

FantasyHSI: Video-Generation-Centric 4D Human Synthesis in Any Scene Through a Graph-Based Multi-Agent Framework

Lingzhou Mu^{1*}, Qiang Wang^{2*}, Fan Jiang^{2†‡}, Mengchao Wang², Mu Xu², Kai Zhang^{1†},

¹Tsinghua University

²AMAP, Alibaba Group

mulz23@mails.tsinghua.edu.cn, {yijing.wq, frank.jf, wangmengchao.wmc, xumu.xm}@alibaba-inc.com, zhangkai@sz.tsinghua.edu.cn

Abstract

Human-Scene Interaction (HSI) seeks to generate realistic human behaviors within complex environments, yet it faces significant challenges in handling long-horizon, high-level tasks and generalizing to unseen scenes. To address these limitations, we introduce FantasyHSI, a novel HSI framework centered on video generation and multi-agent systems that operates without paired data. We model the complex interaction process as a dynamic directed graph, upon which we build a collaborative multi-agent system. This system comprises a scene navigator agent for environmental perception and high-level path planning, and a planning agent that decomposes long-horizon goals into atomic actions. Critically, we introduce a critic agent that establishes a closed-loop feedback mechanism by evaluating the deviation between generated actions and the planned path. This allows for the dynamic correction of trajectory drifts caused by the stochasticity of the generative model, thereby ensuring long-term logical consistency. To enhance the physical realism of the generated motions, we leverage Direct Preference Optimization (DPO) to train the action generator, significantly reducing artifacts such as limb distortion and foot-sliding. Extensive experiments on our custom SceneBench benchmark demonstrate that FantasyHSI significantly outperforms existing methods in terms of generalization, long-horizon task completion, and physical realism.

Code — <https://github.com/Fantasy-AMAP/fantasy-hsi>

Introduction

Human-Scene Interaction (HSI) aims to understand and generate human movements in response to complex environmental contexts. This field has garnered increasing attention in computer vision and graphics due to its significant potential across diverse applications, such as embodied intelligence, virtual reality, and games.

As general intelligent agents, humans can perform a wide range of complex interactive tasks, flexibly respond to observed environmental information, and rapidly adapt to new

surroundings. However, a significant gap remains between current methods and this level of human intelligence. Many approaches (Jiang et al. 2024b; Cen et al. 2024; Chen et al. 2024a; Jiang et al. 2024a) rely on paired human-environment data, which typically requires collecting extensive matched motion capture and scene data within specific environments. Consequently, they lack adaptability when faced with unseen object layouts or dynamic changes, struggling to cover the rich diversity of real-world interactions. While some methods (Li et al. 2024; Li and Dai 2024) attempt to bypass the reliance on paired datasets by leveraging the prior knowledge of Vision-Language Models (VLMs) (Zhu et al. 2023; Zhang et al. 2021) or video diffusion models (VDMs) (Liu et al. 2024a; Wan et al. 2025; Kong et al. 2024) to generate human-environment interaction sequences in a zero-shot manner, these are often limited to low-level, simple actions such as sitting or touching. They are ill-suited for high-level tasks, for instance, exploring a castle. Furthermore, generated motions must also be physically plausible. Any visual artifacts, such as limb deformation or foot-sliding, violate physical laws and severely diminish the realism and practical application of the results.

To address these challenges, we introduce FantasyHSI, a framework that models complex environmental scenes as a dynamic directed graph. By integrating VLM-based multi-agents with VDMs, FantasyHSI achieves effective environmental perception and planning, adjusts human motions based on environmental feedback, generates physically plausible human action sequences, and eliminates the dependency on paired human-environment datasets.

Specifically, we first introduce a dynamic directed graph to represent the human-environment interaction process. In this graph, nodes correspond to timestamped states of the human agent and the 3D scene, while edges encode the topological relationships of continuous action sequences. Building on this structure, we enhance the physical plausibility of VDMs using reinforcement learning, and leverage a motion capture system to acquire 4D spatio-temporal action sequences, which serve as the basis for dynamically generating the graph’s edges.

Subsequently, we design a multi-agents system. This system features a scene navigation agent for environmental per-

*These authors contributed equally.

†Fan Jiang and Kai Zhang are the corresponding authors

‡Project Lead



Figure 1: We introduce **FantasyHSI**, a novel framework that generates dynamic 4D sequences of humans interacting with their 3D environment. As illustrated on the left, FantasyHSI operates based on high-level task instruction, enabling it to autonomously plan paths, traverse obstacles, and execute a variety of complex motions, such as climbing a ladder. Moreover, the right side of the figure illustrates FantasyHSI’s ability to generalize to arbitrary scenes and a variety of actions.

ception and understanding, and a planner agent that performs high-level task decomposition, breaking down long-horizon objectives into primitive actions. Critically, to address the inherent stochasticity of generative models, we introduce a critic agent to form a closed-loop feedback loop, which quantifies the discrepancy between generated actions and the planned trajectory, enabling dynamic correction of deviating node states. This synergistic multi-agent architecture holistically unifies perception, planning, and correction, thereby resolving the issue of trajectory drift stemming from generative randomness and ensuring sustained logical coherence and physical viability in long-term interactions. Our primary contributions can be summarized as follows:

- We propose a novel method of long-horizon human-environment interaction using a dynamic directed graph, which establishes an interpretable foundation for perception, planning, and behavioral refinement.
- We develop a collaborative multi-agent system that integrates environmental perception, path planning, and closed-loop correction to rectify action deviations caused by the inherent stochasticity of generative models.
- We design a controllable, physics-enhanced action generator by optimizing VDMs with reinforcement learning, which significantly improves the physical realism of the generated actions.

Related Work

Human-Scene Interaction Synthesis

The objective of human-scene interaction synthesis is to generate realistic and coherent human motions that naturally interact with elements in a given environment. Early research (Tevet et al. 2022; Chen et al. 2023) predominantly focus on modeling human behavior in isolation, neglecting the critical influence of contextual environmental factors. However, as human actions inherently constitute responses

to external stimuli, recent approaches (Qing et al. 2023; Cen et al. 2024; Chen et al. 2024a; Lou et al. 2024) shift attention to scene-conditioned motion generation, incorporating action labels or textual descriptions as conditional inputs to enhance user controllability. By integrating environmental feedback and constraints, these methodologies achieve improved alignment with the dynamic nature of real-world interactions. Nevertheless, they typically rely on scarce paired action-scene data and exhibit limited generalization capabilities for novel interaction types. Although ZeroHSI (Li et al. 2024) and GenZI (Li and Dai 2024) propose the zero-shot interaction generation paradigm that broadens the scope of motion synthesis applications, current implementations remain constrained to singular simple tasks, coupled with insufficient awareness of physical laws in generated outputs.

Human Video Generation

Human video generation from a single still image of a person, utilizing driving signals such as video (Xie et al. 2024; Wang et al. 2025b), audio (Tian et al. 2024; Cui et al. 2025; Wang et al. 2025a), pose (Hu 2024; Zhu et al. 2024), or text (Chen et al. 2024b). Early approaches (Zhang et al. 2023; Ma et al. 2023) employ generative models like GANs (Goodfellow et al. 2020) or flow-based models (Ho et al. 2019). However, these methods often produce dynamic sequences suffering from artifacts and identity drift, resulting in insufficient realism and naturalness. Recently, diffusion models (Ho, Jain, and Abbeel 2020) gain prominence in this field due to their demonstrated high quality and stability in image and video generation tasks (Rombach et al. 2022; Esser et al. 2024). Trained on large-scale datasets (Nan et al. 2024; Li et al. 2025), diffusion models excel at modeling complex spatiotemporal relationships, yielding video results superior in visual quality and identity preservation. Nevertheless, current research (Wan et al. 2025; Kong et al. 2024) primarily focuses on generating portrait videos in real-world

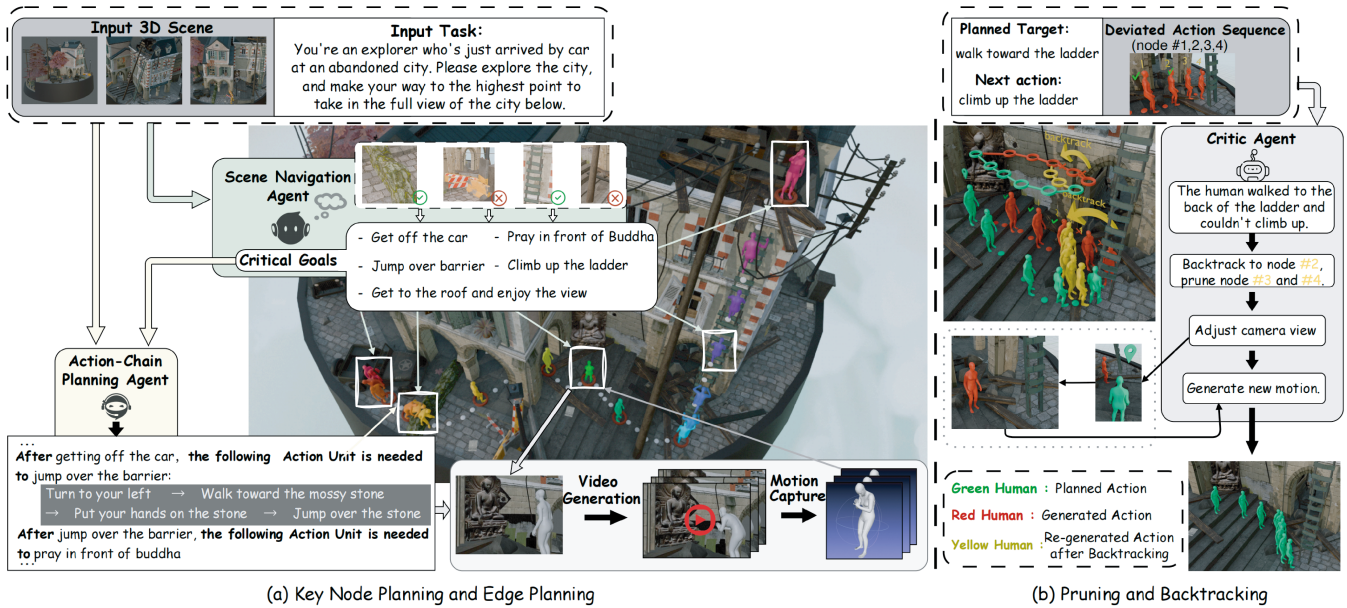


Figure 2: Overview of FantasyHSI.

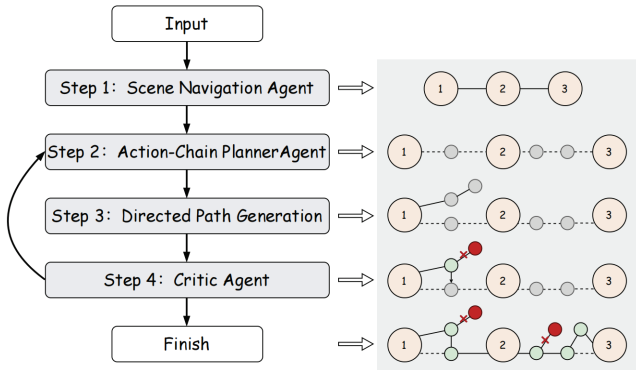


Figure 3: The Multi-Agent Pipeline.

settings. Their performance remains suboptimal when processing inputs like 3D scenes or SMPL-X (Pavlakos et al. 2019) style images, and the generated videos often fail to strictly adhere to physical laws and exhibit deficiencies in character consistency.

Reinforcement Learning from Human Feedback

Reinforcement Learning from Human Feedback (RLHF) (Bai et al. 2022) is a post-training methodology widely employed for large language models (Yuan et al. 2023; Dubey et al. 2024; Mehta et al. 2024; Stiennon et al. 2020) and diffusion models (Xu et al. 2024; Clark et al. 2023; Liu et al. 2025a), aiming to enhance model performance based on human feedback. Within this domain, Direct Preference Optimization (DPO) (Rafailov et al. 2023) is an effective method that trains the model using pairs of generated samples labeled as positive or negative. It enables the model to assign higher probabilities to preferred outputs and lower probab-

ilities to less favored ones. By directly comparing different outputs generated by the model, DPO can more efficiently capture human preference information. DPO has been extensively applied to video diffusion models to improve visual quality, enhance motion coherence, and ensure spatio-temporal consistency, with its effectiveness validated in numerous studies (Liu et al. 2025b; Wu et al. 2025; Wallace et al. 2024). In our work, we employ the DPO approach to optimize video generation models, enabling them to produce content that better adheres to physical laws.

Method

As shown in Figure 2, given the 3D scene and high-level instructions, we first formalize the task as a dynamic directed graph, followed by task decomposition, planning, backtracking, and correction through a VLM-based multi-agents. We employ reinforcement learning to enhance the physical laws of the generator for each edge in the graph. Full details of these components are provided in the corresponding sections.

Dynamic Directed Graph Representation

Graph Construction. To provide an interpretable representation for multi-agents, we model the task as a directed graph $G = (\mathcal{N}, E)$, where \mathcal{N} is the set of nodes and E denotes the set of directed edges. Each node $N_k = \{S_k, \mathcal{H}_k\} \in \mathcal{N}$ represents the state of the human and 3D scene at a specific time point. The human state \mathcal{H} and the scene state S are represented by the 3D mesh. We employ SMPL-X (Pavlakos et al. 2019) to model human pose and movement.

Furthermore, since some nodes represent the attainment of a critical goal, such as reaching the mountaintop or exiting the farm, we partition the node set \mathcal{N} into two categories:

a set of key nodes $\mathcal{K} = \{K_1, K_2, \dots, K_p\} \subset \mathcal{N}$, which indicate signify milestone achievements, and a set of non-key nodes $\mathcal{U} = \{U_1, U_2, \dots, U_q\} = \mathcal{N} \setminus \mathcal{K}$, which correspond to the completion of individual action units but the critical goal remains unachieved. Thus, a sequence of adjacent non-key nodes and corresponding edges can form a directed path that connects two key nodes, thereby representing the process by which a human character completes another critical task within the 3D scene via performing a series of action units, culminating in the desired state denoted by the newly reached key node.

Key Node Definition. A directed edge $E(N_i, N_j) = A_{i,j}$ indicates that the human \mathcal{H}_i in node N_i performs an action $A_{i,j}$, resulting in a new state described by node N_j . We formulate this state transition as $\mathcal{H}_j = \mathcal{H}_i + A_{i,j}$. Here, $A_{i,j}$ denotes an action unit that has complete semantic meaning.

Given that human motion with long duration actually consists of continuous, frame-by-frame action sequences, naively defining human-scene state of each frame as an individual node would lead to excessive graph complexity. Thereby, only the state of the start and the end of an action unit $A_{i,j}$ is defined as node N_i and N_j . In this manner, the human state before and after each action unit — \mathcal{H}_i and \mathcal{H}_j —together with their associated scene states \mathcal{S}_i and \mathcal{S}_j , constituted the adjacent node pair $N_i = \{\mathcal{H}_i, \mathcal{S}_i\}$ and $N_j = \{\mathcal{H}_j, \mathcal{S}_j\}$.

VLM-based Multi-Agents

Overview. As illustrated in Figure 3, upon receiving a high-level task for the human in the scene, the scene navigation agent first analyzes the 3D scene and identify critical sub-goals required to accomplish the high-level task. Then it formulate a comprehensive plan that integrates both spatial trajectories and critical sub-goals to generate the key nodes \mathcal{K} in the graph G . Subsequently, the action-chain planning agent generates a sequence of text described action chain $\{A_{i,i+1}, A_{i+1,i+2}, \dots, A_{j-1,j}\}$, constructing a action chain connecting adjacent key nodes. Next, the generator in Section synthesizes human actions according to realize each action unit planned by the action-chain planning agent to construct a directed path. Due to the inherent randomness of the generative model, this process may introduce new, unplanned nodes into the original directed graph as the generated actions can differ from those initially planned. When such unplanned nodes arise, the critic agent analyses the nodes to route video generator return to the planned key node gradually when generating subsequent actions. This mechanism enables the agent to backtrack to wanted nodes, prune erroneous nodes, while refining viable ones.

Key Node Planning via Scene Navigation Agent. The scene navigation agent is responsible for identifying critical sub-goals required to accomplish the given high-level task and generating a comprehensive plan that contains both trajectories and key events based on natural language description of high-level task T , the initial human position \mathcal{H}_0 , and the initial 3D scene state \mathcal{S}_0 . This plan is represented as a sequence of key nodes \mathcal{K} in the graph G , as discussed in Section .

The detailed reasoning process of the scene navigation agent is provided in Appendix A2.

Edge Planning via Action-Chain Planner Agent. Based on the key nodes and trajectory generated by the scene navigation agent, the action-chain planner agent decomposes the motion required to accomplish each sub-goals into a sequence of action units, each described in natural language. Here, an action unit is defined as a minimal semantic motion unit, representing a semantically coherent movement within three seconds. While complex human motions with long duration can be decomposed into short, meaningful action units, however, these atomic units are not a finite set due to the vast variability of human motion in real-world scenarios. In our work, these action units may represent either complex human behaviors (e.g., “yawning”) or simple mechanical motions (e.g., “turning backward”).

While decomposing the motion into a sequence of more detailed action units, the action-chain planner agent actually extends the graph G by adding intermediate non-key nodes \mathcal{U} and a set of edges E between adjacent key nodes \mathcal{K} . These added non-key nodes and edges can model state transitions between key nodes via a chain of action units, thereby bridging high-level planning with low-level motion execution.

Directed Path Generation. The video generation model enhanced by physical laws in Section serves as a human simulator in our framework, instantiate directed edges between all nodes to form a directed path that guides the agent from its initial state to each key goals planned by the scene navigation agent to accomplish the overall task. Our approach first generates videos for each action unit using a text-conditioned image-to-video model, then lifts the human motion into 3D scene via motion capture (Yin et al. 2024). By iteratively rendering the final 3D state of the captured action and the scene as the initial frame for the next video generation, our approach can construct a continuous, scene-aware action chain, enabling the virtual agent to perform long-horizon, open-ended tasks in arbitrary environments.

Specifically, we first render a snapshot of the human \mathcal{H}_i in scene \mathcal{S}_i of the current node N_i as the first frame of our video generation model, and use detailed description of the action unit $A_{i,j}$ in the graph as a prompt to generate a video clip. We then apply motion capture (Yin et al. 2024) to the generated video, extracting a 3D motion sequence in SMPL-X format. To lift the motion back to the 3D scene, we apply the motion change of each frame to the virtual human in the 3D scene. After completed the action unit $A_{i,j}$, the human-scene state is updated from node N_i to N_j through the instantiated edge $E(N_i, N_j) = A_{i,j}$. This newly updated node N_j is then rendered as the next video’s initial frame, enabling iterative generation of subsequent actions. By repeating this process, we construct a directed path $\{N_0, E_{0,1}, N_1, E_{1,2}, \dots, N_k\}$ within the graph, enabling the virtual human to execute the actions and tasks in 3D space as planned by our agents.

Pruning and Backtracking through Critic Agent. Due to inherent randomness in video generation process and the inevitable ambiguity of language prompt, the generated mo-

Method	P-Score↓	FS↓	CLIP-S↑	CLIP-C↑	Diversity↑	POS↓	RDS↑
TRUMANS	0.083	0.156	0.19	0.9943	0.1703	0.126	0.1806
LINGO	0.085	0.148	0.20	0.9931	0.1510	0.133	0.1592
PedGen	0.063	0.199	0.17	0.9928	0.1102	0.119	0.1290
Ours	0.020	0.102	0.31	<u>0.9937</u>	0.1911	0.063	0.3285
SFT	0.034	0.124	<u>0.27</u>	0.9929	0.1221	<u>0.069</u>	<u>0.3104</u>
Baseline	0.051	0.149	0.23	0.9865	0.1283	0.085	0.2173
w/o multi-agents	0.022	<u>0.103</u>	0.24	0.9932	0.1920	0.070	0.2943
w/o critic agent	<u>0.021</u>	<u>0.105</u>	0.23	0.9935	<u>0.1916</u>	0.082	0.2891

Table 1: Quantitative Results of Comparison and Ablation Studies.

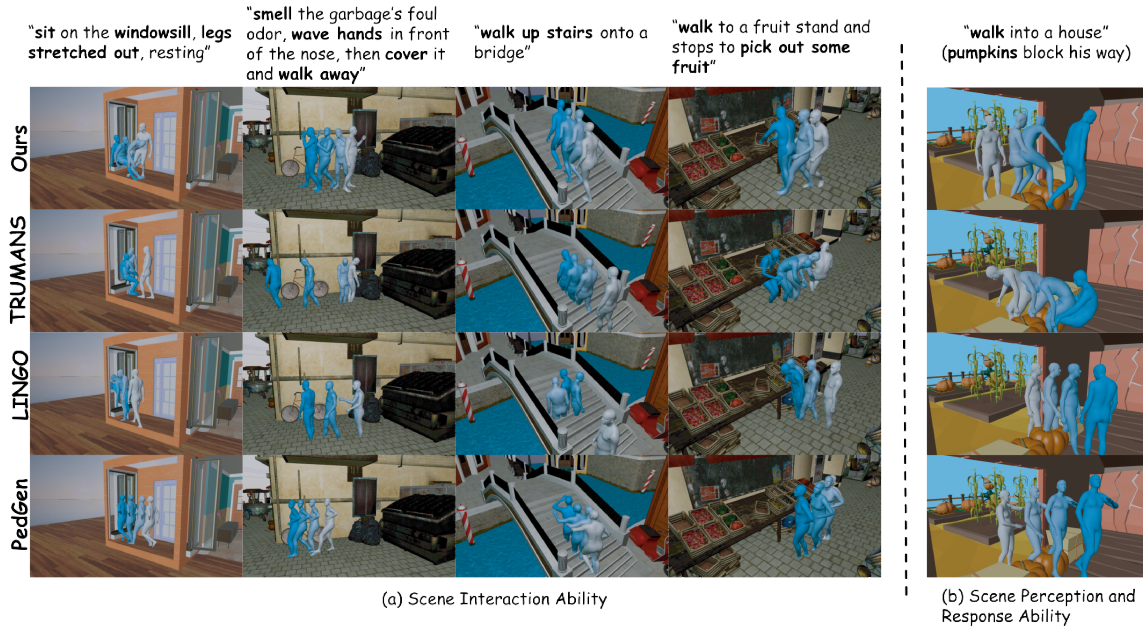


Figure 4: Qualitative Comparison.

tion sequences can sometime lead the virtual human to deviate from the route planned by agents. For example, when generating a video clip of a person walking while enjoying the scenery, the distance traveled and the direction of movement can only be roughly controlled. Moreover, given an text described action like *a sleepy stretch*, the video generation model might generate extra actions like an additional yawn following the stretch to express the person’s sleepiness. Thus, this introduces new nodes into the graph constructed by the scene navigation agent and action-chain planner agent. In some cases, these extra actions and deviations enhance the expressiveness of the overall actions, but they might also be unwanted and disrupt the plan.

To process those newly generated nodes with deviations and unplanned actions, we utilize a critic agent that first assess these new nodes and then apply corrections if needed. Specifically, for each generated and captured motion segment, the critic agent analyzes the corresponding rendered frames, evaluates the motion, and corrects the trajectory and posture. Details of the evaluation and correction process can be found in Appendix A3.

Physical Law Enhancement for Generator

We employ reinforcement learning to enhance the capabilities of video diffusion models in generating physically plausible motion and accurately following instructions. Specifically, to improve performance in instruction following, motion artifacts (including clipping, unnatural motions), limb deformation, and scene inconsistency, we generate samples using four models: VEO (DeepMind 2025), HunYuan-Video (Kong et al. 2024), Runway (Runway AI, Inc. 2025), and Kling (Kling AI Team 2025). Professional annotators then label positive samples x^w and negative samples x^l based on these criteria. We utilize DPO (Rafailov et al. 2023) to train the open-source Wan (Wan et al. 2025) model, thereby improving its ability to generate videos with enhanced physical realism and quality. Following DiffusionDPO (Wallace et al. 2024), the training loss is defined as:

$$\mathcal{L}_{\text{DPO}} = -\mathbb{E} \left[\log \sigma \left(-\frac{\beta}{2} (L(x^w, p) - L(x^l, p)) \right) \right] \quad (1)$$

where β is a temperature coefficient, $L(x^w, p)$ and $L(x^l, p)$ represent the losses for positive and negative parts, re-

spectively. This loss encourages the generation of samples aligned with human preferences, promoting videos of human characters that are more physically plausible and consistent with reality.

Experiment

Implementation Details

For video generator enhanced by physical laws using DPO, we utilize Wan2.1-I2V-14B (Wan et al. 2025) as base model. Then the model was trained on a collected dataset of 10,000 preference pairs for approximately 20 hours using 8 A100 GPUs, with a learning rate set to $1e-5$. The β is set to 5000. We employ Gemini-2.5-Pro (Google DeepMind 2025) as the VLM-based multi-agents.

Evaluation Settings

Evaluate environmental Settings. To systematically evaluate our method, we conduct experiments under two settings, comprising scene interaction evaluation and scene perception and response evaluation. The *scene interaction evaluation* assesses the model’s ability to generate plausible human-scene interactions in static environments, where the scene geometry remains fixed throughout the motion. The *scene perception and response evaluation* measures the model’s ability to perceive and react to the changes and obstacles in the environment. In this setting, we introduce common real-world obstacle, both seen (e.g., chairs, sofas, vases) and novel (e.g., pumpkins, rocks), into the model’s pre-planned path. The model must first detect the obstacle and then react to it accordingly. This evaluates not only robustness to unseen objects but more importantly evaluates how the model perceive the world and react to the world.

Evaluation Dataset. Due to the lack of publicly available HSI benchmarks, systematic evaluation remains challenging. For instance, TRUMANS (Jiang et al. 2024a) only released its training dataset without a standardized test set, and other works such as LINGO (Jiang et al. 2024a) have not yet made their evaluation sets publicly available. Therefore, we introduce SceneBench, an evaluation benchmark comprising diverse 3D environments designed to assess embodied virtual human behavior in indoor and outdoor scene. The scenes are sourced from TRUMANS, and our web-collected scenes from Sketchfab (Sketchfab 2023), resulting in a total of 20 distinct 3D scene, 10 indoor and 10 outdoor, spanning residential spaces (e.g., bedrooms, cowshed, gym), natural landscapes (e.g., grasslands, riversides), urban streets, and rural farms, etc.

For scene interaction evaluation, each scene is annotated with human-scene interaction tasks described in natural language, along with their start and goal positions. This yields 120 text-scene-position pairs as the test instances for evaluation. For scene perception and response evaluation, we sample 15 objects from SceneBench and 15 additional 3D models collected from the Internet. These include both seen household objects from TRUMANS and unseen, out-of-distribution objects. We place them along pre-planned paths to evaluate how models react to them. (See Appendix B1. for more details.)

Baselines. We compare our method with recent HSI methods, as well as a scene-aware 4D pedestrian generation method, including TRUMANS (Jiang et al. 2024b), LINGO (Jiang et al. 2024a) and PedGen (Liu et al. 2024b).

Unlike baseline methods that require predefined path specifications, our multi-agents system operates end-to-end, generating both path and motion plans directly from higher-level natural language instructions. Specifically TRUMANS needs a full trajectory input, while LINGO and PedGen requires start and goal positions. In our experiments, we provide these additional input using ground-truth locations to adapt the baseline methods to our task and thus fully demonstrate the capabilities of the baseline methods.

Metrics. For evaluation, we adopt metrics from prior works (Jiang et al. 2024b; Liu et al. 2024b; Li et al. 2024) to assess both scene interaction and motion quality. To evaluate scene interaction, we employ the Penetration Score (**P-Score**) (Zhao et al. 2023) to measure the percentage of body vertices penetrating the scene, and the Foot Sliding Score (**FS**) (He et al. 2022) to quantify foot sliding. To assess motion quality and diversity, we utilize the **CLIP-Score** (Radford et al. 2021) for measuring the alignment between the generated motion and the input text prompt, alongside **CLIP-Consistency** to evaluate frame-wise temporal coherence, motion **Diversity** to evaluate diversity under same inputs. In addition, to evaluate scene awareness, we introduce two metrics. The Penetration Obstacle Score (**POS**) measures mesh penetration between the human and introduced obstacles, while the Reaction Divergence Score (**RDS**) computes the average per-joint distance between motions generated with and without obstacles.

Evaluation results

Scene Interaction Ability Evaluation. As shown in Figure 4(a), we present qualitative comparison of scene interaction ability between FantasyHSI and baseline approaches on SceneBench. Our results demonstrate that our method generates vivid and expressive motions in various environments, accomplishing diverse high-level human-scene interactions beyond simple walking or touching. For instance, our method can produces highly abstract and human-like behaviors such as fanning one’s nose near a garbage pile, sitting on unusual places like windowsills, and even climb 20 meters up the ladder to reach the rooftop in Figure 1, while all of the other methods failed in these tasks.

Qualitative results indicate that TRUMANS is severely overfitted to its training distribution, as it defaults to generating only sitting motions when encountering novel objects. As shown in the first column of Figure 4(a), it fails to perceive the windowsill’s height, instead generating sitting poses at a standard chair height consistent with its training data. Furthermore, LINGO struggles to perceive surface boundaries in unseen environments, such as the third and fourth columns in Figure 4(a), and present limited scene understanding. As a result, it suffers severe penetration and are unable to produce reasonable motions for highly abstract interaction tasks in column 2. Although PedGen generates temporally coherent walking sequences, its motion diver-

sity is highly limited and lacks the capability to perform meaningful scene interactions. As evidenced in Table 1, our method achieves the highest CLIP Score and diversity with the lowest penetration and FS, outperforming existing approaches across most metrics. This demonstrates that our method generates more semantically aligned and physically plausible motions with better diversity.

Scene Perception and Response Ability Evaluation. For the evaluation of scene perception and response capability, we present qualitative comparison results in Figure 4(b). Among all these methods, only our approach successfully perceived the obstacles (the pumpkins) and generated reasonable reactive behaviors, such as stepping over them. Although TRUMANS and LINGO detect the presence of obstacles via occupancy grids, their perception range is limited to a 1-meter box around the virtual human, represented by point clouds. The point clouds of surrounding objects are truncated by this limited sensing range, preventing the models from perceiving the full structure of the objects and leading to significant loss of semantic information. As a result, LINGO generates a backward glance, while TRUMANS fail to generate reasonable motions, neither successfully avoids or interacts with the obstacle. In contrast, PedGen demonstrates very poor obstacle awareness, simply walking through the pumpkins without any reactive behavior. Align with visual observations, as shown in Table 1, our method also outperforms all methods on Penetration Obstacle Score and Reaction Divergence Score, indicating superior scene understanding and response ability.

Ablation Studies

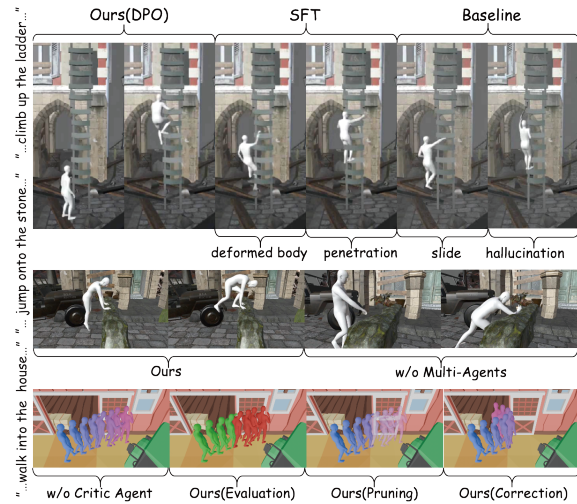


Figure 5: Qualitative Results of Ablation Studies.

Multi-Agents. To evaluate the effectiveness of our multi-agent collaborative framework, we conducted an ablation study in which no agent was involved in neither planning the actions or decomposing complex actions into a chain of action units. In this setting, complex motions were generated directly using video generation model. As shown in

the second row of Figure 5, when prompted to jump on to the fence, the model fails to generate desired actions without multi-agents, since it lacks of detailed motion planning as instructions. By contrast, our approach break down the complex motion into a chain of Action Units. With this detailed plan, the virtual human is first instructed to places his hands on the rock for support, and then the he jumps up and lands with both feet on top of the rock, thereby successfully completing the overall action. Additionally, as shown in Table 1, the significant drop in the CLIP-S score indicates that our method fails to accomplish the task objective without the multi-agent component to decompose the main goal into clear sub-tasks. This result validates the effectiveness of multi-agents.

Critic Agent. To evaluate the effectiveness of the critic agent in our method, we conducted ablation studies comparing results with and without its involvement. As shown in Figure 5, without the evaluation and backtracking mechanisms provided by the critic agent, the model is unable to correct deviations from the intended path, ultimately failing to reach the planned target location. In contrast, when the critic agent is incorporated, the model successfully re-navigates the virtual human to the target position. Furthermore, as evidenced by Table 1, the absence of the critic agent causes a significant drop in CLIP score, indicating that the model struggles to accomplish the specified goals. The observed increase in the Diversity metric is primarily due to the inclusion of deviated motion segments that would otherwise be backtracked and pruned by the critic agent.

DPO for Video Generative Model. We conduct experiments to validate the effectiveness of our DPO optimized model over Supervised Fine-Tuning (SFT) model and the original model. As shown in Figure 5 and Table 1, while the base model and the SFT method demonstrate a certain degree of instruction-following capability, their generated outputs often exhibit physically implausible dynamics. These include artifacts such as characters penetrating the environment, body distortions, and unnatural sliding motions. In contrast, our method, optimized with the DPO, significantly enhances the ability to generate dynamics consistent with real-world physics, thereby achieving superior results.

Conclusion

In this work, we presented FantasyHSI, a novel framework for synthesizing expressive and physically plausible human-scene interactions in complex 3D environments. By reformulating HSI as a dynamic directed graph, we established an interpretable structure for modeling long-horizon interactions. The integrated VLM-based multi-agent collaboration comprise scene understanding, hierarchical planning, and trajectory correction. Furthermore, our reinforcement learning-based optimization of video diffusion models ensures that synthesized motions adhere to physical laws, eliminating artifacts such as foot sliding and body-scene penetration. Experiments show that FantasyHSI surpasses existing methods in generalization to unseen scenes and long-horizon tasks while maintaining motion realism and logical coherence.

References

- Bai, Y.; Jones, A.; Ndousse, K.; Askell, A.; Chen, A.; DasSarma, N.; Drain, D.; Fort, S.; Ganguli, D.; Henighan, T.; et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Cen, Z.; Pi, H.; Peng, S.; Shen, Z.; Yang, M.; Zhu, S.; Bao, H.; and Zhou, X. 2024. Generating human motion in 3d scenes from text descriptions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1855–1866.
- Chen, J.; Hu, P.; Chang, X.; Shi, Z.; Kampffmeyer, M.; and Liang, X. 2024a. Sitcom-crafter: A plot-driven human motion generation system in 3d scenes. *arXiv preprint arXiv:2410.10790*.
- Chen, X.; Jiang, B.; Liu, W.; Huang, Z.; Fu, B.; Chen, T.; and Yu, G. 2023. Executing your commands via motion diffusion in latent space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 18000–18010.
- Chen, X.; Liu, Z.; Chen, M.; Feng, Y.; Liu, Y.; Shen, Y.; and Zhao, H. 2024b. Livephoto: Real image animation with text-guided motion control. In *European Conference on Computer Vision*, 475–491. Springer.
- Clark, K.; Vicol, P.; Swersky, K.; and Fleet, D. J. 2023. Directly fine-tuning diffusion models on differentiable rewards. *arXiv preprint arXiv:2309.17400*.
- Cui, J.; Li, H.; Zhan, Y.; Shang, H.; Cheng, K.; Ma, Y.; Mu, S.; Zhou, H.; Wang, J.; and Zhu, S. 2025. Hallo3: Highly dynamic and realistic portrait image animation with video diffusion transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 21086–21095.
- DeepMind. 2025. Veo: DeepMind’s Generative Video Model. Accessed: 2025-07-01.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv e-prints*, arXiv–2407.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.
- Google DeepMind. 2025. Gemini Pro. <https://deepmind.google/models/gemini/pro/>. Accessed: 2025-07-01.
- He, C.; Saito, J.; Zachary, J.; Rushmeier, H.; and Zhou, Y. 2022. Nemf: Neural motion fields for kinematic animation. *Advances in Neural Information Processing Systems*, 35: 4244–4256.
- Ho, J.; Chen, X.; Srinivas, A.; Duan, Y.; and Abbeel, P. 2019. Flow++: Improving flow-based generative models with variational dequantization and architecture design. In *International conference on machine learning*, 2722–2730. PMLR.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Hu, L. 2024. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8153–8163.
- Jiang, N.; He, Z.; Wang, Z.; Li, H.; Chen, Y.; Huang, S.; and Zhu, Y. 2024a. Autonomous character-scene interaction synthesis from text instruction. In *SIGGRAPH Asia 2024 Conference Papers*, 1–11.
- Jiang, N.; Zhang, Z.; Li, H.; Ma, X.; Wang, Z.; Chen, Y.; Liu, T.; Zhu, Y.; and Huang, S. 2024b. Scaling up dynamic human-scene interaction modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1737–1747.
- Kling AI Team. 2025. Kling AI: Text-to-Video Generation Platform. <https://app.klingai.com/>. Accessed: [Insert Access Date Here].
- Kong, W.; Tian, Q.; Zhang, Z.; Min, R.; Dai, Z.; Zhou, J.; Xiong, J.; Li, X.; Wu, B.; Zhang, J.; et al. 2024. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*.
- Li, H.; Xu, M.; Zhan, Y.; Mu, S.; Li, J.; Cheng, K.; Chen, Y.; Chen, T.; Ye, M.; Wang, J.; et al. 2025. Openhumanvid: A large-scale high-quality dataset for enhancing human-centric video generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 7752–7762.
- Li, H.; Yu, H.-X.; Li, J.; and Wu, J. 2024. Zerohsi: Zero-shot 4d human-scene interaction by video generation. *arXiv preprint arXiv:2412.18600*.
- Li, L.; and Dai, A. 2024. Genzi: Zero-shot 3d human-scene interaction generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20465–20474.
- Liu, J.; Liu, G.; Liang, J.; Yuan, Z.; Liu, X.; Zheng, M.; Wu, X.; Wang, Q.; Qin, W.; Xia, M.; et al. 2025a. Improving video generation with human feedback. *arXiv preprint arXiv:2501.13918*.
- Liu, R.; Wu, H.; Zheng, Z.; Wei, C.; He, Y.; Pi, R.; and Chen, Q. 2025b. Videodpo: Omni-preference alignment for video diffusion generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 8009–8019.
- Liu, Y.; Zhang, K.; Li, Y.; Yan, Z.; Gao, C.; Chen, R.; Yuan, Z.; Huang, Y.; Sun, H.; Gao, J.; et al. 2024a. Sora: A review on background, technology, limitations, and opportunities of large vision models. *arXiv preprint arXiv:2402.17177*.
- Liu, Z.; Lin, J.; Wu, W.; and Zhou, B. 2024b. Learning to Generate Diverse Pedestrian Movements from Web Videos with Noisy Labels. In *The Thirteenth International Conference on Learning Representations*.
- Lou, Z.; Cui, Q.; Wang, T.; Song, Z.; Zhang, L.; Cheng, C.; Wang, H.; Tang, X.; Li, H.; and Zhou, H. 2024. Harmonizing stochasticity and determinism: Scene-responsive diverse human motion prediction. *Advances in Neural Information Processing Systems*, 37: 39784–39811.
- Ma, Y.; Zhang, S.; Wang, J.; Wang, X.; Zhang, Y.; and Deng, Z. 2023. Dreamtalk: When expressive talking head generation meets diffusion probabilistic models. *arXiv preprint arXiv:2312.09767*, 2(3).
- Mehta, S.; Sekhvat, M. H.; Cao, Q.; Horton, M.; Jin, Y.; Sun, C.; Mirzadeh, I.; Najibi, M.; Belenko, D.; Zatloukal, P.; et al. 2024. Openelm: An efficient language model family with open training and inference framework. *arXiv preprint arXiv:2404.14619*.
- Nan, K.; Xie, R.; Zhou, P.; Fan, T.; Yang, Z.; Chen, Z.; Li, X.; Yang, J.; and Tai, Y. 2024. Openvid-1m: A large-scale high-quality dataset for text-to-video generation. *arXiv preprint arXiv:2407.02371*.
- Pavlakos, G.; Choutas, V.; Ghorbani, N.; Bolkart, T.; Osman, A. A.; Tzionas, D.; and Black, M. J. 2019. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10975–10985.
- Qing, Z.; Cai, Z.; Yang, Z.; and Yang, L. 2023. Story-to-motion: Synthesizing infinite and controllable character animation from

- long text. In *SIGGRAPH Asia 2023 technical communications*, 1–4.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36: 53728–53741.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Runway AI, Inc. 2025. Runway.
- Sketchfab. 2023. Sketchfab - Explore, Buy, and Share 3D Models. <https://sketchfab.com/>. Accessed: 2025-06-13.
- Stiennon, N.; Ouyang, L.; Wu, J.; Ziegler, D.; Lowe, R.; Voss, C.; Radford, A.; Amodei, D.; and Christiano, P. F. 2020. Learning to summarize with human feedback. *Advances in neural information processing systems*, 33: 3008–3021.
- Tevet, G.; Raab, S.; Gordon, B.; Shafir, Y.; Cohen-Or, D.; and Bermano, A. H. 2022. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*.
- Tian, L.; Wang, Q.; Zhang, B.; and Bo, L. 2024. Emo: Emote portrait alive generating expressive portrait videos with audio2video diffusion model under weak conditions. In *European Conference on Computer Vision*, 244–260. Springer.
- Wallace, B.; Dang, M.; Rafailov, R.; Zhou, L.; Lou, A.; Purushwalkam, S.; Ermon, S.; Xiong, C.; Joty, S.; and Naik, N. 2024. Diffusion model alignment using direct preference optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8228–8238.
- Wan, T.; Wang, A.; Ai, B.; Wen, B.; Mao, C.; Xie, C.-W.; Chen, D.; Yu, F.; Zhao, H.; Yang, J.; et al. 2025. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*.
- Wang, M.; Wang, Q.; Jiang, F.; Fan, Y.; Zhang, Y.; Qi, Y.; Zhao, K.; and Xu, M. 2025a. Fantasytalking: Realistic talking portrait generation via coherent motion synthesis. *arXiv preprint arXiv:2504.04842*.
- Wang, Q.; Wang, M.; Jiang, F.; Fan, Y.; Qi, Y.; and Xu, M. 2025b. FantasyPortrait: Enhancing Multi-Character Portrait Animation with Expression-Augmented Diffusion Transformers. *arXiv preprint arXiv:2507.12956*.
- Wu, Z.; Kag, A.; Skorokhodov, I.; Menapace, W.; Mirzaei, A.; Gilitschenski, I.; Tulyakov, S.; and Siarohin, A. 2025. DenseDPO: Fine-Grained Temporal Preference Optimization for Video Diffusion Models. *arXiv preprint arXiv:2506.03517*.
- Xie, Y.; Xu, H.; Song, G.; Wang, C.; Shi, Y.; and Luo, L. 2024. X-portrait: Expressive portrait animation with hierarchical motion attention. In *ACM SIGGRAPH 2024 Conference Papers*, 1–11.
- Xu, J.; Huang, Y.; Cheng, J.; Yang, Y.; Xu, J.; Wang, Y.; Duan, W.; Yang, S.; Jin, Q.; Li, S.; et al. 2024. Visionreward: Fine-grained multi-dimensional human preference learning for image and video generation. *arXiv preprint arXiv:2412.21059*.
- Yin, W.; Cai, Z.; Wang, R.; Wang, F.; Wei, C.; Mei, H.; Xiao, W.; Yang, Z.; Sun, Q.; Yamashita, A.; et al. 2024. Whac: World-grounded humans and cameras. In *European Conference on Computer Vision*, 20–37. Springer.
- Yuan, H.; Yuan, Z.; Tan, C.; Wang, W.; Huang, S.; and Huang, F. 2023. Rrhf: Rank responses to align language models with human feedback. *Advances in Neural Information Processing Systems*, 36: 10935–10950.
- Zhang, P.; Li, X.; Hu, X.; Yang, J.; Zhang, L.; Wang, L.; Choi, Y.; and Gao, J. 2021. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5579–5588.
- Zhang, W.; Cun, X.; Wang, X.; Zhang, Y.; Shen, X.; Guo, Y.; Shan, Y.; and Wang, F. 2023. Sadtalker: Learning realistic 3d motion coefficients for stylized audio-driven single image talking face animation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8652–8661.
- Zhao, K.; Zhang, Y.; Wang, S.; Beeler, T.; and Tang, S. 2023. Synthesizing diverse human motions in 3d indoor scenes. In *Proceedings of the IEEE/CVF international conference on computer vision*, 14738–14749.
- Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2023. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.
- Zhu, S.; Chen, J. L.; Dai, Z.; Dong, Z.; Xu, Y.; Cao, X.; Yao, Y.; Zhu, H.; and Zhu, S. 2024. Champ: Controllable and consistent human image animation with 3d parametric guidance. In *European Conference on Computer Vision*, 145–162. Springer.