

# Omni-Effects: Unified and Spatially-Controllable Visual Effects Generation

Fangyuan Mao<sup>1\*</sup>, Aiming Hao<sup>1\*</sup>, Jintao Chen<sup>1,2</sup>, Dongxia Liu<sup>1,3</sup>, Xiaokun Feng<sup>1,4</sup>, Jiashu Zhu<sup>1</sup>,  
Meiqi Wu<sup>1,4</sup>, Chubin Chen<sup>1,3</sup>, Jiahong Wu<sup>1†</sup>, Xiangxiang Chu<sup>1</sup>

<sup>1</sup>AMAP, Alibaba Group

<sup>2</sup>Peking University

<sup>3</sup>Tsinghua University

<sup>4</sup>Institute of Automation, Chinese Academy of Sciences

fangyuanmaocs@gmail.com, {haoaiming.ham, hongxi.wjh}@alibaba.inc.com

## Abstract

Visual effects (VFX) are essential visual enhancements fundamental to modern cinematic production. Although video generation models offer cost-efficient solutions for VFX production, current methods are constrained by per-effect LoRA training, which limits generation to single effects. This fundamental limitation impedes applications that require spatially controllable composite effects, i.e., the concurrent generation of multiple effects at designated locations. However, integrating diverse effects into a unified framework faces major challenges: interference from effect variations and spatial uncontrollability during multi-VFX joint training. To tackle these challenges, we propose *Omni-Effects*, a first unified framework capable of generating prompt-guided effects and spatially controllable composite effects. The core of our framework comprises two key innovations: (1) **LoRA-based Mixture of Experts (LoRA-MoE)**, which employs a group of expert LoRAs, integrating diverse effects within a unified model while effectively mitigating cross-task interference. (2) **Spatial-Aware Prompt (SAP)** incorporates spatial mask information into the text token, enabling precise spatial control. Furthermore, we introduce an Independent-Information Flow (IIF) module integrated within the SAP, isolating the control signals corresponding to individual effects to prevent any unwanted blending. To facilitate this research, we construct a comprehensive VFX dataset *Omni-VFX* via a novel data collection pipeline combining image editing and First-Last Frame-to-Video (FLF2V) synthesis, and introduce a dedicated VFX evaluation framework for validating model performance. Extensive experiments demonstrate that *Omni-Effects* achieves precise spatial control and diverse effect generation, enabling users to specify both the category and location of desired effects.

**Code** — <https://amap-ml.github.io/Omni-Effects.github.io>

**Extended version** — <https://arxiv.org/abs/2508.07981>

## Introduction

Visual effects (VFX) play a crucial role in modern filmmaking, enabling the creation of immersive narratives and fantastical worlds. While traditional VFX pipelines, especially

\*These authors contributed equally.

†Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 1: Capabilities for diverse customized visual effects. *Omni-Effects* supports both (a) single-VFX and (b) multi-VFX generation through pure prompt-guided generation. Integrated with the **Spatial-Aware Prompt**, *Omni-Effects* enables (c) precise spatial VFX control and (d) intricate object-based visual effects with targeted environmental transformations.

for composite effects requiring simultaneous coordination across different spatial locations, are notoriously complex and resource-intensive (Chabanova 2022). Recent rapid advances in video generation technologies (Yang et al. 2024; Bao et al. 2024; Kong et al. 2024; Wan, Wang et al. 2025) are driving a paradigm shift in VFX creation—transitioning from conventional methods to generative model-powered dynamic, efficient synthesis.

The inherent scarcity of VFX data and pronounced variability in dynamic characteristics across effects pose significant challenges to training generative models. Consequently, current methods (Liu et al. 2025b) focus on single-effect generation, employing dedicated Low-Rank Adaptation (LoRA) (Hu et al. 2022) tailored to individual effects. However, this paradigm struggles with multi-VFX scenes, exhibiting two critical limitations: **Limitation 1**. Cross-Adapter Interference where joint multi-LoRA activation (Ding et al. 2023) induces spatial occlusion artifacts (Figure 2 (a)) and shared-subspace hybrid training triggers fidelity-degrading cross-effect confusion via task interfer-

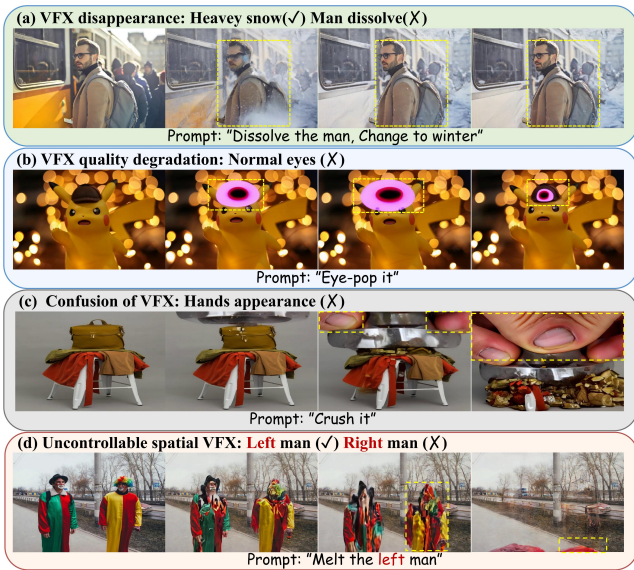


Figure 2: Defects in standard video generation models. (a) VFX disappearance, (b) quality degradation, (c) confusion between VFX elements, and (d) spatial uncontrollability.

ence (Zhang et al. 2025a) (Figure 2 (b, c)). **Limitation 2.** Spatial-Semantic Misalignment wherein the text-pixel space gap prevents precise spatial cue encoding for VFX placement (Figure 2 (d)). These limitations fundamentally constrain conventional video generation adaptation to complex multiple VFX compositions (multi-VFX).

To address these limitations, we propose *Omni-Effects*, a unified framework modeling multi-VFX generation as a multi-condition video generation problem, where textual prompts specify effect categories while spatial masks define their precise locations. First, to tackle Limitation 1, we introduce a **LoRA-based Mixture of Experts (LoRA-MoE)** module (Shazeer et al. 2017; Dou et al. 2023; Zhang et al. 2025c), which partitions effects into specialized subspaces, each optimized by dedicated expert branches, with a gating router dynamically activating relevant subspaces to minimize cross-task interference and enhance effect fidelity. Second, to overcome Limitation 2, we present the **Spatial-Aware Prompt (SAP)**, which integrates explicit mask-based spatial conditioning with textual inputs via full-attention mechanisms for accurate effect placement. To mitigate SAP cross-interference during concurrent application, we introduce the Independent-Information Flow (IIF) module, which isolates condition-specific information flows through IIF Attention Mask, preventing unintended effect blending. Collectively, these innovations enable *Omni-Effects*—to our knowledge, the first VFX framework—to achieve high-fidelity multi-VFX compositions with pixel-level spatial control. (Figure 1).

To advance this research, we construct a high-quality VFX dataset *Omni-VFX* and develop a specialized VFX data pipeline. This pipeline utilises image editing models (Liu et al. 2025a) to generate source image pairs depicting ini-

tial/final effect states, which are then synthesised into VFX videos via the FLV2V framework (built on Wan2.1 (Wan, Wang et al. 2025)). Rigorous manual filtering ensures quality while expanding coverage to 55 distinct effect categories. Further, we introduce an evaluation framework specifically designed for controllable VFX generation tasks. Comprehensive experiments validate the *Omni-Effects* framework’s superior performance across three core capabilities: single-VFX, multi-VFX, and controllable VFX generation. In summary, the major contributions of our work are as follows:

- A unified VFX framework, *Omni-Effects*, is proposed to enable high-fidelity spatial controllable multi-VFX generation through a dual-core architecture: (1) **LoRA-MoE** modules unifying multi-VFX training, and (2) IIF-augmented **SAP** mechanisms enabling independent multi-condition control without interference.
- The most comprehensive VFX dataset *Omni-VFX* is developed with an automated production pipeline to support the generation of diverse high-quality VFX video, complemented by a comprehensive evaluation framework for rigorous controllable VFX assessment.
- Extensive experiments demonstrate that our *Omni-Effects* achieves precise spatial control and enables diverse VFX generation, thereby allowing users to specify both the category and location of desired effects.

## Related Works

### Video Generation Models

Recent advances in diffusion-based video generation (Chen et al. 2023; He et al. 2022; Zeng et al. 2024; Kong et al. 2024; Yang et al. 2024; HaCohen et al. 2024; Bao et al. 2024; Polyak et al. 2024; Wan, Wang et al. 2025; Teng et al. 2025; Seawead et al. 2025; Ling et al. 2025; Chen et al. 2025) have enabled text-to-video (T2V) and image-to-video (I2V) synthesis, where input images establish spatial context. As a critical I2V application, VFX generation creates unrealizable fantastical visuals. However, VFX data scarcity forces reliance on Low-Rank Adaptation (LoRA) (Hu et al. 2022) for limited-data fine-tuning (Liu et al. 2025b). This necessitates separate LoRA models per effect, while combined training causes performance degradation, fundamentally limiting multi-VFX generation within a single video.

*Our architecture unifies multi-VFX training, avoids degradation, and enables concurrent multi-VFX generation.*

### Conditional Video Generation

Condition-guided diffusion models leverage auxiliary inputs for precise output control, falling into two paradigms: spatial-fusion guidance and high-level semantic guidance. Spatial-fusion methods, exemplified by ControlNet (Zhang, Rao, and Agrawala 2023), integrate condition inputs with denoising inputs. These methods (Li et al. 2024; Bai et al. 2024; Bian et al. 2025; Xu et al. 2025; Lei et al. 2025; Jiang et al. 2025) enable fine-grained spatial alignment while preserving the generation quality of pre-trained diffusion models. High-level semantic methods exploit latent interactions between conditions and the denoising process. Techniques

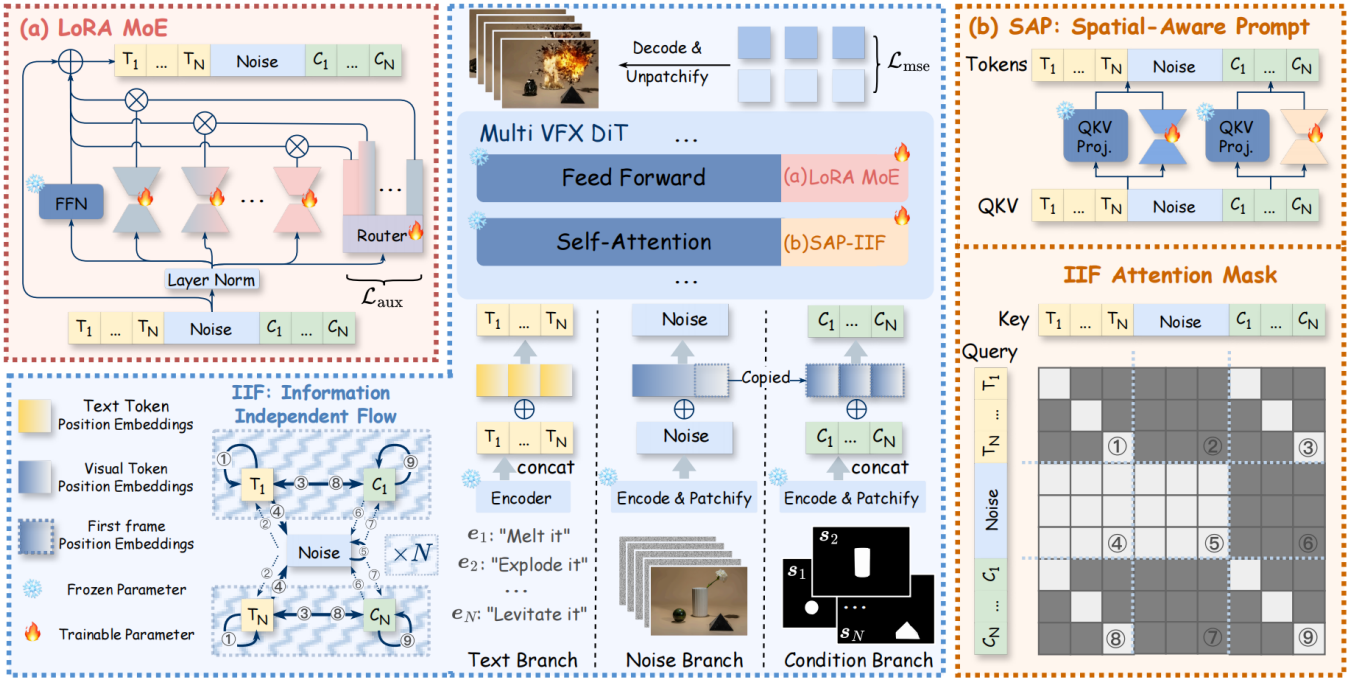


Figure 3: Flowchart of proposed *Omni-Effects*. Given a reference image and composite conditions of arbitrary length, *Omni-Effects* first encodes each input into corresponding tokens. These tokens are concatenated and processed sequentially through downstream DiT blocks. These blocks incorporate two key technologies: (a) **LoRA-MoE**, a MoE plugin replacing standard FFN linear layers to enable collaborative expert task-solving and (b) **SAP**, which fuses effect descriptors with spatial trigger information during the attention stage while mitigating cross-condition information leakage via an IIF mechanism. Note that, in the IIF, dashed lines represent blocked information flow, while solid lines indicate active information transmission.

include cross-attention based mechanisms (Ye et al. 2023; Zhang et al. 2024; Zhou et al. 2025; Yuan et al. 2025) and conditional token concatenation strategies (Wang et al. 2024; Huang et al. 2024a; Tan et al. 2025a,b; Zhang et al. 2025b), dynamically modulate generation through semantic embeddings. Crucially, while these methods effectively handle individual conditions, they struggle to simultaneously and independently control multiple conditions, a critical requirement for professional VFX generation.

*Our model employs IIF-powered SAP control mechanisms to support independent, non-interfering control of multiple conditions within the same video.*

## Method

### Preliminaries

**Video Diffusion Models** Video generation models usually utilize the diffusion paradigm (Ho, Jain, and Abbeel 2020; Lipman et al. 2022), which generates samples from a data distribution  $p(\mathbf{x}_0)$  by progressively denoising samples that are initially drawn from a Gaussian distribution  $p(\mathbf{x}_T)$ . During training, clean samples  $\mathbf{x}_0 \sim p(\mathbf{x}_0)$  undergo iterative corruption through  $T$  diffusion steps:

$$\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad t = 1, \dots, T \quad (1)$$

where  $\alpha_t, \sigma_t > 0$  are scalars that jointly define the Signal-to-Noise Ratio (SNR). The denoiser with parameter  $\theta$  is optimized to predict the target noise  $\epsilon$ . The optimization process

is defined as:

$$\mathcal{L}_{\text{mse}}(\theta) = \mathbb{E}_{\mathbf{x}_t, t, c, \epsilon} \left[ \|\epsilon - \epsilon_\theta(\mathbf{x}_t, t, \tau(c))\|_2^2 \right] \quad (2)$$

where  $c$  is conditions (e.g., text and spatial location), and  $\tau$  denotes the condition encoder. By replacing target noise  $\epsilon$  with  $v$ , which is a weighted combination of  $\mathbf{x}_0$  and  $\epsilon$ , as the prediction target, the  $v$ -prediction (Salimans and Ho 2022) is derived, which is adopted in our *Omni-Effects* framework. Moreover, mainstream video generation leverages Diffusion Transformer (DiT) (Peebles and Xie 2023) architecture by employing attention mechanisms to model spatiotemporal consistency while aligning conditional inputs with visual outputs (Zheng et al. 2024). By integrating diffusion processes with Transformer architectures, the video generation performance is improved, leading to high-quality and accurate video synthesis results.

**Spatially Controllable Multi-VFX Generation** In practical applications, it's often necessary to display different VFX at distinct locations throughout a video. We formalize this task as multi-conditional video generation, wherein video diffusion models take a reference image and a set of  $N$  control signals  $\mathcal{C} = \{c_i\}_{i=1}^N$  as inputs. Each condition  $c_i = (e_i, s_i)$  couples an effect descriptor  $e_i$  with a spatial trigger  $s_i$ , whereby the generated video  $\mathbf{x}_0$  applies effect  $e_i$  at the location specified by  $s_i$ . Specifically, we use a text prompt to describe the VFX, while using a spatial mask




Figure 4: FVD scores for diverse VFX trained with a shared LoRA. VFX performance exhibits an initial improvement followed by progressive degradation with increasing numbers of co-trained effects. This indicates inherent effect clustering: synergistic groups (e.g., Melt-like effects) improve co-training performance, while incompatible combinations (e.g., Deflate + Squish) suffer from mode collapse and underperform relative to compatible sets. Note that, **lower FVD values indicate superior performance**, with optimal VFX results uniformly achieved when the number of co-trained VFX equals 4.

$m \in \mathbb{R}^{H \times W}$  to serve as the spatial trigger. A set of conditions is incorporated into the denoising process, and the denoiser prediction becomes:

$$\hat{v} = \epsilon_{\theta} \left( \mathbf{x}_t, t, \left\{ \tau_e^{(i)}(e_i) \right\}_{i=1}^N, \left\{ \tau_s^{(i)}(s_i) \right\}_{i=1}^N \right), \quad (3)$$

where  $\tau_e$  and  $\tau_s$  denote the text and spatial mask encoder, respectively. Notice that, when  $N = 1$  and the spatial trigger is empty, the above task reduces to the traditional single-VFX generation task (Liu et al. 2025b).

### Omni-Effects

To model the above task, video diffusion models require simultaneous support for multi-VFX inference and spatial control capabilities. We accordingly propose *Omni-Effects*, building upon the CogVideoX (Yang et al. 2024) architecture and incorporating two core components: **LoRA-MoE** and **Spatial-Aware Prompt**. The overview is illustrated in Figure 3, and the details are as follows.

**LoRA-MoE** As mentioned in Figure 2, both multi-LoRA parallel inference and single-LoRA unified training degrade performance. Crucially, we observe the synergistic mechanism in VFX training: *compatible VFX-combination training enhances single-VFX generation quality* (Figure 4). This discovery motivates our adaptive task-space partitioning strategy: inspired by MoE (Shazeer et al. 2017) architectures, we partition distinct effects into specialized subspaces and deploy a gating router for adaptive subspace selection.

Specifically, LoRA-MoE (Dou et al. 2023) integrates MoE with LoRA (Figure 3 (a)), which employs an expert ensemble where each LoRA specializes in distinct VFX manifolds. Formally, for input token  $\mathbf{x} \in \mathbb{R}^d$ , a weight is obtained by a gating network  $G: \mathbb{R}^d \mapsto \mathbb{R}^n$  for each expert, resulting in  $G(\mathbf{x}) = [G(\mathbf{x})_1, G(\mathbf{x})_2, \dots, G(\mathbf{x})_n]$ , where  $n$  represents the number of experts. Each expert  $E_i$  implements

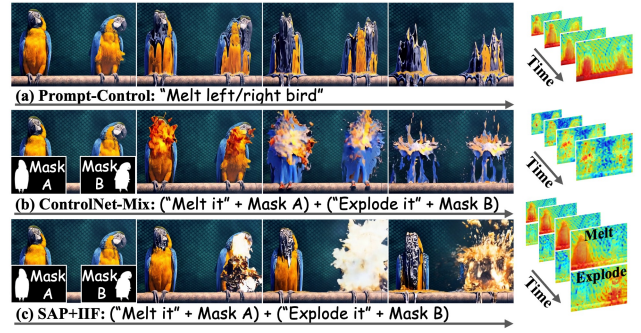


Figure 5: Visualization of controllable VFX performance and attention maps. (a) Position description lacks spatial control; (b) ControlNet faces inter-condition interference, leading to VFX leakage and artifacts; (c) Proposed SAP+IIF achieves precise positional controllability while preventing mutual interference between multi-VFX.

LoRA decomposition:

$$E_i(\mathbf{x}) = \frac{\alpha}{r} \mathbf{x} \mathbf{A}_i \mathbf{B}_i, \quad \mathbf{A}_i \in \mathbb{R}^{d \times r}, \quad \mathbf{B}_i \in \mathbb{R}^{r \times d}, \quad (4)$$

where  $r$  denotes the low-rank and  $\alpha$  is a scaling factor. The final prediction combines base model and expert outputs:

$$\mathbf{y} = \text{Base}(\mathbf{x}) + \sum_{i=1}^n G(\mathbf{x})_i \odot E_i(\mathbf{x}), \quad (5)$$

This MoE-structured plugin replaces standard FFN linear layers, enabling collaborative expert task-solving. During training, Top- $k$  routing strategy ( $k \leq n$ ) is adopted to enforce exactly  $k$  non-zero entries in  $G(\mathbf{x})$ . At inference, all experts are activated to prevent effect suppression caused by Top- $k$  filtering, omitting critical experts, which is essential for multi-VFX combination generation. Moreover, to mitigate workload imbalance caused by the gating network favoring a few experts during training, we also employ a balanced routing auxiliary loss  $\mathcal{L}_{\text{aux}}$  (Fedus, Zoph, and Shazeer 2022). Comprehensive details are provided in Supplement A. The final training objective is expressed as:  $\mathcal{L} = \mathcal{L}_{\text{mse}} + \beta \mathcal{L}_{\text{aux}}$ , where  $\beta$  is hyperparameter.

**Spatial-Aware Prompt** For condition  $c_i = (e_i, s_i)$ , embedding positional descriptors within text prompts proves insufficient for precise spatial control. To investigate this phenomenon, we visualize attention maps across diverse prompts. Crucially, *attention consistently activates identical regions regardless of prompt semantics* (Figure 5 (a)), evidencing textual position cues’ failure to direct activation toward specified targets. Prior work (Liu et al. 2025b; Jiang et al. 2025) mitigates this via ControlNet (Zhang, Rao, and Agrawala 2023) to extract a mask sequence for generation guidance. However, this solution suffers from two critical limitations:

1. **Significant parameter overhead:** ControlNet duplicates a portion of the base model’s parameters (typically half), requiring substantial extra trainable weights;

2. **Severe cross-condition interference:** During multi-VFX generation, parallel ControlNet inference suffers from information leakage, manifesting as erroneous co-occurrence of effect  $e_i$  and  $e_j$  at positions  $s_i$  and  $s_j$  respectively (Figure 5 (b)).

In summary, signals within composite conditions must be integrated while preventing cross-condition interference to ensure robust performance. We address these challenges by proposing the **Spatial-Aware Prompt** to directly inject spatial information into prompts tokens via enhanced spatial-text condition token interactions within attention mechanism, enabling controllable generation with minimal parameter/computational overhead. Building on this, we introduce Information-Independent Flow, which utilizes a designed attention mask to restrict cross-condition information exchange, thereby preventing interference between distinct control streams. Formally, given a set of conditions  $C$ , encoder processing yields text condition tokens  $\{\tau_e^{(i)}(e_i)\}_{i=1}^N$  and spatial condition tokens  $\{\tau_s^{(i)}(s_i)\}_{i=1}^N$ , which are sequentially concatenated with the noisy latent  $x_t$  to form the inputs  $Q$ ,  $K$  and  $V$ . Then we define an attention mask  $M \in \{0, -\infty\}^{l \times l}$  ( $l$  is the total sequence length) to regulate attention flow (Figure 3 (b), details are in Supplement A) that blocks condition-to-condition and noise-to-condition interactions, eliminating cross-condition leakage to prevent effect misalignment or blending. The final output of attention is expressed as:

$$y = \text{Softmax}\left(QK^T/\sqrt{d_k} + M\right)V, \quad (6)$$

where  $d_k$  denotes the feature dimension. To enhance spatial conditioning alignment with noisy latents, we inject positional embeddings from  $x_t$ 's initial frame into  $\tau_s^{(i)}(s_i)$ , coupled with a dedicated Spatial-Condition LoRA. Crucially, all spatial conditions share identical LoRA parameters while maintaining a common base LoRA across other branches, ensuring efficient conditional injection without disrupting pretrained representations. Each text condition is individually processed through text encoder while sharing identical positional encoding. As shown in Figure 5 (c), our SAP+IIF achieves precise VFX targeting in target regions with non-overlapping activation zones.

## Data and Training

### Dataset Collection

VFX fundamentally manifest as radical spatio-temporal state transformations (e.g., explosion). Despite modern techniques like animation and Computer Graphics Interface (CGI) (Chabanova 2022), modeling such dynamics remains challenging. We introduce a novel pipeline: for any input image, Step1X-Edit (Liu et al. 2025a) produces its modified counterpart to establish boundary frames defining a VFX's initial and terminal states. This constraint provides strong transformation priors for generative models. The FLF2V framework (Wan, Wang et al. 2025) then synthesizes the final video by compressing VFX production into boundary-

constrained state-transition path search, significantly reducing modeling complexity. Through curated manual selection, we build a comprehensive dataset *Omni-VFX* spanning **55 distinct VFX** across **instantaneous environmental shifts, artistic styles, human emotions, and so on**, enabling diverse creative applications. For more data details, please refer to Supplement B.

### Training

Since our training dataset contains only single-VFX without multi-VFX data, empirical observations reveal that standard training fails to achieve controllable multi-VFX generation. We overcome this with a tri-level solution. At the **data** level, through random cropping and splicing with two videos, and random temporal freezing, we generate pseudo multi-VFX videos with corresponding masks. At the **scheduler** level, Non-Uniform Sampling prioritizes denoising steps  $\in [900, 1000]$ (early stage) for spatial control learning with increased batch allocation, while dedicating fewer batches to detail refinement in lower steps  $\in [0, 900]$ , motivated by empirical findings that enhanced focus on early denoising accelerates model convergence. At the **training strategy** level, iterative single to multi-VFX ( $N = 2$ ) fine-tuning ensures stable convergence and performance gains. For more training details, please refer to Supplement C.

## Experiments

### Experimental Setup

**Evaluation Metrics** Following previous work (Liu et al. 2025b), for single-VFX evaluation, we employ two established metrics: Fréchet Video Distance (FVD) (Unterthiner et al. 2018) for overall fidelity and Dynamic Degree (Huang et al. 2024b) for motion dynamics. For controllable VFX, we introduce three novel metrics. *Regional Dynamic Degree* (RDD), which utilizes optical flow and masks, quantifies the strength of motion within the target region, thereby quantifies motion strength within target regions to measure visual impression. *Effect Occurrence Rate* (EOR), which is computed by inputting both the video and a given prompt template into Gemini2.5 (Comanici et al. 2025) to obtain the answer, measures intended effect trigger frequency, indicating generation reliability. Building upon EOR, *Effect Controllability Rate* (ECR) assesses spatial precision by verifying VFX confinement to designated areas. Complete metric details appear in Supplement D.

**Implementation Details** Training employs a CogVideoX-5B backbone with LoRA rank of 128 with a total of  $n = 8$  experts, generating  $49 \times 480 \times 720$  resolution videos. For loss,  $\beta$  is set to 0.01. We utilize 8 H20 GPUs (96GB) with a batch size of 1 per GPU. We use AdamW (Loshchilov and Hutter 2017) at a constant  $10^{-4}$  learning rate for 5,000 steps. During inference, DDIM (Nichol and Dhariwal 2021) sampling (50 steps) with CFG (Ho and Salimans 2022) scale 6.0 is applied, which can be perform on a single GPU. Extended details are in Supplement C.

Metrics	Methods	Cake-ify	Crumble	Crush	Decapitate	Deflate	Dissolve	Eye-pop	Harley	Inflate	Levitate	Melt	Squish	Ta-da	Venom	Avg.	Param.#
FVD↓	Single LoRA	2138	2947	1496	<u>1190</u>	913	<u>1770</u>	1995	<u>3576</u>	<u>1505</u>	1401	2827	<u>1415</u>	<b>1053</b>	4146	<u>2026</u>	132.1M
	Mix LoRA	<u>1674</u>	2552	1772	1299	<u>886</u>	1995	<u>1725</u>	4496	2042	<u>1006</u>	<u>2748</u>	<b>1225</b>	1240	<u>3923</u>	2041	9.4M
	LoRA-MoE	<b>1506</b>	<b>1641</b>	<b>1213</b>	<b>1177</b>	<b>839</b>	<b>1118</b>	<b>1460</b>	<b>3330</b>	<b>1304</b>	<b>736</b>	<b>2512</b>	1561	<u>1064</u>	<b>3339</b>	<b>1628</b>	28.5M
Dynamic Degree↑	Single LoRA	0.8	0.8	0.0	0.6	0.0	<b>0.8</b>	0.0	1.0	0.8	0.0	0.6	1.0	1.0	1.0	0.60	132.1M
	Mix LoRA	0.8	0.8	0.0	0.6	0.0	<b>0.8</b>	0.0	1.0	0.8	0.0	0.6	1.0	1.0	1.0	0.60	9.4M
	LoRA-MoE	<b>1.0</b>	<b>1.0</b>	<b>0.6</b>	<b>0.6</b>	0.0	0.4	0.0	<b>1.0</b>	<b>1.0</b>	0.0	<b>0.6</b>	<b>1.0</b>	<b>1.0</b>	<b>1.0</b>	<b>0.66</b>	28.5M

Table 1: Performance comparison on OpenVFX dataset. **Param.#** represents the average training parameters per effect. And the highest metric values are highlighted in **bold**, with the second-best underlined.

Methods	RDD↑				EOR↑				ECR↑			
	Melt	Levitate	Explode	Avg.	Melt	Levitate	Explode	Avg.	Melt	Levitate	Explode	Avg.
<b>CogVideoX</b>	0.91	0.99	1.11	1.00	0.06	0.09	0.11	0.09	0.00	0.00	0.00	0.00
<b>LTX-Video</b>	0.12	0.11	0.14	0.12	0.05	0.02	0.05	0.04	0.00	0.00	0.00	0.00
<b>Wan-2.1</b>	2.06	1.57	2.38	2.00	0.11	0.02	0.03	0.05	0.02	0.00	0.00	0.01
<b>CogV+CN</b>	<b>3.80</b>	<b>2.39</b>	<u>2.09</u>	<u>2.76</u>	<u>0.95</u>	<u>0.80</u>	<u>0.82</u>	<u>0.86</u>	<u>0.56</u>	<u>0.36</u>	<u>0.70</u>	<u>0.54</u>
<b>Ours</b>	<u>2.69</u>	<u>2.22</u>	<b>3.87</b>	<b>2.93</b>	<b>0.99</b>	<b>0.94</b>	<b>0.99</b>	<b>0.97</b>	<b>0.93</b>	<b>0.83</b>	<b>0.89</b>	<b>0.88</b>

Table 2: Quantitative results of Single-VFX control generation. We compare *Omni-Effects* with representative open-source video generation models under three controllable VFX scenarios: Melt, Levitate, and Explode.

## Quantitative Results

In the following, we evaluate the effectiveness of *Omni-Effects* by comparing it with baseline models on unified and controllable VFX generation tasks.

**Unified VFX Generation** We evaluate the LoRA-MoE against VFX-specific training LoRA and mix training LoRA for all VFX on the public OpenVFX dataset as detailed in Table 1. LoRA-MoE achieves the best performance across different types of VFX, while significantly reducing the number of trainable parameters. This demonstrates the effectiveness of the designed VFX task-subspace partitioning strategy. Qualitative results are shown in Supplement E.

**Controllable VFX Generation** To evaluate our model, we perform comprehensive experiments for single- and multi-VFX control, comparing with state-of-the-art methods including (a) CogVideoX (Yang et al. 2024), (b) LTX-Video (HaCohen et al. 2024), (c) Wan2.1 (Wan, Wang et al. 2025), and (d) CogVideoX integrated with ControlNet (CogV+CN). Evaluation targets three spatially localized VFX types—Explode, Melt, and Levitate—to ensure contamination-free assessment.

**Single-VFX Control.** Table 2 demonstrates baseline methods’ fundamental limitations in synthesizing target VFX and achieving precise spatial control. While CogV+CN can synthesize VFX, it exhibits limited controllability. In comparison, *Omni-Effects* achieves the best performance with **0.97** EOR and **0.88** ECR, significantly outperforming all baselines in both generation quality and spatial control precision. This validates that our proposed SAP effectively integrates VFX descriptors with spatial triggers without introducing substantial additional training param-

Methods	Melt+Levitate			Melt+Explode		
	RDD↑	EOR↑	ECR↑	RDD↑	EOR↑	ECR↑
<b>CogVideoX</b>	1.80	0.00	0.00	0.96	0.03	0.00
<b>LTX-Video</b>	0.11	0.00	0.00	0.13	0.00	0.00
<b>Wan-2.1</b>	1.93	0.01	0.00	3.10	0.01	0.00
<b>CogV+CN</b>	<b>3.18</b>	0.09	0.05	<u>3.60</u>	0.08	0.08
<b>Ours</b>	<u>2.63</u>	<b>0.68</b>	<b>0.41</b>	<b>4.59</b>	<b>0.62</b>	<b>0.50</b>

Table 3: Quantitative results of Multi-VFX generation. *Omni-Effects* achieves a high success rate of independent control over Multi-VFX.

ters. Qualitative comparisons are shown in Supplement E.

**Multi-VFX Control.** For multi-VFX generation using two effects combinations, Table 3 shows baseline models consistently failing to generate or spatially control VFX. *Omni-Effects* achieves precise spatial control over simultaneous VFX. Moreover, Figure 6 demonstrates *Omni-Effects*’ superiority: when instructed to melt the left chair while levitating the right chair, CogVideoX erroneously applies melting to both objects; CogV+CN correctly renders melting but fails to generate levitation; whereas *Omni-Effects* simultaneously executes both VFX through spatial condition. This performance validates our proposed IIF’s efficacy in mitigating cross-condition interference. The user study is shown in Supplement E.

**Generalization** Despite being trained with only  $N = 2$  effects, our model generalizes to diverse mask conditions during inference using the shared Spatial-Condition LoRA,




Figure 6: Qualitative comparison of Multi-VFX generation. The desired outcome requires the left chair to melt while the right levitates simultaneously.

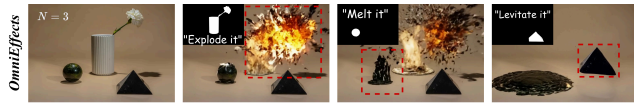


Figure 7: Scalable VFX augmentation. *Omni-Effects* supports inference-time extension to diverse VFX composition.

thereby extending to the generation of more concurrent control VFX ( $N > 2$ ). *Omni-Effects* demonstrates robust extensibility, successfully handling complex effect combinations (Figures 1 (d) and 7), validating its test-time scalable VFX control capability.

## Ablation Studies

**LoRA-MoE** Ablation study on expert count  $n$  and Top- $k$  selection (Table 4) reveals that scaling experts improves generation quality at increased parameter cost. Crucially, our MoE architecture with minimal experts surpasses LoRA baselines (Table 1), demonstrating efficient VFX adaptation through parameter-optimized expert aggregation.

**SAP+IIF** Ablation study on SAP+IIF reveal critical insights in attention mechanisms corresponding to information flow. Removing SAP attention masks from regions  $\{②,⑥,⑦\}$  causes melting artifacts on levitating objects (Figure 8 (b, c)), exposing information leakage, while complete attention induces uncontrolled object melting, demonstrating excessive information interaction degrades control. Strategic masking of  $\{②,⑥,⑦\}$  prevents leakage while preserving independent information flow in target regions. Additional ablation studies are detailed in the Supplement F.

## Conclusion

In this paper, we propose *Omni-Effects*, a unified framework for generating customized VFX videos. It supports the

Metrics	Model	Avg.	Param.#
FVD↓	4 Experts+Top1	1762	18.9
	8 Experts+Top2	<b>1628</b>	28.5
Dynamic Degree↑	4 Experts+Top1	0.65	18.9
	8 Experts+Top2	<b>0.66</b>	28.5

Table 4: Ablation study on LoRA MoE settings. Scaling experts improves generation quality, at the expense of more parameters.



Figure 8: Effect of different attention masks in SAP. Attention Masks are progressively removed while information flow constraints are relaxed from top to bottom.

creation of diverse VFX, ranging from single-VFX, multi-VFX to spatially controllable multi-VFX. To achieve these, our framework integrates two core modules: LoRA-MoE and SAP-IIF. Specifically, the LoRA-MoE module mitigates cross-condition interference arising during mix training of multi-VFX. The SAP module, on the other hand, fuses VFX descriptors with spatial trigger information and tackles cross-condition information leakage via an IIF mechanism. Through the synergistic integration of LoRA-MoE and SAP-IIF, *Omni-Effects* enables precise spatial control and produces high-fidelity multi-VFX composites. We also develop a comprehensive VFX dataset *Omni-VFX* with a specialized data production pipeline and an evaluation framework tailored for controllable VFX generation to further validate our approach. Extensive experiments demonstrate the robustness of *Omni-Effects* across complex, multi-condition VFX generation scenarios. Multi-VFX generation represents a domain of substantial practical value coupled with persistent technical challenges. To the best of our knowledge, this work pioneers the first comprehensive framework explicitly addressing this complex problem. Our methodology substantively advances controllable multi-VFX synthesis capabilities while unlocking novel applications across film production, game development, and advertising creatives.

## References

- Bai, J.; Xia, M.; Wang, X.; Yuan, Z.; Fu, X.; Liu, Z.; Hu, H.; Wan, P.; and Zhang, D. 2024. Syncammaster: Synchronizing multi-camera video generation from diverse viewpoints. *arXiv preprint arXiv:2412.07760*.
- Bao, F.; Xiang, C.; Yue, G.; He, G.; Zhu, H.; Zheng, K.; Zhao, M.; Liu, S.; Wang, Y.; and Zhu, J. 2024. Vidu: a highly consistent, dynamic and skilled text-to-video generator with diffusion models. *arXiv preprint arXiv:2405.04233*.
- Bian, Y.; Zhang, Z.; Ju, X.; Cao, M.; Xie, L.; Shan, Y.; and Xu, Q. 2025. Videopainter: Any-length video inpainting and editing with plug-and-play context control. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, 1–12.
- Chabanova, A. 2022. *VFX—A new frontier: The impact of innovative technology on visual effects*. Ph.D. thesis, University of Westminster.
- Chen, C.; Zhu, J.; Feng, X.; Huang, N.; Wu, M.; Mao, F.; Wu, J.; Chu, X.; and Li, X. 2025. S<sup>2</sup>-Guidance: Stochastic Self Guidance for Training-Free Enhancement of Diffusion Models. *arXiv preprint arXiv:2508.12880*.
- Chen, H.; Xia, M.; He, Y.; Zhang, Y.; Cun, X.; Yang, S.; Xing, J.; Liu, Y.; Chen, Q.; Wang, X.; et al. 2023. Videocrafter1: Open diffusion models for high-quality video generation. *arXiv preprint arXiv:2310.19512*.
- Comanici, G.; Bieber, E.; Schaekermann, M.; Pasupat, I.; Sachdeva, N.; Dhillon, I.; Blistein, M.; Ram, O.; Zhang, D.; Rosen, E.; et al. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Ding, N.; Qin, Y.; Yang, G.; Wei, F.; Yang, Z.; Su, Y.; Hu, S.; Chen, Y.; Chan, C.-M.; Chen, W.; et al. 2023. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature machine intelligence*, 5(3): 220–235.
- Dou, S.; Zhou, E.; Liu, Y.; Gao, S.; Zhao, J.; Shen, W.; Zhou, Y.; Xi, Z.; Wang, X.; Fan, X.; et al. 2023. LoRAMoE: Alleviate world knowledge forgetting in large language models via MoE-style plugin. *arXiv preprint arXiv:2312.09979*.
- Fedus, W.; Zoph, B.; and Shazeer, N. 2022. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120): 1–39.
- HaCohen, Y.; Chiprut, N.; Brazowski, B.; Shalem, D.; Moshe, D.; Richardson, E.; Levin, E.; Shiran, G.; Zabari, N.; Gordon, O.; et al. 2024. Ltx-video: Realtime video latent diffusion. *arXiv preprint arXiv:2501.00103*.
- He, Y.; Yang, T.; Zhang, Y.; Shan, Y.; and Chen, Q. 2022. Latent video diffusion models for high-fidelity long video generation. *arXiv preprint arXiv:2211.13221*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Huang, L.; Wang, W.; Wu, Z.-F.; Shi, Y.; Dou, H.; Liang, C.; Feng, Y.; Liu, Y.; and Zhou, J. 2024a. In-context lora for diffusion transformers. *arXiv preprint arXiv:2410.23775*.
- Huang, Z.; He, Y.; Yu, J.; Zhang, F.; Si, C.; Jiang, Y.; Zhang, Y.; Wu, T.; Jin, Q.; Chanpaisit, N.; et al. 2024b. Vbench: Comprehensive benchmark suite for video generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21807–21818.
- Jiang, Z.; Han, Z.; Mao, C.; Zhang, J.; Pan, Y.; and Liu, Y. 2025. Vace: All-in-one video creation and editing. *arXiv preprint arXiv:2503.07598*.
- Kong, W.; Tian, Q.; Zhang, Z.; Min, R.; Dai, Z.; Zhou, J.; Xiong, J.; Li, X.; Wu, B.; Zhang, J.; et al. 2024. Hunyuan-video: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*.
- Lei, G.; Wang, C.; Zhang, R.; Wang, Y.; Li, H.; and Xu, W. 2025. Animateanything: Consistent and controllable animation for video generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 27946–27956.
- Li, M.; Yang, T.; Kuang, H.; Wu, J.; Wang, Z.; Xiao, X.; and Chen, C. 2024. ControlNet++: Improving Conditional Controls with Efficient Consistency Feedback. In *European Conference on Computer Vision*, 129–147. Springer.
- Ling, X.; Zhu, C.; Wu, M.; Li, H.; Feng, X.; Yang, C.; Hao, A.; Zhu, J.; Wu, J.; and Chu, X. 2025. Vmbench: A benchmark for perception-aligned video motion generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13087–13098.
- Lipman, Y.; Chen, R. T.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2022. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*.
- Liu, S.; Han, Y.; Xing, P.; Yin, F.; Wang, R.; Cheng, W.; Liao, J.; Wang, Y.; Fu, H.; Han, C.; et al. 2025a. Step1x-edit: A practical framework for general image editing. *arXiv preprint arXiv:2504.17761*.
- Liu, X.; Zeng, A.; Xue, W.; Yang, H.; Luo, W.; Liu, Q.; and Guo, Y. 2025b. VFX Creator: Animated Visual Effect Generation with Controllable Diffusion Transformer. *arXiv preprint arXiv:2502.05979*.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, 8162–8171. PMLR.
- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4195–4205.
- Polyak, A.; Zohar, A.; Brown, A.; Tjandra, A.; Sinha, A.; Lee, A.; Vyas, A.; Shi, B.; Ma, C.-Y.; Chuang, C.-Y.; et al. 2024. Movie gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720*.

- Salimans, T.; and Ho, J. 2022. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*.
- Seawead, T.; Yang, C.; Lin, Z.; Zhao, Y.; Lin, S.; Ma, Z.; Guo, H.; Chen, H.; Qi, L.; Wang, S.; et al. 2025. Seaweed-7b: Cost-effective training of video generation foundation model. *arXiv preprint arXiv:2504.08685*.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Tan, Z.; Liu, S.; Yang, X.; Xue, Q.; and Wang, X. 2025a. Ominicontrol: Minimal and universal control for diffusion transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14940–14950.
- Tan, Z.; Xue, Q.; Yang, X.; Liu, S.; and Wang, X. 2025b. Ominicontrol2: Efficient conditioning for diffusion transformers. *arXiv preprint arXiv:2503.08280*.
- Teng, H.; Jia, H.; Sun, L.; Li, L.; Li, M.; Tang, M.; Han, S.; Zhang, T.; Zhang, W.; Luo, W.; et al. 2025. MAGI-1: Autoregressive Video Generation at Scale. *arXiv preprint arXiv:2505.13211*.
- Unterthiner, T.; Van Steenkiste, S.; Kurach, K.; Marinier, R.; Michalski, M.; and Gelly, S. 2018. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*.
- Wan, T.; Wang, A.; et al. 2025. Wan: Open and Advanced Large-Scale Video Generative Models. *arXiv preprint arXiv:2503.20314*.
- Wang, J.; Zhang, Y.; Zou, J.; Zeng, Y.; Wei, G.; Yuan, L.; and Li, H. 2024. Boximator: Generating rich and controllable motions for video synthesis. *arXiv preprint arXiv:2402.01566*.
- Xu, Z.; Yu, Z.; Zhou, Z.; Zhou, J.; Jin, X.; Hong, F.-T.; Ji, X.; Zhu, J.; Cai, C.; Tang, S.; et al. 2025. Hunyuanportrait: Implicit condition control for enhanced portrait animation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 15909–15919.
- Yang, Z.; Teng, J.; Zheng, W.; Ding, M.; Huang, S.; Xu, J.; Yang, Y.; Hong, W.; Zhang, X.; Feng, G.; Yin, D.; Gu, X.; Zhang, Y.; Wang, W.; Cheng, Y.; Liu, T.; Xu, B.; Dong, Y.; and Tang, J. 2024. CogVideoX: Text-to-Video Diffusion Models with An Expert Transformer. *ArXiv*, abs/2408.06072.
- Ye, H.; Zhang, J.; Liu, S.; Han, X.; and Yang, W. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*.
- Yuan, S.; Huang, J.; He, X.; Ge, Y.; Shi, Y.; Chen, L.; Luo, J.; and Yuan, L. 2025. Identity-preserving text-to-video generation by frequency decomposition. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 12978–12988.
- Zeng, Y.; Wei, G.; Zheng, J.; Zou, J.; Wei, Y.; Zhang, Y.; and Li, H. 2024. Make pixels dance: High-dynamic video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8850–8860.
- Zhang, J.; You, J.; Panda, A.; and Goldstein, T. 2025a. Lori: Reducing cross-task interference in multi-task low-rank adaptation. *arXiv preprint arXiv:2504.07448*.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding Conditional Control to Text-to-Image Diffusion Models. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 3813–3824.
- Zhang, Y.; Song, Y.; Liu, J.; Wang, R.; Yu, J.; Tang, H.; Li, H.; Tang, X.; Hu, Y.; Pan, H.; et al. 2024. Ssr-encoder: Encoding selective subject representation for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8069–8078.
- Zhang, Y.; Yuan, Y.; Song, Y.; Wang, H.; and Liu, J. 2025b. Easycontrol: Adding efficient and flexible control for diffusion transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19513–19524.
- Zhang, Z.; Xie, J.; Lu, Y.; Yang, Z.; and Yang, Y. 2025c. In-context edit: Enabling instructional image editing with in-context generation in large scale diffusion transformer. *arXiv preprint arXiv:2504.20690*.
- Zheng, Z.; Peng, X.; Yang, T.; Shen, C.; Li, S.; Liu, H.; Zhou, Y.; Li, T.; and You, Y. 2024. Open-sora: Democratizing efficient video production for all. *arXiv preprint arXiv:2412.20404*.
- Zhou, Y.; Bu, J.; Ling, P.; Zhang, P.; Wu, T.; Huang, Q.; Li, J.; Dong, X.; Zang, Y.; Cao, Y.; et al. 2025. Light-a-video: Training-free video relighting via progressive light fusion. *arXiv preprint arXiv:2502.08590*.