

Hard No-Box Adversarial Attack on Skeleton-Based Human Action Recognition with Skeleton-Motion-Informed Gradient

Zhengzhi Lu^{1,2 †} He Wang³ Ziyi Chang¹ Guoan Yang² Hubert P. H. Shum^{1 ‡}

¹Durham University, UK ²Xi'an Jiaotong University, China ³University College London, UK

lu947867114@stu.xjtu.edu.cn he_wang@ucl.ac.uk ziyi.chang@durham.ac.uk

gayang@mail.xjtu.edu.cn hubert.shum@durham.ac.uk

Abstract

Recently, methods for skeleton-based human activity recognition have been shown to be vulnerable to adversarial attacks. However, these attack methods require either the full knowledge of the victim (i.e. white-box attacks), access to training data (i.e. transfer-based attacks) or frequent model queries (i.e. black-box attacks). All their requirements are highly restrictive, raising the question of how detrimental the vulnerability is. In this paper, we show that the vulnerability indeed exists. To this end, we consider a new attack task: the attacker has no access to the victim model or the training data or labels, where we coin the term *hard no-box attack*. Specifically, we first learn a motion manifold where we define an adversarial loss to compute a new gradient for the attack, named *skeleton-motion-informed (SMI) gradient*. Our gradient contains information of the motion dynamics, which is different from existing gradient-based attack methods that compute the loss gradient assuming each dimension in the data is independent. The SMI gradient can augment many gradient-based attack methods, leading to a new family of no-box attack methods. Extensive evaluation and comparison show that our method imposes a real threat to existing classifiers. They also show that the SMI gradient improves the transferability and imperceptibility of adversarial samples in both no-box and transfer-based black-box settings.

1. Introduction

Deep learning models are vulnerable to adversarial attacks, which compute data perturbations strategically to fool trained networks. Since its discovery [31], a wide variety of models in different tasks have been attacked [1], raising severe concerns as these perturbations are imperceptible to humans. Recently, the adversarial attack in skeleton-based human activity recognition (S-HAR) has attracted at-

Information Accessible	White-Box	Queried Black-Box	Transferred Black-Box	No-Box	Hard No-Box
Model Parameters	✓	×	×	×	×
Queries of Victims	✓	✓	×	×	×
Training Samples	×	×	✓	×	×
Labels	✓	✓	✓	✓	×

Table 1. Comparisons on different settings of adversarial attacks. ✓ and × indicate if a method needs to access the corresponding information.

tention as skeletal data have been widely used in security-critical applications such as sports analysis, bio-mechanics, surveillance, and human-computer interactions [24].

Existing attacks in S-HAR are categorized into white-box and black-box approaches. White-box approaches require prior knowledge of the full details of a victim model [18, 36] while black-box approaches require a large number of queries to the victim model [7] or the access to training data and labels [36]. On the one hand, the victim model details and the training data and labels are unlikely to be available to the attacker in real-world scenarios. On the other hand, making frequent and numerous queries (e.g. tens of thousands) to the victim model is time-consuming and raises suspicion. In other words, the settings of existing S-HAR attacks are overly restrictive. A key to a successful attack is to reduce the required information of the victim model, training data and labels.

In this paper, we introduce a new threat model that requires no access to the victim model, training data or labels. We name the new threat model the *hard no-box attack*, differentiating from the recent no-box attack on images [16] that does not require access to the victim model but still needs access to the labels (i.e. soft no-box attack). Table 1 demonstrates the comparison on different settings of adversarial attacks. Among all attack settings, our hard no-box attack requires the least amount of knowledge, as it can only access the testing data without labels. Designing such an attack is nontrivial and challenging. Without access to the victim model, the attack method cannot rely on the gradient of a classification loss [12], data manipulation during

[†] This work was conducted during the visit to the Durham University.

[‡] Corresponding Author

training [25], and the feedback of a classifier [2]. The challenge is further exacerbated by the requirement of no label and training sample access, where no surrogate model can be trained to attack or estimate the data distribution.

To tackle the challenges, we propose a contrastive learning (CL) [34] solution with a manifold-based no-box adversarial loss. First, we introduce a new application of CL to learn a latent data manifold where similar samples are naturally aggregated while dissimilar samples are dispersed without the need of class labels. It provides a good description of sample similarity that facilitates generating skeletal adversarial samples. CL is suitable for hard no-box attack settings due to its ability to capture the discriminative high-level features under our restricted attack conditions. Second, we compute the perturbation to drag a data sample away from its similar neighbors in the latent space, bounded by a pre-defined budget. In particular, we design a new no-box adversarial loss to maximize each adversary’s dissimilarity with positive samples while minimizing its similarity with negative samples. The loss serves as guidance for the adversary search in our gradient-based attack scheme.

While gradient-based attack methods like I-FGSM [14] are shown to be effective on S-HAR attacks [36, 18], the gradient is computed based on the victim model and the labels, making it unsuitable for hard no-box attacks. Since adversarial samples are likely to lie in or near the motion manifold [7], ideally, we want to explore along the manifold. That is, the computation of adversarial loss gradient should consider the local motion manifold.

To this end, we propose to explicitly model motion dynamics for describing the local manifold around a given motion. Specifically, we introduce the skeleton-motion-informed (SMI) gradient that employs dynamics models (e.g. Markovian and autoregressive) to represent motion dynamics for the loss gradient computation. As a result, while existing methods generally assume each dimension in a data sample to be independent when computing the loss gradient, SMI gradient explicitly considers the dependency between frames in time. Furthermore, the SMI gradient is compatible with existing gradient-based methods including I-FGSM and MI-FGSM [8], allowing us to effectively construct a new family of no-box attack methods.

Extensive experiments show that our method generates effective adversarial samples that successfully attack various victim models across datasets (HDM05, NTU60 and NTU120). Our SMI-gradient based attacks improve the attack transferability in both no-box and transferred black-box settings, with better imperceptibility. Our contributions are summarized as:

- We confirm the S-HAR threat by introducing a new hard no-box attack and proposing the first method to generate adversarial samples without access to the victim model or training data or labels, to the best of our knowledge.

- We propose a new skeleton-motion-informed gradient that guides the adversary search along the motion manifold, explicitly considering the spatial-temporal nature of skeletal motions.
- We present a family of novel gradient-based attack strategies facilitated by the new gradient, improving the transferability and imperceptibility of adversarial samples in no-box and transferred black-box attacks.

2. Related Works

Skeleton-Based Human Action Recognition S-HAR has attracted considerable attention in many applications [24] where deep learning-based approaches have achieved state-of-the-art performance [30, 13]. Recurrent neural networks are employed to model the temporal domain of human motions [9, 29]. Furthermore, unlike images and videos, the skeleton has a graph structure, so graph convolutional networks have shown to be effective in modelling the spatial or spatial-temporal features [28, 41]. The effectiveness is generally achieved by considering the skeleton as a topological graph where the joints and bones correspond to nodes and edges [21]. Improved graph designs and network architectures are subsequently proposed [20, 44, 45, 43].

Adversarial Attacks on Skeletons Adversarial attacks were initially introduced in [31], which showcases the vulnerability of deep neural networks and has been extended to other data types. Generally, adversarial attack is a special technique of data augmentation that aims to reveal the vulnerability of a system by finding new samples, while other data augmentation techniques may have diverse purposes, e.g. training efficiency and inference performance [27]. Recently, the attack on S-HAR has received increasing attention. Wang et al. [36] analyzed the perceptibility of adversarial skeletal samples and proposed a new perceptual loss. Liu et al. [18] focused on GCN-based models and utilized generative adversarial networks to synthesize adversarial examples. Tanaka et al. [32] proposed a new lower-dimensional attack, in which only the length of bones could be perturbed. These methods all require the complete knowledge of victim models, a setting known as the white-box attack. In contrast, Diao et.al [7] introduced the first black-box S-HAR attack method, which searches motion manifolds for adversaries. Still, black-box attacks need to frequently query the victim models, which can be infeasible in real-world systems. In contrast, we consider a more practical threat setting named the hard no-box attack, where an attacker only has access to unlabeled skeletal data.

Gradient-Based Attack Strategies The core component of adversarial attacks is to generate adversarial samples [1]. Gradient-based attack methods have been widely used to introduce perturbations to a sample following the direction of the loss gradient. Goodfellow et al. [12] proposed the fast gradient sign method (FGSM) that perturbs a sam-




Figure 1. The training (left) and attack (right) processes for the hard no-box attack. The trained query encoder in the training process is used for attacks in the attack process.

ple by a single step along the loss gradient. Kurakin et al. [14] proposed I-FGSM by extending FGSM to an iterative process. Dong et al. [8] presented MI-FGSM by adding momentum to the gradient, which boosted the transferability of adversarial samples. Xie et al. [40] applied diversified augmentations to the inputs before each iteration to craft more transferable samples. While these gradient-based strategies are successful in static data and have been adapted to skeletal motions, they neglect the dependency between frames for gradient computation, which is crucial in time series. This motivates us to propose our skeleton-motion-informed attack strategies, which explicitly model the motion dynamics in the temporal domain [39, 35].

3. Hard No-Box Attack for Skeletal Data

Figure 1 shows the overview of our method. The left part is the training process where we adopt contrastive learning to obtain a latent data manifold to distinguish data samples. The attack process is shown on the right-hand side. We first design a new no-box adversarial loss in the trained latent space to guide the adversary search using samples that are dissimilar to the attacked sample. Then we propose a novel skeleton-motion-informed gradient and a new family of attack methods for generating adversarial samples.

3.1. Contrastive Learning for Motion Manifold

While the fundamental idea of adversarial attacks is to perturb a data sample to cross class boundaries, such boundaries cannot be estimated for hard no-box attacks due to the lack of labels. To estimate such boundaries without labels, we present a new application of contrastive learning (CL) [34] to aggregate similar data samples as soft class boundaries in latent space. Such boundaries enable us to adversarially perturb a sample to cross boundaries. We also train an encoder to extract discriminative high-level features for

the motion manifold in latent space. Overall, our CL constructs boundaries in latent space without labels via aggregating similar samples and segregating dissimilar samples. Our attack is guided by the dissimilarity of high-level features between samples for generating adversarial samples, instead of using class boundaries to lead the attack [2].

To incorporate both spatial and temporal information, we train an encoder (Fig. 1 left) based on adaptive graph convolutional network (AGCN) [28]. To force encoders to focus on high-level features, we apply skeleton-specific data augmentations to an input sequence S and obtain two different views S_q and S_k . Augmentations include spatial operations (e.g. pose transformations, joint jittering, etc.) and temporal operations (e.g. temporal crop and resize) [34] (detailed in supplementary material). Then, we feed S_q and S_k into the query encoder f_q and the key encoder f_k respectively for the info-noise-contrastive estimation (InfoNCE) [23]:

$$L_{contrast} = -\log \frac{\exp(f_q(S_q) \cdot f_k(S_k)/\tau)}{\exp(f_q(S_q) \cdot f_k(S_k)/\tau) + \sum_{F_n \sim N} \exp(f_q(S_q) \cdot F_n/\tau)}, \quad (1)$$

where τ is the temperature parameter, N is the dynamic queue storing the features of negative samples F_n obtained in the training process. After training, we use the query encoder f_q , which encodes the motion manifold, for attack.

3.2. Adversarial Loss for Unlabeled Skeletal Data

The adversarial loss of hard no-box attack is significantly different from most existing methods that heavily depend on labels and class boundaries. Since class labels and class boundaries are unavailable in hard no-box attacks, we utilize data samples that are dissimilar to a given sample (i.e. *negative samples*) for defining the adversarial loss. Correspondingly, samples that are similar to the given sample are considered as *positive samples*. We argue when a given sample is perturbed towards its negative samples and away from its positive samples, it tends to become an adversary. This is because the negative samples generally indicate the high-density areas of other classes in the latent space.

The hard no-box adversarial loss is designed as:

$$L_{adv} = -\log \frac{\exp[\text{Sim}(f_q(s), f_q(\tilde{s}))]}{\sum_j \exp[\text{Sim}(f_q(s), f_q(\tilde{s}_j))]}, \quad (2)$$

where Sim is the cosine similarity, s is the adversarial sample to be computed, \tilde{s} is the clean sample regarded as the positive sample, and \tilde{s}_j are the negative samples. Maximizing L_{adv} moves s away from \tilde{s} and towards \tilde{s}_j in the latent manifold. With L_{adv} , gradient-based attacks are employed.

To maximize Eq. 2, the selection of negative samples \tilde{s}_j is crucial and we design a method tailored for no-box attacks. Existing work [10] utilizes cluster-fit [42] to generate

pseudo labels for selecting negative samples during adversarial training, which is less suitable for the no-box attack as it requires another pretrained off-line encoder to obtain pseudo labels. Instead, we adapt K-means, an unsupervised method, to select negatives, removing the need of any pre-training. We discard Q clusters whose cluster centers are the closest to the input sample, mitigating the risk of misleading attacks. The remaining cluster centers are considered as the negative samples \tilde{s}_j .

4. Skeleton-Motion-Informed Gradient

Existing gradient-based attack methods treat each dimension of the data as an independent variable, i.e. raw gradient. Attacks based on raw gradients tend to drag a sample away from the data manifold [11]. With the guidance of class boundaries and a limit on the perturbation budget, the raw gradient can still find deceiving adversaries. However, this setting is infeasible in hard no-box attacks. Without class boundaries, raw gradients that point to the negative samples can drag the adversary far away from the manifold. This is because while the perturbations are towards negative samples, they are not necessarily in a direction orthogonal to the class boundary. Consequently, larger perturbations are needed to cross the boundary, leading to adversaries being far off the manifold. This creates the need to constrain the perturbation within or near the manifold, at least locally. Since the motion manifold is constrained by the motion dynamics [38], we argue that the gradient needs to explicitly capture the dynamics. Therefore, we propose a new gradient named skeleton-motion-informed (SMI) gradient, capturing the manifold information that has been largely ignored by existing methods in loss gradient computation.

4.1. Dynamics in the Gradient Structure

Given a skeletal sequence $S = [S_1, S_2, \dots, S_t]$ and the adversarial loss $J(S)$, a straightforward but effective strategy to craft adversarial perturbations is the gradient-based attack [1]. It utilizes backpropagation of the loss function $\nabla J(S)$ to iteratively change input samples S :

$$\hat{S} = S + \alpha \cdot \text{sign}(\nabla J(S)), \quad (3)$$

where α is the attack step size and \hat{S} is the adversarial samples. In skeletal motions, this attack gradient $\nabla J(S)$ consists of a set of partial derivatives over all frames:

$$\nabla J(S) = \left[\frac{\partial J(S)}{\partial S_1}, \frac{\partial J(S)}{\partial S_2}, \dots, \frac{\partial J(S)}{\partial S_t} \right]. \quad (4)$$

The partial derivative $\frac{\partial J(S)}{\partial S_t}$ assumes each frame is independent, and this is the raw gradient employed in existing methods [36]. However, human motions contain rich dynamics so that the system can be described as $S_t = f(S_{<t})$. So

far, various dynamics models have been attempted to model human motions, such as Markovian models [33, 37], autoregressive models [39], and many-to-many mapping [38], all of which can capture the dynamics at different scales in time. We explore these models to reveal the missing dynamics in the structure of the raw gradient and propose our SMI-gradients that consider motion dynamics.

Markovian Model We assume the motion dynamics can be captured by a Markovian model, i.e. $S_t = f_{d1}(S_{t-1})$. This allows us to derive the 1st-order SMI-gradient:

$$\left(\frac{\partial J(S)}{\partial S_{t-1}} \right)_{d1} = \frac{\partial J(S)}{\partial S_{t-1}} + \frac{\partial J(S)}{\partial S_t} \cdot \frac{dS_t}{dS_{t-1}}, \quad (5)$$

where $\frac{dS_t}{dS_{t-1}}$ is the temporal relationship between two consecutive frames that will be instantiated. Eq. 5 shows that the attack gradient along the motion manifold needs to consider the first-order information in the motion, e.g. velocity.

Autoregressive Model Besides the first-order dynamics, we also model the second-order dynamics by assuming $S_t = f_{d2}(S_{t-1}, S_{t-2})$, as $2nd$ -order dynamics (i.e. joint acceleration) capture the smooth temporal dynamics of skeletal motion [36]. We extend Eq. 5 as:

$$\begin{aligned} \left(\frac{\partial J(S)}{\partial S_{t-2}} \right)_{d2} &= \frac{\partial J(S)}{\partial S_{t-2}} + \frac{\partial J(S)}{\partial S_{t-1}} \cdot \frac{\partial S_{t-1}}{\partial S_{t-2}} \\ &+ \frac{\partial J(S)}{\partial S_t} \cdot \left(\frac{\partial S_t}{\partial S_{t-2}} + \frac{\partial S_t}{\partial S_{t-1}} \cdot \frac{\partial S_{t-1}}{\partial S_{t-2}} \right). \end{aligned} \quad (6)$$

While higher-order models can also be considered, there is an empirical evidence that the first three orders are the most important in skeletal motion adversarial attack [36]. Therefore, we express the SMI gradients of the whole skeletal sequence as $(\nabla J(S))_{d1} = \left[\left(\frac{\partial J(S)}{\partial S_1} \right)_{d1}, \left(\frac{\partial J(S)}{\partial S_2} \right)_{d1}, \dots, \left(\frac{\partial J(S)}{\partial S_t} \right)_{d1} \right]$ and $(\nabla J(S))_{d2} = \left[\left(\frac{\partial J(S)}{\partial S_1} \right)_{d2}, \left(\frac{\partial J(S)}{\partial S_2} \right)_{d2}, \dots, \left(\frac{\partial J(S)}{\partial S_t} \right)_{d2} \right]$.

4.2. Time-Varying Autoregressive Models for Dynamics

We employ explicit models [39] to compute the dynamics-related derivatives in SMI gradients. While implicit models [33] may also be considered, they would require another network to be trained, making them less preferable in hard no-box attacks. To realize f_{d1} and f_{d2} , we use time-varying autoregressive models (TV-AR) [3], which effectively estimates the dynamics of skeleton sequence [39] due to its ability of modelling the temporal non-stationary signals:

$$f_{d1} : S_t = A_t \cdot S_{t-1} + B_t + \gamma_t, \quad (7)$$

$$f_{d2} : S_t = C_t \cdot S_{t-1} + D_t \cdot S_{t-2} + E_t + \gamma_t, \quad (8)$$

where Eq. 7 and Eq. 8 are denoted as TV-AR(1) and TV-AR(2) respectively. The model parameters $\beta_t^1 = [A_t, B_t]$ and $\beta_t^2 = [C_t, D_t, E_t]$ are all time-varying parameters and determined by data-fitting. γ_t is a time-dependent white noise representing the dynamics of stochasticity.

Using Eq. 7 to compute $\frac{\partial S_t}{\partial S_{t-1}}$, Eq. 5 becomes:

$$\left(\frac{\partial J(S)}{\partial S_{t-1}} \right)_{d1} = \frac{\partial J(S)}{\partial S_{t-1}} + \frac{\partial J(S)}{\partial S_t} \cdot A_t. \quad (9)$$

Similarly, using Eq. 8, we can compute $C_t = \frac{\partial S_t}{\partial S_{t-1}}$ and $D_t = \frac{\partial S_t}{\partial S_{t-2}}$. For $\frac{\partial S_{t-1}}{\partial S_{t-2}}$, we use $S_{t-1} = C_{t-1} \cdot S_{t-2} + D_{t-1} \cdot S_{t-3} + E_{t-1} + \gamma_t$ to compute it: $C_{t-1} = \frac{\partial S_{t-1}}{\partial S_{t-2}}$. Then, Eq. 6 becomes:

$$\left(\frac{\partial J(S)}{\partial S_{t-2}} \right)_{d2} = \frac{\partial J(S)}{\partial S_{t-2}} + \frac{\partial J(S)}{\partial S_{t-1}} \cdot C_{t-1} + \frac{\partial J(S)}{\partial S_t} \cdot (D_t + C_t \cdot C_{t-1}). \quad (10)$$

5. Skeleton-Motion-Informed Attack

We construct new gradient-based attack methods based on our novel SMI gradient. Due to its compatibility, our proposed gradient can be integrated with most existing gradient-based methods. We select I-FGSM and MI-FGSM, which have been proven for their efficiency on S-HAR attacks [36, 18]. We augment them with first and second-order SMI-gradients, leading to four new attack methods.

Fast Gradient Sign Methods (FGSM) FGSM [12] is a single-step attack method that generates the adversarial samples $\hat{S} = S + r$ by maximizing the adversarial loss function $J(S)$, where r denotes an adversarial perturbation that is constrained within a budget $\|r\|_p < \epsilon$, where $\|\cdot\|_p$ denotes the l_p -norm. One variant of FGSM is the Iterative Fast Gradient Sign Method (I-FGSM) [14], which extends FGSM to an iterative process:

$$\hat{S}^{i+1} = \hat{S}^i + \alpha \cdot \text{sign}(\nabla_s J(S)), \quad (11)$$

where α is the attack step size and i means iteration. Another variant is MI-FGSM [8], which considers the momentum of the attack to avoid local maxima:

$$g^{i+1} = \mu \cdot g^i + \frac{\nabla_s J(S)}{\|\nabla_s J(S)\|_1} \quad (12)$$

$$\hat{S}^{i+1} = \hat{S}^i + \alpha \cdot \text{sign}(g^{i+1}), \quad (13)$$

where μ is the momentum decay factor and g_i is the gradient in iteration i .

SMI-gradient Based Attacks We replace the original gradient $\nabla J(S)$ in I-FGSM and MI-FGSM with our SMI gradient $(\nabla J(S))_{d1}$ or $(\nabla J(S))_{d2}$. This creates four new

Algorithm 1 S₁I-FGSM and S₂I-FGSM

Input: An encoder k with a loss function J ; a skeletal sequence sample S ; the size of attack step α ; iterations I ; the budget of perturbation ϵ .

Output: An adversarial example \hat{S} with $\|\hat{S} - S\|_p < \epsilon$.

- 1: Initialization: $\hat{S}^0 = S$;
- 2: Fitting S with TV-AR model to obtain the time-varying parameters β_t ;
- 3: **for** $i = 0$ to $I - 1$ **do**
- 4: Input \hat{S}^i to k ;
- 5: Obtain the raw gradient $\nabla J(\hat{S}^i)$ on J ;
- 6: Calculate the SMI gradient $(\nabla J(\hat{S}^i))_{d1}$ with Eq. 9, or $(\nabla J(\hat{S}^i))_{d2}$ with Eq. 10, using β_t and $\nabla J(\hat{S}^i)$;
- 7: Update \hat{S}^{i+1} by applying the sign gradient as:

$$\hat{S}^{i+1} = \hat{S}^i + \alpha \cdot \text{sign}(\nabla J(\hat{S}^i)_{d1}), \text{ or} \quad (14)$$

$$\hat{S}^{i+1} = \hat{S}^i + \alpha \cdot \text{sign}(\nabla J(\hat{S}^i)_{d2}).$$

- 8: **end for**
 - 9: **return** $\hat{S} = \hat{S}^I$
-

dynamic attack strategies: first-order SMI I-FGSM (S₁I-FGSM), second-order SMI I-FGSM (S₂I-FGSM), first-order SMI MI-FGSM (S₁MI-FGSM), and second-order SMI MI-FGSM (S₂MI-FGSM). The processes of SI-FGSM are shown in Algorithm 1. The algorithm of SMI-FGSM can be found in the supplementary material.

6. Experiments

We refer the readers to the supplementary material for extra experimental results.

Datasets We select three widely used skeletal datasets: HDM05 [22] (2,337 sequences of 130 classes performed by 5 actors), NTU60 [26] (56,880 sequences of 60 classes), and NTU120 [19] (114,480 sequences of 120 classes, an extended version of NTU60, one of the largest datasets in the field). We pre-process HDM05 following [9] and both NTU datasets following [28]. The different skeletons are mapped to a standard 25-joint structure as in [38]. For our hard no-box attack, we only use the testing data and do not use the training data, the training labels and the testing labels during attacks.

Target Models We choose multiple state-of-the-art models as victims: ST-GCN [41], 2s-AGCN [28], AS-GCN [15], SGN [44] and MS-G3D [20]. They are trained using the official implementations and following the training protocols. For 2s-AGCN, we attack both the single joint stream model (js-AGCN) and the two-stream model.

Implementation Details We pre-train the CL encoder f_q following [34]. The unsupervised network is trained with a temperature value $\tau = 0.07$ and SGD optimizer for 450 epochs. The learning rate is set to 0.01 with a weight decay of 0.0001. Due to the limitation of hard no-box settings, the attacker cannot access the training samples during the whole process. Therefore, the encoder is trained on the test-

	Victim models	Self-sup Attacker	AGCN Attacker	No-box I-FGSM	No-box S ₁ I-FGSM	No-box S ₂ I-FGSM	No-box MI-FGSM	No-box S ₁ MI-FGSM	No-box S ₂ MI-FGSM
$\epsilon = 0.01$	js-AGCN	11.21%	—	26.05%	28.09%	30.87%	30.68%	34.75%	36.58%
	2s-AGCN	5.36%	—	13.94%	15.04%	16.45%	16.47%	18.23%	19.31%
	ST-GCN	3.57%	12.93%	9.55%	9.86%	9.96%	11.11%	11.36%	11.56%
	MS-G3D	8.39%	35.23%	10.57%	11.14%	11.69%	11.76%	12.85%	14.11%
	SGN	21.25%	26.03%	34.09%	34.46%	35.23%	38.49%	38.75%	38.81%
	ASGCN	5.69%	20.87%	13.92%	14.85%	14.67%	15.95%	16.82%	17.75%
$\epsilon = 0.008$	js-AGCN	10.12%	—	22.84%	24.36%	26.46%	25.70%	29.64%	30.88%
	2s-AGCN	5.04%	—	11.36%	12.19%	13.03%	12.30%	13.85%	15.56%
	ST-GCN	3.26%	10.33%	7.87%	7.99%	8.04%	8.84%	9.07%	9.19%
	MS-G3D	5.19%	31.28%	9.18%	9.51%	9.99%	10.01%	10.26%	12.98%
	SGN	20.95%	23.30%	29.61%	30.20%	30.74%	33.32%	33.63%	34.54%
	ASGCN	5.54%	18.77%	11.42%	12.05%	12.29%	12.77%	13.69%	14.37%
$\epsilon = 0.006$	js-AGCN	7.19%	—	19.70%	20.32%	21.23%	20.04%	22.76%	24.33%
	2s-AGCN	3.70%	—	7.93%	9.80%	10.56%	9.80%	11.26%	12.66%
	ST-GCN	2.46%	7.88%	5.65%	5.85%	5.97%	6.25%	6.54%	6.66%
	MS-G3D	4.46%	23.15%	7.30%	7.84%	7.61%	7.94%	8.05%	8.39%
	SGN	16.76%	19.93%	23.92%	24.57%	25.75%	26.74%	27.04%	27.64%
	ASGCN	3.94%	14.30%	8.79%	9.23%	9.07%	9.71%	10.29%	10.90%

Table 2. The fooling rate of different methods on the target models in NTU60, where ϵ is the perturbation budget.

	Victim models	Self-sup Attacker	AGCN Attacker	No-box I-FGSM	No-box S ₁ I-FGSM	No-box S ₂ I-FGSM	No-box MI-FGSM	No-box S ₁ MI-FGSM	No-box S ₂ MI-FGSM
$\epsilon = 0.01$	js-AGCN	9.97%	—	23.92%	24.64%	25.79%	27.26%	27.93%	29.07%
	2s-AGCN	6.38%	—	20.09%	21.11%	22.06%	23.51%	24.63%	24.96%
	ST-GCN	12.18%	23.84%	23.53%	24.73%	25.77%	26.95%	27.31%	28.76%
	MS-G3D	10.63%	24.63%	19.07%	20.11%	20.20%	21.57%	21.95%	22.73%
	SGN	19.65%	37.90%	31.43%	32.64%	33.75%	38.47%	37.96%	38.85%
	ASGCN	7.29%	24.15%	18.03%	19.29%	20.37%	19.88%	20.22%	21.60%
$\epsilon = 0.008$	js-AGCN	9.04%	—	21.58%	21.77%	22.47%	24.28%	24.62%	25.74%
	2s-AGCN	5.79%	—	15.71%	15.23%	15.94%	18.68%	18.76%	19.57%
	ST-GCN	11.23%	21.35%	20.74%	21.51%	22.22%	23.52%	24.65%	24.87%
	MS-G3D	9.98%	21.73%	17.06%	17.74%	17.92%	19.53%	19.29%	20.33%
	SGN	17.59%	35.88%	27.38%	28.17%	28.65%	32.43%	33.06%	32.55%
	ASGCN	6.87%	21.87%	15.73%	16.93%	17.62%	17.41%	17.75%	18.70%
$\epsilon = 0.006$	js-AGCN	8.23%	—	18.51%	18.88%	18.79%	20.74%	21.88%	22.06%
	2s-AGCN	5.09%	—	7.93%	9.80%	10.56%	9.80%	11.26%	12.66%
	ST-GCN	9.54%	17.51%	17.60%	17.97%	18.67%	18.90%	19.91%	20.34%
	MS-G3D	9.42%	18.07%	15.06%	15.09%	15.20%	16.86%	17.25%	17.56%
	SGN	16.82%	34.10%	22.39%	22.65%	22.76%	25.43%	25.79%	25.92%
	ASGCN	5.39%	19.02%	12.90%	13.45%	14.97%	14.08%	15.26%	16.33%

Table 3. The fooling rate of different methods on the target models in NTU120, where ϵ is the perturbation budget.

ing set. We adopt the l_∞ norm for the perturbation budget ϵ . For clusters of negative samples, the number of clusters is 120, and the number of deleted centers Q is 10.

Evaluation Metrics We employ the fooling rate as a major metric. It is defined as the percentage of data samples whose predicted labels changed after adversarial attacks [18]. Besides, inspired by [36], we define a perceptual deviation indicator to evaluate the imperceptibility of adversarial skeletal samples:

$$\Delta p = \frac{1}{MT} \sum_{n=0}^M \|S - \hat{S}\|_2 + \frac{1}{MT} \sum_{n=0}^M \|B - \hat{B}\|_2 + \frac{1}{MTL} \sum_{n=0}^M \|\ddot{S} - \ddot{\hat{S}}\|_2 \quad (15)$$

where M is the number of adversarial samples, T is the total number of frames, and L is the number of joints in the skeleton. The three terms evaluate the deviations of joint position, bone-length, and acceleration, respectively. A smaller perceptual deviation indicates better imperceptibility.

6.1. Hard No-Box Attack

Baselines We establish two baselines using transfer-based attacks based on a self-supervised and a supervised classifier. This is because our method is the first hard no-box attack and there is no other similar method. The transfer-based attack is the closest setting to ours as they do not require access to the victim model. The first baseline (Self-sup Attacker) is a self-supervised surrogate model with a linear layer appended to our CL encoder. We freeze our CL encoder after the manifold training and then train the linear layer supervisedly [17]. In this way, the need to access labels is only for training the linear layer, not the CL encoder, which is not strictly hard no-box but closer than existing methods. The second baseline is a standard transfer-based attack where we use js-AGCN as the surrogate model and SMART [36] as the white-box attacker (AGCN Attacker). Unlike our method, both baselines still require access to the training data and the training labels during attacks.

For hard no-box attack, different attack strategies are compared including I-FGSM, MI-FGSM, S₁I-FGSM, S₂I-FGSM, S₁MI-FGSM, and S₂MI-FGSM. We run all the attackers for 400 iterations and report the attack perfor-

	Victim models	Self-sup Attacker	AGCN Attacker	No-box I-FGSM	No-box S ₁ I-FGSM	No-box S ₂ I-FGSM	No-box MI-FGSM	No-box S ₁ MI-FGSM	No-box S ₂ MI-FGSM
$\epsilon = 0.01$	js-AGCN	10.12%	—	10.61%	10.98%	11.55%	13.45%	14.20%	14.20%
	2s-AGCN	1.89%	—	5.30%	5.68%	5.68%	6.44%	6.44%	6.82%
	ST-GCN	6.44%	24.26%	9.47%	9.85%	9.47%	8.90%	11.17%	11.74%
	MS-G3D	4.17%	86.95%	41.67%	44.51%	43.94%	53.03%	56.06%	53.41%
	SGN	1.89%	3.98%	2.46%	2.84%	3.03%	3.03%	3.40%	3.59%
	ASGCN	1.89%	27.57%	3.22%	4.36%	3.22%	4.36%	5.68%	5.68%
$\epsilon = 0.008$	js-AGCN	9.46%	—	8.14%	9.09%	9.09%	10.61%	10.98%	11.55%
	2s-AGCN	1.70%	—	4.36%	4.92%	4.73%	4.55%	5.30%	5.87%
	ST-GCN	3.79%	21.69%	8.33%	8.52%	8.52%	8.33%	8.90%	8.33%
	MS-G3D	3.60%	82.17%	25.76%	30.68%	29.17%	36.36%	39.39%	39.02%
	SGN	1.51%	3.60%	0.57%	1.51%	2.08%	2.27%	2.27%	2.46%
	ASGCN	1.70%	22.43%	2.46%	2.84%	2.84%	2.27%	2.84%	3.03%
$\epsilon = 0.006$	js-AGCN	7.19%	—	5.11%	7.01%	6.44%	6.06%	6.63%	7.77%
	2s-AGCN	1.33%	—	2.08%	2.08%	2.27%	3.79%	3.98%	4.17%
	ST-GCN	3.40%	19.85%	6.44%	7.01%	6.63%	7.58%	8.33%	7.77%
	MS-G3D	2.84%	72.43%	11.55%	12.31%	11.92%	15.34%	18.94%	17.99%
	SGN	0.57%	3.40%	0.13%	0.94%	0.57%	1.13%	1.32%	1.51%
	ASGCN	1.52%	17.82%	0.38%	0.76%	1.89%	0.57%	2.27%	2.27%

Table 4. The fooling rate of different methods on the target models in HDM05, where ϵ is the perturbation budget.




Figure 2. Visual comparisons between attack strategies in no-box attacks ($\epsilon=0.006$) with key visual differences highlighted.

mance under different perturbation budgets ϵ in Table 2 for NTU60, Table 3 for NTU120, and Table 4 for HDM05. We omit the results on js-AGCN and 2s-AGCN under AGCN Attacker because AGCN is the surrogate model.

Table 2-4 show that the hard no-box attack poses real threats to a range of S-HAR classifiers. In general, the fooling rate of hard no-box attacks is higher than Self-sup Attackers. This is surprising as the Attacker utilizes training data and training labels while our method does not. Looking further, AGCN Attacker is much better than Self-sup Attacker and the major difference is their feature extraction (i.e. one is a CL encoder and one is a graph network). This shows transfer-based attack heavily relies on the feature extraction ability of the surrogate model that cannot bypass access to the training data and the training labels. Furthermore, when compared with AGCN Attacker, our method achieves similar fooling rates, varying across different victims and datasets. Given that AGCN Attacker requires access to training data, training labels and testing labels, we argue hard no-box attacker achieves superior results and provides a more realistic setting.

In addition, among various no-box attack strategies, S₂MI-FGSM performs the best and often by big margins. All the SMI gradient-based methods generate stronger adversaries compared with baselines I-FGSM and MI-FGSM. The 2nd-order SMI attack method usually outperforms the corresponding 1st-order version. Last, we notice a variance in fooling rate across different victims and datasets. For in-

stance, the multi-stream model (2s-AGCN) significantly enhances the robustness compared to the single-stream model (js-AGCN); the fooling rate of all methods drops by nearly half. This may be because the multi-stream model can ensemble features from different modalities, which improves the robustness. In general, it is still an open question why fooling rate can vary across victims and datasets and we leave the theoretical analysis for future work.

Surrogate Model	Victims	I-FGSM	MI-FGSM	S ₂ I-FGSM (Ours)	S ₂ MI-FGSM (Ours)
2s-AGCN	STGCN	2.10%	2.10%	3.00%	3.01%
	MS-G3D	2.58%	2.59%	2.90%	2.97%
ST-GCN	2sAGCN	2.20%	2.34%	2.44%	2.64%
	MS-G3D	2.00%	2.10%	2.63%	2.92%
MS-G3D	STGCN	1.71%	1.69%	2.65%	2.67%
	2sAGCN	1.76%	1.79%	2.03%	2.07%

Table 5. The fooling rate of different attack strategies in transferred SMART attacks, where attack budgets $\epsilon = 0.01$.

6.2. Transfer-Based Black-Box Attack

SMI gradient not only improves the transferability in hard no-box attacks, but also enhances other gradient-based skeletal attacks. Here, we employ SMART [36], a white-box attacker, as a baseline to compare different attack strategies. In the original SMART settings, I-FGSM is adopted to generate adversarial samples. We replace it with S₂I-FGSM, MI-FGSM, and S₂MI-FGSM to make a comparison. As all strategies achieve similar fooling rates in white-box attacks, we mainly focus on their transferability using

different surrogate models. We utilize SMART to attack 2s-AGCN, ST-GCN, and MS-G3D on NTU60 and transfer the obtained samples to other victim networks.

The performance of the transferred black-box attack is shown in Table 5. S_2 MI-FGSM gives the best performance in transfer-based black-box attacks. S_2 I-FGSM also improves the transferability of adversarial samples compared with baselines. In contrast, MI-FGSM, which succeeds in the image transfer-based attack, struggles in skeletal data. Its performance declines to 1.69% when it attacks STGCN via MS-G3D. Overall, the success rate is low in Table 5 because SMART is sensitive to the chosen surrogate [36]. Transfer-based attacks are proven to suffer from lower fooling rates in S-HAR attack [36] and improving the transferability is still an open problem. Table 5 aims to show our SMI gradient can improve it by incorporating motion dynamics into the attack gradient, compared with other alternative gradients. We will include in future work how to further explore this dynamics for better attack transfer.

6.3. Perceptual Analysis

A key feature of SMI gradient-based attacks is the improvement in the perceptual quality of the adversarial samples due to the consideration of the motion manifold. To verify this, we employ quantitative comparison and qualitative visual analysis on the no-box adversarial samples under various strategies. We compare the perceptual quality Δp on NTU60 in Table 6. We find that S_2 I-FGSM achieves the best imperceptibility and obtains a massive improvement compared with I-FGSM. In contrast, MI-FGSM’s deviation is twice that of S_2 I-FGSM. Although S_2 MI-FGSM does not achieve the best visual performance, it is still slightly better than I-FGSM and achieves a better trade-off between fooling rate and perceptual quality. This is understandable because our method considers dynamics that help to generate more on-manifold adversarial samples. Moreover, the 1st-order SMI attacks outperform baselines but cannot compete with 2nd-order SMI attacks. This demonstrates the importance of considering acceleration in the skeletal attack.

We show the visual comparison of poses under various attack strategies in no-box attacks in Figure 2. The spinal joints demonstrate the most obvious differences. S_2 I-FGSM outperforms the other attack methods and gets the most natural poses, whereas I-FGSM has slight but noticeable joint displacements in the neck and head. S_2 MI-FGSM performs better than its baseline, MI-FGSM, which shows zig-zag bending in the frame $t_0 + 2$. The samples produced by MI-FGSM have the worst imperceptibility, where we can easily find the unnatural jittery movements.

We also evaluate perceptual quality Δp on SMART adversarial samples obtained with different gradients. Results conducted on the NTU60 dataset are shown in Table 7. S_2 I-FGSM reaches the best perceptual performance compared

with all attack strategies. MI-FGSM slightly declines the imperceptibility. S_2 MI-FGSM performs slightly worse in MS-G3D and ST-GCN. The reason is mainly that S_2 MI-FGSM takes more iterations in the white-box attack, leading to late stopping and slightly worse visual performance.

Strategies	$\epsilon = 0.01$	$\epsilon = 0.008$	$\epsilon = 0.006$
I-FGSM	95.87	68.56	43.13
MI-FGSM	131.77	90.44	54.61
S_1 I-FGSM	84.63	60.77	38.15
S_1 MI-FGSM	114.51	78.06	47.04
S_2 I-FGSM	65.60	45.67	27.67
S_2 MI-FGSM	90.93	62.04	37.46

Table 6. The perceptual deviation of different attack strategies in no-box attacks with different budgets ϵ .

Victims	I-FGSM	MI-FGSM	S_2 I-FGSM (Ours)	S_2 MI-FGSM (Ours)
2s-AGCN	1.52	1.62	1.25	1.49
MS-G3D	2.39	2.46	1.69	3.02
ST-GCN	1.13	1.18	1.10	1.47

Table 7. The perceptual deviation of different attack strategies in SMART for different victim models.

7. Conclusions and Discussions

In this paper, we have verified potential threats to S-HAR solutions. A new setting is proposed: the hard no-box attack on skeletal motions without access to the victim model, the training samples or the labels. We validate our setting by proposing the first pipeline for hard no-box attacks. Moreover, as far as we know, we are the first to explore motion dynamics in the adversarial gradient computation, leading to a new SMI gradient compatible with existing gradient-based attacks. By extensive evaluation and comparison, our method has been proven to be threatening and imperceptible, relying on the least prior knowledge.

The SMI gradient also improves the transferability of transferred black-box attacks. Nonetheless, boosting the transferability is still an open problem [36] and we will explore this further with our SMI gradient in the future. We will explore other models (e.g., diffusion models [5]) and other time-series data (e.g. stock price, videos) for the proposed attack. Also, our SMI gradient describes dynamics and may be beneficial for motion synthesis [4].

We call for attention to intensify the S-HAR robustness by considering defences against our hard no-box attack. We validate randomized smoothing [6] as a potential defence method in supplementary materials. Due to the least prior knowledge requirement, security risks posed by our attack can be reduced with such defenses. Otherwise, our attacks become a significantly threat to S-HAR.

Acknowledgments

This research is supported in part by National Natural Science Foundation of China (ref: 61673314, Yang), EP-SRC (ref: EP/X031012/1, NorthFutures, Shum) and EU Horizon 2020 (ref: 899739, CrowdDNA, Wang).

References

- [1] Naveed Akhtar, Ajmal Mian, Navid Kardan, and Mubarak Shah. Advances in adversarial attacks and defenses in computer vision: A survey. *IEEE Access*, 9:155161–155196, 2021.
- [2] Wieland Brendel, Jonas Rauber, and Matthias Bethge. Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. In *International Conference on Learning Representations*, 2018.
- [3] Laura F Bringmann, Ellen L Hamaker, Daniel E Vigo, André Aubert, Denny Borsboom, and Francis Tuerlinckx. Changing dynamics: Time-varying autoregressive models using generalized additive modeling. *Psychological methods*, 22(3):409, 2017.
- [4] Ziyi Chang, Edmund JC Findlay, Haozheng Zhang, and Hubert PH Shum. Unifying human motion synthesis and style transfer with denoising diffusion probabilistic models. *arXiv preprint arXiv:2212.08526*, 2022.
- [5] Ziyi Chang, George A Koulteris, and Hubert PH Shum. On the design fundamentals of diffusion models: A survey. *arXiv preprint arXiv:2306.04542*, 2023.
- [6] Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *international conference on machine learning*, pages 1310–1320. PMLR, 2019.
- [7] Yunfeng Diao, Tianjia Shao, Yong-Liang Yang, Kun Zhou, and He Wang. Basar: Black-box attack on skeletal action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7597–7607, 2021.
- [8] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9185–9193, 2018.
- [9] Yong Du, Wei Wang, and Liang Wang. Hierarchical recurrent neural network for skeleton based action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1110–1118, 2015.
- [10] Lijie Fan, Sijia Liu, Pin-Yu Chen, Gaoyuan Zhang, and Chuang Gan. When does contrastive learning preserve adversarial robustness from pretraining to finetuning? *Advances in Neural Information Processing Systems*, 34:21480–21492, 2021.
- [11] Reuben Feinman, Ryan R Curtin, Saurabh Shintre, and Andrew B Gardner. Detecting adversarial samples from artifacts. *arXiv preprint arXiv:1703.00410*, 2017.
- [12] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [13] Qihong Ke, Mohammed Bennamoun, Senjian An, Ferdous Sohel, and Farid Boussaid. A new representation of skeleton sequences for 3d action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3288–3297, 2017.
- [14] Alexey Kurakin, Ian J Goodfellow, and Samy Bengio. *Adversarial examples in the physical world*, pages 99–112. Chapman and Hall/CRC, 2018.
- [15] Maosen Li, Siheng Chen, Xu Chen, Ya Zhang, Yanfeng Wang, and Qi Tian. Actional-structural graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3595–3603, 2019.
- [16] Qizhang Li, Yiwen Guo, and Hao Chen. Practical no-box adversarial attacks against dnns. *Advances in Neural Information Processing Systems*, 33:12849–12860, 2020.
- [17] Lilang Lin, Sijie Song, Wenhan Yang, and Jiaying Liu. Ms2l: Multi-task self-supervised learning for skeleton based action recognition. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2490–2498, 2019.
- [18] Jian Liu, Naveed Akhtar, and Ajmal Mian. Adversarial attack on skeleton-based human action recognition. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [19] Jun Liu, Amir Shahroudy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex C Kot. Ntu rgb+ d 120: A large-scale benchmark for 3d human activity understanding. *IEEE transactions on pattern analysis and machine intelligence*, 42(10):2684–2701, 2019.
- [20] Ziyu Liu, Hongwen Zhang, Zhenghao Chen, Zhiyong Wang, and Wanli Ouyang. Disentangling and unifying graph convolutions for skeleton-based action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 143–152, 2020.
- [21] Qianhui Men, Edmond S. L. Ho, Hubert P. H. Shum, and Howard Leung. A quadruple diffusion convolutional recurrent network for human motion prediction. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(9):3417–3432, 2021.
- [22] Meinard Müller, Tido Röder, Michael Clausen, Bernhard Eberhardt, Björn Krüger, and Andreas Weber. Mocap database hdm05. *Institut für Informatik II, Universität Bonn*, 2(7), 2007.
- [23] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [24] Bin Ren, Mengyuan Liu, Runwei Ding, and Hong Liu. A survey on 3d skeleton-based action recognition using learning method. *arXiv preprint arXiv:2002.05907*, 2020.
- [25] Aniruddha Saha, Akshayvarun Subramanya, and Hamed Pirsiavash. Hidden trigger backdoor attacks. In *Proceedings of the AAAI conference on artificial intelligence*, pages 11957–11965, 2020.
- [26] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang 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*, pages 1010–1019, 2016.
- [27] Divya Shanmugam, Davis Blalock, Guha Balakrishnan, and John Guttag. Better aggregation in test-time augmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1214–1223, October 2021.

- [28] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12026–12035, 2019.
- [29] Sijie Song, Cuiling Lan, Junliang Xing, Wenjun Zeng, and Jiaying Liu. An end-to-end spatio-temporal attention model for human action recognition from skeleton data. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017.
- [30] Tae Soo Kim and Austin Reiter. Interpretable 3d human action analysis with temporal convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 20–28, 2017.
- [31] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [32] Nariki Tanaka, Hiroshi Kera, and Kazuhiko Kawamoto. Adversarial bone length attack on action recognition. *arXiv preprint arXiv:2109.05830*, 2021.
- [33] Xiangjun Tang, He Wang, Bo Hu, Xu Gong, Ruifan Yi, Qilong Kou, and Xiaogang Jin. Real-time controllable motion transition for characters. *ACM Trans. Graph.*, 41(4), jul 2022.
- [34] Fida Mohammad Thoker, Hazel Doughty, and Cees GM Snoek. Skeleton-contrastive 3d action representation learning. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 1655–1663, 2021.
- [35] He Wang, Yunfeng Diao, Zichang Tan, and Guodong Guo. Defending black-box skeleton-based human activity classifiers. *arXiv preprint arxiv.2203.04713*, 2022.
- [36] He Wang, Feixiang He, Zhexi Peng, Tianjia Shao, Yongliang Yang, Kun Zhou, and David 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*, pages 14656–14665, 2021.
- [37] He Wang, Edmond SL Ho, and Taku Komura. An energy-driven motion planning method for two distant postures. *IEEE transactions on visualization and computer graphics*, 21(1):18–30, 2015.
- [38] He Wang, Edmond SL Ho, Hubert PH Shum, and Zhanxing Zhu. Spatio-temporal manifold learning for human motions via long-horizon modeling. *IEEE transactions on visualization and computer graphics*, 27(1):216–227, 2019.
- [39] Shihong Xia, Congyi Wang, Jinxiang Chai, and Jessica Hodgins. Realtime style transfer for unlabeled heterogeneous human motion. *ACM Transactions on Graphics (TOG)*, 34(4):1–10, 2015.
- [40] Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan L Yuille. Improving transferability of adversarial examples with input diversity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2730–2739, 2019.
- [41] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Thirty-second AAAI conference on artificial intelligence*, 2018.
- [42] Xueting Yan, Ishan Misra, Abhinav Gupta, Deepti Ghadiyaram, and Dhruv Mahajan. Clusterfit: Improving generalization of visual representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6509–6518, 2020.
- [43] Jiahang Zhang, Lilang Lin, and Jiaying Liu. Hierarchical consistent contrastive learning for skeleton-based action recognition with growing augmentations. *arXiv preprint arXiv:2211.13466*, 2022.
- [44] Pengfei Zhang, Cuiling Lan, Wenjun Zeng, Junliang Xing, Jianru Xue, and Nanning Zheng. Semantics-guided neural networks for efficient skeleton-based human action recognition. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1112–1121, 2020.
- [45] Yujie Zhou, Haodong Duan, Anyi Rao, Bing Su, and Jiaqi Wang. Self-supervised action representation learning from partial spatio-temporal skeleton sequences. *arXiv preprint arXiv:2302.09018*, 2023.



To cite this article: Lu, Z., Wang, H., Chang, Z., Yang, G., & Shum, H. P. (in press). Hard No-Box Adversarial Attack on Skeleton-Based Human Action Recognition with Skeleton-Motion-Informed Gradient

Durham Research Online URL:

<https://durham-repository.worktribe.com/output/1715085>

Copyright statement: © 2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.