR (2024)
11. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: Lora: Low-rank adaptation of large language models. ICLR (2022)
12. Hu, L.: Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In: CVPR (2024)
13. Huang, S., Song, Y., Zhang, Y., Guo, H., Wang, X., Shou, M.Z., Liu, J.: Photodoodle: Learning artistic image editing from few-shot pairwise data. arXiv preprint arXiv:2502.14397 (2025)
14. Huang, Z., He, Y., Yu, J., Zhang, F., Si, C., Jiang, Y., Zhang, Y., Wu, T., Jin, Q., Chanpaisit, N., et al.: Vbench: Comprehensive benchmark suite for video generative models. In: CVPR (2024)
15. Jeong, H., Lee, S., Ye, J.C.: Reangle-a-video: 4d video generation as video-to-video translation. In: ICCV (2025)
16. Jiang, Y., Gu, Y., Song, Y., Tsang, I., Shou, M.Z.: Personalized vision via visual in-context learning. arXiv preprint arXiv:2509.25172 (2025)
17. Jiang, Z., Han, Z., Mao, C., Zhang, J., Pan, Y., Liu, Y.: Vace: All-in-one video creation and editing. In: ICCV (2025)
18. Ku, M., Wei, C., Ren, W., Yang, H., Chen, W.: Anyv2v: A tuning-free framework for any video-to-video editing tasks. arXiv preprint arXiv:2403.14468 (2024)
19. Kwon, T., Tekin, B., Stühmer, J., Bogo, F., Pollefeys, M.: H2o: Two hands manipulating objects for first person interaction recognition. In: ICCV (2021)
20. Li, J., Tang, J., Xu, Z., Wu, L., Zhou, Y., Shao, S., Yu, T., Cao, Z., Lu, Q.: Hunyuan-gamecraft: High-dynamic interactive game video generation with hybrid history condition. arXiv preprint arXiv:2506.17201 (2025)

21. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: ICLR (2023)
22. Liu, G., Xia, M., Zhang, Y., Chen, H., Xing, J., Wang, Y., Wang, X., Shan, Y., Yang, Y.: Stylecrafter: Taming artistic video diffusion with reference-augmented adapter learning. ACM TOG (2024)
23. Liu, J.W., Mao, W., Xu, Z., Keppo, J., Shou, M.Z.: Exocentric-to-egocentric video generation. In: NeurIPS (2024)
24. Liu, J.W., Mao, W., Xu, Z., Keppo, J., Shou, M.Z.: Exocentric-to-egocentric video generation. NeurIPS (2024)
25. Liu, Y., Yang, H., Si, X., Liu, L., Li, Z., Zhang, Y., Liu, Y., Yi, L.: Taco: Benchmarking generalizable bimanual tool-action-object understanding. In: CVPR (2024)
26. Lu, R., Zhang, Y., Liu, J., Wang, H., Song, Y.: Easytext: Controllable diffusion transformer for multilingual text rendering. In: AAAI (2026)
27. Luo, H., Zhu, K., Zhai, W., Cao, Y.: Intention-driven ego-to-exo video generation. arXiv preprint arXiv:2403.09194 (2024)
28. Luo, M., Xue, Z., Dimakis, A., Grauman, K.: Put myself in your shoes: Lifting the egocentric perspective from exocentric videos. In: ECCV (2024)
29. Ma, W.D.K., Lewis, J.P., Kleijn, W.B.: Trailblazer: Trajectory control for diffusion-based video generation. In: SIGGRAPH Asia (2024)
30. Ouyang, W., Dong, Y., Yang, L., Si, J., Pan, X.: I2vedit: First-frame-guided video editing via image-to-video diffusion models. In: SIGGRAPH Asia (2024)
31. Park, J., Ye, A.S., Kwon, T.: Egoworld: Translating exocentric view to egocentric view using rich exocentric observations. arXiv preprint arXiv:2506.17896 (2025)
32. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: ICCV (2023)
33. Qu, Y., Chen, D., Li, X., Li, X., Zhang, S., Cao, L., Ji, R.: Drag your gaussian: Effective drag-based editing with score distillation for 3d gaussian splatting. In: SIGGRAPH (2025)
34. Shi, W., Song, Y., Zhang, D., Liu, J., Zou, X.: Fonts: Text rendering with typography and style controls. In: ICCV (2025)
35. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
36. Song, Q., Lin, M., Zhan, W., Yan, S., Cao, L., Ji, R.: Univst: A unified framework for training-free localized video style transfer. IEEE TPAMI (2025)
37. Song, Q., Lin, Z., Zeng, Z., Zhang, Z., Cao, L., Ji, R.: Lightmotion: A light and tuning-free method for simulating camera motion in video generation. arXiv preprint arXiv:2503.06508 (2025)
38. Song, Q., Shen, Y., Chen, M., Sun, H., Lan, J., Zhu, X., Zheng, B., Cao, L.: Fashionchameleon: Towards real-time and interactive human-garment video customization. arXiv preprint arXiv:2605.15824 (2026)
39. Song, Q., Zhou, D., Lin, J., Shen, F., Wang, J., Hu, X., Chen, C., Heng, P.A.: Scenedecorator: Towards scene-oriented story generation with scene planning and scene consistency. In: NeurIPS (2026)
40. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. arXiv preprint arXiv:2011.13456 (2020)
41. Song, Y., Chen, D., Shou, M.Z.: Layertracer: Cognitive-aligned layered svg synthesis via diffusion transformer. In: ICCV (2025)
42. Song, Y., Huang, S., Yao, C., Ye, X., Ci, H., Liu, J., Zhang, Y., Shou, M.Z.: Processpainter: Learn painting process from sequence data. arXiv preprint arXiv:2406.06062 (2024)

43. Song, Y., Liu, C., Jiang, Y., Shou, M.Z.: Streamingeffect: Real-time human-centric video effect generation. arXiv preprint arXiv:2605.17019 (2026)
44. Song, Y., Liu, C., Shou, M.Z.: Makeanything: Harnessing diffusion transformers for multi-domain procedural sequence generation. In: CVPR (2026)
45. Song, Y., Liu, C., Shou, M.Z.: Omniconsistency: Learning style-agnostic consistency from paired stylization data. In: NeurIPS (2026)
46. Song, Y., Yao, W., Wang, H., Shou, M.Z.: Vista: Triplet-supervised video style transfer with diffusion transformers. arXiv preprint arXiv:2605.17312 (2026)
47. Unterthiner, T., Van Steenkiste, S., Kurach, K., Marinier, R., Michalski, M., Gelly, S.: Towards accurate generative models of video: A new metric & challenges. arXiv preprint arXiv:1812.01717 (2018)
48. Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., et al.: Wan: Open and advanced large-scale video generative models. arXiv preprint arXiv:2503.20314 (2025)
49. Wang, G., Fan, S., Liu, H., Song, Q., Wang, H., Xu, J.: Consistent video editing as flow-driven image-to-video generation. In: CVPR (2026)
50. Wang, Z., Yuan, Z., Wang, X., Li, Y., Chen, T., Xia, M., Luo, P., Shan, Y.: Motionctrl: A unified and flexible motion controller for video generation. In: SIGGRAPH (2024)
51. Wang, Z., Zhao, H., Zhou, Q., Lu, X., Li, X., Yang, H., Yang, B., Song, Y.: Diffdecompose: Layer-wise decomposition of alpha-composited images via diffusion transformers. In: CVPR (2026)
52. Wu, J., Li, X., Zeng, Y., Zhang, J., Zhou, Q., Li, Y., Tong, Y., Chen, K.: Motion-booth: Motion-aware customized text-to-video generation. In: NeurIPS (2024)
53. Xie, C., Wang, C., Zhang, Y., Pan, Z., Zou, H., Zhan, Z.: Geometry-aware scene-consistent image generation. arXiv preprint arXiv:2512.12598 (2025)
54. Xu, J., Huang, Y., Pei, B., Hou, J., Li, Q., Chen, G., Zhang, Y., Feng, R., Xie, W.: Egoexo-gen: Ego-centric video prediction by watching exo-centric videos. In: ICLR (2025)
55. Yang, S., Zhou, Y., Liu, Z., Loy, C.C.: Rerender a video: Zero-shot text-guided video-to-video translation. In: SIGGRAPH Asia (2023)
56. Yang, Z., Teng, J., Zheng, W., Ding, M., Huang, S., Xu, J., Yang, Y., Hong, W., Zhang, X., Feng, G., et al.: Cogvideox: Text-to-video diffusion models with an expert transformer. arXiv preprint arXiv:2408.06072 (2024)
57. Ye, Z., Huang, H., Wang, X., Wan, P., Zhang, D., Luo, W.: Stylemaster: Stylize your video with artistic generation and translation. In: CVPR (2025)
58. YU, M., Hu, W., Xing, J., Shan, Y.: Trajectorycrafter: Redirecting camera trajectory for monocular videos via diffusion models. In: ICCV (2025)
59. Zhang, D.J., Paiss, R., Zada, S., Karnad, N., Jacobs, D.E., Pritch, Y., Mosseri, I., Shou, M.Z., Wadhwa, N., Ruiz, N.: Recapture: Generative video camera controls for user-provided videos using masked video fine-tuning. In: CVPR (2025)
60. Zhang, Y., Song, Y., Liu, J., Wang, R., Yu, J., Tang, H., Li, H., Tang, X., Hu, Y., Pan, H., et al.: Ssr-encoder: Encoding selective subject representation for subject-driven generation. In: CVPR (2024)
61. Zhang, Y., Yuan, Y., Song, Y., Wang, H., Liu, J.: Easycontrol: Adding efficient and flexible control for diffusion transformer. In: ICCV (2025)
62. Zhang, Z., Liao, J., Li, M., Dai, Z., Qiu, B., Zhu, S., Qin, L., Wang, W.: Tora: Trajectory-oriented diffusion transformer for video generation. In: CVPR (2025)