

GraphVid: Interactive Graph-Controllable Video Generation

Vedant Shah¹, Onkar Susladkar¹, Tushar Prakash², Kiet A. Nguyen¹, Tianjiao Yu¹, Adheesh Juvekar¹, Muntasir Wahed¹, and Ismini Lourentzou¹

¹University of Illinois Urbana-Champaign

²Sony Research India

{vrshah4,onkarks2,lourent2}@illinois.edu




Fig. 1: GraphVid enables controllable multi-object image-to-video generation through interaction graphs. Starting from an input frame, GraphVid constructs a scene graph that represents entities and their interactions. Users can directly modify these interactions by interacting with the image to specify desired dynamics. The edited interaction graph is then translated into conditioning tokens that guide a frozen video diffusion backbone to generate a coherent video sequence.

Abstract. Controllable video generation remains challenging due to the difficulty of specifying precise multi-object interactions using text prompts or motion-control inputs that primarily constrain pixel movement. In practice, trajectory-based control often requires users to draw accurate tracks for multiple objects, which scales poorly with scene complexity and becomes ambiguous under occlusion or overlap. To enable flexible yet precise multi-subject control, we introduce **GraphVid**, a graph-conditioned image-to-video generation model that enables interactive control through structured interaction graphs. We further curate **GRAPHVID-BENCH**, a large-scale interaction-centric video dataset with structured relational annotations to enable training of interaction-aware video generation models. Despite using substantially less training data and fewer trainable parameters than prior motion-control methods, GraphVid delivers strong controllability and video quality. Compared with Motion-I2V, GraphVid reduces FID by up to 39.9% and FVD by 37.6%, while improving PSNR (9.87→15.98) and SSIM (0.38→0.61). Our

results highlight the potential of structured semantic interfaces as a powerful paradigm for controllable video generation.

 PLAN Lab <https://plan-lab.github.io/graphvid>

Keywords: Controllable Video Generation · Scene Graphs

1 Introduction

Image-to-Video (I2V) generation has attracted increasing attention in generative modeling [5, 34, 42, 43, 45]. Given a single image, the goal is to synthesize a temporally coherent video that preserves visual identity while producing realistic motion. This task is inherently challenging since the model must infer plausible temporal dynamics from a single static frame, where motion, interactions, and physical evolution are fundamentally ambiguous. As a result, generating videos that are both temporally consistent and physically plausible remains difficult.

Recent work on controllable I2V generation has largely focused on trajectory-conditioned motion control [8, 29, 42, 53]. Methods such as Motion Prompting [10], Wan-Move [8], MagicMotion [32], and Motion-I2V [42] guide diffusion-based video models using explicit motion signals, including point trajectories, optical flow, segmentation masks, or sparse bounding boxes. These signals are transformed into dense motion representations and injected into the generative backbone through mechanisms such as ControlNet [58] adapters, latent feature manipulation, or motion-aware attention, enabling user-specified object and camera motion while preserving appearance. Among these, Wan-Move achieves strong performance by directly modifying latent conditioning features of a large-scale Diffusion Transformer (DiT) [37] model and training on millions of videos, improving motion fidelity and generalization. However, these methods treat motion primarily as geometric displacement, overlooking the physical causes governing object dynamics. To address this, a few physics-aware approaches introduce control signals based on forces or velocity fields to model object interactions [12, 39], jointly model visual and latent physical dynamics [41], or incorporate structured physical knowledge by injecting textual descriptions, qualitative categories, and quantitative physical properties through physics-aware attention [49].

Despite recent progress, existing approaches for controllable video generation remain limited in how they model scene dynamics: **(1)** Current physics-aware approaches introduce control signals based on forces, velocities, or structured physical annotations [12, 39, 49]. However, these methods typically encode physics as low-level control fields or predefined categories, which limits their ability to represent complex interactions between multiple entities in open-domain scenes. **(2)** Methods that condition on low-level geometric cues can make control cumbersome in multi-object scenarios, as users must author accurate signals for each object, and small errors in a few controls can yield globally implausible motion (*e.g.*, inconsistent contact, sliding, penetrations, or desynchronized behavior). **(3)** Moreover, both trajectory-based and physics-conditioned methods rely heavily on large-scale curated datasets containing dense motion annotations, trajectories, optical flows, or physics labels, which are expensive to obtain

and difficult to scale across diverse environments. For instance, Wan-Move [8] is trained on 2M samples, Motion Prompting [10] uses 2.2M, and Motion-I2V [42] is trained on 10M samples. This dependence restricts the practicality and generalization of existing controllable video generation pipelines.

To address these limitations and enable flexible interactive multi-object control, we introduce **GraphVid**, a graph-conditioned image-to-video generation framework that represents motion control through directed interaction scene graphs. Starting from an input image, GraphVid constructs a scene graph over detected entities with relational edges describing interactions (Figure 1). Users interact directly with the image to specify desired dynamics, and the edited interaction graph is then translated into conditioning tokens that guide a frozen video diffusion backbone to generate a coherent video sequence. To model interaction-driven dynamics, we introduce Edge-Aware Graph Reasoning, which conditions message passing on directed edge attributes so that object representations explicitly reflect their relational roles before being mapped to diffusion conditioning. By injecting this representation into a pretrained video DiT model and using LoRA modules in attention layers, the model learns to animate the scene using its own generative priors while respecting the specified interaction structure, enabling a more intuitive and interactive mechanism for motion control. To support training, we construct **GRAPHVID-BENCH**, an interaction-centric dataset of $\sim 27\text{K}$ video clips paired with explicit interaction graphs. Despite using $\sim 24\times$ fewer parameters and substantially less training data, GraphVid achieves competitive performance with state-of-the-art motion-control video generation methods, significantly improving perceptual quality and reconstruction fidelity, and reducing FID by up to 39.9% and FVD by 37.6% compared to Motion-I2V, while improving PSNR ($9.87 \rightarrow 15.98$) and SSIM ($0.38 \rightarrow 0.61$). Our contributions are:

- We propose **GraphVid**, to the best of our knowledge, the first framework that enables interactive control of multi-object dynamics through directed interaction scene graphs constructed from a single image. GraphVid’s graph-to-token conditioning interface injects interaction-aware representations into a frozen video diffusion transformer using parameter-efficient LoRA adaptation, enabling structured control without degrading pretrained generative priors.
- We introduce an **Edge-Aware Graph Reasoning module** that conditions graph message passing on directed edge attributes, allowing model representations to capture complex multi-object relations and interaction-driven motion propagation across multiple entities.
- We curate **GRAPHVID-BENCH**, a high-quality interaction-centric video generation dataset of $\sim 27\text{K}$ video clips paired with structured interaction graphs to support training and evaluation of relational video control.
- Despite requiring only 0.6B trainable parameters and achieving the fastest inference time among competing methods, GraphVid delivers substantial quality gains across perceptual and reconstruction metrics, reducing FID by 34.5% ($25.98 \rightarrow 17.02$) and FVD by 7.8% ($107.89 \rightarrow 99.42$), while increasing PSNR by 5.9% ($15.08 \rightarrow 15.98$) and SSIM by 15.0% ($0.53 \rightarrow 0.61$) compared to WISA (physics-centric text inputs), and outperforming trajectory-based baselines

such as Tora (FID -30.3% , FVD -10.0% , PSNR $+41.7\%$, SSIM $+12.9\%$) and Motion-I2V (FID -39.9% , FVD -37.6% , PSNR $+61.9\%$, SSIM $+60.5\%$).

2 Related Work

Video Generation. Recent advances in generative modeling have significantly improved the quality and scale of video generation systems [4, 48, 56]. Early approaches extended image diffusion models to the temporal domain using U-Net backbones and temporal modules to capture motion dynamics [14, 15]. Subsequent work introduced latent video diffusion and motion-conditioned architectures to improve temporal consistency and scalability [22, 55]. More recently, large-scale diffusion transformers (DiTs) have emerged as the dominant architecture for high-quality video synthesis. Models such as CogVideoX [56] and Wan [48] demonstrate that scaling transformer-based diffusion models enables long-range temporal coherence and improved visual fidelity. Recent works start from such strong pretrained video backbones and add a small number of trainable parameters that inject the desired control signal while preserving the backbone’s generative priors [8, 31, 42, 52]. This trend reflects both compute constraints and the empirical observation that large-scale backbones already encode rich appearance and temporal dynamics [10, 12, 29, 40, 44]. GraphVid follows this paradigm by keeping the video backbone frozen and learning lightweight modules that map structured interaction cues into conditioning tokens for generation.

Controllable Video Generation. While general video generation models can synthesize realistic videos, controlling the generated content remains a major challenge. Early controllable generation approaches introduced conditioning signals such as pose, depth, or motion trajectories to guide video synthesis [10, 29, 42]. Other works explored editing-based pipelines that modify existing videos while preserving temporal consistency [11, 28, 33]. More recent methods focus on integrating structured conditioning signals or intermediate representations to enable more precise control. For example, models can condition on spatial layouts [18, 50], motion fields [29, 42], or object trajectories [53, 57] to guide generation while maintaining coherence with pretrained video backbones. Despite their strong performance, trajectory-based methods rely heavily on dense point-tracking supervision [24], making them data-intensive and difficult to scale to higher-level relational reasoning and complex interactions between objects. In contrast, we propose GraphVid, which represents interactions as a directed scene graph and explicitly models multi-object control using an edge-aware GNN to generate more plausible and physically consistent videos.

Scene Graphs for Visual Generation. Scene graphs provide a structured representation of objects and their relations within a scene and have been widely studied for image understanding and generation [7, 9, 25, 30], 3D scene understanding [1, 36, 47], compositional and commonsense reasoning [20, 26], *etc.*, enabling models to reason about compositional structure and object interactions. Recent graph-to-video formulations support compositional generation from object-centric interaction graphs [3]. In contrast, GraphVid targets *interactive* control

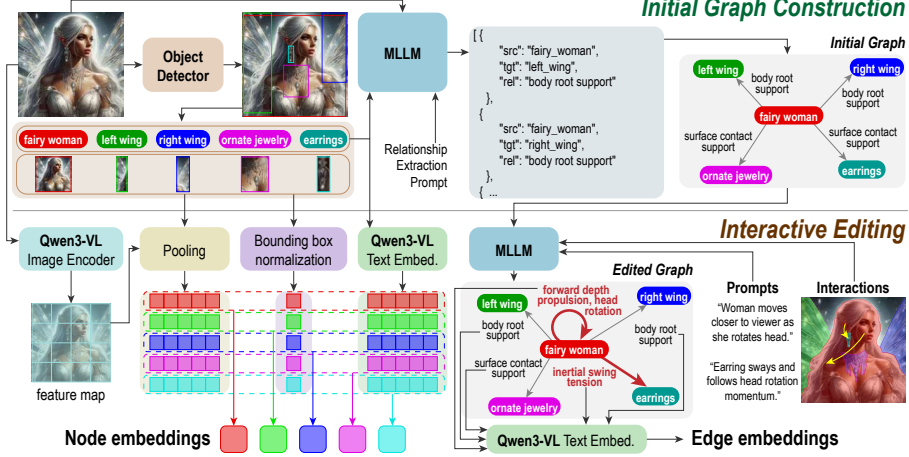


Fig. 2: GraphVid Overview. From an input image, objects are detected and encoded into node embeddings, while an MLLM extracts relational edges to construct an initial interaction graph. Edge-aware graph reasoning converts node and edge representations into conditioning tokens that guide a frozen video diffusion transformer to generate interaction-consistent video dynamics.

with an *interaction-grounded* graph and learns lightweight modules that translate graph cues into conditioning tokens compatible with a video generation backbone, enabling structured control while retaining the backbone’s strong priors.

3 Method

We present **GraphVid**, a framework for controllable image-to-video (I2V) generation from a single conditioning image. The key idea is to represent user intent as a *structured interaction graph* over entities in the scene, and to interface this graph with a large pretrained video diffusion transformer using a lightweight graph-to-token adapter and LoRA-based conditioning. By keeping the backbone frozen and learning only small modules, GraphVid introduces structured, object-centric control while preserving the strong generative priors of the pretrained model. Figure 2 presents an overview.

Given a conditioning image I_0 , our goal is to generate a video sequence $\{I_t\}_{t=1}^T$ that follows a user-specified set of object interactions. To express scene dynamics compositionally, we construct an attributed directed interaction graph $G = (V, E)$, where each node $v_i \in V$ corresponds to an entity detected in I_0 and each directed edge $e_{ij} \in E$ encodes a relation type $r_{ij} \in \mathcal{R}$ from object i to object j (e.g., *push*, *pull*, *lift*, *hold*), describing how entities influence each other, and optional qualitative physical cues (e.g. coarse direction and magnitude), allowing users to describe scene dynamics in intuitive relational terms. Unary directives that involve a single entity are represented as self-loop edges e_{ii} (e.g., *rotate*, *move*, *scale*). We aim to learn a conditional video generator that models $p(\mathbf{I}_{1:T} \mid I_0, G)$ where $\mathbf{I}_{1:T} = \{I_t\}_{t=1}^T$ denotes the T generated frames.




Fig. 3: GraphVid Edge-Aware Graph Reasoning. Node and edge embeddings are encoded into temporal latents and processed by our proposed edge-aware GNN that injects interaction semantics directly into message passing. The resulting graph embeddings capture relational dynamics between entities and are mapped into conditioning tokens that modulate a frozen diffusion video backbone via LoRA adapters.

Node Representations. We extract object entities from the input frame I_0 using a pre-trained vision-language detector. For each detected object v_i , we construct a unified feature vector by fusing: (i) a visual embedding $\mathbf{f}_i^{vis} \in \mathbb{R}^{d_v}$ extracted by encoding the region of I_0 inside the object’s bounding box $\mathbf{b}_i$ and pooling the resulting features, and (ii) a text embedding $\mathbf{f}_i^{txt} \in \mathbb{R}^{d_t}$ encoding the object label. Normalized bounding box coordinates $\mathbf{b}_i \in \mathbb{R}^4$ are concatenated with these embeddings to form a representation $\mathbf{x}_i = [\mathbf{f}_i^{vis} \parallel \mathbf{f}_i^{txt} \parallel \mathbf{b}_i]$ that jointly captures object appearance, semantic identity, and spatial grounding.

Edge Representations. Each directed edge e_{ij} is represented with an open-vocabulary textual descriptor of the intended interaction between the source and target entities. Edges are encoded using a text embedding model to obtain the edge feature $\mathbf{e}_{ij} \in \mathbb{R}^{d_e}$.

3.1 Edge-Aware Graph Reasoning

Standard GCNs treat edges as binary connectivity signals and primarily propagate node features across the graph [17, 27]. However, in physical interaction modeling, the type of relationship between objects directly determines their motion dynamics and implies fundamentally different motion patterns and causal effects. Ignoring these relational attributes reduces the graph to simple spatial adjacency and prevents the model from reasoning about how interactions influence object behavior. To address this, we employ an edge-aware Graph Isomorphism Network (GINEConv [54]) that explicitly incorporates edge semantics during message passing. As illustrated in Figure 3, node and edge features are first projected into a shared hidden space $\mathbf{h}_i^{(0)} = \phi_n(\mathbf{x}_i)$, $\mathbf{a}_{ij} = \phi_e(\mathbf{e}_{ij})$, where

ϕ_n and ϕ_e are learnable MLPs that map node/edge features to a common embedding space $\mathbb{R}^{d_h}$. We then apply L GINEConv layers to propagate relational information across the graph

$$\mathbf{h}_i^{(l+1)} = \text{MLP}^{(l)} \left(\mathbf{h}_i^{(l)} + \sum_{j \in \mathcal{N}(i)} \text{ReLU}(\mathbf{h}_j^{(l)} + \mathbf{a}_{ji}) \right). \quad (1)$$

Unlike standard GCNs, the GINE formulation injects edge attributes directly into the aggregation step, allowing interaction semantics to influence node updates. Consequently, each node representation is conditioned not only on neighboring objects but also on the physical relationships governing their interactions. After L layers, we obtain interaction-aware node embeddings $\mathbf{h}_i \in \mathbb{R}^{d_h}$ that encode both object identity and its dynamic relational state within the scene. These embeddings provide the downstream generative model with a structured representation of which entities should move and how they interact, enabling more physically coherent video generation.

3.2 Graph-to-Video Adapter

Modern Diffusion Transformers (DiTs) trained on large-scale video datasets provide strong spatio-temporal priors for video generation [35, 56]. Training such models from scratch is computationally prohibitive and may compromise the rich generative capabilities learned from large-scale data. Therefore, we leverage a pretrained DiT backbone [16] and keep it frozen, introducing a lightweight adapter that converts graph embeddings into conditioning tokens compatible with the transformer’s input space.

Each interaction-aware node embedding $\mathbf{h}_i$ is therefore projected into the transformer latent space as $\mathbf{z}_i = \text{MLP}(\mathbf{h}_i) \in \mathbb{R}^{d_l}$, where d_l denotes the hidden size of the pretrained DiT. Each $\mathbf{z}_i$ is treated as a semantic conditioning token. Since scenes contain varying numbers of objects, node sequences are padded to a maximum length $N_{\max}$. The resulting tokens, concatenated with edge text tokens, are passed as encoder hidden states. This enables the transformer’s self-attention to reason over entity interactions and translate graph structure into motion during generation.

3.3 Training Objective

Since the pretrained DiT backbone is kept frozen to preserve its strong generative priors, the model requires a lightweight mechanism to adapt to the new graph conditioning modality. We therefore inject Low-Rank Adaptation (LoRA) [21] modules into the $\mathbf{Q}$, $\mathbf{K}$, $\mathbf{V}$, and output projection layers of each transformer block. Modulating these attention projections allows the model to incorporate graph tokens into the attention computation while introducing only a small number of trainable parameters. This enables the transformer to attend to interaction-aware graph embeddings without modifying the pretrained weights.

We train GraphVid using the flow-matching objective employed by the pre-trained diffusion model. Given a target video $\mathbf{x}_0$, we encode it into the latent space and sample a timestep $t \sim p(t)$ from a logit-normal distribution. Noise is then added to obtain the noisy latent $\mathbf{x}_t$. The model is optimized using loss

$$\mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{\mathbf{c} \sim \mathcal{I}_0, G} \mathbb{E}_{t \sim p(t)} \mathbb{E}_{\mathbf{x}_0 \sim p_{\text{data}}(\cdot | \mathbf{c}), \mathbf{x}_1 \sim \mathcal{N}(0, \mathbf{I})} \left[\left\| \mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c}) - (\mathbf{x}_1 - \mathbf{x}_0) \right\|_2^2 \right], \quad (2)$$

where $\mathbf{x}_t = (1 - t)\mathbf{x}_0 + t\mathbf{x}_1$. During training, only the GNN parameters, Graph-to-Adapter MLP layers, and LoRA modules are updated (as θ), while the VAE and DiT backbones remain frozen. This design enables end-to-end learning of a differentiable interactive/symbolic-to-latent mapping from interaction graphs to video dynamics while retaining the pretrained model’s generative capabilities.

4 GRAPHVID-BENCH Dataset

Current video generation benchmarks primarily rely on text prompts [23] or low-level motion cues [8], which provide limited control over object-level interactions. However, many real-world dynamics are fundamentally relational, where actions emerge from how entities interact with one another (*e.g.*, a person turning their head, an object moving relative to its surroundings). To enable structured control over such dynamics, we introduce GRAPHVID-BENCH, a dataset that pairs videos with interaction graphs describing entities and their directed interaction cues.

GRAPHVID-BENCH contains approximately **27K** interaction-centric video clips. To ensure consistent training and evaluation, all clips are standardized to 81 frames at 16 fps. Each video is paired with a directed interaction graph where nodes represent scene entities and edges encode interaction cues between them. Node attributes include visual features, semantic labels, and spatial in-

Table 1: GRAPHVID-BENCH statistics across $\sim 27\text{K}$ curated videos. Interaction complexity highlights the prevalence of multi-entity dynamics.

Metric	Mean	Med.	Std.
Nodes / Graph	7.76	5.0	8.03
Edges / Graph	4.85	4.0	3.46
Interaction Complexity			
No-Interaction ($N=0$)			3881
Single-Interaction ($N=1$)			10982
Multi-Interaction ($N>1$)			12641

formation, while edges capture the type and direction of interactions that influence scene dynamics. As shown in Table 1, the dataset exhibits diverse structural complexity, with an average of 7.76 nodes and 4.85 edges per graph. In total, 10,982 samples contain a single interaction, and 12,641 contain multiple interactions ($N > 1$). As illustrated in Figure 4, GRAPHVID-BENCH covers a broad spectrum of canonical physical interactions. The co-occurrence statistics further reveal that many interactions appear jointly within the same clips, indicating that real-world dynamics frequently involve multiple interacting primitives. We also include 3,881 samples with no explicit interactions ($N = 0$). These scenes contain natural object motion without direct entity-to-entity interactions, providing a baseline for learning general video dynamics. Since **GraphVid** allows




Fig. 4: Interaction statistics in GRAPHVID-BENCH dataset. (a) Frequency of interaction primitives, grouped into four categories (force & manipulation, kinematics & motion, physical contact, and support & position). (b) Interaction co-occurrence matrix captures interaction primitives appearing together within the same video sequences. The co-occurrence of manipulation and motion primitives shows that real-world dynamics are inherently compositional, motivating the use of structured interaction graphs for controllable video generation.

users to add interaction edges at inference time, such samples help the model learn how scene behavior changes when new interactions are introduced. Appendix A provides additional details on the dataset curation process.

5 Experiments

5.1 Implementation Details

GraphVid leverages the pretrained LTX-Video [16] Diffusion Transformer as backbone for video generation. The DiT backbone and the VAE encoder-decoder are kept frozen during training. Object entities are extracted from the conditioning image using Qwen3-VL [2]. For each object, we obtain a $d_v = 8192$ -dimensional visual embedding using mean and max pooling over object crops, a $d_t = 2560$ -dimensional text embedding for the object label, and normalized bounding box coordinates. These are concatenated to form a $d_n = 10756$ -dimensional node feature vector. Physical interactions between objects are encoded using Qwen3-Embedding [59], producing edge attributes of dimension $d_e = 2560$. During edge-aware graph reasoning, node and edge features are projected to a hidden dimension of $d = 512$, producing interaction-aware node embeddings. Node embeddings are projected to the LTX transformer hidden size using an MLP, mapping 512-dimensional graph embeddings to 4096 dimensions. Target videos are encoded into latent space using the pretrained VAE, and noise is injected at timesteps sampled from a logit-normal distribution. We employ LoRA modules of rank 128. During training, only the GNN parameters, graph-to-adaptor projection layers, and LoRA weights are updated.

Table 2: Quantitative comparison on interaction-centric MOVEBENCH subset. GraphVid achieves competitive performance across all metrics while using significantly fewer parameters and training data. **Best** and **second-best** are highlighted.

Method	Train Data	Trainable Params (B)	Inference (s)↓	FID↓	FVD↓	PSNR↑	SSIM↑	EPE↓
Wan-Move [8]	2M	14.5	1800	15.56	82.17	17.21	0.61	2.6
Motion-I2V [42]	10M	1.2	790	28.32	159.32	9.87	0.38	3.9
Tora [60]	630K	5	1200	24.45	110.47	11.27	0.54	3.3
MagicMotion [32]	23K	1.5	750	26.57	105.12	12.04	0.51	3.2
WISA [49]	80K	1	1000	25.98	107.89	15.08	0.53	4.1
FlashMotion [31]	23K	13	677	19.02	104.12	14.04	0.56	3.8
GraphVid (Ours)	27K	0.6	200	17.02	99.42	15.98	0.61	2.9

5.2 Experimental Setup

For evaluation, we establish a MOVEBENCH [8] subset focused on interaction-centric scenarios as the primary evaluation suite. Since GraphVid emphasizes physics-aware relational control rather than explicit pixel-level motion supervision, we prioritize metrics that assess overall video quality, temporal coherence, and structural fidelity. In addition, we evaluate performance on a subset of the distilled MOVEBENCH containing only multi-object interaction scenarios, allowing us to assess the effectiveness of GraphVid in more complex settings where reasoning over interactions between multiple entities is required. We report Fr chet Video Distance (FVD) [46], Fr chet Inception Distance (FID) [19], Peak Signal-to-Noise Ratio (PSNR), and Structural Similarity Index (SSIM) [51] to measure visual quality. We additionally report End-Point Error (EPE) [10], the average displacement between predicted and ground-truth optical flow, which directly measures whether the generated motion follows the intended interactions. Additional details can be found in Appendix B.

5.3 Quantitative Results

Table 2 compares GraphVid with existing controllable video generation methods on the interaction-centric MOVEBENCH benchmark subset, where GraphVid achieves competitive performance across all metrics while requiring substantially fewer training samples, significantly fewer trainable parameters, and faster inference. Among prior approaches, WISA [49] is the most closely related method that explicitly incorporates structured physical knowledge into the video generation process. Compared to WISA, GraphVid achieves a substantial improvement in perceptual quality and reconstruction fidelity. In particular, GraphVid reduces FID from 25.98 to 17.02 and FVD from 107.89 to 99.42, corresponding to a 34.5% and 7.8% improvement, respectively, indicating significantly better visual realism and distribution alignment with ground-truth videos. Similarly, SSIM improves by 15.0%, and PSNR by 5.9%, demonstrating that our graph-based interaction modeling produces frames that are both sharper and structurally more consistent. This advantage also holds for motion accuracy, with GraphVid reducing EPE from 4.1 to 2.9, a 29.3% improvement over WISA, and achieving

Table 3: Multi-object quantitative comparison on interaction-centric MOVEBENCH. GraphVid maintains strong structural coherence compared to larger baselines. **Best** and **second-best** are highlighted.

Method	Train Data	Trainable Params (B)	FID↓	FVD↓	PSNR↑	SSIM↑	EPE↓
Wan-Move [8]	2M	14.5	31.29	252	16.69	0.61	2.2
Tora [60]	630K	5	56.04	369	14.98	0.52	3.5
WISA [49]	80K	1	60.12	341	13.13	0.55	3.9
FlashMotion [31]	23K	13	55.03	311	12.12	0.52	3.9
GraphVid (Ours)	27K	0.6	49.45	291	14.44	0.55	3.0

the lowest EPE among all comparable-scale methods. Since EPE measures how closely the generated optical flow matches the ground truth, this indicates that our graph conditioning produces motion that most faithfully follows the specified interactions. The overall perceptual, reconstruction, and motion gains indicate that explicitly modeling interaction relationships through a directed scene graph provides a stronger inductive bias for controllable video generation than relying on predefined textual (physics-centric) attributes.

We further compare against trajectory-based control methods such as Motion-I2V [42], MagicMotion [32], FlashMotion [31], and Tora [60]. These approaches rely heavily on explicit motion signals such as trajectories, optical flow, or segmentation masks. Compared to Motion-I2V, GraphVid reduces FID by 39.9% (28.32 $\rightarrow$ 17.02) and FVD by 37.6% (159.32 $\rightarrow$ 99.42), while improving PSNR from 9.87 to 15.98 and SSIM from 0.38 to 0.61, indicating substantially better frame reconstruction and perceptual fidelity. Compared to FlashMotion, a recent trajectory-guided video generation method, GraphVid lowers FID by 10.5% (19.02 $\rightarrow$ 17.02) while also achieving stronger reconstruction quality, improving PSNR from 14.04 to 15.98 and SSIM from 0.56 to 0.61. Compared with MagicMotion, GraphVid reduces FID by approximately 36% (26.57 $\rightarrow$ 17.02) and improves temporal consistency. Notably, GraphVid attains a lower EPE (2.9) than all comparable trajectory-based methods (Motion-I2V 3.9, FlashMotion 3.8, Tora 3.3, MagicMotion 3.2), reducing motion error by 9.4–25.6% without relying on any explicit motion supervision. These improvements highlight the effectiveness of modeling scene dynamics through structured interaction reasoning. Wan-Move [8] achieves the best absolute metrics due to its significantly larger scale, being trained on approximately 2M videos and using a 14.5B parameter backbone. In contrast, GraphVid operates with only 0.6B trainable parameters and is trained on a much smaller 27K interaction-centric dataset. Despite using approximately two orders of magnitude less training data and $\sim 24\times$ lower model params, GraphVid achieves competitive performance, outperforming most prior controllable methods on reconstruction and motion-accuracy.

A similar trend can be observed in Table 3, which evaluates performance on the multi-object subset of MOVEBENCH where interaction reasoning becomes particularly important. In these scenarios, GraphVid improves multi-object interaction quality over other lightweight baselines, with clear relative gains, while using substantially fewer resources. When compared to models of similar data

scale, GraphVid achieves the best FID (49.45), improving over the next best (FlashMotion, 55.03) by 10.1%, and the best FVD (291), improving over the next best (FlashMotion, 311) by 6.4%; both indicating better perceptual realism and temporal consistency in multi-entity dynamics. GraphVid also achieves the lowest EPE among comparable-scale methods, confirming that its motion-accuracy advantage carries over to the harder multi-object setting. On reconstruction fidelity, GraphVid remains competitive with Tora (best PSNR) and WISA (best SSIM). Importantly, these gains come with far lower scale. GraphVid uses 0.6B trainable parameters ($\approx 21.7\times$ fewer than FlashMotion and $8.3\times$ fewer than Tora) and trains on 27K clips (comparable to FlashMotion and $23\times$ fewer than Tora), underscoring that interaction-graph conditioning provides strong multi-object coherence without requiring massive models or data. Appendix C presents additional results on our interaction-focused DAVIS [38] subset.

5.4 Efficiency and Scalability Analysis

Figure 5 plots inference throughput against model scale. Beyond accuracy, GraphVid offers strong efficiency advantages, achieving the fastest inference among all compared approaches at 200 seconds, significantly faster than Wan-Move (1800s), Tora (1200s), and WISA (1000s). This efficiency stems from our lightweight interaction graph conditioning and LoRA-based integration into the pre-trained DiT backbone, which avoids the need for heavy auxiliary motion encoders or large-scale motion prediction modules. This confirms that explicitly modeling interaction relationships through a structured scene graph provides a powerful and efficient alternative to trajectory-based motion supervision. By leveraging interaction-aware graph reasoning, GraphVid achieves strong controllability and physical plausibility while operating with substantially lower compute and data.




Fig. 5: Model efficiency comparison. Trade-off between inference throughput and scale. Bubble size reflects model scale.

5.5 Qualitative Analysis

Figure 6 presents representative qualitative results demonstrating that modeling scene dynamics through GraphVid interaction graphs enables precise and interpretable control over complex multi-entity motion. Examples illustrate diverse interaction types, including object-object interactions, human-object interactions, and articulated human motion. Across all scenarios, GraphVid preserves object identity and scene appearance, produces realistic motion trajectories




Fig. 6: Qualitative results of GraphVid for controllable interaction-based video generation. Given a single input image and user-specified interaction context, GraphVid generates temporally coherent video sequences reflecting the desired dynamics. Examples include object-object interactions (geese moving relative to each other), human-object interactions (kicking a ball, holding a cup), and articulated motion (cycling). The generated sequences demonstrate consistent motion, object preservation, and faithful execution of user-specified interaction dynamics.

and interaction-consistent behavior, and maintains stable background geometry. These examples indicate that the interaction-graph interface helps the model bind motion to the correct entities and combine multiple primitives into coherent multi-step dynamics, something that is typically brittle with purely text- or trajectory-only controls. Appendix D presents a human preference study, while Appendix E provides qualitative results across diverse interaction scenarios.

5.6 Ablation Studies

Backbone Generalization. To verify that relational scene graphs serve as a universal conditioning modality rather than an exploit of LTX’s specific latent priors, we replace the 2B LTX backbone with the 5B Wan 2.2 model [48]. When training with this variant, our edge-aware GNN and adapter successfully map the topological graphs into Wan’s latent space, achieving competitive performance (Table 4). This demonstrates that our graph representation is a backbone-agnostic conditioning signal capable of guiding scene dynamics across different foundation models.

Table 4: Backbone ablation.

Backbone	FID↓	FVD↓	PSNR↑	SSIM↑	EPE↓
Ours + LTX	17.02	99.42	15.98	0.61	2.9
Ours + Wan 2.2	17.29	100.01	15.70	0.60	3.1

Graph Sparsity and Attention Dilution. We evaluate the effect of the context window size on the DiT’s self-attention mechanism by varying the maximum node capacity. As shown in Table 5, padding sparse real-world scene graphs to overly large

Table 5: Graph ablation.

#Nodes	FID↓	FVD↓	PSNR↑	SSIM↑
10	17.25	102.04	15.56	0.60
30	17.02	99.42	15.98	0.61
50	17.17	104.97	15.11	0.60
128	17.55	105.45	14.55	0.59
256	17.97	109.92	13.77	0.57

context windows can dilute attention and degrade generation quality. Capping the graph size at 30 nodes consistently outperforms larger capacities (128/256), achieving the best FVD (99.42). This suggests that a spatially tighter, denser node set yields better structural alignment: as the maximum node budget increases (*e.g.*, from 30 $\rightarrow$ 256), FID degrades, likely because the additional capacity is dominated by zero-padding rather than informative structure. The resulting padded tokens inject noise into attention and dilute meaningful relational cues, making it harder for the transformer to concentrate on active interactions. Overall, these results suggest that balanced, densely populated graph representations provide the strongest conditioning signal, and GraphVid therefore uses a maximum node capacity of 30 for all experiments.

Conditioning Modalities and Edge Semantics. To evaluate the importance of structured conditioning and edge semantics, we perform an inference-time ablation comparing three distinct regimes. Table 6 shows (1) the standard LTX baseline, which generates video using a global text prompt and an initial image; (2) GraphVid conditioned on the image and

Table 6: Conditioning modalities.

Variant	FID↓	FVD↓	PSNR↑	SSIM↑	EPE↓
Text + Img (LTX)	20.04	109.34	11.23	0.53	3.6
Graph (no edge text)	17.52	101.32	16.88	0.60	3.0
Graph (full)	17.02	99.42	15.98	0.61	2.9

a topology-only graph, where the user provides interaction arrows and overlays but omits textual edge context; and (3) our full GraphVid approach, utilizing the image and a complete graph equipped with explicit text context for each interaction edge. While the text-based baseline achieves moderate visual quality, it struggles to bind localized motion to specific entities (FID 20.04). Replacing the global text with our topology-only graph improves FID to 17.52, indicating that explicitly graphing spatial boundaries aids motion localization. However, without the semantic richness of the textual interaction embeddings, the model exhibits severe motion ambiguity, failing to distinguish opposing interactions like “pushing” and “pulling”. The full graph conditioning outperforms both baselines, reducing FID to 17.02 and FVD to 99.42. These results suggest that graph topology alone provides useful spatial grounding, while semantic edge labels supply additional interaction intent that helps disambiguate object relationships and improve motion generation.

LoRA Rank and Capacity Trade-offs. We inject low-rank adaptation (LoRA) matrices into the transformer blocks to modulate the pre-trained weights with our graph conditions. In this ab-

Table 7: Ablation on LoRA rank.

Rank	FID↓	FVD↓	PSNR↑	SSIM↑
16	18.34	103.18	13.70	0.58
32	17.89	102.39	14.51	0.60
64	17.53	100.99	15.25	0.60
128	17.02	99.42	15.98	0.61

lation, we study the effect of the LoRA rank on generation quality and reconstruction fidelity. As shown in Table 7, increasing the rank consistently improves performance across all metrics. Moving from rank 16 to 128 reduces FID from 18.34 to 17.02 (−7.2%) and FVD from 103.18 to 99.42 (−3.6%), indicating better perceptual quality and temporal coherence. At the same time, recon-

struction metrics improve substantially, with PSNR increasing from 13.70 to 15.98 (+16.6%) and SSIM from 0.58 to 0.61 (+5.2%). These results suggest that higher-rank adapters provide greater representational capacity to capture complex motion and appearance variations, leading to more faithful and temporally consistent video generation. Rank 128 is used in all the reported experiments.

6 Conclusion

We introduce **GraphVid**, a framework for controllable image-to-video generation that represents scene dynamics through directed interaction graphs. Unlike prior approaches that rely on dense motion signals or low-level physics annotations, GraphVid models motion as structured interactions between entities, enabling intuitive multi-object control directly from a single image. By integrating Edge-Aware Graph Reasoning with a frozen video diffusion transformer through parameter-efficient LoRA conditioning, GraphVid preserves strong pre-trained generative priors while enabling structured control over scene dynamics. To support this paradigm, we also curate **GRAPHVID-BENCH**, a dataset of 27K interaction-centric videos paired with explicit interaction graphs, providing supervision for learning relational video dynamics. Experiments show that GraphVid matches or outperforms prior controllable video generation methods while using far fewer parameters and training data, highlighting interaction graphs as a scalable interface for controllable video generation.

Acknowledgments

This research was partially supported by Google, the Google TPU Research Cloud (TRC) program, the NSF CAREER Award #2542328, the U.S. Defense Advanced Research Projects Agency (DARPA) under award HR001125C0303, and the U.S. Army under contract W5170125CA160. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of Google, NSF, DARPA, the U.S. Army, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein. This work also used the Delta GPUs provided by the National Center for Supercomputing Applications through allocations CIS240808 and CIS250318 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS [6]) program, supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296.

References

1. Armeni, I., He, Z.Y., Gwak, J., Zamir, A.R., Fischer, M., Malik, J., Savarese, S.: 3d scene graph: A structure for unified semantics, 3d space, and camera. In: IEEE Conf. Comput. Vis. Pattern Recog. (2019)

2. Bai, S., Cai, Y., Chen, R., Chen, K., Chen, X., Cheng, Z., Deng, L., Ding, W., Gao, C., Ge, C., et al.: Qwen3-vl technical report. arXiv preprint arXiv:2511.21631 (2025)
3. Bar, A., Herzig, R., Wang, X., Rohrbach, A., Chechik, G., Darrell, T., Globerson, A.: Compositional video synthesis with action graphs. In: Int. Conf. Mach. Learn. (2021)
4. Bar-Tal, O., Chefer, H., Tov, O., Herrmann, C., Paiss, R., Zada, S., Ephrat, A., Hur, J., Liu, G., Raj, A., et al.: Lumiere: A space-time diffusion model for video generation. In: SIGGRAPH Asia Conference Papers (2024)
5. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023)
6. Boerner, T.J., Deems, S., Furlani, T.R., Knuth, S.L., Towns, J.: ACCESS: Advancing Innovation: NSF's Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support. In: Practice and Experience in Advanced Research Computing (PEARC) (2023)
7. Chang, X., Ren, P., Xu, P., Li, Z., Chen, X., Hauptmann, A.: A comprehensive survey of scene graphs: Generation and application. IEEE Transactions on Pattern Analysis and Machine Intelligence **45**(1), 1–26 (2021)
8. Chu, R., He, Y., Chen, Z., Zhang, S., Xu, X., Xia, B., WANG, D., Yi, H., Liu, X., Zhao, H., et al.: Wan-move: Motion-controllable video generation via latent trajectory guidance. In: Adv. Neural Inform. Process. Syst. (2025)
9. Dutta, A., Mehrab, K.S., Sawhney, M., Neog, A., Khurana, M., Fatemi, S., Pradhan, A., Maruf, M., Lourentzou, I., Daw, A., et al.: Open world scene graph generation using vision language models. arXiv preprint arXiv:2506.08189 (2025)
10. Geng, D., Herrmann, C., Hur, J., Cole, F., Zhang, S., Pfaff, T., Lopez-Guevara, T., Aytar, Y., Rubinstein, M., Sun, C., et al.: Motion prompting: Controlling video generation with motion trajectories. In: IEEE Conf. Comput. Vis. Pattern Recog. (2025)
11. Geyer, M., Bar-Tal, O., Bagon, S., Dekel, T.: Tokenflow: Consistent diffusion features for consistent video editing. arXiv preprint arXiv:2307.10373 (2023)
12. Gillman, N., Herrmann, C., Freeman, M., Aggarwal, D., Luo, E., Sun, D., Sun, C.: Force prompting: Video generation models can learn and generalize physics-based control signals. In: Adv. Neural Inform. Process. Syst. (2025)
13. Goyal, R., Ebrahimi Kahou, S., Michalski, V., Materzynska, J., Westphal, S., Kim, H., Haenel, V., Fruend, I., Yianilos, P., Mueller-Freitag, M., et al.: The "something something" video database for learning and evaluating visual common sense. In: IEEE Conf. Comput. Vis. Pattern Recog. (2017)
14. Guo, Y., Yang, C., Rao, A., Agrawala, M., Lin, D., Dai, B.: Sparsectrl: Adding sparse controls to text-to-video diffusion models. In: Eur. Conf. Comput. Vis. (2024)
15. Guo, Y., Yang, C., Rao, A., Liang, Z., Wang, Y., Qiao, Y., Agrawala, M., Lin, D., Dai, B.: Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. Int. Conf. Learn. Represent. (2024)
16. HaCohen, Y., Chiprut, N., Brazowski, B., Shalem, D., Moshe, D., Richardson, E., Levin, E., Shiran, G., Zabari, N., Gordon, O., et al.: Ltx-video: Realtime video latent diffusion. arXiv preprint arXiv:2501.00103 (2024)
17. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. Adv. Neural Inform. Process. Syst. (2017)

18. He, Y., Liu, Z., Chen, J., Tian, Z., Liu, H., Chi, X., Liu, R., Yuan, R., Xing, Y., Wang, W., et al.: Llms meet multimodal generation and editing: A survey. arXiv preprint arXiv:2405.19334 (2024)
19. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Adv. Neural Inform. Process. Syst.* (2017)
20. Holla, M., Lourentzou, I.: Commonsense for zero-shot natural language video localization. In: *AAAI* (2024)
21. 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. *Int. Conf. Learn. Represent.* (2022)
22. Hu, Y., Chen, Z., Luo, C.: Lamd: Latent motion diffusion for image-conditional video generation. *Int. J. Comput. Vis.* (2025)
23. 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: *IEEE Conf. Comput. Vis. Pattern Recog.* (2024)
24. Karaev, N., Rocco, I., Graham, B., Neverova, N., Vedaldi, A., Rupprecht, C.: Cotracker: It is better to track together. In: *Eur. Conf. Comput. Vis.* (2024)
25. Karwande, G., Mbakwe, A.B., Wu, J.T., Celi, L.A., Moradi, M., Lourentzou, I.: Chexrelnet: An anatomy-aware model for tracking longitudinal relationships between chest x-rays. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention* (2022)
26. Khan, M.J., Ilievski, F., Breslin, J.G., Curry, E.: A survey of neurosymbolic visual reasoning with scene graphs and common sense knowledge. *Neurosymbolic Artificial Intelligence* (2025)
27. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016)
28. 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)
29. Lei, G., Wang, C., Zhang, R., Wang, Y., Li, H., Xu, W.: Animateanything: Consistent and controllable animation for video generation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2025)
30. Li, H., Zhu, G., Zhang, L., Jiang, Y., Dang, Y., Hou, H., Shen, P., Zhao, X., Shah, S.A.A., Bennamoun, M.: Scene graph generation: A comprehensive survey. *Neurocomputing* (2024)
31. Li, Q., Xing, Z., Wang, R., Cao, H., Dai, Q., Dong, D., Wu, Z.: Flashmotion: Few-step controllable video generation with trajectory guidance. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2026)
32. Li, Q., Xing, Z., Wang, R., Zhang, H., Dai, Q., Wu, Z.: Magicmotion: Controllable video generation with dense-to-sparse trajectory guidance. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2025)
33. Li, X., Ma, C., Yang, X., Yang, M.H.: Vidtoe: Video token merging for zero-shot video editing. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2024)
34. Li, Y., Wang, X., Zhang, Z., Wang, Z., Yuan, Z., Xie, L., Shan, Y., Zou, Y.: Image conductor: Precision control for interactive video synthesis. In: *AAAI* (2025)
35. Lin, B., Ge, Y., Cheng, X., Li, Z., Zhu, B., Wang, S., He, X., Ye, Y., Yuan, S., Chen, L., et al.: Open-sora plan: Open-source large video generation model. arXiv preprint arXiv:2412.00131 (2024)
36. Ogunleye, M.A., Abdelrahman, E., Lourentzou, I.: 3d-vcd: Hallucination mitigation in 3d-llm embodied agents through visual contrastive decoding. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2026)

37. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: IEEE Conf. Comput. Vis. Pattern Recog. (2023)
38. Pont-Tuset, J., Perazzi, F., Caelles, S., Arbeláez, P., Sorkine-Hornung, A., Van Gool, L.: The 2017 davis challenge on video object segmentation. arXiv preprint arXiv:1704.00675 (2017)
39. Romero, D., Bermudez, A., Li, H., Pizzati, F., Laptev, I.: Learning to generate rigid body interactions with video diffusion models. arXiv preprint arXiv:2510.02284 (2025)
40. Shen, Y., Liu, J., Li, X., Liu, Y., Li, B., Yang, H., Jia, W., Li, Y., Yu, T., Rehg, J.M., Cao, X., Lourentzou, I.: Egoforge: Goal-directed egocentric world simulator. arXiv preprint arXiv:2603.20169 (2026)
41. Shen, Y., Xiong, J., Yu, T., Lourentzou, I.: Phantom: Physics-infused video generation via joint modeling of visual and latent physical dynamics. In: IEEE Conf. Comput. Vis. Pattern Recog. (2026)
42. Shi, X., Huang, Z., Wang, F.Y., Bian, W., Li, D., Zhang, Y., Zhang, M., Cheung, K.C., See, S., Qin, H., et al.: Motion-i2v: Consistent and controllable image-to-video generation with explicit motion modeling. In: SIGGRAPH Conference Papers (2024)
43. Singer, U., Polyak, A., Hayes, T., Yin, X., An, J., Zhang, S., Hu, Q., Yang, H., Ashual, O., Gafni, O., et al.: Make-a-video: Text-to-video generation without text-video data. arXiv preprint arXiv:2209.14792 (2022)
44. Susladkar, O., Prakash, T., Juvekar, A., Nguyen, K.A., Jang, D.H., Dhillon, I.S., Lourentzou, I.: Pyratok: Language-aligned pyramidal tokenizer for video understanding and generation. In: IEEE Conf. Comput. Vis. Pattern Recog. (2026)
45. Susladkar, O., Sen Gupta, J., Sehgal, C., Mittal, S., Singhal, R.: Motionaura: Generating high-quality and motion consistent videos using discrete diffusion. In: Int. Conf. Learn. Represent. (2025)
46. Unterthiner, T., Van Steenkiste, S., Kurach, K., Marinier, R., Michalski, M., Gelly, S.: Fvd: A new metric for video generation. Int. Conf. Learn. Represent. Worksh. DeepGenStruct (2019)
47. Wald, J., Dhano, H., Navab, N., Tombari, F.: Learning 3d semantic scene graphs from 3d indoor reconstructions. In: IEEE Conf. Comput. Vis. Pattern Recog. (2020)
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, J., Ma, A., Cao, K., Zheng, J., Feng, J., Zhang, Z., Pang, W., Liang, X.: Wisa: World simulator assistant for physics-aware text-to-video generation. In: Adv. Neural Inform. Process. Syst. (2025)
50. Wang, X., Yuan, H., Zhang, S., Chen, D., Wang, J., Zhang, Y., Shen, Y., Zhao, D., Zhou, J.: Videocomposer: Compositional video synthesis with motion controllability. Adv. Neural Inform. Process. Syst. (2023)
51. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE Int. Conf. Image Process. (2004)
52. Wei, Y., Zhang, S., Yuan, H., Gong, B., Tang, L., Wang, X., Qiu, H., Li, H., Tan, S., Zhang, Y., et al.: Dreamrelation: Relation-centric video customization. In: IEEE Conf. Comput. Vis. Pattern Recog. (2025)
53. Wu, W., Li, Z., Gu, Y., Zhao, R., He, Y., Zhang, D.J., Shou, M.Z., Li, Y., Gao, T., Zhang, D.: Draganything: Motion control for anything using entity representation. In: Eur. Conf. Comput. Vis. (2024)

54. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? arXiv preprint arXiv:1810.00826 (2018)
55. Yang, X., He, C., Ma, J., Zhang, L.: Motion-guided latent diffusion for temporally consistent real-world video super-resolution. In: Eur. Conf. Comput. Vis. (2024)
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. In: Int. Conf. Learn. Represent. (2025)
57. Yin, S., Wu, C., Liang, J., Shi, J., Li, H., Ming, G., Duan, N.: Dragnuwa: Fine-grained control in video generation by integrating text, image, and trajectory. arXiv preprint arXiv:2308.08089 (2023)
58. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: IEEE Conf. Comput. Vis. Pattern Recog. (2023)
59. Zhang, Y., Li, M., Long, D., Zhang, X., Lin, H., Yang, B., Xie, P., Yang, A., Liu, D., Lin, J., et al.: Qwen3 embedding: Advancing text embedding and reranking through foundation models. arXiv preprint arXiv:2506.05176 (2025)
60. 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: IEEE Conf. Comput. Vis. Pattern Recog. (2025)

A GRAPHVID-BENCH Details and Curation Pipeline

Dataset Provenance and Aggregation. We construct our GRAPHVID-BENCH training data from three complementary open-source datasets to ensure a broad distribution of actions and interaction types. Specifically, we use WISA-80K [49] for rigid-body physics-driven dynamics, SOMETHING-SOMETHING v2 [13] for diverse commonsense human-object and object-object interactions, and Magic-Data [32] for complex real-world foreground motion and multi-interaction scenarios. For evaluation, we curate dedicated test splits distilled from DAVIS 2017 [38] and MOVEBENCH [8] to assess generalization on established benchmarks.

Automated Semantic Filtering. To remove confounding motion factors, we introduce an automated semantic filtering stage using a vision-language model (Qwen3-VL [2]). For each video, lightweight proxy clips are extracted and provided to the vision-language model to classify the dominant motion source into categories such as camera motion, ego-motion, ambient environmental dynamics, or discrete object interactions. We retain only videos classified as exhibiting discrete object interactions and discard sequences dominated by camera panning, background motion, or fluid/gaseous environmental effects. This filtering step ensures that the final dataset emphasizes interaction-driven scene dynamics rather than camera-induced motion.

Spatio-Temporal Standardization and Motion-Energy Windowing. To align the dataset with the latent resolution constraints of the video diffusion backbone and to maintain consistent training inputs, all videos are standardized spatially and temporally. Each video is resized to a fixed resolution of 512×288 . Instead of center-cropping, which often truncates key interactions near frame boundaries, we apply dynamic letterboxing to preserve the original aspect ratio and avoid geometric distortion of interacting entities. The temporal dimension is standardized to 81 frames at 16 FPS. Since many source videos exceed this duration, naive random cropping may capture segments where the interaction is inactive. To address this, we introduce a sliding motion-energy window algorithm. The preprocessing pipeline computes a motion-energy signal by summing the pixel-wise absolute differences between consecutive grayscale frames. An 81-frame window is then slid across this signal, and the segment with the highest aggregated motion energy is selected. This ensures that the extracted clip corresponds to the most active phase of the interaction, providing stronger supervisory signals for training. During training, a random subset of each batch is upsampled to different backbone-compatible resolutions, encouraging the model to generate high-quality videos across multiple output resolutions.

A.1 Test Set Distillation

To ensure fair and consistent evaluation, the test datasets undergo the same spatio-temporal standardization pipeline described above.

DAVIS 2017: Raw frame sequences are first converted into video tensors. Sequences slightly shorter than 81 frames (*e.g.*, 78–79 frames) are edge-padded using the final frame to satisfy the temporal requirement. Longer sequences are




Fig. 7: Category distribution of the MOVEBENCH dataset and the curated evaluation subset. The figure shows the number of samples per semantic category in the original MOVEBENCH dataset (light bars) and the subset selected for evaluation in our experiments (dark bars). Categories are ordered by decreasing subset frequency. Our curated evaluation split preserves broad semantic coverage while reducing extreme long-tail categories with few samples. Importantly, the subset retains 53 out of 54 unique categories present in MOVEBENCH (98.15% coverage), ensuring that evaluation remains representative of the full dataset while focusing on interaction-rich scenarios.

processed using the motion-energy windowing procedure to extract the most dynamic segment. The resulting clips are then filtered using the VLM motion classifier to remove sequences dominated by camera or ego-motion.

MOVEBENCH: Videos are first filtered by the VLM to retain only clips exhibiting discrete object-level physical interactions. The retained videos are then resized to 512×288 and standardized to 81 frames to match the visual input requirements of the downstream model. Figure 7 shows the category distribution of both the original MOVEBENCH dataset and our distilled subset.

B Additional Experimental Details

We train our Edge-Aware GNN [54] with a hidden dimension of 512. The graph representation is projected into the 4096-dimensional latent space of the LTX 2B backbone [16] through a 3-layer MLP adapter, while the backbone itself remains frozen throughout training. To enable parameter-efficient adaptation, we insert LoRA modules into the attention projections, namely `to_q`, `to_k`, `to_v`, and `to_out.0`, with rank $r = 128$ and scaling factor $\alpha = 32$. Training is performed on 8 NVIDIA A100 GPUs for approximately 2.5 days using `bf16` mixed precision and the Accelerate library for memory-efficient distributed optimization. We use AdamW with a learning rate of 1×10^{-4} . Unless otherwise stated, inference is conducted at a spatial resolution of 512×288 , with 81 frames at 16 FPS.

Table 8: Quantitative comparison on the DAVIS motion benchmark. We evaluate controllable video generation quality across perceptual (FID, FVD) and reconstruction (PSNR, SSIM) metrics. GraphVid achieves strong overall performance while using significantly fewer trainable parameters. Highlighted **best** and **second-best** among the parameter-efficient architectures.

Method	Trainable Params (B)	FID↓	FVD↓	PSNR↑	SSIM↑	EPE↓
Wan-Move [8]	14.5	15.04	86	12.37	0.57	2.5
Tora [60]	5.0	24.56	107	10.21	0.49	3.5
Motion-I2V [42]	1.2	26.12	114	9.98	0.43	3.8
WISA [49]	1.0	25.05	110	10.45	0.47	3.3
MagicMotion [32]	1.5	20.11	100	11.32	0.49	3.5
FlashMotion [31]	13	17.09	107	9.92	0.49	3.4
GraphVid (Ours)	0.6	16.20	98	12.85	0.55	2.9

For sampling, we adopt the `FlowMatchEulerDiscreteScheduler` with continuous time-step sampling. For fair comparison with prior methods, FlashAttention is disabled during inference.

Environmental Impact. We estimate the training carbon footprint using a uniform hardware-level accounting protocol for both GraphVid and Wan-Move. Assuming the maximum thermal design power of an NVIDIA A100 GPU (400 W), our training setup with 8 A100 GPUs for approximately 2.5 days (60 hours) corresponds to a GPU-only energy consumption of $8 \times 0.4 \times 60 = 192$ kWh. Under the same assumption, Wan-Move, which uses 64 A100 GPUs for 10 days (240 hours), corresponds to $64 \times 0.4 \times 240 = 6144$ kWh.

C Additional Experiments

In this section, we provide extended empirical evaluations to further validate GraphVid’s efficiency, generalization capability, and architectural robustness.

Evaluation on the filtered DAVIS benchmark. Table 8 evaluates GraphVid on our interaction-centric DAVIS subset, where GraphVid achieves highly competitive performance while utilizing a fraction of the trainable parameters (0.6B) and training data (27K videos). Compared to WISA [49], GraphVid substantially improves both temporal coherence (FVD: 110 $\rightarrow$ 98) and perceptual quality (FID: 25.05 $\rightarrow$ 16.20). This confirms that explicitly modeling interaction relationships through a directed scene graph provides a much stronger inductive bias than relying on predefined textual attributes. Furthermore, GraphVid consistently achieves better structural fidelity (PSNR: 12.85, SSIM: 0.55), perceptual realism, and temporal consistency (FVD: 98) than trajectory-conditioned methods (*i.e.*, Motion-I2V [42], Tora [60], MagicMotion [32], and FlashMotion [31]). GraphVid excels in overall frame reconstruction, highlighting the advantage of structured interaction reasoning over geometric motion cues alone. Finally, while the heavyweight Wan-Move (14.5B parameters, 2M training videos) establishes the best performance, GraphVid achieves state-of-the-art perceptual

Table 9: Training-time graph representation ablation. We compare graph variants that use only node features against our GraphVid full model, which additionally incorporates semantic edge embeddings through message passing.

Variant	Nodes	Edges	FID↓	FVD↓	PSNR↑
Node-only Graph Encoder	✓	✗	19.55	103.00	15.24
GraphVid (Ours)	✓	✓	17.02	99.42	15.98

and structural fidelity among all parameter-efficient baselines despite an order-of-magnitude reduction in model capacity and training scale.

Edge-Aware GNN Training Robustness. We conduct an architectural ablation to isolate the contribution of relational edge messaging within our GNN. We train a GraphVid baseline variant utilizing only isolated node features (extracted bounding boxes and visual object embeddings) and omitting the semantic edge connections and message-passing layers. We compare this against our full Edge-Aware GNN implementation, which inherently integrates the relational dynamics through its explicit edge-routing architecture. As detailed in Table 9, while the node-only variant converges and provides a baseline level of control (FVD 103), it fails to fully resolve the complex physical dependencies between interacting entities. The integration of explicit edge features during training substantially improves generation quality across all metrics, driving FVD down to 99.42 and improving structural reconstruction (PSNR 15.98). This confirms that learning the physical causality of a scene fundamentally requires modeling the relationships between objects.

D Human Preference Study

We conduct a human evaluation study to assess how well the generated videos satisfy the intended interaction semantics described in the prompt. The study includes a diverse group of 60 participants evaluating 10 distinct interaction scenarios. Participants are presented with the text prompt describing the intended interaction, the first-frame conditioning image, and a short motion guidance video illustrating the desired trajectory. They are then shown four anonymized candidate videos and asked two separate questions to select the result that (a) best satisfies the semantic intent of the prompt and (b) has the best visual quality. For each scenario, participants provide preference judgments, yielding $60 \times 10 \times 2 = 1200$ individual judgments.

In Figure 8, we observe that GraphVid achieves a high Win Rate of 60% and 90% for semantic intent and visual quality, respectively. Figure 8 also summarizes the aggregated preference rates for models that were preferred for both metrics. Across the 10 diverse evaluated scenarios, **GraphVid** wins the majority vote for *both* semantic intent and visual quality in 6 out of 10 test cases, with 3 ties and only 1 loss. Moreover, GraphVid receives the highest semantic preference, capturing 39.5% of all votes. This substantially outperforms Wan-Move (22.3%), Tora (19.3%), and WISA (18.9%), representing a relative improvement of over




Fig. 8: User preference study across different dimensions. GraphVid achieves the highest Selection Rate (Wins) compared to all baselines. Here, “Wins” denotes the percentage of times a model was chosen as the absolute best for a given metric, “Ties” indicates scenarios where multiple models were deemed equally preferred, and “Losses” indicates the model was not selected.

77% compared to the strongest baseline. The strong preference for GraphVid confirms that users perceive the generated videos as better capturing the core intent of the prompt, bridging the gap between abstract control signals and visually realistic video generation. These consistent results indicate that GraphVid’s semantic interaction representation enables a more faithful realization of the prompt across diverse contexts without sacrificing visual fidelity.

E Qualitative Results

Figure 9 presents qualitative comparisons of multi-object interaction control across several baseline methods and GraphVid. Given an input image and user-specified interaction cues, prior approaches often struggle to faithfully follow the intended interaction dynamics. In the ship scenario (top), baseline methods either fail to separate the motion trajectories of the two ships or introduce visual artifacts during generation. In contrast, GraphVid produces coherent diverging trajectories consistent with the user-specified interaction. In the jogging scenario (bottom), competing methods exhibit unstable motion patterns, including subjects disappearing from the frame, incorrect action generation (walking instead of jogging), or noticeable motion blur. GraphVid maintains stable subject identities and produces temporally consistent motion while correctly preserving the relative dynamics between the lead and follow joggers. These examples highlight the advantage of structured interaction graph conditioning for modeling coordinated multi-object dynamics.

Figure 10 presents additional qualitative comparisons between GraphVid and prior controllable video generation methods. Given the same input image and user-specified interaction cues, baseline approaches often struggle to maintain coherent motion dynamics across time. In the handshake scenario, competing models exhibit unstable interaction behavior, including abrupt termination of




Fig. 9: Qualitative comparison on multi-object interaction control. Given an input image and user-specified interaction cues (left), different methods generate video sequences over time (center). Highlighted regions indicate common failure modes in prior approaches. In the ship scenario (top), baseline methods struggle to follow the specified motion directions or introduce artifacts, whereas GraphVid produces coherent diverging trajectories consistent with the interaction cues. In the jogging scenario (bottom), competing methods show unstable motion (*e.g.*, subject disappearance, walking instead of jogging, or blur), while GraphVid maintains stable identities and consistent multi-entity dynamics.

the motion trajectory or blurred hand dynamics, indicating difficulty in sustaining consistent multi-entity interactions. In contrast, GraphVid produces a stable handshake sequence with smooth temporal evolution and consistent contact between the two subjects. In the metro scenario, baseline methods either terminate the motion prematurely or introduce structural hallucinations such as incorrect wheel formations, reflecting weak grounding between motion generation and object structure. GraphVid preserves object geometry while generating a steady forward motion along the intended track. These examples further demonstrate that interaction graph conditioning improves temporal consistency and structural reliability in multi-object video generation.

F Failure Case Analysis

While GraphVid significantly improves controllable video generation, certain limitations remain in semantic granularity and object grounding. A common failure mode occurs when the VLM fails to isolate small or symbolic elements during scene graph construction. As shown in Figure 11, the user attempts to apply

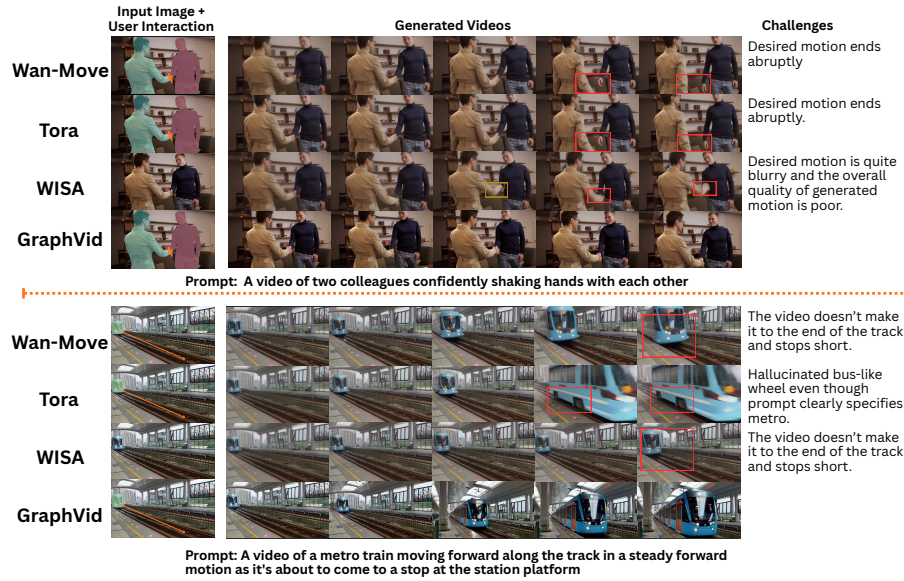


Fig. 10: Qualitative comparison on interaction-driven video generation.

Given an input image and user-specified interaction cues (left), different models generate video sequences over time (center). Highlighted regions mark common failure modes of prior methods. In the handshake scenario (top), baseline models exhibit abrupt motion termination or blurry hand dynamics, whereas GraphVid produces a consistent handshake with smooth temporal continuity. In the metro scenario (bottom), competing methods stop prematurely or hallucinate structural artifacts (*e.g.*, bus-like wheels), while GraphVid preserves object structure and generates stable forward motion along the intended trajectory.

a rotational transformation to the small “Zzz” symbol above the sleeping bear. However, the VLM does not segment this low-salience text as an independent node. As a result, the interaction vector is incorrectly associated with the nearest high-confidence object—the bear itself. Consequently, the generation model applies a rigid-body rotation to the entire subject rather than the intended text element, producing severe geometric distortion across frames. This example highlights a limitation of current zero-shot detection pipelines for fine-grained interaction control and suggests that future systems may benefit from explicit node specification or interactive segmentation to ensure correct grounding of small semantic elements.

G Broader Impacts

Controllable video generation has the potential to significantly expand the accessibility of visual content creation. By enabling users to specify scene dynamics through intuitive interaction graphs rather than complex motion annotations, GraphVid may lower the technical barrier for generating high-quality anima-




Fig. 11: Failure case: semantic grounding error for small symbolic objects. The user interaction specifies a rotational transformation applied to the “Zzz” symbol, indicating sleep (left). However, the detector fails to isolate the text as an independent node during scene graph construction. Consequently, the motion transformation is incorrectly applied to the dominant object (the bear), producing severe geometric distortion and unrealistic deformation across frames. This example highlights a limitation of current vision-language detection when grounding small semantic elements that lack strong visual salience.

tions and simulations. This capability could benefit applications such as creative media production, virtual prototyping, educational visualization of physical phenomena, and simulation environments for robotics and embodied AI research. In particular, representing dynamics through structured interactions may also support improved interpretability in generative systems by making motion control more explicit and modular.

At the same time, advances in controllable video synthesis also raise concerns regarding misuse. The ability to generate realistic videos from minimal input could potentially contribute to the creation of deceptive or manipulated media, including misinformation or impersonation. While GraphVid focuses on structured interaction control rather than photorealistic identity synthesis, similar generative methods may still be misused if deployed irresponsibly. We encourage future work to explore safeguards such as watermarking, provenance tracking, and detection methods to mitigate harmful applications.



Fig. 12: Additional qualitative results on diverse interaction scenarios. Each row shows a different example from our evaluation set. The first column provides the input image, while the remaining columns show frames sampled from the generated video. GraphVid produces temporally coherent motion across a wide range of activities, including object manipulation, human actions, and articulated motion such as cooking, sculpting, dancing, playing instruments, and animal movement. These examples demonstrate that structured interaction conditioning enables consistent object dynamics while preserving scene appearance and identity across frames.

Prompt for Motion Source Classification

You are an Expert in Kinematics and Video Motion Analysis.

You are given:

1. A short video clip
2. A caption describing the video: "{caption}"

Your task is to isolate true, discrete object-to-object physics from camera-induced motion and ambient environmental motion.

Motion Categories

- **discrete_object_physics**: Solid, distinct objects (*e.g.*, people, cars, bats, balls) moving due to explicit physical forces, collisions, or locomotion. The motion can be easily tracked with a bounding box.
- **ambient_environmental**: The primary motion is fluid, gaseous, or driven by the environment. Examples: fire burning, water flowing, waves crashing, smoke billowing, or static trees swaying in the wind.
- **camera_motion**: The camera is panning, tilting, zooming, or shaking. Key indicator: the static background elements (walls, ground) shift uniformly across the screen.
- **ego_motion**: First-person perspective moving through space (*e.g.*, POV walking, driving).
- **static**: Little to no meaningful motion.

Strict Evaluation Rules

1. **The Ambient Rule (CRITICAL)**: If the dominant motion in the video is fire, water, smoke, or wind, you **MUST** classify it as **ambient_environmental**. Scene graphs cannot easily model fluids.
2. **Background First**: Always look at the static background elements first. If the background is shifting uniformly, classify as **camera_motion**.
3. **The Tracking Rule**: If a camera pans heavily to track a moving subject (blurring the background), classify as **camera_motion** or **ego_motion**, because pixel-level motion is dominated by the camera.

Return **STRICT JSON only**. You must follow this exact key order to ensure proper reasoning:

```
{
  "background_and_ambient_analysis": "Briefly check for shifting
backgrounds or fluid/environmental motion (fire, water, wind).",
  "object_analysis": "Briefly describe what discrete,
solid objects are moving.",
  "primary_motion_source": "discrete_object_physics |
ambient_environmental | camera_motion | ego_motion | static",
  "confidence": 0.0-1.0
}
```

Prompt for Initial Scene Graph Generation

ROLE

You are a Physics-Informed Vision-Language Expert. Your goal is to analyze a static image and its corresponding Object JSON to construct a *Physics Scene Graph*.

TASK

Generate a directed graph of physical interactions between the provided objects. Your output must be a strict JSON object containing a list of edges.

Interaction Rules (The Physics Engine)

1. **Visible Evidence Only:** Do not hallucinate interactions.
2. **Implicit Proximity:** If objects are not touching but are spatially related (*e.g.*, ball near bat), infer a `spatial_proximity` or `trajectory_threat` edge. **Do not return an empty list.**
3. **Newtonian Terminology:** Use physics-grounded terms such as Friction, Tension, Normal Force, Gravity, Torque, or Applied Force.
4. **Full Sentence Relations:** The `relation_description` must be grammatically complete:
[Spatial Preposition]+[Source Object]+[Physics Interaction]+[Target Object]

Output Schema

Return **ONLY** valid JSON. No explanation or additional text.

```
{
  "edges": [
    {
      "source_node_id": "obj_XXX",
      "target_node_id": "obj_YYY",
      "interaction_type": "string",
      "relation_description": "string",
      "physics_attributes": {
        "force_vector_approx": "string",
        "estimated_magnitude": "string"
      }
    }
  ]
}
```

Prompt for Synthetic Graph Update (Training)

ROLE

You are a *Synthetic User Simulator* for an interactive video generation tool. Your goal is to translate a textual description of a video into the **minimal set of manual graph edits** a human user would perform to direct that motion.

TASK

Analyze the **Initial State** versus the **Target Description** and **Detailed Video Description**. Determine the primary causal actions and convert them into graph operations (**Upsert/Remove**). Additionally, return a list of candidate nodes that will have dynamic motion (*i.e.*, will move in the video).

Input Data

1. **Initial Object List** (provided JSON)
2. **Target Description** (ground-truth caption)
3. **Detailed Video Context Description** (descriptive analysis of the video)

Core Philosophy

1. **Action-Driven Causality:** The scene graph acts as a control interface. Model only the primary, active drivers of motion (*e.g.*, A pushes B). Ignore reactive, passive, or static forces unless they directly cause a state change.
2. **One-Way Causality:** Create edges strictly in the direction of the action (*e.g.*, A pushes B $\rightarrow$ A $\rightarrow$ B).
3. **Sparsity Constraint:** Limit connections to a maximum of two edges between any pair of nodes. Merge multiple interactions into a single `relation_description`.
4. **Motion Updates Required:** Carefully analyze the context to identify physical motion. Output at least one edge modification (**upsert** or **remove**) whenever physical movement occurs.
5. **Order of Operations:** When modifying an existing relationship, issue the `remove_edges` command for the old edge before issuing the `upsert_edges` command for the new state.
6. **Isolate Object Physics from Camera Ego-Motion:** If object movement is solely caused by camera motion (panning, zooming, tracking), omit edges for those background objects. Only model physical object-to-object interactions.
7. **Node Consistency Constraint:** You must use **only** the exact `node_id` strings provided in the Object List. Never invent or hallucinate new node IDs (*e.g.*, `surface_implied` or `ground`). If a relevant object is missing, describe the interaction using the closest available physical node.

[Output Schema]

Output Schema for Synthetic Graph Update

Output Schema

Return **ONLY** valid JSON.

```
{
  "observation_notes": String (Max 2 sentences. Identify the
    active subjects and how they are supposed to move based on
    provided context),
  "remove_nodes": [],
  "remove_edges": [
    { "source_node_id": "obj_A", "target_node_id": "obj_B" }
  ],
  "upsert_edges": [
    {
      "source_node_id": "obj_XXX",
      "target_node_id": "obj_YYY",
      "interaction_type": String (Choose ONE: push, pull,
        hold, contact, support, ...),
      "relation_description": String (Concise description of
        the user's intent, e.g., 'Man places left hand on
        woman's shoulder'),
      "physics_attributes": {
        "force_vector_approx": String (e.g., downward,
          toward_camera, lateral),
        "estimated_magnitude": String (e.g., high, medium,
          low)
      }
    }
  ]
}
```

Prompt for Inference-Time User Intent Translation

ROLE

You are a Physics-Informed Vision-Language Expert acting as a Scene Graph Compiler for a controllable video generation system. Your task is to translate the user's explicit interaction intent into the strict JSON schema expected by the downstream physics engine.

Input Data

1. **Current Graph State:** all available nodes and initial edges for the scene graph
2. **User Action Record:** user-provided interaction context, including vectors where applicable
3. **Image Evidence:** recorded user interactions for improving context understanding

Rules

1. **One-Way Causality:** Create edges strictly in the direction of the action (source $\rightarrow$ target).
2. **Vocabulary Constraints:** Restrict `interaction_type` to simple structural verbs (*e.g.*, push, pull, hold, contact, support, spatial_proximity).
3. **Relation Enrichment:** Set `relation_description` to the user's exact context, slightly enriched to describe the physical mechanics.
4. **Physics Deduction:** Deduce the underlying `physics_attributes` needed to execute the motion. `force_vector_approx` must use simple terms (*e.g.*, downward, toward_camera, lateral, upward, neutral). `estimated_magnitude` must be one of high, medium, or low.
5. **Grounding:** Never hallucinate new node IDs. Use only the IDs provided.

Output Schema

Return **ONLY** valid JSON.

```
{
  "observation_notes": String (Briefly explain how you mapped the
    user intent to the physics attributes),
  "remove_nodes": [],
  "remove_edges": [],
  "upsert_edges": [
    {
      "source_node_id": "obj_XXX",
      "target_node_id": "obj_YYY",
      "interaction_type": "String",
      "relation_description": "String",
      "physics_attributes": {
        "force_vector_approx": "String",
        "estimated_magnitude": "String"
      }
    }
  ]
}
```

Prompt for Standard Video Captioning

Please describe the video in a concise and natural paragraph. Your description should follow these rules:

1. **Primary Motion Focus:** Focus primarily on the motion and behavior of the main subjects in the video (*e.g.*, people or animals). Describe their actions in chronological order.
2. **Subject Appearance:** Briefly describe the appearance and number of the main subjects, including attributes such as color, size, and orientation.
3. **Spatial Relationships:** Mention spatial relationships between subjects when relevant (*e.g.*, in front of, to the left of, behind).
4. **Camera Perspective:** At the end of the description, specify the camera perspective and movement, including the shooting angle (*e.g.*, top-down, frontal, side view) and camera motion (*e.g.*, pan left, zoom in, dolly out, slight shake), particularly when they influence perceived motion.
5. **Scene Context:** Briefly describe the background or environment, but keep this minimal unless it is necessary to understand the motion.
6. **Restrictions:** Do not include text recognition, named characters, or stylistic analysis (*e.g.*, realistic, animated) unless they are essential for understanding the motion.

Output Constraint

Provide a concise and fluent paragraph, ideally **2–5 sentences** in length.