

# MoReGen: Multi-Agent Motion-Reasoning Engine for Code-based Text-to-Video Synthesis

Xiangyu Bai<sup>\*,1</sup>

He Liang<sup>\*,2</sup>

Bishoy Galoaa<sup>1</sup>

Utsav Nandi<sup>1</sup>

Shayda Moezzi<sup>1</sup>

Yuhang He<sup>†,3</sup>

Sarah Ostadabbas<sup>†,1</sup>




Figure 1. Samples generated from selected prompts in our dataset, MoReSet, using current commercial text-to-video models. Left: OpenAI’s Sora 2 [11]; Middle: Google DeepMind’s Veo3 [18]; Right: X AI’s Grok Imagine [49]. The first row illustrates failures in object counting and momentum conservation; the second row highlights misinferred Newtonian forces; the third shows incorrect velocity and pressure computations. Even the most advanced text-to-video models lack physical reasoning abilities and are unable to follow physics rules<sup>1</sup>.

## Abstract

While text-to-video (T2V) generation has achieved remarkable progress in photorealism, generating intent-aligned videos that faithfully obey physics principles remains a core challenge. In this work, we systematically study Newtonian motion-controlled text-to-video generation and evaluation, emphasizing physical precision and motion coherence. We introduce MoReGen, a motion-aware, physics-grounded T2V framework that integrates multi-agent large language models (LLMs), physics simulators, and renderers to generate reproducible, physically accurate videos from text prompts in the code domain. To quantitatively assess physical validity, we propose object-trajectory correspondence as a direct evaluation metric and present MoReSet, a benchmark of 1,275 human-annotated videos spanning nine classes of Newtonian phenomena with scene descriptions, spatiotemporal relations, and ground-truth trajectories. Using MoReSet, we conduct experiments on existing T2V models, evaluating their physical validity through both our MoRe metrics and existing physics-based evaluators. Our results reveal that state-of-the-art models struggle to maintain physical validity, while MoReGen establishes a

*principled direction toward physically coherent video synthesis.*

## 1. Introduction

Generative AI has gained unparalleled momentum in recent years, with text-to-video (T2V) models emerging as one of its most intuitive and rapidly advancing applications. These systems translate natural language prompts into rich visual narratives, powering both creative and scientific domains through automated video synthesis [9, 12, 32, 46, 48]. Despite remarkable progress in visual fidelity, existing T2V approaches remain predominantly appearance-driven rather than physics-grounded. They frequently generate videos that appear realistic yet violate fundamental principles of motion, dynamics, and causality, revealing a critical gap between visual realism and physical plausibility. Addressing this gap requires generative systems that reason about motion and obey physical constraints, not merely replicate static visual patterns (see Figure 1).

In this paper, we introduce **MoReGen**, a multi-agent motion-reasoning engine that redefines text-to-video generation through a motion-aware physics-grounded framework in the code domain. MoReGen unifies multi-agent large language models (LLMs), physics simulators, and rendering engines to generate reproducible and physically accurate videos directly from textual descriptions (see Figure 2).

\* Equal contribution. †Equal contribution corresponding authors.

<sup>1</sup>Department of Electrical and Computer Engineering, Northeastern University <sup>2</sup>College of Computer Science, University of Oxford <sup>3</sup>Microsoft Research. Code, dataset and more samples are available at our [GitHub page](#).




Figure 2. Overview of our multi-agent motion-reasoning engine (MoReGen) for physics-grounded text-to-video synthesis. MoReGen focuses on achieving high-precision Newtonian motion through coordinated multi-agent reasoning. Given a natural language prompt, the *text-parser* agent  $\mathcal{A}_{\text{text}}$  extracts physical parameters and motion descriptors, which are then translated into executable simulation code by the *code-writer* agent  $\mathcal{A}_{\text{coder}}$ . The resulting code is executed within a sandboxed *video-renderer* agent  $\mathcal{A}_{\text{render}}$  to produce a physically plausible video. This video serves as the basis for evaluator feedback, guiding enhancements in code robustness and physical fidelity for subsequent iterations. By leveraging open-source LLMs, few-shot tuning of  $\mathcal{A}_{\text{text}}$  and multi-modal evaluator, MoReGen enables accurate and reproducible Newtonian motion synthesis from natural language instructions.

In contrast to conventional diffusion-based pipelines that rely on probabilistic denoising, MoReGen directly generates executable code for physics simulators and rendering engines, enabling the synthesis of Newtonian motion dynamics from natural language instructions via a coder LLM agent. To ensure both physical validity and motion coherence, MoReGen employs an iterative multi-agent feedback loop, where motion-understanding models assess temporal consistency and dynamic plausibility, while vision-language models (VLMs) evaluate semantic and visual alignment. This interplay among agents fosters context-aware refinement, leading to video generations that are simultaneously physically consistent and visually faithful.

Recent advances in T2V have been largely driven by vision-transformer-based architectures. Researchers have extended pre-trained image diffusion backbones by inflating 2D transformers into 3D modules to integrate temporal dynamics [4]. Similar to the ongoing “arms race” in LLMs, the dominant paradigm emphasizes architectural scaling through massive datasets and high-cost computational resources [27]. While this approach has undeniably improved visual quality, with larger models consistently producing high-fidelity, intent-aligned videos from natural language prompts, their success remains primarily superficial [14].

Recent evaluation efforts [1, 20, 35] have revealed that, despite impressive aesthetics, most T2V systems still fail to reason about physical reality. Generated videos exhibit spatial and temporal inconsistencies, along with physically impossible motion trajectories [24]. The purely data-driven design of transformer-based architectures dictates that optimization favors pattern memorization over causal reasoning. Instead of inferring the underlying laws that govern motion, these models reproduce statistical correlations in the training data, producing outputs that mirror what was seen rather than reasoning about what should occur [17, 29]. The problem becomes particularly pronounced under out-of-distribution (OOD) prompts, where generated videos deviate sharply from physical reality as Figure 1 shows.

A parallel limitation arises in evaluation methodology. Most existing T2V models rely on appearance-based ob-

jectives such as Fréchet inception distance (FID) [51] and Fréchet video distance (FVD) [44] to assess performance. While these metrics quantify the distributional similarity of generated pixels to ground-truth data, they provide no indication of whether the generated videos obey underlying physical laws or maintain temporal coherence. As a result, researchers have turned to human-based evaluations [7, 9]; however human perceptual judgments are inherently qualitative and insufficient for evaluating precise physical relationships, limiting their utility in domains where quantitative accuracy is essential [8, 36, 43]. Although recent studies acknowledge this gap, explicit modeling of physical validity remains largely unexplored. Existing physics-informed evaluation methods rely on scene reconstruction metrics such as peak signal-to-noise ratio (PSNR) [28], which overlook temporal consistency, or on distributional statistics such as semantic adherence [6] and average Jaccard [1] derived from fine-tuned data encoders [35], which are themselves vulnerable to OOD failures.

To address these challenges, MoReGen introduces a paradigm shift toward physics-aware, motion-centric generative modeling. Beyond generating visually plausible videos, MoReGen enforces dynamic coherence, physical validity, and semantic alignment through the interaction of reasoning agents and physics simulators. The framework employs a multi-agent generative architecture in which a text-parser agent converts natural language descriptions into structured physical specifications, a code-writer agent translates those specifications into executable simulation code, a renderer agent generates motion-consistent videos under a physics simulator and an evaluator that analyzes generated videos using temporal trajectories, VLMs and LLMs. The text-parser is supervisedly fine-tuned on phenomenon-specific datasets to learn robust mappings between linguistic cues and physical parameters. This modular, feedback-driven design enables MoReGen to synthesize dynamically coherent and physically faithful motion sequences that can be quantitatively evaluated.

To quantitatively assess physical validity, we further propose object-trajectory correspondence as a direct evalua-

tion metric. Complementing this, we introduce **MoReSet**, a benchmark suite comprising 1,275 videos across nine canonical classes of Newtonian phenomena. Each video is human-annotated with descriptions and ground-truth object trajectories. To our knowledge, no existing dataset provides such a detailed pairing of physically grounded annotations with rendered videos. For our experiments, in addition to adopting the physics-based and data-centric metrics proposed by [1] and [7], we introduce new motion-centric measures, called **MoRe Metrics** that explicitly quantify trajectory fidelity and physical coherence. Through extensive experiments, we show that MoRe metrics provide a faithful quantitative assessment of physical validity, whereas conventional data-centric metrics fail to capture the motion-level physics accuracy in generated videos. Our key contributions are summarized as follows:

- We propose MoReGen, a multi-agent text-to-video generation pipeline that directly produces executable code for physics engines, synthesizing physics-grounded videos of Newtonian motion from natural language instructions.
- We design a modular multi-agent coding and rendering framework that transforms natural language into executable Newtonian simulations, enabling reproducible and physically valid video generation.
- We introduce MoReSet, a benchmark suite containing 1,275 videos demonstrating nine fundamental Newtonian physics, each paired with detailed scene annotations and ground-truth object motion trajectories.
- We develop and apply MoRe metrics, a novel motion-centric evaluation metric that reveals a critical research gap between trajectory fidelity and physical coherence in current T2V models, offering a principled foundation for physics-based generative evaluation.

## 2. Related Work

We review recent research on T2V models and evaluation, emphasizing why data-centric architectures and aesthetics-focused metrics fail to ensure or measure the physical validity. We also highlight the unique advantage of our framework and evaluation over traditional approaches.

### 2.1. Transformer-based Video Generation

Breakthroughs in denoising diffusion probabilistic models (DDPMs) [22] and transformer architectures [45] have revolutionized video generation. The former enabled realistic synthesis of realistic videos from Gaussian noise [41], while the latter introduced scalable modeling principles that improve generalization with increasing data and compute [14]. As a result, diffusion-based transformer (DiT) architectures [37] have come to dominate both industry and academia, setting new standards in open-source generative video modeling [21, 23, 46, 54]. Despite their success in visual fidelity, DiT-based models exhibit several fundamental limitations. Their dependence on massive datasets and high-

cost computational infrastructure makes large-scale training inaccessible to most researchers, where limited resources pose a significant barrier [38]. More critically, these models remain appearance-drive rather than physics-aware, they lack an understanding of motion dynamics and physical causality because of their memorization of their training distributions [3, 25]. When confronted with OOD scenarios, they often produce visually compelling yet physically implausible videos [20], revealing a persistent disconnect between aesthetic realism and genuine physical reasoning [6]. This limitation stems from the inductive biases inherent to transformer architectures, which favor statistical pattern memorization over grounded physical modeling and hinder their generalization beyond the training distribution [50].

### 2.2. Physics-aided Generation and Evaluation

Recently, video generation research has begun shifting from purely data-centric and aesthetic alignment approaches towards physics-informed modeling, recognizing that physical laws inherently govern spatial and temporal dynamics in all videos. Bastek et al. [8] proposed a framework that incorporates domain-specific physical guidance into diffusion backbones; Liu et al. [30] used simulators to generate motion based on video priors, and Zhang et al. [53] employed LLMs to penalize physically implausible outputs. Lv et al. [33] further advanced this direction by integrating LLM-based planning and simulation to condition diffusion models to enhance motion validity. While these efforts represent meaningful progress toward realism, they remain constrained by the limitations of diffusion backbones, which lack explicit causal reasoning and struggle to maintain full physical consistency. To better evaluate such physics-aware models, several benchmarks and metrics have been introduced. Kang et al. [25] introduced a 2D simulation testbed for assessing object movement and collisions; Bansal et al. [6] proposed VideoPhy, a benchmark evaluating physical validity through physical commonsense reasoning; and Guo et al. [20] presented T2VPhysBench, where human evaluators assess model adherence to twelve core physical laws. Although valuable, these approaches often rely on LLM- or VLM-based captioning and scoring pipelines, which introduce risks of hallucination, bias, and lack of grounded understanding, limiting their quantitative reliability.

In this paper, we address these challenges through MoReGen, a physics-grounded video generation framework that integrates reasoning agents with simulation engines to ensure physically coherent synthesis. Complementing it, we introduce MoReSet, a benchmark capturing a broad spectrum of Newtonian phenomena with dense annotations of motion trajectories and spatiotemporal relations. Together with MoRe Metrics, a suite of motion-centric evaluation metrics grounded in physical trajectory, our framework establishes a principled approach to both generating and evaluating physics-valid T2V models.

### 3. Introducing MoReGen

The MoReGen framework is designed to achieve physics-precise T2V generation by explicitly modeling Newtonian motion through multi-agent collaboration. Given a natural language prompt, MoReGen produces executable simulation code that accurately represents the Newtonian dynamics. MoReGen consists of three cooperative agents and an evaluator: a *text-parser* agent  $\mathcal{A}_{\text{text}}$ , a *code-writer* agent  $\mathcal{A}_{\text{coder}}$ , a *video-render* agent  $\mathcal{A}_{\text{render}}$  and multi-component evaluator  $\mathcal{E}$  (see Figure 2). All three agents are LLM-based (Qwen in our case). As illustrated in Algorithm 1,  $\mathcal{A}_{\text{text}}$  parses the raw description into a structured Newtonian specification, and  $\mathcal{A}_{\text{coder}}$  converts this specification into executable physics simulation code, which is run inside a sandboxed environment to obtain object configurations and trajectories.  $\mathcal{A}_{\text{render}}$  then produces a rendering script that consumes these trajectories to generate the video and the evaluator iteratively refines the video by providing feedback to  $\mathcal{A}_{\text{coder}}$ . By leveraging the reasoning and code-generation capabilities of LLMs, MoReGen can generate high-fidelity Newtonian motion videos that are physically consistent, requiring only few-shot training  $\mathcal{A}_{\text{text}}$ .

#### 3.1. Supervised Fine-Tuned Text-Parser Agent

Given a free-form textual description of a Newtonian phenomenon  $x$  (e.g., “a ball is launched at  $45^\circ$  with velocity 10 m/s”), the text-parser agent  $\mathcal{A}_{\text{text}}$  (green block in Figure 2) produces a structured simulation specification containing all required objects, parameters, and initial conditions. This task is challenging because (1) natural-language descriptions are often incomplete with missing parameters and conditions that are not explicitly stated, and (2) different physical processes rely on distinct governing variables—such as ramp inclination for slope, and launch velocities for projectiles. Without explicit supervision, generic prompts rarely specify sufficient constraints, leading untrained LLMs to merely infer physical parameters and generate incomplete or inconsistent representations.

To provide precise supervision for training, we first construct a phenomenon-specific structured schema for each class of Newtonian motion. These schemas serve as the target output format during supervised fine tuning (SFT) and specify (1) valid parameter ranges and physical units, (2) the essential physical entities (e.g., objects, anchors, constraints), (3) geometric and mechanical consistency requirements, and (4) physically coherent initial conditions. Each schema is instantiated through a task-specific prompt, which reliably produces high-quality structured specifications from raw textual descriptions. All generated outputs are manually validated to ensure physical correctness, parameter consistency, and conformity to the schema; the resulting (text, specification) pairs form the SFT training data.

We adopt SFT to enable  $\mathcal{A}_{\text{text}}$  to internalize the mapping

from natural language to structured physics specifications, instead of relying on the task-specific prompts or schema templates. After fine-tuning, the agent no longer requires phenomenon-specific prompts: a single general instruction suffices to generate complete and physically consistent specifications across all phenomena. During training, the model learns lightweight linguistic reasoning rules (e.g., interpreting “first two balls” or “pushed from the right” into object indices and force directions) and acquires the ability to infer missing parameters while avoiding hallucinated values. We fine-tune Qwen2.5-Coder-14B [40] on 1200 curated text–specification pairs using AdamW (learning rate  $1 \times 10^{-5}$ ) for five epochs. The model is optimized with supervised next-token loss:

$$\mathcal{L}_{\text{SFT}} = -\mathbb{E}_{(x,S) \sim \mathcal{D}} \left[ \sum_{t=1}^{|S|} \log p_{\theta}(S_t | x, S_{<t}) \right], \quad (1)$$

where  $S$  is the target structured specification paired with input description  $x$  sampled from all training pairs  $\mathcal{D}$ .

The fine-tuned model thus generates complete and physically coherent structured specifications, even under under-specified or diversely phrased descriptions. These structured outputs then serve as the authoritative input for subsequent code synthesis and video rendering.

#### 3.2. Code-Writer and Video-Render Agents

The middle stage of our pipeline (blue block in Figure 2) is responsible for converting the structured specification into an executable physics simulation and a rendered video. This stage is implemented by the code-writer agent  $\mathcal{A}_{\text{coder}}$ , which synthesizes the simulation code, and the video-render agent  $\mathcal{A}_{\text{render}}$ , which generates the rendering script and produces the final video from simulated trajectories.

**Code-Writer Agent ( $\mathcal{A}_{\text{coder}}$ ).** The code-writer agent receives the structured specification from  $\mathcal{A}_{\text{text}}$  as input and generates executable simulation code  $C_t$  that simulates the underlying Newtonian motion process. The code is based on a unified prompt that provides a general Python class framework defining the key simulation components, including space initialization, object creation, constraint setup, and simulation loop. Within this framework, the agent reads each field from the structured specification, assigns corresponding numerical parameters, and selects appropriate open-source physics engine APIs to instantiate bodies, shapes, and constraints. The generated code saves all object configurations and records the system states at every simulation step  $\Delta t$ , including position, velocity, and orientation, forming a complete temporal trajectory for each object. By leveraging such physics engines (e.g. Pymunk [39], Blender [10] or Manim [34]), the generated code models Newtonian motion physics precisely, ensuring that the resulting simulation adheres to real-world mechanics rather than statistical approximations.

---

**Algorithm 1** MoReGen Pipeline

---

**Require:** Raw text prompt  $x$ **Ensure:** Rendered video  $v_t$  and telemetry  $\mathcal{T}$ 

```
1: Text Parsing
2:    $\mathcal{S} \leftarrow \mathcal{A}_{\text{text}}(x)$  // produce structured specification
3: Code Synthesis
4:    $C_t \leftarrow \mathcal{A}_{\text{coder}}(\mathcal{S})$  // generate physics simulation code
5: Physics Simulation in Sandbox
6: /* get scene  $\mathcal{C}$  and trajectories  $\mathcal{T}$  from simulation */
7:    $(\mathcal{C}, \mathcal{T}) \leftarrow \text{Simulate}(C_t)$ 
8: Video Rendering
9:    $v_t \leftarrow \mathcal{A}_{\text{render}}(\mathcal{C}, \mathcal{T})$  //render video from simulation
10: return  $v_t, \mathcal{T}$ 
```

---

**Video-Render Agent** ( $\mathcal{A}_{\text{render}}$ ). The rendering agent executes the code  $C_t$  within a sandboxed environment and produces two synchronized outputs: (1) *telemetry data* recording each asset’s spatial position, velocity, and orientation over time; (2) *rendered video*  $v_t$  that visually depicts the simulated motion. The rendering stage is jointly synthesized by  $\mathcal{A}_{\text{render}}$  through a single general-purpose prompt, requiring no task-specific adaptation. Since visualization essentially involves reconstructing simulated geometries and replaying state trajectories, the same rendering logic generalizes naturally across all physical phenomena. The generated module employs `pygame`—which integrates seamlessly with `Pymunk`—to interpret the simulation configuration, rebuild scene geometry, and replay the recorded object states in real time. This unified rendering pipeline produces videos  $v_t$  that not only reproduce the simulated Newtonian motion with high temporal and spatial fidelity but also ensure visual coherence across diverse motion types. Besides, this decoupled design between trajectory generation and video rendering allows the framework to be easily extended to more realistic 3D rendering engines such as Unreal Engine, Blender, or Unity, enabling future integration with photorealistic simulation environments.

### 3.3. Evaluator Design

We introduce a multi-iteration refinement process for MoReGen to enhance the robustness and physical fidelity of the generated code and video. The evaluator (colored purple in Figure 2) analysis output videos ( $v_t$ ) and provide feedback to  $\mathcal{A}_{\text{coder}}$  to guide the refinement process. First, we detect object locations in the video using GroundedDINO [31] informed by object descriptors extracted from  $\mathcal{A}_{\text{text}}$ . These locations guide CoTracker3 [26] to estimate normalized object trajectories  $\mathcal{T}_{\text{est}}$ , after which an LLM (Qwen in this case) evaluates the similarity and divergence from trajectory  $\mathcal{T}$  obtained from simulation process. We also use Qwen2.5-VL [2] to assesses the video from two perspectives: physical plausibility based on a set of physics rules, and alignment with the intended prompt  $x$ . Finally, we use

LLM to synthesize feedback  $\mathcal{F}$  by summarizing trajectory alignment, physical correctness, and prompt fidelity:

$$\mathcal{F} = LLM(\text{“Summarize”}, < LLM(\text{“Traj”}, \mathcal{T}_{\text{est}}, \mathcal{T}), VLM(\text{“Phys”}, v_t), VLM(\text{“Intent”}, v_t, x) >). \quad (2)$$

This data is provided to  $\mathcal{A}_{\text{coder}}$  to guide improvements in both code generation and video synthesis for next iteration.

## 4. Experimental Design and Results

This section presents the MoReSet benchmark and MoRe metrics, a trajectory-centric evaluation framework quantifying the physical validity of generated videos beyond visual realism. We demonstrate that existing benchmarks and metrics fail to capture fundamental Newtonian principles and show that MoReGen consistently preserves both user intent and underlying physical relationships. Finally, we conduct ablation studies to assess the contribution of fine-tuning, model scale, and evaluator feedback to the physics accuracy of generated videos.

### 4.1. MoReSet Benchmark

We introduce MoReSet, a novel open-source benchmark for evaluating T2V models through the lens of physics fidelity. The dataset comprises of 1,275 videos spanning nine fundamental categories of Newtonian physics experiments: gravity, acceleration, collision, oscillation, momentum, buoyancy, inertia, pendular motion and pulley mechanics. The training set includes 1,200 Blender-generated simulations each paired with a free-form textual description and a high-quality structured JSON specification  $\mathcal{S}$  capturing complete physical parameters. The test set includes 75 real-world laboratory videos, each human-annotated with: a natural-language description detailing the scene composition, physical relationships between objects and their environment, and chronological dynamics; object-level labels and environmental attributes; and the camera perspective. To enable precise motion-based evaluation, every test video is annotated with key object identities and full pixel-level trajectories, extracted automatically via CoTracker3 [26], supervised and corrected by human annotators. Figure 3 shows representative samples, and Table ?? compares MoReSet with existing benchmarks for physics-enabled T2V model evaluations. Unlike prior datasets that target limited physics scenarios or rely on synthetic motion heuristics, MoReSet offers broad physical coverage and explicit trajectory annotations, providing a standardized testbed for quantitative, physics-grounded T2V evaluation.

### 4.2. Evaluation Metrics

Existing physics-aware metrics remain inadequate for assessing physical validity, as they emphasize perceptual similarity or in-distribution priors over causal motion accuracy. **MoRe Metrics.** Pixel-level similarity metrics (FID, FVD, PSNR, and LPIPS) or latent-space features from pretrained

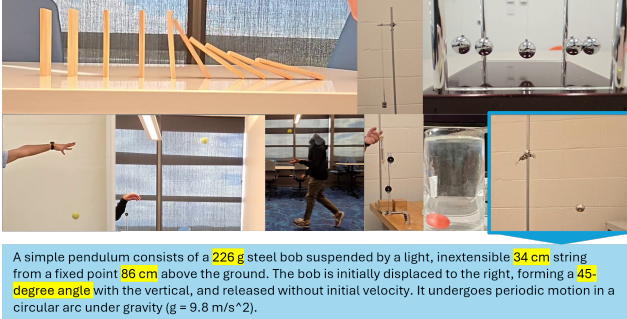


Figure 3. Sample frames and corresponding text prompts from our MoReSet dataset. Each image (extracted from videos) illustrates a distinct Newtonian physics phenomenon. The provided annotation for the pendulum corresponds to the rightmost video in the second row, with highlighted text emphasizing the numerical relationships depicted in the scene.

Table 1. Comparison of object motion trajectories across state-of-the-art T2V models, evaluated using MoRe metrics. To ensure comprehensive object coverage, the initial object detection stage was manually supervised by a human annotator. For objects that did not appear in videos, a random pixel is selected. DTW-N stands for normalized DTW distance, where trajectories are first centered and then scaled down to unit arc length. All videos are downsampled to 480p resolution at 10 frames per second. **Bold text** indicates best performance.

T2V Models	MoRe Metrics		
	DTW ↓	DTW-N ↓	Procrustes ↓
Wan2.2-Ti2V-5B [46]	17.18 ± 15.88	0.10 ± 0.09	0.65 ± 0.26
LTXV-2B-Distilled [21]	16.99 ± 16.05	0.11 ± 0.11	0.71 ± 0.28
CogVideoX-5B [23]	12.94 ± 13.56	0.09 ± 0.08	0.62 ± 0.26
Veo3 [18]	13.35 ± 12.96	0.07 ± 0.06	0.57 ± 0.27
Grok [49]	13.00 ± 11.58	0.08 ± 0.08	0.55 ± 0.31
Sora2 [11]	11.21 ± 12.76	0.08 ± 0.07	0.55 ± 0.30
NewtonGen [52]	17.88 ± 14.04	0.11 ± 0.09	0.62 ± 0.34
WISA [47]	12.67 ± 9.22	0.09 ± 0.08	0.68 ± 0.27
<b>MoReGen (Ours)</b>	<b>8.93 ± 9.61</b>	<b>0.06 ± 0.07</b>	<b>0.48 ± 0.30</b>

video encoders (Trajan, VideoPhy) are previously used to evaluate generated videos. Inspired by the recent success of point-tracking models [15, 16, 26], we propose MoRe metrics, a trajectory-based suite comprised of dynamic time warping (DTW), Normalized DTW (DTW-N) and Procrustes Analysis. These metrics directly measure the trajectory fidelity of key objects in motion. Given a structured specification  $\mathcal{S}$ , we identify the key object in motion, localize it in the initial frame via text-grounded object detection, and tracks its trajectory throughout the video. We then compute: (1) DTW [42] to align estimated trajectories with ground-truth sequences of varying temporal lengths by projecting them into a non-linear space; (2) DTW-N to normalize distance across clips of different scales and durations; and (3) Procrustes Analysis Score [19] to assess geometric and structural alignment between predicted and ground-truth trajectories.

**Data-Driven Physics Metrics.** To maintain comparability with community baselines, we also report results using two state-of-the-art physics-based evaluation frameworks from Trajan [1] and VideoPhy2 [6]. Trajan employs a point-track autoencoder trained to reconstruct spatiotemporal trajectories from BootsTAP [13], reporting average Jaccard (AJ

and Occlusion accuracy (OA) to measure temporal consistency and object visibility, respectively. VideoPhy2 provides Semantic Adherence (SA) and Physical Commonsense (PC) scores via its AutoEvaluator, a pretrained video-language model tuned on VidCon [5] for physics understanding.

Together, these methods allow us to evaluate not only visual fidelity and semantic alignment but also the underlying physical coherence of generated sequences.

### 4.3. Evaluation with MoRe Metrics

To capture the breadth of model capabilities, we evaluate three categories of T2V models using MoRe metrics on MoReSet test prompts: (1) medium-sized open-source models (Wan2.2 [46], LTX Video [21] and CogVideoX [23]), representing academically driven, general-purpose T2V approaches; (2) large-scale commercial models (Veo3 [18], Sora2 [11] and Grok Imagine [49]), reflecting the current state-of-the-art; and (3) physics-based models that support text-only input (NewtonGen [52] and WISA [47]), used as a baseline for comparison with MoReGen to expose their limitations and underscore our superior physics consistency. All models are evaluated using identical prompts from the MoReSet test set. As seen in Table 2, MoReGen outperforms existing state-of-the-art T2V engines by a large margin, indicating better alignment with ground truth object trajectories. These results align with qualitative analyses and human observations, enabled by the integration of a physics engine that produces realistic, physically accurate object motion. Additionally, we observe that while commercial T2V engines significantly outperform smaller models in visual aesthetics and overall quality, this advantage does not extend to motion trajectory accuracy. Notably, CogVideoX-5B and WISA, both of which are medium-sized open-source model, achieves comparable performance to OpenAI’s Sora2 [11] when evaluated using MoRe metrics. Similarly, NewtonGen and Wan2.2 exhibit comparable accuracy—an expected outcome given that NewtonGen and WISA are derived from Wan and CogVideo. This suggests that their approach to fine-tuning existing T2V models for physics validity is ineffective. Another notable observation is the poor stability of existing models when prompted with varying physics principles, as evidenced by large standard deviations in their performance. This variability suggests an imbalance in their training data, with different models excelling in different physical domains. In contrast, MoReGen maintains consistent performance across principles, owing to its rule-driven design and physics-grounded generation pipeline.

### 4.4. Evaluation with Data-Driven Physics Metrics

We further evaluate MoReGen and competing models on the MoreSet test set using data-driven physics metrics adopted from Trajan and VideoPhy2, to reveal their lim-

Table 2. Comparison of object motion trajectories across state-of-the-art T2V models, evaluated using MoRe metrics. To ensure comprehensive object coverage, the initial object detection stage was manually supervised by a human annotator. For objects that did not appear in videos, a random pixel is selected. DTW-N stands for normalized DTW distance, where trajectories are first centered and then scaled down to unit arc length. All videos are downsampled to 480p resolution at 10 frames per second. **Bold text** indicates best performance.

T2V Models	MoRe Metrics		
	DTW ↓	DTW-N ↓	Procrustes ↓
Wan2.2-TI2V-5B [46]	17.18 ± 15.88	0.10 ± 0.09	0.65 ± 0.26
LTXV-2B-Distilled [21]	16.99 ± 16.05	0.11 ± 0.11	0.71 ± 0.28
CogVideoX-5B [23]	12.94 ± 13.56	0.09 ± 0.08	0.62 ± 0.26
Veo3 [18]	13.35 ± 12.96	0.07 ± 0.06	0.57 ± 0.27
Grok [49]	13.00 ± 11.58	0.08 ± 0.08	0.55 ± 0.31
Sora2 [11]	11.21 ± 12.76	0.08 ± 0.07	0.55 ± 0.30
NewtonGen [52]	17.88 ± 14.04	0.11 ± 0.09	0.62 ± 0.34
WISA [47]	12.67 ± 9.22	0.09 ± 0.08	0.68 ± 0.27
<b>MoReGen (Ours)</b>	<b>8.93 ± 9.61</b>	<b>0.06 ± 0.07</b>	<b>0.48 ± 0.30</b>

itations. As shown in table 3, both frameworks fail to provide reliable assessments under OOD conditions, consistently assigning lower scores to our simulator-rendered videos due to their divergence from the metrics’ training distributions. Trajan fails to capture either motion correctness or perceptual realism as it awards disproportionately high scores to low-fidelity models such as LTX video that struggle to generate objects that even loosely resemble the user prompt, let alone produce accurate motion. The performance of LTX video according to this metric even surpasses larger commercial models (Veo3, Grok Imagine) that exhibit clearer physical plausibility. For VideoPhy2, the Semantic Adherence (SA) metric strongly correlates with visual quality rather than physical validity, ranking visually appealing but physically inconsistent models, namely Sora 2 and Grok Imagine, highest. While the Physical Commonsense (PC) metric better reflects actual dynamics, correctly favoring MoreGen and other simulator-based models, it still misranks CogVideoX above more physically accurate systems. In our evaluation, we observed that both Sora 2 and Grok Imagine produce more realistic physical motions than CogVideoX, however the PC metric fails to capture this. These observations confirm that current “physics-based” metrics fail to robustly capture the physical validity of generated videos.

#### 4.5. Qualitative Analysis

In Figure 4, we compare our generation against state-of-the-art open-source and commercial T2V models using Newton’s cradle, which illustrates collision and momentum. Qualitatively, the following observations can be made: first, most data-driven T2V models struggle with object counting, often generating an incorrect number of balls. When requested with five balls, LTX Video [21], Sora2 [11], Veo3 [18], NewtonGen [52] and WISA [47] all created four balls in the beginning, and while CogVideoX [23] and Wan2.2 [46] created the corrected number of balls, Wan2.2 introduced a sixth ball mid-sequence, and CogVideoX produced a largely static video. Secondly, many models misinterpret

Table 3. Comparison of physical validity and consistency across state-of-the-art T2V models (Wan2.2-TI2V-5B [46], LTXV-2B-Distilled [21], CogVideoX-5B [23], Veo3 [18], Grok Imagine [49], Sora2 [11], NewtonGen [52] and WISA [47]) using physics-based video evaluation metrics. For Trajan, videos were resized to 256×256 and 128 points were sampled. AJ denotes average Jaccard, and OA denotes occlusion accuracy. For VideoPhy, each video was rated from 1 to 5 based on semantic adherence (SA) and physics consistency (PC). We report the mean and standard deviation for all evaluated T2V models. **Bold text** indicates best performance.

T2V Models	Trajan [1]		VideoPhy2 [7]	
	AJ ↑	OA ↑	SA ↑	PC ↑
Wan2.2	0.38 ± 0.02	0.94 ± 0.02	2.88 ± 0.46	3.77 ± 0.68
LTXV	<b>0.79 ± 0.03</b>	<b>0.99 ± 0.02</b>	2.95 ± 0.28	3.75 ± 0.68
CogVideoX	0.17 ± 0.02	0.72 ± 0.04	2.83 ± 0.47	4.07 ± 0.59
Veo3	0.15 ± 0.02	0.82 ± 0.02	2.95 ± 0.33	3.49 ± 0.50
Grok	0.11 ± 0.02	0.72 ± 0.05	<b>3.04 ± 0.34</b>	3.76 ± 0.56
Sora2	0.45 ± 0.02	0.92 ± 0.02	3.01 ± 0.35	3.91 ± 0.52
NewtonGen	0.56 ± 0.02	0.96 ± 0.02	2.77 ± 0.42	4.21 ± 0.52
WISA	0.41 ± 0.02	0.86 ± 0.03	2.92 ± 0.45	4.01 ± 0.66
<b>MoReGen (Ours)</b>	0.10 ± 0.03	0.64 ± 0.05	2.73 ± 0.44	<b>4.53 ± 0.69</b>

the collision dynamics and momentum transfer as object addition or removal. This is particularly evident in LTX Video, Wan2.2, Veo3, and Grok Imagine [49], where a ball is inserted into the cradle perpendicular to the intended motion axis (left-right). Lastly, although NewtonGen and WISA failed to capture the correct number of balls and concept of Newton’s cradle, their physics-optimized design enables motion trajectories that appear physically plausible to the human eye; other models consistently introduced motion artifacts and physically inconsistent behaviors in their video outputs: In Wan2.2, the ball merges with the others and then splits again; in Sora, the system halts after the first ball strike; and in Veo3, the ball is lifted up while the string leans unnaturally to the right. In contrast, videos produced by our model demonstrate superior physical understanding, yielding a realistic Newton’s cradle with the correct number of balls and physics-valid motion. Notably, it captures subtle oscillations of the middle balls due to imperfect impact—details typically observed only in real-world scenarios.

#### 4.6. Ablation Study

To assess the contributions of individual model components, we conduct experiments across varying configurations of the Coder Agent and the presence or absence of supervised fine-tuning (SFT) and evaluator feedback (see Table 4). We compare the performance of a medium-sized coding model, Qwen2.5-Coder-14B, against GPT-5, and demonstrate that incorporating SFT significantly enhances the physical accuracy of generated videos—enabling the 14B model to outperform GPT-5 according to our evaluation metrics. We also found that incorporating even a single iteration of evaluator feedback significantly improves model performance. Initially, the 14B variant was unable to generate code executable by the rendering client. However, after adding feedback, Qwen successfully produced runnable code. Feedback also led to consistent performance gains for GPT-5 across all evaluation metrics. Finally, consistent

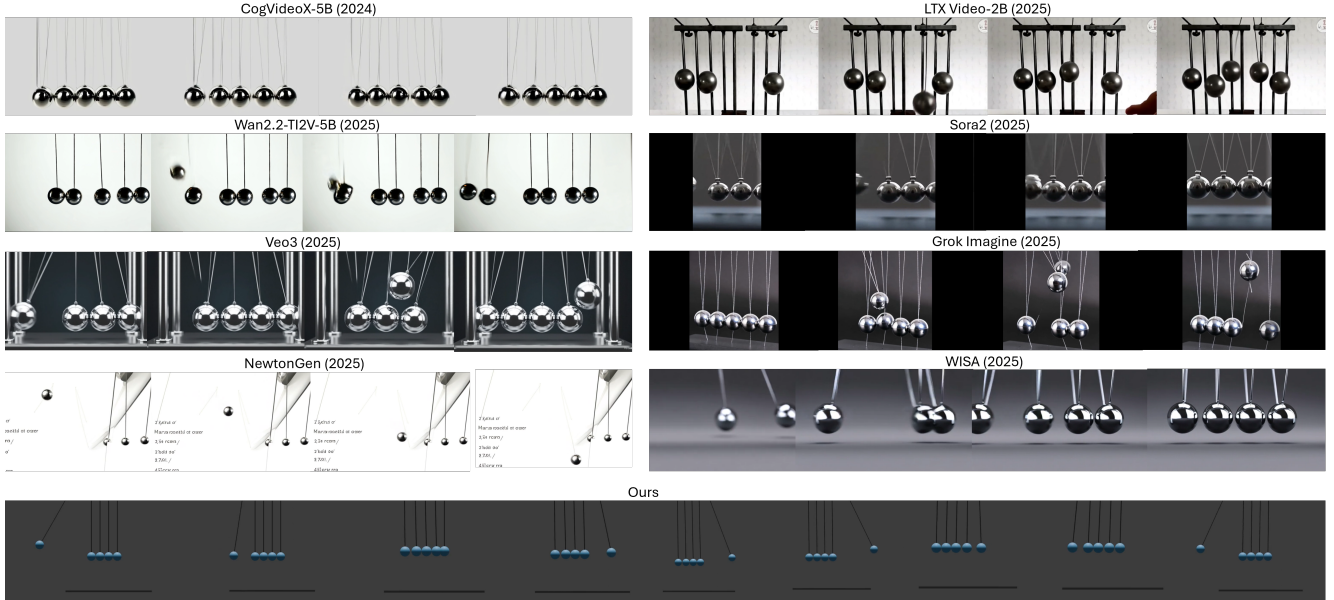


Figure 4. Qualitative comparison of our model with recent open-source and commercial models, prompted to generate a video of Newton’s cradle. We used the same prompt across the board: “Generate a video that showcase the following scene: Five shiny metal balls of a newton’s cradle is visible, along with parts of a single vertical string for each metal ball respectively. These strings keeps their respective metal ball suspended. The top part of the newton’s cradle is not visible. The camera faces all the five metal balls. The first and leftmost ball is at an angle of 30 degrees from the cradle and released. Due to gravity, the ball comes and strikes the second ball from the left. This causes momentum to be transferred to the fifth and the right most ball which is launched at a slightly lesser angle, having lost some momentum. This process keeps repeating itself till the rightmost ball has lost a lot of momentum when the video ends.” For Grok Imagine, we always select the first video; for WISA, we use Qwen3-4B to generate asset .json file from our prompt.

Table 4. Comparison of model performance gains from supervised fine-tuning (SFT) and evaluator feedback. Evaluation was conducted using the designated evaluation set from our dataset, with tests performed on Qwen2.5 Coder and GPT-5. Feedback involves one iteration of processing the video to extract object trajectories, comparing them with ground truth, and obtaining an overall evaluation from Qwen2.5-VL. GPT-5 then summarizes this evaluation into a list of actionable improvements, which are fed into the Coder Agent. If the code contains syntax error, GPT-5 is asked to correct the issue without providing the full code. **Bold text** indicates best performance.

Ablations			MoReGen (Ours)			Trajan			VideoPhy	
Coder Agent	SFT	Feedback	DTW ↓	DTW-N ↓	Procrustes ↓	AJ ↑	OA ↑	SA ↑	PC ↑	
Qwen2.5-Coder-14B	×	×				Unable to generate most cases due to syntax error in provided code by the agent.				
Qwen2.5-Coder-14B	×	✓	18.01 ± 16.53	0.08 ± 0.15	0.70 ± 0.52	0.09 ± 0.02	0.61 ± 0.06	2.51 ± 0.51	4.49 ± 0.57	
Qwen2.5-Coder-14B	✓	×	<b>8.93 ± 9.61</b>	<b>0.06 ± 0.07</b>	<b>0.48 ± 0.30</b>	0.10 ± 0.03	0.64 ± 0.05	<b>2.73 ± 0.44</b>	4.53 ± 0.69	
GPT-5	×	×	15.47 ± 14.60	0.07 ± 0.07	0.58 ± 0.29	0.32 ± 0.06	0.70 ± 0.04	2.60 ± 0.49	4.53 ± 0.55	
GPT-5	×	✓	14.13 ± 14.20	0.07 ± 0.07	0.51 ± 0.28	<b>0.38 ± 0.06</b>	<b>0.77 ± 0.04</b>	2.52 ± 0.50	<b>4.57 ± 0.52</b>	

with our observations in Section 4.4, Trajan and VideoPhy2 show disagreement with our metrics on the physics validity of the videos. However, physics consistency (PC) remains strong across all code-driven entries, consistently outperforming data-driven T2V models. Despite the accuracy of the physics engine, semantic adherence (SA), average Jac-card (AJ), and occlusion accuracy (OA) remain consistently low across all models, highlighting a bias in these metrics toward data- and aesthetics-driven approaches.

## 5. Conclusion

In this paper, we introduced MoReGen, a motion-aware, physics-grounded text-to-video (T2V) framework that unifies multi-agent LLMs, simulators, and renderers to generate reproducible and physically accurate videos from text prompts in the code domain. Additionally, we proposed MoRe Metrics as a direct evaluation approach

and presented MoReSet, a benchmark comprising human-annotated videos spanning nine classes of Newtonian phenomena with scene descriptions, spatiotemporal relations, and ground-truth trajectories for quantitative assessment of physics validity. Finally, we demonstrated that state-of-the-art models struggle with physical consistency, while our framework provides a principled path toward physically coherent video synthesis. MoReGen marks a new direction for physics-grounded T2V research, and we plan to extend it toward photorealistic generation in future work.

## References

- [1] Kelsey Allen, Carl Doersch, Guangyao Zhou, Mohammed Suhail, Danny Driess, Ignacio Rocco, Yulia Rubanova, Thomas Kipf, Mehdi SM Sajjadi, Kevin Murphy, et al. Direct motion models for assessing generated videos. *arXiv preprint arXiv:2505.00209*, 2025. 2, 3, 6, 7

- [2] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 5
- [3] Xiangyu Bai, Yedi Luo, Le Jiang, Aniket Gupta, Pushyami Kaveti, Hanumant Singh, and Sarah Ostadabbas. An Evaluation Platform to Scope Performance of Synthetic Environments in Autonomous Ground Vehicles Simulation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023. 3
- [4] Xiangyu Bai, Yedi Luo, Le Jiang, and Sarah Ostadabbas. Dual-conditioned temporal diffusion modeling for long driving video generation. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2025. 2
- [5] Hritik Bansal, Yonatan Bitton, Idan Szpektor, Kai-Wei Chang, and Aditya Grover. Videocon: Robust video-language alignment via contrast captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13927–13937, 2024. 6
- [6] Hritik Bansal, Zongyu Lin, Tianyi Xie, Zeshun Zong, Michal Yarom, Yonatan Bitton, Chenfanfu Jiang, Yizhou Sun, Kai-Wei Chang, and Aditya Grover. Videophy: Evaluating physical commonsense for video generation. *arXiv preprint arXiv:2406.03520*, 2024. 2, 3, 6
- [7] Hritik Bansal, Clark Peng, Yonatan Bitton, Roman Goldenberg, Aditya Grover, and Kai-Wei Chang. Videophy-2: A challenging action-centric physical commonsense evaluation in video generation. *arXiv preprint arXiv:2503.06800*, 2025. 2, 3, 7
- [8] Jan-Hendrik Bastek, WaiChing Sun, and Dennis M Kochmann. Physics-informed diffusion models. *arXiv preprint arXiv:2403.14404*, 2024. 2, 3
- [9] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22563–22575, 2023. 1, 2
- [10] Blender Foundation. Blender – a 3D Modelling and Rendering Package. <https://www.blender.org>, 2025. Version 4.2, accessed October 31, 2025. 4
- [11] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. Video generation models as world simulators, 2024. Accessed: 2025-11-06. 1, 6, 7
- [12] Muh Sulaiman Dadiono. Case study of representation of the rise of social criticism of the police through google veo 3 ai videos. *Academic Letters*, 1(4):53–58, 2025. 1
- [13] Carl Doersch, Pauline Luc, Yi Yang, Dilara Gokay, Skanda Koppula, Ankush Gupta, Joseph Heyward, Ignacio Rocco, Ross Goroshin, Joao Carreira, et al. Bootstap: Bootstrapped training for tracking-any-point. In *Proceedings of the Asian Conference on Computer Vision*, pages 3257–3274, 2024. 6
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2, 3
- [15] Bishoy Galoaa, Xiangyu Bai, Shayda Moezzi, Utsav Nandi, Sai Siddhartha Vivek Dhir Rangoju, Somaieh Amraee, and Sarah Ostadabbas. Look around and pay attention: Multi-camera point tracking reimagined with transformers. In *Thirteenth International Conference on 3D Vision*, 2026. 6
- [16] Bishoy Galoaa, Xiangyu Bai, Utsav Nandi, Sai Siddhartha Vivek Dhir Rangoju, Somaieh Amraee, and Sarah Ostadabbas. Unitrack: Differentiable graph representation learning for multi-object tracking. In *International Conference on Learning Representations (ICLR)*, 2026. 6
- [17] Bishoy Galoaa, Shayda Moezzi, Xiangyu Bai, and Sarah Ostadabbas. Motion-o: Trajectory-grounded video reasoning. *arXiv preprint arXiv:2603.18856*, 2026. 2
- [18] Google DeepMind. Veo: A text-to-video generation system. Technical report, Google DeepMind, 2024. Accessed: 2025-11-06. 1, 6, 7
- [19] John C Gower. Generalized procrustes analysis. *Psychometrika*, 40(1):33–51, 1975. 6
- [20] Xuyang Guo, Jiayan Huo, Zhenmei Shi, Zhao Song, Jiahao Zhang, and Jiale Zhao. T2vphysbench: A first-principles benchmark for physical consistency in text-to-video generation. *arXiv preprint arXiv:2505.00337*, 2025. 2, 3
- [21] Yoav HaCohen, Nisan Chiprut, Benny Brazowski, Daniel Shalem, Dudu Moshe, Eitan Richardson, Eran Levin, Guy Shiran, Nir Zabari, Ori Gordon, et al. Ltx-video: Realtime video latent diffusion. *arXiv preprint arXiv:2501.00103*, 2024. 3, 6, 7
- [22] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 3
- [23] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *arXiv preprint arXiv:2205.15868*, 2022. 3, 6, 7
- [24] Hyeonho Jeong, Chun-Hao P Huang, Jong Chul Ye, Niloy J Mitra, and Duygu Ceylan. Track4gen: Teaching video diffusion models to track points improves video generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 7276–7287, 2025. 2
- [25] Bingyi Kang, Yang Yue, Rui Lu, Zhijie Lin, Yang Zhao, Kaixin Wang, Gao Huang, and Jiashi Feng. How far is video generation from world model: A physical law perspective. *arXiv preprint arXiv:2411.02385*, 2024. 3
- [26] Nikita Karaev, Yuri Makarov, Jianyuan Wang, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Co-tracker3: Simpler and better point tracking by pseudo-labelling real videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6013–6022, 2025. 5, 6
- [27] Aonian Li, Bangwei Gong, Bo Yang, Boji Shan, Chang Liu, Cheng Zhu, Chunhao Zhang, Congchao Guo, Da Chen, Dong Li, et al. Minimax-01: Scaling foundation models with lightning attention. *arXiv preprint arXiv:2501.08313*, 2025. 2

- [28] Jingyun Liang, Yuchen Fan, Kai Zhang, Radu Timofte, Luc Van Gool, and Rakesh Ranjan. Movideo: Motion-aware video generation with diffusion model. In *European Conference on Computer Vision*, pages 56–74. Springer, 2024. 2
- [29] Minghui Lin, Xiang Wang, Yishan Wang, Shu Wang, Fengqi Dai, Pengxiang Ding, Cunxiang Wang, Zhengrong Zuo, Nong Sang, Siteng Huang, et al. Exploring the evolution of physics cognition in video generation: A survey. *arXiv preprint arXiv:2503.21765*, 2025. 2
- [30] Shaowei Liu, Zhongzheng Ren, Saurabh Gupta, and Shenglong Wang. Physgen: Rigid-body physics-grounded image-to-video generation. In *European Conference on Computer Vision*, pages 360–378. Springer, 2024. 3
- [31] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European conference on computer vision*, pages 38–55. Springer, 2024. 5
- [32] Yixin Liu, Kai Zhang, Yuan Li, Zhiling Yan, Chujie Gao, Ruoxi Chen, Zhengqing Yuan, Yue Huang, Hanchi Sun, Jianfeng Gao, et al. Sora: A review on background, technology, limitations, and opportunities of large vision models. *arXiv preprint arXiv:2402.17177*, 2024. 1
- [33] Jiayi Lv, Yi Huang, Mingfu Yan, Jiancheng Huang, Jianzhuang Liu, Yifan Liu, Yafei Wen, Xiaoxin Chen, and Shifeng Chen. Gpt4motion: Scripting physical motions in text-to-video generation via blender-oriented gpt planning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1430–1440, 2024. 3
- [34] Manim Community Developers. Manim: Mathematical Animation Engine. <https://www.manim.community/>, 2025. Version 0.18.1, accessed October 31, 2025. 4
- [35] Fanqing Meng, Jiaqi Liao, Xinyu Tan, Wenqi Shao, Quanfeng Lu, Kaipeng Zhang, Yu Cheng, Dianqi Li, Yu Qiao, and Ping Luo. Towards world simulator: Crafting physical commonsense-based benchmark for video generation. *arXiv preprint arXiv:2410.05363*, 2024. 2
- [36] Van Phi Nguyen, Tri Nhan Luong Ha, Huy Hieu Pham, and Quoc Long Tran. Training-free condition video diffusion models for single frame spatial-semantic echocardiogram synthesis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 670–680. Springer, 2024. 2
- [37] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023. 3
- [38] Xiangyu Peng, Zangwei Zheng, Chenhui Shen, Tom Young, Xinying Guo, Binluo Wang, Hang Xu, Hongxin Liu, Mingyan Jiang, Wenjun Li, et al. Open-sora 2.0: Training a commercial-level video generation model in \$200 k. *arXiv preprint arXiv:2503.09642*, 2025. 3
- [39] Pymunk Developers. Pymunk: A 2D Physics Library for Python. <https://www.pymunk.org>, 2025. Version 6.6.0, accessed October 31, 2025. 4
- [40] Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025. 4
- [41] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3
- [42] Pavel Senin. Dynamic time warping algorithm review. *Information and Computer Science Department University of Hawaii at Manoa Honolulu, USA*, 855(1-23):40, 2008. 6
- [43] Naichen Shi, Hao Yan, Shenghan Guo, and Raed Al Kontar. Multi-physics simulation guided generative diffusion models with applications in fluid and heat dynamics. *arXiv preprint arXiv:2407.17720*, 2024. 2
- [44] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphaël Marinier, Marcin Michalski, and Sylvain Gelly. Fvd: A new metric for video generation. *OpenReview*, 2019. 2
- [45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 3
- [46] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. 1, 3, 6, 7
- [47] Jing Wang, Ao Ma, Ke Cao, Jun Zheng, Zhanjie Zhang, Jiasong Feng, Shanyuan Liu, Yuhang Ma, Bo Cheng, Dawei Leng, et al. Wisu: World simulator assistant for physics-aware text-to-video generation. *arXiv preprint arXiv:2503.08153*, 2025. 6, 7
- [48] Jun Wang, Xijuan Zeng, Chunyu Qiang, Ruilong Chen, Shiyao Wang, Le Wang, Wangjing Zhou, Pengfei Cai, Jiahui Zhao, Nan Li, et al. Kling-foley: Multimodal diffusion transformer for high-quality video-to-audio generation. *arXiv preprint arXiv:2506.19774*, 2025. 1
- [49] xAI. Grok imagine, 2025. Accessed: 2025-11-06. 1, 6, 7
- [50] Tianyi Xie, Yiwei Zhao, Ying Jiang, and Chenfanfu Jiang. Physanimator: Physics-guided generative cartoon animation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 10793–10804, 2025. 3
- [51] Yu Yu, Weibin Zhang, and Yun Deng. Frechet inception distance (fid) for evaluating gans. *China University of Mining Technology Beijing Graduate School*, 3(11), 2021. 2
- [52] Yu Yuan, Xijun Wang, Tharindu Wickremasinghe, Zeeshan Nadir, Bole Ma, and Stanley H Chan. Newtongen: Physics-consistent and controllable text-to-video generation via neural newtonian dynamics. *arXiv preprint arXiv:2509.21309*, 2025. 6, 7

- [53] Ke Zhang, Cihan Xiao, Yiqun Mei, Jiacong Xu, and Vishal M Patel. Think before you diffuse: Llms-guided physics-aware video generation. *arXiv preprint arXiv:2505.21653*, 2025. [3](#)
- [54] Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. Open-sora: Democratizing efficient video production for all. *arXiv preprint arXiv:2412.20404*, 2024. [3](#)