

Quality-Preserving Imperceptible Adversarial Attack on Skeleton-based Human Action Recognition

Ziyi Chang^{1b}, Kanglei Zhou^{2b}, Xiaohui Liang^{3b}, Hubert P. H. Shum^{1b}, *Senior Member, IEEE*

Abstract—Adversarial attacks on skeletal human action recognition have received significant attention. However, existing methods typically introduce noise-like perturbations that degrade motion quality post-attack, and thereby are inherently perceptible with recent advancements in S-HAR systems. We discover that this degradation stems from the gap between empirical and true risks during the optimization process of previous adversarial attacks. To address this issue, we propose an attack where adversarial motions are obtained without compromising their motion quality. To minimize the risk gap and preserve motion quality, we propose a distribution-based adversarial attack method without introducing noise-like perturbations. To faithfully evaluate the motion quality, we propose a new metric that aligns with human perception on real-world naturalness. Experiments have been conducted on the state-of-the-art S-HAR methods across two datasets, demonstrating the superiority of our method in both the attack success rate and the post-attack motion quality through qualitative and quantitative analyses. The success of our quality-preserving attack application and distribution-based method raises serious concerns about the robustness of action recognizers, highlighting the need for further enhancements in this domain.

Index Terms—Adversarial Attack, Human Action Recognition, Human Skeletal Motion, Diffusion Model, Motion Quality.

I. INTRODUCTION

ADVERSARIAL attacks on skeletal human action recognition (S-HAR) have received significant attention due to concerns regarding the robustness of S-HAR systems. These systems have been deployed to recognize human actions from skeletal inputs and have been deployed in various security-critical and human-centric domains such as medical care, action assessment, and safety surveillance [1]. The adversarial motions pose substantial threats to such life-critical applications by misleading the S-HAR systems [2]. Therefore, obtaining adversarial motions is important to enhance the robustness of these S-HAR systems and boost trustworthy and responsible artificial intelligence [3], [4].

Manuscript received June 10, 2026. This research is supported in part by the EPSRC NorthFutures project (ref: EP/X031012/1). (*Corresponding author: Hubert P. H. Shum.*)

Ziyi Chang and Hubert P. H. Shum are with the Department of Computer Science, Durham University, Durham, DH1 4FL, UK (e-mail: {ziyi.chang, hubert.shum}@durham.ac.uk).

Kanglei Zhou is with the Department of Psychology and Cognitive Science, Tsinghua University, Beijing 100084, China (e-mail: zhokan-glei@tsinghua.edu.cn).

Xiaohui Liang is with the State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, Beijing 100191, China, and also with the Zhongguancun Laboratory, Beijing 100190, China (e-mail: liang_xiaohui@buaa.edu.cn).

Previous adversarial attack methods [5], [6] typically introduce noise-like perturbations to skeletal inputs to deceive S-HAR algorithms, resulting in reduced post-attack motion quality. These perturbations inherently conflict with human perceptions of natural motion [7], [8] due to the quality decline, and thereby diminish the imperceptibility of existing attacks. Nonetheless, this vulnerability is often overlooked due to inherent dataset noise [9], [10], [11] and limited classifier expressiveness [12], [13], [14]. Furthermore, existing metrics that compare paired pre-attack and post-attack motions inadequately assess post-attack motion quality because motions are sparsely distributed and their neighborhoods may not be natural and plausible [15]. With advancements enabling high-quality skeletal inputs for S-HAR systems [16], [17], [18], the attacked motions become increasingly perceptible and undermine the imperceptibility of adversarial attack.

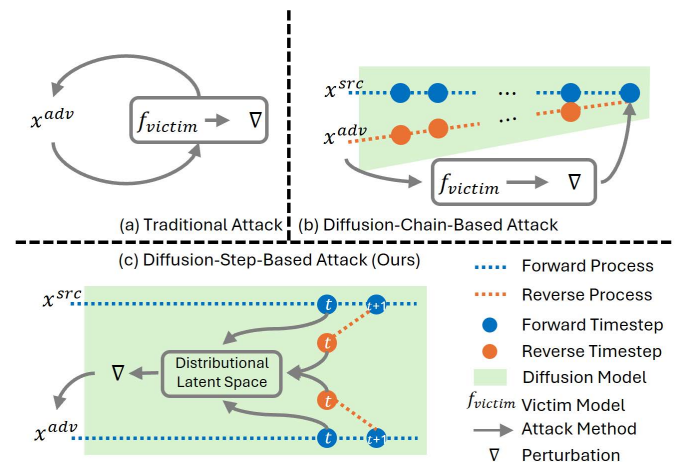


Fig. 1: The schema of our proposed method when compared with existing attack methods. Instead of relying on the diffusion chain and the gradient from external victims, our method constructs a distributional latent space to leverage the inherent distribution knowledge learned by diffusion models for attack.

We identify the decline in the quality of adversarial motions resulting from the optimization process in previous attack methods as shown in Fig. 1. Firstly, existing methods typically rely on paired loss to calculate perturbations where the empirical risk for optimization significantly deviates from the true risk. The empirical risk typically refers to the loss defined over a large number of observations, while the true risk is defined as the expected loss across the implicit data distribution [19]. However, in the adversarial attack context, only an individual motion is involved for optimizing a motion to be adversarial,

which leads to a large gap between the empirical and true risks [20]. Secondly, the optimization process relies heavily on the gradient signals from a classifier, whose decision boundaries often extend beyond the true data manifold. As a result, the adversarial modification is driven in directions that do not necessarily align with the underlying distribution of natural motions. This classifier-dependent gradient guidance frequently pushes the motion off the manifold, leading to unnatural, noise-like patterns [21] that ultimately degrade the quality and plausibility of the adversarial motions.

To tackle the decline in the quality of adversarial motions, we propose a novel attack application where adversarial motions are obtained without sacrificing quality. This application demonstrates a new task towards imperceptible adversarial attack where the post-attack motion quality is explicitly constrained rather than the paired-motion constraints in previous methods. To minimize the risk gap in the optimization process of adversarial attack, we propose a distribution-based S-HAR attack method where a diffusion model is employed to learn the data distribution, as shown in Fig.2. The proposed distribution-based method relies on the data distribution instead of paired pre-attack and post-attack motions to optimize the given motions. By reducing the risk gap, our optimization does not introduce noise-like perturbations and thereby our adversarial motions are obtained without sacrificing the motion quality. Furthermore, to faithfully quantify the quality of adversarial motions, we propose a new metric to measure the naturalness based on physiological analysis of human movements [22], [23]. This new metric measures the inherent physiological naturalness rather than relying on the paired individuals comparison in existing adversarial metrics, aligning with the human perceptions of natural motions in the real world.

Extensive experiments show that our method generates adversarial motions with the least quality decrease for attacking against the four latest S-HAR classifiers across both a high-quality dataset, i.e., 100STYLE [24] and two commonly used datasets, i.e., HDM05 [25] and NTU60 [9] when compared with state-of-the-art S-HAR adversarial attack methods. Our user study further validates that our adversarial motions are the least perceivable by humans. We also conduct ablation studies on hyper-parameters of diffusion models and the specification of generative models. Codes are available in <https://github.com/mrzzzy2021/QualityPreservingAttack>. Our contributions are summarized as follows:

- We discover a critical yet previously overlooked vulnerability in existing attack methods where their noise-like adversarial perturbations are inherently perceptible to humans. To address this, we propose a novel attack framework that imperceptible adversarial attack is fulfilled without sacrificing post-attack motion quality.
- To preserve post-attack motion quality, we propose a distribution-based attack method where imperceptible adversarial motions are generated through a pre-trained diffusion model by minimizing the gap between empirical and true risks in our optimization of adversarial attack.
- To faithfully quantify the quality of adversarial motions, we propose a new metric based on existing physiologi-

cal analysis of real-world human motions, enabling the assessment of previously overlooked vulnerabilities that existing metrics fail to capture.

The remainder of this paper is organized as follows: We first review related work in Section II. Then, we formulate the proposed quality-preserving adversarial attack application in Section III. Subsequently, we present our proposed distribution-based adversarial attack method with quality preservation in Section IV. We formulate the proposed metric in Section V to faithfully evaluate the motion quality after attack. Finally, we demonstrate the superiority of our method in Section VI and provide a summary in Section VII with potential future work.

II. RELATED WORK

A. General Adversarial Attack

Adversarial attacks are first introduced in [26], highlighting the vulnerability of deep neural networks and subsequently extending to various data types. Generally, adversarial attacks serve as a specialized form of data augmentation aimed at exposing system vulnerabilities by generating new samples, whereas other data augmentation techniques may focus on objectives such as enhancing training efficiency and improving inference performance [27]. Deep neural networks remain susceptible to meticulously crafted adversarial attacks despite that significant achievements have been witnessed. Noise-like perturbations have been applied to input data to easily deceive these high-performing neural networks and people raise concerns on the trustworthiness and reliability of deployed neural networks [3], [4]. In response to these concerns, researchers have extensively explored adversarial attacks across different data modalities, including 2D images [28], [29], videos [30], [31], 3D objects [32], physical objects [33], and graph data [34]. While adversarial research on different modalities has been widely investigated, attacks on time-series data have recently gained attention [35], with relatively less focus on 3D skeletal motions characterized by spatio-temporal structures.

B. Adversarial Attacks on S-HAR

Adversarial attacks on skeletal human action recognition (S-HAR) aim to perturb 3D human skeletons to deceive classifiers, mirroring objectives from the image domain. Skeletal motion is extensively utilized in HAR to mitigate challenges such as lighting variations, occlusions, and diverse view angles [1]. Consequently, the vulnerability of skeleton-based classifiers to adversarial attacks has received increasing attention. With a significantly lower degree of freedom in motions [36], [37], previous S-HAR attack methods focus on improving imperceptibility through different perspectives. [6] targets GCN-based models by employing generative adversarial networks as a discriminator to achieve imperceptibility in terms of the anthropomorphic plausibility. [38], [39] propose a new attack method where only the length of bones could be perturbed to enhance the imperceptibility from the perspective of bone length consistency. [5] analyzes the perceptibility of adversarial skeletal samples and proposes constraints between paired pre-attack and post-attack motions, integrating




Fig. 2: Overview of the proposed distribution-based imperceptible adversarial S-HAR attack where post-attack motion quality is preserved. We optimize the motions along the gradient provided by a pretrained diffusion model based on the distributional latent space to keep the post-attack motion quality.

the observations from [6] and [38] to further improve imperceptibility. [40] proposes to find perturbation directions on the manifold of motions to pursue imperceptibility. However, while these methods effectively achieve to fool the S-HAR system, they inherently introduce noise-like perturbations that degrade post-attack motion quality, and thereby undermine the imperceptibility of post-attack motions.

C. Diffusion Models

Diffusion models [41] have attracted extensive attention for their powerful generative capabilities. These models utilize numerous timesteps to transform data distributions into a standard Gaussian distribution through a forward process, while a denoising network is trained to recover the less noisy data distribution [42]. A large number of timesteps are typically required for smooth distribution transitions and simultaneously provide a hierarchical latent space with rich semantics.

Some have leveraged the generative ability of diffusion models for adversarial attack. [43], [44] directly learns adversarial generation using paired adversarial samples, though this dependency may not be practical for real-world applications. [45] introduces a label-based generative method, training diffusion models on auxiliary datasets. [46] leverages diffusion models with modified denoising process to generate adversarial samples. [47] utilizes diffusion models for high-quality image steganography. [48] focuses on purification as a defense approach by the generative ability of diffusion models rather than the adversarial attack. Some rely on DDIM inversion to achieve adversarial attack [29], [49] or defence [50] through the full denoising chain of diffusion models. Despite the applications mentioned above, the previous diffusion-based methods still rely heavily discriminative classifier to calculate the perturbation and do not leverage the inherent distributional knowledge contained by the diffusion model itself.

III. IMPERCEPTIBILITY OF S-HAR ADVERSARIAL ATTACK VIA MOTION QUALITY PRESERVATION

In this section, we formulate our proposed attack application, which conducts imperceptible adversarial attack against skeleton-based human action recognition (S-HAR) systems without compromising post-attack motion quality.

a) Inherent Vulnerability of Imperceptibility: Existing adversarial attacks achieve imperceptibility by introducing small, noise-like perturbations that inherently and detrimentally affect post-attack motion quality. Motion quality encompasses naturalness and plausibility [51], as human motions adhere to physical and biomechanical constraints. The human brain has specialized neural mechanisms for perceiving biological motion [52], [53] and is highly sensitive to unnatural and implausible kinematics [7], [8]. Given the highly nonlinear and articulated spatial-temporal structures of human motions [5], even subtle noise-like perturbations significantly degrade quality. Such perturbations disrupt motion dynamics and physical constraints, and deviate motions from the natural and plausible distribution. Thus, perturbed motions suffer from the decline in motion quality and become perceptible.

b) Task Formulation: Nonetheless, the inherent perceptible vulnerability resulting from the introduced noise-like perturbations is often overlooked by previous S-HAR attack methods due to the intrinsic noise in datasets [9], [10], [11] and the limited expressiveness capacity of classifiers [12], [13], [14]. As advancements [16], [17], [18] have widely facilitated cutting-edge S-HAR systems to utilize skeletons with high quality as inputs for decision-making, human supervisors readily perceive suspicious noise-perturbed skeletal inputs.

Furthermore, previous constraints during their optimization process, which usually focus on paired pre-attack and post-attack motions, fail to faithfully reflect the quality of adversarial motions. Since motions are sparsely distributed [15], attacked motions that remain within the neighborhood of their pre-attack motions usually may not be natural and plausible, especially when considering the motion dynamics. Therefore, adversarial motions suffering from post-attack quality become increasingly perceptible due to the added noise-like perturbations, which urgently need to be resolved to improve the imperceptibility of adversarial attacks.

To achieve imperceptibility by preserving the quality of adversarial samples instead of introducing noise-like perturbations, we propose an adversarial task, in which post-attack motions preserve pre-attack motion quality. Specifically, our task explicitly requires the quality of the post-attack motions as the optimization constraint instead of individual motion pairs,

and aims to achieve imperceptibility by preserving post-attack motion quality. We define our adversarial task as follows:

$$\begin{aligned} & \underset{\mathbf{x}^{\text{adv}}}{\text{minimize}} \quad \mathcal{L}_{\text{adv}}(C(\mathbf{x}^{\text{adv}})) \\ & \text{subject to} \quad Q(\mathbf{x}^{\text{adv}}) = Q(\mathbf{X}^{\text{src}}), \end{aligned} \quad (1)$$

where $C(\cdot)$ is an S-HAR classifier, $\mathbf{x}^{\text{adv}} \in \mathbb{R}^{T \times J \times 3}$ represents the attacked motion, $\mathbf{X}^{\text{src}} \in \mathbb{R}^{N \times T \times J \times 3}$ represents the set of natural and plausible motions, and $Q(\cdot)$ represents the motion quality. Compared with previous adversarial attacks, we propose optimizing an individual motion sample, \mathbf{x}^{adv} , with respect to the entire set of natural and plausible motions rather than rely on a pair of pre-attack and post-attack motions.

IV. DISTRIBUTION-BASED S-HAR ATTACK METHOD

In this section, we demonstrate the proposed distribution-based adversarial attack method against skeleton-based human action recognition (S-HAR) systems. First, we identify the origin of noise-like perturbations introduced in existing attack methods and propose integrating the data distribution into adversarial attacks by constructing a generative diffusion distributional latent, as described in Section IV-A. Then, we present our attack strategy to leverage the constructed latent for a quality-preserving adversarial attack in Section IV-B.

A. The Diffusion Latent for Quality Preservation

a) The Risk Gap of Previous Optimization: The noise-like perturbations in previous attacks are attributed to the gap between empirical and true risks in their optimization processes. Empirical risk typically refers to the loss defined over a large number of observations, while true risk is defined as the expected loss across the implicit data distribution [19]. However, in previous adversarial attacks, the optimization process involves only an individual motion sample, leading to a large gap between empirical and true risks [20]. When optimizing a motion with respect to a single pre-attack motion observation, noise-like perturbations are usually introduced [21] and decrease the quality of adversarial motions.

b) Constructing Generative Diffusion Latent: To minimize the risk gap in the optimization process of adversarial attacks, we propose leveraging distributional latents from generative diffusion models for our quality-maintaining adversarial attack. As shown in Fig. 3, diffusion models have semantically meaningful latent spaces that capture different feature patterns with respect to various choices of timesteps [54]. When compared to discriminative models [55], they are capable of modelling motion distributions that are critical for quality preservation. Compared with other generative models like VAEs [56], their hierarchical structures are more comprehensive to represent underlying patterns to modify motions.

Our attack leverages diffusion distributional latents as proxies for the given pre-attack motion individual. Previous S-HAR adversarial attack methods typically optimize motions directly in the original data space. However, the original space is sparse where natural motions being sparsely distributed [15]. This sparsity leads to difficulties in modifying motions within the original data space while maintaining their quality. In contrast,

we build our attack method on the stochastic latent space of diffusion models. The stochastic distributional latents are located in an approximately smooth space [57] and encode different levels of semantics for adversarial modification. By transforming the given pre-attack motion into latent distributions, the empirical risk of our optimization involves infinite samples within distributions rather than individual motions.

Specifically, we define our latent for adversarial attack to include two parts: a posterior mean for its rich semantics encoded in the stochastic process of diffusion models and a stochastic latent for its regularization. The posterior mean combines both the forward and reverse processes in a diffusion model, thereby comprehensively and hierarchically representing the motion. First, we leverage the forward process to convert motions into a distribution, which is defined as:

$$p(\mathbf{x}_{t+1}|\mathbf{x}_0) = \mathcal{N}(\sqrt{\alpha_{t+1}}\mathbf{x}_0, (1 - \alpha_{t+1})\mathbf{I}), \quad (2)$$

where \mathbf{x}_0 is a pre-attack motion, and \mathbf{x}_{t+1} is a stochastic sample at the timestep $t+1$ from a Gaussian distribution. Then, to achieve finer-grained latent manipulation for adversarial attack, we propose to construct our latent proxy for adversarial attack in a single timestep of diffusion reverse process. The posterior distribution in the reverse process is defined as:

$$p(\mathbf{x}_t|\mathbf{x}_{t+1}, \mathbf{x}_0) = \mathcal{N}(\boldsymbol{\mu}, \sigma_t\mathbf{I}), \quad (3)$$

where $\mathbf{x}_{t+1} \sim p(\mathbf{x}_{t+1}|\mathbf{x}_0)$, σ_t is the posterior variance and $\boldsymbol{\mu}$ is the posterior mean of the distribution. Since the posterior variance is pre-defined in noise schedule, $\boldsymbol{\mu}$ represents the denoising direction according to the class label. To show the relationship between the posterior mean and data distribution, we explicitly derive the posterior mean based on Tweedie's formula [58] and formulate it as:

$$\boldsymbol{\mu}_t = \gamma_t\mathbf{x}_0 + \lambda_t\epsilon_t + \delta_t\nabla \log p_{\theta}(\mathbf{x}_{t+1}, \mathbf{y}), \quad (4)$$

where \mathbf{x}_0 is the motion in original data space, \mathbf{x}_{t+1} is sampled from the latent distribution, i.e., $\mathbf{x}_{t+1} \sim p(\mathbf{x}_{t+1}|\mathbf{x}_0)$ and $\epsilon_t \sim \mathcal{N}(0, \mathbf{I})$. The posterior mean is a latent that encapsulates the prior knowledge of the entire dataset's distribution. The first term, \mathbf{x}_0 ensures content preservation during the attack, while the distributional term is responsible for aligning the data distribution to remain post-attack motion quality.

Different from the previous methods that regularize the attack process using point-wise difference, i.e., with respect to \mathbf{x}_0 , we add the regulation over the latent $\mathbf{x}_t \sim p(\mathbf{x}_t|\mathbf{x}_0)$ to regularize the data distribution during attack.

In summary, the motions during attack are represented by

$$\boldsymbol{\kappa}_t = \boldsymbol{\mu}_t + \mathbf{x}_t, \quad (5)$$

which will be leveraged later in the attack process as discussed in Section IV-B.

c) Benefits of Distributional Latent: As shown in Eq. 5, the constructed latent, i.e., the posterior mean with regularization, is derived from the data distribution rather than from a single sample. Compared with deterministic latents [59], the stochastic latents are more efficient and have much higher capacity [54]. Moreover, the posterior mean represents a direction that conforms to the data distribution toward high-density




Fig. 3: The visualization of diffusion latents at different timesteps. As shown, the earlier timesteps maintain more low-level details, and the later timesteps focus on high-level structures until latents become pure noise.

areas associated with the given labels. When we optimize over this latent for the adversarial attack, our empirical risk is based on all data samples from the approximated distribution, instead of a single sample. Additionally, instead of computing over the entire diffusion chain, we focus on only one denoising step to minimize changes in a single optimization iteration. Through sampling different timesteps, the constructed latent captures different features that encode underlying semantics.

B. The Attack Strategy for Quality-Preserving Attack

a) *Overview of Strategy:* To leverage the diffusion latent space for adversarial attacks, we propose perturbing the source motion observation by randomly sampling different class labels rather than relying on the gradient of a specific classifier. This is not only because obtaining the gradient of a specific classifier is challenging in real-world applications [2], but also because the gradient of a classifier may not be reliable for preserving motion quality. As classifiers focus on the label distribution rather than data distribution, their gradient may point to out-of-distribution regions. While following this gradient achieves the shortest trajectory of deceiving classifiers, it leads to a decline in post-attack motion quality and undermines the imperceptibility. As a result, we design our method to only rely on the classifier’s decisions to determine when to stop and drive the motion towards adversarial samples by sampling adversarial labels for the conditional diffusion model.

b) *Optimization Objective:* Specifically, we denote the stochastic distributional latents of the source motion observation and the adversarial motion as κ_t^{src} and κ_t^{adv} . The source motion and the adversarial motion are mapped to the latent space by the diffusion forward process and then the desired posterior mean are obtained within a single timestep denoising conditioned on the ground truth label \mathbf{y}^{src} and the randomly sampled the adversarial label \mathbf{y}^{adv} from the set of all possible labels excluding the ground truth label, respectively. The regularizer is obtained through the forward process. We define our objective function as follows:

$$\begin{aligned} \mathcal{L}_{\kappa_t} &:= 0.5 \times \mathbb{E}_{t, \epsilon} \left[\|\kappa_t^{\text{adv}} - \kappa_t^{\text{src}}\|_2^2 \right], \\ &= 0.5 \times \mathbb{E}_{t, \epsilon} \left\| \underbrace{\mu_t^{\text{adv}} - \mu_t^{\text{src}}}_{\text{latent}} + \underbrace{(\mathbf{x}_t^{\text{adv}} - \mathbf{x}_t^{\text{src}})}_{\text{regularization}} \right\|_2^2 \end{aligned} \quad (6)$$

where κ_t^{src} and κ_t^{adv} are our defined latents and serve as proxies for adversarial attack. The first term, κ_0^{adv} , represents

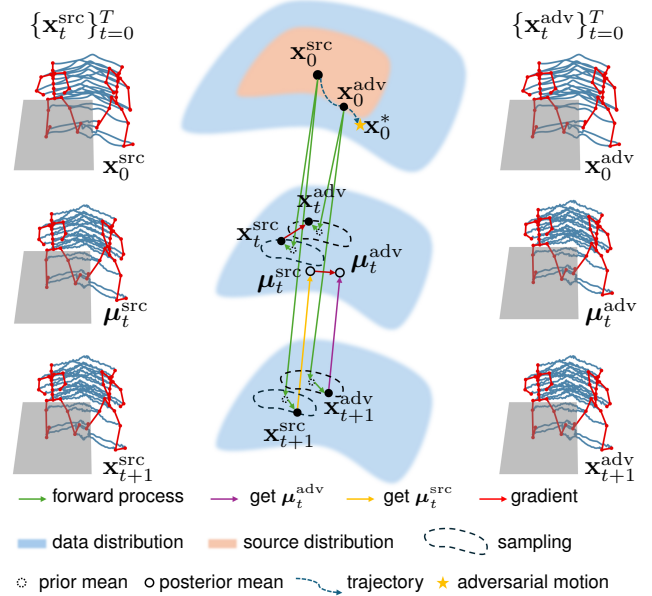


Fig. 4: The illustration of the attack strategy. We illustrate an intermediate calculation at the timestep t during the optimization of achieving the final adversarial motion $\mathbf{x}_0^{\text{adv}}$.

the direction towards being adversarial, while the second term, κ_0^{src} , represents the direction of maintaining the representative semantics within the original class. Our optimization objective is an expectation over the diffusion timesteps t and the stochastic distributional latent, indicated by the randomly sampled noise ϵ in the forward process where semantics are encoded into the latent space of the diffusion model.

By minimizing the defined object, our optimization aligns the stochastic distributional latents of the source and the adversarial motions. Specifically, we calculate the gradient of \mathcal{L}_{κ_t} with respect to κ_t^{adv} and obtain our adversarial gradient:

$$\begin{aligned} \text{grad} &:= \nabla \mathcal{L}_{\kappa_t} \\ &= \mathbb{E}_{t, \epsilon} \left[\kappa_t^{\text{adv}} - \kappa_t^{\text{src}} \right], \end{aligned} \quad (7)$$

on which we rely to iteratively update the $\mathbf{x}_0^{\text{adv}}$. The Eq. 7 facilitates the adversarial effect and quality preservation. Instead of directly requiring the paired pre-attack and post-attack motions to be close to each other in data space, our optimization aligns the latent distribution of pre-attack and

Algorithm 1 Diffusion-based Quality-Preserving Adversarial Motion Attack on S-HAR

Require: Diffusion model θ , a classifier φ , a motion $\mathbf{x}_0^{\text{src}}$ with label $\mathbf{y} \in \mathbf{Y}$, maximum iteration I , diffusion timesteps T

- 1: $\mathbf{x}_0^{\text{adv}} \leftarrow \mathbf{x}_0^{\text{src}}$
- 2: $i \leftarrow 0$
- 3: $\mathbf{y}^{\text{adv}} \sim \mathbf{Y} / \mathbf{y}^{\text{src}} \triangleright$ Randomly sample an adversarial label
- 4: **while** $i \leq I$ **and** $\mathbf{y}^{\text{pred}} \neq \mathbf{y}^{\text{src}}$ **do**
- 5: $t \sim [1, T]$
- 6: $\kappa_t^{\text{src}} = \mu_t(\mathbf{x}_0^{\text{src}}, \mathbf{y}^{\text{src}}; \theta) + \mathbf{x}_t^{\text{src}} \triangleright$ Eq. 5
- 7: $\kappa_t^{\text{adv}} = \mu_t(\mathbf{x}_0^{\text{adv}}, \mathbf{y}^{\text{adv}}; \theta) + \mathbf{x}_t^{\text{adv}} \triangleright$ Eq. 5
- 8: $\text{grad} = \kappa_t^{\text{adv}} - \kappa_t^{\text{src}} \triangleright$ Eq. 7
- 9: $\mathbf{x}^{\text{adv}} = \mathbf{x}^{\text{adv}} + \text{grad}$
- 10: $\mathbf{y}^{\text{pred}} = \arg \max p_\varphi(\mathbf{y} | \mathbf{x}^{\text{adv}}) \triangleright$ Get classifier decision
- 11: **end while**

post-attack motions. Constraining the two distributional latents κ_t^{adv} and κ_t^{src} promotes that the latent distributions of $\mathbf{x}_0^{\text{adv}}$ and $\mathbf{x}_0^{\text{src}}$ remain closely aligned rather than merely examine the pre-attack and the post-attack motions and rely on the gradients from an external classifier. Our optimization represents a single timestep examination, which is different from a multi-timestep generative process conditioned on \mathbf{y}^{adv} . By optimizing over the expectation of randomly sampled timesteps, the trajectory defined by the posteriors is expected to be closely aligned with each other. Consequently, our attack strategy facilitates the generation of $\mathbf{x}_0^{\text{adv}}$ that corresponds with \mathbf{y}^{adv} to deceive the target model, while simultaneously preserving the quality of $\mathbf{x}_0^{\text{src}}$. Detailed adversarial attack methodology is provided in Algorithm 1 and Fig. 4.

c) Relationship with Previous Optimization: We further demonstrate that our proposed method implicitly integrates previous approaches while additionally offer the advantage of distributional prior knowledge provided by a pre-trained diffusion model to minimize the risk gap. The Eq. 7 is equivalently represented using the input motion and the learned distribution, from which the following detailed formulation is derived:

$$\begin{aligned}
 \text{grad} &:= \mathbb{E}_{t, \epsilon_t} [\kappa_t^{\text{adv}} - \kappa_t^{\text{src}}] \\
 &= \mathbb{E}_{t, \epsilon_t} \left[\underbrace{\psi(t) (\mathbf{x}_0^{\text{adv}} - \mathbf{x}_0^{\text{src}})}_{\text{Distance Function in Original Space}} \right. \\
 &\quad \left. + \chi(t) \left(\underbrace{(\nabla \log p(\mathbf{x}_t^{\text{adv}}) - \nabla \log p(\mathbf{x}_t^{\text{src}}))}_{\text{Distribution Constraint}} \right) \right. \\
 &\quad \left. + \left(\underbrace{\nabla \log p(\mathbf{y}^{\text{adv}} | \mathbf{x}_t^{\text{adv}})}_{\text{Adversarial Gradient}} - \underbrace{\nabla \log p(\mathbf{y}^{\text{src}} | \mathbf{x}_t^{\text{src}})}_{\text{Representiveness of the Given Input}} \right) \right], \quad (8)
 \end{aligned}$$

where we decompose the probability with Bayes' theorem, $\psi(t) = \sqrt{\alpha_{t+1}}$, and $\chi(t) = 1 - \alpha_{t+1}$. The four terms serve distinct yet complementary purposes in our objective function. The first term demonstrates a similar measurement to previous attack methods, ensuring that the adversarial sample does not deviate excessively from the input sample. However, our ability of quality preservation, which distinguishes

our method from previous methods, is ensured through the remaining terms. Unlike traditional adversarial optimization, which considers only an individual motion pair as shown in the first term, our perturbation strategy provides the modification gradient based on data distributions that are composed of infinite motions. The second term measures the distribution density and ensures that the adversarial motions remain not only neighborhood but also in high-density areas. The third term considers the adversarial gradient while the fourth term involves the representativeness of the observed sample as a consideration for the range of pre-attack motion neighborhood by evaluating how representative the given input is with respect to the label. These considerations leverage the distributional prior knowledge and enables the quality preservation under adversarial modifications.

V. HUMAN PERCEPTION ALIGNED NATURALNESS METRIC

In this section, we examine previous metrics and formulate the proposed metric to faithfully measure the quality of generated adversarial motions in terms of naturalness.

a) Unfaithfulness of Existing Metrics: The evaluation of existing metrics fails to accurately measure quality. Previous metrics rely on paired comparisons between the pre-attack and post-attack motions to assess quality, especially naturalness [5]. However, remaining within the neighborhood of a pre-attack motion does not ensure motion quality comparable to clean motions because natural and plausible motions are sparsely distributed [15]. Consequently, these metrics cannot reliably determine whether adversarial motions are sufficiently natural due to their misalignment with human perception.

b) Human Perception Aligned Metric: Natural human motions are subject to biomechanical constraints where movements are driven by muscle activation. Every action results from the brain sending electrical signals to nerves, which in turn contract muscles to move joints [60]. The human brain develops prior knowledge of natural motions based on real-world observations of biomechanically validated musculoskeletal movements [61] rather than paired comparison. Therefore, human perception is inherently biased toward the naturalness self-embedded in the observed motion.

To faithfully reflect naturalness, we propose a novel metric that evaluates the physiological naturalness of generated adversarial motions. Since real-world movements are created by muscle activation, our metric is grounded in the bio-activity of muscles. Specifically, the muscle activation patterns of natural human motions are constrained by physiological characteristics [22], [23], resulting in smooth activation curves. We introduce a new metric to measure quality from the biomechanical perspective of muscle activation:

$$\text{naturalness} = \frac{1}{T} \int_0^T \frac{d^4 J(t)}{dt^4} dt, \quad (9)$$

where T is the temporal length of a motion, $J(t)$ is the joint position at time t . BASAR [62] mainly enforces naturalness as source-relative manifold plausibility, whereas our work targets distribution-level quality preservation and evaluates motion naturalness as an intrinsic biomechanical property of the attacked motion itself.

TABLE I: Generated Adversarial Motion Quality Comparison on 100STYLE dataset [24].

Victim	Method	Success Rate \uparrow	FID \downarrow	MMD \downarrow	Physiological Naturalness \downarrow	Foot Skating \downarrow	Bone Variation \downarrow
Style [63]	I-FGSM	81.72%	194.95	0.053	129.42	0.147	10.96
	MI-FGSM	82.08%	195.22	0.053	131.33	0.147	11.16
	MIG	70.61%	238.23	0.071	264.46	0.234	23.22
	SMART	42.65%	242.12	0.072	97.73	0.119	4.00
	Ours	100%	18.91	0.011	16.05	0.075	2.94
STTFormer [64]	I-FGSM	80.29%	162.55	0.043	190.98	0.196	19.28
	MI-FGSM	80.29%	162.78	0.043	191.10	0.197	19.29
	MIG	72.76%	190.75	0.050	280.56	0.241	27.63
	SMART	29.39%	195.06	0.053	103.56	0.167	3.92
	Ours	100%	20.28	0.013	16.42	0.077	3.09
SkateFormer [65]	I-FGSM	68.46%	79.63	0.019	43.90	0.053	4.53
	MI-FGSM	68.82%	79.06	0.019	44.70	0.053	4.61
	MIG	60.22%	102.77	0.026	70.72	0.069	7.30
	SMART	31.90%	125.46	0.030	29.32	0.047	1.60
	Ours	100%	20.16	0.014	16.33	0.079	2.99
FR-Head [66]	I-FGSM	86.38%	226.80	0.068	312.90	0.276	25.65
	MI-FGSM	86.74%	226.67	0.068	316.77	0.280	26.08
	MIG	78.14%	242.85	0.075	489.37	0.396	38.12
	SMART	32.97%	262.75	0.100	214.63	0.301	6.33
	Ours	100%	19.13	0.011	16.75	0.084	3.23
Average	I-FGSM	79.21%	165.98	0.046	169.3	0.168	15.11
	MI-FGSM	79.48%	165.93	0.046	170.98	0.169	15.29
	MIG	69.34%	193.65	0.056	276.28	0.235	24.07
	SMART	34.23%	206.35	0.064	111.31	0.159	3.96
	Ours	100%	19.62	0.012	16.39	0.079	3.06

VI. EXPERIMENTS

In this section, we begin by outlining the experimental settings. Subsequently, we quantitatively and qualitatively analyze the performance of our proposed method and existing adversarial techniques. Finally, we perform a user study to empirically evaluate the imperceptibility of post-attack motions and also ablation studies to validate our configurations.

A. Experimental Settings

a) Datasets: To evaluate our imperceptible adversarial attacks via quality preserving, we select the 100STYLE [24] dataset due to its noise-free and inherently high-quality characteristics. The 100STYLE dataset is collected using a motion capture system and comprises 100 classes of different styles. We represent the skeletons using Cartesian coordinates for 23 joints. This dataset is pre-processed by segmenting long sequences into several segments according to the valid periods provided by [24]. Additionally, we employ the HDM05 [25] dataset and NTU60 [9] to evaluate our method. HDM05 and NTU60 are commonly used for action recognition tasks with lower quality compared to 100STYLE. We adhere to the pre-processing procedure outlined in [5]. The HDM05 dataset includes 65 classes of different human actions, and the hip joint is fixed to the origin after pre-processing.

b) Evaluated Models: Given the significant advancements in the field of human action recognition and defense methods, we adopt the latest S-HAR models as victim classifiers with randomized smoothing [67] as a defense method to effectively evaluate the performance of adversarial attack methods against advanced classifiers. Specifically, we select the Style classifier [63], STTFormer [64], Skateformer [65], and FR-Head [66] as victim models. These models encompass

both transformer-based and graph-based architectures. We utilize their publicly available codebases to train their models.

c) Evaluation Metrics: Human motions are governed by physical and biomechanical constraints, thereby necessitating the assessment of visual motion quality in terms of naturalness and plausibility [51]. To evaluate the motion quality in adversarial attacks, we introduce the physiological naturalness metric, as discussed in Section V, which is grounded in the physiological characteristics of natural human movements. Additionally, we report the Frechet Inception Distance (FID) and Maximum Mean Discrepancy (MMD) based on acceleration to measure the naturalness by assessing the distributional similarity between pre-attack and post-attack motions. As for plausibility, we report the foot skating ratio and bone length variations between frames. Following [68], foot skating is quantified by the consistency between foot velocity and foot height. A high foot skating ratio indicates a significant violation in terms of physical constraints. Bone length variation [69] measures the consistency of bone lengths across frames. Higher deviations indicate distortions in the skeleton structure. Finally, we report the success rate of the adversarial attacks to evaluate the threatfulness.

d) Attacking Methods: We compare our method with the state-of-the-art (SOTA) S-HAR attack technique, i.e. SMART [5], as well as other adversarial attack methods including I-FGSM [70], MI-FGSM [71], and MIG [72]. To ensure a fair comparison, we execute 2000 iterations for each attack method, allowing all methods to explore a broader solution space in their pursuit of effective adversarial motions. Since our method requires a pre-trained diffusion model, we adopt the diffusion model proposed by [73] and follow the prescribed settings to pre-train it on the two datasets.

TABLE II: Generated Adversarial Motion Quality Comparison on HDM05 dataset.

Victim	Method	Success Rate \uparrow	FID \downarrow	MMD \downarrow	Physiological Naturalness \downarrow	Bone Variation \downarrow
Style [63]	I-FGSM	93.91%	129.78	0.086	1226.33	126.39
	MI-FGSM	93.91%	129.79	0.086	1226.98	126.46
	MIG	92.47%	194.33	0.170	2273.09	238.34
	SMART	68.46%	142.39	0.116	1052.44	41.00
	Ours	100%	22.91	0.012	187.28	32.15
STTFormer [64]	I-FGSM	86.74%	162.98	0.124	1488.08	144.97
	MI-FGSM	86.74%	163.18	0.125	1487.92	144.95
	MIG	76.34%	202.24	0.188	1945.83	189.95
	SMART	75.63%	174.86	0.129	974.33	33.30
	Ours	100%	21.89	0.011	185.73	30.31
Skateformer [65]	I-FGSM	83.87%	36.66	0.022	521.52	54.50
	MI-FGSM	84.23%	36.83	0.022	523.93	54.83
	MIG	78.14%	57.59	0.037	692.53	78.87
	SMART	65.59%	50.84	0.028	364.26	13.54
	Ours	100%	18.67	0.011	181.89	28.09
FR-Head [66]	I-FGSM	96.06%	169.34	0.124	1902.39	177.91
	MI-FGSM	96.77%	169.96	0.126	1920.98	179.56
	MIG	92.47%	215.88	0.206	2757.84	243.96
	SMART	81.72%	211.88	0.173	1574.95	55.46
	Ours	99.28%	20.80	0.013	182.96	31.07
Average	I-FGSM	90.15%	124.69	0.089	1284.58	125.94
	MI-FGSM	90.41%	124.94	0.090	1289.95	126.45
	MIG	84.86%	167.51	0.150	1917.32	187.78
	SMART	72.85%	144.99	0.112	991.50	35.83
	Ours	99.82%	21.07	0.012	184.47	30.41

TABLE III: Generated Adversarial Motion Quality Comparison on NTU60 dataset.

Method	Success Rate \uparrow	FID \downarrow	MMD \downarrow	Physiological Naturalness \downarrow	Bone Variation \downarrow
I-FGSM	95.07%	3.73	0.002	398.99	40.33
MI-FGSM	95.43%	3.60	0.002	400.57	40.46
MIG	93.10%	8.10	0.003	412.59	44.25
SMART	67.57%	432.67	0.377	403.91	36.94
Ours	100%	10.50	0.005	395.78	39.93

B. Adversarial Motion Quality Evaluation

We qualitatively and quantitatively evaluate the performance of adversarial attack in terms of deceitfulness, motion quality including naturalness and plausibility, and human imperceptibility. Deceitfulness measures the effectiveness of an adversarial method. In addition to deceitfulness, motion quality assesses whether the adversarial motions are plausible and natural, and serves as an indicator for imperceptibility. Beyond analytical measurements, we empirically evaluate human imperceptibility. Imperceptibility demonstrates the possibility that humans cannot distinguish pre-attack and post-attack motions when they are mixed together as S-HAR inputs. For simplicity, we present only representative results in this paper.

a) Deceitfulness: The success rates presented in Table I indicate that our method is the most deceitful compared with other methods on the 100STYLE dataset. Our method consistently achieves an average success rate of 100% even in the face of a randomized smoothing defense strategy. Our attack effectively modifies nearly every input to become adversarial without relying on the gradient of victim model within a limited number of iterations. In contrast, other methods struggle to generate adversarial motions within iteration constraints. This demonstrates that the stochastic latent features convey comprehensive semantics [74] about the underlying dynamics via different timesteps.

b) Motion Quality: Regarding naturalness, the FID and MMD scores in Table I show that the distribution of our generated motions closely resembles that of natural motions, whereas other methods exhibit significant distributional deviations. Our adversarial motions maintain proximity to the ground truth distribution of motion dynamics by a considerable margin. The lowest FID and MMD scores indicate that our adversarial motions are indistinguishable from those in the ground truth dataset. In other words, the post-attack motion quality is well preserved because our method leverages the distributional knowledge introduced by the pre-trained diffusion model. By utilizing data distribution, our method minimizes the gap between empirical and true risks by assessing deviations between the modified motions and the data distribution. Conversely, other methods can only access single source motions, and disrupt the naturalness of post-attack motions, resulting in much higher FID and MMD scores.

Physiological naturalness reflects whether adversarial motions conform to natural movements from the biomechanical perspective of muscle activation. Our method achieves the best performance by a large margin, as shown in Table I. This indicates that even after adversarial attacks, our generated motions retain plausible movements that real-world humans can physically perform. To further evaluate naturalness, as illustrated in Fig. 5, we present the mean power spectral density

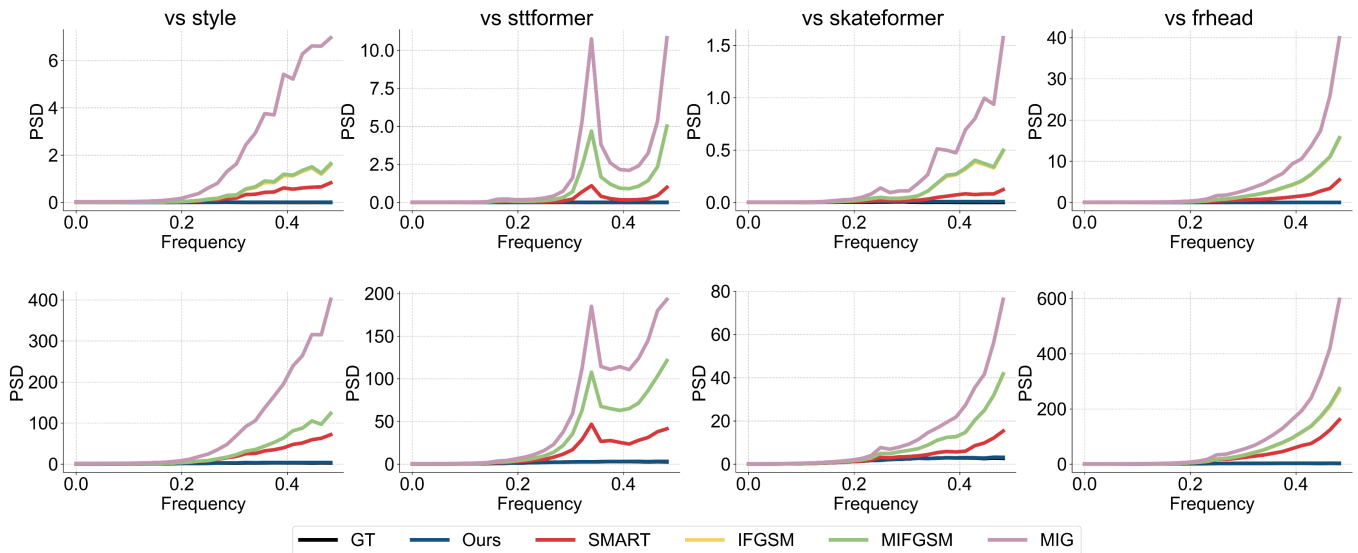


Fig. 5: The mean power spectral density of adversarial samples found on 100STYLE (upper row) and HDM05 (lower row) against four classifiers.




Fig. 6: The visualization of acceleration changes.

(PSD) of adversarial samples. The spectral density of our post-attack motions in the high-frequency domain is significantly lower than that of other adversarial motions. Besides, our spectral density closely aligns with the ground truth curve. This suggests that our adversarial motions suffer from few noise-like perturbations because we minimize empirical risk over the data distribution rather than a single sample.

We further compare the effect of explicitly regulating motion dynamics with respect to a given sample in SMART and implicitly regulating motion dynamics through data distribution in our diffusion-based method. Fig. 6 displays the acceleration changes of a sample for comparison. Our method introduces modifications that are not only smaller in magnitude but also smoother and more consistent, whereas SMART does not perturb the motion coherently and consistently. These results imply that minimizing empirical risk over only a given sample leads to deviations from the natural motion distribution and results in the post-attack quality decline.

In terms of physical plausibility, the foot sliding ratio and bone length variation in Table I demonstrate that our generated adversarial motions exhibit superior movement coherency and skeletal consistency compared to other methods. The foot skating ratio indicates the coherence of movements concerning foot contact with the ground. Our method achieves the lowest foot skating ratio, suggesting that post-attack movements are

still well synchronized. This is attributed to our method’s ability to minimize the risk gap through the integration of data distribution rather than relying on a single sample. Additionally, our adversarial motions exhibit the lowest variation in bone lengths, maintaining consistent skeletal structures across frames. By leveraging access to the data distribution rather than individual samples, our method ensures that modifications preserve cross-frame skeletal integrity.

We also achieve the best performance on the HDM05 dataset, as shown in Table II. Our method successfully fools all systems with an average success rate of 99.82%. This further validates the efficacy of our attack method even when the data quality and amount are not as high as those of the 100STYLE dataset, demonstrating its threatfulness to such systems. Moreover, the distribution of our adversarial motions closely aligns with the ground truth dataset, evidenced by the lowest FID and MMD scores. This validates that our adversarial motions are as natural as pre-attack motions. Besides, physiological naturalness indicates that our generated motions adhere to real-world biomechanical constraints. Since the HDM05 dataset data have been centered to the origin of the coordinates, we do not report the foot skating ratio for this dataset. Additionally, our method achieves superior motion plausibility with the lowest bone length variation. Both the best physiological and physical plausibility demonstrate that our method preserves the quality of adversarial motions.

The NTU60 dataset has also been tested with a randomized smoothing defence strategy as shown in Table III. Our method still outperforms other methods in terms of success rate and physiological naturalness. However, the motions in this dataset suffer heavily from inherent noise. The adversarial noise introduced by other methods cannot be reflected by the FID and MMD due to the inherent noise from data.

c) *Human Imperceptibility*: We visualize the trajectories of joints in adversarial motions generated by our method and previous methods in Fig. 7 against the four victim models. In real-world natural movements, joint trajectories are typically

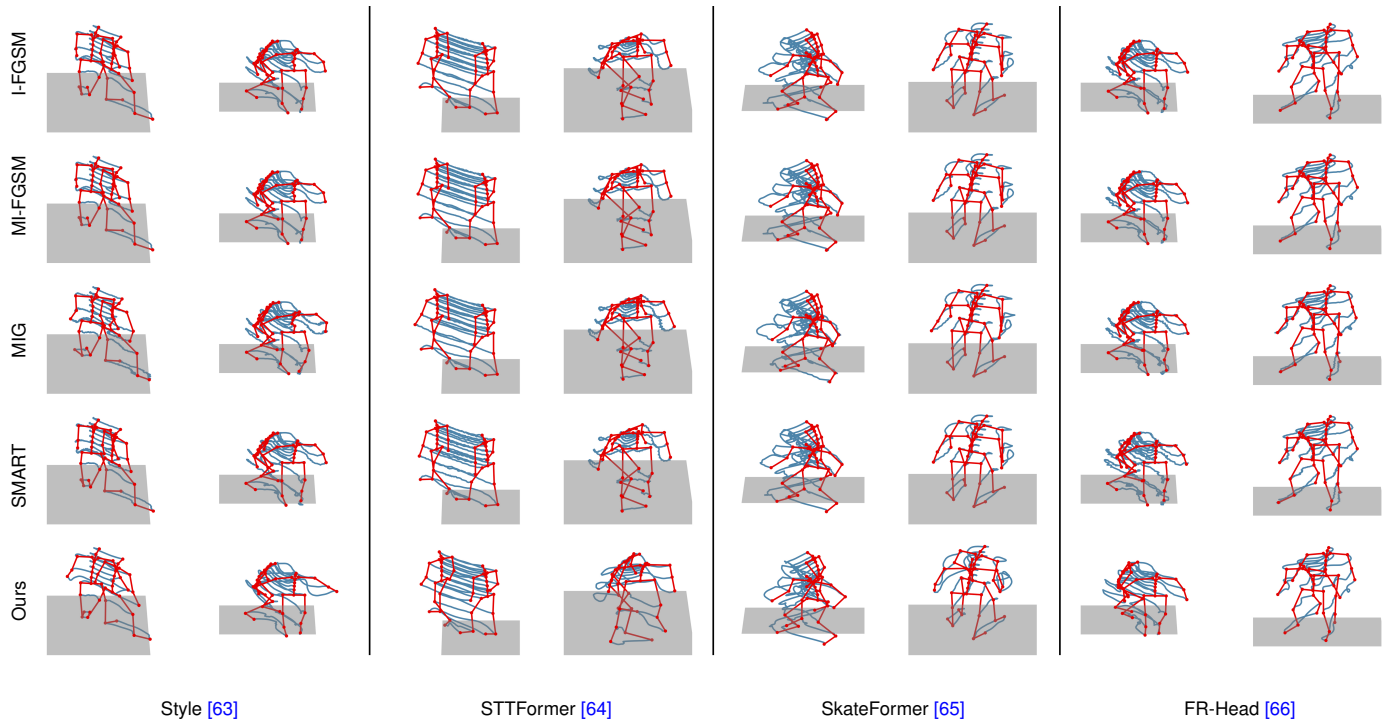


Fig. 7: Visual comparison among the adversarial motions generated by different attack methods against victim models. We visualize the starting and the ending poses in red, the trajectories of all joints in blue, and the ground floor in grey. Our adversarial motions exhibit the most natural and stable trajectories.

smooth and stable. As shown in Fig. 7, our method produces the most stable and smooth trajectories for all joints in the adversarial motions. By minimizing the risk gap over the data distribution rather than a single motion, our optimization results do not introduce noise-like perturbations, whereas other methods suffer from significant gaps by optimizing over a single input. Consequently, the noise-like perturbations lead to a decline in the post-attack motion quality with observable unstable trajectories and undermine imperceptibility.

Additionally, we recruited volunteers without visual impairments from diverse backgrounds and a balanced gender representation. They participate in questionnaires to evaluate the imperceptibility of our adversarial samples. We provided participants with batches of motions from the same label, which pre-attack motions and post-attack motions generated by our method are mixed together. Participants are asked to select the motions they consider most likely to have potentially been attacked without any time constraints. As illustrated in Fig. 8, our method produces adversarial motions that are the least perceivable by humans.

C. Black-Box Settings

Our method does not require gradient access to the victim classifier during optimization. To further evaluate its effectiveness under restricted-access conditions, we additionally report black-box experiments. For the gradient-based baselines, we adopt a transfer-based black-box protocol: for each victim model, the other three models are used as substitute models in turn, and we report the averaged transfer performance. The




Fig. 8: Perceptibility comparison across different methods.

same maximum attack budget as in the main paper is used, namely 2000 iterations. By contrast, our method does not rely on surrogate models or victim gradients during optimization. In the current black-box evaluation, the victim model is only checked after the full attack procedure to determine whether the generated adversarial motion succeeds. Randomized smoothing is applied consistently across all reported results. Note that, unlike the transfer-based baselines, our method does not consult the victim model during optimization and therefore cannot exploit intermediate victim-side stopping decisions. Table IV shows the results on 100STYLE dataset in the black-box setting. Our method shows superiority on both deceitfulness and naturalness. Notably, our presented

TABLE IV: Generated Black-box Adversarial Motion Quality Comparison on 100STYLE dataset.

Victim	Method	Success Rate \uparrow	FID \downarrow	MMD \downarrow	Physiological Naturalness \downarrow	Bone Variation \downarrow
Style [63]	I-FGSM	49.46%	281.00	0.097	79.29	11.43
	MI-FGSM	51.61%	283.70	0.098	87.51	12.61
	MIG	67.38%	307.14	0.105	414.82	66.07
	SMART	49.46%	240.27	0.071	57.87	4.83
	TASAR	51.97%	182.23	0.060	46.87	4.96
	Ours	96.42%	57.50	0.020	21.74	5.90
STTFormer [64]	I-FGSM	36.20%	194.05	0.046	62.67	11.21
	MI-FGSM	39.07%	200.65	0.048	69.48	12.47
	MIG	65.95%	297.48	0.104	293.88	52.50
	SMART	35.13%	193.07	0.049	57.76	4.80
	TASAR	39.78%	162.31	0.039	43.21	4.90
	Ours	93.90%	62.01	0.023	25.42	6.10
Skateformer [65]	I-FGSM	41.22%	134.05	0.036	18.77	3.70
	MI-FGSM	42.65%	138.07	0.037	20.23	4.02
	MIG	55.20%	240.37	0.066	97.03	18.34
	SMART	39.07%	111.17	0.028	17.74	1.91
	TASAR	42.30%	99.25	0.023	15.21	2.01
	Ours	99.64%	19.51	0.012	10.97	2.86
FR-Head [66]	I-FGSM	39.78%	275.96	0.098	111.98	16.22
	MI-FGSM	45.52%	284.80	0.103	124.31	18.07
	MIG	70.97%	336.28	0.156	552.29	77.84
	SMART	43.37%	266.16	0.099	123.33	8.56
	TASAR	48.75%	243.41	0.079	104.72	7.42
	Ours	100%	20.02	0.012	11.21	3.06
Average	I-FGSM	41.67%	221.27	0.069	68.18	10.64
	MI-FGSM	44.71%	226.81	0.072	70.38	11.79
	MIG	64.88%	295.32	0.108	339.51	53.69
	SMART	41.76%	202.67	0.062	64.18	5.03
	TASAR	45.70%	171.80	0.050	52.50	4.82
	Ours	97.49%	39.76	0.017	17.34	4.48

TABLE V: The quality of adversarial motions generated via different variants and configurations.

Variants	Configuration	Success Rate \uparrow	FID \downarrow	MMD \downarrow	Physiological Naturalness \downarrow	Foot Skating \downarrow	Bone Variation \downarrow
Timestep	[1, 20]	100%	14.01	0.007	13.28	0.090	3.22
	[980, 1000]	100%	26.14	0.017	30.37	0.050	1.62
Latent	κ_t	100%	18.91	0.011	16.05	0.075	2.94
	\tilde{x}_0	100%	19.26	0.011	56.70	0.083	3.31
Architecture	Diffusion [73]	100%	18.91	0.011	16.05	0.075	2.94
	VAE [56]	96.06%	198.00	0.056	1058.08	0.083	214.75

attack method does not involve explicit constraints such as the bone variation while SMART relies on a series of heuristic constraints for the optimization. The experiment shows that our method is versatile to be applied for both white-box and black-box settings without any modification.

D. Ablation Study

Given that our method is based on the latent space of a diffusion model, we first perform an ablation study to investigate the influence of the chosen timesteps in constructing the stochastic latent space. Secondly, we examine the impact of different latents used for adversarial attacks to determine which latent is the most suitable for the quality-preserving imperceptible attack task. Finally, we validate the choice of diffusion models by exploring alternative generative models for quality-preserving adversarial attack.

1) *Timestep of Stochastic Latent Space*: We investigate the impact of different timestep ranges used to map motions from data space to diffusion latent space for adversarial attacks. It has been shown that the semantics encapsulated by the latents

have smooth transitions with respect to the diffusion timesteps [75], [76], we conduct experiments using two representative timestep ranges: the earliest timestep range ($t \sim [1, 20]$) and the latest timestep range ($t \sim [980, 1000]$) for comparison.

A trade-off exists between motion naturalness and plausibility, as shown in Table V. This trade-off effect arises from the different emphasis on semantics encoded in the latents. Generally, latents derived from earlier timesteps primarily capture low-level details, while those from later timesteps focus on high-level structural patterns [77]. Utilizing earlier timesteps to construct the stochastic distributional latents is biased toward detailed movements, with a reduced understanding of global coherency and consistency. This bias enhances naturalness by conveying motion dynamics. However, it also exacerbates foot skating and bone length variation due to the lack of global rationality in the latents. Conversely, constructing the distributional latents using later timesteps undermines naturalness, as the generated adversarial motions deviate from the natural distribution, leading to higher FID and MMD scores.

TABLE VI: Generated Adversarial Motion Quality Comparison on 100STYLE dataset against Style [63]. Expectation over Transformation (EoT) [33] has been integrated on I-FGSM, MI-FGSM, MIG, and SMART.

Method	Success Rate \uparrow	FID \downarrow	MMD \downarrow	Physiological Naturalness \downarrow	Foot Skating \downarrow	Bone Variation \downarrow
I-FGSM	81.72%	185.09	0.051	117.48	0.129	9.91
MI-FGSM	81.72%	185.58	0.052	116.92	0.128	9.86
MIG	70.97%	241.43	0.066	281.46	0.222	24.96
SMART	81.00%	247.50	0.079	341.36	0.306	32.63
Ours	100%	18.91	0.011	16.05	0.075	2.94

2) *Alternative Latents for Attack*: We examine the influence of different latents used for our quality-preserving adversarial attacks. Specifically, we compare the constructed latent with the predicted pre-attack data as the latent for adversarial attack. The constructed latent κ evaluates the required perturbations based on a single timestep of denoising with regularization, and allows for step-by-step motion modifications. In contrast, the predicted pre-attack data \hat{x}_0 represents modifications along the entire chain of previous timesteps with approximation. We conduct experiments to determine whether the constructed latent is more effective to preserve post-attack motion quality than leveraging the predicted pre-attack data.

As shown in Table V, using the predicted pre-attack data as the latent leads to a consistent decline in all performances of adversarial attack. Although both latent choices achieve the same success rate, the quality of adversarial motions deteriorates significantly when switching from the constructed latent to the predicted data. This deterioration stems from the approximation error when calculating the required latents. The predicted \hat{x}_0 is derived from the x_{t+1} by approximating all previous timesteps $\{i\}_{i=1}^t$ collectively, thereby ignoring information from intermediate timesteps. Conversely, the constructed latent κ_t , also derived from the x_{t+1} , considers only the desired changes within a single timestep. This finer-grained modification enhances the quality preservation during adversarial attacks by integrating fewer approximation errors.

3) *Alternative Distribution Modeller*: Given that our attack method is based on generative models, we explore the use of alternative generative architectures. Similarly, we utilize their latent spaces to modify motions adversarially. Specifically, we conduct experiments using either diffusion models or variational autoencoders (VAEs). We train the VAE with label conditions following the architecture and training configurations outlined in [56]. To maintain consistency with our use of the latent distribution, we employ the latent distribution in the VAE. All other configurations remain consistent with those used in our diffusion-based method.

As shown in Table V, the performance of adversarial attacks using a VAE model significantly declines compared to our diffusion-based method. It indicates that effective motion modification requires a comprehensive representation of underlying patterns. In contrast to diffusion models, VAEs utilize a single latent feature rather than a hierarchy of features. Consequently, using a single latent feature in VAEs results in ambiguity regarding motion dynamics and leads to under-expressed movement coherency and consistency.

4) *Adaptive Strategy*: Table VI reports the results after equipping the gradient-based baselines with EoT under the randomized smoothing setting. Compared with the original

comparison, adaptive gradient estimation improves the effectiveness of some baselines, most notably SMART, whose success rate increases substantially to 81.00%, indicating that part of its previous performance degradation indeed came from the stochastic defense rather than the attack formulation itself. However, even after this stronger adaptive evaluation, our method still consistently outperforms all baselines by a large margin in both attack success and post-attack motion quality. In particular, our method achieves a 100% success rate while maintaining dramatically lower FID, MMD, physiological naturalness, foot skating, and bone variation, showing that the advantage of our approach does not rely on a non-adaptive evaluation protocol. These results suggest that although EoT can partially recover the effectiveness of gradient-based attacks against randomized models, such methods still tend to introduce noticeable quality degradation, whereas our distribution-based attack remains substantially more effective at preserving the naturalness and plausibility of adversarial motions. A possible explanation is that EoT makes the estimated gradient more consistent with the randomized victim, which improves attack success, but it does not alter the fact that these baselines still optimize adversarial perturbations in the original sample space. As a result, the recovered adversarial strength comes at the expense of larger motion distortion, whereas our method constrains the attack through the learned motion distribution and therefore preserves motion quality much more effectively.

TABLE VII: Runtime comparison between our method and I-FGSM. We report the average runtime per attack and the average runtime per iteration.

Method	Per attack (s)	Per iteration (s)
I-FGSM	18.05	0.06
Ours	1.72	0.07

5) *Time Cost Analysis*: We further evaluate the computational efficiency of different attack methods by reporting two runtime statistics. The first is the average runtime per attack, computed as the total time of all attack attempts divided by the number of attacks, which reflects the effective time cost required to obtain one adversarial motion. The second is the average runtime per iteration, which measures the computational overhead of a single optimization step. As shown in Table VII, our method requires 1.72 seconds per attack on average, compared with 18.05 seconds for I-FGSM, while the per-iteration runtime is 0.07 seconds for our method and 0.06 seconds for I-FGSM. These results show that the additional denoising module introduces only a marginal per-iteration overhead, whereas the overall runtime is substantially reduced in terms of attack generation. The iteration speed is

comparable because our method does not rely on iterative gradient back-propagation through the victim classifier and only performs a single-step denoising computation over two adjacent diffusion timesteps rather than traversing the full diffusion chain.

VII. CONCLUSION

We propose a novel attack application that imperceptible adversarial motions are achieved without compromising post-attack motion quality. Additionally, we introduce a distribution-based adversarial attack method targeting skeleton-based human action recognition (S-HAR) systems by minimizing the optimization gap inherent in previous approaches. Our method integrates a generative diffusion model, wherein the posterior mean of single timestep denoising with regularization is constructed as the proxy to fulfill our attack strategy. To faithfully assess the naturalness of adversarial motions, we develop a new metric aligned with human perception of natural, real-world human movements. We evaluate the quality of adversarial motions in terms of threatfulness, motion quality, and imperceptibility, demonstrating that our adversarial motions achieve superior performance across these metrics. The success of our proposed quality-preserving attack application and distribution-based attack method raises significant concerns regarding the robustness of action recognizers, highlighting the necessity for further enhancements in this area.

While achieving natural adversarial motions, our approach presents opportunities for future research. As shown in our experiment, there exists a trade-off in motion quality with respect to the timesteps of the diffusion model. Future work can potentially investigate the influence of trading-off post-attack motion quality on different S-HAR systems. It is also valuable to investigate noise-aware or disentangled diffusion priors that explicitly separate motion dynamics from acquisition noise, especially for lower-fidelity datasets where the learned data distribution may itself contain substantial noise. Our proposed physiological naturalness focuses on the physiological perspective of motion naturalness. Future work can potentially integrate the research field of action quality assessment [78], [79], [80] and character animation [81] for more comprehensive motion quality measurements as well as constraints.

REFERENCES


- [1] B. Ren, M. Liu, R. Ding, and H. Liu, "A survey on 3d skeleton-based action recognition using learning method," *Cyborg and Bionic Systems*, vol. 5, p. 0100, 2024.
- [2] Y. Diao, B. Wu, R. Zhang, A. Liu, X. Wei, M. Wang, and H. Wang, "Tasar: Transferable attack on skeletal action recognition," *arXiv preprint arXiv:2409.02483*, 2024.
- [3] T. Bai, J. Luo, J. Zhao, B. Wen, and Q. Wang, "Recent advances in adversarial training for adversarial robustness," in *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, Z.-H. Zhou, Ed. International Joint Conferences on Artificial Intelligence Organization, 8 2021, pp. 4312–4321, survey Track.
- [4] B. Chander, C. John, L. Warriar, and K. Gopalakrishnan, "Toward trustworthy artificial intelligence (tai) in the context of explainability and robustness," *ACM Computing Surveys*, 2024.
- [5] H. Wang, F. He, Z. Peng, T. Shao, Y.-L. Yang, K. Zhou, and D. Hogg, "Understanding the robustness of skeleton-based action recognition under adversarial attack," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 14 656–14 665.
- [6] J. Liu, N. Akhtar, and A. Mian, "Adversarial attack on skeleton-based human action recognition," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 33, no. 4, pp. 1609–1622, 2020.
- [7] N. F. Troje, "Decomposing biological motion: A framework for analysis and synthesis of human gait patterns," *Journal of vision*, vol. 2, no. 5, pp. 2–2, 2002.
- [8] S. Shimada and K. Oki, "Modulation of motor area activity during observation of unnatural body movements," *Brain and cognition*, vol. 80, no. 1, pp. 1–6, 2012.
- [9] A. Shahroudy, J. Liu, T.-T. Ng, and G. Wang, "Ntu rgb+ d: A large scale dataset for 3d human activity analysis," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1010–1019.
- [10] W. Kay, J. Carreira, K. Simonyan, B. Zhang, C. Hillier, S. Vijayanarasimhan, F. Viola, T. Green, T. Back, P. Natsev *et al.*, "The kinetics human action video dataset," *arXiv preprint arXiv:1705.06950*, 2017.
- [11] J. Liu, A. Shahroudy, M. Perez, G. Wang, L.-Y. Duan, and A. C. Kot, "Ntu rgb+ d 120: A large-scale benchmark for 3d human activity understanding," *IEEE transactions on pattern analysis and machine intelligence*, vol. 42, no. 10, pp. 2684–2701, 2019.
- [12] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in *International Conference on Learning Representations*, 2018.
- [13] A. Ilyas, S. Santurkar, D. Tsipras, L. Engstrom, B. Tran, and A. Madry, "Adversarial examples are not bugs, they are features," *Advances in neural information processing systems*, vol. 32, 2019.
- [14] A. Li, Y. Wang, Y. Guo, and Y. Wang, "Adversarial examples are not real features," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [15] K. Karunratanakul, K. Preechakul, E. Aksan, T. Beeler, S. Suwajanakorn, and S. Tang, "Optimizing diffusion noise can serve as universal motion priors," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 1334–1345.
- [16] L. Kovács, B. M. Bódis, and C. Benedek, "Lidpose: Real-time 3d human pose estimation in sparse lidar point clouds with non-repetitive circular scanning pattern," *Sensors*, vol. 24, no. 11, p. 3427, 2024.
- [17] J. Xu, Y. Guo, and Y. Peng, "Finepose: Fine-grained prompt-driven 3d human pose estimation via diffusion models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 561–570.
- [18] S. Shin, J. Kim, E. Halilaj, and M. J. Black, "Wham: Reconstructing world-grounded humans with accurate 3d motion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 2070–2080.
- [19] V. Vapnik, "Principles of risk minimization for learning theory," *Advances in neural information processing systems*, vol. 4, 1991.
- [20] L. Devroye, L. Györfi, and G. Lugosi, *A probabilistic theory of pattern recognition*. Springer Science & Business Media, 2013, vol. 31.
- [21] G. Claeskens and N. L. Hjort, "Model selection and model averaging," *Cambridge books*, 2008.
- [22] Y. Ueyama, "Costs of position, velocity, and force requirements in optimal control induce triphasic muscle activation during reaching movement," *Scientific Reports*, vol. 11, no. 1, p. 16815, 2021.
- [23] Y. P. Ivanenko, R. E. Poppele, and F. Lacquaniti, "Five basic muscle activation patterns account for muscle activity during human locomotion," *The Journal of physiology*, vol. 556, no. 1, pp. 267–282, 2004.
- [24] I. Mason, S. Starke, and T. Komura, "Real-time style modelling of human locomotion via feature-wise transformations and local motion phases," *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, vol. 5, no. 1, may 2022.
- [25] M. Müller, T. Röder, M. Clausen, B. Eberhardt, B. Krüger, and A. Weber, "Documentation mocap database hdm05," Universität Bonn, Tech. Rep. CG-2007-2, June 2007.
- [26] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.
- [27] D. Shanmugam, D. Blalock, G. Balakrishnan, and J. Guttag, "Better aggregation in test-time augmentation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 1214–1223.
- [28] M. Xie, Y. He, Z. Qin, and M. Fang, "Retouchuaa: Unconstrained adversarial attack via realistic image retouching," *IEEE Transactions*

- on *Circuits and Systems for Video Technology*, vol. 35, no. 3, pp. 2586–2602, 2025.
- [29] Y. Gan, C. Wu, D. Ouyang, S. Tang, M. Ye, and T. Xiang, “Lesep: Boosting adversarial transferability via latent encoding and semantic embedding perturbations,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 35, no. 4, pp. 3368–3382, 2025.
- [30] K. Jiang, Z. Chen, J. Fu, L. Hong, J. Li, and W. Zhang, “Videopure: Diffusion-based adversarial purification for video recognition,” *IEEE Transactions on Circuits and Systems for Video Technology*, pp. 1–1, 2025.
- [31] J. Huang, X. Yuan, C.-T. Lam, S.-K. Im, F. Lei, and X. Bi, “Transhfc: Joints hypergraph filtering convolution and transformer framework for temporal forgery localization,” *IEEE Transactions on Circuits and Systems for Video Technology*, 2025.
- [32] J. Wang, F. Li, and L. He, “A unified framework for adversarial patch attacks against visual 3d object detection in autonomous driving,” *IEEE Transactions on Circuits and Systems for Video Technology*, 2025.
- [33] A. Athalye, L. Engstrom, A. Ilyas, and K. Kwok, “Synthesizing robust adversarial examples,” in *International conference on machine learning*. PMLR, 2018, pp. 284–293.
- [34] Y. Li, W. Jin, H. Xu, and J. Tang, “Deeprobust: A pytorch library for adversarial attacks and defenses,” *arXiv preprint arXiv:2005.06149*, 2020.
- [35] F. Karim, S. Majumdar, and H. Darabi, “Adversarial attacks on time series,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 10, pp. 3309–3320, 2020.
- [36] Z. Lu, H. Wang, Z. Chang, G. Yang, and H. P. Shum, “Hard no-box adversarial attack on skeleton-based human action recognition with skeleton-motion-informed gradient,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 4597–4606.
- [37] Y. Diao, H. Wang, T. Shao, Y. Yang, K. Zhou, D. Hogg, and M. Wang, “Understanding the vulnerability of skeleton-based human activity recognition via black-box attack,” *Pattern Recognition*, vol. 153, p. 110564, 2024.
- [38] N. Tanaka, H. Kera, and K. Kawamoto, “Adversarial bone length attack on action recognition,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 2, 2022, pp. 2335–2343.
- [39] Y. Cao, K. Ye, D. Wang, M. Xue, H. Ge, C. Qian, and J. S. Dong, “Bones of contention: Exploring query-efficient attacks against skeleton recognition systems,” *IEEE Transactions on Information Forensics and Security*, vol. 21, pp. 183–196, 2025.
- [40] Z. Kang, Y. Zhang, R. Zhang, Y. Jiang, and H. Xia, “Qesar: Query effective decision-based attack on skeletal action recognition,” in *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*. Springer, 2023, pp. 417–429.
- [41] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” in *International Conference on Machine Learning*. PMLR, 2015, pp. 2256–2265.
- [42] Z. Chang, G. A. Koulouris, and H. P. Shum, “On the design fundamentals of diffusion models: A survey,” *arXiv preprint arXiv:2306.04542*, 2023.
- [43] R. Liu, W. Zhou, T. Zhang, K. Chen, J. Zhao, and K.-Y. Lam, “Boosting black-box attack to deep neural networks with conditional diffusion models,” *IEEE Transactions on Information Forensics and Security*, 2024.
- [44] X. Chen, X. Gao, J. Zhao, K. Ye, and C.-Z. Xu, “Advdiffuser: Natural adversarial example synthesis with diffusion models,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 4562–4572.
- [45] R. Liu, D. Wang, Y. Ren, Z. Wang, K. Guo, Q. Qin, and X. Liu, “Unstoppable attack: Label-only model inversion via conditional diffusion model,” *IEEE Transactions on Information Forensics and Security*, 2024.
- [46] C. Hu, Y. Li, Z. Feng, and X. Wu, “Towards transferable attack via adversarial diffusion in face recognition,” *IEEE Transactions on Information Forensics and Security*, 2024.
- [47] X. Hu, S. Li, Q. Ying, W. Peng, X. Zhang, and Z. Qian, “Establishing robust generative image steganography via popular stable diffusion,” *IEEE Transactions on Information Forensics and Security*, 2024.
- [48] Y. Chen, X. Li, X. Wang, P. Hu, and D. Peng, “Diffilter: Defending against adversarial perturbations with diffusion filter,” *IEEE Transactions on Information Forensics and Security*, 2024.
- [49] J. Yang, Z. Guan, J. Li, Z. Shi, and X. Liu, “Diffusion patch attack with spatial-temporal cross-evolution for video recognition,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 34, no. 12, pp. 13 190–13 200, 2024.
- [50] K. Jiang, Z. Chen, J. Fu, L. Hong, J. Li, and W. Zhang, “Videopure: Diffusion-based adversarial purification for video recognition,” *IEEE Transactions on Circuits and Systems for Video Technology*, 2025.
- [51] W. Zhu, X. Ma, D. Ro, H. Ci, J. Zhang, J. Shi, F. Gao, Q. Tian, and Y. Wang, “Human motion generation: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [52] S.-J. Blakemore and J. Decety, “From the perception of action to the understanding of intention,” *Nature reviews neuroscience*, vol. 2, no. 8, pp. 561–567, 2001.
- [53] E. Grossman, M. Donnelly, R. Price, D. Pickens, V. Morgan, G. Neighbor, and R. Blake, “Brain areas involved in perception of biological motion,” *Journal of cognitive neuroscience*, vol. 12, no. 5, pp. 711–720, 2000.
- [54] C. H. Wu and F. De la Torre, “A latent space of stochastic diffusion models for zero-shot image editing and guidance,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 7378–7387.
- [55] P. Jaini, K. Clark, and R. Geirhos, “Intriguing properties of generative classifiers,” in *The Twelfth International Conference on Learning Representations*, 2024.
- [56] M. Petrovich, M. J. Black, and G. Varol, “Action-conditioned 3d human motion synthesis with transformer vae,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10985–10995.
- [57] K. Preechakul, N. Chatthee, S. Wizadwongsa, and S. Suwajanakorn, “Diffusion autoencoders: Toward a meaningful and decodable representation,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 10 619–10 629.
- [58] K. Kim and J. C. Ye, “Noise2score: tweedie’s approach to self-supervised image denoising without clean images,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 864–874, 2021.
- [59] J. Song, C. Meng, and S. Ermon, “Denosing diffusion implicit models,” in *International Conference on Learning Representations*, 2021.
- [60] M. Chiquier and C. Vondrick, “Muscles in action,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 22 091–22 101.
- [61] D. Schneider, S. Reiß, M. Kugler, A. Jaus, K. Peng, S. Sutschet, M. S. Sarfraz, S. Matthiesen, and R. Stiefelwagen, “Muscles in time: Learning to understand human motion in-depth by simulating muscle activations,” *Advances in Neural Information Processing Systems*, 2025.
- [62] Y. Diao, T. Shao, Y.-L. Yang, K. Zhou, and H. Wang, “Basar: Black-box attack on skeletal action recognition,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, Conference Proceedings, pp. 7597–7607.
- [63] L. Zhong, Y. Xie, V. Jampani, D. Sun, and H. Jiang, “Smoodi: Stylized motion diffusion model,” in *European Conference on Computer Vision*. Springer, 2025, pp. 405–421.
- [64] H. Qiu, B. Hou, B. Ren, and X. Zhang, “Spatio-temporal tuples transformer for skeleton-based action recognition,” *arXiv preprint arXiv:2201.02849*, 2022.
- [65] J. Do and M. Kim, “Skateformer: Skeletal-temporal transformer for human action recognition,” in *European Conference on Computer Vision*. Springer, 2025.
- [66] H. Zhou, Q. Liu, and Y. Wang, “Learning discriminative representations for skeleton based action recognition,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 10 608–10 617.
- [67] J. Cohen, E. Rosenfeld, and Z. Kolter, “Certified adversarial robustness via randomized smoothing,” in *international conference on machine learning*. PMLR, 2019, pp. 1310–1320.
- [68] C. Guo, S. Zou, X. Zuo, S. Wang, W. Ji, X. Li, and L. Cheng, “Generating diverse and natural 3d human motions from text,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 5152–5161.
- [69] C. Duan, Z. Zhang, X. Liu, Y. Dang, and J. Yin, “Physics-constrained attack against convolution-based human motion prediction,” *Neurocomputing*, vol. 575, p. 127272, 2024.
- [70] A. Kurakin, I. J. Goodfellow, and S. Bengio, “Adversarial examples in the physical world,” in *Artificial intelligence safety and security*. Chapman and Hall/CRC, 2018, pp. 99–112.
- [71] Y. Dong, F. Liao, T. Pang, H. Su, J. Zhu, X. Hu, and J. Li, “Boosting adversarial attacks with momentum,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 9185–9193.
- [72] W. Ma, Y. Li, X. Jia, and W. Xu, “Transferable adversarial attack for both vision transformers and convolutional networks via momentum integrated gradients,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 4630–4639.

- [73] G. Tevet, S. Raab, B. Gordon, Y. Shafir, D. Cohen-or, and A. H. Bermano, "Human motion diffusion model," in *The Eleventh International Conference on Learning Representations*, 2023.
- [74] Y.-H. Park, M. Kwon, J. Choi, J. Jo, and Y. Uh, "Understanding the latent space of diffusion models through the lens of riemannian geometry," *Advances in Neural Information Processing Systems*, vol. 36, pp. 24 129–24 142, 2023.
- [75] A. Sclocchi, A. Favero, and M. Wyart, "A phase transition in diffusion models reveals the hierarchical nature of data," *arXiv preprint arXiv:2402.16991*, 2024.
- [76] Y. Huang, J. Wang, Y. Shi, B. Tang, X. Qi, and L. Zhang, "Dreamtime: An improved optimization strategy for diffusion-guided 3d generation," in *The Twelfth International Conference on Learning Representations*, 2023.
- [77] M. Kwon, J. Jeong, and Y. Uh, "Diffusion models already have a semantic latent space," in *The Eleventh International Conference on Learning Representations*, 2023.
- [78] K. Zhou, L. Wang, X. Zhang, H. P. H. Shum, F. W. B. Li, J. Li, and X. Liang, "Magr: Manifold-aligned graph regularization for continual action quality assessment," in *Proceedings of the 2024 European Conference on Computer Vision*, ser. ECCV '24. Springer, 2024.
- [79] K. Zhou, R. Cai, Y. Ma, Q. Tan, X. Wang, J. Li, H. P. Shum, F. W. Li, S. Jin, and X. Liang, "A video-based augmented reality system for human-in-the-loop muscle strength assessment of juvenile dermatomyositis," *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 5, pp. 2456–2466, 2023.
- [80] K. Zhou, H. P. H. Shum, F. W. B. Li, X. Zhang, and X. Liang, "Phi: Bridging domain shift in long-term action quality assessment via progressive hierarchical instruction," *IEEE Transactions on Image Processing*, vol. 34, pp. 3718–3732, 2025.
- [81] H. Wang, E. S. Ho, H. P. Shum, and Z. Zhu, "Spatio-temporal manifold learning for human motions via long-horizon modeling," *IEEE transactions on visualization and computer graphics*, vol. 27, no. 1, pp. 216–227, 2019.



Hubert P. H. Shum (Senior Member, IEEE) is a Professor of Visual Computing and the Director of Research in the Department of Computer Science at Durham University, specialising in modelling spatio-temporal information with responsible AI. He is also a Co-Founder and the Co-Director of Durham University Space Research Centre. Before this, he was an Associate Professor at Northumbria University and a Postdoctoral Researcher at RIKEN Japan. He received his PhD degree from the University of Edinburgh. He chaired conferences such as Pacific Graphics, BMVC and SCA. He has authored over 200 research publications in the fields of Computer Vision, Computer Graphics and AI in Healthcare, underpinned by Responsible AI designs and algorithms.



Ziyi Chang is a PhD student in the Department of Computer Science at Durham University. His research focuses on diffusion models with human motions. Specifically, his research involves diffusion models for styled skeleton-based human motion synthesis, skeleton-based human motion analysis, skeleton-based human interaction modelling. He also has interest in 3D surface reconstruction and domain adaptation. He received MSc degree from the University of Edinburgh in 2020 and BSc degree from Renmin University of China in 2019.



Kanglei Zhou is currently a Postdoctoral Researcher with the Department of Psychology and Cognitive Science, Tsinghua University. He received his Ph.D. degree in Computer Science and Engineering from Beihang University in 2025 and his B.E. degree from Henan Normal University in 2020. In 2024, he was a Visiting Student at Durham University. His research interests include human motion analysis and continual learning.



Xiaohui Liang received his Ph.D. degree in computer science and engineering from Beihang University, China. He is currently a Professor, working in the School of Computer Science and Engineering at Beihang University. His main research interests include computer graphics and animation, visualization, and virtual reality.