



UNIVERSITY OF LEEDS

This is a repository copy of *BASAR:Black-box Attack on Skeletal Action Recognition*.

White Rose Research Online URL for this paper:

<http://eprints.whiterose.ac.uk/171782/>

Version: Accepted Version

Proceedings Paper:

Yunfeng, D, Tianjia, S, Yongliang, Y et al. (2 more authors) (Accepted: 2021)

BASAR:Black-box Attack on Skeletal Action Recognition. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition. The

Conference on Computer Vision and Pattern Recognition, 14-19 Jun 2021, Online. IEEE . (In Press)

Protected by copyright. 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.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

BASAR:Black-box Attack on Skeletal Action Recognition

Yunfeng Diao^{1,2*}, Tianjia Shao^{3†}, Yong-Liang Yang⁴, Kun Zhou³, He Wang^{1‡}

¹University of Leeds, UK ²Southwest Jiaotong University, China

³State Key Lab of CAD&CG, Zhejiang University, China ⁴ University of Bath, UK

dyf@my.swjtu.edu.cn, tjshao@zju.edu.cn, y.yang@cs.bath.ac.uk, kunzhou@zju.edu.cn, h.e.wang@leeds.ac.uk

Abstract

Skeletal motion plays a vital role in human activity recognition as either an independent data source or a complement [33]. The robustness of skeleton-based activity recognizers has been questioned recently [29, 50], which shows that they are vulnerable to adversarial attacks when the full-knowledge of the recognizer is accessible to the attacker. However, this white-box requirement is overly restrictive in most scenarios and the attack is not truly threatening. In this paper, we show that such threats do exist under black-box settings too. To this end, we propose the first black-box adversarial attack method BASAR. Through BASAR, we show that adversarial attack is not only truly a threat but also can be extremely deceitful, because on-manifold adversarial samples are rather common in skeletal motions, in contrast to the common belief that adversarial samples only exist off-manifold [18]. Through exhaustive evaluation and comparison, we show that BASAR can deliver successful attacks across models, data, and attack modes. Through harsh perceptual studies, we show that it achieves effective yet imperceptible attacks. By analyzing the attack on different activity recognizers, BASAR helps identify the potential causes of their vulnerability and provides insights on what classifiers are likely to be more robust against attack.

1. Introduction

Deep learning methods have been proven to be vulnerable to carefully devised data perturbations since first identified in [45]. This causes major concerns especially in safety and security [1], as the perturbations are *imperceptible* to humans but *destructive* to machine intelligence. Consequently, how to detect and defend attacks has also been investigated [6]. While the attack on static data (e.g. images, texts, graphs) has been widely studied, the attack on time-

series data has only been recently explored [23, 16]. In this paper, we look into a specific yet important type of time series data, skeletal motions, under adversarial attack.

Skeletal motion is crucial in activity recognition as it increases the robustness by mitigating issues such as lighting, occlusion, view angles, etc [37]. Therefore, the vulnerability of skeleton-based classifiers under adversarial attack has recently drawn attention [29, 61, 50]. Albeit identifying a key issue that needs to be addressed, their methods are essentially *white-box* methods. The attempt on black-box attack is via surrogate models, i.e. attack a classifier in a white-box manner then use the results to attack the target classifier. While white-box attack requires the full knowledge of the attacked model which is unlikely to be available in real-world scenarios, black-box attack via surrogate models cannot guarantee success due to its heavy dependence on the choice of the surrogate model [51]. In this paper, we propose BASAR, the *very first* black-box attack method on skeletal action recognition to our best knowledge.

A skeletal motion has unique features that distinguish itself from other data under adversarial attack. First, a skeleton usually has less than 100 Degrees of freedom (Dofs), much smaller than previously attacked data such as images/meshes. This low dimensionality leads to low-redundancy [46], restricting possible attacks within small subspaces. Second, *imperceptibility* is a prerequisite for any successful attack, but its evaluation on skeletal motions is under-explored. Different from the attack where visual imperceptibility has high correlations mainly with the perturbation magnitude (e.g. images), a skeletal motion has dynamics that are well-recognized by human perceptual systems. More specifically, any sparse attack, e.g. on individual joints or individual frames, albeit small would break the dynamics and therefore be easily perceptible. In contrast, coordinated attacks on all joints and frames can provide better imperceptibility even when perturbations are relatively large [50]. As a result, using the perturbation magnitude alone (as in most existing methods) is not a reliable metric for skeletal motion. Last but not least, prior methods mainly assume that adversarial samples are off the data manifold

*The research was conducted during the visit to the University of Leeds.

† Corresponding author

‡ <https://youtu.be/PjWgnAkV8g>

[18]. As we will show, skeletal motion is one real-world example where on-manifold adversarial samples not only exist but are rather common. This raises a serious concern for human activity recognition solutions as these on-manifold adversarial samples are *implementable*.

Given a motion \mathbf{x} with class label $C_{\mathbf{x}}$, BASAR aims to find \mathbf{x}' that is close to \mathbf{x} (measured by some distance function) and can fool the black-boxed classifier such that $C_{\mathbf{x}'} \neq C_{\mathbf{x}}$. BASAR formulates it as a constrained optimization problem, aiming to find \mathbf{x}' that is just outside $C_{\mathbf{x}}$ while still on the data manifold. The optimization is highly non-linear due to the complexity of the classification boundaries and the data manifold. The former dictates that any greedy search (e.g. gradient-based) near the boundaries will suffer from local minima; while the latter means that not all perturbation directions result in equal visual quality (in-manifold perturbation tends to be better than off-manifold perturbation). Consequently, there are often conflicts between these two spaces when searching for \mathbf{x}' . To reconcile the conflicts, we propose a method called *guided manifold walk* (GMW). GMW consists of three sub-routines: aimed probing, random exploration, and manifold projection. It starts from a random position (untargeted attack) outside $C_{\mathbf{x}}$, or a random sample within a specific class (targeted attack with the specific class as the targeted class). It then can approach \mathbf{x} by aimed probing attempting to find a sample which is close to the boundary of $C_{\mathbf{x}}$, or by random explorations to overcome local minima to find samples that are closer to \mathbf{x} , or by manifold projection to find the closest point on the data manifold. The above sub-routines are iteratively executed driven by the quality of the adversarial sample until a satisfactory \mathbf{x}' is found or the maximum number of iterations is reached.

We extensively evaluate BASAR on several state-of-the-art methods using multiple datasets in both untargeted and targeted attack tasks. The results show that not only is BASAR successful across models and datasets, it can also find on-manifold adversarial samples, in contrast to the common assumption that adversarial samples only exist off-manifold [18]. On par with very recent work that also found on-manifold samples in images [43], we show, for the first time, the existence and commonality of such samples in skeletal motions. We also comprehensively compare BASAR with other methods, showing the superiority of BASAR by large margins. Finally, since the perturbation magnitude alone is not enough to evaluate the attack quality, we conduct harsh perceptual studies to evaluate the naturalness, deceitfulness, and indistinguishability of the attack.

Formally, we demonstrate that adversarial attack is truly a threat to skeleton-based activity recognition. To this end, we propose the first black-box attack method and comprehensively evaluate the vulnerability of several state-of-the-art activity recognition methods. We show the existence of

on-manifold adversarial samples in various skeletal motion datasets and provide key insights on what classifiers tend to resist on-manifold adversarial samples.

2. Related Work

2.1. Skeleton-based Activity Recognition

Early research has well explored a variety of hand-crafted features [48, 17, 13] to recognize human actions from skeletal motions. Benefiting from the ability of neural networks, recent deep learning based methods have achieved state-of-the-art performance. Among them, RNN-based methods treat the motion data as a time series of joint coordinates, and employ common RNN architectures such as LSTMs to conduct the motion classification [15, 30, 41, 57]. CNN-based methods try to utilize image classification methods to recognize the motion data [31, 25, 42]. To this end, they transform the motion data into 2D pseudo-images, where each column stores a frame of skeleton joint coordinates. Graph convolutional network (GCN) based methods consider the skeleton as a topological graph where the nodes correspond to joints and edges correspond to the bones connecting neighboring joints. They perform convolution operations on graphs to effectively recognize the actions [56, 28, 40, 39, 10, 60]. For example, to capture long-range joint relations, Liu et al. [32] utilize disentangled multi-scale graph convolutions, which contains a unified spatial-temporal GCN operator for capturing complex spatial-temporal dependencies. Zhang et al. [58] introduce joint semantics to the GCN model, resulting in a lightweight yet effective method.

Our work is complementary to existing research by exploring their vulnerability to adversarial attacks and suggesting potential improvements. Based on the codes shared by the authors, we extensively evaluated BASAR on several state-of-the-art methods, demonstrating that even the latest methods with remarkable successes are still vulnerable to adversarial attacks.

2.2. Adversarial Attack

2.2.1 White-box Attack

A large variety of adversarial attack methods have been explored in different tasks [1, 44, 59]. Most of them consider the white-box setting, where the full-knowledge of the classifier is known by the attacker. Apart from common vision-based tasks such as classification [19, 26, 5], adversarial attacks on time series and action recognition have attracted attention recently. For general time series, multivariate time-series attack has been conducted [29, 20] based on adversarial transformation networks (ATN) [2]. Untargeted attack is proposed for optical flow based action recognition [22]. For skeletal action recognition, a one-step and an iterative method are proposed [29] based on FGSM and PGD [26]

respectively to attack GNN-based models. Wang et al. [50] achieve better attacking results using a novel perceptual loss that minimizes the dynamics difference between the original and the adversarial motion. Moreover, the performance is proven by extensive perceptual studies. Zheng et al. [61] formulate adversarial attack as an optimization problem with joint angle constraints in the Euclidean space, then solve it with Alternating Direction Method of Multipliers.

Although white-box attacks achieve satisfactory performance based on the full access to the attacked models, they are not very applicable in real-world scenarios since the details of classifiers are not usually exposed to the attacker. Hence, it is unclear whether the classifier vulnerability under (white-box) adversarial attack is a real threat. We on the other hand show that black-box attack is possible and is more threatening.

2.2.2 Black-box Attack

The difficulty of applying white-box attack in the real-world motivates the study of black-box attack, in which the attackers cannot access any underlying information of the attacked model but only make queries. A simple approach is transfer-based attack, which generates adversarial samples from one surrogate model via white-box attack [36, 35]. Existing black-box methods on skeletal motion [29, 50, 61] all rely on such a method, and cannot guarantee success due to the heavy dependence on the surrogate model [50]. Another type of methods is to use the predicted scores or soft labels (e.g. the softmax layer output) for attack [8, 47]. However, this is not exactly a black-box setting since the attacker can still access the model.

In a truly black-box setting, only the final class labels (hard-labels) can be used. Brendel et al. [3] perform the first hard-label attack by a random walk along the decision boundary. Optimization-based approaches select the search direction by estimating the gradient using the information computed at the decision boundary [7, 12, 11]. Dong et al. [14] model the local geometry of the search directions through an evolutionary attack algorithm. Li et al. [27] explore the subspace optimization methods from three different representative subspaces including spatial, frequency, and intrinsic component.

Hard-label black-box attack on time-series data is much less explored. In the black-box setting in [20, 23], to replicate/approximate the target model, the authors train a local substitute student network to mimic the final decision of the teacher classifier, to attack general time series. To attack videos, Wei et al. [55] heuristically attack the key frames and regions based on [12] due to the high dimensionality. To our best knowledge, there is no prior work on hard-label black-box attack on skeletal activity recognition.




Figure 1. 2D illustration of BASAR on a single frame. x is the attacked frame. x' is the ideal adversarial sample. \mathcal{M} (black line) is the natural pose manifold and ∂C_x (blue line) is the class boundary of C_x . x'_1 is the result of last iteration. x'_2 , x'_3 and x'_4 are the three intermediate results of the current iteration.

3. Methodology

We denote a motion with n frames as $\mathbf{x} = \{x^1, \dots, x^n\}$, each frame $x^t = \{q_1, \dots, q_m\}$ including m Dofs. These Dofs are usually joint positions or angles, depending on the data. Existing action recognizers take motions as input and output a class label. Specifically, given a trained recognizer G , G maps a motion to a probabilistic distribution over classes, $G: \mathbf{x} \rightarrow R^k$ where k is the total number of action classes. The class label C_x then can be derived e.g. via *softmax*. To find an adversarial sample \mathbf{x}' of motion \mathbf{x} , we start with the formulation in [5] but augment it for modeling motions:

$$\begin{aligned} & \text{minimize} && L(\mathbf{x}, \mathbf{x}') \\ & \text{subject to} && C_{\mathbf{x}'} = c \text{ and } \mathbf{x}' \in [0, 1]^{m \times n}, \end{aligned} \quad (1)$$

where L is the Euclidean distance. c is the targeted class. Note that the constraint can also be replaced by $C_{\mathbf{x}'} \neq C_x$ for untargeted attack. However, simply applying this formulation to skeletal motions is not sufficient because it only restricts the adversarial sample \mathbf{x}' in a hyper-cube $[0, 1]^{m \times n}$. Given that human poses lie in a natural pose manifold \mathcal{M} [53, 54], \mathbf{x}' can easily contain off-manifold frames which are unnatural/implausible poses and easily perceptible. We therefore add another constraint $\mathbf{x}' \in \mathcal{M}$:

$$\begin{aligned} & \text{minimize} && L(\mathbf{x}, \mathbf{x}') \\ & \text{subject to} && \mathbf{x}' \in [0, 1]^{m \times n}, \mathbf{x}' \in \mathcal{M} \\ & && C_{\mathbf{x}'} = c \text{ (targeted) or } C_{\mathbf{x}'} \neq C_x \text{ (untargeted)}. \end{aligned} \quad (2)$$

In practice, we find that $\mathbf{x}' \in [0, 1]^{m \times n}$ is less restrictive than other constraints and always satisfied. Eq. 2 is highly nonlinear and cannot be solved analytically. It thus requires a numerical solution.

3.1. Guided Manifold Walk

We propose a new method called *Guided Manifold Walk* (GMW) to solve Eq. 2. For simplicity, we start with a single frame x and a 2D illustration of GMW on x shown in

Fig. 1, where x' is the ideal adversarial sample which is on-manifold and close to x . Given the non-linearity of the classification boundary and the data manifold, BASAR aims to exploit the properties of both simultaneously, inspired by decision-boundary-based attack [3]. In each iteration, BASAR conducts three sub-routines: *random exploration*, *aimed probing*, and *manifold projection*. Random exploration is to explore the vicinity of the current adversarial sample x'_1 to find a random sample x'_2 . Aimed probing is to find a sample x'_3 in proximity to $\partial C_{\mathbf{x}}$ and is closer to x , in the direction from x'_2 to x . Finally, manifold projection is to project x'_3 onto \mathcal{M} to obtain x'_4 . Although GMW is only explained for a single frame x , it can be applied to a motion \mathbf{x} . The algorithm overview for perturbing the whole motion is given in Alg. 1, where λ and β are parameters and l is a distance function. They are defined in Eq. 3, 4 and 7 respectively. Next, we give details of random exploration, aimed probing, and manifold projection.

3.2. Random Exploration

We first extend the image-based random exploration in [3] to motions with dynamic weighting. Random exploration is an operation to explore in proximity to the classification boundary, by making a small step towards a random direction. The random perturbation is calculated by Eq. 3:

$$\begin{aligned} \tilde{\mathbf{x}} &= \mathbf{x}' + \mathbf{W}\Delta, \\ \text{where } \Delta_* &= \mathbf{R}_* - (\mathbf{R}_*^T \mathbf{d}_*) \mathbf{d}_*, \mathbf{d}_* = \frac{\mathbf{x}_* - \mathbf{x}'_*}{\|\mathbf{x}_* - \mathbf{x}'_*\|}, \\ \mathbf{R}_* &= \lambda \frac{\mathbf{r}}{\|\mathbf{r}\|} \|\mathbf{x}_* - \mathbf{x}'_*\|, r \in N(0, \mathbf{I}), \end{aligned} \quad (3)$$

where $\tilde{\mathbf{x}}$ is the new perturbed sample, \mathbf{x} and \mathbf{x}' are the attacked motion and current adversarial sample. We use joint positions and the subscript $*$ indicates either the x , y , or z joint coordinate. Unlike [3], the update on \mathbf{x}' is Δ weighted by \mathbf{W} - a diagonal matrix with joint weights. This is because we observe that perturbations on different joints are not equally effective and imperceptible, e.g. perturbation on the spinal joints cause larger visual distortion but is less effective. We therefore weight them differently. Δ_* controls the direction and magnitude of the update, and depends on two variables \mathbf{R}_* and \mathbf{d}_* . \mathbf{d}_* is the directional vector from \mathbf{x}' to \mathbf{x} . \mathbf{R}_* is a random directional vector sampled from a Normal distribution $N(0, \mathbf{I})$ where \mathbf{I} is an identity matrix, $\mathbf{I} \in R^{z \times z}$, $z = mn/3$, m is the number of Dofs in one frame and n is total frame number. This directional vector is scaled by $\|\mathbf{x}_* - \mathbf{x}'_*\|$ and λ .

In Eq. 3, Δ is *orthogonal* to \mathbf{d} . The random exploration essentially explores in the directions that are orthogonal to the direction towards \mathbf{x} , within the local area around \mathbf{x}' : $\mathbf{x}' + \mathbf{W}\Delta$. This mainly aims to avoid local minima along the $\partial C_{\mathbf{x}}$ when approaching \mathbf{x} . If $\tilde{\mathbf{x}}$ is not adversarial, we

Algorithm 1: Overview of the GMW

Input: attacked motion \mathbf{x} ; random sample $\tilde{\mathbf{x}}_0$,
 $C_{\tilde{\mathbf{x}}_0} = c$ (targeted) or $C_{\tilde{\mathbf{x}}_0} \neq C_{\mathbf{x}}$ (untargeted);
maximum number of iterations K ; threshold ϵ ;
Initialization: $\mathbf{x}'_0 = \text{AimedProbing}(\tilde{\mathbf{x}}_0, \mathbf{x})$,
ensuring that \mathbf{x}'_0 is adversarial;
for $k = 1$ **to** K **do**
 repeat
 | $\tilde{\mathbf{x}} = \text{RandomExploration}(\mathbf{x}'_{k-1}, \mathbf{x})$;
 until $\tilde{\mathbf{x}}$ is adversarial or $\lambda < 10^{-10}$;
 if $\lambda \geq 10^{-10}$ **then** $\mathbf{x}'_k = \tilde{\mathbf{x}}$ **else break**;
 repeat
 | $\tilde{\mathbf{x}} = \text{AimedProbing}(\mathbf{x}'_k, \mathbf{x})$;
 until $\tilde{\mathbf{x}}$ is adversarial or $\beta < 10^{-10}$;
 if $\beta \geq 10^{-10}$ **then** $\mathbf{x}'_k = \tilde{\mathbf{x}}$ **else break**;
 $\hat{\mathbf{x}} = \text{ManifoldProjection}(\mathbf{x}'_k, \mathbf{x})$;
 while $\hat{\mathbf{x}}$ is not adversarial **do**
 | $\hat{\mathbf{x}} = \text{AimedProbing}(\mathbf{x}'_k, \hat{\mathbf{x}})$;
 end
 if $l(\hat{\mathbf{x}}, \mathbf{x}) \geq \epsilon$ **then** $\mathbf{x}'_k = \hat{\mathbf{x}}$ **else break**;
end

decrease λ and perform the random exploration again; otherwise, we increase λ , then enter the next sub-routine.

3.3. Aimed Probing

We adopt Aimed probing from [3], aiming to find a new adversarial sample between the perturbed motion and the original, so that the new adversarial sample is closer to the attacked motion:

$$\tilde{\mathbf{x}} = \mathbf{x}' + \beta(\mathbf{x} - \mathbf{x}'), \quad (4)$$

where β is a forward step size. β is decreased to conduct the aimed probing again if $\tilde{\mathbf{x}}$ is not adversarial; otherwise, we increase β , then enter the next sub-routine.

3.4. Manifold Projection

After aimed probing and random exploration, the perturbed motion is often off the manifold, resulting in implausible and unnatural poses. We thus project them back to the manifold. Natural pose manifold can be obtained in two ways: explicit modeling [52] or implicit learning [53, 9]. Using implicit learning would require to train a data-driven model then use it for projection [53], breaking BASAR into a two-step system. Therefore we employ explicit modeling. Specifically, we replace the constraint $\mathbf{x}' \in \mathcal{M}$ in Eq. 2 with hard constraints on bone lengths and joint rotations.

We also constrain the dynamics of \mathbf{x}' to be similar to \mathbf{x} :

$$\begin{aligned} & \text{minimize } L(\mathbf{x}, \mathbf{x}') + wL(\ddot{\mathbf{x}}, \ddot{\mathbf{x}}') \\ & \text{subject to } B'_i = B_i \text{ and } \theta_i^{\min} \leq \theta'_i \leq \theta_i^{\max} \\ & C_{\mathbf{x}'} = c \text{ (targeted) or } C_{\mathbf{x}'} \neq C_{\mathbf{x}} \text{ (untargeted)}, \end{aligned} \quad (5)$$

where $\ddot{\mathbf{x}}$ and $\ddot{\mathbf{x}}'$ are the $2nd$ -order derivatives of \mathbf{x} and \mathbf{x}' , w is a weight. Matching the $2nd$ -order derivatives is proven to be important for visual imperceptibility in adversarial attack [50]. L is the Euclidean distance. B_i and B'_i are the i -th bone's lengths of the attacked and adversarial motion respectively. When the bone lengths change from frame to frame in the original data, we impose the bone-length constraint on each frame. θ'_i is the i -th joint in every frame of \mathbf{x}' and subject to joint limits bounded by θ_i^{\min} and θ_i^{\max} .

Eq. 5 is difficult to solve, especially to satisfy both the bone length and joint limit constraints in the joint position space [61]. We therefore solve Eq. 5 in two steps. First, we solve it without any constraints by Inverse Kinematics [52] in the joint angle space, which automatically preserves the bone lengths. Next, Eq. 5 is solved in the joint angle space:

$$\begin{aligned} & \text{minimize } L(\theta, \theta') + wL(\ddot{\theta}, \ddot{\theta}') \\ & \text{subject to } \theta_i^{\min} \leq \theta'_i \leq \theta_i^{\max}, \\ & C_{\mathbf{x}'} = c \text{ (targeted) or } C_{\mathbf{x}'} \neq C_{\mathbf{x}} \text{ (untargeted)}. \end{aligned} \quad (6)$$

Note that the objective function in Eq. 6 is designed to match the joint angles and the joint angular acceleration. We use a primal-dual interior-point method [21, 49] to solve Eq. 6. After solving for θ' , the joint positions of the adversarial motion are computed using Forward Kinematics. Please refer to the supplementary materials for the details of mathematical deduction, implementation, and performance.

4. Experiments

4.1. Target Models and Datasets

We select three state-of-the-art action recognition models to attack, STGCN [56], MS-G3D [32] and SGN [58]. STGCN [56] is one of the first GCN-based methods. MS-G3D [32] and SGN [58] are both state-of-the-art methods. The three methods are sufficiently representative. We choose three frequently used benchmark datasets: HDM05 [34], NTU60 [38] and Kinetics [24]. HDM05 dataset [34] has 130 action classes, 2337 sequences from 5 non-professional actors. Due to its high quality, any obvious difference between the adversarial motion and original motion can be easily perceived, making HDM05 suitable for our perceptual study. Since our target models are not trained on HDM05 in the original papers, we process HDM05 following [15] and train the target models strictly following the protocols in the papers, achieving 87.2%, 94.4%, 99.1% on STGCN, MS-G3D and SGN respectively. NTU RGB+D

60 [38] includes 56568 skeleton sequences with 60 action classes, performed by 40 subjects. Due to the large intra-class and viewpoint variations, it is ideal for verifying the effectiveness and generalizability of our approach. Kinetics 400 [24] is a large but highly noisy human action video dataset taken from different YouTube Videos. The skeletons are extracted from Openpose [4], consisting of over 260000 skeleton sequences. Since SGN is not trained on it, we follow [58] to preprocess Kinetics skeletons and then train the model. Because some target models are big and making queries to them become slow, it is impractical to attack all motions in a dataset. So we randomly sample motions to attack. We gradually increase the number of motions to attack until all evaluation metrics (explained below) stabilize, so that we know the attacked motions are sufficiently representative in the dataset.

4.2. Evaluation Metrics

We employ the success rate as a major indicator for evaluation. In addition, to further numerically evaluate the quality of the adversarial samples, we also define evaluation metrics between the original motion \mathbf{x} and its adversarial sample \mathbf{x}' , including the averaged joint position deviation l , the averaged joint acceleration deviation Δa , the averaged joint angular acceleration deviation $\Delta \alpha$, and the averaged bone-length deviation percentage $\Delta B/B$:

$$\begin{aligned} l &= \frac{1}{nN} \sum_{j=0}^N \|\mathbf{x}^{(j)} - \mathbf{x}'^{(j)}\|_2 \\ \Delta a &= \frac{1}{nON} \sum_{j=0}^N \|\ddot{\mathbf{x}}^{(j)} - \ddot{\mathbf{x}}'^{(j)}\|_2 \\ \Delta \alpha &= \frac{1}{nON} \sum_{j=0}^N \|\ddot{\theta}^{(j)} - \ddot{\theta}'^{(j)}\|_2 \\ \Delta B/B &= \frac{\sum_{j=0}^N \sum_{i=0}^T ((B_i^{(j)} - B_i'^{(j)})/B_i^{(j)})}{TN} \end{aligned} \quad (7)$$

where N is the number of adversarial samples. O and T are the total number of joints and bones in a skeleton. n is the number of frames in a motion. In addition, we also investigate the percentage of on-manifold (OM) adversarial motions after the attack. We regard a motion as on-manifold if all its frames respect bone-length and joint limit constraints. Finally, since Kinetics has missing joints from time to time, it is impossible to attack it in the joint angle space. So we only attack it in the joint position space. Consequently, $\Delta \alpha$ and OM cannot be computed on Kinetics.

4.3. Untargeted Attack

To initialize for untargeted attack, we randomly sample a motion \mathbf{x}' for a target motion \mathbf{x} where $C_{\mathbf{x}'} \neq C_{\mathbf{x}}$.

Models		$l\downarrow$	$\Delta a\downarrow$	$\Delta\alpha\downarrow$	$\Delta B/B\downarrow$	OM \uparrow
STGCN	MP	0.13	0.05	0.11	0.00%	99.55%
	No MP	0.10	0.04	0.34	0.66%	0.00%
MSG3D	MP	0.76	0.12	0.49	1.78%	0.13%
	No MP	0.70	0.09	0.82	1.81%	0.00%
SGN	MP	11.53	1.92	6.70	9.60%	60.52%
	No MP	7.93	2.00	14.36	39.64%	0.00%
STGCN	MP	0.08	0.02	0.07	4.82%	4.68%
	No MP	0.10	0.02	0.09	5.57%	1.82%
MSG3D	MP	0.08	0.03	0.12	8.14%	0.86%
	No MP	0.12	0.03	0.17	10.02%	0.57%
SGN	MP	0.28	0.08	0.21	11.11%	28.95%
	No MP	0.30	0.10	0.42	28.00%	4.55%
STGCN	MP	0.05	0.0057	n/a	2.54%	n/a
	No MP	0.07	0.0062	n/a	3.53%	n/a
MSG3D	MP	0.10	0.011	n/a	5.16%	n/a
	No MP	0.10	0.012	n/a	5.69%	n/a
SGN	MP	0.12	0.020	n/a	4.23%	n/a
	No MP	0.13	0.022	n/a	6.93%	n/a

Table 1. Untargeted attack on HDM05 (top), NTU (middle) and Kinetics (bottom). All attacks have a 100% success rate. l : averaged joint position deviation; Δa : averaged joint acceleration deviation; $\Delta\alpha$: averaged joint angular acceleration deviation; $\Delta B/B$: averaged bone-length deviation percentage; on-manifold sample percentage (OM). MP means Manifold Projection.

For HDM05, we randomly select 700 motions to attack STGCN, MSG3D and SGN. For NTU and Kinetics, we randomly sample 1200 and 500 motions respectively. The maximum iterations are 500, 1000 and 2000 on HDM05, Kinetics and NTU respectively.

The results are shown in Tab. 1. Note that BASAR achieves 100% success in all tasks. Here we also conduct ablation studies (MP/No MP) to show the effects of manifold projection. First, the universal successes across all datasets and models demonstrate the effectiveness of BASAR. The manifold projection directly affects the OM results. BASAR can generate as high as 99.55% on-manifold adversarial samples. As shown in the perceptual study later, the on-manifold samples are very hard to be distinguished from the original motions even under very harsh visual comparisons. Detailed confusion matrices, visual results and ablation studies can be found in the supplementary materials. We leave the performance variation analysis across models to Section 4.5.

4.4. Targeted Attack

For targeted attack, we randomly select the same number of motions from each dataset as in untargeted attack, with maximum iteration set to 3000. To initiate a targeted attack on motion \mathbf{x} , we randomly select a motion \mathbf{x}' where $C_{\mathbf{x}'} = c$ and c is the targeted class.

The results are shown in Tab. 2. All attacks achieve

Models		$l\downarrow$	$\Delta a\downarrow$	$\Delta\alpha\downarrow$	$\Delta B/B\downarrow$	OM \uparrow
STGCN	MP	4.97	0.10	0.65	3.44%	56.98%
	No MP	6.25	0.09	0.92	5.85%	0.00%
MSG3D	MP	4.34	0.12	0.71	4.51%	1.64%
	No MP	4.35	0.11	1.01	5.08%	0.00%
SGN	MP	16.31	1.28	6.97	12.29%	20.96%
	No MP	16.13	1.63	13.28	29.86%	0.00%
STGCN	MP	0.37	0.03	0.25	9.73%	0.63%
	No MP	0.38	0.04	0.16	11.55%	0.16%
MSG3D	MP	0.36	0.05	0.24	15.43%	0.00%
	No OP	0.40	0.06	0.27	17.72%	0.00%
SGN	MP	1.28	0.09	0.38	28.24%	2.63%
	No MP	1.35	0.10	0.53	39.43%	0.00%
STGCN	MP	0.63	0.03	n/a	29.10%	n/a
	No MP	0.67	0.03	n/a	31.48%	n/a
MSG3D	MP	0.56	0.05	n/a	27.26%	n/a
	No MP	0.57	0.07	n/a	28.35%	n/a
SGN	MP	1.51	0.18	n/a	68.45%	n/a
	No MP	1.54	0.19	n/a	72.09%	n/a

Table 2. Targeted attack on HDM05 (top), NTU (middle) and Kinetics (bottom). All attacks have a 100% success rate.

100% success. The targeted attack is more challenging than the untargeted attack [50], because the randomly selected label often has completely different semantic meanings from the original one. Attacking an ‘eating’ motion to ‘drinking’ is much easier than to ‘running’. This is why the targeted attack in general has worse results than untargeted attack under every metric. Even under such harsh settings, BASAR can still produce as high as 56.98% on-manifold adversarial samples. The performance variation across models is consistent with the untargeted attack. Additional details can be found in the supplementary materials.

4.5. Classifier Robustness

We notice that BASAR’s performance varies across target models consistently in both the untargeted and targeted attack. In other words, action recognizers are not equally gullible. SGN, in general, is the hardest to fool across all datasets. When looking at the results, the perturbation always needs to be larger to fool SGN, shown by larger l , Δa , $\Delta\alpha$ and $\Delta B/B$ on every dataset compared with STGCN and MSG3D. This includes both joint-angle and joint-position attack. We speculate that it has to do with the features that SGN uses. Unlike STGCN and MSG3D which use raw joints and bones and rely on networks to learn good classification features, SGN also employs *semantic* features, where different joint types are encoded in learning their patterns and correlations. This type of semantic information requires large perturbations to bring the motion out of its pattern. Hence, semantic information improves its robustness against attacks, as larger perturbations are more likely to be perceptible. Note that SGN sometimes has a higher

OM percentage. However, we find that these OM motions have large visual differences from the original motions. In other words, unlike STGCN and MSG3D where a higher OM percentage indicates more visually indistinguishable adversarial samples, the OM samples of SGN look natural, can fool the classifier and probably can fool humans when being observed independently, but are unlikely to survive strict side-by-side comparisons with the original motions. Next, MSG3D is slightly harder to fool than STGCN. Although both use joint positions, MSG3D explicitly uses the bone information, which essentially recognizes the relative movements of joints. The relative movement pattern of joints helps resist attack by requiring the perturbation to be large enough to break the patterns. Finally, the robustness of a specific model varies across datasets. Meanwhile, the pattern of robustness variations across datasets also differs from model to model. It is hard to theoretically identify the cause and we will leave it to future research. Additional detailed findings can be found in the supplementary materials.

4.6. Perceptual Study

Numerical evaluation alone is not sufficient to evaluate the imperceptibility of adversarial attack on skeletal motions, because they cannot accurately indicate whether the attack is perceptible to humans [50]. We, therefore, conduct three perceptual studies: Deceitfulness, Naturalness, Indistinguishability, following [50]. Since we have 36 scenarios (models vs datasets vs attack types vs MP/No MP), it is unrealistic to exhaustively cover all conditions. We choose HDM05 and untargeted attack for our perceptual studies. We exclude NTU and Kinetics as they contain severe and noticeable noises. Our preliminary study shows that it is hard for people to tell if a motion is attacked therein. The quality of HDM05 is high where perturbations can be easily identified. In total, we recruited 50 subjects (age between 20 and 54). We briefly explain our study setting here and refer the readers to the supplementary materials for details.

We first test whether BASAR visually changes the meaning of the motion and whether the meaning of the original motion is clear to the subjects (**Deceitfulness**). In each user study, we randomly choose 45 motions (15 from STGCN, MSG3D and SGN respectively) with the ground-truth label and after-attack label for 45 trials. In each trial, the video is played for 6 seconds then the user is asked to choose which label best describes the motion with no time limits. We also perform ablation studies to test whether on-manifold adversarial samples look more natural than off-manifold samples (**Naturalness**). In each trial, we randomly select 120 motions for 60 trials. Each trial includes one motion attacked with manifold projection and one without. The two motions are played side-by-side for 6 seconds twice, then the user needs to pick which one looks more natural, or cannot tell the difference, with no time limits. Finally,




Figure 2. Normalized user preferences on Naturalness. MP/No MP refers to with/without manifold projection. The vertical axis is the percentage of user preference. The horizontal axis is the attacked models.

we conduct the strictest test to see whether the adversarial samples from BASAR can survive side-by-side scrutiny (**Indistinguishability**). In each trial, two motions are displayed side by side. The left motion is always the original and the user is told so. The right one can be the original motion or the attacked motion. The user is asked whether he/she can see any visual difference after the motions are played for 6 seconds twice.

4.6.1 Results

The average success rate of **Deceitfulness** is 79.64% across three models, with 88.13% , 83.33%, 67.47% on STGCN, MS-G3D and SGN respectively. This is consistent with our prediction because SGN requires larger perturbations, thus is more likely to lead to the change of the motion semantics. Next, Fig. 2 shows the results of **Naturalness**. It is very clear that subjects regard on-manifold samples as more natural. This is understandable as manifold projection not only makes sure the poses are on the manifold, but also enforces the similarity of the dynamics between the attacked and original motion. Finally, the results of **Indistinguishability** are 89.90% on average. BASAR even outperforms the white-box attack (80.83%) in [50]. We further look into on-manifold vs off-manifold. Both samples are tested in Indistinguishability, but 94.63% of the on-manifold samples fooled the users; while 84.69% of the off-manifold samples fooled the users, showing that on-manifold samples are more deceitful. The off-manifold samples which successfully fool the users contain only very small deviations.

4.7. Comparison

Since BASAR is the very first black-box adversarial attack method on skeletal motions, there is no baseline for comparison. We therefore employ several methods that are closet to our approach. The first is SMART [50] which also attacks skeletal motions but is a white-box approach. Although it can deliver a black-box attack, it needs a surrogate model. According to their work, we choose HRNN [15] and 2SAGCN [40] as the surrogate models. The second baseline method is MTS [20] which is a black-box method but only on general multivariate time-series. It is the most

similar method to BASAR but does not model the data manifold. Another baseline method is BA [3] which attacks image data. Since it is also a black-box approach and focuses on the boundary attack, we also include it in the comparison. We choose HDM05 and NTU for comparisons. For each comparison, we randomly select 1200 and 490 motions from NTU and HDM05. Since MTS is not designed for untargeted attack, we only compare BASAR with it on the targeted attack.

Tab. 3 lists the success rates of all methods. BASAR performs the best and often by big margins. In the targeted attack, the highest among the baseline methods is merely 30.31% which is achieved by SMART on HDM05/MSG3D. BASAR achieves 100%. In the untargeted attack, the baseline methods achieve higher performances but still worse than BASAR. SMART achieves as high as 99.33% on NTU/MSG3D. However, its performance is not reliable as it highly depends on the chosen surrogate model, which is consistent with [50]. In addition, we further look into the results and find that SMART’s results are inconsistent. When the attack is transferred, the result labels are often different from the labels obtained during the attack.

We find that BA can also achieve 100% success. However, BA is designed to attack image data and does not consider the data manifold. We therefore compare detailed metrics and show the results in Tab. 4. BA is in general worse than BASAR under every metric. The worst is the bone-length constraint violation. Visually, the skeletal structure cannot be observed at all. This happens for both the untargeted and the targeted attack across all datasets and models. This is understandable because BA’s assumption is that perturbations in all directions have the same effect, while BASAR assumes that in-manifold perturbations provide better visual quality.

5. Discussion and Conclusion

To systematically investigate the vulnerability of state-of-the-art action recognizers, we proposed the very first black-box adversarial attack method which gives strong performance across datasets, models and attack modes. Quantitative evaluations demonstrate the high quality of the attack. Harsh perceptual studies show the imperceptibility of BASAR. Through comprehensive comparisons, BASAR outperforms existing methods by large margins in the success rate and quality. More broadly, we show, for the first time, the wide existence of on-manifold adversarial samples in skeletal motions. Based on the attack, we discuss what classifiers or features tend to be robust against attack. BASAR raises a serious yet unaddressed security concern that we expect will inform the future research. In future, we will investigate the perceptually-guided adversarial attack. We found that imperceptible perturbations form a small subspace. The interplay among the classification boundary,

Targeted	Attack Method	HDM05	NTU
STGCN	BASAR	100.00%	100.00%
	MTS	3.27%	12.00%
	SMART(HRNN)	3.20%	2.27%
	SMART(2SAGCN)	2.33%	0.19%
MSG3D	BASAR	100.00%	100.00%
	MTS	2.18%	12.90%
	SMART(HRNN)	30.31%	1.16%
	SMART(2SAGCN)	2.46%	0.00%
SGN	BASAR	100.00%	100.00%
	MTS	2.91%	0.00%
	SMART(HRNN)	29.69%	1.42%
	SMART(2SAGCN)	3.28%	1.83%
Untargeted	Attack Method	HDM05	NTU
STGCN	BASAR	100.00%	100.00%
	SMART(HRNN)	66.87%	89.25%
	SMART(2SAGCN)	86.10%	12.88%
MSG3D	BASAR	100.00%	100.00%
	SMART(HRNN)	86.88%	99.33%
	SMART(2SAGCN)	88.73%	3.08%
SGN	BASAR	100.00%	100.00%
	SMART(HRNN)	89.25%	98.25%
	SMART(2SAGCN)	0.41%	97.75%

Table 3. Attack success rate comparison with baseline methods.

Models		$l\uparrow$	$\Delta a\uparrow$	$\Delta\alpha\uparrow$	$\Delta B/B\uparrow$	$OM\downarrow$
STGCN	UA	1.44	0.65	4.74	10.60%	0.00%
	TA	8.83	0.17	1.60	8.56%	0.00%
MSG3D	UA	1.17	0.36	2.81	6.00%	0.00%
	TA	7.93	0.10	1.07	7.49%	0.00%
SGN	UA	13.35	3.45	21.96	75.11%	0.00%
	TA	15.40	1.52	11.77	29.45%	0.00%
STGCN	UA	1.04	0.47	1.97	235.10%	0.00%
	TA	0.41	0.08	0.37	35.00%	0.00%
MSG3D	UA	1.24	1.73	2.38	911.7%	0.00%
	TA	0.27	0.07	0.34	25.72%	0.00%
SGN	UA	0.22	0.28	1.26	125.57%	3.60%
	TA	0.42	0.15	0.66	65.31%	0.17%

Table 4. Boundary Attack (BA) on HDM05 (top), NTU (bottom). UA/TA refers Untargeted/Targeted Attack.

the data manifold and the imperceptible subspace is likely to lead to adversarial samples with higher quality. The interplay is also likely to shed light on effective defense strategies against imperceptible attack.

Acknowledgements: This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 899739 CrowdDNA, EPSRC (EP/R031193/1), NSF China (No. 61772462, No. U1736217), RCUK grant CAMERA (EP/M023281/1, EP/T014865/1), China Scholarship Council (201907000157), and the 100 Talents Program of Zhejiang University.

References

- [1] Naveed Akhtar and Ajmal Mian. Threat of adversarial attacks on deep learning in computer vision: A survey. *IEEE Access*, 6:14410–14430, 2018. 1, 2
- [2] Shumeet Baluja and Ian Fischer. Adversarial transformation networks: Learning to generate adversarial examples. *arXiv preprint arXiv:1703.09387*, 2017. 2
- [3] Wieland Brendel, Jonas Rauber, and Matthias Bethge. Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. In *6th International Conference on Learning Representations*, 2018. 3, 4, 8
- [4] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7291–7299, 2017. 5
- [5] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. *IEEE Symposium on Security and Privacy*, 2017. 2, 3
- [6] Anirban Chakraborty, Manaar Alam, Vishal Dey, Anupam Chattopadhyay, and Debdeep Mukhopadhyay. Adversarial Attacks and Defences: A Survey. *arXiv:1810.00069 [cs, stat]*, 2018. arXiv: 1810.00069. 1
- [7] Jianbo Chen, Michael I Jordan, and Martin J Wainwright. Hopskipjumpattack: A query-efficient decision-based attack. In *IEEE Symposium on Security and Privacy (SP)*, pages 1277–1294, 2020. 3
- [8] Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, pages 15–26, 2017. 3
- [9] Wenheng Chen, He Wang, Yi Yuan, Tianjia Shao, and Kun Zhou. Dynamic future net: Diversified human motion generation. In *Proceedings of the 28th ACM International Conference on Multimedia*, 2020. 4
- [10] Ke Cheng, Yifan Zhang, Xiangyu He, Weihan Chen, Jian Cheng, and Hanqing Lu. Skeleton-based action recognition with shift graph convolutional network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [11] Minhao Cheng, Thong Le, Pin-Yu Chen, Huan Zhang, Jinfeng Yi, and Cho-Jui Hsieh. Query-efficient hard-label black-box attack: An optimization-based approach. In *International Conference on Learning Representations, 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. 3
- [12] Minhao Cheng, Simranjit Singh, Patrick H. Chen, Pin-Yu Chen, Sijia Liu, and Cho-Jui Hsieh. Sign-opt: A query-efficient hard-label adversarial attack. In *8th International Conference on Learning Representations, ICLR, 2020*. 3
- [13] Maxime Devanne, Hazem Wannous, Stefano Berretti, Pietro Pala, Mohamed Daoudi, and Alberto Del Bimbo. 3-d human action recognition by shape analysis of motion trajectories on riemannian manifold. *IEEE transactions on cybernetics*, 45(7):1340–1352, 2015. 2
- [14] Y. Dong, H. Su, Baoyuan Wu, Z. Li, W. Liu, Tong Zhang, and Jun Zhu. Efficient decision-based black-box adversarial attacks on face recognition. *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7706–7714, 2019. 3
- [15] Yong Du, Wei Wang, and Liang Wang. Hierarchical recurrent neural network for skeleton based action recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1110–1118, 2015. 2, 5, 7
- [16] Hassan Ismail Fawaz, G. Forestier, Jonathan Weber, L. Idoumghar, and P. Muller. Adversarial Attacks on Deep Neural Networks for Time Series Classification, 2019. 1
- [17] Basura Fernando, Efstratios Gavves, José Oramas M., Amir Ghodrati, and Tinne Tuytelaars. Modeling video evolution for action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5378–5387. IEEE Computer Society, 2015. 2
- [18] Justin Gilmer, Luke Metz, Fartash Faghri, Samuel S. Schoenholz, Maithra Raghu, Martin Wattenberg, and Ian Goodfellow. Adversarial Spheres. *arXiv:1801.02774 [cs]*, 2018. arXiv: 1801.02774. 1, 2
- [19] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *3rd International Conference on Learning Representations*, 2015. 2
- [20] Samuel Harford, Fazle Karim, and Houshang Darabi. Adversarial attacks on multivariate time series. *arXiv preprint arXiv:2004.00410*, 2020. 2, 3, 7
- [21] John D Hedengren, Reza Asgharzadeh Shishavan, Kody M Powell, and Thomas F Edgar. Nonlinear modeling, estimation and predictive control in apmonitor. *Computers & Chemical Engineering*, 70:133–148, 2014. 5
- [22] Nathan Inkawhich, Matthew Inkawhich, Yiran Chen, and Hai Li. Adversarial attacks for optical flow-based action recognition classifiers. *arXiv:1811.11875 [cs]*, 2018. arXiv: 1811.11875. 2
- [23] Fazle Karim, Somshubra Majumdar, and Houshang Darabi. Adversarial attacks on time series. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 1, 3
- [24] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset, 2017. 5
- [25] Qihong Ke, Mohammed Bennamoun, Senjian An, Ferdous Ahmed Sohel, and Farid Boussaïd. A new representation of skeleton sequences for 3d action recognition. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4570–4579, 2017. 2
- [26] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. *ICLR Workshop*, 2017. 2
- [27] Huichen Li, Xiaojun Xu, X. Zhang, S. Yang, and B. Li. Qeba: Query-efficient boundary-based blackbox attack. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1218–1227, 2020. 3
- [28] Maosen Li, Siheng Chen, Xu Chen, Ya Zhang, Yanfeng Wang, and Qi Tian. Actional-structural graph convolutional networks for skeleton-based action recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2

- [29] Jian Liu, Naveed Akhtar, and Ajmal Mian. Adversarial Attack on Skeleton-based Human Action Recognition. *arXiv:1909.06500 [cs]*, 2019. arXiv: 1909.06500. **1, 2, 3**
- [30] Jun Liu, Amir Shahroudy, Dong Xu, and Gang Wang. Spatio-temporal lstm with trust gates for 3d human action recognition. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 816–833, 2016. **2**
- [31] Mengyuan Liu, Hong Liu, and Chen Chen. Enhanced skeleton visualization for view invariant human action recognition. *Pattern Recogn.*, 68(C):346–362, 2017. **2**
- [32] Ziyu Liu, Hongwen Zhang, Zhenghao Chen, Zhiyong Wang, and Wanli Ouyang. Disentangling and unifying graph convolutions for skeleton-based action recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. **2, 5**
- [33] L. Minh Dang, Kyungbok Min, Hanxiang Wang, Md. Jalil Piran, Cheol Hee Lee, and Hyeonjoon Moon. Sensor-based and vision-based human activity recognition: A comprehensive survey. *Pattern Recognition*, 108, 2020. **1**
- [34] Meinard Müller, Tido Röder, Michael Clausen, Bernhard Eberhardt, Björn Krüger, and Andreas Weber. Documentation mocap database hdm05. Technical Report CG-2007-2, Universität Bonn, 2007. **5**
- [35] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z. Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM Asia Conference on Computer and Communications Security*, pages 506–519, 2017. **3**
- [36] Nicolas Papernot, Patrick D. McDaniel, and Ian J. Goodfellow. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *CoRR*, abs/1605.07277, 2016. **3**
- [37] Bin Ren, Mengyuan Liu, Runwei Ding, and Hong Liu. A Survey on 3D Skeleton-Based Action Recognition Using Learning Method. *arXiv:2002.05907 [cs]*, 2020. arXiv: 2002.05907. **1**
- [38] 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. **5**
- [39] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Skeleton-based action recognition with directed graph neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7912–7921, 2019. **2**
- [40] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. **2, 7**
- [41] 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 Thirty-First AAAI Conference on Artificial Intelligence, AAAI’17*, pages 4263–4270, 2017. **2**
- [42] 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 (CVPR) Workshops*, 2017. **2**
- [43] David Stutz, M. Hein, and B. Schiele. Disentangling adversarial robustness and generalization. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6969–6980, 2019. **2**
- [44] Lichao Sun, Yingdong Dou, Carl Yang, Ji Wang, Philip S Yu, and Bo Li. Adversarial attack and defense on graph data: A survey. *arXiv preprint arXiv:1812.10528*, 2018. **2**
- [45] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv:1312.6199 [cs]*, 2014. arXiv: 1312.6199. **1**
- [46] Florian Tramèr, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. The Space of Transferable Adversarial Examples. *arXiv:1704.03453 [cs, stat]*, 2017. arXiv: 1704.03453. **1**
- [47] Chun-Chen Tu, Pai-Shun Ting, Pin-Yu Chen, Sijia Liu, Huan Zhang, Jinfeng Yi, Cho-Jui Hsieh, and Shin-Ming Cheng. Autozoom: Autoencoder-based zeroth order optimization method for attacking black-box neural networks. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI*, pages 742–749, 2019. **3**
- [48] Raviteja Vemulapalli, Felipe Arrate, and Rama Chellappa. Human action recognition by representing 3d skeletons as points in a lie group. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 588–595, 2014. **2**
- [49] Andreas Wächter and Lorenz T Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1):25–57, 2006. **5**
- [50] He Wang, Feixiang He, Zhexi Peng, Yongliang Yang, Tianjia Shao, Kun Zhou, and David Hogg. SMART: Skeletal Motion Action Recognition aTtack. *arXiv:1911.07107*, 2020. **1, 3, 5, 6, 7, 8**
- [51] He Wang, Feixiang He, Zhexi Peng, Yong-Liang Yang, Tianjia Shao, Kun Zhou, and David Hogg. Understanding the robustness of skeleton-based action recognition under adversarial attack. *arXiv:2103.05347*, 2021. **1**
- [52] 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. **4, 5**
- [53] He Wang, Edmond S. L. Ho, Hubert P. H. Shum, and Zhanxing Zhu. Spatio-temporal manifold learning for human motions via long-horizon modeling. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–1, 2019. **3, 4**
- [54] He Wang, Kirill A Sidorov, Peter Sandilands, and Taku Komura. Harmonic parameterization by electrostatics. *ACM Transactions on Graphics (TOG)*, 32(5):155, 2013. **3**
- [55] Zhipeng Wei, Jingjing Chen, Xingxing Wei, Linxi Jiang, Tat-Seng Chua, Fengfeng Zhou, and Yu-Gang Jiang. Heuristic black-box adversarial attacks on video recognition models. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI*, pages 12338–12345, 2020. **3**
- [56] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, AAAI’18*, 2018. **2, 5**

- [57] Pengfei Zhang, Cuiling Lan, Junliang Xing, Wenjun Zeng, Jianru Xue, and Nanning Zheng. View adaptive neural networks for high performance skeleton-based human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8):1963–1978, 2019. [2](#)
- [58] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [2](#), [5](#)
- [59] Wei Emma Zhang, Quan Z. Sheng, Ahoud Alhazmi, and Chenliang Li. Adversarial attacks on deep-learning models in natural language processing: A survey. *ACM Transactions on Intelligent Systems and Technology*, 11(3):1–41, 2020. [2](#)
- [60] Xikun Zhang, Chang Xu, and Dacheng Tao. Context aware graph convolution for skeleton-based action recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [2](#)
- [61] Tianhang Zheng, Sheng Liu, Changyou Chen, Junsong Yuan, Baochun Li, and Kui Ren. Towards Understanding the Adversarial Vulnerability of Skeleton-based Action Recognition. *arXiv:2005.07151 [cs]*, 2020. arXiv: 2005.07151. [1](#), [3](#), [5](#)