# PhyGDPO: Physics-Aware Groupwise Direct Preference Optimization for Physically Consistent Text-to-Video Generation
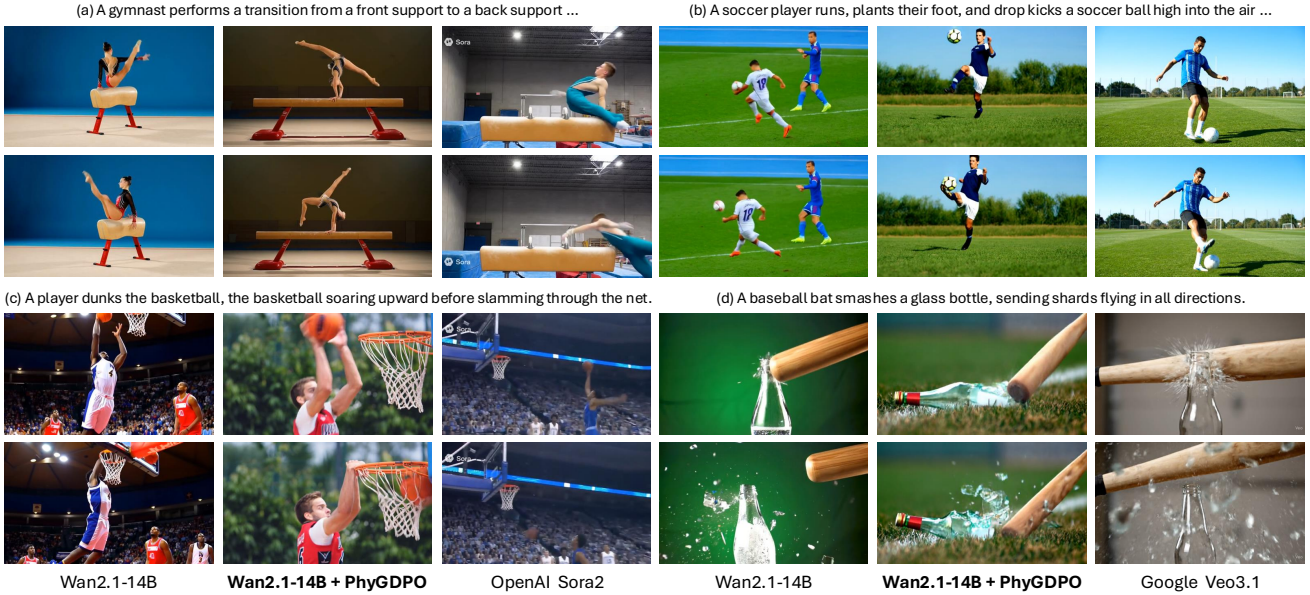
**Yuanhao Cai**[1,2,*], **Kunpeng Li**[1], **Menglin Jia**[3], **Jialiang Wang**[1], **Junzhe Sun**[1], **Feng Liang**[1], **Weifeng Chen**[1], **Felix Juefei-Xu**[1], **Chu Wang**[1], **Ali Thabet**[1], **Xiaoliang Dai**[1], **Xuan Ju**[4], **Alan Yuille**[2,†], **Ji Hou**[1,†]

[1]Meta Superintelligence Labs, [2]Johns Hopkins University, [3]Meta BizAI, [4]CUHK
[*]Work was done while Yuanhao Cai was an intern in Meta Superintelligence Labs, [†]Equal Advising

Recent advances in text-to-video (T2V) generation have achieved good visual quality, yet synthesizing videos that faithfully follow physical laws remains an open challenge. Existing methods mainly based on graphics or prompt extension struggle to generalize beyond simple simulated environments or learn implicit physical reasoning. The scarcity of training data with rich physics interactions and phenomena is also a problem. In this paper, we first introduce a Physics-Augmented video data construction Pipeline, PhyAugPipe, that leverages a vision–language model (VLM) with chain-of-thought reasoning to collect a large-scale training dataset, PhyVidGen-135K. Then we formulate a principled Physics-aware Groupwise Direct Preference Optimization, PhyGDPO, framework that builds upon the groupwise Plackett–Luce probabilistic model to capture holistic preferences beyond pairwise comparisons. In PhyGDPO, we design a Physics-Guided Rewarding (PGR) scheme that embeds VLM-based physics rewards to steer optimization toward physical consistency. We also propose a LoRA-Switch Reference (LoRA-SR) scheme that eliminates memory-heavy reference duplication for efficient training. Comprehensive experiments show that our method significantly outperforms state-of-the-art open-source methods on the PhyGenBench and VideoPhy2 datasets.

**Project Page:** https://caiyuanhao1998.github.io/project/PhyGDPO

∞ Meta



(a) A gymnast performs a transition from a front support to a back support ...

(b) A soccer player runs, plants their foot, and drop kicks a soccer ball high into the air ...

(c) A player dunks the basketball, the basketball soaring upward before slamming through the net.

(d) A baseball bat smashes a glass bottle, sending shards flying in all directions.

Wan2.1-14B    **Wan2.1-14B + PhyGDPO**    OpenAI Sora2    Wan2.1-14B    **Wan2.1-14B + PhyGDPO**    Google Veo3.1

**Figure 1** Text-to-video generation on four challenging action categories: (a) gymnastics, (b) soccer, (c) basketball, and (d) glass smashing. When using our post-training method, PhyGDPO, on Wan2.1-T2V-14B [58], the model yields more physically plausible results than OpenAI Sora2 [7] and Google Veo3.1 [20] by generating well-structured human bodies and reasonable physical interactions such as foot kicking the soccer, basketball passing through the hoop, and glass shattering.

## 1  Introduction

With the development of computing power and the increasing scale of training data, text-to-video (T2V) generation [7, 11, 58] has witnessed significant progress. However, accurately and consistently modeling the physics in the generated video still remains challenging and less explored. Improving the physics reasoning capability of video generation models can make them closer to real-world simulators, which can benefit a wide range of applications such as video gaming [10, 55], autonomous driving [61, 64, 69], robotics [1, 2, 16], film making [51, 60, 65], *etc.*

To improve the physics modeling, graphics-based methods [32] rely on simulation engines to specify physical parameters of simple scenarios such as perfectly elastic collisions and basic rigid body dynamics. Nonetheless, it is impractical to apply them in real-world scenes as the environments are too complex to parameterize.

Another technical route [67, 72] is based on prompt extension with a large language model (LLM). These methods adopt an LLM to extend the prompts with explicit physics laws and phenomena, and then use the extended prompts as input to simulate the physics by iteratively generating the video or finetuning the model on a small subset. These methods simply follow the LLM-augmented prompts and lack the ability of thinking in physics. Instead, they use the LLM as their surrogate brain and outsource the reasoning process to it. Even so, the prompt following ability of current T2V models is still limited and current LLM's physics reasoning ability is also weak and often erroneous, which may in turn mislead the T2V model when it follows such guidance to generate videos.

To learn the implicit physics, most current T2V foundation models are trained on massive collections of high-quality text–video data pairs. However, collecting and annotating such data is extremely costly and labor-intensive. Moreover, models obtained through such supervised fine-tuning (SFT) or training still exhibit limited physics reasoning ability. For instance, OpenAI Sora2 [7] and Google DeepMind Veo3.1 [20] often fail in complex human motions or physical phenomena, as shown in Fig. 1, 4, 5, and 6. The key reason is that there are no negative training data to provide contrastive signals discouraging physically inconsistent generations.

The emergence of direct preference optimization [52] (DPO) may provide a potential solution to address the limitations of SFT. However, directly using DPO may encounter three problems: (i) Lack of training data pairs that comprehensively capture physical activities, interactions, and phenomena. (ii) The supervision is mainly based on the Bradley–Terry (BT) probabilistic model, which only compares pairwise data samples. Such binary win/lose feedback lacks the guidance of physics-aware rewarding and struggles to capture inherently holistic global preference signals, such as physical plausibility and motion smoothness. (iii) Vanilla DPO duplicates the training model as a fixed reference, which consumes substantial GPU memory and significantly decreases computational efficiency.

To address these issues, we firstly propose a data construction pipeline, PhyAugPipe, that exploits a vision-language model (VLM) to filter text-video data pairs capturing rich physical interactions and phenomena from a large T2V data pool by parsing the entities and reasoning their actions and outcomes using our designed chain-of-thought (CoT) [63] rule for each data sample. Then we use PhyAugPipe to collect a training dataset, PhyVidGen-135K, containing over 135K text-video data pairs. Secondly, we formulate a principled Physics-aware Groupwise Direct Preference Optimization (PhyGDPO) framework for physically consistent T2V generation. Our PhyGDPO is based on the groupwise Plackett-Luce (PL) probabilistic model, which stimulates the probability distribution over a ranked list of candidate videos and thus enables holistic preference adaptation rather than being constrained to isolated pairwise judgments. Plus, the original DPO training uses the model itself as reward [52] but the model itself can not provide sufficiently good signals and knowledge of physics. Hence, we propose to embed more physics reward during training. Unlike GRPO [54], which suffers from slow speed and unstable optimization, we propose a Physics-Guided Rewarding (PGR) scheme that leverages a physics-aware VLM to guide the DPO data sampling favor more challenging physical actions and allow physics-violating samples to exert stronger influence during preference optimization. To improve the DPO training efficiency and stability, we also propose a LoRA [28]-Switch Reference (LoRA) mechanism that does not need to copy the full model as the reference, which occupies redundant GPU memory. Instead, we attach a trainable LoRA module with a switch to the frozen backbone. Benefit from our dataset and proposed techniques, PhyGDPO significantly improves the physical plausibility of the T2V model, Wan2.1-T2V-14B [58], and yields better results than the two recent best commercial T2V models, OpenAI Sora2 [7] and Google DeepMind Veo3.1 [20] on some challenging action categories, as shown in Fig. 1, 4, 5, and 6.

Our contributions can be summarized as follows:

- We formulate a principled DPO framework, PhyGDPO, for physically consistent text-to-video generation.

- We design PGR to guide data sampling and preference optimization favor challenging physical actions. We propose LoRA-SR scheme to reduce GPU memory occupancy for more efficient and stable DPO training.

- We present PhyAugPipe to construct physics-rich text-video pairs. We use PhyAugPipe to collect a training dataset, PhyVidGen-135K, with over 135K data pairs for studying physically consistent T2V generation.

- Experiments demonstrate that our PhyGDPO outperforms SOTA methods on the PhyGenBench and VideoPhy2 datasets while yielding higher human preference in the user study of physical realism.
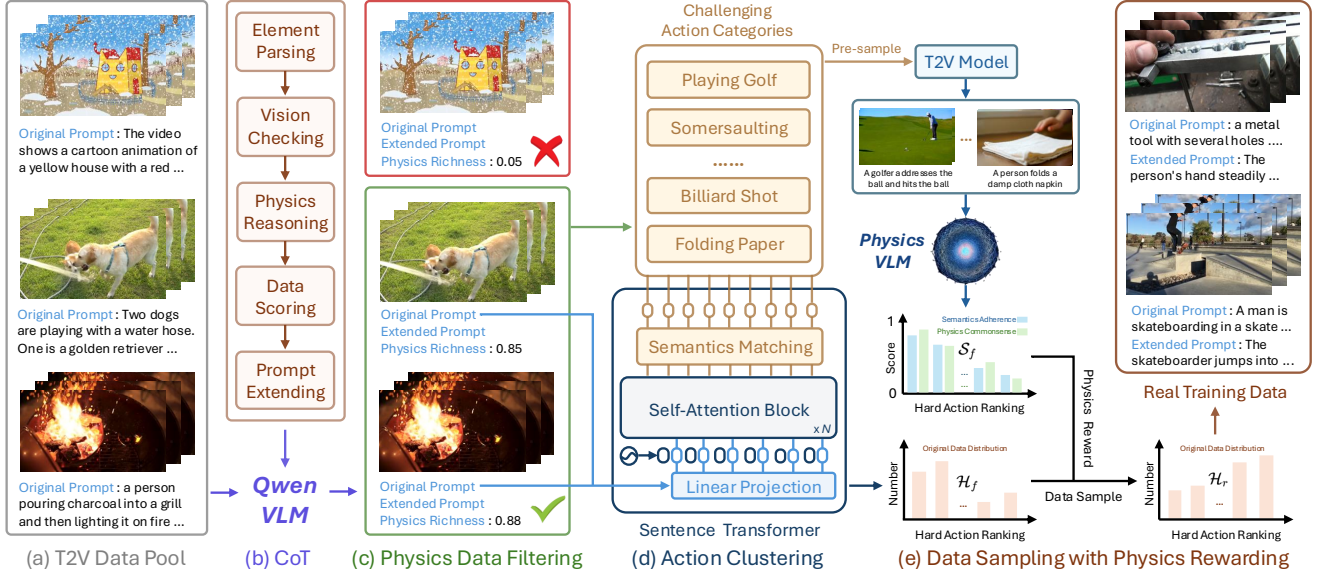
**Figure 2** Our physics-augmented video data construction pipeline (PhyAugPipe) first adopts a VLM, Qwen-2.5-72B-Instruct [68], following our designed chain-of-thought (CoT) rule in (b) to select text-video data pairs that contain rich physics interactions and phenomena from a large-scale high-quality text-video data pool in (a). Then in (d), we perform action clustering on the filtered data pairs from (c) through the semantics matching via a sentence Transformer [53]. Subsequently, in (e), we adopt a physics-aware VLM, VideoCon-Physics [3], to evaluate the difficulty of different action categories and then sample the text-video pairs accordingly as the wining cases of our training data for preference optimization.

## 2 Method

### 2.1 Physics-Augmented Video Data Construction

Our physics augmented video data construction pipeline (PhyAugPipe) is shown in Fig. 2. We use a VLM, Qwen-2.5-72B-Instruct [68], with our chain-of-thoughts (CoT) in (b) to filter videos with rich physics interactions and phenomena in (c) from a T2V data pool in (a). Then in (d), we cluster the filtered data according to their action categories and sample the data with the rewards of physics-aware VLM, VideoCon-Physics [3], in (e).

**Data Filtering with CoT.** As shown in Fig. 2 (b), Qwen2.5 with our CoT first parses the elements from the given prompts and video frames, including the physical objects with materials, actions and forces between them. Based on the parsed elements, Qwen2.5 reasons how the entities interact and what results, and rates the physics richness of the data with a score from 0 to 1. Eventually, Qwen2.5 extends the prompt with explicit physics reasoning. Although we explore the implicit physics reasoning ability of T2V models, but our dataset still include the extended prompts to support other research purposes, such as LLM-guided T2V. Please refer to Alg. 1 for the details of our CoT rules. Then in Fig. 2 (c), we threshold the data according to the estimated physics richness.

**Action Clustering via Semantics Matching.** After data filtering, the retained samples may exhibit distributional imbalance on different actions. To check this issue, we categorize the filtered data into semantically coherent action clusters. Specifically, in Fig. 2 (d), we first compile a list containing $K_a$ challenging action categories. Then we feed the original text prompt of each filtered sample and the action list into a sentence Transformer [53] to perform fuzzy semantics matching, thereby determining the action category to which each data sample belongs. We count the number of data samples in each action category to obtain the distribution histogram $\mathcal{H}_f$ for the filtered data.

**Data Sampling with Physics Rewarding.** As different action categories vary significantly in physical complexity thus leading to different generation difficulties for the T2V model, we propose to balance the training data accordingly. In particular, we first evaluate how well the T2V model performs on each action category. Based on the performance distribution, we adjust the sampling ratio by allocating more samples to the categories where the model performs poorly. As shown in Fig. 2 (d), we first select the top-$n_c$ samples with the highest semantics matching scores within each category as its representatives. Then we employ a physics-aware VLM, VideoCon-Physics [3], to evaluate each video representative. VideoCon-Physics outputs a semantics adherence score and a physics commonsense score $\in [0, 1]$ to measure the physics plaucibility of the video. We average the two scores to obtain the overall score for each video. Then we compute the mean score of all representatives within each category to obtain the category score. In this way, we derive the score distribution histogram $\mathcal{S}_f$. Then we sample the data according

---
**Algorithm 1** Brief Pipeline of our Chain-of-Thoughts Prompts
---
**Require:** Original prompt $p$, video frames $V$
**Ensure:** JSON object containing: original, parse, reason, extended, physics_richness, physics_label

1: **Step 1: Element Parsing**
2: Extract entities, actions, forces, and outcomes from $(p, V)$ using the VLM.
3: Ensure all agents are included and avoid unsupported or speculative items.
4: Save results into the dictionary *parse*.

5: **Step 2: Vision Checking**
6: Compare *parse* with video frames and the original prompt.
7: Remove hallucinated elements and add missing visible entities or interactions.
8: Update *parse* accordingly.

9: **Step 3: Physics Reasoning**
10: Produce a concise explanation describing how the parsed entities interact through physical forces and lead to observed outcomes.
11: Store the explanation as *reason*.

12: **Step 4: Data Scoring**
13: Compute physics_richness $\in [0, 1]$ based on:
   – number and interaction of entities;
   – presence of explicit forces and outcomes;
   – causal clarity in *reason*;
   – penalties from camera motion, stylization, or static-aftermath filters.
14: Set physics_label $= 1$ if physics_richness $\geq 0.60$, otherwise 0.

15: **Step 5: Prompt Extending**
16: Extend the original prompt by inserting causal physical details based on *reason*, without adding new entities, forces, or sensory descriptions.
17: Limit the extension to at most 100 words and store it as *extended*.

18: **return** JSON object containing all computed fields.
---

to the difficulty of each action category. More specifically, we define the difficulty of the $k$-th action category as $d_k = 1 - \mathcal{S}_f(k)$ and assign a sampling weight following an exponential form as $r_k = \exp(\tau d_k)$, where $\tau$ is a hyperparameter that controls how strongly the sampling favors difficult action categories. Given the total sampling budget $N$, the number of data pairs sampled from the $k$-th action category is formulated as

$$\mathcal{H}_r(k) = \min \left( \mathcal{H}_f(k), \, N \cdot \frac{r_k}{\sum_{j=1}^{K_a} r_j} \right), \tag{1}$$

where $\mathcal{H}_r$ denotes the distribution histogram after data sampling with physics rewarding. This sampling strategy allocates more data to challenging actions, encouraging the model to learn more complex physics during training.

## 2.2 Physics-Aware Groupwise Direct Preference Optimization

Existing DPO algorithms are based on Bradley–Terry (BT) probabilistic model, which compares pairwise data samples, showing limitations in capturing global and holistic preference signals and aligning with human feedback. Besides, previous DPO [39, 57] mainly use the T2V model itself as the reward [52], but the model itself exhibits weak physics reasoning ability as it usually generates physically inconsistent videos. To address these issues, we formulate a Physics-Aware Groupwise Direct Preference Optimization (PhyGDPO), as shown in Fig. 3.

**Groupwise Probability.** We denote the reward as $r(c, x_0)$ with the generation $x_0$ and condition $c$. Different from the normal DPO, we start from the groupwise Plackett-Luce (PL) probabilistic model, as shown in Fig. 3 (c). We adopt the real video as the winning case $x_0^w$ because it always follows physical laws and a set of generated videos as

(a) Textual Prompts  (b) LoRA-Switch Reference  (c) Groupwise Direct Preference Optimization with Physics-Guided Rewarding
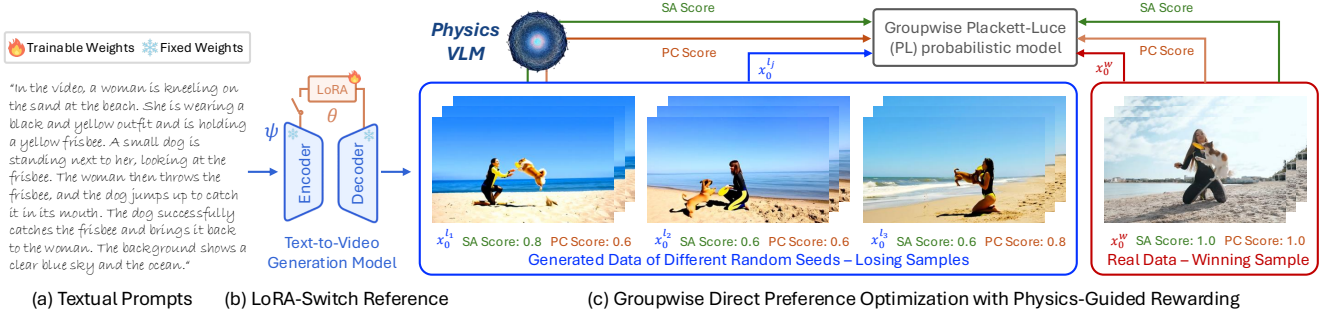
**Figure 3** Overview of our PhyGDPO framework. The text prompts without physics reasoning extension in (a) are fed into the T2V model. We propose a LoRA-switch reference scheme in (b) to save the GPU memory and increase the training stability. In (c), PhyGDPO is based on the groupwise Plackett-Luce (PL) probabilistic model and adopts a physics-aware VLM, VideoCon-Physics [3], to reward the DPO training with the real video as the wining sample as it has perfect physics.

the losing cases $\mathcal{G}^l(c) = \{x_0^{l_1}, \ldots, x_0^{l_m}\}$. The preference probability of the PL model is formulated as

$$p_{\text{PL}}(x_0^w \mid \mathcal{G}^l(c), c) = \frac{\exp(r(c, x_0^w))}{\sum_{j=1}^m \exp(r(c, x_0^{l_j}))}. \tag{2}$$

$r(c, x_0)$ can be parameterized by a neural network $\phi$ [52] and estimated via maximum likelihood training as

$$\mathcal{L}_{\text{PL}}(\phi) = -\mathbb{E}_{c, \mathcal{G}^l(c)}\left[\log \frac{\exp(r_\phi(c, x_0^w))}{\sum_{j=1}^m \exp(r_\phi(c, x_0^{l_j}))}\right]. \tag{3}$$

Vanilla DPO [39, 57] represents $r(c, x_0)$ by the T2V model itself as

$$r(c, x_0) = \beta \log \frac{p_\theta^*(x_0|c)}{p_\psi(x_0|c)} + \beta \log Z(c), \tag{4}$$

where $p_\theta$ and $p_\psi$ denote the conditional probabilities of trained model $\theta$ and reference model $\psi$. $p_\theta^*$ and $Z(c)$ are the unique global optimal solution and the partition function. We plug Eq. (4) into Eq. (2) and drop the group-constant term $\beta \log Z(c)$ as it does not affect the optimization direction to obtain the groupwise DPO (GDPO) loss as
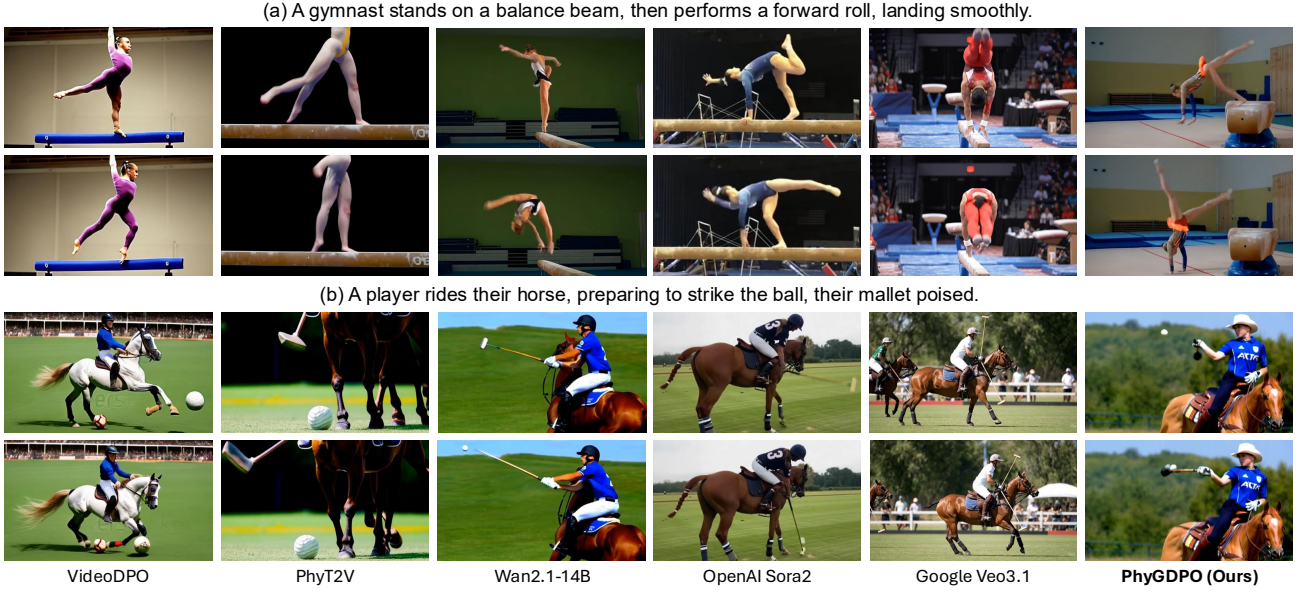
$$
\begin{aligned}
\mathcal{L}_{\text{GDPO}}(\theta) &= -\mathbb{E}_{c, \mathcal{G}^l(c)}\left[\log \frac{\exp(\beta f_\theta(x_0^w, c))}{\sum_{j=1}^m \exp(\beta f_\theta(x_0^{l_j}, c))}\right] \\
&= \mathbb{E}_{c, \mathcal{G}^l(c)}\left[\log \frac{\sum_{j=1}^m \exp(\beta f_\theta(x_0^{l_j}, c))}{\exp(\beta f_\theta(x_0^w, c))}\right] \\
&= \mathbb{E}_{c, \mathcal{G}^l(c)}\left[\log\left(\sum_{j=1}^m \exp(\beta(f_\theta(x_0^{l_j}, c) - f_\theta(x_0^w, c)))\right)\right].
\end{aligned}
\tag{5}
$$

where we define the function $f_\theta(x_0, c)$ as the log-likelihood ratio as

$$\log \frac{p_\theta(x_0|c)}{p_\psi(x_0|c)} = \sum_{k=1}^T \log \frac{p_\theta(x_{t_{k-1}} \mid x_{t_k}, c)}{p_\psi(x_{t_{k-1}} \mid x_{t_k}, c)} \triangleq \sum_{k=1}^T \Delta_k. \tag{6}$$

Here the terminal distribution is inherently represented as the joint transition process over multiple timesteps $t_k$ by the definition of diffusion or flow matching models. Subsequently, we reformulate Eq. (5) and estimate its upper bound using the Jensen's inequality into a single timestep for more efficient training as

$$
\begin{aligned}
\mathcal{L}_{\text{GDPO}}(\theta) &= \mathbb{E}_{c, \mathcal{G}^l(c)}\left[\log\left(\sum_{j=1}^m \exp(\beta T \mathbb{E}_k[\Delta_k^{l_j} - \Delta_k^w])\right)\right] \\
&\leq \mathbb{E}_{c, \mathcal{G}^l(c), k}\left[\log\left(\sum_{j=1}^m \exp(\beta T[\Delta_k^{l_j} - \Delta_k^w])\right)\right].
\end{aligned}
\tag{7}
$$

(a) A gymnast stands on a balance beam, then performs a forward roll, landing smoothly.

(b) A player rides their horse, preparing to strike the ball, their mallet poised.

| VideoDPO | PhyT2V | Wan2.1-14B | OpenAI Sora2 | Google Veo3.1 | **PhyGDPO (Ours)** |

**Figure 4** Results of two challenging actions (gymnastics and polo) on VideoPhy2. Our method generates more physically consistent videos, showing coherent, deformation-free gymnastic movements and realistic ball–mallet striking interactions.

Although the upper bound in Eq. (7) allows single-timestep training, it needs to infer the model $2m + 2$ times for each group, which requires a long time and large GPU memory. To handle this issue, we exploit an inequality as

$$\sum_{j=1}^{m} e^{x_j} \leq \prod_{j=1}^{m} \left(1 + e^{\alpha_j x_j}\right)^{\gamma_j}, \quad 0 < \alpha_j \leq 1, \ \gamma_j \geq 1/\alpha_j. \tag{8}$$

Please refer to the supplementary for the detailed proof process of this inequality. Subsequently, we can further formulate the upper bound of $\mathcal{L}_{\mathrm{GDPO}}(\theta)$ into a single sample within each group for efficient training as

$$
\begin{aligned}
\mathcal{L}_{\mathrm{GDPO}}(\theta) &\leq \mathbb{E}_{c,\mathcal{G}^l(c),k}\Big[\sum_{j=1}^{m} \gamma_j \log\Big(1 + \exp\big(-\alpha_j \beta T[\Delta_k^w - \Delta_k^{l_j}]\big)\big)\Big] \\
&= \mathbb{E}_{c,\mathcal{G}^l(c),k,j}\Big[-\gamma_j \log \sigma\big(-\alpha_j \beta T(\Delta_k^w - \Delta_k^{l_j})\big)\Big].
\end{aligned}
\tag{9}
$$

This upper bound allows efficient training with single data pair sampling in a single timestep for each iteration.
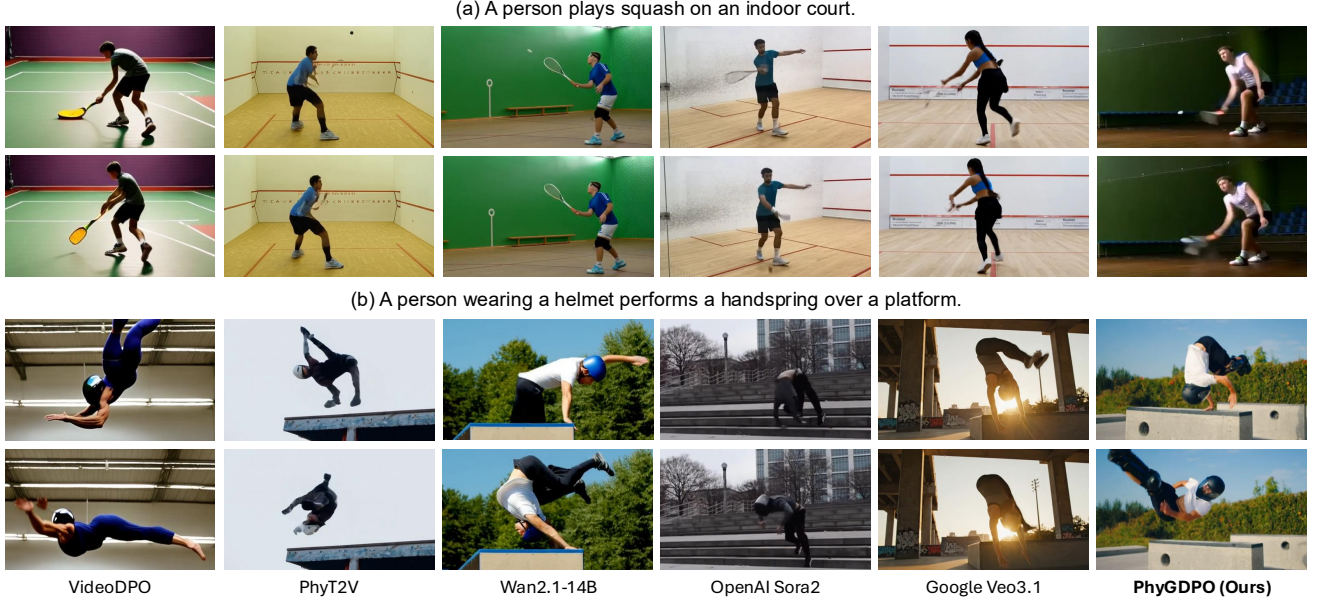
**Physics-Guided Rewarding.** As aforementioned, normal DPO framework uses the T2V model itself as the reward like Eq. (4), which cannot provide sufficiently good physics knowledge. To embed more physics guidance into our DPO framework, we parameterize $\gamma_j$ and $\alpha_j$ in Eq. (9) by the normalized semantics adherence score $s_j^{sa}$ and physics commonsense score $s_j^{pc}$ predicted by a physics-aware VLM, VideoCon-Physics [3], as

$$
v_j = 1 - \frac{s_j^{sa} + s_j^{pc}}{2}, \quad \gamma_j = \frac{1 + \lambda \cdot \sigma\big(\kappa_\gamma(v_j - b_\gamma)\big)}{\alpha_{\min}},
$$
$$
\alpha_j = \alpha_{\min} + (1 - \alpha_{\min}) \cdot \tanh\big(\kappa_\alpha(v_j - b_\alpha)\big),
\tag{10}
$$

where $v_j$ measures the physics difficulty and mainly modulates $\gamma_j$ and $\alpha_j$, $\alpha_{\min}$ sets a lower bound to avoid vanishing gradients, $\tanh(\cdot)$ provides a smooth nonlinear mapping that gradually sharpens the preference decision boundary as the gap increases, and the sigmoid function $\sigma(\cdot)$ ensures a bounded and smooth adjustment of $\gamma_j$, stabilizing the optimization. $\alpha_j$ adaptively adjusts the sharpness of the preference comparison, while $\gamma_j$ amplifies the physics-guided reward for samples with lower physical plausibility. They are controlled by the hyperparameters $\lambda$, $\kappa_{\alpha,\gamma}$, $b_{\alpha,\gamma}$, and $\alpha_{\min}$. Our designed physics rewards $(\alpha_j, \gamma_j)$ balances the training stability with adaptivity, allowing physics-violating samples to exert stronger influence during optimization.

**Flow Matching Probabilistic Modeling.** Our method is based on the standard rectified flow matching as

$$x_t = (1 - t)\, x_0 + t\, x_1, \quad x_1 \sim \mathcal{N}(0, I), \tag{11}$$

(a) A person plays squash on an indoor court.

(b) A person wearing a helmet performs a handspring over a platform.

| VideoDPO | PhyT2V | Wan2.1-14B | OpenAI Sora2 | Google Veo3.1 | **PhyGDPO (Ours)** |

**Figure 5** Comparison on two challenging random user-input actions (squash and handspring). Our method generates more physically plausible videos, capturin racket–ball interactions in squash and well-coordinated body motion in handspring.

which induces the oracle velocity $u^*(x_t, t \mid c) = x_1 - x_0$. Discretizing time with step size $h$ (so $t \mapsto t - h$), the backward update is $x_{t-h} = x_t - h\, u^*(x_t, t \mid c)$. Since the rectified flow one-step reverse transition has no intermediate noise, we follow a vanishing-noise regularization and approximate $p_\theta(x_{t-h} \mid x_t, c)$ and $p_\psi(x_{t-h} \mid x_t, c)$ as

$$p_\theta(x_{t-h} \mid x_t, c) \approx \mathcal{N}\big(x_t - h\, v_\theta(x_t, t, c),\ \varepsilon I\big),$$
$$p_\psi(x_{t-h} \mid x_t, c) \approx \mathcal{N}\big(x_t - h\, v_\psi(x_t, t, c),\ \varepsilon I\big), \tag{12}$$

then let $\varepsilon \to 0$ and use the log-ratio between Gaussians with the identity $(x_{t-h} - x_t)/h = -u^*(x_t, t \mid c)$, we obtain

$$\log \frac{p_\theta(x_{t-h} \mid x_t, c)}{p_\psi(x_{t-h} \mid x_t, c)} \approx -\frac{h}{2\varepsilon}\Big(\ell_\theta(x_t, t) - \ell_\psi(x_t, t)\Big),$$
$$\ell_\theta(x_t, t) = \|v_\theta(x_t, t, c) - u^*(x_t, t \mid c)\|_2^2, \tag{13}$$
$$\ell_\psi(x_t, t) = \|v_\psi(x_t, t, c) - u^*(x_t, t \mid c)\|_2^2,$$

and reformulate the upper bound of Eq. (9) into the final overall training objective as

$$\mathcal{L} = \mathbb{E}_{c, \mathcal{G}^l(c), k, j}\Big[ -\gamma_j \log \sigma\big( -\alpha_j \beta T[(\ell_\theta^w - \ell_\psi^w) - (\ell_\theta^{l_j} - \ell_\psi^{l_j})]\big)\Big], \tag{14}$$

where any global constant scale factor such as $h/2\varepsilon$ in Eq. (13) can be absorbed into the hyperparameter $\beta$. For simplicity, we denote $\ell_\theta^w = \ell_\theta(x_{t_k}^w \mid t_k)$, $\ell_\psi^w = \ell_\psi(x_{t_k}^w \mid t_k)$, $\ell_\theta^{l_j} = \ell_\theta(x_{t_k}^{l_j} \mid t_k)$, and $\ell_\psi^{l_j} = \ell_\psi(x_{t_k}^{l_j} \mid t_k)$ in Eq. (14).

**LoRA-Switch Reference.** Previous DPO methods usually copy the full model and fix it as the reference, which occupies redundant GPU memory and degrades computational efficiency. In addition, this full-copy strategy often leads to unstable and less effective DPO training because this strategy updates the entire set of the backbone weights and may cause the trainable model to quickly and dramatically deviate from the frozen reference model.

To address this issue, we design a LoRA-Switch Reference (LoRA-SR) scheme. As shown in Fig. 3 (b), we freeze the backbone as the reference model $\psi$ and attach trainable LoRA modules to $\psi$ as the trained model $\theta$. Our LoRA-SR enables the reference and trainable models to share the same heavy backbone while flexibly switching the lightweight LoRA parameters for the training and reference modes. Thus, we can evaluate the trainable and reference terms without storing or loading an extra full model. This significantly increases computational efficiency and prevents the model from drifting too far from the reference, thus stabilizing the training process.

| Methods | Hard | Activity | Interaction | Overall |
|---|---|---|---|---|
| Vcrafter2 [12] | 0.0222 | 0.1071 | 0.0943 | 0.1034 |
| Wan2.1-14B [58] | 0.0111 | 0.1190 | 0.1572 | 0.1288 |
| Hunyuan [36] | 0.0222 | 0.1286 | 0.1572 | 0.1356 |
| VideoDPO [39] | 0.0167 | 0.1310 | 0.1572 | 0.1373 |
| PhyT2V [67] | 0.0389 | 0.1405 | 0.1698 | 0.1492 |
| Sora2 [7] | 0.0389 | 0.1429 | 0.1698 | 0.1508 |
| Veo3.1 [20] | 0.0444 | 0.1405 | 0.1887 | 0.1525 |
| PhyGDPO (Ours) | 0.0500 | 0.1571 | 0.1761 | 0.1627 |

**(a)** Quantitative results on VideoPhy2

| Methods | Mechanics | Optics | Thermal | Material | Avg |
|---|---|---|---|---|---|
| Lavie [62] | 0.40 | 0.44 | 0.38 | 0.32 | 0.36 |
| Wan2.1-14B [58] | 0.36 | 0.53 | 0.36 | 0.33 | 0.40 |
| Open-Sora [76] | 0.43 | 0.50 | 0.44 | 0.37 | 0.44 |
| Pika [50] | 0.35 | 0.56 | 0.43 | 0.39 | 0.44 |
| Vchitect-2.0 [21] | 0.41 | 0.56 | 0.44 | 0.37 | 0.45 |
| PhyT2V [67] | 0.45 | 0.55 | 0.43 | 0.53 | 0.50 |
| VideoDPO [39] | 0.48 | 0.60 | 0.47 | 0.58 | 0.54 |
| PhyGDPO (Ours) | 0.55 | 0.60 | 0.58 | 0.47 | 0.55 |

**(b)** Quantitative results on PhyGenBench

| Compared Methods | Preference of Ours |
|---|---|
| Vcrafter2 [12] | 94.2 % |
| VideoDPO [39] | 89.4 % |
| PhyT2V [67] | 88.5 % |
| Wan2.1-14B [58] | 86.5 % |
| Hunyuan [36] | 82.7 % |
| Sora2 [7] | 67.3 % |
| Veo3.1 [20] | 64.4 % |
| PhyGDPO (Ours) | – |

**(c)** User preference of our method

**Table 1** Quantitative comparison and user preference (%) of our method on VideoPhy2 [4] and PhyGenBench [45] datasets.

(a) A wooden pencil is carefully dipped into a glass of crystal-clear water, …



(c) A small burning ball of paper was thrown into a pile of dry paper.



| VideoDPO | PhyT2V | Wan2.1-14B | OpenAI Sora2 | Google Veo3.1 | **PhyGDPO (Ours)** |
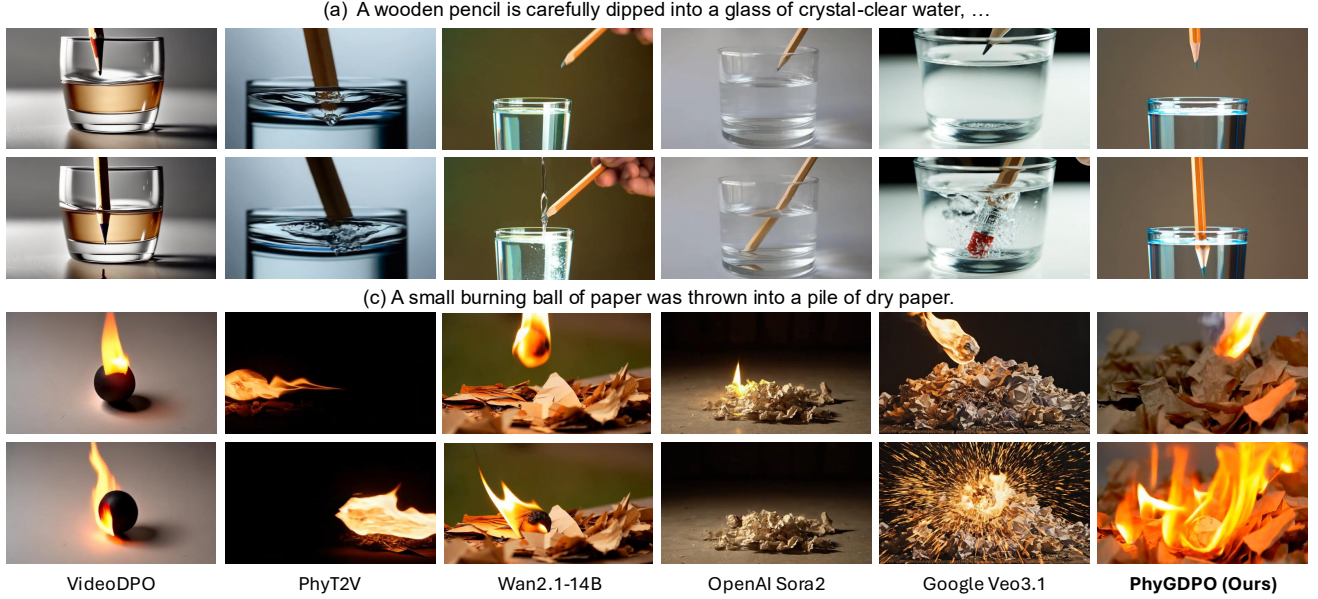
**Figure 6** Results of two challenging physics phenomena (water refraction and paper combustion) on PhyGenBench [45]. Our method can model the magnifying convex-lens effect and the realistic light refraction caused by the curved water surface, as well as the physically correct ignition and flame propagation on paper, demonstrating stronger physics reasoning ability.

# 3 Experiment

**Dataset.** We first use PhyAugPipe to process one million high-quality text-video pairs. After physics data filtering in Fig. 2 (c), we obtain 135K data pairs with sufficient physics interaction and phenomena. Then we perform action clustering and data resampling with physics rewarding to derive 17K data pairs from the 135K filtered data samples. Subsequently, we use the pre-trained T2V model Wan2.1-14B to generate videos for the 17K text prompts with different random seeds and use VideoCon-Physics [3] to score them. We adopt the short unextended prompts from the two datasets, VideoPhy2 [4] and PhyGenBench [45], for evaluation.

**Implementation Details.** We implement our PhyGDPO by pytorch [48] based on a foundation T2V model Wan2.1-14B. Our model is finetuned for 10K steps in total at a batch size of 8 on 8 H100 GPUs for 6 days. To save GPU memory, we adopt mixed-precision training [47] with BF16 and sublinear memory training [15]. We adopt the AdamW optimizer [41] ($\beta_1 = 0.9$, $\beta_2 = 0.999$) with a weight decay of 0.01. The learning rate is initially set as $1e^{-5}$ and decays to $1e^{-6}$ using cosine annealing [40] algorithm. The training and inference spatial resolution of the video is 480×832. We set $\tau = 3$, $N = 100$, $\alpha_{\min} = 0.5$, $k_\gamma = 2.0$, $b_\gamma = 0.4$, $\lambda = 0.6$, $k_\alpha = 5.0$, $b_\alpha = 0.5$. The rank of LoRA is set to 48. We follow the same metrics of VideoPhy2 and PhyGenBench to evaluate the T2V generation models.

## 3.1 Comparison with State-of-the-Art Methods

**Quantitative Results.** We compare PhyGDPO with state-of-the-art (SOTA) methods in Tab. 1a and 1b. PhyGDPO outperforms SOTA methods by large margins. **(i)** Our method surpasses the two recent strongest closed-source models, Sora2 [7] and Veo3 [20] across the tracks of hard actions, activities and sports, and the overall score on VideoPhy2. Especially on the hard actions, PhyGDPO achieves 4.5× higher score than the base model, Wan2.1-14B,

| Method | Hard | Activity | Interaction | Overall |
|---|---|---|---|---|
| Baseline | 0.0333 | 0.1405 | 0.1635 | 0.1475 |
| + Chain-of-Thought | 0.0389 | 0.1452 | 0.1698 | 0.1525 |
| + Action Clustering | 0.0500 | 0.1524 | 0.1698 | 0.1575 |
| + Physics Rewarding | 0.0500 | 0.1571 | 0.1761 | 0.1627 |

**(a)** Break-down ablation of PhyAugPipe on Wan2.1-14B

| Method | Hard | Activity | Interaction | Overall |
|---|---|---|---|---|
| Baseline | 0.0111 | 0.1190 | 0.1572 | 0.1288 |
| + LoRA-SR | 0.0278 | 0.1381 | 0.1635 | 0.1458 |
| + Groupwise Model | 0.0389 | 0.1500 | 0.1698 | 0.1559 |
| + Physics Rewarding | 0.0500 | 0.1571 | 0.1761 | 0.1627 |

**(b)** Break-down ablation study of PhyGDPO on Wan2.1-14B

| Method | Hard | Activity | Interaction | Overall |
|---|---|---|---|---|
| Baseline | 0.0118 | 0.1136 | 0.1447 | 0.1232 |
| Flow-DPO [38] | 0.0296 | 0.1247 | 0.1342 | 0.1283 |
| VideoDPO [39] | 0.0278 | 0.1262 | 0.1509 | 0.1305 |
| PhyGDPO (Ours) | 0.0444 | 0.1357 | 0.1635 | 0.1407 |

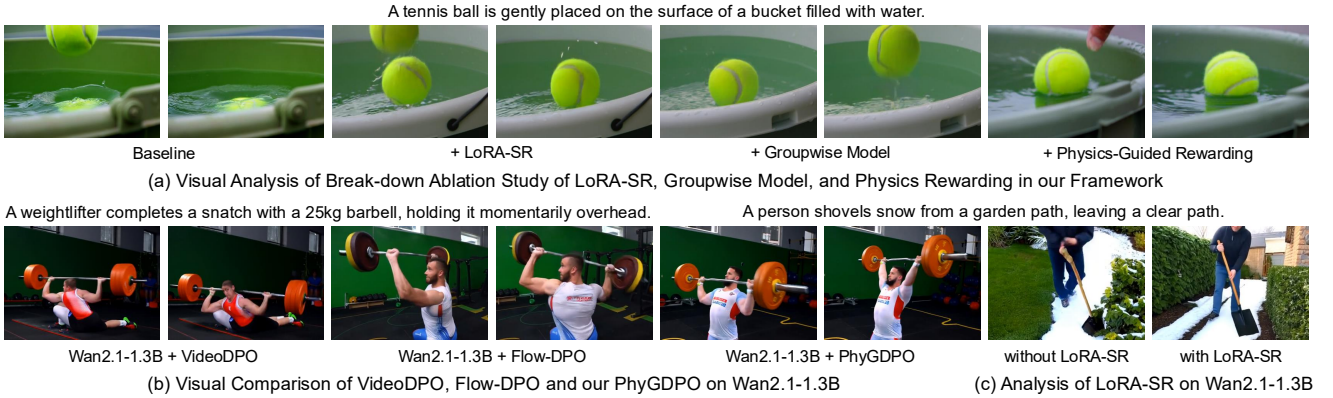**(c)** Comparison of PhyGDPO and VideoDPO on Wan2.1-1.3B

| Method | GPU Memory | Storage Space | Hard Score | Overall Score |
|---|---|---|---|---|
| Baseline | - | - | 0.0118 | 0.1232 |
| LoRA-SFT | 24.7GB | 84MB | 0.0167 | 0.1283 |
| w/o LoRA-SR | 48.7GB | 5.3GB | 0.0389 | 0.1373 |
| with LoRA-SR | 25.3GB | 84MB | 0.0444 | 0.1407 |

**(d)** Ablation of LoRA-SR mechanism on Wan2.1-1.3B

**Table 2** Ablation study on VideoPhy2 [4]. (a) and (b) study the components of PhyAugPipe and PhyGDPO towards better performance. (c) compares VideoDPO with PhyGDPO with the same settings. (d) studies the effect of LoRA-SR scheme.

A tennis ball is gently placed on the surface of a bucket filled with water.



Baseline      + LoRA-SR      + Groupwise Model      + Physics-Guided Rewarding

(a) Visual Analysis of Break-down Ablation Study of LoRA-SR, Groupwise Model, and Physics Rewarding in our Framework

A weightlifter completes a snatch with a 25kg barbell, holding it momentarily overhead.      A person shovels snow from a garden path, leaving a clear path.



Wan2.1-1.3B + VideoDPO      Wan2.1-1.3B + Flow-DPO      Wan2.1-1.3B + PhyGDPO      without LoRA-SR      with LoRA-SR

(b) Visual Comparison of VideoDPO, Flow-DPO and our PhyGDPO on Wan2.1-1.3B      (c) Analysis of LoRA-SR on Wan2.1-1.3B

**Figure 7** Visual analysis. (a) When using LoRA-SR, groupwise probabilistic modeling, and physics-guided rewarding, the generated tennis ball motion becomes more physically realistic, better adhering to the laws of buoyancy. (b) Using our PhyGDPO can generate more physically coherent actions with well-structured human body motion than VideoDPO and Flow-DPO. (c) Applying our LoRA-SR can generate more accurate and plausible hand–shovel contact in snow shoveling.

and is 29% and 13% higher than Sora2 and Veo3. **(ii)** Compared to the SOTA DPO algorithm for video generation (VideoDPO) and SOTA method for physically consistent video generatiion (PhyT2V), our method surpasses VideoDPO and PhyT2V and by 200% and 29% on the hard action score of VideoPhy2. Our method is also 22%/15% and 35%/23% higher than PhyT2V/VideoDPO on the mechanics and thermal tracks of PhyGenBench.

**User Study.** Since the auto-evaluation tools of VideoPhy2 and PhyGenBench are based on VLMs, their assessment of physical plausibility may be inaccurate. Therefore, we conduct a user study with 104 participants. Each participant views the videos generated by our PhyGDPO and a random competing method, along with the randomly selected text prompt from VideoPhy2 and PhyGenBench. The participants are asked to pick up the video that better follows the physics laws. Tab. 1c reports the user preference (%) of our method over competing entries. PhyGDPO is consistently preferred by the human evaluators than all SOTA methods.

**Qualitative Results.** The visual comparisons on challenging physical actions or phenomena including gymnastics, soccer, basketball, glass smashing, polo, squash, handspring, refraction, and combustionare shown in Fig. 1, 4, 5, and 6. Our method generates videos with more realistic physical dynamics, deformation-free body motion, and accurate object interactions, effectively capturing phenomena such as force transfer, material deformation, light refraction, and flame propagation. Compared to existing open- and closed-source models, PhyGDPO shows superior generalization from human activities to complex physical events. Please refer to the project page for video results.

## 3.2 Ablation Study

**Ablation of PhyAugPipe.** We conduct a break-down ablation to study the effect of each component of PhyAugPipe in Tab. 2a. The baseline is directly training PhyGDPO on the randomly selected text-video data. When we keep the number of selected training data the same and apply the data filtering with CoT, action clustering with semantics matching, and data sampling with physics-guided rewarding, all track scores are improved. Especially the score on

the hard actions gains by over 50%. These results suggest the effectiveness of our data construction techniques.

**Ablation of PhyGDPO.** Tab. 2b shows a break-down ablation for PhyGDPO. When progressively using LoRA-SR, PL groupwise probabilistic model, and physics-guided rewarding, the performance steadily gains by large margins. The score on the hard action categories yields $4.5\times$ improvement, demonstrating the effectiveness of our designed methods in improving T2V physics plausibility. In addition, we conduct a visual analysis in Fig. 7 (a) for PhyGDPO. When gradually using our proposed techniques, the tennis in the generated video no longer presents ghosting artifacts, and its floating motion on the water surface better conforms to the physics laws of fluid buoyancy.

**PhyGDPO vs. SOTA DPO.** For fair comparison with two SOTA DPO methods: Flow-DPO [38] and VideoDPO, we use them to finetune Wan2.1-1.3B with the same data and settings. As reported in Tab. 2c, PhyGDPO surpasses Flow-DPO and VideoDPO across all tracks by large margins, especially on the hard action track, where it yields 50% improvement. We conduct a visual comparison in Fig. 7 (b) on the challenging weightlifting action. Flow-DPO and VideoDPO generate distorted or unstable body poses. In contrast, PhyGDPO generates coherent body motion with stable shapes and balanced force dynamics, highlighting its advantage in modeling complex physical actions.

**Analysis of LoRA-SR.** We conduct an ablation study of LoRA-SR in Tab. 2d. In our PhyGDPO training, using LoRA-SR reduces GPU memory consumption by 44% and achieves over $60\times$ compression in storage size, while still improving the score on hard actions by 14% compared to the version without LoRA-SR. We also conduct a visual analysis in Fig. 7 (c), using our LoRA-SR in DPO training can generate more plausible physical interactions between human hands and the snow shovel. Besides, we also compare with SFT using the same LoRA in Tab. 2d. Our PhyGDPO with LoRA-SR uses almost the same GPU memory but surpasses LoRA-SFT by large margins, especially on the hard-action categories, where the score achieves $2.7\times$ improvement.

# 4    Related Work

## 4.1    Text–to–Video Generation Models

T2V generation [5–7, 18, 25–27, 30, 62, 71, 74, 77] has witnessed significant progress. Many existing works follow DiT [49] to adopt a Transformer [56] to predict the noise in diffusion [9, 22, 24, 35, 37, 42–44, 51] or estimate the velocity field in flow matching [8, 14, 19, 23, 29, 31, 34, 58, 59, 73, 75]. Although high visual quality is achieved, accurately modeling the underlying physics-related effects remains challenging. Two recent works, DiffPhy [72] and PhyT2V [67], adopt an LLM to extend the text prompt with explicit physics laws and phenomena, and then iteratively generates the video or finetunes a T2V model with extended prompts. Yet, these prompt extension-based methods are easily misled by the mistakes of LLM and struggle to learn the implicit physics. Different from these works, we aim to explore the physics reasoning potential of the T2V model itself.

## 4.2    Direct Preference Optimization

Direct Preference Optimization (DPO) [52] is proposed to align LLMs with human preferences and has been adopted in image [17, 33, 46, 57] and video [13, 38, 39, 66, 70] generation. For example, DiffusionDPO [57] finetunes a text-to-image diffusion model with human-annotated preference image pairs. VideoDPO [39] adapts DiffusionDPO into T2V generation. However, these DPO frameworks mainly improve the visual quality and aesthetics. In addition, these DPO methods suffer from a severe low-efficiency problem because they need to copy a full model as the reference and they also lack sufficient guidance of physics, which makes the preference signal less reliable and limits the improvement on physical consistency. Our work aims to cope with these research problems.

# 5    Conclusion

In this paper, we focus on studying a challenging problem, physically consistent T2V generation, without using LLM for prompt extension in inference because our goal is to explore the implicit physics reasoning ability of the video generation model. To this end, we first propose a data construction method, PhyAugPipe, to collect a training dataset PhyVidGen-135K containing over 135K text-video data pairs with rich physics interaction and phenomena. Subsequently, we formulate a principled groupwise DPO framework, PhyGDPO, with two technical designs PGR and LoRA-SR. PhyGDPO efficiently post-trains Wan2.1-T2V-14B model on PhyVidGen-135K to boost the physics plausibility in video generation. Comprehenseiv experiments show that our method quantitatively and qualitatively outperforms SOTA algorithms and achieves higher human preference in the user study.

# References

[1] Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, Yin Cui, Yifan Ding, et al. Cosmos world foundation model platform for physical ai. *arXiv preprint arXiv:2501.03575*, 2025.

[2] Hassan Abu Alhaija, Jose Alvarez, Maciej Bala, Tiffany Cai, Tianshi Cao, Liz Cha, Joshua Chen, Mike Chen, Francesco Ferroni, Sanja Fidler, et al. Cosmos-transfer1: Conditional world generation with adaptive multimodal control. *arXiv preprint arXiv:2503.14492*, 2025.

[3] 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. In *ICLR*, 2025.

[4] 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. In *ICMLW*, 2025.

[5] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.

[6] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.

[7] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, et al. Video generation models as world simulators. *OpenAI Blog*, 2024.

[8] Yuanhao Cai, He Zhang, Xi Chen, Jinbo Xing, Yiwei Hu, Yuqian Zhou, Kai Zhang, Zhifei Zhang, Soo Ye Kim, Tianyu Wang, et al. Omnivcus: Feedforward subject-driven video customization with multimodal control conditions. In *NeurIPS*, 2025.

[9] Yuanhao Cai, He Zhang, Kai Zhang, Yixun Liang, Mengwei Ren, Fujun Luan, Qing Liu, Soo Ye Kim, Jianming Zhang, Zhifei Zhang, Yuqian Zhou, Yulun Zhang, Xiaokang Yang, Zhe Lin, and Alan Yuille. Baking gaussian splatting into diffusion denoiser for fast and scalable single-stage image-to-3d generation and reconstruction. In *ICCV*, 2025.

[10] Haoxuan Che, Xuanhua He, Quande Liu, Cheng Jin, and Hao Chen. Gamegen-x: Interactive open-world game video generation. In *ICLR*, 2025.

[11] Haoxin Chen, Menghan Xia, Yingqing He, Yong Zhang, Xiaodong Cun, Shaoshu Yang, Jinbo Xing, Yaofang Liu, Qifeng Chen, Xintao Wang, et al. Videocrafter1: Open diffusion models for high-quality video generation. *arXiv preprint arXiv:2310.19512*, 2023.

[12] Haoxin Chen, Yong Zhang, Xiaodong Cun, Menghan Xia, Xintao Wang, Chao Weng, and Ying Shan. Videocrafter2: Overcoming data limitations for high-quality video diffusion models. In *CVPR*, 2024.

[13] Harold Haodong Chen, Haojian Huang, Qifeng Chen, Harry Yang, and Ser-Nam Lim. Hierarchical fine-grained preference optimization for physically plausible video generation. In *NeurIPS*, 2025.

[14] Shoufa Chen, Chongjian Ge, Yuqi Zhang, Yida Zhang, Fengda Zhu, Hao Yang, Hongxiang Hao, Hui Wu, Zhichao Lai, Yifei Hu, et al. Goku: Flow based video generative foundation models. In *CVPR*, 2025.

[15] Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*, 2016.

[16] Xiaoyu Chen, Junliang Guo, Tianyu He, Chuheng Zhang, Pushi Zhang, Derek Cathera Yang, Li Zhao, and Jiang Bian. Igor: Image-goal representations are the atomic control units for foundation models in embodied ai. *arXiv preprint arXiv:2411.00785*, 2024.

[17] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, Nicu Sebe, and Mubarak Shah. Curriculum direct preference optimization for diffusion and consistency models. In *CVPR*, 2025.

[18] Aram Davtyan, Sepehr Sameni, and Paolo Favaro. Efficient video prediction via sparsely conditioned flow matching. In *ICCV*, 2023.

[19] Aram Davtyan, Sepehr Sameni, and Paolo Favaro. Efficient video prediction via sparsely conditioned flow matching. In *ICCV*, 2023.

[20] Google DeepMind. Veo-3: A scalable and controllable text-to-video model. https://storage.googleapis.com/deepmind-media/veo/Veo-3-Tech-Report.pdf, 2025. Technical Report, Google DeepMind.

[21] Weichen Fan, Chenyang Si, Junhao Song, Zhenyu Yang, Yinan He, Long Zhuo, Ziqi Huang, Ziyue Dong, Jingwen He, Dongwei Pan, et al. Vchitect-2.0: Parallel transformer for scaling up video diffusion models. *arXiv preprint arXiv:2501.08453*, 2025.

[22] Peng Gao, Le Zhuo, Dongyang Liu, Ruoyi Du, Xu Luo, Longtian Qiu, Yuhang Zhang, Chen Lin, Rongjie Huang, Shijie Geng, et al. Lumina-t2x: Transforming text into any modality, resolution, and duration via flow-based large diffusion transformers. *arXiv preprint arXiv:2405.05945*, 2024.

[23] Yu Gao, Haoyuan Guo, Tuyen Hoang, Weilin Huang, Lu Jiang, Fangyuan Kong, Huixia Li, Jiashi Li, Liang Li, Xiaojie Li, et al. Seedance 1.0: Exploring the boundaries of video generation models. *arXiv preprint arXiv:2506.09113*, 2025.

[24] Agrim Gupta, Lijun Yu, Kihyuk Sohn, Xiuye Gu, Meera Hahn, Fei-Fei Li, Irfan Essa, Lu Jiang, and José Lezama. Photorealistic video generation with diffusion models. In *ECCV*, 2024.

[25] Yingqing He, Tianyu Yang, Yong Zhang, Ying Shan, and Qifeng Chen. Latent video diffusion models for high-fidelity long video generation. *arXiv preprint arXiv:2211.13221*, 2022.

[26] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022.

[27] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *NeurIPS*, 2022.

[28] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *ICLR*, 2022.

[29] Yang Jin, Zhicheng Sun, Ningyuan Li, Kun Xu, Hao Jiang, Nan Zhuang, Quzhe Huang, Yang Song, Yadong Mu, and Zhouchen Lin. Pyramidal flow matching for efficient video generative modeling. In *ICLR*, 2025.

[30] Yang Jin, Zhicheng Sun, Ningyuan Li, Kun Xu, Kun Xu, Hao Jiang, Nan Zhuang, Quzhe Huang, Yang Song, Yadong Mu, and Zhouchen Lin. Pyramidal flow matching for efficient video generative modeling. In *ICLR*, 2025.

[31] Xuan Ju, Tianyu Wang, Yuqian Zhou, He Zhang, Qing Liu, Nanxuan Zhao, Zhifei Zhang, Yijun Li, Yuanhao Cai, Shaoteng Liu, et al. Editverse: Unifying image and video editing and generation with in-context learning. *arXiv preprint arXiv:2509.20360*, 2025.

[32] 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. In *ICML*, 2025.

[33] Shyamgopal Karthik, Huseyin Coskun, Zeynep Akata, Sergey Tulyakov, Jian Ren, and Anil Kag. Scalable ranked preference optimization for text-to-image generation. In *ICCV*, 2025.

[34] Lei Ke, Haohang Xu, Xuefei Ning, Yu Li, Jiajun Li, Haoling Li, Yuxuan Lin, Dongsheng Jiang, Yujiu Yang, and Linfeng Zhang. Proreflow: Progressive reflow with decomposed velocity. In *CVPR*, 2025.

[35] Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Grant Schindler, Rachel Hornung, Vighnesh Birodkar, Jimmy Yan, Ming-Chang Chiu, et al. Videopoet: A large language model for zero-shot video generation. In *ICML*, 2024.

[36] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*, 2024.

[37] Xin Li, Wenqing Chu, Ye Wu, Weihang Yuan, Fanglong Liu, Qi Zhang, Fu Li, Haocheng Feng, Errui Ding, and Jingdong Wang. Videogen: A reference-guided latent diffusion approach for high definition text-to-video generation. *arXiv preprint arXiv:2309.00398*, 2023.

[38] Jie Liu, Gongye Liu, Jiajun Liang, Ziyang Yuan, Xiaokun Liu, Mingwu Zheng, Xiele Wu, Qiulin Wang, Wenyu Qin, Menghan Xia, et al. Improving video generation with human feedback. *arXiv preprint arXiv:2501.13918*, 2025.

[39] Runtao Liu, Haoyu Wu, Ziqiang Zheng, Chen Wei, Yingqing He, Renjie Pi, and Qifeng Chen. Videodpo: Omni-preference alignment for video diffusion generation. In *CVPR*, 2025.

[40] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *ICLR*, 2017.

[41] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.

[42] Haoyu Lu, Guoxing Yang, Nanyi Fei, Yuqi Huo, Zhiwu Lu, Ping Luo, and Mingyu Ding. Vdt: General-purpose video diffusion transformers via mask modeling. In *ICLR*, 2024.

[43] Xin Ma, Yaohui Wang, Gengyun Jia, Xinyuan Chen, Ziwei Liu, Yuan-Fang Li, Cunjian Chen, and Yu Qiao. Latte: Latent diffusion transformer for video generation. *TMLR*, 2024.

[44] Willi Menapace, Aliaksandr Siarohin, Ivan Skorokhodov, Ekaterina Deyneka, Tsai-Shien Chen, Anil Kag, Yuwei Fang, Aleksei Stoliar, Elisa Ricci, Jian Ren, et al. Snap video: Scaled spatiotemporal transformers for text-to-video synthesis. In *CVPR*, 2024.

[45] 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. In *ICML*, 2025.

[46] Sanghyeon Na, Yonggyu Kim, and Hyunjoon Lee. Boost your human image generation model via direct preference optimization. In *CVPR*, 2025.

[47] Sharan Narang, Gregory Diamos, Erich Elsen, Paulius Micikevicius, Jonah Alben, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, et al. Mixed precision training. In *ICLR*, 2018.

[48] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.

[49] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *ICCV*, 2023.

[50] Pika. Pika art. https://pika.art/, 2025. Accessed: 2025-05-04.

[51] Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, et al. Movie gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720*, 2024.

[52] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *NeurIPS*, 2023.

[53] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *EMNLP*, 2019.

[54] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

[55] Dani Valevski, Yaniv Leviathan, Moab Arar, and Shlomi Fruchter. Diffusion models are real-time game engines. In *ICLR*, 2025.

[56] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, 2017.

[57] Bram Wallace, Meihua Dang, Rafael Rafailov, Linqi Zhou, Aaron Lou, Senthil Purushwalkam, Stefano Ermon, Caiming Xiong, Shafiq Joty, and Nikhil Naik. Diffusion model alignment using direct preference optimization. In *CVPR*, 2024.

[58] Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.

[59] Duomin Wang, Wei Zuo, Aojie Li, Ling-Hao Chen, Xinyao Liao, Deyu Zhou, Zixin Yin, Xili Dai, Daxin Jiang, and Gang Yu. Universe-1: Unified audio-video generation via stitching of experts. *arXiv preprint arXiv:2509.06155*, 2025.

[60] Qinghe Wang, Yawen Luo, Xiaoyu Shi, Xu Jia, Huchuan Lu, Tianfan Xue, Xintao Wang, Pengfei Wan, Di Zhang, and Kun Gai. Cinemaster: A 3d-aware and controllable framework for cinematic text-to-video generation. In *SIGGRAPH*, 2025.

[61] Xiaofeng Wang, Zheng Zhu, Guan Huang, Xinze Chen, Jiagang Zhu, and Jiwen Lu. Drivedreamer: Towards real-world-drive world models for autonomous driving. In *ECCV*, 2024.

[62] Yaohui Wang, Xinyuan Chen, Xin Ma, Shangchen Zhou, Ziqi Huang, Yi Wang, Ceyuan Yang, Yinan He, Jiashuo Yu, Peiqing Yang, et al. Lavie: High-quality video generation with cascaded latent diffusion models. *IJCV*, 2025.

[63] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*, 2022.

[64] Yuqing Wen, Yucheng Zhao, Yingfei Liu, Fan Jia, Yanhui Wang, Chong Luo, Chi Zhang, Tiancai Wang, Xiaoyan Sun, and Xiangyu Zhang. Panacea: Panoramic and controllable video generation for autonomous driving. In *CVPR*, 2024.

[65] Weijia Wu, Mingyu Liu, Zeyu Zhu, Xi Xia, Haoen Feng, Wen Wang, Kevin Qinghong Lin, Chunhua Shen, and Mike Zheng Shou. Moviebench: A hierarchical movie level dataset for long video generation. In *CVPR*, 2025.

[66] Ziyi Wu, Anil Kag, Ivan Skorokhodov, Willi Menapace, Ashkan Mirzaei, Igor Gilitschenski, Sergey Tulyakov, and Aliaksandr Siarohin. Densedpo: Fine-grained temporal preference optimization for video diffusion models. *arXiv preprint arXiv:2506.03517*, 2025.

[67] Qiyao Xue, Xiangyu Yin, Boyuan Yang, and Wei Gao. Phyt2v: Llm-guided iterative self-refinement for physics-grounded text-to-video generation. In *CVPR*, 2025.

[68] An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zhihao Fan. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.

[69] Jiazhi Yang, Shenyuan Gao, Yihang Qiu, Li Chen, Tianyu Li, Bo Dai, Kashyap Chitta, Penghao Wu, Jia Zeng, Ping Luo, et al. Generalized predictive model for autonomous driving. In *CVPR*, 2024.

[70] Xiaomeng Yang, Zhiyu Tan, and Hao Li. Ipo: Iterative preference optimization for text-to-video generation. *arXiv preprint arXiv:2502.02088*, 2025.

[71] Yongsheng Yu, Ziyun Zeng, Haitian Zheng, and Jiebo Luo. Omnipaint: Mastering object-oriented editing via disentangled insertion-removal inpainting. In *ICCV*, 2025.

[72] 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.

[73] Shilong Zhang, Wenbo Li, Shoufa Chen, Chongjian Ge, Peize Sun, Yida Zhang, Yi Jiang, Zehuan Yuan, Binyue Peng, and Ping Luo. Flashvideo: Flowing fidelity to detail for efficient high-resolution video generation. *arXiv preprint arXiv:2502.05179*, 2025.

[74] Shiwei Zhang, Jiayu Wang, Yingya Zhang, Kang Zhao, Hangjie Yuan, Zhiwu Qin, Xiang Wang, Deli Zhao, and Jingren Zhou. I2vgen-xl: High-quality image-to-video synthesis via cascaded diffusion models. *arXiv preprint arXiv:2311.04145*, 2023.

[75] Yifu Zhang, Hao Yang, Yuqi Zhang, Yifei Hu, Fengda Zhu, Chuang Lin, Xiaofeng Mei, Yi Jiang, Bingyue Peng, and Zehuan Yuan. Waver: Wave your way to lifelike video generation. *arXiv preprint arXiv:2508.15761*, 2025.

[76] 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.

[77] Daquan Zhou, Weimin Wang, Hanshu Yan, Weiwei Lv, Yizhe Zhu, and Jiashi Feng. Magicvideo: Efficient video generation with latent diffusion models. *arXiv preprint arXiv:2211.11018*, 2022.

# Appendix

## Mathematical proof for Ineq. (8)

In the main paper, we use an inequality (8) as

$$\sum_{j=1}^{m} e^{x_j} \leq \prod_{j=1}^{m} \left(1 + e^{\alpha_j x_j}\right)^{\gamma_j}, \quad 0 < \alpha_j \leq 1, \ \gamma_j \geq 1/\alpha_j. \tag{15}$$

Here we prove it. For each $j$, since $0 < \alpha_j \leq 1$, the function $f_j(t) = t^{\alpha_j}$ is concave on $[0, +\infty)$. For any $u, v \geq 0$, let $s = u + v$. If $s = 0$ the inequality is trivial. Assume $s > 0$ and set $\lambda = \frac{u}{s} \in [0, 1]$, so that $u = \lambda s$ and $v = (1 - \lambda)s$. By the concavity of $f_j$, we can derive the following inequality:

$$\begin{aligned} f_j(u) &= f_j(\lambda s + (1 - \lambda) \cdot 0) \\ &\geq \lambda f_j(s) + (1 - \lambda) f_j(0) = \lambda f_j(s), \end{aligned} \tag{16}$$

and similarly, we derive the inequality for $f_j(v)$ as

$$f_j(v) = f_j\big((1 - \lambda)s + \lambda \cdot 0\big) \geq (1 - \lambda) f_j(s). \tag{17}$$

Adding the two inequalities Ineq. (16) and Ineq. (17) yields

$$\begin{aligned} u^{\alpha_j} + v^{\alpha_j} &= f_j(u) + f_j(v) \\ &\geq \lambda f_j(s) + (1 - \lambda) f_j(s) \\ &= f_j(s) = (u + v)^{\alpha_j}, \end{aligned} \tag{18}$$

this is,

$$(u + v)^{\alpha_j} \leq u^{\alpha_j} + v^{\alpha_j}. \tag{19}$$

Applying Ineq. (19) to $u = 1$ and $v = e^{x_j}$ yields

$$\left(1 + e^{x_j}\right)^{\alpha_j} \leq 1 + e^{\alpha_j x_j}, \tag{20}$$

which is equivalent to

$$1 + e^{x_j} \leq \left(1 + e^{\alpha_j x_j}\right)^{1/\alpha_j}. \tag{21}$$

Since $\gamma_j \geq 1/\alpha_j$ and the base $\left(1 + e^{\alpha_j x_j}\right) \geq 1$, we have

$$\left(1 + e^{\alpha_j x_j}\right)^{1/\alpha_j} \leq \left(1 + e^{\alpha_j x_j}\right)^{\gamma_j}, \tag{22}$$

and thus from Ineq. (21), we can easily derive:

$$1 + e^{x_j} \leq \left(1 + e^{\alpha_j x_j}\right)^{\gamma_j} \implies e^{x_j} \leq \left(1 + e^{\alpha_j x_j}\right)^{\gamma_j} - 1. \tag{23}$$

We define $a_j \triangleq \left(1 + e^{\alpha_j x_j}\right)^{\gamma_j} - 1 \geq 0$. Then Ineq. (23) gives $e^{x_j} \leq a_j$ for all $j$, and hence

$$\sum_{j=1}^{m} e^{x_j} \leq \sum_{j=1}^{m} a_j. \tag{24}$$

On the other hand, for any nonnegative $\{a_j\}_{j=1}^{m}$ we have

$$\prod_{j=1}^{m} (1 + a_j) = 1 + \sum_{j=1}^{m} a_j + (\text{higher-order terms in } \{a_j\}), \tag{25}$$

where all higher-order terms are nonnegative, so

$$\sum_{j=1}^{m} a_j \leq \prod_{j=1}^{m} (1 + a_j). \tag{26}$$

Combining Ineq. (24) and Ineq. (26) and substituting the definition of $a_j$ yields

$$\sum_{j=1}^{m} e^{x_j} \le \sum_{j=1}^{m} a_j \le \prod_{j=1}^{m} (1 + a_j)$$

$$= \prod_{j=1}^{m} \left(1 + \left(1 + e^{\alpha_j x_j}\right)^{\gamma_j} - 1\right) \tag{27}$$

$$= \prod_{j=1}^{m} \left(1 + e^{\alpha_j x_j}\right)^{\gamma_j},$$

which is exactly the claimed Ineq (15) used in the main paper.