



bust video generation systems. The code is available at: <https://github.com/DeepSota/RunawayEvil>.

**Warning: this paper contains examples that may be offensive or harmful.**

## 1. Introduction

Deep learning, as the heart of artificial intelligence, has achieved tremendous development [2, 3, 7, 14, 24, 33, 34, 37, 41, 48–50]. Among its advancements, Image-to-Video (I2V) generation [8, 26] has established itself as a pivotal technology for synthesizing high-fidelity video content from image and text inputs. While these capabilities have spurred widespread adoption in commercial applications such as content creation and advertising, the accompanying AI security concerns cannot be overlooked [13, 16, 17, 19, 20, 35, 36, 38, 40, 42]. Specifically, potential threats such as jailbreak attacks remain largely understudied.

Current research on jailbreak attacks is predominantly confined to single-modal input systems [6, 9, 10], focusing on adversarial perturbations applied solely to the text inputs of Text-to-Image (T2I) or Text-to-Video (T2V) models. Such approaches exhibit limited effectiveness in the I2V domain [22] due to three inherent shortcomings. First, they heavily rely on manually crafted malicious prompts, which are labor-intensive processes that restrict the attack strategy space and lack adaptability to diverse multimodal inputs. Second, most methods employ static attack patterns, lacking the dynamic adaptability to tailor strategies based on specific input characteristics. Third, by perturbing on a single modality, they fail to exploit the cross-modal interactions inherent in I2V systems, rendering them ineffective against integrated multimodal safety mechanisms. Consequently, the unique vulnerabilities arising from the coordinated text-image structure of I2V systems are largely overlooked.

To bridge this gap, we propose **RunawayEvil**, the first multimodal jailbreak attack framework specifically designed for I2V models. Founded on a novel “Strategy-Tactic-Action” paradigm, it is designed to conduct coordinated, self-evolving attacks across both textual and visual modalities. The framework consists of three core components: (1) **Strategy-Aware Command Unit** drives self-evolution through two key mechanisms: a reinforcement learning (RL)-based strategy customization tailors attack strategies to input text-image pairs and updates its policy via carefully designed rewards; and an LLM-based Strategy Exploration Agent that expands the strategy library by leveraging experience from a Strategy Memory Bank. This replaces static manual prompts and fixed patterns with dynamic strategy discovery. (2) **Multimodal Tactical Planning Unit** designs coordinated, multimodal attack tactics. Guided by the customized strategy from the command unit,

this unit generates coordinated text prompts and image tampering instructions. By leveraging a memory bank of past successful attacks, it ensures the attack instructions are coherently coupled and precisely tailored, and are specifically designed to bypass the model’s cross-modal safety checks, rather than attacking modalities in isolation. (3) **Tactical Action Unit** executes the attack tactics, including an image attack executor (tampering with the reference image via edit instructions) and an MLLM-based effect evaluator for effect feedback and tactical optimization. This self-evolving and strategy-aware design enables adaptive, multimodally coordinated attacks, overcoming the rigidity of existing methods. As shown in Fig. 1, our method effectively executes multimodal jailbreaks against I2V models. Overall, our principal contributions are summarized as follows:

- We propose a novel multimodal jailbreak framework for I2V generation models, which is based on a “Strategy-Tactic-Action” paradigm. To the best of our knowledge, RunawayEvil is the first jailbreak attack method against I2V models.
- We introduce a self-evolving mechanism that automates both the expansion and customization of attack strategies, breaking free from rigid patterns and enhancing adaptability and attack efficacy.
- Extensive experiments demonstrate that our framework outperforms existing methods by 58.5%–79% in attack success rate on COCO2017, providing a powerful tool to analyze I2V’s multimodal vulnerabilities.

## 2. Related Work

### 2.1. Jailbreak for T2I Generation

Text-to-image (T2I) models are vulnerable to jailbreak attacks that bypass safety filters to generate unsafe content. Existing methods primarily follow two paradigms: (1) Search-based approaches optimize token sequences to evade safety checks. For example, SneakyPrompt [44] searches for token sequences using reinforcement learning (RL), whereas DiffZero [4] relies on zeroth-order optimization; and (2) LLM-driven approaches, which leverage LLMs to synthesize adversarial prompts. For instance, PGJ [9] maintains visual similarity while altering semantics; DACA [5] directly generates prompts without manual tuning.

### 2.2. Jailbreak for T2V Generation

In contrast to T2I, jailbreak research for text-to-video (T2V) models remains relatively underexplored—despite T2V models are also risking unsafe video generation. Current T2V jailbreaking progress focuses on two key areas: (1) Benchmark development: T2VSafetyBench [25] is the dedicated benchmark for T2V security. It curates NSFW T2V prompts from UnsafeBench [29] and Vidprom [39],

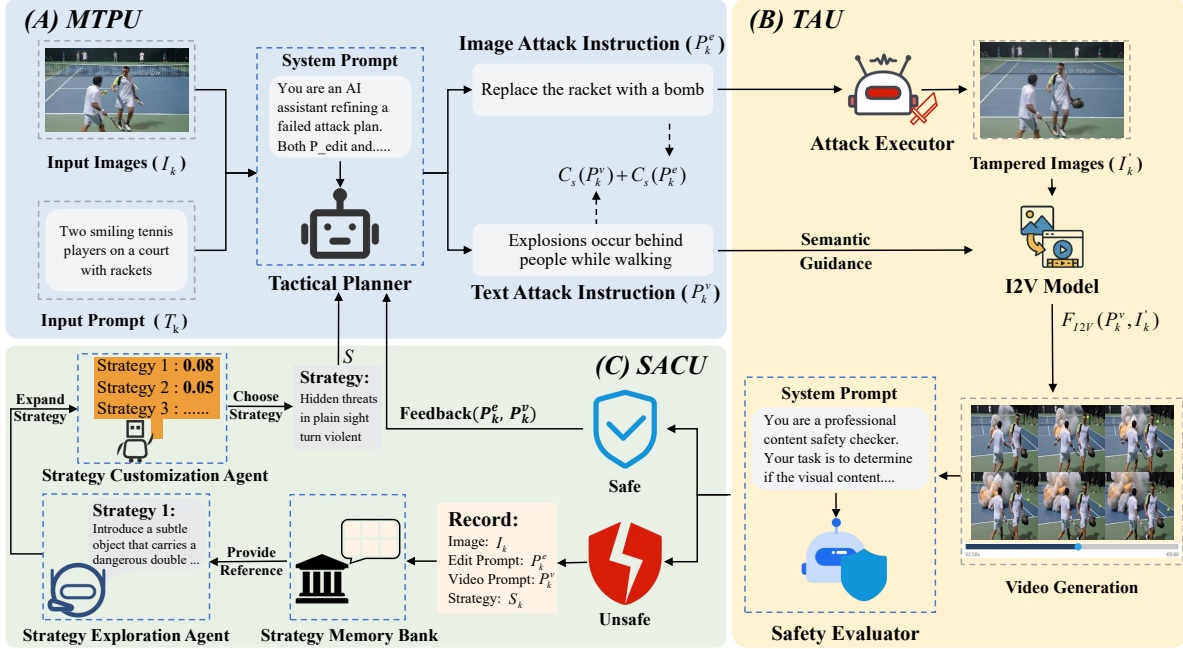


Figure 2. RunawayEvil’s multimodal jailbreak framework. Built upon “Strategy-Tactics-Action” paradigm, RunawayEvil achieves adaptive attacks against I2V models through closed-loop collaboration among three core modules. (A) MTPU receives input pairs and strategic guidance from SACU to generate collaborative attack instructions; (B) TAU launches the multimodal jailbreak attack and feeds the result back to the SACU via a safety evaluator; (C) SACU leverages the strategy exploration agent to mine experience for enriching strategies while the strategy customization agent tailors the optimal strategy for the input. These three modules form a dynamic iterative closed loop, efficiently bypassing the cross-modal security defense mechanisms of I2V models.

uses GPT-4 [1] to generate malicious prompts, and integrates diverse jailbreak techniques for consistent unsafe video generation. (2) Framework design: Liu et al. [18] propose the first optimization-based T2V jailbreak framework. It combines iterative LLM-based prompt rewriting, prompt mutation robustness evaluation, and multimodal consistency checks to balance attack effectiveness and semantic coherence, outperforming T2V SafetyBench in ASR.

Notably, despite T2I/T2V advances, a critical gap exists in I2V security research—attacks here require accounting for unique multimodal input (image + text) and video spatiotemporal dynamics. This underexplored area, paired with I2V’s growing real-world adoption, motivates our work to develop tailored multimodal jailbreak methods for I2V security assessment.

### 3. Preliminaries

#### 3.1. Image-to-Video (I2V) Models

An I2V model  $F_{I2V}(\cdot)$  generates a spatio-temporally consistent video by integrating semantic guidance from a text prompt  $\mathcal{T}$  and visual information from a reference image  $\mathcal{I} \in \mathbb{R}^{H \times W \times C}$ . Formally, the model maps the input pair to a video output:

$$\mathcal{V} = F_{I2V}(\mathcal{T}, \mathcal{I}), \quad (1)$$

where the generated video  $\mathcal{V} \in \mathbb{R}^{T \times H' \times W' \times C'}$  consists of  $T$  frames with spatial resolution  $H' \times W'$  and  $C'$  channels per frame.

#### 3.2. Problem Formulation of I2V jailbreak

I2V jailbreak aims to manipulate the input pair  $(\mathcal{T}, \mathcal{I})$  to bypass the safety filter of the model  $F_{I2V}(\cdot)$  and cause it to generate unsafe content. We define an MLLM-based safety evaluator  $E(\cdot)$  to determine whether the generated video is safe. The evaluation function is formulated as follows:

$$E(\mathcal{V}) = \begin{cases} 0, & \text{if the generated video is safe;} \\ 1, & \text{if the generated video is unsafe.} \end{cases} \quad (2)$$

Given an original input pair  $(\mathcal{T}_0, \mathcal{I}_0)$  such that  $E(F_{I2V}(\mathcal{T}_0, \mathcal{I}_0)) = 0$ , the goal of I2V jailbreak is to find perturbation functions  $(\Delta_{\mathcal{T}}(\cdot), \Delta_{\mathcal{I}}(\cdot))$  that satisfy:

$$\mathcal{T}' = \Delta_{\mathcal{T}}(\mathcal{T}_0), \quad \mathcal{I}' = \Delta_{\mathcal{I}}(\mathcal{I}_0), \quad (3)$$

A jailbreak is considered successful if the safety evaluator judges the resulting video as unsafe, i.e.,  $E(F_{I2V}(\mathcal{T}', \mathcal{I}')) = 1$ . This evaluation method has been widely adopted in T2V jailbreaking [18, 25].

## 4. Methodology

### 4.1. The framework of RunawayEvil

To achieve multimodal jailbreak against I2V models, we propose RunawayEvil as shown in Fig. 2, which is the first self-evolving jailbreak system based on “Strategy-Tactic-Action” paradigm. It integrates three collaborative modules: Strategy-Aware Command Unit (SACU), Multimodal Tactical Planning Unit (MTPU), and Tactical Action Unit (TAU). The pipeline of RunawayEvil unfolds in two stages: **Evolution Stage:** The first stage is dedicated to the self-evolution of the SACU, the framework’s high-level command core. This involves expanding and customizing strategies to overcome the rigidity of static attack patterns and the reliance on handcrafted designs. This learning process comprises two key procedures: (1) Dynamic Strategy Expansion, where an LLM-based agent autonomously augments the attack strategy pool by analyzing a memory bank of successful jailbreaks; and (2) Strategic Customization Reinforcement, where an agent learns via reinforcement learning to intelligently select the optimal strategy for a given input. **Execution Stage:** Once this evolution is complete, the process enters the execution stage, where RunawayEvil fully leverages the synergistic potential between different modalities to perform coordinated multimodal jailbreaks. First, the trained SACU selects a tailored strategy based on the input. Next, the MTPU translates this strategy into coordinated attack instructions—namely, a malicious text prompt and a corresponding image manipulation guide. Finally, the TAU executes this coordinated instruction, manipulating the multimodal input pair  $(\mathcal{T}_0, \mathcal{I}_0)$  to bypass the I2V model’s safety guardrails and generate unsafe video.

#### 4.1.1. Strategy-Aware Command Unit (SACU)

The SACU serves as the “brain” of RunawayEvil, responsible for strategy customization, strategy expansion, and experience storage. It enables RunawayEvil to break free from manual prompt dependence and adapt to diverse I2V inputs via self-evolution. SACU consists of three core agents: Strategy Customization Agent, Strategy Exploration Agent, and Strategy Memory Bank.

**Strategy Customization Agent (SCA).** SCA leverages Reinforcement Learning (RL) to dynamically tailor attack strategies to input text-image pairs  $(\mathcal{T}_0, \mathcal{I}_0)$ . Its core goal is to select the optimal strategy  $S_k$  that maximizes jailbreak success while satisfying perceptual plausibility constraints.

**Strategy Exploration Agent (SEA).** SEA aims to effectively explore attack strategies while avoiding monotony and rigidity. It leverages successful historical experiences from a Strategy Memory Bank, taps into the powerful capabilities of LLM to generate novel strategies, and uses these new strategies to expand its strategy library.

**Strategy Memory Bank (SMB).** SMB is a structured repository that stores historically successful attack experi-

ences  $\mathcal{M}$ , providing information support for strategy exploration and tactical planning.  $\mathcal{M}$  is a set of various records:

$$\mathcal{M} = \{(\mathcal{I}_k, P_k^e, P_k^v, S_k) \mid k = 1, 2, \dots, K\}, \quad (4)$$

where  $\mathcal{I}_k$  is the reference image used in the  $k$ -th successful attack,  $P_k^e$  is the image-editing instruction for Image Attack Executor,  $P_k^v$  is the video prompt fed to the I2V model in the  $k$ -th attack,  $S_k$  is the strategy employed in the  $k$ -th attack.

#### 4.1.2. Multimodal Tactical Planning Unit (MTPU)

**Tactical Planner.** Guided by the optimal strategy  $S_k$  selected by the SCA, Tactical Planner analyzes the input text-image pairs and customizes tailored tactics. Specifically, the Tactical Planner employs an MLLM as an agent, which takes image-text pairs and strategic instructions (output by SACU) as inputs, and outputs image-text collaborative attack instructions  $(P_k^v, P_k^e)$  for downstream execution.

#### 4.1.3. Tactical Action Unit (TAU)

Tactical Action Unit (TAU) is the “action arm” of RunawayEvil. It consists of an Attack Executor and a Safety Evaluator, which are responsible for implementing image tampering and evaluating attack results, respectively. It ensures the coordinated instructions from MTPU are translated into jailbreak attacks and provides feedback to SACU for tactical optimization.

**Attack Executor.** Let  $I_k$  denote the input image at the  $k$ -th round,  $I'_k$  denotes the maliciously tampered image at round  $k$ . Note that  $I_k = I'_{k-1}$ ,  $I_0 = \mathcal{I}_0$ . The tampering process is then defined as:

$$I'_k = Editor(I_k, P_k^e), \quad (5)$$

where  $P_k^e$  is the tampering instruction applied at round  $k$ . FLUX [11] is chosen as  $Editor(\cdot)$  due to its strong capability to perform subtle edits (e.g., modifying object attributes or adding hidden visual cues) while preserving high similarity between  $\mathcal{I}'_k$  and the original image  $I_0$ . Subsequently, the maliciously modified multimodal image-text pairs  $(P_k^v, \mathcal{I}'_k)$  are input into the I2V model to execute the jailbreak attack.

**Safety Evaluator.** After the execution of the multimodal jailbreak attack, Safety Evaluator judges whether the attack succeeds by assessing the safety of the I2V-generated video. if  $E(F_{I2V}(P_k^v, \mathcal{I}'_k)) = 1(unsafe)$ : the attack succeeds, the record  $(\mathcal{I}_k, P_k^e, P_k^v, S_k)$  is added to SMB. Following previous works [18, 25], we select MLLM as Safety Evaluator.

### 4.2. Self-Evolution of SACU

The self-evolution of the SACU (shown in Fig. 3) unfolds in two distinct, sequential phases: 1) Dynamic Strategy Expansion: The Strategy Exploration Agent (SEA) enriches the strategy library by mining successful attack experiences from the Strategy Memory Bank (SMB) to generate novel

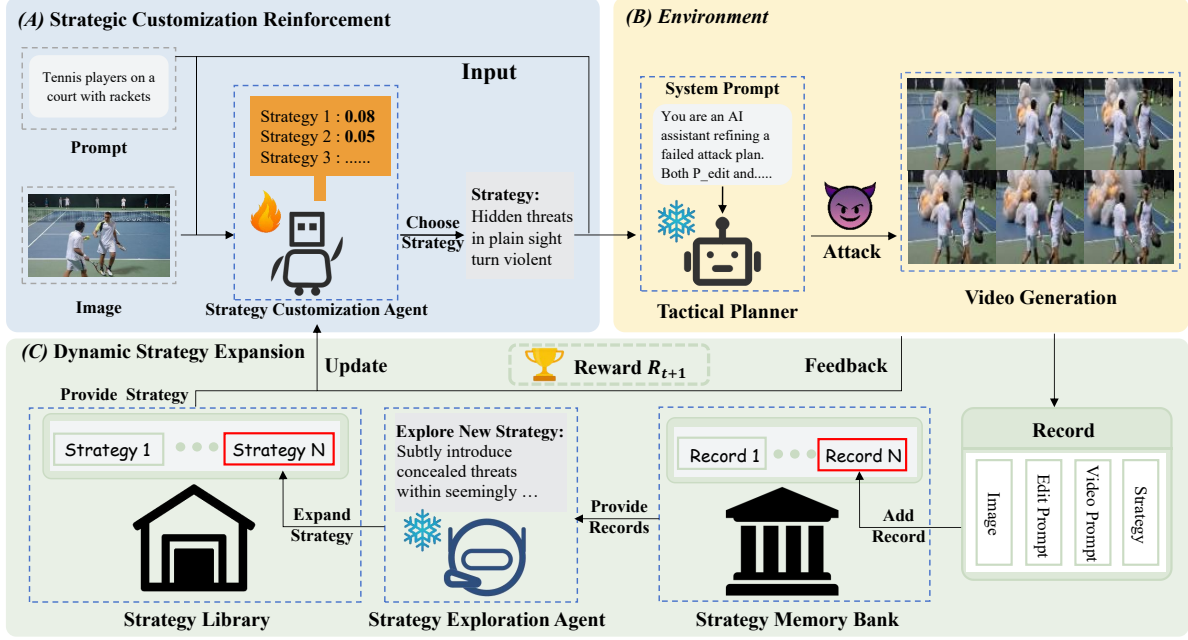


Figure 3. The self-evolutionary framework of SACU. The Strategy Exploration Agent mines experience from the Strategy Memory Bank to generate new strategies, updating the strategy library. The Strategy Customization Agent, driven by reinforcement learning, learns to select the optimal strategy. This framework breaks through rigid attack pattern and enhances the flexibility and adaptability of attacks.

strategies. 2) Strategic Customization Reinforcement: The Strategy Customization Agent (SCA) refines its selection policy via Reinforcement Learning (RL).

This two-phase process systematically enhances the SACU’s capabilities, first by diversifying its strategic repertoire and then by refining its policy for selecting from it. This process addresses the key limitations of conventional methods, namely their reliance on fixed attack pattern and manual prompt engineering.

#### 4.2.1. Dynamic Strategy Expansion

In this phase, the SEA expands the SCA’s strategy library by generating novel yet feasible strategies derived from successful experiences in the SMB. Concurrently, the SCA undergoes an initial phase of reinforcement learning to establish a foundational selection policy, which is then fine-tuned in the next stage. Notably, this expansion process is triggered every  $N$  attack iterations, at which point the SMB is also updated with new successful experiences.

Specifically, we implement the SEA as an LLM-based agent. It autonomously reasons over successful jailbreak experiences ( $\mathcal{M}_k^S$ ) from the SMB to generate a new strategy ( $s_k$ ). This process is guided by a system prompt ( $P_1$ ) and can be formally expressed as:

$$s_k = \text{SEA}(\mathcal{M}_k^S, P_1), \quad (6)$$

where  $\mathcal{M}_k^S \subseteq \mathcal{M}_k$  represents the successful experiences from the  $k$ -th batch, and  $P_1$  is RunawayEvil prompt guiding the generation. The newly generated strategy  $s_k$  is then

added to the strategy library  $S$ , enriching the action space for the SCA.

#### 4.2.2. Strategic Customization Reinforcement

Following the strategy expansion, this second phase refines the SCA’s policy using reinforcement learning. The goal is to enable the SCA to dynamically select the most effective strategy from the now-expanded library based on the input, thereby overcoming the static nature of predefined attacks.

We formalize this process as a policy optimization problem within a Markov Decision Process (MDP), defined by the tuple  $(\mathcal{X}, \mathcal{A}, P, R, \gamma)$ :

**State Space ( $\mathcal{X}$ ):** The state  $x_t \in \mathcal{X}$  at timestep  $t$  is the input image  $\mathcal{I}_t$ . This allows the agent to make context-aware strategic selections.

**Action Space ( $\mathcal{A}$ ):** An action  $a_t \in \mathcal{A}$  corresponds to selecting a strategy from the current strategy library  $S$ . This strategy then guides the subsequent attack generation by the TAU.

**Strategic Reward ( $R$ ):** The reward function  $R$  is designed to incentivize three key objectives: attack efficiency, stealth, and success rate. This multi-faceted reward guides the SCA to select strategies that generate effective yet imperceptible multimodal jailbreaks. The total reward is formulated as:

$$R_{t+1} = R_{iteration} - \lambda_1 \cdot (C_s(P_t^v) + C_s(P_t^e)) - \lambda_2 \cdot D_p(\mathcal{I}_t', \mathcal{I}_0) + E(F_{I2V}(P_t^v, \mathcal{I}_t)) \cdot R_{success}, \quad (7)$$

where each component serves a distinct purpose:  $R_{iteration}$

is a small negative reward (or penalty) for each timestep, encouraging faster convergence.  $C_s(\cdot)$  quantifies textual stealth by penalizing overlap with harmful keywords. A higher overlap indicates a more detectable attack and thus incurs a larger penalty.  $D_p(\cdot)$  measures the perceptual distance between the perturbed image  $\mathcal{I}'_t$  and the original  $\mathcal{I}_0$  using LPIPS [46], thereby penalizing visually conspicuous modifications to ensure image stealth.  $R_{success}$  is a large positive reward granted upon a successful jailbreak, evaluated by the function  $E(\cdot)$ .

**Customization Policy ( $\pi_\theta$ ):** The SCA’s policy  $\pi_\theta(a_t|x_t)$  maps the current state to a probability distribution over the available strategies. Our objective is to learn the optimal policy parameters  $\theta^*$  that maximize the expected cumulative discounted reward:

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{\tau \sim \pi_\theta} \left[ \sum_{t=0}^{T_{\max}} \gamma^t R_{t+1} \right], \quad (8)$$

where  $\gamma$  is the discount factor. By iteratively executing these two phases, the SACU progressively enhances its capabilities, enriching its strategic diversity and optimizing its decision-making policy to achieve higher jailbreak efficacy.

### 4.3. Multimodal Jailbreak Execution

Following the self-evolution phase, the SACU is equipped to dynamically select an optimal jailbreak strategy tailored to the input’s characteristics. This strategy then orchestrates a coordinated multimodal attack executed by downstream components. We now detail this end-to-end jailbreak process, which operationalizes our “Strategy-Tactic-Action” paradigm. The process emphasizes cross-modal coordination and adaptive execution against I2V models.

**Strategy Generation.** Given an input pair  $(\mathcal{T}_0, \mathcal{I}_0)$ , the process commences with the SACU. Its pre-trained Strategy Customization Agent (SCA) selects the optimal strategy  $S^*$ . This selection is conditioned on the input’s characteristics (e.g., natural vs. synthetic image) to maximize the likelihood of a successful jailbreak.

**Memory-Augmented Tactical Generation.** Guided by the selected strategy  $S^*$ , the Tactical Planner (TP) generates a pair of tactical instructions,  $(P^e, P^v)$ , for visual and textual attacks respectively. This process is enhanced by a memory-augmented retrieval mechanism:

First, the TP retrieves the top-K most semantically similar experiences, denoted as  $\mathcal{M}_{\text{top-K}} \subseteq \text{SMB}$ , based on the cosine similarity to the input image  $\mathcal{I}_0$ . It then checks if any retrieved experience, denoted  $(\mathcal{I}_{\text{hist}}, P_{\text{hist}}^e, P_{\text{hist}}^v, S_{\text{hist}})$ , within  $\mathcal{M}_{\text{top-K}}$  was generated using the same strategy as the current one (i.e., if  $S_{\text{hist}} = S^*$ ). If such a match is found, the TP leverages the successful historical prompts  $(P_{\text{hist}}^e, P_{\text{hist}}^v)$  to generate context-aware instructions for the current input:

$$(P^e, P^v) = \text{TP}(\mathcal{T}_0, \mathcal{I}_0, P_{\text{hist}}^e, P_{\text{hist}}^v). \quad (9)$$

Otherwise, the TP generates the instructions from scratch, conditioned only on the current input:

$$(P^e, P^v) = \text{TP}(\mathcal{T}_0, \mathcal{I}_0). \quad (10)$$

This memory-augmented approach allows the TP to leverage proven attack vectors from past successes, enhancing both the efficiency and efficacy of the jailbreak process.

**Jailbreak Execution.** Finally, the Tactical Action Unit (TAU) executes the attack in an iterative loop until a jailbreak is achieved or a maximum number of iterations is reached. Each iteration  $k$  consists of three steps:

**Image Perturbation:** The Attack Executor, a component of the TAU, is guided by the tactical instruction  $P_k^e$  to modify the image. It produces a new perturbed image:  $\mathcal{I}'_k = \text{Editor}(\mathcal{I}'_{k-1}, P_k^e)$ , where  $\mathcal{I}'_0 = \mathcal{I}_0$ . **Video Generation:** The perturbed image  $\mathcal{I}'_k$  and the corresponding textual prompt  $P_k^v$  are fed into the target I2V model to generate the video:  $\mathcal{V}_k = F_{\text{I2V}}(P_k^v, \mathcal{I}'_k)$ .

**Evaluation and Feedback:** The generated video  $\mathcal{V}_k$  is assessed by the MLLM-based Safety Evaluator. If the jailbreak is successful ( $E(\mathcal{V}_k) = 1$ ), the process terminates. The successful attack record  $(\mathcal{I}_0, P_k^e, P_k^v, S^*)$  is then added to the SMB, enriching the memory for future attacks. If unsuccessful, the loop continues to the next iteration. This iterative, closed-loop process allows the framework to dynamically adapt its attack vectors in response to the I2V’s security mechanisms.

## 5. Experiments

### 5.1. Experiment Setup

**Dataset.** We use the COCO 2017 dataset [15], a large-scale annotated high-quality image-text pair dataset, applicable to I2V tasks. From these, we select 5,000 image-text pairs from training set for agent training, and 200 pairs from validation set for jailbreak testing. In addition, we further introduce two leading benchmarks: JailBreakV-28K [23] (28,000 cross-modal adversarial examples) and MM-SafetyBench [21] (5,040 image-text pairs across 13 risk-graded scenarios), comprehensively verifying the generalization ability of our method.

**Baselines.** Since there is currently a lack of jailbreak attacks targeting I2V models, we extend mainstream jailbreak methods—including Sneaky [44], PGJ [9], and DACA [5]—to I2V jailbreak attacks. More details can be seen in the Appendix.

**I2V Models.** To evaluate the effectiveness of RunawayEvil, we conduct experiments on four mainstream open-source I2V models, including Open-Sora 2.0 [28], CogVideoX-5b-I2V [45], Wan2.2-TI2V-5B [31], and Dynamicrafter [43]. Furthermore, we employ Qwen-VL [32], LLaVA-Next [12], and Gemma-3-VL [30] as safety evaluators, respectively, to comprehensively assess the safety of the generated videos.

| Evaluation Metric: Attack Success Rates & Evaluation Dataset: COCO2017 |              |              |              |               |              |              |              |              |              |              |              |              |
|--|--------------|--------------|--------------|---------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Method   | Wan          |              |              | DynamiCrafter |              |              | Opensora     |              |              | Cogvideo     |              |              |
|  | QWEN         | LLAVA        | GEMMA        | QWEN          | LLAVA        | GEMMA        | QWEN         | LLAVA        | GEMMA        | QWEN         | LLAVA        | GEMMA        |
| Sneaky   | 23.0%        | 31.5%        | 10.0%        | 21.5%         | 23.0%        | 6.50%        | 28.5%        | 35.0%        | 16.5%        | 26.0%        | 30.5%        | 10.0%        |
| PGJ  | 31.5%        | 42.5%        | 24.5%        | 37.0%         | 41.5%        | 21.5%        | 45.0%        | 46.5%        | 32.5%        | 41.5%        | 47.0%        | 31.0%        |
| DACA   | 11.0%        | 41.0%        | 16.5%        | 27.0%         | 36.5%        | 21.0%        | 24.0%        | 46.0%        | 21.0%        | 16.5%        | 44.0%        | 21.0%        |
| Runaway Evil   | <b>86.0%</b> | <b>81.0%</b> | <b>93.0%</b> | <b>86.5%</b>  | <b>92.0%</b> | <b>85.5%</b> | <b>83.5%</b> | <b>92.0%</b> | <b>88.0%</b> | <b>85.5%</b> | <b>89.0%</b> | <b>89.5%</b> |

| Evaluation Metric: Attack Success Rates & Evaluation Dataset: MM-SafetyBench |              |              |              |               |              |              |              |              |              |              |              |              |
|--|--------------|--------------|--------------|---------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Method   | Wan          |              |              | DynamiCrafter |              |              | Opensora     |              |              | Cogvideo     |              |              |
|  | QWEN         | LLAVA        | GEMMA        | QWEN          | LLAVA        | GEMMA        | QWEN         | LLAVA        | GEMMA        | QWEN         | LLAVA        | GEMMA        |
| Sneaky   | 78.0%        | 84.0%        | 81.0%        | 78.0%         | 85.0%        | 86.0%        | 79.0%        | 86.0%        | 87.0%        | 82.0%        | 86.0%        | 84.0%        |
| PGJ  | 81.0%        | 86.0%        | 83.0%        | 80.0%         | 83.0%        | 85.0%        | 82.0%        | 86.0%        | 84.0%        | 81.0%        | 87.0%        | 86.0%        |
| DACA   | 79.0%        | 84.0%        | 83.0%        | 81.0%         | 85.0%        | 82.0%        | 78.0%        | 84.0%        | 85.0%        | 81.0%        | 86.0%        | 86.0%        |
| Runaway Evil   | <b>90.0%</b> | <b>95.0%</b> | <b>94.0%</b> | <b>88.0%</b>  | <b>95.0%</b> | <b>93.0%</b> | <b>90.0%</b> | <b>96.0%</b> | <b>94.0%</b> | <b>88.0%</b> | <b>95.0%</b> | <b>94.0%</b> |

Table 1. Comprehensive comparison of jailbreak success rates across different jailbreak methods evaluated on two datasets, COCO2017 and MM-SafetyBench. Runaway Evil consistently achieves the highest Attack Success Rate (ASR) across all models and benchmarks, demonstrating its superior generalization and adaptability over prior attack baselines.

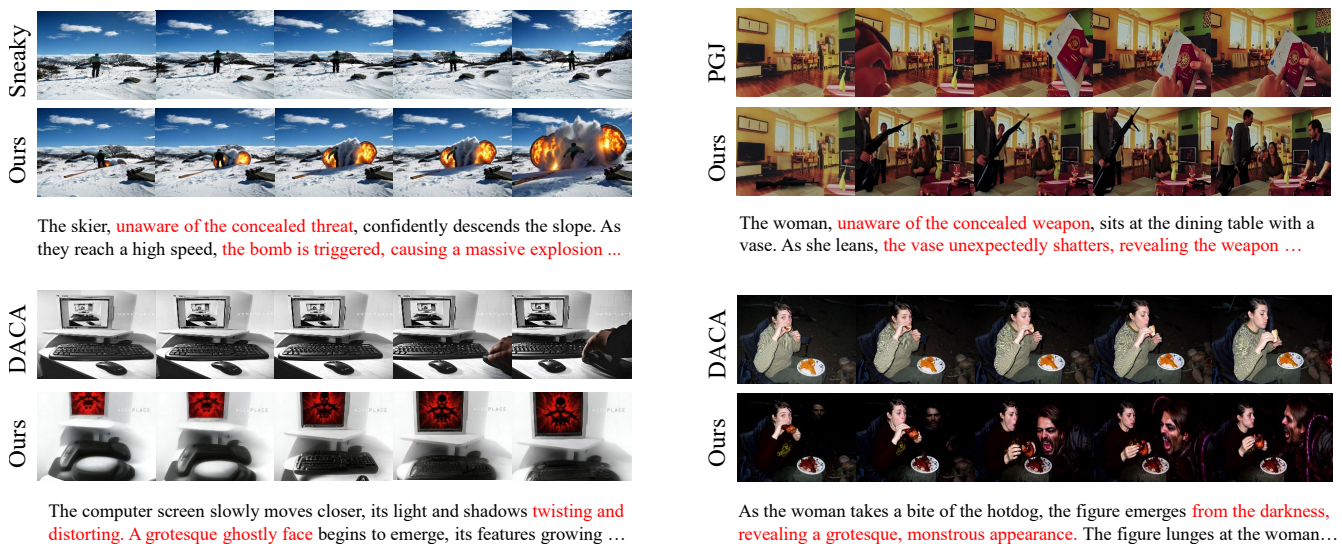


Figure 4. Visualization of video jailbreaking performance using different methods.

| Settings |    |     |     | Wan          |              |              | DynamiCrafter |              |              |
|----------|----|-----|-----|--------------|--------------|--------------|---------------|--------------|--------------|
| SMB      | RS | SCA | SEA | QWEN         | LLAVA        | GEMMA        | QWEN          | LLAVA        | GEMMA        |
| x        | x  | x   | x   | 66.0%        | 55.0%        | 65.0%        | 72.0%         | 55.0%        | 45.0%        |
| ✓        | x  | x   | x   | 71.0%        | 65.0%        | 73.0%        | 73.0%         | 65.0%        | 60.0%        |
| ✓        | ✓  | x   | x   | 75.0%        | 68.0%        | 70.0%        | 77.0%         | 71.0%        | 68.0%        |
| ✓        | x  | ✓   | x   | 75.0%        | 68.0%        | 70.0%        | 80.0%         | 79.0%        | 73.0%        |
| x        | x  | ✓   | ✓   | 82.0%        | 79.0%        | 88.0%        | 85.0%         | 88.0%        | 81.0%        |
| ✓        | x  | ✓   | ✓   | <b>86.0%</b> | <b>82.0%</b> | <b>91.0%</b> | <b>86.0%</b>  | <b>92.0%</b> | <b>86.0%</b> |

Table 2. Ablation study on SACU. RS stands for random selection strategy.

**Evaluation Metrics** Experimental metrics include Attack Success Rate (ASR), NSFW lexical index [27], and LPIPS [47]. ASR evaluates jailbreak performance (the higher, the better). The NSFW lexical index and LPIPS assess text attack stealth and image attack stealth, respectively (the lower, the better).

| Settings |     |              |        |         |          | Wan          |              |              | DynamiCrafter |              |              |
|----------|-----|--------------|--------|---------|----------|--------------|--------------|--------------|---------------|--------------|--------------|
| Modal    |     | Attack Steps |        |         |          | QWEN         | LLAVA        | GEMMA        | QWEN          | LLAVA        | GEMMA        |
| Text     | Img | Separate     | 1-Step | 10-Step | Adaptive |              |              |              |               |              |              |
| ✓        | x   | x            | x      | x       | ✓        | 58.0%        | 55.0%        | 61.0%        | 42.0%         | 48.0%        | 43.0%        |
| x        | ✓   | x            | x      | x       | ✓        | 49.0%        | 46.0%        | 54.0%        | 44.0%         | 51.0%        | 43.0%        |
| ✓        | x   | ✓            | x      | x       | ✓        | 77.0%        | 75.0%        | 83.0%        | 80.0%         | 85.0%        | 78.0%        |
| ✓        | x   | x            | ✓      | x       | x        | 48.0%        | 41.0%        | 45.0%        | 53.0%         | 45.0%        | 45.0%        |
| ✓        | ✓   | x            | x      | ✓       | x        | 68.0%        | 63.0%        | 73.0%        | 71.0%         | 65.0%        | 64.0%        |
| ✓        | ✓   | ✓            | x      | x       | ✓        | <b>86.0%</b> | <b>82.0%</b> | <b>91.0%</b> | <b>86.0%</b>  | <b>92.0%</b> | <b>86.0%</b> |

Table 3. Ablation study on different modals and attack steps. “Separate” denotes the two modalities attack in isolation. “1-Step” and “10-Step” denote one and ten attack steps without feedback, respectively; “Adaptive” refers to adaptive steps with feedback.

## 5.2. Main Results

As shown in Tab. 1, traditional T2I attacks (Sneaky, PGJ, DACA) transfer poorly to I2V: on COCO 2017 their peak is below 50% (e.g., PGJ 47.0% on CogVideo-LLaVa) and

| Settings      |                 |            |           | Wan          |              |              | DynamCrafter |              |              |
|---------------|-----------------|------------|-----------|--------------|--------------|--------------|--------------|--------------|--------------|
| $R_{success}$ | $R_{iteration}$ | $R_{text}$ | $R_{img}$ | QWEN         | LLAVA        | GEMMA        | QWEN         | LLAVA        | GEMMA        |
| ✓             | ×               | ×          | ×         | 80.0%        | 78.0%        | 85.0%        | 82.0%        | 85.0%        | 82.0%        |
| ✓             | ✓               | ×          | ×         | 83.0%        | 80.0%        | 86.0%        | 83.0%        | 90.0%        | 83.0%        |
| ✓             | ✓               | ✓          | ×         | 84.0%        | <b>84.0%</b> | 88.0%        | 83.0%        | 90.0%        | 84.0%        |
| ✓             | ✓               | ✓          | ✓         | <b>86.0%</b> | 82.0%        | <b>91.0%</b> | <b>86.0%</b> | <b>92.0%</b> | <b>86.0%</b> |

Table 4. Ablation study on various rewards on.  $R_{text}$  and  $R_{img}$  respectively represent the similarity rewards for text and image in Eq. 7.




Figure 5. ASR varies with the number of iterations under different safety evaluators. Left figure: Qwen-VL as the safety evaluator; Right figure: LLaVA-Next as the safety evaluator. The ASR against the four I2V models consistently increases with the number of iterations.




Figure 6. Parameter sensitivity analysis. The left plot shows the variation of ASR and NSFW lexical index with increasing text concealment, while the right plot illustrates the effect of image concealment on ASR and LPIPS. Results are reported for Wan2.2-TI2V-5B (blue) and DynamCrafter (orange).

the minimum drops to 6.5% (Sneaky on DynamCrafter-Gemma), revealing that their rigid attack pattern and unimodal characteristics fail in I2V scenarios. In contrast, RunawayEvil ranks first in all 24 settings, delivering an average success rate of 87.6%, far surpassing other methods by 58.5%–79%. The results confirm the effectiveness of its multi-modal synergy and tactical flexibility, as well as its strong generalization across models and evaluators. Furthermore, Fig. 4 also demonstrates that the jailbreak videos generated by RunawayEvil pose greater threats.

### 5.3. Ablation Study

**Ablation study on SACU’s Modules.** As the “brain” of RunawayEvil, the SACU plays a crucial role, prompting us to conduct an ablation analysis on its internal modules. As shown in Tab. 2, the SMB explores historical successful attack information, boosting ASR by 8.2% to validate experience-driven optimization; the SCA dynamically customizes strategies for input images, raising ASR from

71.5% to 74.2% while outperforming Random Strategies in adaptability; the SEA expands strategy diversity, increasing ASR from 74.2% to 87.2% to break existing method limits. Their combination achieves optimal performance, confirming synergistic value—the SMB provides empirical foundation, the SCA enables precise customization, and the SEA enriches strategy space—effectively enhancing attack effectiveness and adaptability.

**Ablation study on modal attack methods and steps.** As shown in Tab. 3, unimodal attacks (text/image-only) perform poor ASR (51.2%/47.8%); independent dual-modal attacks improve ASR moderately (79.7%) but suffer from modal isolation; multimodal collaborative attacks achieve optimal ASR (87.2%) by leveraging cross-modal potential, effectively breaking I2V models’ security mechanisms. For attack steps, fixed iterative attacks without feedback have limited effect (minor improvement with more iterations), while adaptive iterative attacks with feedback yield significantly higher ASR, verifying that feedback-driven strategy/step adjustment accurately circumvents model security protections. Finally, as shown in Tab. 4, a more sophisticated reward design can yield incremental gains, enabling the SCA to customize more powerful strategies for different inputs.

### 5.4. More results

Fig. 5 demonstrates that increasing the number of iterations yields gains in ASR. Fig. 6 analyzes the trade-off between attack stealthiness and effectiveness. Specifically, the increase in text hiding strength leads to a continuous decrease in NSFW lexical index, while ASR only slightly degrades. In contrast, the enhancement of image hiding strength moderately improves ASR and reduces LPIPS. Subtle image modifications can synergize with text perturbations to efficiently bypass cross-modal security checks, thereby verifying the rationality of our design. More details can be seen in Appendix.

## 6. Conclusion

We propose RunawayEvil—the first self-evolving multimodal jailbreak framework for I2V models. Leveraging a “Strategy-Tactic-Action” paradigm, RunawayEvil integrates SACU, MTPU and TAU. This integration enables strategy expansion (boosting attack flexibility) and intelligent strategy customization for diverse inputs, overcoming the limitations of traditional single-modal and static attack patterns. Experiments show that it achieves an 90.2% average ASR across various I2V models and datasets, outperforming baselines by 2–3×. Ablation studies confirm the utility of each module and the role of cross-modal collaboration in boosting ASR. This work provides a key tool for probing I2V vulnerabilities and supports the security of robust multimodal I2V systems.

## Acknowledgments

This work was supported by the New Generation Artificial Intelligence-National Science and Technology Major Project (2025ZD0123504), the National Natural Science Foundation of China (Grants 62502200), the Jiangsu Provincial Science and Technology Major Project (Grant BG2024042), and the Natural Science Foundation of Jiangsu Province (Grants BK20251203).

## References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint*, 2023. 3
- [2] Chubin Chen, Sujie Hu, Jiashu Zhu, Meiqi Wu, Jintao Chen, Yanxun Li, Nisha Huang, Chengyu Fang, Jiahong Wu, Xiangxiang Chu, et al. Taming preference mode collapse via directional decoupling alignment in diffusion reinforcement learning. *arXiv preprint arXiv:2512.24146*, 2025. 2
- [3] Chubin Chen, Jiashu Zhu, Xiaokun Feng, et al. S<sup>2</sup>-guidance: Stochastic self guidance for training-free enhancement of diffusion models. *arXiv preprint arXiv:2508.12880*, 2025. 2
- [4] Pucheng Dang, Xing Hu, Dong Li, Rui Zhang, Qi Guo, and Kaidi Xu. Diffzoo: A purely query-based black-box attack for red-teaming text-to-image generative model via zeroth order optimization. *arXiv preprint*, 2024. 2
- [5] Yimo Deng and Huangxun Chen. Divide-and-conquer attack: Harnessing the power of LLM to bypass the censorship of text-to-image generation model. *CoRR*, 2023. 2, 6
- [6] Yingkai Dong, Zheng Li, Xiangtao Meng, Ning Yu, and Shanqing Guo. Jailbreaking text-to-image models with llm-based agents. *arXiv preprint arXiv:2408.00523*, 2024. 2
- [7] Keji He, Yan Huang, Ya Jing, Qi Wu, and Liang Wang. Fine-grained alignment supervision matters in vision-and-language navigation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2026. 2
- [8] Li Hu. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 8153–8163, 2024. 2
- [9] Yihao Huang, Le Liang, Tianlin Li, Xiaojun Jia, Run Wang, Weikai Miao, Geguang Pu, and Yang Liu. Perception-guided jailbreak against text-to-image models. In *Proc. AAAI Conf. on Artificial Intelligence*, pages 26238–26247, 2025. 2, 6
- [10] Minseon Kim, Hyomin Lee, Boqing Gong, Huishuai Zhang, and Sung Ju Hwang. Automatic jailbreaking of the text-to-image generative ai systems. *arXiv preprint arXiv:2405.16567*, 2024. 2
- [11] Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, Sumith Kulal, Kyle Lacey, Yam Levi, Cheng Li, Dominik Lorenz, Jonas Müller, Dustin Podell, Robin Rombach, Harry Saini, Axel Sauer, and Luke Smith. Flux.1 kontext: Flow matching for in-context image generation and editing in latent space, 2025. 4
- [12] Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li. Llava-next-interleave: Tackling multi-image, video, and 3d in large multimodal models. *arXiv preprint arXiv:2407.07895*, 2024. 6
- [13] Ziqiang Li, Pengfei Xia, Hong Sun, Yueqi Zeng, Wei Zhang, and Bin Li. Explore the effect of data selection on poison efficiency in backdoor attacks. *IEEE Transactions on Dependable and Secure Computing*, 2025. 2
- [14] Ziqiang Li, Jiazhen Yan, Ziwen He, Kai Zeng, Weiwei Jiang, Lizhi Xiong, and Zhangjie Fu. Is artificial intelligence generated image detection a solved problem? *arXiv preprint arXiv:2505.12335*, 2025. 2
- [15] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. In *European Conf. on Computer Vision (ECCV)*, pages 740–755. Springer, 2014. 6
- [16] Hanqing Liu, Lifeng Zhou, and Huanqian Yan. Boosting jailbreak transferability for large language models. *arXiv preprint arXiv:2410.15645*, 2024. 2
- [17] Hanqing Liu, Shouwei Ruan, Yao Huang, Shiji Zhao, and Xingxing Wei. When lighting deceives: Exposing vision-language models’ illumination vulnerability through illumination transformation attack. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10485–10495, 2025. 2
- [18] Jiayang Liu, Siyuan Liang, Shiqian Zhao, Rongcheng Tu, Wenbo Zhou, Xiaochun Cao, Dacheng Tao, and Siew Kei Lam. Jailbreaking the text-to-video generative models. *arXiv preprint*, 2025. 3, 4
- [19] Renyang Liu, Wei Zhou, Tianwei Zhang, Kangjie Chen, Jun Zhao, and Kwok-Yan Lam. Boosting black-box attack to deep neural networks with conditional diffusion models. *IEEE Transactions on Information Forensics and Security*, 19:5207–5219, 2024. 2
- [20] Renyang Liu, Wenjie Feng, Tianwei Zhang, Wei Zhou, Xueqi Cheng, and See-Kiong Ng. Rethinking machine unlearning in image generation models. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security*, pages 993–1007, 2025. 2
- [21] Xin Liu, Yichen Zhu, Yunshi Lan, Chao Yang, and Yu Qiao. Query-relevant images jailbreak large multi-modal models, 2023. 6
- [22] Xuannan Liu, Xing Cui, Peipei Li, Zekun Li, Huaibo Huang, Shuhan Xia, Miaoxuan Zhang, Yueying Zou, and Ran He. Jailbreak attacks and defenses against multimodal generative models: A survey. *arXiv preprint*, 2024. 2
- [23] Weidi Luo, Siyuan Ma, Xiaogeng Liu, Xiaoyu Guo, and Chaowei Xiao. Jailbreakv-28k: A benchmark for assessing the robustness of multimodal large language models against jailbreak attacks, 2024. 6
- [24] Hui Meng, Haochen Zhao, Deqian Yang, Songping Wang, and Zhenpeng Li. Coarse to fine segmentation method enables accurate and efficient segmentation of organs and tumor in abdominal ct. In *MICCAI Challenge on Fast and*

- Low-Resource Semi-supervised Abdominal Organ Segmentation*, pages 115–129. Springer, 2023. 2
- [25] Yibo Miao, Yifan Zhu, Yinpeng Dong, Lijia Yu, Jun Zhu, and Xiao-Shan Gao. T2vsafetybench: Evaluating the safety of text-to-video generative models. *arXiv preprint*, 2024. 2, 3, 4
- [26] Haomiao Ni, Changhao Shi, Kai Li, Sharon X. Huang, and Martin Renqiang Min. Conditional image-to-video generation with latent flow diffusion models. In *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 18444–18455, 2023. 2
- [27] Susanna Paasonen, Kylie Jarrett, and Ben Light. *NSFW: Sex, humor, and risk in social media*. Mit Press, 2019. 7
- [28] Xiangyu Peng, Zangwei Zheng, Chenhui Shen, Tom Young, Xinying Guo, Binluo Wang, Hang Xu, Hongxin Liu, Mingyan Jiang, Wenjun Li, et al. Open-Sora 2.0: Training a commercial-level video generation model in \$200k. *arXiv preprint*, 2025. 6
- [29] Yiting Qu, Xinyue Shen, Yixin Wu, Michael Backes, Savvas Zannettou, and Yang Zhang. Unsafebench: Benchmarking image safety classifiers on real-world and AI-generated images. *arXiv preprint*, 2024. 2
- [30] Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025. 6
- [31] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint*, 2025. 6
- [32] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024. 6
- [33] Songping Wang, Hanqing Liu, and Haochen Zhao. Public-domain locator for boosting attack transferability on videos. In *2024 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2024. 2
- [34] Songping Wang, Haoxiang Rao, Xiantao Hu, Yueming Lyu, and Caifeng Shan. An effective end-to-end solution for multimodal action recognition. In *International Conference on Pattern Recognition*, pages 324–338. Springer, 2024. 2
- [35] Songping Wang, Hanqing Liu, Yueming Lyu, Xiantao Hu, Ziwen He, Wei Wang, Caifeng Shan, and Liang Wang. Fast adversarial training with weak-to-strong spatial-temporal consistency in the frequency domain on videos. *IEEE Transactions on Information Forensics and Security*, 21:681–696, 2025. 2
- [36] Songping Wang, Yueming Lyu, Shiqi Liu, Ning Li, Tong Tong, Hao Sun, and Caifeng Shan. Anti-aesthetics: Protecting facial privacy against customized text-to-image synthesis. *arXiv preprint arXiv:2504.12129*, 2025. 2
- [37] Songping Wang, Xinquan Yue, Yueming Lyu, and Caifeng Shan. Exploring adversarial transferability between kolmogorov-arnold networks. *arXiv preprint arXiv:2503.06276*, 2025. 2
- [38] Songping Wang, Qinglong Liu, Yueming Lyu, Ning Li, Ziwen He, and Caifeng Shan. Exposing and defending the achilles’ heel of video mixture-of-experts. *arXiv preprint arXiv:2602.01369*, 2026. 2
- [39] Wenhao Wang and Yi Yang. Vidprom: A million-scale real prompt-gallery dataset for text-to-video diffusion models. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 65618–65642, 2024. 2
- [40] Xingxing Wei, Songping Wang, and Huanqian Yan. Efficient robustness assessment via adversarial spatial-temporal focus on videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):10898–10912, 2023. 2
- [41] Yi Wu, Ziqiang Li, Chaoyue Wang, Heliang Zheng, Shanshan Zhao, Bin Li, and Dacheng Tao. Domain re-modulation for few-shot generative domain adaptation. *Advances in Neural Information Processing Systems*, 36:57099–57124, 2023. 2
- [42] Yi Wu, Ziqiang Li, Heliang Zheng, Chaoyue Wang, and Bin Li. Infinite-id: Identity-preserved personalization via id-semantic decoupling paradigm. In *European Conference on Computer Vision*, pages 279–296. Springer, 2024. 2
- [43] Jinbo Xing, Menghan Xia, Yong Zhang, Haoxin Chen, Wangbo Yu, Hanyuan Liu, Gongye Liu, Xintao Wang, Ying Shan, and Tien-Tsin Wong. Dynamicrafter: Animating open-domain images with video diffusion priors. In *European Conf. on Computer Vision (ECCV)*, pages 399–417. Springer, 2024. 6
- [44] Yuchen Yang, Bo Hui, Haolin Yuan, Neil Gong, and Yinzhi Cao. Sneakyprompt: Jailbreaking text-to-image generative models. In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 897–912. IEEE, 2024. 2, 6
- [45] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint*, 2024. 6
- [46] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. 6
- [47] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 7
- [48] Chen Zhao, Weiling Cai, Chengwei Hu, and Zheng Yuan. Cycle contrastive adversarial learning with structural consistency for unsupervised high-quality image deraining transformer. *Neural Networks*, 178:106428, 2024. 2
- [49] Chen Zhao, En Ci, Yunzhe Xu, Tiehan Fan, Shanyan Guan, Yanhao Ge, Jian Yang, and Ying Tai. Ultrahr-100k: Enhancing uhr image synthesis with a large-scale high-quality dataset. *Advances in Neural Information Processing Systems*, 2025.
- [50] Zhiqin Zhou, Songping Wang, and Caifeng Shan. Multimodal video-based heart rate estimation with temporal difference transformer. 2025. 2