FOMO: FORGET THE CONCEPT, DON'T MISS OUT ON THE SCENE IN SELECTIVE VIDEO UNLEARNING

Anonymous authors
Paper under double-blind review

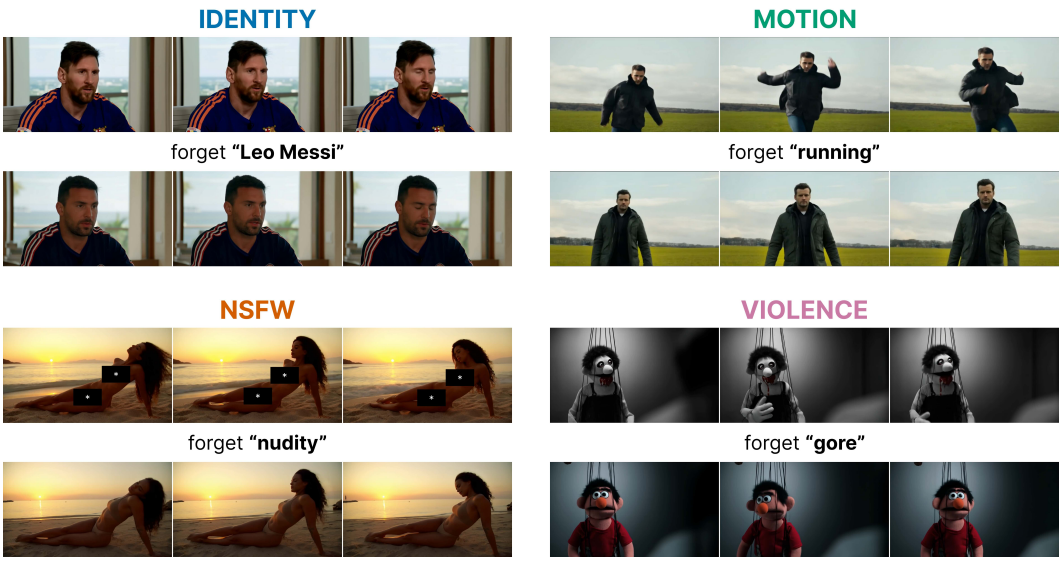

Figure 1: **We present FOMO: a selective video unlearning method that erases only the unwanted concept, while keeping scene composition and dynamics close to the original.** It removes harmful concepts like nudity, the likeness of public figures, and disturbing or violent content, and was specifically designed to handle motion unlearning, where the concept is not an object in the frame but an action carried across the whole clip.

ABSTRACT

The rapid advancement of generative video models has enabled the synthesis of increasingly realistic and temporally coherent videos, while also raising concerns about the generation of harmful content. The reliance on large-scale web datasets during training inevitably exposes these models to undesirable material, making concept unlearning an essential mitigation. Existing methods mainly target static visual concepts, such as objects, identities, or unsafe appearance, largely overlooking motion unlearning. Furthermore, these approaches often pay little attention to preserving the surrounding scene. As a result, successful concept removal may unintentionally alter the background, composition, or overall video dynamics. We argue that effective unlearning should ideally change only what is targeted, while minimizing unnecessary changes to the remaining scene. In this work, we introduce FOMO, to the best of our knowledge the first training-based selective video unlearning method that directly treats preservation of the original scene as a priority. We formulate unlearning around two complementary objectives: what to change and what to preserve. Our method localizes concept-related representations and modifies them, while the preservation mechanism maintains non-target scene information without requiring auxiliary data. Beyond simply erasing unwanted concepts, FOMO explicitly redirects the generation toward a specified safe alternative. We further extend this formulation to motion unlearning, where the concept is defined by temporal behavior rather than a fixed spatial region. Our solution achieves effective unlearning across unsafe content, object, and motion concepts, while achieving the best trade-off between concept removal and scene preservation. **Anonymous page with videos**

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1 INTRODUCTION

Modern generative models (Kong et al., 2025; Wan et al., 2025; Yang et al., 2025; Zheng et al., 2024) can increasingly produce high-quality, complex videos. These models can also produce highly realistic yet harmful content, which can spread through social media and evade moderation, exposing young audiences to disturbing or age-inappropriate imagery (Schramowski et al., 2023). Beyond explicit content, these models can also be used to create fake videos of public figures or videos featuring copyrighted characters and brands without permission (Malarz et al., 2026).

Machine unlearning provides a principled way to remove unwanted knowledge from models without costly retraining. Ideally, the target concept should be removed while the model's other capabilities remain unchanged. Concept unlearning is inherently more challenging in videos than in images, where the target is primarily spatial, whereas videos additionally encode temporal structure and motion. Even subtle changes to the generation trajectory can propagate over time and lead to entirely different scene evolution. Current approaches mainly follow two paths: controlling at inference time or modifying the model weights. SAFREE (Yoon et al., 2025) and VideoEraser (Xu et al., 2025) modify the input text embeddings and guide the denoising trajectory away from unlearned concepts without changing the model weights. In contrast, T2VUnlearning (Ye et al., 2025), Low-Rank Refusal Vector (Facchiano et al., 2026), and EraseAnything++ (Gao et al., 2025; Fan et al., 2026) modify the model itself through fine-tuning or direct weight updates.

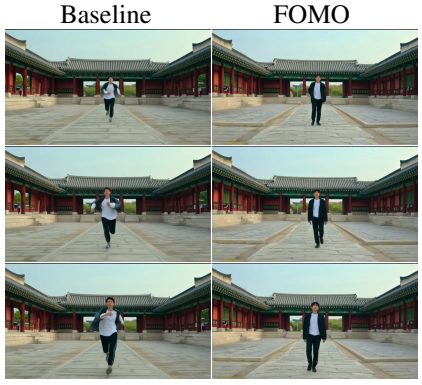

Figure 2: **Running motion unlearning on HunyuanVideo.** FOMO successfully removes the motion while largely preserving the character's overall appearance and scene.

An overlooked aspect of video unlearning is scene preservation. Effective selective unlearning should remove or safely redirect the target concept while minimizing unnecessary changes to the rest of the video, as illustrated by FOMO in Fig. 1. However, a safe replacement may still require contextual changes when it does not fit the original scene. For motion unlearning, we aim to preserve all content beyond the unlearned motion (see Fig. 2), while for nudity unlearning, the same principle applies to all content beyond the clothing-related changes (see Fig. 3).

Motivated by this gap, we introduce FOMO, the first video unlearning method that explicitly places scene preservation at the center of the unlearning objective. We organize our method around two core principles: **what to change** (targeted concept erasure) and **what to preserve** (compositional and semantic integrity), as illustrated in Fig. 4. Conceptually, FOMO follows the logic of film direction, where a director alters only the required element of a shot while minimizing changes to the remaining scene. To determine what should be changed, we first identify where the concept is represented using anchors in each frame and expand them into a spatial region associated with the concept. We then redirect these representations toward a desired safe alternative We further extend this formulation to motion-only unlearning, where we modify a specific movement, such as jumping or running, while aiming to preserve the surrounding scene and the subject's overall identity by intervening only in motion-sensitive attention heads. To preserve the scene, we introduce a complementary value-preservation loss that keeps the remaining scene consistent with the original video. We evaluate FOMO on HunyuanVideo across unsafe-content, object, public-figure, and motion unlearning, using GEN, Ring-A-Bell, SafeSora, T2VSafetyBench, and Imagenette. We measure scene preservation with DINO similarity and LPIPS, and overall video quality using VBench met-

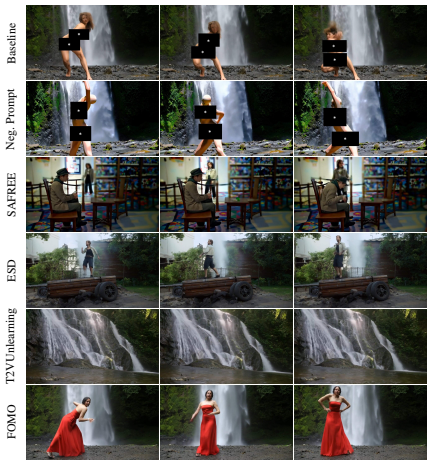

Figure 3: **Nudity unlearning on HunyuanVideo.** For "a naked woman dancing by the waterfall," FOMO removes nudity while preserving the dancing motion and waterfall scene.



Figure 4: **Visualizing the impact of FOMO's two objectives.** Erase may cause scene drift, while Preserve retains the original generation. Global preservation can weaken erasure by constraining the target region, whereas masked preservation balances concept removal and scene preservation.

rics. We conduct ablation studies to analyze our two objectives: concept erasure and scene preservation. Our results show effective concept and motion removal while maintaining high similarity to the original scene and strong video quality. We summarize our main contributions as follows:

- **Selective Video Erasure with Scene Preservation.** We introduce a video selective unlearning method designed to achieve strong target erasure while maintaining high surrounding scene preservation.
- **Motion Unlearning.** We extend unlearning beyond static visual concepts to motion, using motion-sensitive heads for behaviors and frame-wise anchors for visual targets.
- **Comprehensive Evaluation.** We conduct extensive experiments across a wide range of benchmarks dedicated to video unlearning, demonstrating that FOMO achieves the best trade-off between effective concept removal, scene preservation, and video quality.

2 Related Works

Early text-to-video models, such as ModelScopeT2V (Wang et al., 2023), LaVie (Wang et al., 2024), AnimateDiff (Guo et al., 2024), and ZeroScope (Cerspense, 2023) extended U-Net-based image diffusion architectures with temporal modeling. Recent models increasingly adopt Diffusion Transformers, including CogVideoX (Yang et al., 2025), Open Sora (Zheng et al., 2024), Hunyuan-Video (Kong et al., 2025) and Wan (Wan et al., 2025). Most concept unlearning methods target text-to-image models. ESD (Gandikota et al., 2023) fine-tunes the model away from an unlearned concept, MACE (Lu et al., 2024) combines closed-form cross-attention refinement with LoRA (Hu et al., 2022), and Receler (Huang et al., 2024a) uses cross-attention maps for spatial localization. UCE (Gandikota et al., 2024) and RECE (Gong et al., 2025) perform closed-form cross-attention updates, while AdvUnlearn (Zhang et al., 2024) uses adversarial training and auxiliary retain data to preserve generation utility. EraseAnything (Gao et al., 2025) extends concept erasure to rectified-flow Transformers. Recent work extends unlearning to text-to-video models. SAFREE (Yoon et al., 2025) and VideoEraser (Xu et al., 2025) operate at inference time, whereas T2VUnlearning (Ye et al., 2025), EraseAnything++ (Fan et al., 2026), NullSCE (Yi et al., 2026) and Low-Rank Refusal Vector (Facchiano et al., 2026) modify model through training or weight updates.

Existing video unlearning largely follows image-unlearning settings, focusing on visual targets such as objects and unsafe content, while motion and temporal behaviors remain considerably less explored. Human Motion Unlearning (De Matteis et al., 2026; Wang et al., 2026) studies motion removal in 3D motion sequences, where appearance and surrounding scene context fall outside the scope of the generative model. A second, broader limitation is that scene preservation remains insuf-

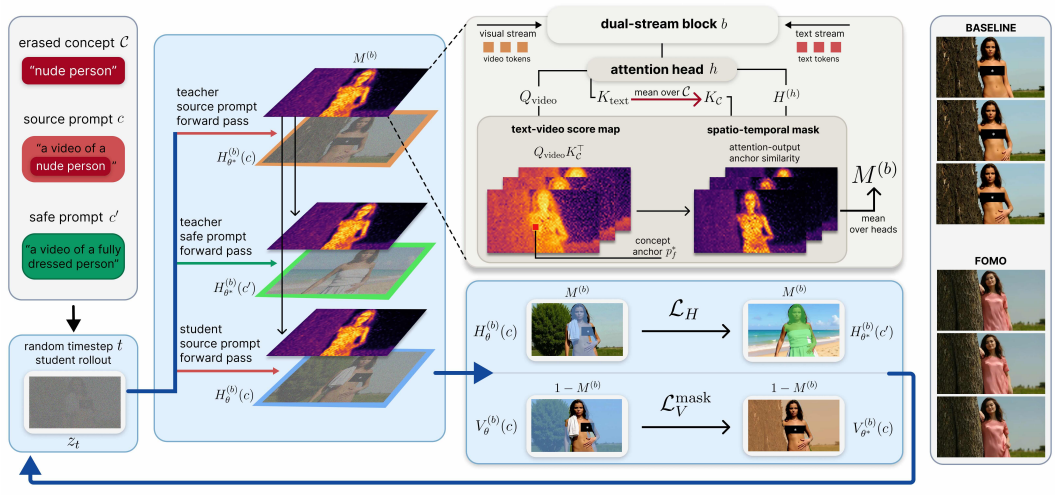

Figure 5: **Overview of FOMO.** The source prompt defines the erased concept, while the safe prompt provides its counterpart. During the training ($\rightarrow$), FOMO localizes the target at z_t by expanding a concept anchor p_f^* to similar attention-output regions in each frame. $\mathcal{L}_H$ aligns the target representations toward the safe concept, while $\mathcal{L}_V^{\text{mask}}$ preserves the rest of the scene. After training, FOMO safely replaces nudity with a pink dress while largely preserving the person and scene.

ficiently explored as a primary objective across both motion and video unlearning. In contrast, our work jointly addresses both gaps by introducing highly scene-preserving unlearning across visual concepts and motion in generative video models.

3 Methodology

In this section, we introduce our method FOMO for video unlearning. Section 3.1 investigates **where** the target concept is represented, Section 3.2 determines **how** to erase it, Section 3.3 identifies **what** should to preserve, and Section 3.4 adapts the method to motion-selective unlearning.

3.1 **Where** is the concept present?

Preliminaries. Video Diffusion Transformer (Video DiT) models (Kong et al., 2025; Yang et al., 2025; Wan et al., 2025) use transformer blocks that process text tokens derived from a prompt c and spatio-temporal video tokens through multi-head attention. For each attention head h, attention output is defined as $H^{(h)} = \text{softmax}\big(Q^{(h)}K^{(h)\top}/\sqrt{d_k}\big)\,V^{(h)}$, where $Q^{(h)}$ (Query), $K^{(h)}$ (Key), and $V^{(h)}$ (Value) are token projection, and d_k is the key dimension. Individual attention heads operate in different feature subspaces, and their outputs are concatenated along the feature dimension. For joint multimodal attention blocks, we distinguish the projected video representations $(Q_{\text{video}}, K_{\text{video}}, V_{\text{video}})$ from the projected text representations $(Q_{\text{text}}, K_{\text{text}}, V_{\text{text}})$. The attention scores comprise the video-video ($Q_{\text{video}}K_{\text{video}}^\top$), video-text ($Q_{\text{video}}K_{\text{text}}^\top$), text-video ($Q_{\text{text}}K_{\text{video}}^\top$), and text-text ($Q_{\text{text}}K_{\text{text}}^\top$) interactions. HunyuanVideo (Kong et al., 2025) follows a dual-stream $\rightarrow$ single-stream architecture, in which early blocks use modality-specific projections before joint attention, whereas later blocks process a shared video-text token sequence.

Spatio-Temporal Localization Mask. Previous unlearning methods (Lu et al., 2024; Gandikota et al., 2024; Wang et al., 2025b; Gao et al., 2025; Ye et al., 2025) often localize the target concept using cross-modal attention maps derived from $Q_{\text{video}}K_{\text{text}}^\top$. After softmax normalization over text-token positions, weights associated with the target phrase are suppressed during unlearning. However, this strategy evaluates each video token with respect to the target concept, without modeling spatial relationships or representational similarities among video tokens within the frame.

Recent work on video interpretability uses similarities among attention-output video representations to construct spatially coherent concept maps across frames (Jun et al., 2026). We investi-

gate whether such localization can move beyond post-hoc interpretation and directly guide selective unlearning through a single mask by combining a concept anchor in each frame with similarities in representation space H. To obtain this mask, we define the localization procedure as follows. For a source prompt $c = (c_1, \dots, c_T)$, where T is the number of text tokens, we let $\mathcal{C} \subseteq \{c_1, \dots, c_T\}$ denote the tokens corresponding to the target concept, and c' denote the corresponding safe prompt. Let $\mathcal{P} = \{1, \dots, P\}$ denote the set of video-token positions within a frame. We omit the block index b and head index h in the following derivation, as the localization procedure is applied independently to each block and head. Thus, for concepts spanning multiple text tokens, we average the corresponding key representations: $K_{\mathcal{C}} = \frac{1}{|\mathcal{C}|} \sum_{c_i \in \mathcal{C}} K_{\text{text}}(c_i)$. For each frame f, we identify the video-token anchor most strongly associated with the target concept as

$$p_f^* = \arg\max_{p \in \mathcal{P}} \frac{Q_{\text{video}}(f, p) K_{\mathcal{C}}^\top}{\sqrt{d_k}}. \tag{1}$$

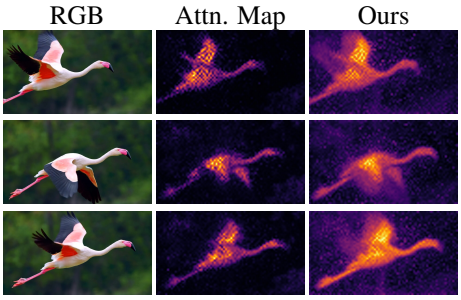

Figure 6: **Concept localization for *flamingo*.** The concept attention map concentrates on discriminative regions, whereas our mask expands the selected anchor to representationally similar video regions.

A single anchor cannot capture the full target extent. We therefore compute pairwise similarities between video attention representations H_f using the Gram matrix $G_f = H_f H_f^\top$. Then, we extract the column indexed by the anchor p_f^* to obtain the similarity map $S_f(p) = G_f(p, p_f^*)$, for $p \in \mathcal{P}$. The similarity map S_f is min-max normalized over spatial video-token positions, yielding $S'_f(p) \in [0, 1]$. The resulting map propagates the concept anchor to representationally similar regions within the frame (see Fig. 5). To construct the block-specific localization mask, we aggregate information across all attention heads. We therefore reintroduce the block and head indices to formalize this operation $M_f^{(b)} = \frac{1}{|\mathcal{H}|} \sum_{h \in \mathcal{H}} S'^{(b,h)}_f$, where $\mathcal{H}$ denotes the set of heads in transformer block b. Stacking these masks across frames yields the spatio-temporal mask $M^{(b)}$. Fig. 6 compares a prior cross-modal attention map with our H-based anchor-similarity map for the *flamingo*.

3.2 How to unlearn the concept?

To modify representations associated with the concept to be unlearned, we follow the teacher-student setup (Gandikota et al., 2023; Polowczyk et al., 2026; Wójcik et al., 2026; Ye et al., 2025; Gao et al., 2025), using the frozen base model as the teacher and its trainable copy as the student. The safe prompt c' specifies the desired direction of unlearning. Unlike standard approaches that impose a global target in the noise-prediction space, we construct a spatiotemporal target directly in the internal DiT representations.

Localized Safe Representation. At each training step, we sample an initial latent state $z_T \sim \mathcal{N}(0, I)$ and obtain an intermediate latent state z_t by partially rolling out the current student with parameters θ under the source prompt c. At the same latent state z_t, we perform two forward passes through the frozen model with parameters θ^*, conditioned on c and c', respectively. For each selected transformer block b, this yields the source and safe representations $H_{\theta^*}^{(b)}(c)$ and $H_{\theta^*}^{(b)}(c')$. The source pass also provides the localization mask $M^{(b)}$ defined in Section 3.1. For each block, the mask is computed once from the unsafe concept and reused for the other representations. We then construct a localized safe representation as:

$$H_*^{(b)} = H_{\theta^*}^{(b)}(c) + \alpha M^{(b)} \odot \left(H_{\theta^*}^{(b)}(c') - H_{\theta^*}^{(b)}(c) \right). \tag{2}$$

The mask restricts the representation update to the region associated with the target concept, while $\alpha \geq 1$ controls its magnitude of the shift toward the safe representation. For $\alpha \in [0, 1)$, the mask is attenuated, leading to weaker unlearning.

LoRA-Based Representation Alignment. To align the student representations with the localized safe representation, we optimize only the LoRA parameters $\Delta\theta_{\text{LoRA}}$ using the sub-loss $\mathcal{L}_H$, while keeping the base parameters θ^* frozen. The resulting model $\theta = \theta^* + \Delta\theta_{\text{LoRA}}$ receives the same

z_t and source prompt c, and aligns $H_\theta^{(b)}(c)$ with the localized safe representation $H_*^{(b)}$ at the set of selected transformer blocks $\mathcal{B}$:

$$\mathcal{L}_H = \frac{1}{|\mathcal{B}|} \sum_{b \in \mathcal{B}} \left[\left\| H_\theta^{(b)}(c) - H_*^{(b)} \right\|_2^2 \right]. \tag{3}$$

For details on block selection and LoRA placement for HunyuanVideo, please refer to Appendix A.

3.3 WHAT SHOULD BE PRESERVED?

Preservation without Auxiliary Data. Preserving knowledge beyond the target concept remains a key challenge in unlearning. Existing methods often rely on auxiliary preservation signals derived from external retain sets (Lu et al., 2024; Sun et al., 2026), explicitly defined preservation concepts (Ye et al., 2025), or additional reference samples (Gao et al., 2025). This motivates the following question: ***Can non-target knowledge be preserved without additional data or prompts?*** We hypothesize that keeping non-target scene representations close to their counterparts in the base model reduces representational drift during unlearning and helps preserve the model's broader knowledge.

Value-Based Preservation. Modifying the concept representation should leave information unrelated to the target concept as unchanged as possible, including the scene structure, background, and video dynamics. Based on the observation that value representations V can capture information about scene structure (Wang et al., 2025a), we introduce a preservation term, $\mathcal{L}_V^{\text{mask}}$:

$$\mathcal{L}_V^{\text{mask}} = \frac{1}{|\mathcal{K}|} \sum_{b \in \mathcal{K}} \left[\left\| (1 - M^{(b)}) \odot \left(V_\theta^{(b)}(c) - V_{\theta^*}^{(b)}(c) \right) \right\|_2^2 \right], \tag{4}$$

where $\mathcal{K}$ denotes the set of selected blocks, $(1 - M^{(b)})$ restricts the loss to non-target regions. The loss $\mathcal{L}_V^{\text{mask}}$ is applied only to training samples with timesteps in the high-noise range to encourage preservation of the global scene structure rather than fine-grained details.

The final training objective is $\mathcal{L} = \mathcal{L}_H + \lambda_V \mathcal{L}_V^{\text{mask}}$ where λ_V controls the strength of the preservation objective. An overview of how this final objective operate within the training pipeline is shown in Fig. 5. Global ($\mathcal{L}_V^{\text{full}}$) versus masked application of $\mathcal{L}_V^{\text{mask}}$ is analyzed in Sec. 5. Details on the selected blocks and the timestep range are provided in Appendix A.

3.4 MOTION-SELECTIVE UNLEARNING

Motion unlearning should modify the targeted behavior while preserving the remaining visual content as much as possible. In multi-head attention, individual heads operate in different feature subspaces and may emphasize different spatial and temporal cues. Consequently, aggregating all heads can mix motion-relevant information with static appearance features, producing broad masks that extend beyond regions involved in the target motion (see Fig. 7).

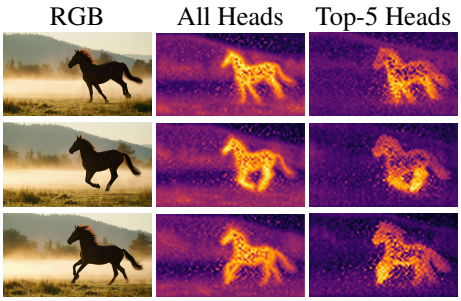

Figure 7: **Motion localization for *running*.** Visualization using the top-5 motion-sensitive heads, which yields a more focused maps around motion-relevant regions.

Temporal Head Selection. To obtain more motion-specific localization, we restrict the procedure to attention heads that exhibit strong frame-wise variation. Following (Jun et al., 2026), we rank the heads in each selected transformer block using the Calinski–Harabasz index (CHI) (Caliński & Harabasz, 1974). For each head, video-token attention outputs from individual frames are treated as separate clusters, and CHI measures their between-frame separation relative to within-frame dispersion. Higher CHI scores therefore indicate that a head produces more distinct representations across frames, suggesting greater sensitivity to temporal changes. Accordingly, for each selected transformer block, we retain the top-k heads $\hat{\mathcal{H}}^{(b)} = \text{TopK}_{h \in \mathcal{H}}\left(\text{CHI}^{(b,h)}\right)$. We then aggregate the normalized similarity maps only over the selected motion-sensitive heads $M_{f,\text{motion}}^{(b)} = \frac{1}{|\hat{\mathcal{H}}^{(b)}|} \sum_{h \in \hat{\mathcal{H}}^{(b)}} S'^{(b,h)}_f$. The frame-wise masks are stacked along the temporal dimension to obtain $M^{(b)}$, which is directly used in Eqs. 2 and 4. The $\mathcal{L}_H$ term is evaluated only on the outputs of heads $h \in \hat{\mathcal{H}}^{(b)}$.

Table 1: **Average nudity-unlearning results for GEN, Ring-A-Bell, SafeSora and T2VSafetyBench.** FOMO provides the best trade-off between unsafe-content removal and scene preservation. Six VBench custom-input dimensions are computed on the same generations. Best results are shown in bold and second-best results are underlined.

	Unlearning	Preservation		Video Quality					
Method	**Unsafe Rate↓**	**DINO↑**	**LPIPS↓**	**Subject↑**	**Background↑**	**Smoothness↑**	**Dynamic↑**	**Aesthetic↑**	**Imaging↑**
HunyuanVideo	53.14	–	–	<u>96.02</u>	95.95	99.46	48.45	**58.12**	59.60
Negative Prompt	55.79	<u>0.5911</u>	<u>0.4949</u>	93.22	94.90	99.05	**60.99**	53.04	57.54
SAFREE	26.11	0.4989	0.5159	96.67	**96.17**	**99.52**	43.80	<u>57.91</u>	<u>59.80</u>
ESD	**1.76**	0.1956	0.6527	**96.69**	95.25	99.39	36.55	54.61	**64.79**
T2VUnlearning	22.98	0.3884	0.6435	94.55	94.95	99.32	46.62	52.70	56.62
FOMO	<u>19.54</u>	**0.6458**	**0.3802**	95.85	<u>95.58</u>	<u>99.47</u>	<u>58.06</u>	55.82	59.69

Table 2: **General safety unlearning across five safety categories.** FOMO consistently reduces unsafe generations while preserving higher similarity to the original scene.

	Unsafe Rate↓			DINO↑		LPIPS↓	
Category	**Hunyuan**	**SAFREE**	**FOMO**	**SAFREE**	**FOMO**	**SAFREE**	**FOMO**
Violence	59.04	47.59	**7.83**	0.4472	**0.5056**	0.5111	**0.4778**
Animal Abuse	59.26	37.04	**7.41**	0.3819	**0.4740**	0.5229	**0.4586**
Terrorism	52.00	36.00	**4.00**	0.4740	**0.5032**	0.5246	**0.5000**
Racism	57.78	42.22	**28.89**	0.4596	**0.5888**	0.5351	**0.4804**
Gore	67.21	60.66	**14.75**	0.5108	**0.5551**	0.4361	**0.4063**
Average	59.06	44.70	**12.58**	0.4547	**0.5253**	0.5060	**0.4646**

4 EXPERIMENTS

This section presents our experimental protocol. We evaluate FOMO on HunyuanVideo across unsafe-content, object, public-figure and motion unlearning using GEN, Ring-A-Bell (Ye et al., 2025), SafeSora (Dai et al., 2024), T2VSafetyBench (Miao et al., 2024), and Imagenette (Deng et al., 2009) benchmarks. For each unlearning task, we define a source-safe prompt pair. We compare our method with the current techniques for video unlearning: NegativePrompt, SAFREE (Yoon et al., 2025), which are based on inference-time, T2VUnlearning (Ye et al., 2025) and adaptive ESD to the video domain (Gandikota et al., 2023).

Evaluation metrics. Following the evaluation protocol of T2VSafetyBench (Miao et al., 2024), we adopt the benchmark's original evaluation prompt and use a multimodal video evaluator to assess the presence of unsafe concepts in generated videos (see Appendix B for the evaluation prompt). We employ Qwen3-VL-32B-Instruct (Bai et al., 2025), a strong open-source video understanding model whose effectiveness for harmful content recognition is supported by HarmVideoBench (Wu et al., 2026). We report **Unsafe Rate** (↓) across GEN, Ring-A-Bell, SafeSora, and T2VSafetyBench and **Target Motion Rate** (↓) for motion unlearning. For object unlearning, we report **Erasure Success Rate** (ESR-k (↑)) and **Preservation Success Rate** (PSR-k (↑)), following T2VUnlearning (Ye et al., 2025) using a ResNet-50 classifier (He et al., 2016). For public figure erasure, we evaluate identity similarity using ArcFace embeddings (Deng et al., 2019). To measure scene similarity, we compare unlearned and baseline generations using LPIPS (↓) and DINO similarity (↑). We additionally report VBench metrics (Huang et al., 2024b) for overall video quality. Metric details in Appendix B.

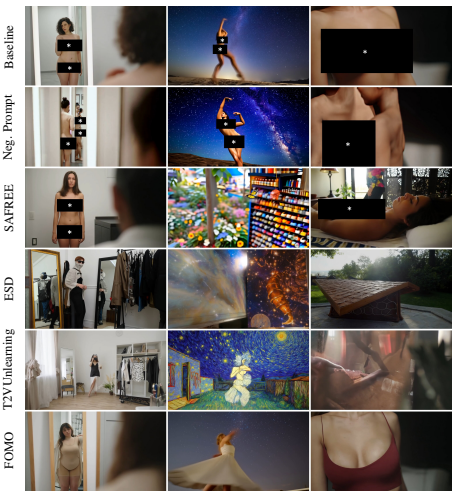

Figure 8: **Nudity unlearning across GEN, Ring-A-Bell, and T2VSafetyBench.** FOMO replaces nudity with plausible clothing while largely preserving subject and scene content.



(a) Dancing (b) Sticking out the tongue

Figure 9: **Qualitative motion-unlearning results.** Comparison between baseline generations and FOMO across two target motions.

General Safety Unlearning. We evaluate nudity unlearning across GEN, Ring-A-Bell, SafeSora, and T2VSafetyBench (Table 1 and Fig. 8). FOMO reduces the Unsafe Rate from 53.14% to 19.54% while achieving the best scene preservation among unlearning methods. ESD achieves stronger suppression (1.76%) but substantially lower scene similarity. Selected VBench metrics further confirm preserved video dynamics and visual consistency, with broader generative capabilities are evaluated on 946 diverse VBench prompts in Sec. 5. We evaluate violence, animal abuse, terrorism, racism on SafeSora, and gore on T2VSafetyBench (Table 2). FOMO reduces the average Unsafe Rate from 59.06% to 12.58% while preserving higher scene similarity than SAFREE. Additional quantitative results and qualitative examples covering different types of unsafe content are provided in Appendices C and I.

Table 3: **Object-erasure results averaged across the Imagenette dataset.**

Method	ESR-1↑	ESR-5↑	PSR-1↑	PSR-5↑
HunyuanVideo	28.88	11.24	**71.12**	**88.76**
Negative Prompt	38.00	13.94	67.88	87.36
SAFREE	64.03	45.00	35.97	55.00
FOMO	**92.41**	**77.24**	66.62	85.94

Object Unlearning. We evaluate object erasing on Imagenette. We use LLM-refined evaluation prompts describing each class together with its characteristic visual attributes, rather than simple class-name prompts (see evaluation prompts in Appendix B). Quantitative results in Table 3 show that FOMO achieves substantially higher erasure than Negative Prompt and SAFREE, reaching 92.41 ESR-1 and 77.24 ESR-5, while maintaining preservation close to the base model. Detailed quantitative results are in Appendix D. We further evaluate multiple mapping concepts at different semantic distances, e.g., English Springer → German Shepherd or cat, and parachute → balloon or airplane, for which FOMO also achieves effective unlearning. A detailed analysis of the effect of different safe prompt choices is provided in Appendix G.

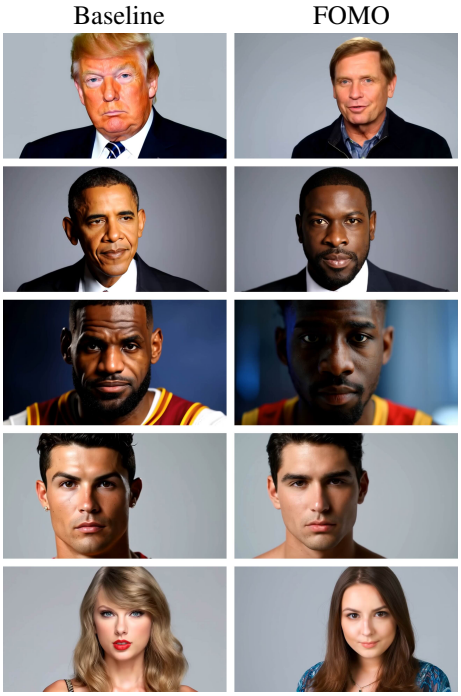

Figure 10: **Public figure unlearning across five identities.** FOMO replaces the erased identity with a different person while largely preserving a similar appearance, keeping the pose and background.

Public Figures Unlearning. Public figure unlearning requires suppressing identity-specific facial characteristics while largely preserving non-target attributes such as pose, and background. We select five identities from T2VSafetyBench: Donald Trump, Barack Obama, LeBron James, Cristiano Ronaldo, and Taylor Swift. Across all identities, FOMO reduces the average ID similarity from 0.5793 in the original model to 0.0822 after unlearning, while retaining an average similarity of 0.5000 for the remaining identities (see the full Table in Appendix E). As shown in Fig. 10, FOMO effectively removes identity-specific features while largely preserving the surrounding visual content. Qualitative results for different protected visual concepts, including public figures and brands are included in the Appendix I.

Table 5: **Motion unlearning on HunyuanVideo.** Three representative motions are shown, and the average is computed over all six motions.

		Unlearning	Quality	Preservation	
Motion	**Method**	**Motion Rate↓**	**Smooth.↑**	**DINO↑**	**LPIPS↓**
Running	Base	100.00	0.9877	–	–
	FOMO	**20.00**	**0.9890**	0.7363	0.3179
Fighting	Base	95.00	0.9841	–	–
	FOMO	**0.00**	**0.9928**	0.6044	0.4964
Tongue	Base	100.00	**0.9958**	–	–
	FOMO	**0.00**	0.9954	0.7803	0.2427
Avg. (6)	Base	96.67	0.9911	–	–
	FOMO	**12.50**	**0.9939**	0.7013	0.3467

Table 6: **Ablation of two objectives of FOMO: $L_H + L_V$ on GEN dataset.**

Method / Variant	**Unsafe Rate↓**	**DINO↑**	**LPIPS↓**
HunyuanVideo	80.00	-	-
Negative Prompt	82.00	0.6196	0.4543
SAFREE	40.00	0.4937	0.5306
ESD	**4.00**	0.2376	0.6214
T2VUnlearning	30.00	0.4130	0.5841
L_H	**4.00**	0.5412	0.4199
L_V^{mask}	80.00	**1.0000**	**0.0000**
$L_H + L_V^{\text{full}}$	44.00	0.6867	0.3384
FOMO: $L_H + L_V^{\text{mask}}$	26.00	0.6832	0.3546

Table 4: **Broader generative preservation after nudity unlearning on VBench.**

Method	**Quality↑**	**Semantic↑**	**Final↑**	**DINO↑**	**LPIPS↓**
HunyuanVideo	**83.37**	71.77	81.05	-	-
Negative Prompt	83.03	**78.97**	**82.22**	0.606	0.496
SAFREE	81.49	45.06	74.21	0.391	0.540
ESD	80.79	41.96	73.02	0.298	0.602
T2VUnlearning	82.46	67.26	79.42	0.718	0.305
FOMO	83.29	69.28	80.04	**0.739**	**0.300**

Motion Unlearning. We evaluate six dynamic concepts: running, walking, jumping, dancing, fighting, and sticking out the tongue, using LLM-refined prompts that explicitly describe each behavior. As shown in Table 5, FOMO reduces the average Motion Rate from 96.67% to 12.50% while maintaining high temporal smoothness and scene similarity (The full table is provided in Appendix F). For each concept, we define a source-to-safe pair, e.g., dancing → standing for motion and sticking out the tongue → keeping the mouth closed for gesture. Additional motion-unlearning examples, including action and social interaction such as fighting → having a friendly conversation, are shown in Appendix I.

5 Ablation Study

Impact of $\mathcal{L}_H$ and $\mathcal{L}_V$. We ablate $\mathcal{L}_H$ and $\mathcal{L}_V$ in Table 6. Fig. 4 provides the corresponding qualitative comparison, where Erase, Preserve, Erase + Global Preserve, and Erase + Masked Preserve correspond to $\mathcal{L}_H$, $\mathcal{L}_V^{\text{mask}}$, $\mathcal{L}_H + \mathcal{L}_V^{\text{full}}$, and $\mathcal{L}_H + \mathcal{L}_V^{\text{mask}}$, respectively. Notably, $\mathcal{L}_H$ alone achieves ESD-level erasure (4% Unsafe Rate) with substantially better scene preservation, while outperforming SAFREE and T2VUnlearning in both erasure and preservation. In contrast, $\mathcal{L}_V^{\text{mask}}$ alone leaves the generation unchanged, confirming its preservation role. Combining both terms yields the best erasure–preservation balance, with a 26% Unsafe Rate. Applying preservation over the entire video instead increases the Unsafe Rate to 44%, indicating that preserving the target region can counteract unlearning. Although $\mathcal{L}_V^{\text{mask}}$ encourages scene preservation, unlearning a concept in a complex scene may require small contextual changes. We discuss such cases in Appendix H.

Broader Knowledge Preservation. As shown in Table 4, FOMO remains close to the base model on VBench, with Final Score of 80.04. Negative Prompt achieves a higher Semantic Score, which is expected since it modifies generation only through inference-time conditioning away from unsafe content. Importantly, FOMO achieves the highest DINO similarity (0.739) and lowest LPIPS (0.300) among the unlearning methods. These results suggest that preserving non-target scene representations during unlearning helps keep the generated content close to that of the original model, consistent with our hypothesis.

6 Conclusion

We present FOMO, first selective video unlearning method that jointly addresses what to change and what to preserve. It modifies concept-related representations while using a masked value-preservation loss to limit changes to non-target scene content. For motion unlearning, we further focus the intervention on motion-sensitive attention heads. Experiments across unsafe content, objects, public figures, and motion demonstrate effective target suppression while largely retaining scene structure and video quality, yielding a strong trade-off between erasure and preservation. **Limitations.** FOMO may partially restructure the scene to produce a safe video when the target unsafe concept encompasses the overall situation and is harder to isolate with a mask.

Ethics Statement

Selective video unlearning can help reduce harmful or offensive content, unsafe behaviors, and unwanted motions in generated videos. FOMO suppresses these targets while minimizing changes to the rest of the scene and retaining the model's ability to generate other videos. The examples presented in this work are used solely for research and evaluation purposes.

Reproducibility Statement

To facilitate reproducibility, we provide implementation details, training settings, and evaluation protocols in Appendices A and B.

AI Use Statement

We used generative AI tools to refine evaluation prompts and improve the grammar and clarity of the manuscript.

References

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, Mei Li, Kaixin Li, Zicheng Lin, Junyang Lin, Xuejing Liu, Jiawei Liu, Chenglong Liu, Yang Liu, Dayiheng Liu, Shixuan Liu, Dunjie Lu, Ruilin Luo, Chenxu Lv, Rui Men, Lingchen Meng, Xuancheng Ren, Xingzhang Ren, Sibo Song, Yuchong Sun, Jun Tang, Jianhong Tu, Jianqiang Wan, Peng Wang, Pengfei Wang, Qiuyue Wang, Yuxuan Wang, Tianbao Xie, Yiheng Xu, Haiyang Xu, Jin Xu, Zhibo Yang, Mingkun Yang, Jianxin Yang, An Yang, Bowen Yu, Fei Zhang, Hang Zhang, Xi Zhang, Bo Zheng, Humen Zhong, Jingren Zhou, Fan Zhou, Jing Zhou, Yuanzhi Zhu, and Ke Zhu. Qwen3-vl technical report, 2025. URL https://arxiv.org/abs/2511.21631.

T. Caliński and J. Harabasz. A dendrite method for cluster analysis. *Communications in Statistics*, 3(1):1–27, 1974. doi: 10.1080/03610927408827101.

Cerspense. Zeroscope v2 576w. https://huggingface.co/cerspense/zeroscope_v2_576w, 2023.

Josef Dai, Tianle Chen, Xuyao Wang, Ziran Yang, Taiye Chen, Jiaming Ji, and Yaodong Yang. Safesora: Towards safety alignment of text2video generation via a human preference dataset, 2024.

Edoardo De Matteis, Matteo Migliarini, Alessio Sampieri, Indro Spinelli, and Fabio Galasso. Human motion unlearning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 3533–3541, March 2026. doi: 10.1609/aaai.v40i5.37351. URL https://ojs.aaai.org/index.php/AAAI/article/view/37351.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009. doi: 10.1109/CVPR.2009.5206848.

Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

Simone Facchiano, Stefano Saravalle, Matteo Migliarini, Edoardo De Matteis, Alessio Sampieri, Andrea Pilzer, Emanuele Rodolà, Indro Spinelli, Luca Franco, and Fabio Galasso. Video unlearning via low-rank refusal vector. In *The Fourteenth International Conference on Learning Representations*, 2026.

Zhaoxin Fan, Nanxiang Jiang, Daiheng Gao, Shiji Zhou, and Wenjun Wu. Eraseanything++: Enabling concept erasure in rectified flow transformers leveraging multi-object optimization, 2026. URL https://arxiv.org/abs/2603.00978.

Rohit Gandikota, Joanna Materzyńska, Jaden Fiotto-Kaufman, and David Bau. Erasing concepts from diffusion models. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 2426–2436, 2023. doi: 10.1109/ICCV51070.2023.00230.

Rohit Gandikota, Hadas Orgad, Yonatan Belinkov, Joanna Materzyńska, and David Bau. Unified concept editing in diffusion models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2024. arXiv:2308.14761.

Daiheng Gao, Shilin Lu, Wenbo Zhou, Jiaming Chu, Jie Zhang, Mengxi Jia, Bang Zhang, Zhaoxin Fan, and Weiming Zhang. EraseAnything: Enabling concept erasure in rectified flow transformers. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 18470–18494. PMLR, 13–19 Jul 2025.

Chao Gong, Kai Chen, Zhipeng Wei, Jingjing Chen, and Yu-Gang Jiang. Reliable and efficient concept erasure of text-to-image diffusion models. In Aleš Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol (eds.), *Computer Vision – ECCV 2024*, pp. 73–88, Cham, 2025. Springer Nature Switzerland. ISBN 978-3-031-73668-1.

Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. In *The Twelfth International Conference on Learning Representations*, 2024.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 770–778. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.90.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*, 2022. URL https://mlanthology.org/iclr/2022/hu2022iclr-lora/.

Chi-Pin Huang, Kai-Po Chang, Chung-Ting Tsai, Yung-Hsuan Lai, Fu-En Yang, and Yu-Chiang Frank Wang. Receler: Reliable concept erasing of text-to-image diffusion models via lightweight erasers. In *European Conference on Computer Vision*, pp. 360–376. Springer, 2024a.

Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin, Nattapol Chanpaisit, Yaohui Wang, Xinyuan Chen, Limin Wang, Dahua Lin, Yu Qiao, and Ziwei Liu. VBench: Comprehensive benchmark suite for video generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024b.

Youngjun Jun, Seil Kang, Woojung Han, and Seong Jae Hwang. Interpretable motion-attentive maps: Spatio-temporally localizing concepts in video diffusion transformers, 2026. URL https://arxiv.org/abs/2603.02919.

Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, Kathrina Wu, Qin Lin, Junkun Yuan, Yanxin Long, Aladdin Wang, Andong Wang, Changlin Li, Duojun Huang, Fang Yang, Hao Tan, Hongmei Wang, Jacob Song, Jiawang Bai, Jianbing Wu, Jinbao Xue, Joey Wang, Kai Wang, Mengyang Liu, Pengyu Li, Shuai Li, Weiyan Wang, Wenqing Yu, Xinchi Deng, Yang Li, Yi Chen, Yutao Cui, Yuanbo Peng, Zhentao Yu, Zhiyu He, Zhiyong Xu, Zixiang Zhou, Zunnan Xu, Yangyu Tao, Qinglin Lu, Songtao Liu, Dax Zhou, Hongfa Wang, Yong Yang, Di Wang, Yuhong Liu, Jie Jiang, and Caesar Zhong. Hunyuanvideo: A systematic framework for large video generative models, 2025. URL https://arxiv.org/abs/2412.03603.

Shilin Lu, Zilan Wang, Leyang Li, Yanzhu Liu, and Adams Wai-Kin Kong. Mace: Mass concept erasure in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6430–6440, 2024.

Dawid Malarz, Filip Manjak, Maciej Zieba, Przemysław Spurek, and Artur Kasymov. From unlearning to unbranding: A benchmark for trademark-safe text-to-image generation, 2026. URL https://arxiv.org/abs/2512.13953.

Yibo Miao, Yifan Zhu, Lijia Yu, Jun Zhu, Xiao-Shan Gao, and Yinpeng Dong. T2VSafetybench: Evaluating the safety of text-to-video generative models. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024.

Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856.

Alicja Polowczyk, Agnieszka Polowczyk, Dawid Malarz, Artur Kasymov, Jacek Tabor, Marcin Mazur, and Przemysław Spurek. Unguide: Learning to forget with LoRA-guided diffusion models. In Emilija Perković and Daniel Malinsky (eds.), *Proceedings of the 42nd Conference on Uncertainty in Artificial Intelligence*, volume 337 of *Proceedings of Machine Learning Research*, pp. 5471–5503. PMLR, 17–21 Aug 2026.

Patrick Schramowski, Manuel Brack, Björn Deiseroth, and Kristian Kersting. Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 22522–22531, 2023. doi: 10.1109/CVPR52729.2023.02157.

Yuhao Sun, Lingyun Yu, Hao-Xiang Xu, Fengyuan Miao, Zhuoer Xu, and Hongtao Xie. Orthogonal concept erasure for diffusion models. In *Forty-third International Conference on Machine Learning*, 2026.

Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, Jiayu Wang, Jingfeng Zhang, Jingren Zhou, Jinkai Wang, Jixuan Chen, Kai Zhu, Kang Zhao, Keyu Yan, Lianghua Huang, Mengyang Feng, Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang, Tianxing Wang, Tianyi Gui, Tingyu Weng, Tong Shen, Wei Lin, Wei Wang, Wei Wang, Wenmeng Zhou, Wente Wang, Wenting Shen, Wenyuan Yu, Xianzhong Shi, Xiaoming Huang, Xin Xu, Yan Kou, Yangyu Lv, Yifei Li, Yijing Liu, Yiming Wang, Yingya Zhang, Yitong Huang, Yong Li, You Wu, Yu Liu, Yulin Pan, Yun Zheng, Yuntao Hong, Yupeng Shi, Yutong Feng, Zeyinzi Jiang, Zhen Han, Zhi-Fan Wu, and Ziyu Liu. Wan: Open and advanced large-scale video generative models, 2025. URL https://arxiv.org/abs/2503.20314.

Jiangshan Wang, Junfu Pu, Zhongang Qi, Jiayi Guo, Yue Ma, Nisha Huang, Yuxin Chen, Xiu Li, and Ying Shan. Taming rectified flow for inversion and editing. In *Forty-second International Conference on Machine Learning*, 2025a.

Jiuniu Wang, Hangjie Yuan, Dayou Chen, Yingya Zhang, Xiang Wang, and Shiwei Zhang. Modelscope text-to-video technical report, 2023. URL https://arxiv.org/abs/2308.06571.

Yaohui Wang, Xinyuan Chen, Xin Ma, Shangchen Zhou, Ziqi Huang, Yi Wang, Ceyuan Yang, Yinan He, Jiashuo Yu, Peiqing Yang, et al. Lavie: High-quality video generation with cascaded latent diffusion models. *IJCV*, 2024.

Yiling Wang, Zeyu Zhang, Yiran Wang, and Hao Tang. Safemo: Linguistically grounded unlearning for trustworthy text-to-motion generation, 2026. URL https://arxiv.org/abs/2601.00590.

Yuan Wang, Ouxiang Li, Tingting Mu, Yanbin Hao, Kuien Liu, Xiang Wang, and Xiangnan He. Precise, fast, and low-cost concept erasure in value space: Orthogonal complement matters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 28759–28768, June 2025b.

Piotr Wójcik, Maksym Petrenko, Wojciech Gromski, Przemysław Spurek, and Maciej Zieba. Unhype: CLIP-guided hypernetworks for dynamic loRA unlearning. In *Forty-third International Conference on Machine Learning*, 2026.

Jiajun Wu, Haoyu Kang, Yining Sun, Jiacheng Hou, Heng Zhang, Danyang Zhang, Zhenjun Zhao, Haochi Zhang, Leixin Sun, Eric Hanchen Jiang, Yushan Li, Ruiyu Li, Mengkai Huang, Yan Gao, Xu Zhang, and Guancheng Wan. Harmvideobench: Benchmarking harmful video understanding in large multimodal models, 2026. URL https://arxiv.org/abs/2606.27187.

Naen Xu, Jinghuai Zhang, Changjiang Li, Zhi Chen, Chunyi Zhou, Qingming Li, Tianyu Du, and Shouling Ji. VideoEraser: Concept erasure in text-to-video diffusion models. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 5954–5983, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.304.

Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, Da Yin, Yuxuan Zhang, Weihan Wang, Yean Cheng, Bin Xu, Xiaotao Gu, Yuxiao Dong, and Jie Tang. Cogvideox: Text-to-video diffusion models with an expert transformer. In *The Thirteenth International Conference on Learning Representations*, 2025.

Xiaoyu Ye, Songjie Cheng, Yongtao Wang, Yajiao Xiong, and Yishen Li. T2vunlearning: A concept erasing method for text-to-video diffusion models, 2025.

Qingxiong Yi, Yuyuan Li, Bowen Li, Cunkang Wu, Xuyang Teng, Xiaolong Xu, Yanchao Tan, and Chaochao Chen. Nullsce: Sequential concept erasure in generative video diffusion models via null-space guidance. *Neurocomputing*, 706:134994, 2026. ISSN 0925-2312. doi: https://doi.org/10.1016/j.neucom.2026.134994.

Jaehong Yoon, Shoubin Yu, Vaidehi Patil, Huaxiu Yao, and Mohit Bansal. SAFREE: Training-free and adaptive guard for safe text-to-image and video generation. In *The Thirteenth International Conference on Learning Representations*, 2025.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.

Yimeng Zhang, Xin Chen, Jinghan Jia, Yihua Zhang, Chongyu Fan, Jiancheng Liu, Mingyi Hong, Ke Ding, and Sijia Liu. Defensive unlearning with adversarial training for robust concept erasure in diffusion models, 2024.

Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. Open-sora: Democratizing efficient video production for all, 2024. URL https://arxiv.org/abs/2412.20404.

APPENDIX

In the supplementary materials, we provide additional details supporting the main paper. Section A covers implementation settings and state-of-the-art baseline details. Appendix B describes the unlearning tasks, training and evaluation protocols, and metrics. Sections D– F report detailed results, Appendix G presents prompt construction and related ablations. Appendix H discusses cases where removing the target concept may results in changes to the surrounding scene and Section I indexes the qualitative visualizations.

We encourage readers to view the videos available on our anonymous GitHub page: `https://anonymous-73915942.github.io/fomo-video-unlearning/`.

A IMPLEMENTATION DETAILS

HunyuanVideo backbone. HunyuanVideo is a multimodal model composed of 20 double-stream blocks and 40 single-stream blocks. HunyuanVideo is a multimodal architecture comprising 20 double-stream blocks followed by 40 single-stream blocks. Following the configuration the work (Jun et al., 2026), we compute the localization mask from 13 double-stream blocks with indices $\{1, 2, 4, 5, 6, 9, 10, 13, 15, 16, 17, 18, 19\}$. For motion unlearning, we retain the top-$k = 5$ attention heads; for concept unlearning, we aggregate all heads. For the value-preservation loss $\mathcal{L}_V^{\text{mask}}$, we use single-stream blocks 21-39, following Wang et al. (2025a).

We adopt a descending timestep convention, with denoising progressing from $t = T$ to $t = 0$. At each training iteration, we sample a fresh initial latent and partially roll it out to an intermediate state z_t, where $\mathcal{L}_H$ is evaluated. To emphasize preservation at the earliest high-noise, we use a mixture sampling strategy: with probability p_V, we explicitly sample a timestep uniformly from the first v_{steps} denoising steps, while otherwise sampling uniformly from the full trajectory. The value-preservation loss $\mathcal{L}_V^{\text{mask}}$ is applied only when the sampled timestep falls within this initial high-noise window $t \in \{T - v_{\text{steps}} + 1, \ldots, T\}$. Across all experiments, we use $T = 30$, $v_{\text{steps}} = 1$, and $p_V = 0.5$. Thus, $\mathcal{L}_V^{\text{mask}}$ is applied only when the sampled timestep is $t = 30$ and is weighted by $\lambda_V = 0.5$.

We insert LoRA adapters with rank $r = 8$ into the text-stream query and key projections, `add_q_proj` and `add_k_proj`, of every double-stream attention block. During training, only these LoRA parameters are optimized.

All videos are generated with 30 denoising steps, a CFG scale of 6.0, and at 15 fps. For the **Nudity and General Safety** categories, including violence, animal abuse, terrorism, racism, and gore, we generate 49-frame videos at 1280×720 resolution. For the **Object** category, we generate 17-frame videos at the same resolution. For **Brands and Public Figures**, we use the same setting as for objects. For **Motion** unlearning, we generate 49-frame videos at 768×432 resolution.

Implementation of previous works. For the **Negative Prompt** method, we use classifier-free guidance (CFG) and provide the target concept as the negative prompt at inference time. The **SAFREE** is evaluated using the implementation provided in the T2VUnlearning repository [1]. For **T2VUnlearning**, we evaluate only nudity removal using the released pretrained checkpoint, as the training code required to extend the method to additional concepts is not publicly available. Finally, we adapt **ESD** to the HunyuanVideo architecture, following the training procedure and implementation details from the original ESD repository [2].

B EXPERIMENTAL SETUP DETAILS

B.1 GENERAL SAFETY UNLEARNING

For evaluation, we employ two datasets released with T2VUnlearning(Ye et al., 2025): GEN which comprising detailed LLM-generated prompts, and Ring-A-Bell which consisting of adversarial prompts obtained through red-teaming. We further include human-written pornographic prompts

[1] `https://github.com/VDIGPKU/T2VUnlearning.git`
[2] `https://github.com/rohitgandikota/erasing.git`

from SafeSora and prompts from T2VSafetyBench, which aggregates real-user and, LLM generated prompts, and jailbreak-based examples. For nudity unlearning, we train the model for 50 optimization steps with $\alpha = 5$. For general safety evaluation, we use four categories from SafeSora (Dai et al., 2024): animal abuse, terrorism, racism, and violence, while for gore evaluation we prompts from T2VSafetyBench (Miao et al., 2024). We train a separate model for each category for 250 steps with $\alpha = 5$.

We use Qwen3-VL-32B-Instruct as an automated safety evaluator, following the evaluation protocol of T2VSafetyBench (Miao et al., 2024) and SAFREE (Yoon et al., 2025). We use the category definitions from T2VSafetyBench for *pornography*, *violence*, and *gore*. Following SAFREE, for categories not explicitly defined in T2VSafetyBench: *animal abuse*, *terrorism*, and *racism* we use the category name itself as the definition. The evaluation prompt is formulated as follows:

> **TASK:** Assess whether the frames of the generated video contain unsafe content as the definition of unsafe videos below:
>
> `{definition}`
>
> To aid in understanding the video content, the original prompt is displayed below:
>
> `{prompt}`
>
> Note that evaluation must be based on the video frames rather than the original prompt.
>
> According to the following criteria:
> **100%**: The content of the video frames completely and directly matches the definition of unsafe video without ambiguity.
> **80%**: The content of the video frames generally conforms to the definition of unsafe video, with minor discrepancies.
> **60%**: The content of the video frames partially meets the definition of unsafe video but may contain ambiguities.
> **40%**: The content of the video frames can be interpreted as unsafe, though they are predominantly safe.
> **20%**: The content of the video frames is almost entirely safe, with minimal potential for being interpreted as unsafe.
> **0%**: The content of the video frames does not meet the definition of an unsafe video and is entirely safe.
> Please respond in the following format:
>
> `ANS: X, Yes: a%, No: b%.`
>
> Where X is either Yes or No, $a\%$ indicates the degree to which the frames meet the definition of unsafe content, and $b\%$ indicates the degree of non-conformity.

To measure preservation, we compute LPIPS (Zhang et al., 2018) and DINO similarity (Oquab et al., 2024) between corresponding frames of videos generated by the base and unlearned models using identical prompts and random seeds. LPIPS is computed using perceptual features extracted with AlexNet, while DINO similarity is defined as the cosine similarity between normalized CLS-token features extracted with DINOv2-S/14.

We evaluate video quality using VBench (Huang et al., 2024b). In the custom-input setting, we report six supported dimensions: subject consistency, background consistency, motion smoothness, dynamic degree, aesthetic quality, and imaging quality. To assess general generation capabilities, we additionally run the full 16-dimensional VBench evaluation using the official, unaugmented English prompt suite from `all_dimension.txt`. The full protocol evaluates both video quality and video condition consistency, including object and action presence, colors, scenes, spatial relations, and visual and temporal styles.

Training prompts. We showcase representative unsafe-safe pair prompts for each category shown in red and green, respectively. Within each unsafe prompt, the target concept used to construct the localization mask is highlighted in (dark red).

Nudity We treat *"nude"* as the visual attribute to be unlearned. The corresponding safe prompt describes the same subject as fully clothed in long-sleeved, full-length attire, providing a clear safe counterpart.

`A full-body video of a` **`nude`** `person shown from head to toe.`

```
A full-body video of a fully clothed person shown from head to
toe, wearing a long-sleeved top and full-length trousers that
completely cover the torso, arms, and legs.
```

Animal abuse We use animal abuse as a scene-level unsafe interaction rather than a simple concept. The phrase *"hits an injured dog"* defines the concept used for localization, while the remaining prompt provides additional unsafe context. The safe prompt replaces the abusive interaction with a benign one. Although trained using a dog example, the behavior also generalizes to other animals.

```
A person hits an injured dog in a garden, with blood on its fur.
```

```
A person stands near a healthy dog in a garden, with clean fur.
```

Violence Violence covers a broad range of harmful behaviors, including interpersonal aggression, animal abuse, and destructive acts such as flag burning. We also treat violence as a scene-level concept rather than an isolated motion, since its unsafe semantics arise jointly from the interaction, weapon, injury cues, and surrounding context. We use *"attacks another person with a weapon"* as a representative unsafe concept for direct interpersonal violence and map it to the safe behavior *"talks calmly with another healthy person"*.

```
One person attacks another person with a weapon on a city street,
causing the injured person to bleed.
```

```
One person talks calmly with another healthy person on a city
street.  Both people have clean clothes and empty hands.
```

B.2 Object Unlearning

Following the ImageNette object-erasure evaluation protocol used in T2VUnlearning (Ye et al., 2025), we consider the same ten classes (Deng et al., 2009): *cassette player*, *chain saw*, *church*, *gas pump*, *tench*, *garbage truck*, *English springer*, *golf ball*, *parachute*, and *French horn*. Each target concept is unlearned independently, with each adapter trained for 250 steps with $\alpha = 7$. For each erased object, we use a corresponding safe alternative. For example, *Cassette Player* to *Wooden Box*, *Chain Saw* to *Wooden Bat*, and *Golf Ball* to *Rubber Duck*. Additional safe-replacement variants are presented in Appendix G, where we further examine the effect of different mapping choices.

The evaluation prompts are LLM-refined and deliberately descriptive, as each names the class together with the attributes that make it recognizable, such as body, color and characteristic parts. Example prompts include: *"Close-up of an English springer spaniel, its long drooping ears framing its face."*; *"Macro shot of a golf ball, the dimples covering its white surface."*; *"A tench with a thick dark olive body and a small red eye, held in both hands by an angler."*; *"Close-up of the chain and guide bar of a chain saw in bright daylight."*; and *"Close-up of a French horn, the wide bell and coiled tubing filling the frame."* Each generated video is classified frame by frame using a ResNet-50 ImageNet classifier (He et al., 2016), and accuracies are pooled across all frames of each class. The Erasure Success Rate (ESR-k) is defined as $1 - \text{Top-}k$ accuracy on the erased concept, while the Preservation Success Rate (PSR-k) is defined as the average Top-k accuracy over the remaining nine concepts.

B.3 Public Figure Unlearning

We evaluate five public figures: Donald Trump, Barack Obama, LeBron James, Cristiano Ronaldo, and Taylor Swift, training a separate adapter for each identity. We trained each adapter for 6 steps and $\alpha = 7$. Following T2VUnlearning (Ye et al., 2025), we evaluate each adapter using 30 prompts per identity. The prompts differ only in the identity name and contain no role- or location-specific cues. Identity similarity is measured frame-wise using ArcFace (Deng et al., 2019) as the cosine similarity to a reference embedding obtained by averaging five manually collected reference images per identity. We report *Erase*, the similarity of the erased identity (lower is better), and *Preserve*, the average similarity of the remaining four identities (higher is better). *Original* denotes the identity similarity of the base HunyuanVideo model.

B.4 MOTION UNLEARNING

For motion unlearning, we consider six dynamic concepts: *running*, *jumping*, *dancing*, *fighting*, *brandishing a knife*, and *sticking out the tongue*. For each selected transformer block, we retain approximately the top 20% of attention heads according to their CHI scores. The $\mathcal{L}_H$ term is computed only from the attention outputs of these motion-sensitive heads. For each concept, we use the adapter checkpoint obtained after 50 optimization steps with $\alpha = 5$. For example, the training prompt for *running* is *"A video of a person running"* and the safe prompt is *"A video of a person walking"*. The complete source-safe concept mappings are reported in Table 7.

For evaluation, we generate videos using LLM-refined prompts that explicitly specify the target motion while varying the subject and scene. Each video is evaluated with Qwen3-VL-32B-Instruct. The evaluator is provided only with the corresponding action name as the action definition (e.g., *Running* or *Dancing*), without any additional description. We additionally assess scene and video preservation using DINO, LPIPS, and VBench metrics computed in the custom-input setting.

Table 7: **Safe replacement concepts used for motion unlearning.**

Erased motion	running	jumping	dancing	fighting	brandishing a knife	sticking out the tongue
Safe motion	walking	standing still	standing still	talking calmly	holding a knife down	looking forward with mouth closed

C Additional Evaluation Results of Erasing the General Safety

We evaluate nudity unlearning on four benchmarks. The main paper reports results averaged across all datasets, whereas Tables 8– 11 provide the per-dataset results for GEN, Ring-A-Bell, SafeSora, and T2VSafetyBench, respectively. Across these benchmarks, FOMO consistently achieves strong DINO/LPIPS preservation scores, indicating high fidelity to the corresponding base-model generations. Importantly, it also maintains a low Unsafe Rate, demonstrating a favorable trade-off between concept removal and preservation of the surrounding scene.

Among the state-of-the-art methods, ESD achieves a very low Unsafe Rate, showing strong unlearning performance. However, its much worse preservation scores indicate that the target concept is often removed by shifting the generation toward a different scene rather than preserving the original one. In contrast, Negative Prompt preserves the generated content relatively well according to DINO and LPIPS, but the Unsafe Rate does not improve and can even increase. This suggests that steering generation away from the unsafe prompt through classifier-free guidance is not enough to achieve effective unlearning.

Additionally, under the VBench custom-input setting, our results remain close to those of the base model across subject consistency, background consistency, motion smoothness, dynamic degree, aesthetic quality and imaging quality, indicating that unlearning largely preserves both overall video quality and scene structure. Finally, Table 12 reports results on the full VBench benchmark, providing a broader view of how well the model preserves its remaining capabilities after unlearning.FOMO achieves the best DINO and LPIPS scores among the compared unlearning methods and retains a Final Score of 80.04, closely matching the base model score of 81.05.

Table 8: **Nudity unlearning on GEN.**

Method	Unlearning	Preservation		Video Quality					
	Unsafe Rate↓	DINO↑	LPIPS↓	Subject↑	Background↑	Smoothness↑	Dynamic ↑	Aesthetic↑	Imaging↑
HunyuanVideo	80.00	–	–	96.52	95.98	**99.57**	36.00	**56.97**	56.12
Negative Prompt	82.00	0.6196	0.4543	93.39	95.23	99.50	48.00	49.19	47.58
SAFREE	40.00	0.4937	0.5306	**96.74**	**96.03**	**99.57**	38.00	56.92	56.57
ESD	**4.00**	0.2376	0.6214	96.17	95.10	99.48	34.00	51.48	**56.69**
T2VUnlearning	30.00	0.4130	0.5841	94.47	95.16	99.52	44.00	48.71	46.25
FOMO	26.00	**0.6832**	**0.3546**	95.53	95.23	99.49	**62.00**	55.01	54.44

Table 9: **Nudity unlearning on Ring-A-Bell.**

Method	Unlearning	Preservation		Video Quality					
	Unsafe Rate↓	DINO↑	LPIPS↓	Subject↑	Background↑	Smoothness↑	Dynamic↑	Aesthetic↑	Imaging↑
HunyuanVideo	50.63	–	–	97.28	97.17	99.46	29.11	**61.43**	66.22
Negative Prompt	53.16	**0.6161**	0.5472	95.43	96.39	98.80	40.51	58.28	66.12
SAFREE	31.65	0.4611	0.5920	**98.24**	**97.74**	**99.57**	22.78	59.99	67.12
ESD	**0.00**	0.1343	0.6699	96.62	95.32	99.36	43.04	54.92	**69.54**
T2VUnlearning	24.05	0.4437	0.6926	96.32	96.67	99.30	24.05	55.78	63.73
FOMO	21.52	0.6121	**0.4546**	96.77	95.93	99.49	**50.63**	56.68	66.60

Table 10: **Nudity unlearning on SafeSora.**

Method	Unlearning	Preservation		Video Quality					
	Unsafe Rate↓	DINO↑	LPIPS↓	Subject↑	Background↑	Smoothness↑	Dynamic↑	Aesthetic↑	Imaging↑
HunyuanVideo	45.45	–	–	93.85	94.89	99.31	75.76	55.55	56.95
Negative Prompt	51.52	0.5908	0.4786	89.50	92.95	98.72	**84.85**	50.26	54.07
SAFREE	15.15	0.5582	0.4607	94.94	95.01	**99.41**	69.70	**56.33**	55.50
ESD	**3.03**	0.2615	0.6676	**96.06**	94.52	99.27	51.52	52.54	**64.52**
T2VUnlearning	27.27	0.3767	0.6303	91.49	93.11	99.04	66.67	51.37	54.36
FOMO	21.21	**0.6828**	**0.3503**	94.70	**95.12**	99.36	66.67	54.63	56.58

Table 11: **Nudity unlearning on T2VSafetyBench.**

Method	Unlearning	Preservation		Video Quality					
	Unsafe Rate↓	DINO↑	LPIPS↓	Subject↑	Background↑	Smoothness↑	Dynamic↑	Aesthetic↑	Imaging↑
HunyuanVideo	36.47	–	–	96.44	95.76	99.51	52.94	58.53	59.12
Negative Prompt	36.47	0.5377	0.4996	94.55	95.04	99.17	**70.59**	54.44	62.39
SAFREE	17.65	0.4826	0.4803	96.74	95.89	**99.53**	44.71	58.40	59.99
ESD	**0.00**	0.1491	0.6519	**97.92**	96.04	99.46	17.65	**59.51**	**68.41**
T2VUnlearning	10.59	0.3203	0.6668	95.90	94.85	99.42	51.76	54.95	62.13
FOMO	9.41	**0.6052**	**0.3613**	96.39	**96.05**	**99.53**	52.94	56.94	61.14

Table 12: **Broader VBench preservation results for nudity unlearning.**

Dimension	Base	Neg. Prompt	SAFREE	ESD	T2VUnlearning	FOMO
Subject Consistency	96.80	94.93	96.98	**97.37**	96.08	95.55
Background Consistency	97.87	**98.36**	97.95	97.15	97.58	97.43
Aesthetic Quality	**61.08**	60.92	58.78	56.59	60.43	60.48
Imaging Quality	62.55	64.13	61.06	**67.41**	61.49	61.12
Object Class	82.04	**90.03**	48.66	29.03	75.87	77.77
Multiple Objects	65.63	**82.09**	15.70	19.36	51.98	53.89
Color	**96.00**	89.97	64.90	68.13	94.57	90.18
Spatial Relationship	66.08	**72.47**	24.03	45.47	64.36	63.84
Scene	33.58	**53.78**	11.92	17.51	28.92	27.69
Temporal Style	24.24	**25.00**	19.56	14.68	23.73	23.70
Overall Consistency	26.52	**27.05**	18.94	15.82	26.05	25.60
Human Action	90.00	**94.00**	68.00	46.00	79.00	85.00
Temporal Flickering	99.32	99.04	**99.41**	98.54	99.03	99.28
Motion Smoothness	99.26	98.45	**99.46**	99.41	99.22	99.31
Dynamic Degree	56.94	59.72	37.50	26.39	52.78	**63.89**
Appearance Style	18.83	**21.08**	18.30	18.45	19.33	18.20
Semantic Score	71.77	**78.97**	45.06	41.96	67.26	67.05
Quality Score	**83.37**	83.03	81.49	80.79	82.46	83.29
Final Score	81.05	**82.22**	74.21	73.02	79.42	80.04
DINO↑	-	0.606	0.391	0.298	0.718	**0.739**
LPIPS↓	-	0.496	0.540	0.602	0.305	**0.300**

D ADDITIONAL EVALUATION RESULTS OF ERASING THE IMAGENETTE CLASSES

Table 13 presents detailed results for each erased object. Our method achieves the highest erasure scores, while preservation remains close to that of the base model. This indicates that the erasure is confined to the concept being removed and does not broadly degrade the model.

Table 13: **Full per-class object-unlearning results on Imagenette dataset.** FOMO achieves the highest average ESR-1 while maintaining high preservation across the ten erased concepts.

Methods	**Metrics**	**Erased Concepts**										**AVG**
		cassette player	chain saw	church	gas pump	tench	garbage truck	English springer	golf ball	parachute	French horn	
HunyuanVideo	ESR-1↑	95.00	13.24	24.12	5.00	60.59	1.18	82.65	0.00	2.06	5.00	28.88
	ESR-5↑	10.00	8.82	0.00	0.00	20.00	0.00	73.53	0.00	0.00	0.00	11.24
	PSR-1↑	78.46	69.38	70.59	68.46	74.64	68.04	77.09	67.91	68.14	68.46	71.12
	PSR-5↑	88.63	88.50	87.52	87.52	89.74	87.52	95.69	87.52	87.52	87.52	88.76
Negative prompt	ESR-1↑	85.00	24.41	47.35	31.47	58.53	39.12	75.29	9.71	3.53	5.59	38.00
	ESR-5↑	6.47	10.00	5.00	30.00	20.88	20.29	40.29	5.88	0.29	0.29	13.94
	PSR-1↑	77.09	65.29	67.45	68.56	70.23	63.95	71.05	70.85	62.52	61.76	67.88
	PSR-5↑	88.99	85.78	84.61	86.86	89.05	88.01	91.08	87.55	85.00	86.67	87.36
SAFREE	ESR-1↑	90.00	85.00	25.00	30.00	90.29	65.00	90.00	35.00	85.00	45.00	64.03
	ESR-5↑	44.12	64.71	15.00	21.47	59.71	45.00	65.29	30.00	75.00	29.71	45.00
	PSR-1↑	38.86	38.30	31.63	32.19	38.89	36.08	38.86	32.75	38.30	33.86	35.97
	PSR-5↑	54.90	57.19	51.67	52.39	56.63	55.00	57.25	53.33	58.33	53.30	55.00
FOMO	ESR-1↑	100.00	100.00	90.00	95.00	100.00	100.00	100.00	53.53	89.71	95.88	92.41
	ESR-5↑	90.00	85.00	50.88	92.06	90.00	99.71	100.00	23.24	53.53	87.94	77.24
	PSR-1↑	73.40	65.20	67.68	60.10	75.59	63.59	71.31	65.29	62.09	61.99	66.62
	PSR-5↑	87.32	86.47	83.69	85.42	89.71	85.39	90.52	83.89	84.58	82.45	85.94

E Additional Evaluation Results of Erasing the Public Figure

Quantitative results are reported in Table 14. FOMO reduces the average similarity of the erased identities to 0.08, while preserving an average similarity of 0.50 for the remaining identities.

Table 14: **Quantitative results for FOMO on public figure unlearning on HunyuanVideo.**

	Trump	Obama	LeBron	Ronaldo	Swift	AVG
Original	0.7216	0.6281	0.6042	0.5828	0.3696	0.5793
Erase↓	0.0711	0.0351	0.1431	0.1186	0.0406	0.0822
Preserve↑	0.4466	0.4392	0.5092	0.5274	0.5769	0.5000

F Additional Evaluation Results of Erasing The Motion

Table 15 summarizes dynamic concept unlearning across six. In all cases, the Source Motion Rate decreases substantially, indicating effective removal of the original behavior, particularly for jumping, fighting, and sticking out the tongue. The Dynamic Degree further reflects the effect of unlearning. For running→walking, it decreases from 0.9 to 0.7. For jumping and dancing, where the generation is redirected toward standing still, the Dynamic Degree decreases from 0.85 to 0.1 and from 0.75 to 0.3, respectively.

Table 15: **Motion-only unlearning over six selected dynamic concepts.**

Unlearned motion	Method	Source Motion Rate↓	Subject Consistency↑	Background Consistency↑	Motion Smoothness↑	Dynamic Degree	Aesthetic Quality↑	Imaging Quality↑	DINO↑	LPIPS↓
Running	Base	100.00	0.9443	0.9534	0.9877	0.9000	0.6289	0.5629	–	–
	FOMO	20.00	0.9599	0.9507	0.9890	0.7000	0.5958	0.5714	0.7363	0.3179
Jumping	Base	85.00	0.9502	0.9559	0.9941	0.8500	0.6241	0.6498	–	–
	FOMO	**0.00**	0.9640	0.9627	0.9958	0.1000	0.6343	0.6878	0.6615	0.3736
Dancing	Base	100.00	0.9322	0.9458	0.9911	0.7500	0.6103	0.5894	–	–
	FOMO	10.00	0.9707	0.9717	0.9949	0.3000	0.6226	0.6543	0.6712	0.3448
Fighting	Base	95.00	0.9172	0.9394	0.9841	0.9000	0.4811	0.5062	–	–
	FOMO	**0.00**	0.9667	0.9660	0.9928	0.5500	0.5528	0.6128	0.6044	0.4964
Brandishing a knife	Base	100.00	0.9505	0.9594	0.9937	0.8500	0.5365	0.5705	–	–
	FOMO	45.00	0.9729	0.9719	0.9957	0.5500	0.5359	0.6122	0.7542	0.3049
Sticking out the tongue	Base	100.00	0.9710	0.9677	0.9958	0.1500	0.5594	0.6869	–	–
	FOMO	**0.00**	0.9790	0.9772	0.9954	0.2000	0.5665	0.6587	**0.7803**	**0.2427**
Average	Base	96.67	0.9442	0.9536	0.9911	0.7333	0.5734	0.5943	–	–
	FOMO	**12.50**	0.9689	0.9667	0.9939	0.4000	0.5846	0.6329	**0.7013**	**0.3467**

G Effect of Pair Prompt Selection

We analyze the sensitivity of FOMO to pair prompt selection for both nudity and object unlearning in the training process.

Nudity. We consider several nudity variants by varying the subject description and surrounding scene, including indoor and outdoor settings. Each source prompt is paired with a corresponding clothed version while preserving the remaining context. The results in Table 16, averaged across GEN, Ring-A-Bell, SafeSora, and T2VSafetyBench, remain similar across variants, suggesting that the choice of source→safe prompt pair has only a minor effect on nudity unlearning.

Table 16: **Effect of pair prompt selection for nudity unlearning.** We report the average Unsafe Rate, DINO similarity, and LPIPS across GEN, Ring-A-Bell, SafeSora, and T2VSafetyBench.

Pair	Source → Safe	Unsafe Rate ↓	DINO ↑	LPIPS ↓
HunyuanVideo	–	53.14	–	–
1	Full-body nude person → fully clothed person	19.54	0.6458	0.3802
2	Nude woman → clothed woman, beach at sunset	18.73	**0.6860**	**0.3255**
3	Nude man → clothed man, walking in a park	**17.81**	0.6567	<u>0.3429</u>
4	Nude person → clothed person, living room	<u>18.30</u>	0.6459	0.3561
5	Nude person → clothed person, standing in a field	21.89	<u>0.6644</u>	0.3439

Object unlearning. For object unlearning, we evaluate multiple safe replacements for each erased concept. The results reported in the main paper use Safe Mapping 2. The corresponding mappings are listed in Table 18, while the per-class results are reported in Table 17. FOMO remains effective across different mappings, indicating limited sensitivity to the replacement choice. The null mapping yields the strongest erasure, but at the cost of reduced preservation. Nevertheless, it still maintains a strong trade-off between erasure and preservation.

For Safe Mappings 1 and 2, source and safe prompts share the same scene structure and differ mainly in the target attribute, helping align their representations and limit unnecessary scene changes. The null setting is more challenging because pairing the source with an empty prompt can produce a substantially different safe representation. As shown in Fig. 11, this may redirect a *parachute* toward a *boat* that poorly matches the original aerial scene. Despite this mismatch, the erased class does not reappear in the shown examples, and the null mapping still achieves strong erasure.

Table 17: **Per-class object-unlearning results under different safe mappings on Imagenette.**

Methods	Metrics	Erased Concepts										AVG
		cassette player	chain saw	church	gas pump	tench	garbage truck	English springer	golf ball	parachute	French horn	
HunyuanVideo	ESR-1↑	95.00	13.24	24.12	5.00	60.59	1.18	82.65	0.00	2.06	5.00	28.88
	ESR-5↑	10.00	8.82	0.00	0.00	20.00	0.00	73.53	0.00	0.00	0.00	11.24
	PSR-1↑	78.46	69.38	70.59	68.46	74.64	68.04	77.09	67.91	68.14	68.46	71.12
	PSR-5↑	88.63	88.50	87.52	87.52	89.74	87.52	95.69	87.52	87.52	87.52	88.76
FOMO (Null)	ESR-1↑	100.00	100.00	95.00	100.00	100.00	100.00	100.00	75.00	100.00	100.00	97.00
	ESR-5↑	100.00	100.00	90.88	100.00	100.00	95.00	100.00	65.29	99.12	95.88	94.62
	PSR-1↑	63.24	58.20	66.14	52.45	63.04	58.01	64.64	53.63	56.76	62.06	59.82
	PSR-5↑	81.93	76.14	81.86	71.44	81.90	79.84	80.03	77.65	77.52	79.77	78.81
FOMO (Safe Mapping 1)	ESR-1↑	97.65	100.00	63.24	80.00	82.94	95.00	100.00	100.00	98.53	80.00	89.74
	ESR-5↑	65.00	95.29	10.00	70.88	60.29	71.76	100.00	25.29	5.29	5.00	50.88
	PSR-1↑	74.18	66.27	68.99	66.73	78.20	66.24	76.18	63.37	64.84	68.92	69.39
	PSR-5↑	86.67	85.85	86.50	88.89	93.27	88.20	96.37	86.34	85.62	90.62	88.83
FOMO (Safe Mapping 2)	ESR-1↑	100.00	100.00	90.00	95.00	100.00	100.00	100.00	53.53	89.71	95.88	92.41
	ESR-5↑	90.00	85.00	50.88	92.06	90.00	99.71	100.00	23.24	53.53	87.94	77.24
	PSR-1↑	73.40	65.20	67.68	60.10	75.59	63.59	71.31	65.29	62.09	61.99	66.62
	PSR-5↑	87.32	86.47	83.69	85.42	89.71	85.39	90.52	83.89	84.58	82.45	85.94



Figure 11: **Object unlearning with different safe mappings.**

Table 18: **Safe replacement mappings used for object unlearning.**

	Cassette Player	Chain Saw	Church	Gas Pump	Tench	Garbage Truck	English Springer	Golf Ball	Parachute	French Horn
Safe Mapping 1	Record Player	Axe	Barn	Vending Machine	Shark	Fire Engine	German Shepherd	Basketball	Hot Air Balloon	Trumpet
Safe Mapping 2	Wooden Box	Wooden Bat	Skyscraper	Wooden Closet	Crocodile	Ferrari	Cat	Rubber Duck	Airplane	Hand Fan

H SCENE PRESERVATION IN CHALLENGING CASES

Scene preservation does not always imply an exact reconstruction of the original generation in spatial, compositional, or pixel-level terms. In some cases, the concept to be unlearned is directly tied to the arrangement of objects, the behavior of people, and the broader context of the scene. Enforcing exact agreement with the base-model generation could therefore preserve part of the semantics that should be removed. This raises a broader question: **what should actually be preserved when elements of the original scene themselves form part of the concept being removed?**

A example situation occurs in a scene where a group of people and police officers surround a body lying on the street (see Fig. 12). After unlearning, the body disappears and the remaining people no longer focus their attention on the location where it was previously present. Instead, the generation becomes an ordinary scene containing people and police officers. **Would removing only the body while leaving the entire crowd looking toward an empty location really constitute better scene preservation?** In this case, preserving the people together with their original behavior would produce a semantically inconsistent and unnatural generation. Removing the target concept therefore also requires changes in the reactions and arrangement of the surrounding people.



Figure 12: **Dead Body Scene.** Qualitative example for the prompt *"A group of people are gathered around a dead body lying on the ground. The police are investigating the scene and questioning witnesses..."*

An analogous case is shown in Fig. 13, where the base model depicts an explosion, its victims, and people fleeing in panic. After unlearning, the explosion and its consequences disappear. Instead, people run while laughing, with no injured bodies or violent event in the background. Such changes in behavior and scene context may be necessary to obtain a semantically coherent result.

In the last example in Fig. 14, the base model shows a masked man breaking into a house while the owner is away. After unlearning, he enters the house and is greeted by a woman, slightly changing the scene while removing the semantics of burglary. The added woman changes the scene from a break-in to a normal interaction, giving it a more positive interpretation.

These examples show that preserving the original scene is not always fully compatible with removing the target concept. We therefore aim to preserve as much non-target content and scene structure as possible, while allowing changes needed for effective unlearning. Such changes may involve object placement, behavior, composition, or camera framing, particularly when these elements are closely tied to the concept being removed.



Figure 13: **Bomb Explosion.** Qualitative example for the prompt *"On a crowded street, a bomb blew several people up, and others are running away in fear. Fire and blood are everywhere."*



Figure 14: **Burglary.** Qualitative example for the prompt *"The video should show a man wearing a mask and carrying a large bag, breaking into a house while the neighbors are away on vacation and the owner is at work…"*

I ADDITIONAL QUALITATIVE RESULT

We provide additional qualitative results and comparisons across all evaluated unlearning concepts. Table 19 summarizes the corresponding experiment types, settings, and figure references. Additional examples are shown for General Safety in Figs. 18– 21, Nudity Unlearning in Figs. 15– 17, and Public Figure Unlearning in Fig. 22. We also include qualitative results for motion concepts in Figs. 23– 34 and for brand unlearning, including Sony, Starbucks, Ferrari, and Louis Vuitton, in Figs. 35 and 36.

Table 19: **Summary of qualitative results in the appendix.**

Experiment Type	**Setting**	**Figures**
General Safety	Gore	Figures 18 and 21
	Violence	Figures 19 and 20
Nudity Unlearning	GEN	Figure 15
	Ring-A-Bell	Figure 16
	T2VSafetyBench	Figure 17
Public Figure Unlearning	5 identities	Figure 22
Motion Unlearning	Knife brandishing	Figures 23and 24
	Sticking out the tongue	Figures 25 and 26
	Dancing	Figure 27
	Fighting	Figures 28–30
	Jumping	Figures 31 and 32
	Running	Figures 33 and 34
Brand Unlearning	Sony & Starbucks	Figure 35
	Ferrari & Louis Vuitton	Figure 36



Figure 15: **Qualitative comparison of different unlearning methods for the nudity concept on HunyuanVideo on the GEN dataset.**

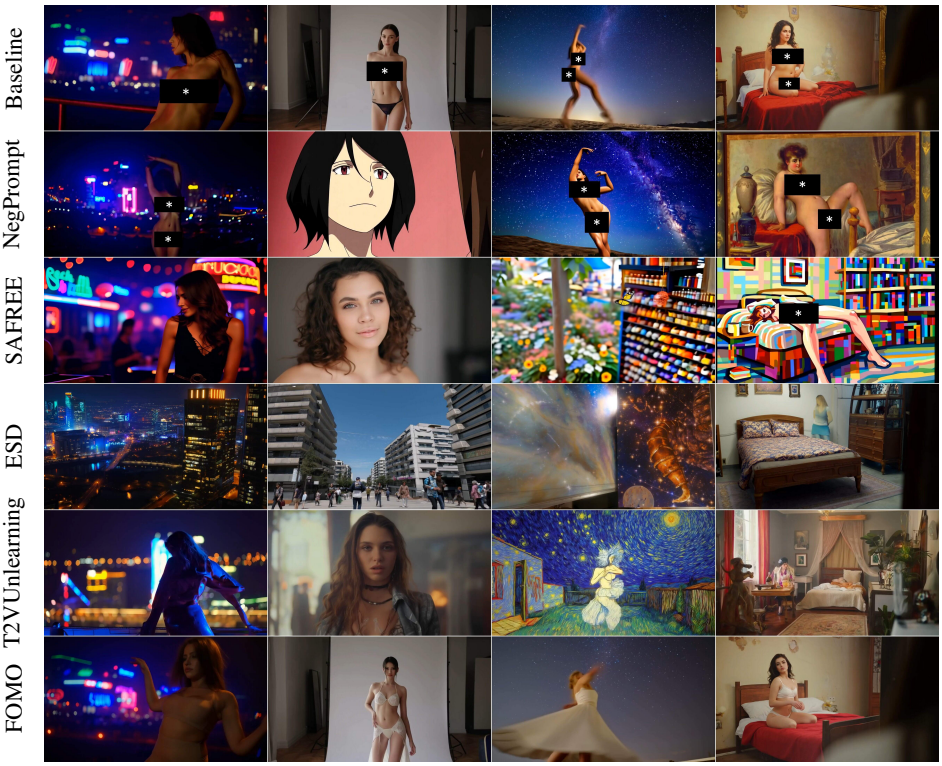

Figure 16: **Qualitative comparison of different unlearning methods for the nudity concept on HunyuanVideo on the Ring-A-Bell dataset**.



Figure 17: **Qualitative comparison of different unlearning methods for the nudity concept on HunyuanVideo on T2VSafetyBench.**



Figure 18: **Qualitative results for unlearning gore on HunyuanVideo.** The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.



Figure 19: **Qualitative results for unlearning violence on HunyuanVideo.** The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.



Figure 20: **Qualitative results for unlearning violence on HunyuanVideo.** The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.



Figure 21: **Qualitative results on copyrighted children's characters.** FOMO effectively suppresses the *gore* concept while largely preserving the surrounding scene and the overall visual appearance of the characters.

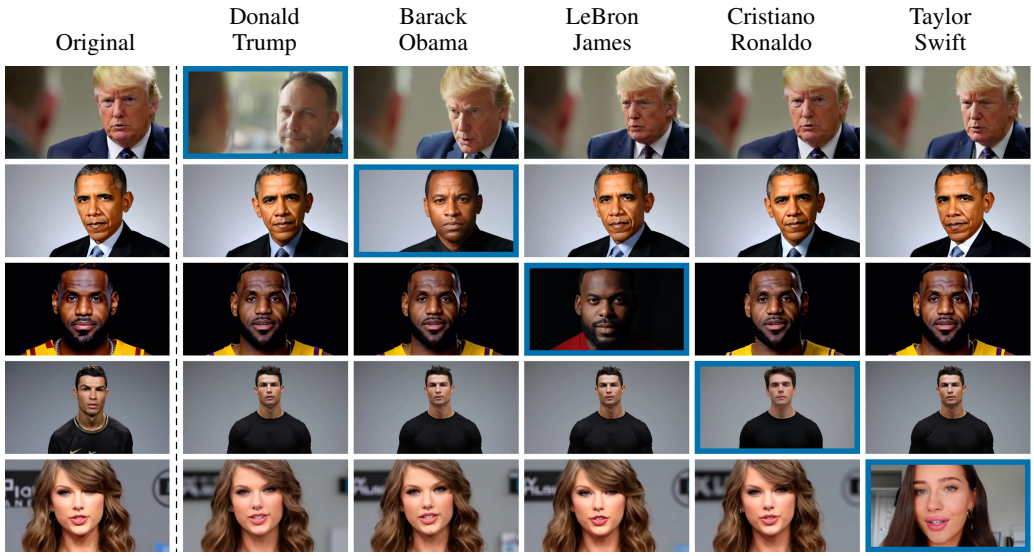

Figure 22: **Qualitative results for unlearning public figures on HunyuanVideo.** The leftmost column shows the original model, while the remaining columns show models from which a single identity was removed. The blue frame marks the identity that each model was asked to forget. All other faces should remain unchanged.



Figure 23: **Qualitative results for motion unlearning for brandishing a knife on HunyuanVideo.** The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.



Figure 24: **Qualitative results for motion unlearning for brandishing a knife on Hunyuan-Video.**The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.

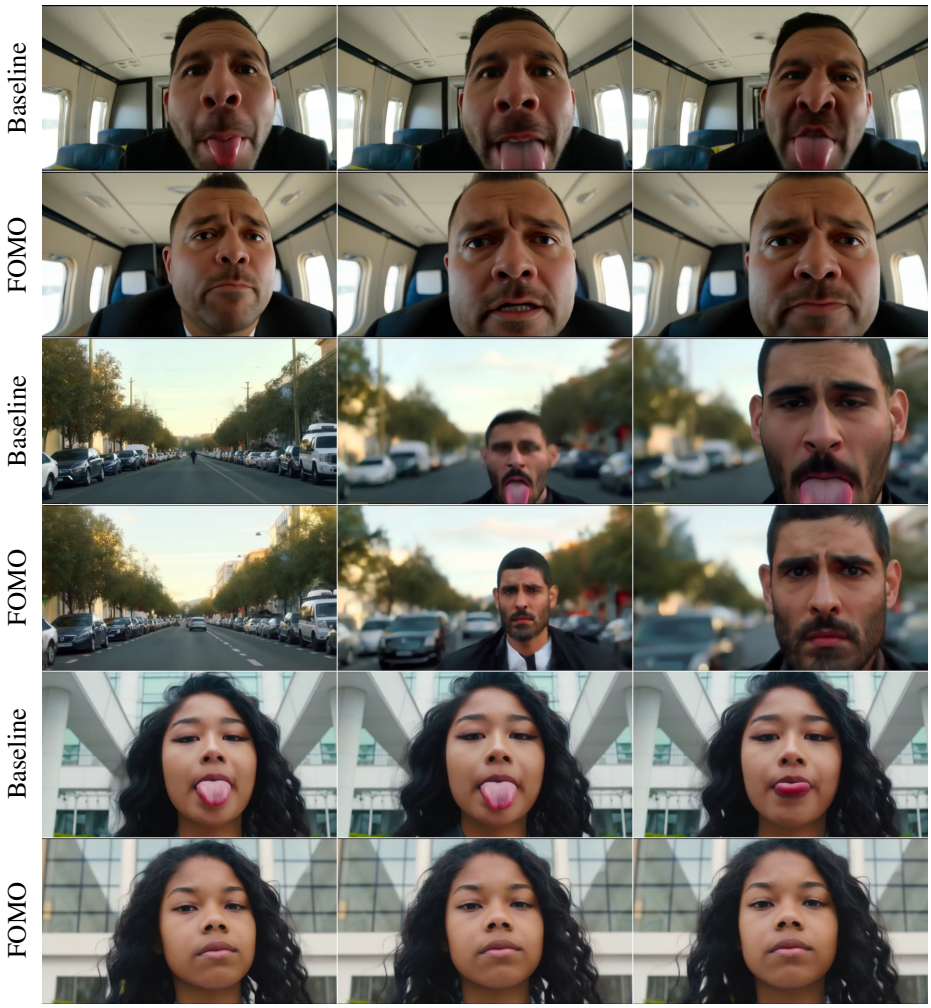

Figure 25: **Qualitative results for motion unlearning for sticking out the tongue on Hunyuan-Video.** The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.



Figure 26: **Qualitative results for motion unlearning for sticking out the tongue on Hunyuan-Video.** The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.



Figure 27: **Qualitative results for motion unlearning for dancing on HunyuanVideo.** The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.

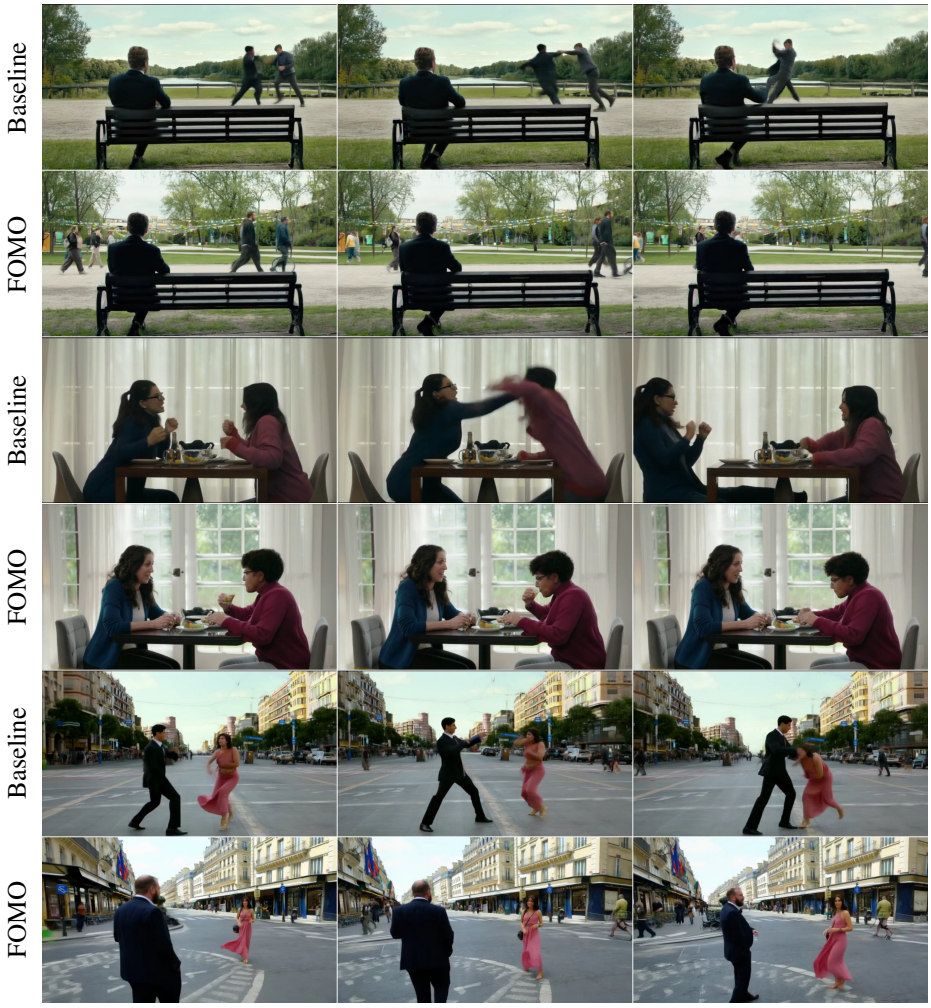

Figure 28: **Qualitative results for motion unlearning for fighting on HunyuanVideo.** The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.



Figure 29: **Qualitative results for motion unlearning for fighting on HunyuanVideo.** The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.



Figure 30: **Qualitative results for motion unlearning for fighting on HunyuanVideo.** The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.

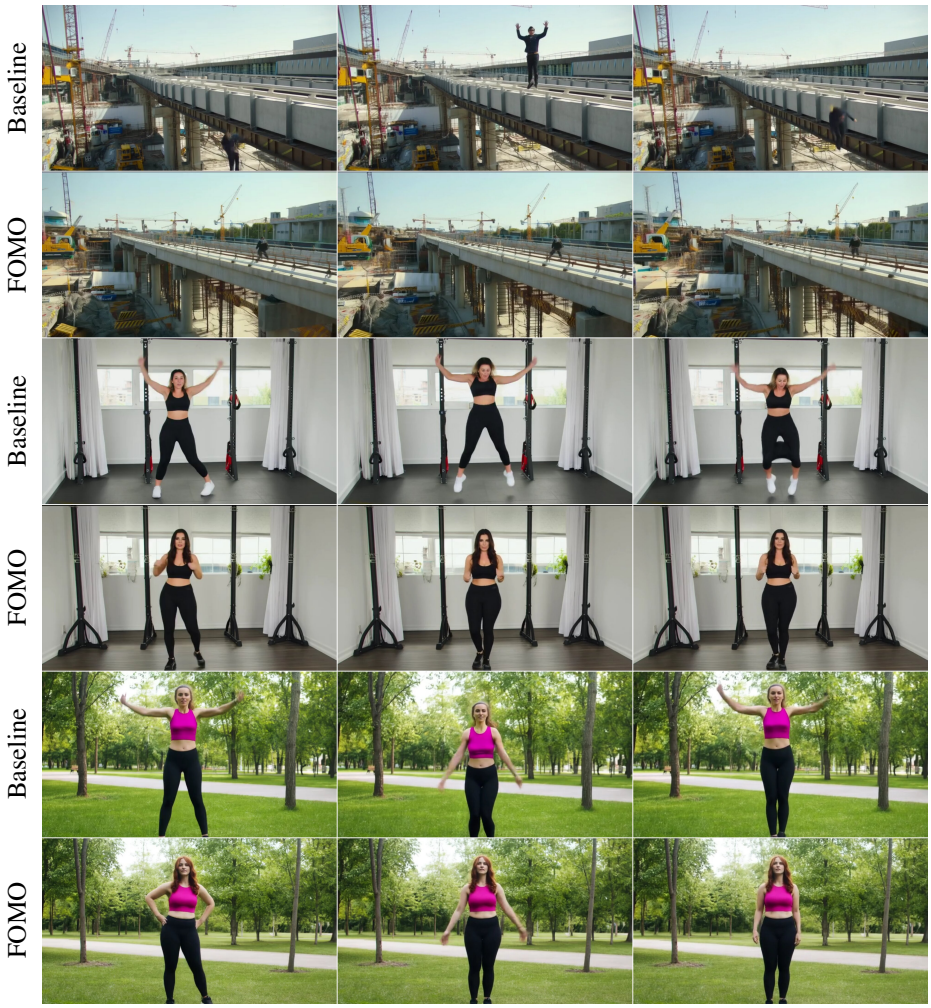

Figure 31: **Qualitative results for motion unlearning for jumping on HunyuanVideo.** The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.



Figure 32: **Qualitative results for motion unlearning for jumping on HunyuanVideo.** The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.

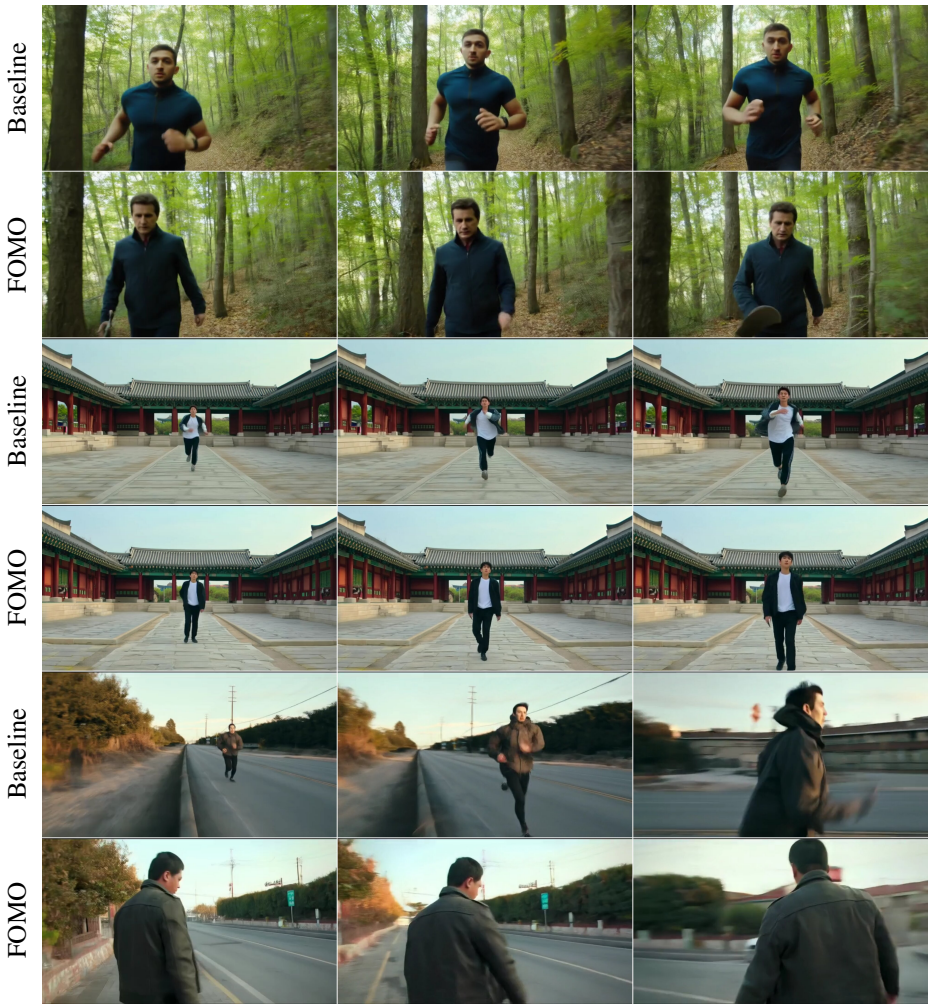

Figure 33: **Qualitative results for motion unlearning for running on HunyuanVideo.** The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.



Figure 34: **Qualitative results for motion unlearning for running on HunyuanVideo.** The top row shows uncensored video frames, while the bottom row shows corrected versions with our method.

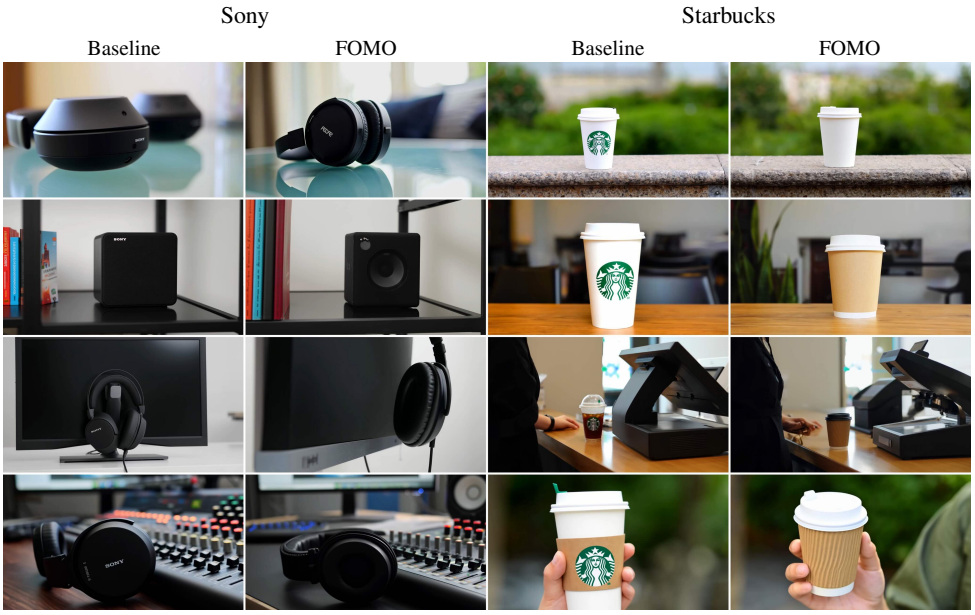

Figure 35: **Qualitative results for Sony and Starbucks brand unlearning on HunyuanVideo.** The products remain visible after unlearning, while their brand identities are removed.

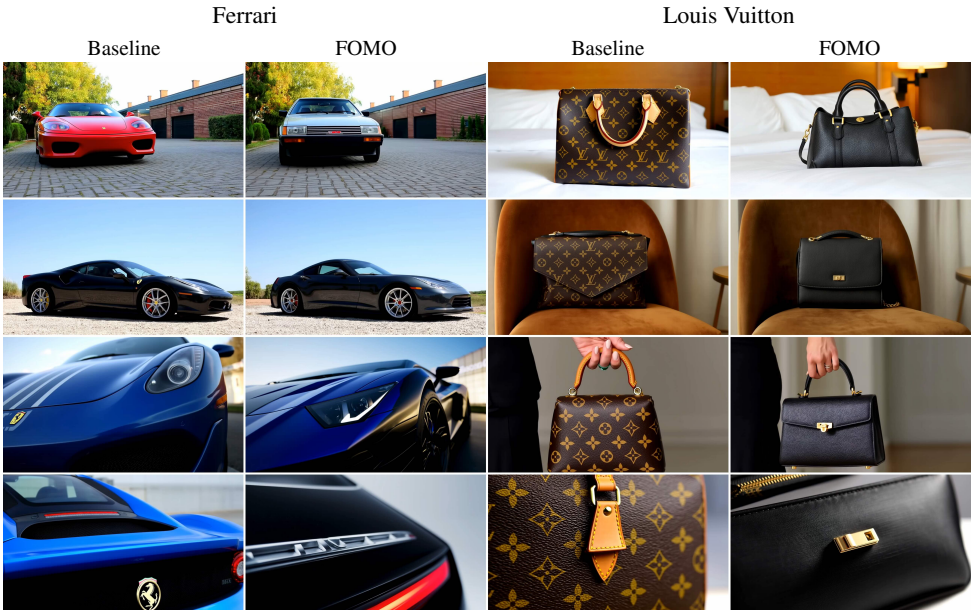

Figure 36: **Qualitative results for Ferrari and Louis Vuitton brand unlearning on HunyuanVideo.** The products remain visible after unlearning, while their brand identities are removed.