SafeMo: Linguistically Grounded Unlearning for Trustworthy Text-to-Motion Generation

Anonymous ACL submission

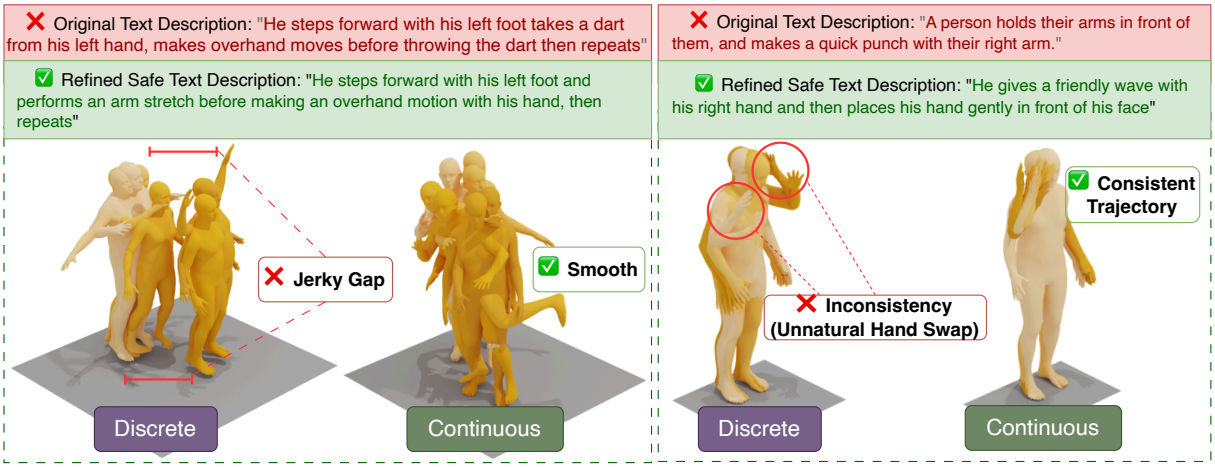

Figure 1: We propose a linguistically grounded framework that first utilizes **SafeMoEngine** to rewrite harmful prompts into safe alternatives (green boxes). Compared to discrete motion method which suffers from quantization artifacts and piecewise transitions, continuous-space approach maintains smooth trajectories and kinematic consistency while strictly adhering to the refined safe intent.

Abstract

Text-to-motion (T2M) generation have achieved impressive realism but pose significant safety risks by faithfully executing harmful textual prompts. Existing safety measures face dual challenges: brittle keyword filtering and discrete codebook manipulation, which degrade benign generation quality and introduce jerky transitions. To address these challenges, we propose **SafeMo**, a linguistically grounded framework that leverages Large Language Models (LLMs) to align motion generation with safety constraints. First, we introduce **SafeMoEngine**, an LLM-agent pipeline that autonomously classifies harmful intents and performs semantic rewriting to construct **SafeMoVAE-29K**, the first safety-aligned text-to-motion dataset. Second, we propose **Minimal Motion Unlearning (MMU)**, a continuous-space unlearning strategy that projects harmful semantic concepts into a task vector for precise negation, avoiding the quantization artifacts of discrete methods. Experiments demonstrate effective unlearning performance of SafeMo by showing strengthened forgetting on unsafe prompts, reaching 2.5× and 14.4× higher forget-set FID on HumanML3D and Motion-X respectively, compared to the previous SOTA human motion unlearning method LCR, with benign performance on safe prompts being better or comparable.

1 Introduction

Generative models thrive across domains, including texts (Brown et al., 2020; Chowdhery et al., 2023; Touvron et al., 2023; Qin et al.), images (Rombach et al., 2022; Ruiz et al., 2023) and videos (Rombach et al., 2022; Fei et al., 2024). Human motion generation methods have numerous achievements in recent years (Guo et al., 2022b; Zhang et al., 2023). Diffusion-based text-to-motion (T2M) models produce compelling human motions conditioned on natural language (Chen et al., 2023; Tevet et al., 2023, 2025). Recent benchmarks such as HumanML3D (Guo et al., 2022a) and Motion-

X (Lin et al., 2023) enable large-scale training and evaluation. However, these methods can memorize and produce harmful motions (e.g. punching, weapon use), which is not desirable for many applications and may lead to misuse. Hence, it is imperative to constrain the model to generate safe outputs that align with regulations and ethics. Machine unlearning is a good strategy to address the safety generation issue, which has been extensively studied on LLMs (Yao et al., 2024) and images (Gandikota et al., 2024; Gong et al., 2024; Lu et al., 2024). It enables the model to forget unsafe samples and undesired behaviors obtained in the training process. Existing human motion unlearning method (De Matteis et al., 2025) notably redirect the generation process away from harmful patterns by replacing the codebook entries in discrete latent space.

However, exiting motion unlearning methods suffer from several issues, as shown in Figure 1: *(i)* codebook coupling problem and smoothness loss, which are resulted from operating VQ tokens reused by benign prompts in discrete code space, perturbing learned token distribution, introducing jerky transitions and behavior drifts on safe prompts; *(ii)* lack of a trustworthy text-to-motion (T2M) dataset for human motion unlearning, with fine-grained safe rewritten text prompts and corresponding refined safe motion.

Hence, to address the first challenge, we propose a Minimal Motion Unlearning (MMU) strategy for human motion unlearning on top of DiP transformer, which isolates the harmful capability in a low-rank subspace and then subtracts it by the needed scale. We first train LoRA adapters using motion-aware objectives to push the model along with the unsafe generation, together with a negative preservation divergence that deliberately pushes the model away from the performance of the frozen base model on benign tasks to obtain a harmful task vector (Ilharco et al., 2022), enabling the later subtraction of this increment not only to erase the model's capability to generate unsafe motion but also to restore the utility on everyday tasks. After that, a LoRA scaling negation is performed at inference, instantly removing the learned unsafe task vector to obtain the trustworthy safe motion generation model.

Furthermore, to address the second challenge, we design and present the first safe text-to-motion dataset on top of HumanML3D, with fine-grained LLM agent rewritten safe text prompts and refined trustworthy human motion in both discrete and continuous versions, namely SafeMoVQ-29K and SafeMoVAE-29K, respectively. Compared to existing methods' keyword-based trimming strategy, our designed LLM-based classify-then-rewrite SafeMoEngine fundamentally mitigates the editing brittleness issue. To obtain refined texts for unsafe prompts, prior works rely on handcrafted keyword lists, where toxic intents are merely removed, distorting semantics. In contrast, our proposed method ensures higher fidelity in linguistic meaning and broader coverage against implicit toxicity and covers both continuous and discrete forms, accommodating different model architectures and ensuring broad usability.

Our contributions can be summarized as follows:

- We propose SafeMo, an trustworthy text-to-motion generative framework equipped with a powerful two-stage selective harmful motion unlearning method, MMU, which enables effective erasure of undesirable behaviors while preserving model utility on benign inputs.
- We design and release the first safe text-to-motion dataset, SafeMoVAE-29K, with rewritten safe text prompts and trustworthy motion, along with corresponding discrete version, via an LLM-assisted pipeline. This dataset fills the critical gap of lacking safe T2M datasets, overcomes the brittleness of keyword-based refinement, and provides broad applicability across different model architectures.
- SafeMo demonstrates stronger empirical unlearning performance than LCR (De Matteis et al., 2025), achieving forget set +150.5% FID and -35.3% R@1 on HumanML3D, and $14.4\times$ FID on Motion-X, with better or comparable performance on benign tasks.

2 Related Work

Text-to-motion generation. Text-driven 3D human motion generation has progressed rapidly (Zhang et al., 2024d), following two lines: (i) discrete token-based sequence modeling (Zhang et al., 2024b,a, 2025b) and (ii) continuous-space generative modeling (Zhang et al., 2024c). Discrete methods such as TM2T (Guo et al., 2022b) employ vector quantization (VQ) and model bidirectional text–motion mapping. T2M-GPT (Zhang et al., 2023) combines Vector Quantized Variational Autoencoders (VQ-VAEs) with autoregressive

transformers for strong text–motion alignment, while MoMask (Guo et al., 2024) adopts hierarchical residual VQ, improves precision and enables finer details. Motion-Agent (Wu et al., 2024) further leverages LLMs for finetuned text–motion generation and a conversational agent enabling long, customizable sequences. These VQ-based approaches offer efficient sampling and long-range structure, but may incur information loss, error accumulation, and stitching artifacts. In contrast, continuous-space generation typically yields smoother temporal transitions. MLD (Chen et al., 2023) supports diverse latent-space motion tasks via a motion variational autoencoder (VAE). MotionGPT3 (Zhu et al., 2025) adopts a bimodal motion–language framework inspired by Mixture-of-Transformers (MoT), modeling motion in a continuous latent space by separate model parameters, enabling cross-modal interaction and multimodal scaling.

Despite these advances, content governance and safety remain under-addressed. Most works assume benign inputs and do not sanitize harmful intents. PhysDiff (Yuan et al., 2023) emphasize physical plausibility. ReinDiffuse (Han et al., 2025) uses reinforcement learning enhanced diffusion to constrain realism and safety. Recent efforts begin to target safety explicitly. Method (Bao et al., 2025) integrates a vision–language model (VLM) with confidence-based structured prompting and fallback strategies for socially appropriate real-time motion. Latent Code Replacement (LCR) (De Matteis et al., 2025) is a training-free unlearning approach that censors unsafe behaviors by replacing toxic-correlated entries in the discrete VQ codebook, without updating model weights. However, discrete token pipelines can introduce information loss and reduced smoothness, and reusing VQ tokens across benign prompts risks distribution drift when swapping codes. In contrast, our method operates in a continuous latent space and selectively forgets unsafe motion knowledge via two-stage unlearning, mitigating distribution shift while preserving benign performance.

Machine unlearning and trustworthy AI. Machine unlearning removes the influence of specified data from trained models (Cao and Yang, 2015), with *exact* (Bourtoule et al., 2021) and *approximate* (Pan et al., 2023; Liu et al., 2024a) variants. For diffusion safety, the method (Gandikota et al., 2023) finetunes to erase targeted visual concepts. Recent advance (Chen et al., 2025) removes target subspaces leveraging the model's embedding space. These approaches cast unsafe content mitigation as concept erasure via model editing, reducing the ability to produce disallowed outputs.

3 The Proposed Method

3.1 Overview

SafeMo compromises *(i) SafeMoEngine*, an LLM-agent pipeline that curates a prompt toxicity classifier, a safety-aware dual-strategy rewriter, producing safe text–motion datasets, and *(ii) Minimal Motion Unlearning* (MMU), a two-stage unlearning scheme on a continuous DiP (Tevet et al., 2025) backbone that localizes unsafe capability into a LoRA (Hu et al., 2022) subspace and removes it via inference-time negation. Upon a general (a mix of safe and unsafe intents) dataset, MMU finetunes a general-purpose T2M model using LoRA to selectively suppress unsafe motion semantics while preserving benign text-to-motion utility.

3.2 SafeMoEngine: LLM-Agent Pipeline for Safety Refinement

SafeMoEngine synthesizes safety-refined text–motion pairs via classify–rewrite–regenerate, as shown in Fig. 2. First, a violence-aware text classifier agent divided prompts into three different levels: *(i)* level 1: safe, not harmful content, which means the texts do not have any semantic toxic intent related to violence, crime, etc.; *(ii)* level 2: risky, partially harmful content, those containing violence-related motion in parts of its description; *(iii)* level 3: unsafe, which are toxic or violence, crime-related content as a whole. We then create a level-based strategy to alter the texts using separate rule-enhanced few-shot guided rewriting agents to intently positive ones, while keeping the altered descriptions with similar semantics. For example, *a man punches someone with his right fist*, will be modified to a description like *a man waves friendly with his right hand.* For level 2, we apply a partial rewriting strategy: only alter the semantically toxic parts while keeping the other parts semantically unchanged. For the level 3 content, we apply a stronger prompt that the agent needs to modify the content to a whole new, positive one.

The rewritten texts are fed into two text-to-motion generators: a continuous domain based one, MotionFlow Transformer (Guo et al., 2025), and a discrete VQ-token one, MotionAgent (Wu et al.,

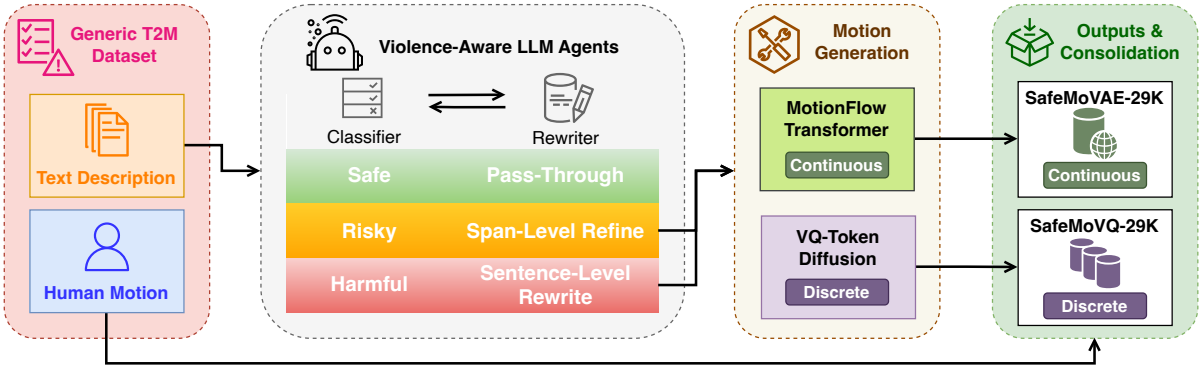

Figure 2: **Overview of the SafeMoEngine.** We first classify and rewrite harmful texts (Level 2 & 3), route Level 1 texts to original motions, compose text conditions and syhthesize motions via two generative models, to construct SafeMoVAE-29K and SafeMoVQ-29K, respectively.

2024), which produce non-violent substitute motions conditioned on the rewritten descriptions. We convert the generated motions to the standard HumanML3D (Guo et al., 2022a) representation and replace only the unsafe subset with these substitutes. After that, we obtain two versions of safe motion datasets, *SafeMoVQ-29K* and *SafeMoVAE-29K*, being in discrete and continuous fashion, respectively.

To the best of our knowledge, our datasets SafeMoVAE-29K, along with its discrete version, is the first text-to-motion dataset explicitly constructed for the *safe text-to-motion unlearning* setting, providing aligned continuous and discrete revised text–motion pairs through targeted rewriting of unsafe descriptions and regeneration of corresponding motions.

3.3 Minimal Motion Unlearning

We propose a Minimal Motion Unlearning (MMU) method in text-to-motion diffusion models, with a diffusion planner (DiP) (Tevet et al., 2025) as the backbone. This method consists of two stages, as shown in Figure 3: *(i)* Harmful Knowledge Absorption isolates unsafe behaviors by amplifying unsafe generation capability while intentionally reducing benign performance, yielding a harmful task vector that captures unsafe generation capability with minimal benign utility; and *(ii)* Harmful Knowledge Negation subtracts this harmful update from the base model at inference with a prompt-class-aware scaling factor α, suppressing unsafe semantics while preserving benign generation.

DiP backbone and notation. The DiP model is a text-conditioned auto-regressive diffusion with a transformer decoder backbone. The DiP can predict the clean motion x^{pred} from a prefix x^{prefix} and the noisy motion prediction x_t^{pred}, along with the diffusion step t, and a text prompt as a condition, at each step $t \in [0, T]$. The model also supports optional target-location conditioning, but we disable it in this work to avoid confounding control signals and ensure fair comparison with baselines. The text tokens $C_{\text{text}} \in \mathbb{R}^{N_{\text{tokens}} \times d}$ are first encoded by a fixed instance of DistilBERT (Sanh et al., 2019), and then be coordinated dimensions by a learned linear layer, after which they are injected through the cross-attention blocks in all transformer layers. We denote the model parameters by θ, the *base model* by θ_0, and the harm-tuned model, *bad model*, by θ_{bad}, and the obtained *safe model* by θ_{safe}. Sampling follows DDPM-style iterative denoising (Ho et al., 2020) with a single-step prediction head of $\hat{x_0}^{\text{pred}}$ per step. In our method, we use the LLM-based classifier agent in SafeMoEngine to split the text prompts into a safe set (level 1) and an unsafe set (level 2 and level 3), which are denoted by $\mathcal{S}$ and $\mathcal{U}$ respectively.

Harmful knowledge absorption. This stage updates only LoRA parameters to obtain a harmful task vector $\Delta\theta$ in a low-rank subspace, which we later negate at inference. Inspired by the Selective Knowledge negation Unlearning (SKU) technique (Liu et al., 2024b) on LLMs, we design a synchronized two-stream training process: an unsafe stream optimizing the harmful loss $\mathcal{L}_{\text{harm}}$ in *guided distortion* module (GD), and the random decoupling loss $\mathcal{L}_{\text{dec}}$ in *random decoupling* module (RD), and a safe stream optimizing through negative preservation divergence $\mathcal{L}_{\text{pres}}$ in *preservation divergence* module (PD). Given an unsafe batch with length mask $m \in \{0, 1\}^{B \times T}$, the model predicts

Table 1: Statistics of compared motion–language datasets and our dataset. “Quantity" reports counts of motion clips and text descriptions. “Supported Tasks" specifies the applicable downstream tasks. Abbreviations: T2M (Text-to-Motion), A2M (Audio-to-Motion), SMU (Safe Motion Unlearning), PE (Pose Estimation), MR (Mesh Recovery). “Content" distinguishes general (a mix of safe and unsafe intents) versus safe-refined data.

Dataset	Quantity		Supported Tasks	Content			
	Motion	Text		General Motion	General Text	Refined Safe Motion	Refined Safe Text
HumanML3D (Guo et al., 2022a)	14.6K	44.9K	T2M	✓	✓	✗	✗
KIT-ML (Plappert et al., 2016)	3.9K	6.3K	T2M	✓	✓	✗	✗
Motion-X (Lin et al., 2023)	81.1K	81.1K	T2M, MR	✓	✓	✗	✗
Motion-X++ (Zhang et al., 2025a)	120.5K	120.5K	T2M, MR, PE, A2M	✓	✓	✗	✗
SafeMoVQ-29K	17.2K	46.2K	T2M, SMU	✓	✓	✓	✓
SafeMoVAE-29K	17.2K	46.2K	T2M, SMU	✓	✓	✓	✓

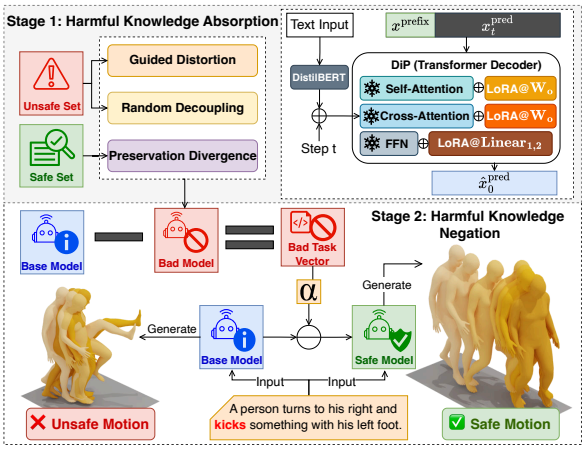

Figure 3: **Overview of SafeMo.** *Stage 1* (top): the unsafe stream optimizes through a harmful motion-specific loss and a random decoupling strategy, while the safe stream applies a negative preservation divergence. Only LoRA adapters on DiP are updated to obtain the pure harmful task vector. *Stage 2* (bottom): we negate the learned harmful task vector via a prompt-class-aware α, such that the model suppresses unsafe behaviors on unsafe prompts while preserving performance on safe prompts.

the clean motion $\hat{\mathbf{x}}_0 = \{\hat{\mathbf{p}}_t\}_{t=1}^{T} = f_\theta(x_t, t, \mathcal{C}_{\text{text}})$ to match the target motion $\mathbf{x}_0 = \{\mathbf{p}_t\}_{t=1}^{T}$, where $\mathbf{p}_t, \hat{\mathbf{p}}_t \in \mathbb{R}^{3J}$ stack the J joint 3D coordinates. The motion-specific harmful loss combines kinematic terms and a text-motion alignment term in GD:

$$\begin{aligned}\mathcal{L}_{\text{harm}} = \lambda_{\text{mpjpe}}\ &\text{MPJPE}(\hat{\mathbf{x}}_0, \mathbf{x}_{\text{tgt}}; m) \\ &+ \lambda_{\text{vel}}\ \mathcal{L}_{\text{vel}}(\hat{\mathbf{x}}_0, \mathbf{x}_{\text{tgt}}; m) \\ &+ \lambda_{\text{acc}}\ \mathcal{L}_{\text{acc}}(\hat{\mathbf{x}}_0, \mathbf{x}_{\text{tgt}}; m) \\ &+ \lambda_{\text{foot}}\ \mathcal{L}_{\text{foot}}(\hat{\mathbf{x}}_0; m) \\ &+ \lambda_{\text{text}}\ \mathcal{L}_{\text{text}\leftrightarrow\text{mo}}(\hat{\mathbf{x}}_0, \mathcal{C}_{\text{text}}).\end{aligned} \tag{1}$$

Instead of the token-level cross-entropy in SKU, we optimize a weighted sum of motion-specific objectives for harmful knowledge absorption. In the guided distortion (GD) module, the harmful loss (Eq. 1) combines kinematic constraints and text–motion alignment: the masked reconstruction term is $\text{MPJPE}(\hat{\mathbf{x}}_0, \mathbf{x}_{\text{tgt}}; m) = \big(\sum_t m_t \|\hat{\mathbf{p}}_t - \mathbf{p}_t\|_2\big)/\big(\sum_t m_t + \varepsilon\big)$; temporal smoothness is enforced by derivative matching with additional spectral emphasis, $\mathcal{L}_{\text{vel}} = \big(\sum_t m_t \|\Delta\hat{\mathbf{p}}_t - \Delta\mathbf{p}_t\|_2\big)/\big(\sum_t m_t + \varepsilon\big) + \mathcal{S}_{\text{vel}}$ and $\mathcal{L}_{\text{acc}} = \big(\sum_t m_t \|\Delta^2\hat{\mathbf{p}}_t - \Delta^2\mathbf{p}_t\|_2\big)/\big(\sum_t m_t + \varepsilon\big) + \mathcal{S}_{\text{acc}}$, where $\Delta\mathbf{p}_t = \mathbf{p}_t - \mathbf{p}_{t-1}$ and $\Delta^2\mathbf{p}_t = \mathbf{p}_t - 2\mathbf{p}_{t-1} + \mathbf{p}_{t-2}$ (and analogously for $\hat{\mathbf{p}}_t$). To discourage foot sliding, we use $\mathcal{L}_{\text{foot}} = \big(\sum_t m_t\ \text{mean}_{j\in\mathcal{F}} |\hat{\mathbf{p}}_t^{(j)} - \hat{\mathbf{p}}_{t-1}^{(j)}|\big)/\big(\sum_t m_t + \varepsilon\big)$, and we align text and motion embeddings via the symmetric contrastive objective from T2M (Zhang et al., 2023), $\mathcal{L}_{\text{text}\leftrightarrow\text{mo}} = \frac{1}{2}\big[\text{CE}(e_t e_m^\top/\tau, \text{Id}) + \text{CE}(e_m e_t^\top/\tau, \text{Id})\big]$. Notably, $\mathcal{S}_{\text{vel}}$ emphasizes higher-frequency errors in the velocity residual $r_t = \Delta\hat{\mathbf{p}}_t - \Delta\mathbf{p}_t$ by applying an rFFT over t and weighting magnitudes with a logarithmic prior, $\mathcal{S}_{\text{vel}} = \text{mean}\big(|\mathcal{F}(r)| \cdot \log(1 + 9\nu)\big)$, where $\nu \in [0, 1]$ denotes normalized frequency bins broadcasted over joints and the mean is taken over valid time bins under m. $\mathcal{S}_{\text{acc}}$ is defined analogously using $r_t = \Delta^2\hat{\mathbf{p}}_t - \Delta^2\mathbf{p}_t$. The contributions of the above loss terms are examined via auxiliary analyses and ablation studies in Appendix E.

In the RD module, we adopt the idea of misalignment from the SKU for LLMs, but design a sequence perturbation strategy for our task, applied at the sequence level. To broaden the harmful prototypes without heavy data editing, we adopt temporal segments shuffling or time-reversing to each unsafe motion sequence to obtain a decoupled motion $\tilde{\mathbf{x}}_{\text{tgt}}$. The corresponding prefix in the condition is synchronously replaced to remain consistent with the decoupled target. A single forward pass then computes

$$\begin{aligned}\mathcal{L}_{\text{dec}} = \lambda_{\text{mpjpe}}\ &\text{MPJPE}(\hat{\mathbf{x}}_0^{\text{mix}}, \tilde{\mathbf{x}}_{\text{tgt}}; m) \\ &+ \lambda_{\text{vel}}\ \mathcal{L}_{\text{vel}}(\hat{\mathbf{x}}_0^{\text{mix}}, \tilde{\mathbf{x}}_{\text{tgt}}; m) \\ &+ \lambda_{\text{acc}}\ \mathcal{L}_{\text{acc}}(\hat{\mathbf{x}}_0^{\text{mix}}, \tilde{\mathbf{x}}_{\text{tgt}}; m),\end{aligned} \tag{2}$$

where a mixed noisy batch $\mathbf{x}_t^{\text{mix}}$ is constructed by replacing only the unsafe instances with their decoupled-noised targets and synchronously replace the corresponding prefixes. Resulting predictions for the unsafe instances are $\hat{\mathbf{x}}_0^{\text{mix}}$.

On safe batches we encourage the model to diverge from a frozen baseline snapshot f_{θ_0} at a pooled representation level in PD module. Unlike SKU for LLMs where a token-level negative KL is natural, continuous motion diffusion lacks a comparable discrete likelihood target; we therefore use a stable representation-level proxy, $\mathcal{L}_{\text{pres}}$, and validate it with ablations in Appendix E. Let $\mathbf{z}_{\text{cur}} = \text{Pool}(f_\theta(\mathbf{x}_t, t, \mathcal{C}))$, and $\mathbf{z}_{\text{base}} = \text{Pool}(f_{\theta_0}(\mathbf{x}_t, t, \mathcal{C}))$, where $\text{Pool}(\cdot)$ denotes temporal-averages joint features. To make the divergence robust to light temporal perturbations, we design a safe-only decoupling term. For each safe sequence $\mathbf{x}_0$ we create a decoupled target $\tilde{\mathbf{x}}_0$ by either segment permutation or time reversal at the sequence level, uniformly chosen. Using the same diffusion timestep t and noise ε as the main safe batch, we form $\mathbf{x}_t^{\text{dec}} = q(\tilde{\mathbf{x}}_0, t, \varepsilon)$ and $\mathcal{C}^{\text{dec}} = \text{SyncPrefix}(\mathcal{C}; \tilde{\mathbf{x}}_0)$, to obtain the decoupled features $\mathbf{z}_{\text{cur}}^{\text{dec}} = \text{Pool}\big(f_\theta(\mathbf{x}_t^{\text{dec}}, t, \mathcal{C}^{\text{dec}})\big)$, $\mathbf{z}_{\text{base}}^{\text{dec}} = \text{Pool}\big(f_{\theta_0}(\mathbf{x}_t^{\text{dec}}, t, \mathcal{C}^{\text{dec}})\big)$, where $\text{SyncPrefix}(\cdot)$ replaces the motion prefix so that the condition matches the decoupled target. The negative preservation divergence is then

$$\begin{aligned}\mathcal{L}_{\text{pres}} = &-\gamma\|\mathbf{z}_{\text{cur}} - \mathbf{z}_{\text{base}}\|_2^2 \\ &- (1-\gamma)\|\mathbf{z}_{\text{cur}}^{\text{dec}} - \mathbf{z}_{\text{base}}^{\text{dec}}\|_2^2,\end{aligned} \tag{3}$$

where $\gamma \in [0, 1]$. This negative term makes minimizing $\mathcal{L}_{\text{pres}}$ result in deviations from the baseline on benign prompts, including their decoupled variants, enforcing deliberate deviation on benign prompts during the first stage, enabling the negation in the next stage to restore utility. Let $\mathcal{U}_t$ and $\mathcal{S}_t$ denote unsafe and safe sets in a batch at step t respectively. The overall stage 1 objective is then formed by

$$\begin{aligned}\theta_{t+1} \leftarrow \theta_t - \eta\nabla_\theta\big(&W_{\text{harm}}\mathcal{L}_{\text{harm}}(\mathcal{U}_t) \\ &+ W_{\text{dec}}\mathcal{L}_{\text{dec}}(\mathcal{U}_t) + W_{\text{pres}}\mathcal{L}_{\text{pres}}(\mathcal{S}_t)\big).\end{aligned} \tag{4}$$

LoRA subspace and injection points. We replace selected linear layers by LoRA modules with rank r, scaling α within a dropout rate $p_{\text{LoRA}}^{\text{dropout}}$. We attach rank-r LoRA adapters to the attention output and the FFN in and out projections; trainable parameters are only the LoRA matrices $A \in \mathbb{R}^{r\times d}$ and $B \in \mathbb{R}^{d_{\text{out}}\times r}$, while all backbone parameters are frozen; hence, updates are confined to the low-rank subspace.

Harmful knowledge negation. The harmful task vector is obtained by $\Delta\theta = \theta_{\text{bad}} - \theta_0$ in the LoRA subspace. Let the α denote the scaling weight for $\mathcal{U}$ and $\mathcal{S}$, the updated safe model is then obtained by $\theta_{\text{safe}} = \theta_0 - \alpha\,\Delta\theta$ at inference. We design *SafeMo-Static* and *SafeMo-Gated* with different scaling strategies. SafeMo-Static applies a fixed α for all text prompts without any external agent model, providing a light and offline-fashioned method for balanced performance on both the unsafe set and safe set, similar to the selective knowledge unlearning strategy (Liu et al., 2024b). SafeMo-Gated applies a larger α on unsafe prompts and smaller α on safe prompts with SafeMoEngine's classifier agent's decision on the toxicity of the input text prompt, maximizing the effect of toxic motion unlearning while minimizing the influence on benign tasks. The algorithm of MMU can be found in Appendix A.

4 Experiments

4.1 Experiment Setup

Datasets and Evaluation Metrics. We evaluate our model's performance on the HumanML3D (Guo et al., 2022a) and Motion-X (Lin et al., 2023). We report standard T2M metrics: R-precision (R@k) for text–motion retrieval accuracy in a shared embedding space, Fréchet Inception Distance (FID) to quantify distributional similarity between generated and ground-truth motion features, and Diversity as the average pairwise distance among generated features to reflect intra-set variability. Additional details in Appendix B.

Implementation details. Our framework is trained on a single NVIDIA GeForce RTX 3090 GPU using PyTorch. The LLM agents used in SafeMoEngine are on top of the Qwen2.5-7B-Instruct (Bai et al., 2023) with few-shot rule-enhanced prompt templates. We adopted DiP (Tevet et al., 2025) as our base text-to-motion model in the MMU stage, which is an 8-layer transformer decoder with a latent dimension size of 512 and 4 attention heads. The text encoder is a fixed instance of DistilBERT (Sanh et al., 2019). We

Table 2: **Results on HumanML3D dataset.** Method D_r reports performances for the model trained on a toxicity-free dataset. Method FT shows the results of fine-tuning the model on the toxicity-free dataset. **SafeMo-Static** denotes our fixed-α model without an external classifier, lightweight and offline. **SafeMo-Gated** denotes the classifier-agent-guided α-gating model. Diversity is reported for reference. *Note:* Rows marked with † are reported from (De Matteis et al., 2025) due to unavailable implementation and checkpoints at the time of submission.

	Forget Set			Retain Set		
	FID ↑	Diversity	R@1 ↓	FID ↓	Diversity	R@1 ↑
MoMask $D_r^\dagger$	$13.644_{\pm.365}$	$7.611_{\pm.088}$	$0.129_{\pm.004}$	$0.093_{\pm.003}$	$10.059_{\pm.080}$	$0.291_{\pm.001}$
MoMask† (Guo et al., 2024)	$0.956_{\pm.084}$	$6.146_{\pm.092}$	$0.159_{\pm.005}$	$0.064_{\pm.002}$	$10.143_{\pm.081}$	$0.290_{\pm.001}$
MoMask $FT^\dagger$	$1.589_{\pm.116}$	$6.439_{\pm.088}$	$0.148_{\pm.006}$	$0.088_{\pm.002}$	$10.143_{\pm.097}$	$0.280_{\pm.001}$
MoMask w/ UCE†	$25.039_{\pm.442}$	$8.693_{\pm.067}$	$0.105_{\pm.005}$	$0.395_{\pm.008}$	$10.194_{\pm.091}$	$0.257_{\pm.001}$
MoMask w/ RECE†	$58.487_{\pm.899}$	$8.591_{\pm.073}$	$0.069_{\pm.004}$	$12.557_{\pm.092}$	$9.612_{\pm.147}$	$0.121_{\pm.001}$
MoMask w/ LCR†	$12.434_{\pm.303}$	$6.580_{\pm.066}$	$0.133_{\pm.004}$	$0.077_{\pm.002}$	$10.106_{\pm.086}$	$0.287_{\pm.001}$
BAMM $D_r^\dagger$	$15.604_{\pm.334}$	$7.688_{\pm.074}$	$0.122_{\pm.005}$	$0.566_{\pm.015}$	$10.092_{\pm.093}$	$0.279_{\pm.001}$
BAMM† (Pinyoanuntapong et al., 2024)	$1.353_{\pm.107}$	$6.202_{\pm.089}$	$0.164_{\pm.007}$	$0.135_{\pm.004}$	$10.118_{\pm.100}$	$0.302_{\pm.001}$
BAMM $FT^\dagger$	$1.443_{\pm.118}$	$6.224_{\pm.098}$	$0.161_{\pm.006}$	$0.163_{\pm.005}$	$10.109_{\pm.086}$	$0.301_{\pm.002}$
BAMM w/ UCE†	$57.953_{\pm.893}$	$9.482_{\pm.073}$	$0.074_{\pm.003}$	$4.296_{\pm.063}$	$9.654_{\pm.080}$	$0.184_{\pm.001}$
BAMM w/ RECE†	$34.367_{\pm.484}$	$7.740_{\pm.073}$	$0.061_{\pm.004}$	$13.310_{\pm.118}$	$8.470_{\pm.094}$	$0.122_{\pm.001}$
BAMM w/ LCR†	$9.712_{\pm.214}$	$6.502_{\pm.077}$	$0.136_{\pm.005}$	$0.140_{\pm.005}$	$10.068_{\pm.102}$	$0.299_{\pm.001}$
DiP D_r	$3.002_{\pm.108}$	$7.272_{\pm.106}$	$0.249_{\pm.005}$	$0.301_{\pm.028}$	$9.248_{\pm.091}$	$0.476_{\pm.009}$
DiP (Tevet et al., 2025)	$0.440_{\pm.046}$	$7.331_{\pm.100}$	$0.308_{\pm.007}$	$0.250_{\pm.025}$	$9.274_{\pm.089}$	$0.482_{\pm.006}$
DiP FT	$1.399_{\pm.100}$	$7.527_{\pm.093}$	$0.271_{\pm.010}$	$0.207_{\pm.024}$	$9.337_{\pm.073}$	$0.459_{\pm.007}$
SafeMo-Static	$10.288_{\pm.055}$	$6.993_{\pm.072}$	$0.168_{\pm.002}$	$2.224_{\pm.002}$	$8.606_{\pm.176}$	$0.335_{\pm.003}$
SafeMo-Gated	$31.147_{\pm.221}$	$4.986_{\pm.084}$	$0.086_{\pm.004}$	$0.407_{\pm.003}$	$9.404_{\pm.401}$	$0.386_{\pm.002}$

follow the base model's setting for generation, with 10 diffusion steps, prefix length $N_p = 20$ and generation length $N_g = 40$.

4.2 Results and Analysis

Baselines and comparisons. We compare SafeMo with the prior state-of-the-art text-to-motion unlearning method LCR (De Matteis et al., 2025), together with their motion-adapted UCE (Gandikota et al., 2024) and RECE (Gong et al., 2024) baselines introduced in the same work, on HumanML3D (Guo et al., 2022a) and Motion-X (Lin et al., 2023) datasets. For comparisons with LCR and UCE/RECE (De Matteis et al., 2025), we follow their keyword-based partitioning protocol to build the forget split for unsafe prompts and the retain split for safe prompts for a fair evaluation. As official implementations and checkpoints for LCR and their T2M implementations of UCE and RECE were unavailable at the time of submission, we report the corresponding baseline numbers from (De Matteis et al., 2025), marked by † in Table 2. Unlike LCR, which measures detoxification by how closely the forget-split performance approaches a clean-data model, we evaluate unlearning as capability suppression on unsafe prompts while retaining on safe prompts: on the forget split, stronger unlearning reduces text–motion retrieval R@1 and increase distributional distance to the original unsafe motions FID, while on the retain split we use standard T2M quality criteria FID ↓, R@1 ↑. We report Diversity as a sanity check.

Quantitative results. SafeMo-Static uses a fixed scaling factor $\alpha_{\text{static}} = 1.0$ for all prompts and requires no external classifier. SafeMo-Gated uses the SafeMoEngine toxicity classifier to apply $\alpha_{\text{gated}}^{\text{unsafe}} = 2.0$ for unsafe prompts and $\alpha_{\text{gated}}^{\text{safe}} = 0.05$ for benign prompts. α-sweep sensitivity analysis is in Appendix E.3. Results on HumanML3D are shown in Table 2. On the retain set, SafeMo-Static achieves R@1 $= 0.335$, outperforming MoMask/BAMM w/ LCR (0.287/0.299), and SafeMo-Gated further improves retain utility to R@1 $= 0.386$. On the forget set, SafeMo-Gated enforces substantially stronger suppression than LCR, increasing FID by +150.5%/+220.7% and reducing R@1 by -35.3%/-36.8% compared to MoMask/BAMM w/ LCR, respectively, while maintaining high retain-set utility. We note that absolute FID values can vary across discrete and diffusion backbones under the same evaluator (e.g., continuous DiP may exhibit higher retrieval but

Table 3: **Ablation study of three modules in MMU stage-1.** Results on HumanML3D. On unsafe sets, higher FID and lower retrieval (R@K) indicate stronger forgetting; on the safe set, lower FID and higher retrieval indicate better utility. Diversity is reported for reference.

	Unlearned Unsafe Set					**Unseen** Unsafe Set					**Unseen** Safe Set				
	FID↑	Div.	R@1↓	R@2↓	R@3↓	FID↑	Div.	R@1↓	R@2↓	R@3↓	FID↓	Div.	R@1↑	R@2↑	R@3↑
SafeMo ($\alpha = 0.0$)	1.7197	7.3746	0.2517	0.3914	0.4969	2.3050	7.5191	0.2365	0.3896	0.5052	0.5232	9.3375	0.3755	0.5599	0.6732
SafeMo-Static	8.0235	6.8083	0.2016	0.3164	0.4043	9.0499	6.8880	0.1958	0.3167	0.3937	2.5539	8.7060	0.3172	0.4935	0.6052
SafeMo-Static w/o GD	5.2830	6.9963	0.2188	0.3449	0.4377	5.7963	6.9634	0.2104	0.3333	0.4208	1.4697	8.8189	0.3347	0.5144	0.6295
SafeMo-Static w/o RD	5.7285	7.1058	0.2195	0.3378	0.4307	6.7159	7.2363	0.1990	0.3375	0.4375	1.9663	9.0015	0.3432	0.5263	0.6379
SafeMo-Static w/o PD	8.9693	6.6601	0.1960	0.3092	0.3962	10.5409	6.7565	0.1896	0.3000	0.3740	2.9845	8.6333	0.3178	0.4878	0.6040
SafeMo-Gated	28.0806	5.0169	0.0947	0.1630	0.2168	28.0574	4.8520	0.0865	0.1542	0.2104	0.5355	9.3224	0.3775	0.5628	0.6769
SafeMo-Gated w/o GD	46.9030	2.8490	0.0544	0.1055	0.1543	45.4955	2.7078	0.0615	0.1083	0.1542	0.5248	9.3258	0.3760	0.5624	0.6768
SafeMo-Gated w/o RD	21.3313	5.6771	0.1449	0.2351	0.3002	21.3717	5.5564	0.1469	0.2292	0.2760	0.5385	9.3204	0.3775	0.5625	0.6742
SafeMo-Gated w/o PD	31.0764	4.7098	0.0926	0.1497	0.2020	31.3418	4.5750	0.1042	0.1688	0.2073	0.5380	9.3241	0.3783	0.5631	0.6761

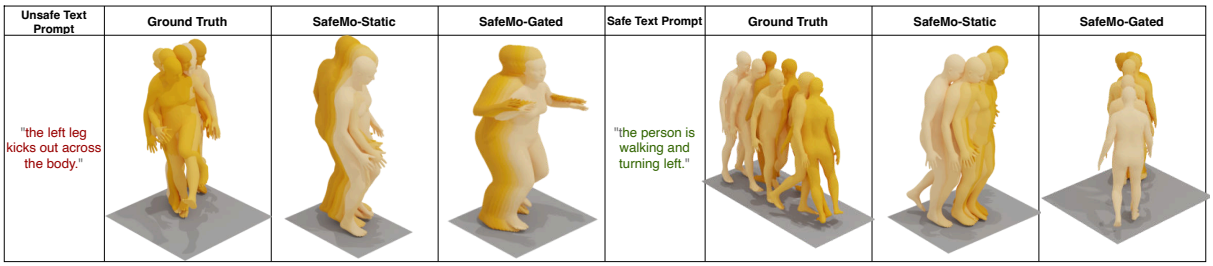

Figure 4: **Qualitative results.** More results can be found in Appendix D.

higher FID than discrete models such as MoMask and BAMM), so we emphasize the forget–retain trade-off and within-backbone references. To contextualize architecture differences, we report DiP D_r/FT trained on the toxicity-free subset following LCR (De Matteis et al., 2025)'s practice for clean-data references, showing that simply training on filtered data under the same split does not fully eliminate unsafe semantics. The same trend holds on Motion-X, results are in Appendix C.

Qualitative results. Figure 4 shows qualitative results, illustrating stronger forgetting on unsafe intents and preserved fidelity on benign prompts. More qualitative results, discussions can be found in Appendix D.

4.3 Ablation Studies

We ablate the three modules, Guided Distortion (GD), Random Decoupling (RD), and Preservation Divergence (PD), to assess their roles in the safety-utility tradeoff. Results are reported in Table 3. Following SKU (Liu et al., 2024b), we disentangle in-distribution forgetting and generalization by evaluating on unlearned and unseen unsafe prompts. For the gated regime, toxicity is detected at inference by the SafeMoEngine classifier to select α. Removing GD weakens forgetting in the static regime, with unsafe FID -34.2% on unlearned, and -36.0% on unseen data, indicating GD is the primary driver for aligning the edited direction with harmful motion patterns. The gated regime yields higher unsafe FID with reduced diversity suggesting potential over-suppression and near-identical safe-set utility, showing gating can dominate the forgetting strength. Removing RD reduces unsafe FID in both static and gated settings, while slightly improving safe-set metrics in the static regime, reflecting a utility–forgetting trade-off. PD mainly protects benign utility without external gating. Without PD, safe-set FID increases 16.9% and retrieval degrades under static while the impact under gating is marginal. Additional ablations on loss-terms, LoRA ranks are provided in Appendix E.

5 Conclusion

In this work, we introduce an innovative continuous domain-based T2M unlearning framework, SafeMo, for trustworthy motion generation. We introduce the first safety-aligned T2M dataset, SafeMoVAE-29K, along with its discrete version, to facilitate future research and standardized benchmarking in human motion unlearning. The proposed absorb-then-negate machine unlearning strategy designed for text-to-motion models, Minimal Motion Unlearning, enables selective knowledge unlearning on unsafe motions while preserving benign utility on safe prompts. Extensive experiments on HumanML3D and Motion-X datasets demonstrate that our model achieves SOTA performance on human motion unlearning.

6 Limitations

To the best of our knowledge, SafeMo is the first continuous latent space framework to investigate text-to-motion unlearning. On unsafe prompts, our goal is *semantic removal*, i.e., preventing the model from expressing the unsafe motion semantics, rather than producing a high-fidelity safe substitute motion. However, this safety-first operating point may lead to over-suppression for some unsafe prompts, e.g., stationary-like patterns, and can exacerbate kinematic artifacts such as foot-skating, especially under larger negation gating scales. Additional discussions on failure modes and future work are provided in Appendix G.

7 Ethical Considerations

This work contributes to the T2M safety alignment. The proposed datasets are constructed by identifying harmful prompts from HumanML3D and systematically rewriting them into safe alternatives using LLMs, followed by regenerating them into safe motions by both discrete and continuous methods. While the dataset includes the original harmful prompts to serve as a baseline for unlearning tasks, its primary purpose is to facilitate the development if safety constraints and remediation strategies. We strictly limit the use of the datasets to research purposes.

References

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, and 1 others. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.

Lingfan Bao, Yan Pan, Tianhu Peng, Dimitrios Kanoulas, and Chengxu Zhou. 2025. Hierarchical intention-aware expressive motion generation for humanoid robots. *arXiv preprint arXiv:2506.01563*.

Lucas Bourtoule, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. 2021. Machine unlearning. In *2021 IEEE symposium on security and privacy (SP)*, pages 141–159. IEEE.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Yinzhi Cao and Junfeng Yang. 2015. Towards making systems forget with machine unlearning. In *2015 IEEE symposium on security and privacy*, pages 463–480. IEEE.

Huiqiang Chen, Tianqing Zhu, Linlin Wang, Xin Yu, Longxiang Gao, and Wanlei Zhou. 2025. Safe and reliable diffusion models via subspace projection. *arXiv preprint arXiv:2503.16835*.

Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. 2023. Executing your commands via motion diffusion in latent space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18000–18010.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, and 1 others. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113.

Edoardo De Matteis, Matteo Migliarini, Alessio Sampieri, Indro Spinelli, and Fabio Galasso. 2025. Human motion unlearning. *arXiv preprint arXiv:2503.18674*.

Hao Fei, Shengqiong Wu, Wei Ji, Hanwang Zhang, and Tat-Seng Chua. 2024. Dysen-vdm: Empowering dynamics-aware text-to-video diffusion with llms. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7641–7653.

Rohit Gandikota, Joanna Materzynska, Jaden Fiotto-Kaufman, and David Bau. 2023. Erasing concepts from diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2426–2436.

Rohit Gandikota, Hadas Orgad, Yonatan Belinkov, Joanna Materzyńska, and David Bau. 2024. Unified concept editing in diffusion models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5111–5120.

Chao Gong, Kai Chen, Zhipeng Wei, Jingjing Chen, and Yu-Gang Jiang. 2024. Reliable and efficient concept erasure of text-to-image diffusion models. In *European Conference on Computer Vision*, pages 73–88. Springer.

Chuan Guo, Yuxuan Mu, Muhammad Gohar Javed, Sen Wang, and Li Cheng. 2024. Momask: Generative masked modeling of 3d human motions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1900–1910.

Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. 2022a. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5152–5161.

Chuan Guo, Xinxin Zuo, Sen Wang, and Li Cheng. 2022b. Tm2t: Stochastic and tokenized modeling for the reciprocal generation of 3d human motions and texts. In *European Conference on Computer Vision*, pages 580–597. Springer.

Ziyan Guo, Zeyu Hu, De Wen Soh, and Na Zhao. 2025. Motionlab: Unified human motion generation and editing via the motion-condition-motion paradigm. *arXiv preprint arXiv:2502.02358*.

Gaoge Han, Mingjiang Liang, Jinglei Tang, Yongkang Cheng, Wei Liu, and Shaoli Huang. 2025. Reindiffuse: Crafting physically plausible motions with reinforced diffusion model. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2218–2227. IEEE.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.

Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. 2022. Editing models with task arithmetic. *arXiv preprint arXiv:2212.04089*.

Jing Lin, Ailing Zeng, Shunlin Lu, Yuanhao Cai, Ruimao Zhang, Haoqian Wang, and Lei Zhang. 2023. Motion-x: A large-scale 3d expressive whole-body human motion dataset. *Advances in Neural Information Processing Systems*.

Zheyuan Liu, Guangyao Dou, Eli Chien, Chunhui Zhang, Yijun Tian, and Ziwei Zhu. 2024a. Breaking the trilemma of privacy, utility, and efficiency via controllable machine unlearning. In *Proceedings of the ACM Web Conference 2024*, pages 1260–1271.

Zheyuan Liu, Guangyao Dou, Zhaoxuan Tan, Yijun Tian, and Meng Jiang. 2024b. Towards safer large language models through machine unlearning. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 1817–1829, Bangkok, Thailand. Association for Computational Linguistics.

Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. 2015. Smpl: a skinned multi-person linear model. *ACM Transactions on Graphics (TOG)*, 34(6):1–16.

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

Chao Pan, Eli Chien, and Olgica Milenkovic. 2023. Unlearning graph classifiers with limited data resources. In *Proceedings of the ACM Web Conference 2023*, pages 716–726.

Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. 2019. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10975–10985.

Ekkasit Pinyoanuntapong, Muhammad Usama Saleem, Pu Wang, Minwoo Lee, Srijan Das, and Chen Chen. 2024. Bamm: Bidirectional autoregressive motion model. In *European Conference on Computer Vision*, pages 172–190. Springer.

Matthias Plappert, Christian Mandery, and Tamim Asfour. 2016. The kit motion-language dataset. *Big data*, 4(4):236–252.

Chengwei Qin, Aston Zhang, Zhuosheng Zhang, Jiaao Chen, Michihiro Yasunaga, and Diyi Yang. Is chatgpt a general-purpose natural language processing task solver? In *The 2023 Conference on Empirical Methods in Natural Language Processing*.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695.

Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22500–22510.

Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.

Guy Tevet, Sigal Raab, Setareh Cohan, Daniele Reda, Zhengyi Luo, Xue Bin Peng, Amit Haim Bermano, and Michiel van de Panne. 2025. CLoSD: Closing the loop between simulation and diffusion for multi-task character control. In *The Thirteenth International Conference on Learning Representations*.

Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations*.

Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. 2023. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations*.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, and 1 others. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Qi Wu, Yubo Zhao, Yifan Wang, Xinhang Liu, Yu-Wing Tai, and Chi-Keung Tang. 2024. Motion-agent: A conversational framework for human motion generation with llms. *arXiv preprint arXiv:2405.17013*.

Yuanshun Yao, Xiaojun Xu, and Yang Liu. 2024. Large language model unlearning. *Advances in Neural Information Processing Systems*, 37:105425–105475.

Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. 2023. Physdiff: Physics-guided human motion diffusion model. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16010–16021.

Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Yong Zhang, Hongwei Zhao, Hongtao Lu, Xi Shen, and Ying Shan. 2023. Generating human motion from textual descriptions with discrete representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14730–14740.

Yuhong Zhang, Jing Lin, Ailing Zeng, Guanlin Wu, Shunlin Lu, Yurong Fu, Yuanhao Cai, Ruimao Zhang, Haoqian Wang, and Lei Zhang. 2025a. Motion-x++: A large-scale multimodal 3d whole-body human motion dataset. *arXiv preprint arXiv:2501.05098*.

Zeyu Zhang, Hang Gao, Akide Liu, Qi Chen, Feng Chen, Yiran Wang, Danning Li, Rui Zhao, Zhenming Li, Zhongwen Zhou, and 1 others. 2024a. Kmm: Key frame mask mamba for extended motion generation. *arXiv preprint arXiv:2411.06481*.

Zeyu Zhang, Akide Liu, Qi Chen, Feng Chen, Ian Reid, Richard Hartley, Bohan Zhuang, and Hao Tang. 2024b. Infinimotion: Mamba boosts memory in transformer for arbitrary long motion generation. *arXiv preprint arXiv:2407.10061*.

Zeyu Zhang, Akide Liu, Ian Reid, Richard Hartley, Bohan Zhuang, and Hao Tang. 2024c. Motion mamba: Efficient and long sequence motion generation. In *European Conference on Computer Vision*, pages 265–282. Springer.

Zeyu Zhang, Yiran Wang, Wei Mao, Danning Li, Rui Zhao, Biao Wu, Zirui Song, Bohan Zhuang, Ian Reid, and Richard Hartley. 2025b. Motion anything: Any to motion generation. *arXiv preprint arXiv:2503.06955*.

Zeyu Zhang, Yiran Wang, Biao Wu, Shuo Chen, Zhiyuan Zhang, Shiya Huang, Wenbo Zhang, Meng Fang, Ling Chen, and Yang Zhao. 2024d. Motion avatar: Generate human and animal avatars with arbitrary motion. *arXiv preprint arXiv:2405.11286*.

Bingfan Zhu, Biao Jiang, Sunyi Wang, Shixiang Tang, Tao Chen, Linjie Luo, Youyi Zheng, and Xin Chen. 2025. Motiongpt3: Human motion as a second modality. *Preprint*, arXiv:2506.24086.

A Minimal Motion Unlearning (MMU) Algorithm

The complete procedure of Minimal Motion Unlearning (MMU) is outlined in Algorithm 1.

B Motion Representations, Datasets and Evaluation Metrics

B.1 Motion representations.

We follow MDM (Tevet et al.) and use the HumanML3D motion representation. At each frame n, a pose $p_n \in \mathbb{R}^F$ is $p_n = \big(r^a,\, r^x,\, r^z,\, r^y,\, j^p,\, j^r,\, j^v,\, f\big)$, where $r^a \in \mathbb{R}$ is the root angular velocity along the Z-axis, $r^x, r^z \in \mathbb{R}$ are the root linear velocities on the XY-plane, and $r^y \in \mathbb{R}$ is the root height. $j^p \in \mathbb{R}^{3(J-1)}$, $j^r \in \mathbb{R}^{6(J-1)}$, and $j^v \in \mathbb{R}^{3J}$ denote, respectively, the local joint positions, rotations (in the 6D continuous form), and velocities, all defined with respect to the root. $f \in \mathbb{R}^4$ are binary foot-contact labels for four foot joints (two per leg).

B.2 Datasets

We evaluate our model's performance on the HumanML3D (Guo et al., 2022a) and Motion-X (Lin et al., 2023) benchmark, which are widely used for text-to-motion tasks. HumanML3D contains 14.6k human motion sequences and 44.9k detailed text descriptions with pos tagging. Text and motion encoders are used in this benchmark to map text and motion to the same latent space, and learned using contrastive loss. Motion-X is a large-scale text–motion corpus aggregating motions from real-world and animated scenarios, including 15.6M whole-body poses and 81.1K motion clips annotations. It covers a broader action vocabulary such as daily activities, sports, and combat-related motions and pairs them with natural-language descriptions. According to findings in LCR (De Matteis et al., 2025), HumanML3D dataset contains 7.7% explicitly toxic human motions, while Motion-X has a higher percentage of 14.9%.

B.3 Metrics

We use R-precision, Fréchet Inception Distance (FID), Diversity to measure the effectiveness of

Algorithm 1 Minimal Motion Unlearning (MMU)

Require: Unsafe set $\mathcal{U}$, Safe set $\mathcal{S}$; base DiP model f_{θ_0}; LoRA config $(r, \alpha_{\text{LoRA}}, p_{\text{dropout}})$; loss weights $\{\lambda, \mu, \beta\}$; diffusion schedule.

Ensure: Task vector $\Delta\theta$ and safe model f_{θ^*}

Initialize $\theta \leftarrow \theta_0$; insert LoRA adapters at attention/FFN; freeze non-LoRA params.

Stage 1: Harmful Knowledge Absorption (training)

repeat

  Sample unsafe batch $\mathcal{U}_b$, safe batch $\mathcal{S}_b$, timesteps t, noise ε.

  for $i \in \mathcal{U}_b$ **do**

    Generate noisy $x_t^{(i)} \sim q(x_0^{(i)}, t, \varepsilon)$

    Predict $\hat{x}_0^{(i)} \leftarrow f_\theta(x_t^{(i)}, t, C^{(i)})$

    Compute $\mathcal{L}_{\text{harm}}^{(i)}$ using Eq. 1

    Build decoupled $\tilde{x}_{\text{tgt}}^{(i)}$ (segment shuffle / reverse)

    Sync prefix $\tilde{C}^{(i)} \leftarrow \text{SyncPrefix}(C^{(i)}, \tilde{x}_{\text{tgt}}^{(i)})$

    Predict $\hat{x}_0^{\text{mix}} \leftarrow f_\theta(x_t^{(i)}, t, \tilde{C}^{(i)})$

    Compute $\mathcal{L}_{\text{dec}}^{(i)}$ via Eq. 2

  end for

  for $j \in \mathcal{S}_b$ **do**

    Generate $x_t^{(j)} \sim q(x_0^{(j)}, t, \varepsilon)$

    Extract pooled features $z_{\text{cur}}^{(j)}, z_{\text{base}}^{(j)}$

    Create decoupled $\tilde{x}_0^{(j)}$ and synced prefix C^{dec}

    Extract decoupled pooled features $z_{\text{cur}}^{\text{dec}}, z_{\text{base}}^{\text{dec}}$

    Compute $\mathcal{L}_{\text{pres}}^{(j)}$ via Eq. 3

  end for

  Form stage-1 objective $\mathcal{L}_{\text{Stage1}}$ by Eq. 4

  Update only LoRA parameters with $\nabla_\theta \mathcal{L}_{\text{Stage1}}$

until convergence

$\theta_{\text{bad}} \leftarrow \theta$, $\Delta\theta \leftarrow \theta_{\text{bad}} - \theta_0$

Stage 2: Harmful Knowledge Negation (inference)

for prompt $(\text{text}, \text{class})$ **do**

  if Static then

    $\alpha \leftarrow \alpha_{\text{Static}}$

  else if Gated then

    if class = `unsafe` **then** $\alpha \leftarrow \alpha_{\text{unsafe}}$ **else** $\alpha \leftarrow \alpha_{\text{safe}}$

  end if

  Set $\theta^* \leftarrow \theta_0 - \alpha\,\Delta\theta$

  Generate motion via DDPM-style denoising with f_{θ^*}

end for

our model on this dataset. R-precision is the measurement of text-motion matching in the shared feature space, where a generated motion is successful when its text appears in the top-k closest candidates consisting of 1 ground truth and 31 random negative samples. FID computes the Fréchet distance between Gaussian fits of motion features from generated results and ground truths, measuring the distance of the generated motion distribution to the ground truth distribution. Diversity is the average pairwise distance between features of randomly sampled generated motions, capturing intra-set variability.

R-Precision. Following t2m (Zhang et al., 2023), a shared text–motion feature space is used for

retrieval. R-Precision reports top-k accuracy when each generated motion is queried against 1 ground-truth caption and 31 mismatched captions (R@1/2/3).

FID. FID computes the Fréchet distance between two Gaussians fitted to motion features from generated and real samples, capturing distributional discrepancy. This is measured by the L2-loss between their latent feature representations.

Diversity. Diversity estimates intra-set variability by splitting the generated set into two equal halves $\{m_1, \ldots, m_{M_d}\}$ and $\{m'_1, \ldots, m'_{M_d}\}$ and averaging cross-set feature distances:

$$\text{Diversity} = \frac{1}{M_d} \sum_{i=1}^{M_d} \|m_i - m'_i\|_2. \quad (5)$$

C Results on Motion-X Dataset

Results on Motion-X (Lin et al., 2023) dataset are shown in Table 4. SafeMo-Static and SafeMo-Gated yield forget-set FID increases of +372.8% and +1344.5%, with R@1 degradations of -48.4% and -73.5% vs. MoMask w/ LCR. On the retain set, SafeMo-Gated attains a lower FID (-56.0%), while SafeMo-Static remains comparable. In summary, SafeMo-Gated, with prompt-toxicity-awareness, has strong forgetting capability on unsafe prompts while maintaining high fidelity on the benign prompts, while SafeMo-Static acts as an external-agent-free, offline variant still pushing the unsafe generative distribution away at a comparable level with modest degradation on retain quality. On the forget split, our model exhibits neutralization: FID substantially increases and R@1 sharply drops, indicating effective removal of unsafe semantics. Meanwhile, crucially, on the retain set, FID and R@1 remain at a comparable level with base model's result, largely preserving normal utility.

D Qualitative Results

We present results on unsafe prompts in Figure 5 and Figure 6, and results on safe prompts in Figure 7.

From the results on unsafe prompts in Figure 5 and Figure 6, we observe that both SafeMo-Static and SafeMo-Gated can effectively erase the toxic motion semantics, which aligns with the design and aim of our unlearning strategy. SafeMo-Static erases toxic motion semantics effectively, while SafeMo-Gated tends to generate stationary or repeated pattern-like motion, which demonstrates a stronger tone of unlearning.

However, some limitations and flaws are observed in our qualitative results. Firstly, although it is of a high success rate that the model's generated results are not aligned with the unsafe text prompts, we observe some suboptimal results in certain scenarios, e.g., very long and detailed descriptions will cause some atomic semantics to be omitted, or being in a stationary-like pattern.

Secondly, we observe foot sliding and skating artifacts as a byproduct of the unlearning, with increased occurrence when applying a larger alpha to the text vector. We also observed that from Table 5, terms like $\mathcal{L}_{\text{foot}}$ are not in a linear-mapping fashion, i.e., with a larger alpha applied, the performance gaps are not changing in a linear manner. This indicates that we may need to explore a more complex relationships between the unlearning effect and each term of our designed method to further improve the model's performance in future work.

As the results shown in Figure 7, both SafeMo-Static and SafeMo-Gated can generate semantically aligned results on safe prompts. In some cases, SafeMo-Static favors lower-amplitude, more conservative kinematics. Additionally, foot-sliding and skating artifacts as a byproduct of the unlearning are also observed. SafeMo-Static is more susceptible to this byproduct than SafeMo-Gated because of the larger α weighted task vector negation applied on it on safe prompts than on SafeMo-Gated.

E Ablation Study

E.1 Loss Subterm Removal

We conducted ablation experiments by iteratively removing the loss terms, MPJPE, $\mathcal{L}_{\text{vel}}$, $\mathcal{L}_{\text{acc}}$, $\mathcal{L}_{\text{foot}}$, and $\mathcal{L}_{\text{text}\leftrightarrow\text{mo}}$, to demonstrate each term's significance in the model learning process. As shown in Table 5, the removal of MPJPE drastically destroys the model's unlearning performance on the unsafe set. The $\mathcal{L}_{\text{vel}}$ and $\mathcal{L}_{\text{acc}}$ both play an important role in the model's unlearning of unsafe patterns as well, with lower FID and higher R precision on unsafe prompts after removing them. $\mathcal{L}_{\text{foot}}$ plays a role in enhancing the model's understanding of the motion semantics and helps generate more physically aligned results, with slight degradation in unlearning on unsafe prompts after removing it. $\mathcal{L}_{\text{text}\leftrightarrow\text{mo}}$ has a significant influence on the model's understanding of unsafe motion since we observe

Table 4: **Results on Motion-X dataset.** Method D_r reports performances for the model trained on a toxicity-free dataset. Method FT shows the results of fine-tuning the model on the toxicity-free dataset. Diversity is reported for reference. *Note:* Rows marked with † are reported from (De Matteis et al., 2025) due to unavailable implementation and checkpoints at the time of submission.

	Forget Set			Retain Set		
	FID ↑	Diversity	R@1 ↓	FID ↓	Diversity	R@1 ↑
MoMask $D_r^\dagger$	$8.435_{\pm.295}$	$15.721_{\pm.255}$	$0.119_{\pm.007}$	$4.508_{\pm.103}$	$19.560_{\pm.332}$	$0.332_{\pm.002}$
MoMask† (Guo et al., 2024)	$2.028_{\pm.127}$	$15.884_{\pm.219}$	$0.289_{\pm.008}$	$2.686_{\pm.045}$	$19.366_{\pm.214}$	$0.344_{\pm.001}$
MoMask $FT^\dagger$	$2.072_{\pm.099}$	$15.855_{\pm.050}$	$0.280_{\pm.001}$	$3.325_{\pm.060}$	$19.405_{\pm.228}$	$0.347_{\pm.002}$
MoMask w/ UCE†	$10.522_{\pm.223}$	$6.648_{\pm.112}$	$0.033_{\pm.001}$	$3.740_{\pm.041}$	$6.243_{\pm.059}$	$0.046_{\pm.008}$
MoMask w/ RECE†	$12.704_{\pm.327}$	$6.241_{\pm.132}$	$0.031_{\pm.003}$	$14.287_{\pm.133}$	$6.342_{\pm.062}$	$0.029_{\pm.001}$
MoMask w/ LCR†	$2.218_{\pm.159}$	$15.606_{\pm.210}$	$0.283_{\pm.007}$	$2.656_{\pm.043}$	$19.260_{\pm.216}$	$0.335_{\pm.001}$
SafeMo-Static	$10.487_{\pm.102}$	$6.066_{\pm.166}$	$0.146_{\pm.001}$	$3.470_{\pm.003}$	$7.429_{\pm.043}$	$0.231_{\pm.002}$
SafeMo-Gated	$32.038_{\pm.026}$	$4.603_{\pm.046}$	$0.075_{\pm.003}$	$1.168_{\pm.007}$	$8.468_{\pm.159}$	$0.261_{\pm.001}$

lower FIDs on both versions of the model on unsafe prompts.

E.2 LoRA Rank Ablation

Unless otherwise specified, we use LoRA with rank $r = 16$ as a prior default, which offers a stable capacity-regularization trade-off in our decoder-only DiP backbone. To evaluate the effect of different LoRA (Hu et al., 2022) rank, we conduct an ablation study on different LoRA rank values. The results are shown in Table 6. We evaluate the same checkpoints after training on the MMU strategy for 20K steps with different LoRA ranks and the same LoRA alpha as the LoRA rank, while keeping all other hyperparameters the same. Across different sets, $r = 16$ yields the most balanced forgetting-retention effect: unsafe FID increases while unsafe R@k is reduced or comparable, and performance on the safe set is maintained to an acceptable level with the best balanced results on applying the same checkpoint for SafeMo-Static and SafeMo-Gated. We hereby unfold our findings to support our choice. While $r = 8$ sometimes produces slightly higher FID shifts on unsafe subsets, it frequently exhibits higher R@k, which indicates a more sense of geometric displacement and a lower level of commensurate semantic forgetting on unsafe prompts. We also observe that with the same replication times, the confidence intervals (CI) of $r = 8$ are consistently larger than on $r = 16$, which is a sign of instability and inadequate capability of obtaining the exact knowledge we want during the first stage. Conversely, model with $r = 32$ tends to under-forget on the unsafe sets with relatively large performance gaps on FID and R precisions on both the Gated and the Static settings. Apart from that, in the Gated setting, it greatly harms the safe set's performance even with a small value of α in the Gated setting, making it undesirable.

E.3 Alpha Scaling Ablation

We study how the task-vector scale $\boldsymbol{\alpha}$ controls the trade-off between retaining benign capability and forgetting harmful behaviors. The results are shown in Figure 8. The curves reveal a clear effective-unlearning window on $[0.05, 1.2]$, where the unsafe split deteriorates markedly with mild changes happening on safe prompts. Beyond $\alpha = 1.2$, the safe curves also degrade steeply indicating over-forgetting, which is undesirable for benign prompts. We reckon that $\alpha = 1.0$ is the sweet spot for SafeMo-Static, with acceptable degradation on benign tasks (FID = 2.51, R@1 = 0.33) and good unlearning performance on unsafe prompts (FID = 9.55, R@1 = 0.18). These observations demonstrate the large selective deterioration on unsafe prompts before the knee, validating that our unlearning is effective. These also motivate the SafeMo-Gated setting for a more flexible and accurate control utilizing the proposed LLM-base agent in *SafeMoEngine*.

F Implementation Details

Forget and retain set partitioning. To compare our method with LCR (De Matteis et al., 2025) despite the fact that they have not made their implementation open-source, we adopt the same keyword-based partitioning paradigm they describe. We construct a curated list of harmful

Table 5: **Ablation study of loss subterms in MMU stage-1.** Results on HumanML3D. On unsafe sets, higher FID and lower retrieval (R@K) indicate stronger forgetting; on the safe set, lower FID and higher retrieval indicate better utility. Diversity is reported for reference.

	Unlearned Unsafe Set					**Unseen** Unsafe Set					**Unseen** Safe Set				
	FID↑	Div.	R@1↓	R@2↓	R@3↓	FID↑	Div.	R@1↓	R@2↓	R@3↓	FID↓	Div.	R@1↑	R@2↑	R@3↑
SafeMo ($\alpha = 0.0$)	1.7197	7.3746	0.2517	0.3914	0.4969	2.3050	7.5191	0.2365	0.3896	0.5052	0.5232	9.3375	0.3755	0.5599	0.6732
SafeMo-Static	8.0235	6.8083	0.2016	0.3164	0.4043	9.0499	6.8880	0.1958	0.3167	0.3937	2.5539	8.7060	0.3172	0.4935	0.6052
SafeMo-Static w/o MPJPE	4.6513	7.3699	0.2243	0.3654	0.4587	4.7538	7.3980	0.2271	0.3729	0.4677	1.7615	9.0878	0.3548	0.5368	0.6493
SafeMo-Static w/o $\mathcal{L}_{\text{vel}}$	6.9403	6.9088	0.2108	0.3318	0.4221	7.8048	7.0443	0.1885	0.3229	0.4167	2.1361	8.8828	0.3346	0.5123	0.6286
SafeMo-Static w/o $\mathcal{L}_{\text{acc}}$	7.7847	6.8950	0.2074	0.3266	0.4147	8.8772	6.9484	0.3073	0.4479	0.5434	2.3508	8.8585	0.3266	0.5078	0.6199
SafeMo-Static w/o $\mathcal{L}_{\text{foot}}$	7.5536	6.8465	0.2031	0.3260	0.4109	8.5371	6.9620	0.1854	0.3198	0.4094	2.3039	8.7675	0.3262	0.5037	0.6139
SafeMo-Static w/o $\mathcal{L}_{\text{text}\leftrightarrow\text{mo}}$	7.0224	6.7054	0.2039	0.3295	0.4172	7.9530	6.7763	0.1979	0.3125	0.4094	2.2494	8.6139	0.3253	0.5086	0.6201
SafeMo-Gated	28.0806	5.0169	0.0947	0.1630	0.2168	28.0574	4.8520	0.0865	0.1542	0.2104	0.5355	9.3224	0.3775	0.5628	0.6769
SafeMo-Gated w/o MPJPE	11.7242	6.8297	0.1962	0.3154	0.4026	11.9482	6.9063	0.2115	0.3167	0.4052	0.5600	9.3109	0.3743	0.5611	0.6761
SafeMo-Gated w/o $\mathcal{L}_{\text{vel}}$	25.0236	5.1915	0.1238	0.2037	0.2600	25.4520	5.1998	0.1292	0.2167	0.2875	0.5330	9.4388	0.3713	0.5641	0.6751
SafeMo-Gated w/o $\mathcal{L}_{\text{acc}}$	25.2587	5.1204	0.1252	0.2076	0.2681	25.3217	5.1177	0.1229	0.2125	0.2771	0.5327	9.4913	0.3814	0.5701	0.6802
SafeMo-Gated w/o $\mathcal{L}_{\text{foot}}$	27.1904	4.9923	0.1152	0.1858	0.2411	27.6051	4.9489	0.1208	0.2021	0.2615	0.5327	9.4370	0.3719	0.5649	0.6768
SafeMo-Gated w/o $\mathcal{L}_{\text{text}\leftrightarrow\text{mo}}$	22.3646	4.7543	0.1076	0.1825	0.2429	22.7699	4.7297	0.1031	0.1948	0.2521	0.5263	9.4303	0.3726	0.5659	0.6764

Table 6: **Ablation study of LoRA rank in MMU stage-1.** Results on HumanML3D. On unsafe sets, higher FID and lower retrieval (R@K) indicate stronger forgetting; on the safe set, lower FID and higher retrieval indicate better utility. Diversity is reported for reference.

	Unlearned Unsafe Set					**Unseen** Unsafe Set					**Unseen** Safe Set				
	FID↑	Div.	R@1↓	R@2↓	R@3↓	FID↑	Div.	R@1↓	R@2↓	R@3↓	FID↓	Div.	R@1↑	R@2↑	R@3↑
SafeMo ($\alpha = 0.0$)	1.7197	7.3746	0.2517	0.3914	0.4969	2.3050	7.5191	0.2365	0.3896	0.5052	0.5232	9.3375	0.3755	0.5599	0.6732
SafeMo-Static	8.0235	6.8083	0.2016	0.3164	0.4043	9.0499	6.8880	0.1958	0.3167	0.3937	2.5539	8.7060	0.3172	0.4935	0.6052
SafeMo-Static (LoRA r=8)	8.1880	6.8448	0.2022	0.3225	0.4117	8.9847	6.9679	0.1958	0.3250	0.4042	2.4408	8.6262	0.3252	0.5031	0.6127
SafeMo-Static (LoRA r=32)	5.6007	7.7033	0.2068	0.3372	0.4325	4.9216	7.6619	0.1854	0.3333	0.4396	2.4861	8.7952	0.2947	0.4700	0.5841
SafeMo-Gated	28.0806	5.0169	0.0947	0.1630	0.2168	28.0574	4.8520	0.0865	0.1542	0.2104	0.5355	9.3224	0.3775	0.5628	0.6769
SafeMo-Gated (LoRA r=8)	31.5091	4.1975	0.1080	0.1800	0.2313	32.7289	4.1776	0.1052	0.1750	0.2417	0.5328	9.2650	0.3781	0.5666	0.6785
SafeMo-Gated (LoRA r=32)	19.6348	6.7620	0.1292	0.2184	0.2924	17.7689	6.7304	0.1365	0.2344	0.3052	2.5434	9.0249	0.3222	0.4983	0.6139

action lemmas as described in their method, lemmatize captions, and perform exact lemma-level matching with phrase-first priority. A sample is assigned to the *forget set* if *any* of its captions hits the list; otherwise, it belongs to the *retain set*.

Motion-X. Motion-X (Lin et al., 2023) provides SMPL-X (Pavlakos et al., 2019) data including hand and facial feature, which is not aligned with our setup. We process it using the official code from Motion-X (Lin et al., 2023), converting SMPL-X features to SMPL (Loper et al., 2015) representations. To ensure the comparability with LCR, we preprocess the text prompts using HumanML3D (Zhang et al., 2023)'s Semantic Role Labeling method, and train feature extractors following official t2m (Zhang et al., 2023) implementation for 300 epochs, as the same as in LCR.

G Limitation and Future Work

This appendix expands the brief limitations stated in the main paper, with concrete failure cases and future directions for both SafeMoEngine and MMU.

Safe T2M dataset. Despite that we design the LLM-based agent in a classify-then-rewrite fashion with explicit few-shot, rule-enhanced prompt engineering, and effective, level-aligned rewriting strategy, we acknowledge several limitations. Firstly, we have observed that the classifier agent tends to be slightly oversensitive to toxic semantics. In level-2 prompts, we found a few sport-related or dancing-related prompts, e.g., golf or dancing with crazy legs. Secondly, although we apply two recent advances, one continuous and the other discrete token-based, for generating new refined safe motion for unsafe prompts, some generated results can be suboptimal due to the base model's respective limitations.

Minimal motion unlearning. Although our results demonstrate the effectiveness of unsafe motion unlearning, we have found several kinds of suboptimal cases or failed cases. First, when the context is too long and contains many details, the generated motion can omit some atomic semantics, or even, at worst, become a stationary-like pattern. We reckon that this is the base model's limited understanding capability of long and complex prompts. In future work, we aim to address this problem by using a more semantic-accurate base model or designing a text prompt reasoning helper, e.g., methods similar to Motion-agent (Wu et al., 2024). Secondly, task-vector negation can exacerbate foot-skating artifacts, which are particularly noticeable for SafeMo-Gated under larger gating scales. Future work may mitigate this byprod-

uct by designing more physics-guided constraints and integrating physics-based trackers, such as in CLoSD (Tevet et al., 2025), to further improve physically plausible contacts and environment interactions. Thirdly, Operating in continuous motion space alleviates discrete codebook stitching artifacts (Figure 1), but it may trade off some standard T2M fidelity metrics compared to strong VQ-token pipelines.

H User Study

We conduct a comprehensive user study to evaluate the overall quality and unlearning performance of motion sequences generated by our methods. A total of 50 participants completed a Google Forms survey designed to assess the physical plausibility, unlearning outcomes on unsafe prompts, and benign performance on safe prompts.

As illustrated in Figure 9, section 1 and 2 displays different generated results on unsafe prompts by SafeMo-Static and SafeMo-Gated respectively, followed by section 3 and 4 showing SafeMo-Gated and SafeMo-Static's generated results on safe prompts, with 3 different prompts and 2 different questions each section. Section 5 and 6 then compare SafeMo-Static and SafeMo-Gated's results on the same safe text prompt and the same unsafe prompt respectively. Participants are asked to rate each motion at a 5-point Likert scale (where 1 represents low and 5 represents high) based on motion naturalness, degree of unlearning performance on unsafe prompts and quality of text alignment on safe prompts.

This study aims to evaluate not only the unlearning performance of our model, but also the benign performance and quality distinctions between SafeMo-Static and SafeMo-Gated.

The results of the user study can be summarized as follows: *(i)* Our method achieved an overall motion quality of 4.24 on SafeMo-Static, and 4.82 on SafeMo-Gated. *(ii)* On unsafe prompts, 98% of participants agreed that SafeMo-Gated effectively removed unsafe components in the motion, and 86% on SafeMo-Static. *(iii)* On safe prompts, our method's generated results on safe set scored 4.64 on text-motion visual alignment rating. *(iv)* 56% of the participants preferred SafeMo-Gated method on unsafe prompts, and 84% preferred SafeMo-Gated method on safe prompts.

I LLM Use Declaration

Large Language Models (ChatGPT) were used exclusively to improve the clarity and fluency of English writing. They were not involved in research ideation, experimental design, data analysis, or interpretation. The authors take full responsibility for all content.

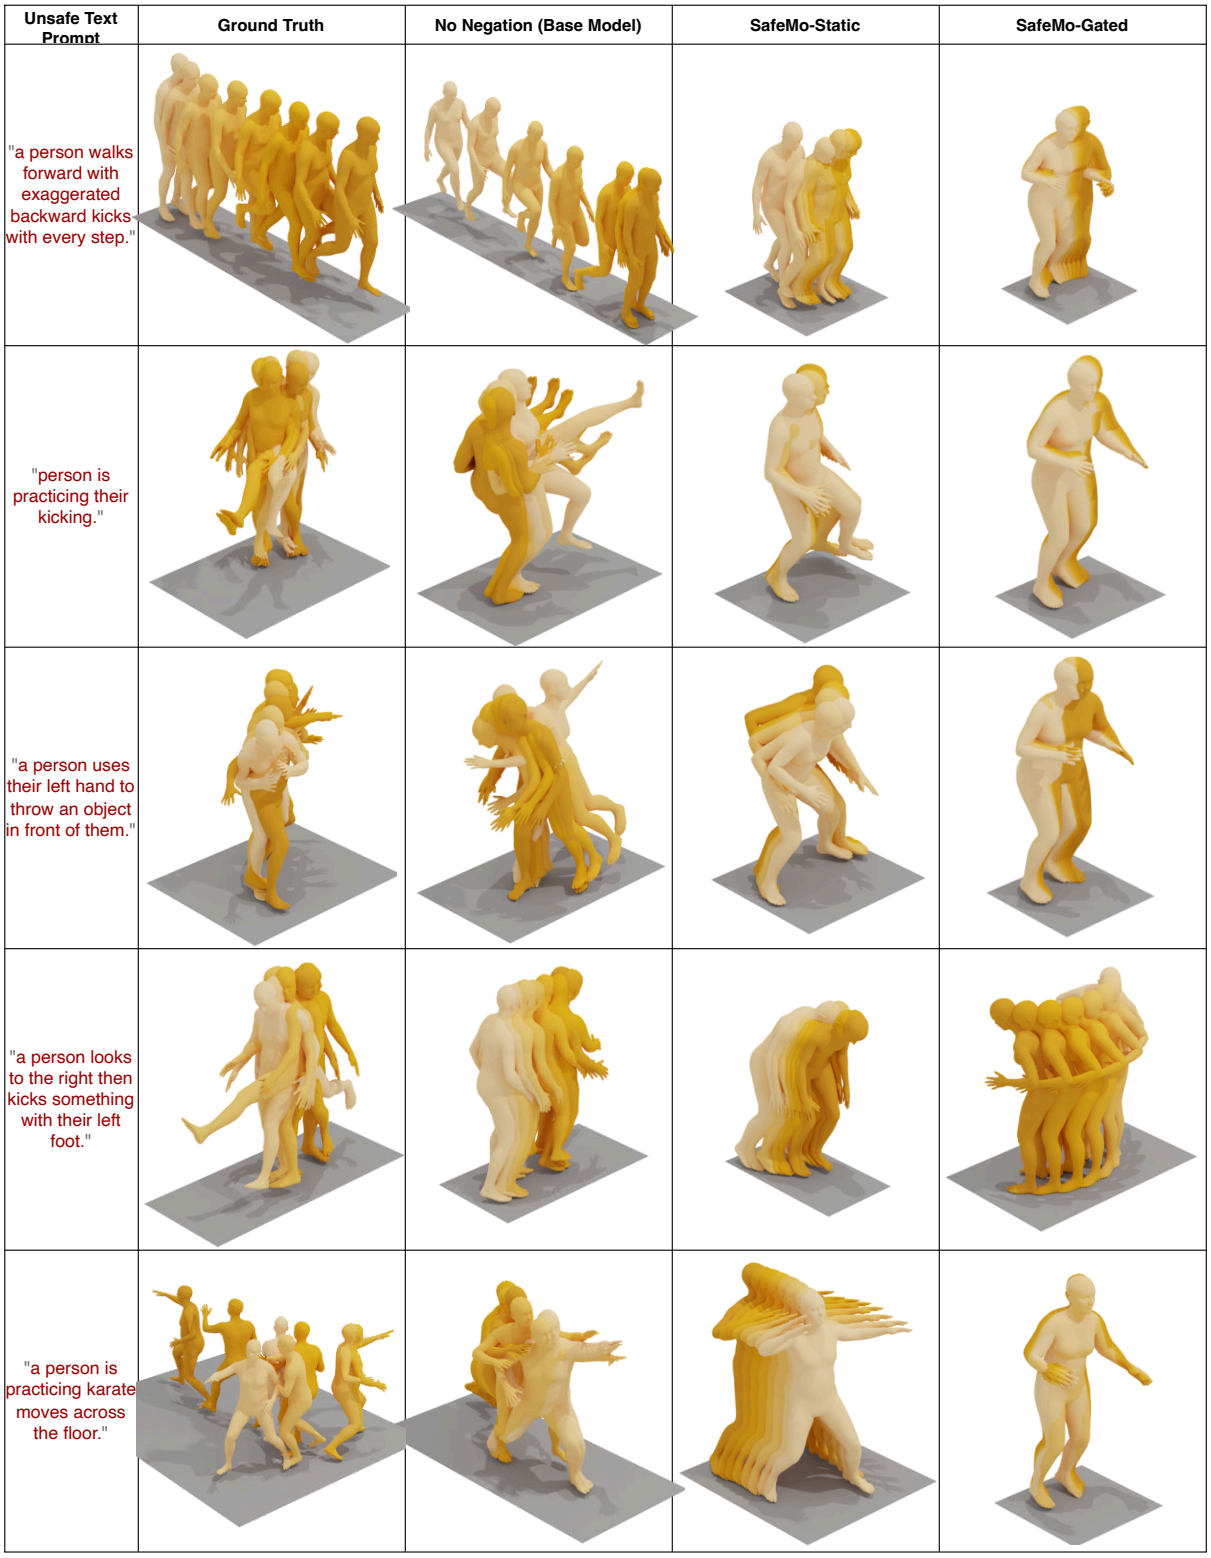

Figure 5: Qualitative results of generated motions on unsafe prompts (Part I).

Unsafe Text Prompt	Ground Truth	No Negation (Base Model)	SafeMo-Static	SafeMo-Gated
"a person turns to the right and brings both hands together while kicking slightly to the right with the left foot."	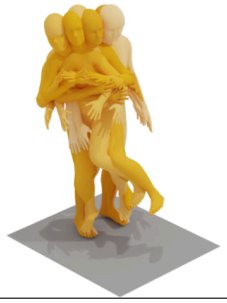			
"a person kicks with their right leg twice, and then once with their left."				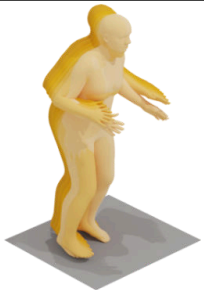
"a person grabbing something in front of them, swinging it around to the side then throwing it overhead."				
"a person looks to the left then kicks something with their right foot."				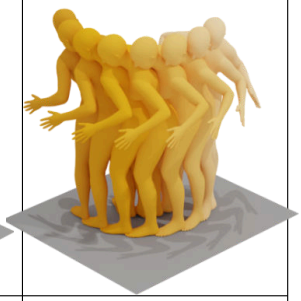
"a person preparing for and then throwing something similar to how a quarterback throws a football."				

Figure 6: Qualitative results of generated motions on unsafe prompts (Part II).

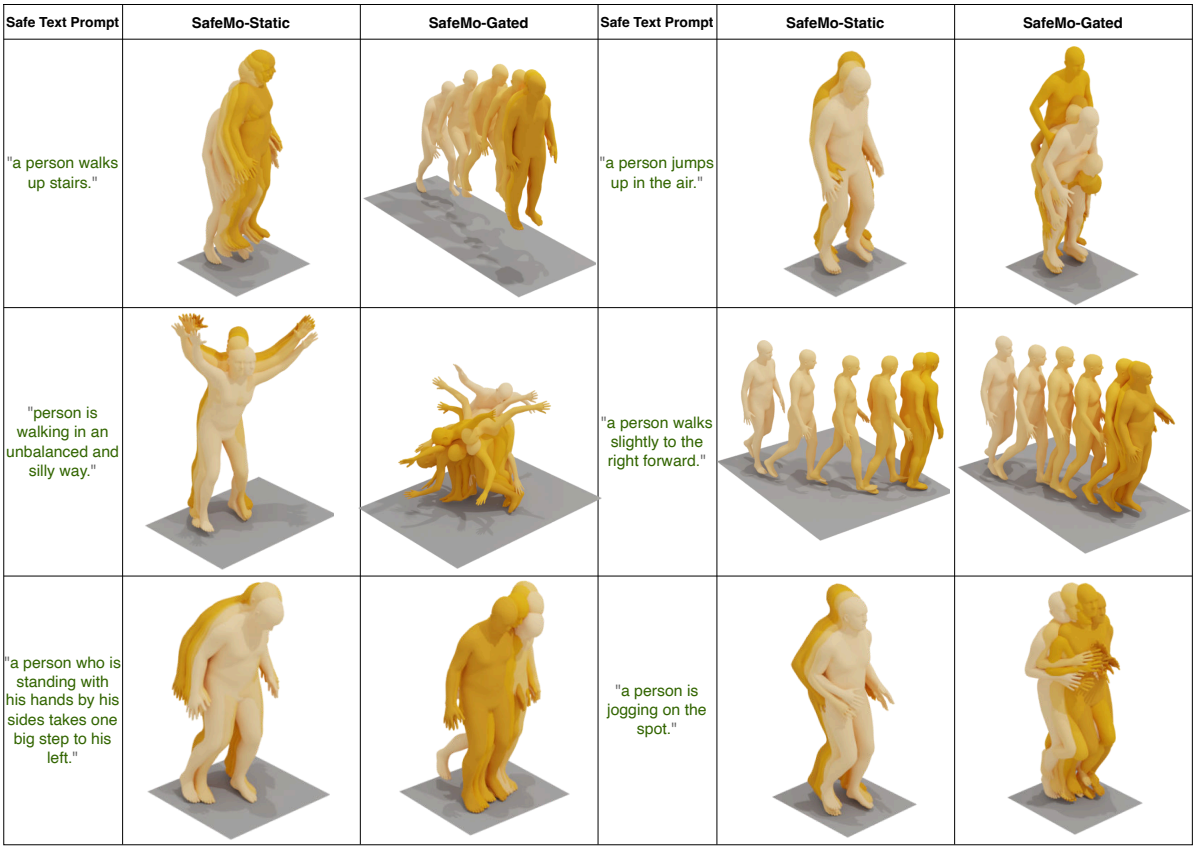

Figure 7: Qualitative comparison of generated motions on safe prompts.

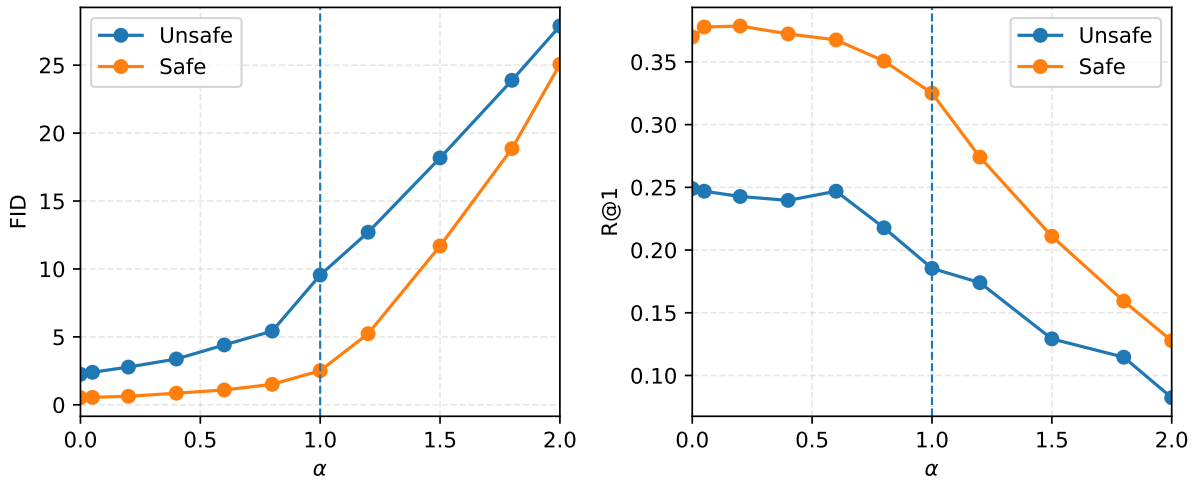

Figure 8: **Effect of the task-vector scaling α on safe and unsafe prompts.** Left: FID (lower is better on safe prompts; higher indicates stronger forgetting on unsafe prompts). Right: R@1 (higher is better on safe prompts; lower indicates stronger forgetting on unsafe prompts). The dashed line marks $\alpha = 1.0$.

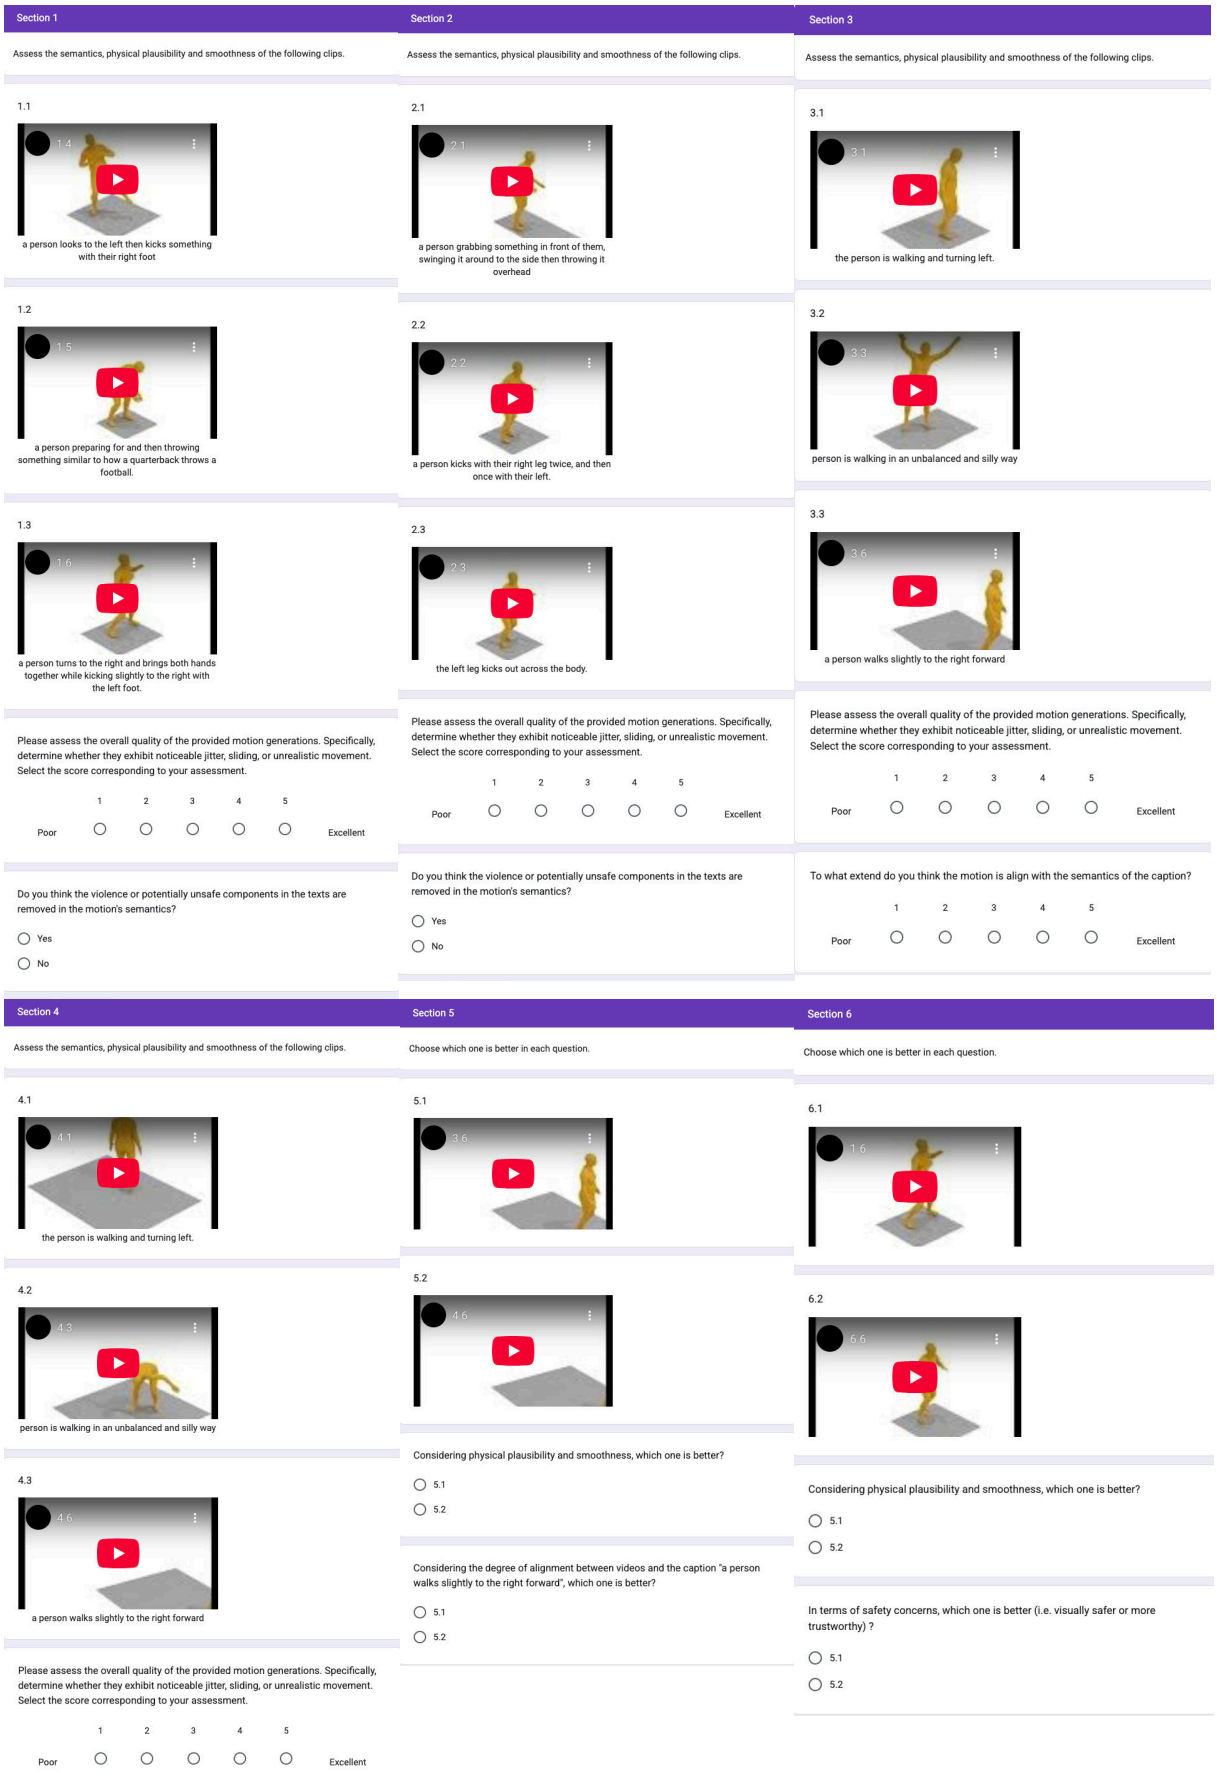

Figure 9: **User study Google Forms.** The User Interface (UI) used in our user study.