

# Adar: Adversarial Activity Recognition in Wearables

Ramesh Kumar Sah, and Hassan Ghasemzadeh  
Washington State University, Pullman, WA 99164-2752  
{ramesh.sah,hassan.ghasemzadeh}@wsu.edu

## ABSTRACT

Recent advances in machine learning and deep neural networks have led to the realization of many important applications in the area of personalized medicine. Whether it is detecting activities of daily living or analyzing images for cancerous cells, machine learning algorithms have become the dominant choice for such emerging applications. In particular, the state-of-the-art algorithms used for human activity recognition (HAR) using wearable inertial sensors utilize machine learning algorithms to detect health events and to make predictions from sensor data. Currently, however, there remains a gap in research on whether or not and how activity recognition algorithms may become the subject of adversarial attacks. In this paper, we take the first strides on (1) investigating methods of generating adversarial example in the context of HAR systems; (2) studying the vulnerability of activity recognition models to adversarial examples in feature and signal domain; and (3) investigating the effects of adversarial training on HAR systems. We introduce Adar<sup>1</sup>, a novel computational framework for optimization-driven creation of adversarial examples in sensor-based activity recognition systems. Through extensive analysis based on real sensor data collected with human subjects, we found that simple evasion attacks are able to decrease the accuracy of a deep neural network from 95.1% to 3.4% and from 93.1% to 16.8% in the case of a convolutional neural network. With adversarial training, the robustness of the deep neural network increased on the adversarial examples by 49.1% in the worst case while the accuracy on clean samples decreased by 13.2%.

## KEYWORDS

mobile health, activity recognition, machine learning, adversarial examples, wearable sensors, security

## ACM Reference Format:

Ramesh Kumar Sah, and Hassan Ghasemzadeh. 2019. Adar: Adversarial Activity Recognition in Wearables. In *ICCAD '19: International Conference On Computer Aided Design, November 04–07, 2019 Westminster, CO*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

<sup>1</sup>Software code and experimental data for Adar are available online at <https://github.com/rameshKrSah/Adar>.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

ICCAD '19, November 04–07, 2019, Westminster, CO

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00  
<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

## 1 INTRODUCTION

Human activity recognition (HAR) is a major research area in the field of mobile and ubiquitous computing. Its applications in health monitoring, patients rehabilitation, assessment of performance in athletes, and gaming industry [6] are the contributing factors for the active research in this area. Moreover, advances in sensor technology and detection algorithms allow for real-time continuous detection of activities with battery-powered devices of small form factor to be used in daily life. In particular, the state-of-the-art algorithms used for HAR using wearable inertial sensors utilize machine learning algorithms to detect various biomarkers and to provide real-time and continuous clinical interventions based on the sensor data.

Recent studies have found that machine learning systems are often vulnerable to adversarial perturbations. Even the addition of a small amount of carefully computed perturbations to the clean samples degrade the performance of machine learning systems significantly [5, 14, 23, 27]. What distinguishes adversarial perturbations from random noise is that adversarial examples are misclassified far more often than samples that have been perturbed by random noise, even if the magnitude of noise is much larger compared to the adversarial perturbation [27]. Adversarial inputs were first formally described in [11], in which the researchers studied the techniques used by spammers to circumvent spam filters. In recent years, research on adversarial examples has witnessed tremendous growth and in particular have been investigated in great details in computer vision and audio processing researches [8, 14, 17].

Human activity recognition shares many challenges common to other areas such as computer vision and natural language processing but it also has its own unique challenges and requires dedicated set of computational methods [7]. For example, the data used for the classification of activities often comes from a collection of heterogeneous sensors with different characteristics unlike fixed and well-defined sensing modules such as cameras and microphones. Furthermore, unlike other areas where the problem is well defined (e.g., "Is this an image of a dog?" or "Is the word Harry present in this sentence?"), HAR algorithms need to deal with new challenges due to (1) complexity of human movements captured from individual body joints and observed partially with wearable sensors; and (2) lack of concrete definitions, languages or structure of human activities [7]. For starter, human activity is highly complex and diverse and the sensor readings for an activity can be very different even if the activity is performed by the same person under similar conditions compared to, for example, image classification where an image of a dog is always a dog independent of the presentation and context.

All this leads to the conclusion that HAR though sharing many common points with the other research areas also have unique sets of problems and challenges with socially unique outcomes. In

fact, the issue of adversarial examples in HAR is mostly an unexplored topic despite the great magnitude of the consequences of adversarial machine learning in behavioral medicine. For example, imagine that someone falls and the detection algorithm fails to detect this event or the system used to estimate clinical measures such as activity level, eating, medication, etc. is underestimating or overestimating due to an adversary. These events can have fatal outcomes and therefore requires that we define and assess the effects of adversarial examples in mobile health and behavioral medicine. In this paper, we study the consequences of adversarial examples for activity recognition models and propose an optimization-driven framework called Adar for generating adversarial examples in activity recognition systems. The contributions of this paper are as follows.

- We introduce Adar framework for adversarial examples generation for human activity recognition systems.
- We successfully investigate optimization methods of generating adversarial examples in the context of HAR system.
- We study the vulnerability of activity recognition models to adversarial examples with a comprehensive set of experiments.
- We compare the robustness of a HAR system in feature and signal domain. We also present the result after adversarial training of an activity recognition model.

Note that most of the study focusing on adversarial examples stems from the field of computer vision and speech recognition and it remains to be shown the effect of adversarial examples in other fields such as ubiquitous computing and mobile health. In general, machine learning models that are both highly accurate and robust to adversarial examples remains an open research problem to the research community.

The rest of the paper is structured as follows: In section 2 we explain the Adar framework and we also briefly review the human activity recognition pipeline and explain different adversarial attack methods and defense strategies. Following it, in section 3 we describe our experiments and results.

## 2 ADAR FRAMEWORK DESIGN

Before discussing the details of the Adar framework, we will briefly discuss human activity recognition pipeline and different types of adversarial attacks.

### 2.1 Human Activity Recognition Pipeline

The problem of human activity recognition can be defined as: Given a set  $W = \{W_0, \dots, W_{m-1}\}$  of  $m$  equally sized temporal window of sensor readings, such that each window  $W_i$  contains a set of sensor reading  $S = \{S_{i,0}, \dots, S_{i,k-1}\}$ , and a set  $A = \{a_0, \dots, a_n - 1\}$  of  $n$  activity labels, the goal is to find a mapping function  $f : S_i \rightarrow A$  that can be evaluated for all possible values of  $S_i$  [19]. As shown in Figure 1 the activity recognition system generally consists of sensing, signal processing, signal segmentation, feature extraction and selection, and classification stages [20]. Raw data from various sensors, such as accelerometer, gyroscope, magnetometer, etc. are collected and passed into the signal processing stage, where filtering, noise removal, etc. are applied to the sensor signals. This is followed by a segmentation stage, where a continuous stream

of a signal is divided into temporal windows. For segmentation the sliding window method is used most widely due to its simplicity and real-time performance. After segmentation, statistical and structural features are extracted from each window segment. In feature extraction, the size of the sliding window has significant effects on the accuracy and speed of the system. Larger window size can support the detection of complex activities but result in slower processing whereas smaller window supports faster processing at the expense of accuracy [12]. Usually, a window size of 1-2 second with 50% overlap between the successive windows is considered a good choice[3]. Following the feature extraction, feature selection may be used to select out the best features from the collection of features. Finally, a machine learning classifier is trained on the feature data and the trained classifier is used to make the inference on the future unseen data. Recently, with the success of convolutional neural networks (CNN) several works have been published [4, 15] where a CNN is used for activity detection. What separates the CNN approach from other machine learning approaches is that CNN's completely remove the need to compute features from raw sensor signal for classification. The input to the CNN model is usually the raw sensor signal with or without segmentation, and the CNN model learns the features and the classifier simultaneously during the training process.

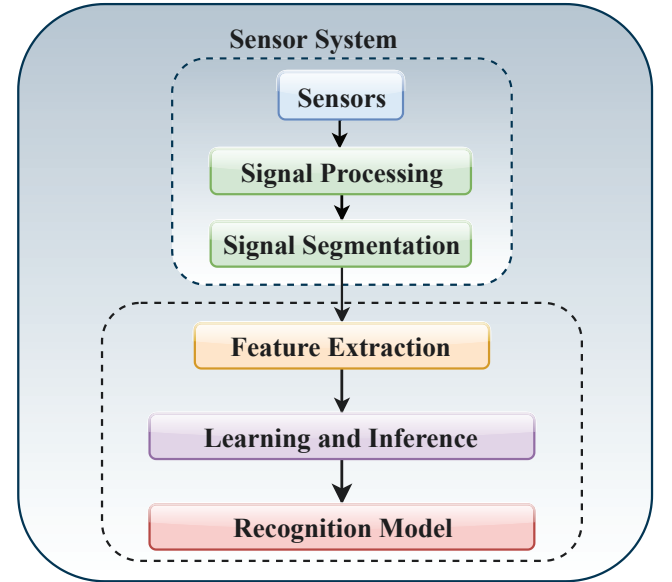


Figure 1: The general framework of human activity recognition system.

### 2.2 Adversarial Attacks

Adversarial examples are inputs formed by applying small but intentional perturbation to the inputs from the dataset such that the perturbed inputs are almost indistinguishable from the true inputs and results in the model outputting an incorrect answer with high confidence [14]. The objective of adversarial learning is to find a perturbation  $\delta$  which when added to the true inputs  $X$  i.e.,  $X^{adv} = X + \delta$  changes the output of the model. The noise level of the perturbation is constrained by the  $l_\infty$  norm denoted by  $\epsilon$  such that the added perturbation is sufficiently small. In general, an adversary can attack a machine learning system in three ways:

- (1) Confidence reduction or poisoning attack: In confidence reduction or poisoning the adversary attempts to degrade the system performance by inducing false connection during training process by corrupting the training data.
- (2) Evasion attack: There are two types of evasion attack: (1) Misclassification attack, and (2) Source/target misclassification attack. In misclassification attack, the adversary computes adversarial examples which are intended to be classified into any class other than the true class. This is also called untargeted attack. In source/target misclassification attack the adversary defines the target class in which it intends to have the model misclassify an input belonging to any class. This is also called targeted attack.
- (3) Exploratory attack: In exploratory attack, an adversary tries to gain as much knowledge as possible about the learning algorithm of the target system and the pattern in its training data.

Furthermore, the strength of an adversary is characterized by its ability to operate either in a white-box setting or in a black-box setting [13]. An adversary with access to the target system's architecture, parameters, and data is said to operate in a white-box setting and an adversary with no knowledge of the target system is said to operate in a black-box setting. A large number of evasion attack methods has been proposed that can be used to compute adversarial examples. In our experiments, we have used the simplest attack methods because of their low computational cost. Now, we will briefly explain these attack methods by using the following notations.

- $X$ , is a true input taken from the dataset. In our case  $X$  can be either a feature vector or a window of raw sensor values.
- $X^{adv}$ , is an adversarial sample. Its type will be the same as the type of  $X$ .
- $y_{true}$ , is the true class label for the input  $X$ .
- $J_\theta(X, y)$ , is the cross-entropy cost function of the neural network with parameters  $\theta$ , given input  $X$  and label  $y$ . Note that after training the network, the parameters of the network remains fixed.
- $Clip_{X, \epsilon}(A)$ , denotes the element-wise clipping of  $A$ , with  $A_{i,j}$  clipped in the range of  $[X_{i,j} - \epsilon, X_{i,j} + \epsilon]$ .

**2.2.1 Fast gradient sign method (FGSM).** Fast gradient sign method is one of the simplest and computationally efficient technique of generating adversarial examples. Proposed by Goodfellow et al. [14], this method is motivated by the linearizing of the cost function and solving for perturbations bounded by  $\epsilon$  that maximizes the cost subject to the  $L_\infty$  norm.

$$X^{adv} = X + \epsilon \text{sign}(\nabla_X J_\theta(X, y_{true})) \quad (1)$$

**2.2.2 One step target class method.** One step target class method tries to maximize the probability  $p(y_{target}|X)$  of some specific target class  $y_{target}$ , which is unlikely to be the true class for the input  $X$ . This is different from the FGSM, in which we try to find adversarial perturbations that increase loss for all classes[17].

$$X^{adv} = X - \epsilon \text{sign}(\nabla_X J_\theta(X, y_{target})) \quad (2)$$

**2.2.3 Basic iterative method.** This method is a simple extension of the fast gradient sign method and applies FGSM multiple times

with small step size.

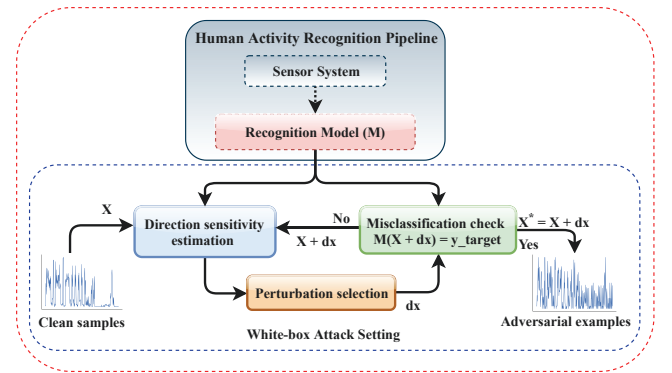
$$\begin{aligned} X_0^{adv} &= X, \\ X_{N+1}^{adv} &= Clip_{X, \epsilon} \{X_N^{adv} + \alpha \text{sign}(\nabla_X J_\theta(X, y_{true}))\} \end{aligned} \quad (3)$$

$\alpha$  is the step size and is set to  $\epsilon/\text{number of iterations}$  in our experiments. The number of iterations is set to 10 for all iterative methods.

**2.2.4 Iterative least-likely class method.** Similar to the basic iterative method which applies FGSM multiple times the iterative least-likely class method applies the one step target class method multiple times with target class set to the least likely class for the given input.

$$\begin{aligned} X_0^{adv} &= X, \\ X_{N+1}^{adv} &= Clip_{X, \epsilon} \{X_N^{adv} - \alpha \text{sign}(\nabla_X J_\theta(X, \end{aligned} \quad (4)$$

## 2.3 Generating Adversarial Examples



**Figure 2: Adar framework for adversarial examples generation in human activity recognition systems.**

Figure 2 shows the proposed Adar framework for computing adversarial examples in human activity recognition systems. This is a general framework and is based on the white-box attack setting, in which an adversary has complete knowledge of the target system. The framework has two components: (1) the human activity recognition pipeline, and (2) the gradient-based attack in a white-box setting. After the HAR pipeline, the trained model is used to compute adversarial examples using gradient-based attacks methods such as fast gradient sign method, basic iterative method or one step target class method. Direction sensitivity estimation finds dimensions of inputs  $X$  that will produce the desired adversarial behavior with the smallest perturbation. Perturbation selection uses this knowledge to select perturbations  $\delta$  affecting sample  $X$ 's classification [24]. If the resulting sample  $X + \delta$  is misclassified by the model  $M$  into the target class  $y_{target}$ , an adversarial example  $X^*$  has been found. If not, the steps can be repeated with updated input  $X + \delta$ . It has been shown that an adversarial example that was designed to be misclassified by model  $M_1$  is often also misclassified by model  $M_2$  [27]. This is called the transferability property of adversarial examples and it means that we can generate adversarial examples and perform misclassification attack on a system without having any knowledge of the system [23]. Therefore an adversary operating in a black-box setting can attack a machine

learning system by building a substitute model and then computing the adversarial examples using the substitute model. This approach has been successfully demonstrated in [22]. Due to this reason and for the low computational cost of generating adversarial examples the Adar framework is based on the white-box setting.

## 2.4 Defense Strategies

There has been a lots of work to build practical defenses against adversarial examples, but a defense technique that is robust against the entire spectrum of adversarial attack is still an open problem. Most of the current defense methods are not adaptive to all types of attacks and one method may block one kind of attack but leaves another vulnerability open for other kinds of attacks [9]. The existing defense mechanism can be categorized into the following types based on their method of implementation.

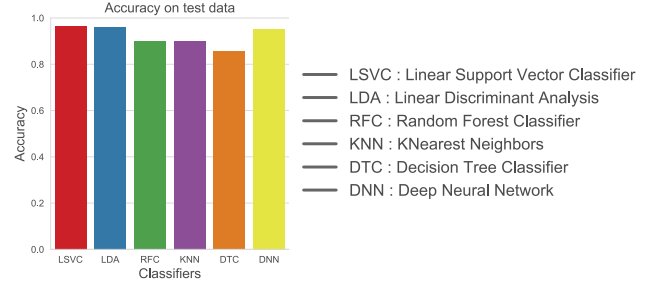
- (1) Adversarial training: The general idea of adversarial training is to augment the training data with adversarial examples during the training process. Adversarial training increases the robustness of the model against white-box attacks when the attack method used to augment the training set is the same as the method used by the attacker for the attack [10]. In a black-box setting and for an attack method different from that used for adversarial training this approach performs poorly.
- (2) Gradient hiding: Gradient hiding is a natural defense against gradient-based attacks, and is based on hiding information about the model's gradient from an adversary. However, this defense is easily fooled by learning a substitute model and using the gradients of the substitute model to craft adversarial examples [9].
- (3) Defensive distillation: Distillation is a way to transfer knowledge from a large neural network to a smaller one. In [24] the authors used distillation to formulate a defense against adversarial attacks. In defensive distillation, the output probability of a neural network  $M$  is used as the label for training a newer network  $M^*$  of same architecture on the same training data.

## 3 EXPERIMENTS AND RESULTS

### 3.1 Dataset

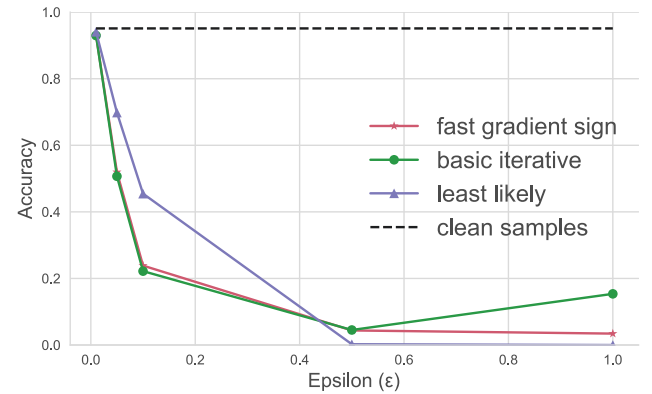
In our experiments, we have used the UCI dataset [2] compiled from a group of 30 participants each wearing a smartphone on their waist and performing 6 different activities: standing, sitting, laying, walking, walking upstairs, and walking downstairs. The sensor data consists of a 3-axial accelerometer and a 3-axial gyroscope sampled at a frequency of 50 Hz. The sensor readings were filtered to remove noise and then segmented into windows of size 2.56 seconds (i.e. 128 readings per window) with 50% overlap. The acceleration values were separated into gravity and body component using a Butterworth low pass filter with a cutoff frequency of 0.3 Hz. From each window, 561 features was computed from time and frequency domain. There are 10299 samples in the dataset with 7352 training samples and 2947 test samples. The availability of raw sensor reading in the dataset allows us to explore adversarial examples in both feature and signal domains.

## 3.2 Adversarial examples in feature domain



**Figure 3: The accuracy of different classifiers computed on 2947 test samples in the dataset.**

In the first experiment, we trained multiple classifiers on the full (561) features data. Figure 3 shows the accuracy of these classifiers on the test data. A deep neural network (DNN) of architecture (64, 32, 6) with l2-regularization (0.001) on the first and second layer gave the best results in terms of accuracy and complexity. We settled on this configuration after testing multiple DNNs with different configurations. *TensorFlow* [1] was used to build the DNN with training parameters: 50 epoch, mini-batch size of 32, Adam optimizer, and sparse categorical cross-entropy loss. All other classifiers were trained with their default settings using the *sklearn* [25] library. We used the *CleverHans* [21] library to compute adversarial examples in all cases.



**Figure 4: The accuracy of the DNN model for the untargeted attack. The value of  $\epsilon$  varies from 0.01 to 1.0 and for iterative methods the number of iteration was set to 10.**

Figure 4 shows the accuracy of the DNN model at different values of  $\epsilon \in [0.01, 0.05, 0.1, 0.5, 1]$ , for the untargeted attack. With an increase in the value of  $\epsilon$  the success rate of the adversarial examples also increased (as shown by the decrease in accuracy) for all methods. Initially, the least likely method was the least effective but with the increase in  $\epsilon$  it was able to exceed other methods. The fast gradient sign method and the basic iterative method performed almost the same at all values of  $\epsilon$  except at  $\epsilon = 1$ . At  $\epsilon = 1$  the success rate of the basic iterative method decreased compared to the fast gradient sign method. We experimented with different values for the number of iterations and found that with a lower number of iterations, the basic iterative method was able to achieve a similar



success rate as that of fast gradient sign method. This is due to the nature of data, which is normalized between  $[-1.0, 1.0]$ , such that at higher values of  $\epsilon$ , the gradient of the cost function starts overshooting and the iterative method fails to find the optimum perturbations.

For the transferability of attack, we tested the adversarial examples computed for the untargeted attack on the DNN model on all other classifiers. Figure 6, shows the accuracy of these classifiers. We can see the transferability of attack exists and the accuracy of all classifiers decreased with increase in the value of  $\epsilon$ . The k-nearest neighbors (kNN) classifier was found to be the most robust compared to other classifiers in all cases. This is consistent with the observation in [28], where the authors showed that a kNN classifier is usually more robust to adversarial examples and its robustness depends on the properties of the data distribution and on the value of  $k$ . However, it has been shown that with complex attack methods adversarial examples can be computed that can fool a kNN model with high success rate [26].

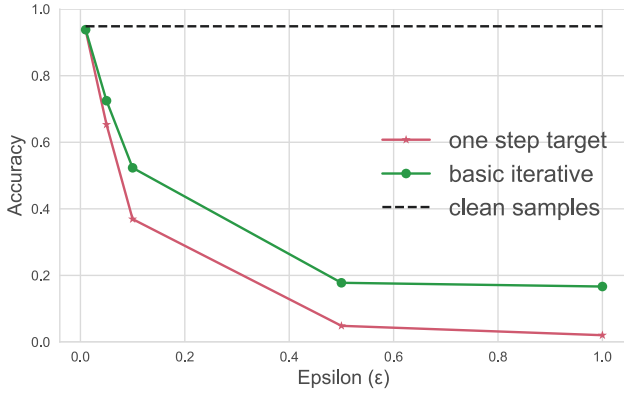


Figure 5: The accuracy of the DNN model for targeted attack. The value of  $\epsilon$  varies from 0.01 to 1.0 and for iterative methods the number of iteration was set to 10.

Figure 5 shows the accuracy of the DNN model at different values of  $\epsilon$  for the targeted attack. One step target class method and the basic iterative method was used to compute the adversarial examples with **Sit** activity class as the target class. Figure 7 shows the confusion matrix of the DNN model on the clean test samples and adversarial examples. In the figure, WalkU and WalkD represent walking upstairs and walking downstairs respectively. It is clear that the one step target class method is more competent at untargeted misclassification attack than the basic iterative method, and the basic iterative method is more successful at targeted misclassification attack. This is due to the fact that one step target method adds  $\epsilon$ -scaled noise to each sample whereas the iterative methods exploit much finer perturbations and therefore is able to find perturbations that cause the classifier to classify an input into a particular class.

### 3.3 Adversarial examples in signal domain

One another approach of building human activity recognition system is by using data-driven methods such as hidden Markov model (HMM), hidden conditional random fields (HCRF), and convolutional neural network (CNN). These methods learn the features

and classifier simultaneously from the raw sensor signal without any prior knowledge of sensor signals and eliminate the need to calculate hand-crafted features for classification. In recent years, methods involving CNN's have shown very good results compared to other methods. This is because the signal representation learned by CNN models capture the local dependency and scale invariance in the signal exceptionally well.

In this section, we will discuss adversarial examples in the signal domain. The UCI dataset has sensor readings from an accelerometer and a gyroscope, making it a multi-modal dataset. For multi-modal data 2D convolutional neural networks are considered to work best because of their ability to capture local dependency among the temporal dimension as well as spatial dimension [16]. For the purpose of simplicity, we have used a 1D kernel in the following experiments and we are confident that similar results will be obtained with a 2D kernel. We have used a ConvNet architecture composed of two 1D convolutional layers with 64, 32 filters and kernel size of 5 and 3 respectively. This is followed by  $3 \times 1$  max-pooling layer and then three dense layers of sizes 64, 32, and 6.

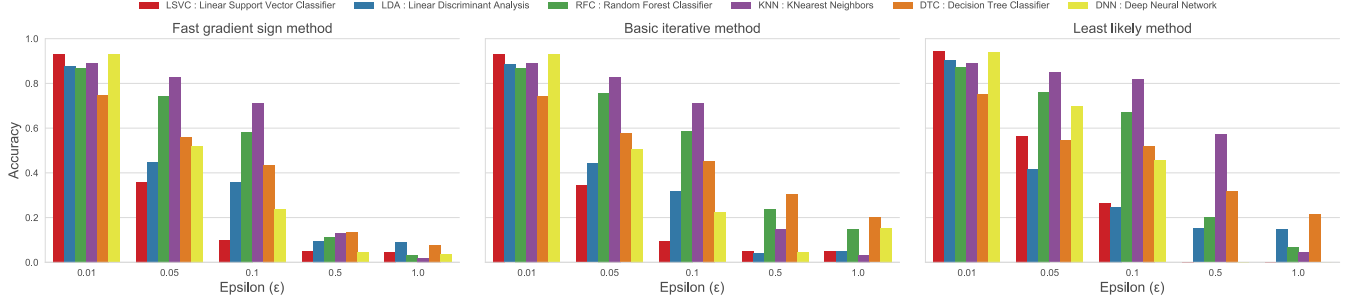
Table 1: The accuracy of the CNN model for the untargeted attack with different attack methods and at different values of  $\epsilon$ . The accuracy was computed on all 2947 test samples in the dataset.

Epsilon ( $\epsilon$ )	Attack methods		
	Fast gradient sign	Basic iterative	Least likely
$\epsilon = 0.01$	90.77%	90.63%	92.09%
$\epsilon = 0.05$	70.44%	64.43%	75.63%
$\epsilon = 0.1$	44.18 %	34.06%	39.02%
$\epsilon = 0.5$	26.63%	3.49%	1.69%
$\epsilon = 1$	16.83%	1.42%	1.01%

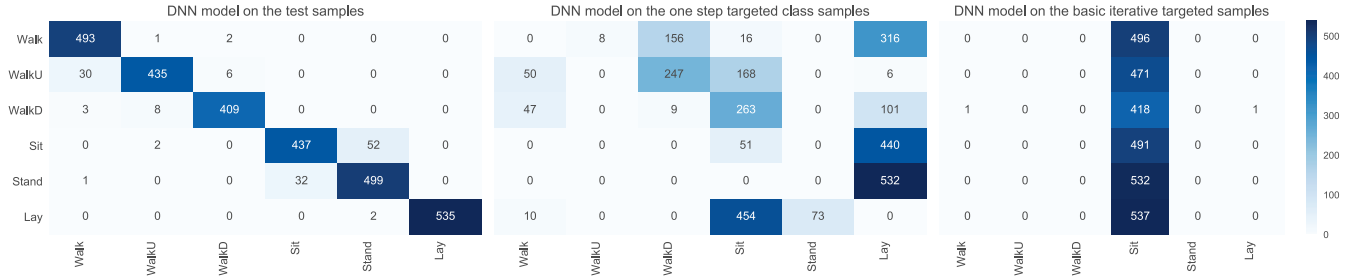
Table 2: The success rate of the targeted adversarial attack on the CNN model at different values of  $\epsilon$ . The values represent the number of samples that were successfully classified into the target class over the total number of samples. The target activity class was **Sit**.

Epsilon ( $\epsilon$ )	Attack methods	
	One step target class	Basic iterative
$\epsilon = 0.01$	17.71%	17.78%
$\epsilon = 0.05$	22.12%	23.31%
$\epsilon = 0.1$	30.50 %	36.81%
$\epsilon = 0.5$	4.64%	95.01%
$\epsilon = 1$	0%	90.19%

Table 1 shows the accuracy of the CNN model for untargeted misclassification attack. With the increase in the value of  $\epsilon$ , the accuracy of the model decreases for all attack methods. The decrease in accuracy for the same value of  $\epsilon$  is largest for the basic iterative method. This is because the basic iterative method is better suited to search the model's input space and find perturbations that can fool the target system with higher confidence. The similar pattern follows for the targeted attack as seen in Table 2. It shows the success rate of the targeted attack which is defined as the ratio of the number of adversarial examples that are classified into the target



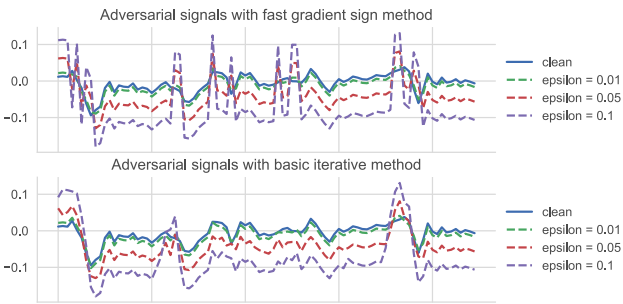
**Figure 6: The accuracy of different classifiers on the adversarial examples computed using fast gradient sign, basic iterative, and least likely methods for the untargeted attack on DNN model. The default settings were used for training all classifiers, except the deep neural network which we described ahead.**



**Figure 7: Confusion matrix of the DNN model on true test samples and targeted adversarial examples computed using one step target class method and basic iterative method at  $\epsilon = 1.0$ . The y-axis represents the true label and the x-axis represents the predicted label.**

class (Sit activity) to the total number of samples. One interesting thing to note here is that at higher values of  $\epsilon$  one step target class method performs poorly compared to at lower values of  $\epsilon$ . This is because at higher values of  $\epsilon$  the input samples are completely destroyed by the addition of  $\epsilon$ -scaled noise to them. Since the adversarial examples computed for the CNN model is in the signal domain, and we call these **adversarial signals**.

Figure 8 shows the x-channel of the body acceleration signal at different values of  $\epsilon$  for fast gradient sign method and basic iterative method. As we can see, the fast gradient sign method is just adding  $\epsilon$ -scaled noise to the clean signal whereas the basic iterative method is following the signal pattern more closely and the added noise preserves the temporal characteristics of the signal.



**Figure 8: Adversarial signal for the x-channel of body acceleration at different values of  $\epsilon$ .**

### 3.4 Adversarial training as a defense

Adversarial training, which augments the training data with adversarial examples, is a well-known defense used to improve the robustness of a machine learning system against adversarial attacks. In this section, we will discuss the effect of adversarial examples on a deep neural network (DNN) trained using adversarial training. We also compare the performance of this model with a baseline mode i.e., a DNN model without adversarial training. We used the algorithm defined in [18] for adversarial training of the DNN model. In all cases we trained the model with a mini-batch size of 32, in which half the samples were clean and the other half were adversarial. It has been shown that if a fixed value of  $\epsilon$  is used to compute adversarial examples during adversarial training then the trained model become robust only to that specific value of  $\epsilon$  [18]. Hence, in our experiments at each epoch a random value of  $\epsilon$ , taken from a truncated normal distribution defined in interval  $[0, 1]$  with underlying normal distribution  $N(\mu = 0, \sigma = 0.5)$  was used to compute adversarial examples. In particular, we tested three cases of adversarial training for the DNN model. In the first case, we used the fast gradient sign method to compute adversarial examples during the adversarial training. After training, the robustness of the trained model was tested on the adversarial examples computed using fast gradient sign method at different values of  $\epsilon$ . Table 3 shows the accuracy of the baseline model and an adversarially trained model. As we can see adversarial training increased the robustness of the baseline model at all values of  $\epsilon$  but also resulted in decrease in accuracy by 13.88% on clean test samples.

**Table 3: Accuracy of the baseline and an adversarially trained models on clean test samples and adversarial examples. Fast gradient sign method was used to compute adversarial examples for both training and testing.**

Epsilon ( $\epsilon$ )	Classifiers	
	Baseline	Adversarial training
clean	95.11%	81.23%
$\epsilon = 0.01$	92.97%	90.12%
$\epsilon = 0.05$	51.95%	57.31%
$\epsilon = 0.1$	23.85%	46.31%
$\epsilon = 0.5$	4.37%	50.69%
$\epsilon = 1$	3.42%	52.49%

**Table 4: Accuracy of the baseline and an adversarially trained models on clean test samples and adversarial examples. One step target class method with the least likely class as the target class was used to compute adversarial examples for both training and testing.**

Epsilon ( $\epsilon$ )	Classifiers	
	Baseline	Adversarial training
clean	95.11%	81.23%
$\epsilon = 0.01$	92.53%	94.70%
$\epsilon = 0.05$	48.18%	69.05%
$\epsilon = 0.1$	22.97%	55.24%
$\epsilon = 0.5$	1.9%	23.04%
$\epsilon = 1$	1.15%	16.38%

**Table 5: Accuracy of the baseline and an adversarially trained models on clean test samples and adversarial examples computed using the basic iterative method at different  $\epsilon$  values. The target classes for the adversarial examples were set to the least likely class. Here the first adversarially trained model  $M1$  uses fast gradient sign method to compute adversarial examples and model  $M2$  uses one step target class method with the least likely class as the target label to compute adversarial examples during adversarial training.**

Epsilon ( $\epsilon$ )	Classifiers		
	Baseline	Adversarial training fgsm	Adversarial training one step target
clean	95.11%	81.23%	81.23%
$\epsilon = 0.01$	94.09%	94.97%	95.38%
$\epsilon = 0.05$	69.83%	85.91%	80.59%
$\epsilon = 0.1$	45.47%	82.21%	67.05%
$\epsilon = 0.5$	0.23%	44.14%	4.3%
$\epsilon = 1$	0%	3.35%	0.61%

In the second case, we used the one step target class method with the least likely class as the target class to compute adversarial examples for both training and testing. Table 4 shows the accuracy of the baseline model and an adversarially trained model. Compared to the adversarial training using the FGSM method, the one step target class method has improved the model performance at smaller values of  $\epsilon$  but at larger values, the performance has decreased significantly.

Finally, in the third case, we tested the robustness of adversarially trained models  $M1$  (trained using fast gradient sign method) and  $M2$  (trained using one step target class method) on the adversarial examples computed using the basic iterative method with least likely class as the target class. Table 5 shows the result. As expected the iterative method was able to degrade the performance of both  $M1$  and  $M2$  severely at higher values of  $\epsilon$ . This shows that iterative methods are more competent at finding adversarial examples that can fool a model with high confidence.

## 4 CONCLUSIONS AND FUTURE WORK

In this paper, we introduced the Adar computational framework for computing adversarial examples for inertial sensor-based activity recognition systems. We investigated different gradient-based evasion attack methods and showed that even computationally simple attack methods can significantly degrade the performance of the activity recognition models. In particular, we found that the accuracy of the deep neural networks trained on the features data decreased from 95.2% to 3.4% and the accuracy of the convolutional neural network trained on the window segments of sensor signal decreased from 93.1% to 16.8%. We also showed the transferability of untargeted attacks in the feature domain and found that the k-nearest neighbor and decision tree classifiers were more robust compared to the other classifiers at all values of  $\epsilon$ . We observed that adversarial training provides robustness to adversarial examples generated using one-step attack methods, but fails against iterative attack methods. In the worst case with adversarial training the robustness of the deep neural networks increased by 49.1% but at the same time suffered a loss in accuracy by 13.2% on clean test samples.

Our results in this paper have extremely important implications in behavioral medicine and mobile health, where detecting human behavior is at the center of clinical interventions. The vulnerability of activity recognition models to adversarial examples not only jeopardize the validity of behavioral interventions that rely on accurate detection of physical and behavioral context such as dietary intake, medication adherence, opioid dependence, smoking behavior, hydration status, and physical activity but also opens up risk points that can have serious impacts on people lives. In this work, we only scratched the surface studying the consequences of adversarial examples in human activity recognition systems, and there remains a need to fully understand and define adversarial attacks for HAR systems in a much greater extent. From the lack of standardization to the nature of the HAR systems, the topic of adversarial example has many facets in activity recognition models. The results of our work also ask many important questions. Should we consider the CNN model more suitable for activity recognition acknowledging the fact that the CNN models were more robust compared to DNN models and also has an advantage of learning features and classifier directly from the sensor data? We also showed the existence of adversarial signals using CNN models. This opens up newer dimensions in the study of adversarial attacks because the ability to generate adversarial examples at the raw signal level implies that many other data processing blocks such as segmentation and pre-processing are also vulnerable to adversarial attacks. This suggests that the study of adversarial attacks in the context of sensor-based

systems need to extend beyond conventional feature-level attacks. What are ramifications of adversarial signals and how can we use the information from the adversarial signals to impact the sensor system remains an open question. Some of the limitations of our work are as follows:

- (1) Since there are no standard datasets and activity recognition models we have carried out our experiments with simple models built from scratch and trained on a single dataset. A follow up on our work can show the effects of adversarial examples on activity recognition models with multiple datasets.
- (2) In our experiments, we have only considered the white-box attack setting. A follow up on our work will be to design adversarial attacks on activity recognition model in the black-box setting.

## ACKNOWLEDGMENTS

This work was supported in part by the United States National Science Foundation under grant CNS-1750679. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding organizations.

## REFERENCES

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <http://tensorflow.org/> Software available from tensorflow.org.
- [2] Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra Perez, and Jorge Luis Reyes Ortiz. 2013. A public domain dataset for human activity recognition using smartphones. In *21th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, ESANN 2013*. 437–442.
- [3] Oresti Banos, Juan-Manuel Galvez, Miguel Damas, Hector Pomares, and Ignacio Rojas. 2014. Window Size Impact in Human Activity Recognition. *Sensors (Basel, Switzerland)* 14, 4 (April 2014), 6474–6499. <https://doi.org/10.3390/s140406474>
- [4] Antonio Bevilacqua, Kyle MacDonald, Aamina Rangarej, Venessa Widjaya, Brian Caulfield, and Tahar Kechadi. 2019. Human Activity Recognition with Convolutional Neural Networks. In *Machine Learning and Knowledge Discovery in Databases*. Vol. 11053. Springer International Publishing, Cham, 541–552. [https://doi.org/10.1007/978-3-030-10997-4\\_33](https://doi.org/10.1007/978-3-030-10997-4_33)
- [5] Battista Biggio, Igino Corona, Davide Maiorca, Blaine Nelson, Nedim Srndic, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. 2013. Evasion Attacks against Machine Learning at Test Time. *arXiv:1708.06131 [cs]* 7908 (2013), 387–402. [https://doi.org/10.1007/978-3-642-40994-3\\_25](https://doi.org/10.1007/978-3-642-40994-3_25)
- [6] Judit Bort-Roig, Nicholas D. Gilson, Anna Puig-Ribera, Ruth S. Contreras, and Stewart G. Trost. 2014. Measuring and influencing physical activity with smartphone technology: a systematic review. *Sports Medicine (Auckland, N.Z.)* 44, 5 (May 2014). <https://doi.org/10.1007/s40279-014-0142-5>
- [7] Andreas Bulling, Ulf Blanke, and Bernt Schiele. 2014. A Tutorial on Human Activity Recognition Using Body-worn Inertial Sensors. *ACM Comput. Surv.* 46, 3 (Jan. 2014), 33:1–33:33. <https://doi.org/10.1145/2499621>
- [8] Nicholas Carlini, Pratyush Mishra, Tavish Vaidya, Yuankai Zhang, Micah Sherr, Clay Shields, David Wagner, and Wenchao Zhou. 2016. Hidden Voice Commands. 513–530. <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/carlini>
- [9] Anirban Chakraborty, Manaar Alam, Vishal Dey, Anupam Chattopadhyay, and Debdeep Mukhopadhyay. 2018. Adversarial Attacks and Defences: A Survey. *arXiv:1810.00069 [cs, stat]* (Sept. 2018). <http://arxiv.org/abs/1810.00069> arXiv: 1810.00069.
- [10] Ting-Jui Chang, Yukun He, and Peng Li. 2018. Efficient Two-Step Adversarial Defense for Deep Neural Networks. *arXiv:1810.03739 [cs, stat]* (Oct. 2018). <http://arxiv.org/abs/1810.03739> arXiv: 1810.03739.
- [11] Nilesh Dalvi, Pedro Domingos, Mausam, Sumit Sanghai, and Deepak Verma. 2004. Adversarial classification. In *Proceedings of the 2004 ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '04*. ACM Press, 99. <https://doi.org/10.1145/1014052.1014066>
- [12] G. Ding, J. Tian, J. Wu, Q. Zhao, and L. Xie. 2018. Energy efficient human activity recognition using wearable sensors. In *2018 IEEE Wireless Communications and Networking Conference Workshops (WCNCW)*. 379–383. <https://doi.org/10.1109/WCNCW.2018.8368980>
- [13] Ian Goodfellow, Patrick McDaniel, and Nicolas Papernot. 2018. Making Machine Learning Robust Against Adversarial Inputs. *Commun. ACM* 61, 7 (June 2018), 56–66. <https://doi.org/10.1145/3134599>
- [14] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *CoRR* (2014). <https://doi.org/abs/1412.6572>
- [15] S. Ha and S. Choi. 2016. Convolutional neural networks for human activity recognition using multiple accelerometer and gyroscope sensors. In *2016 International Joint Conference on Neural Networks (IJCNN)*. 381–388. <https://doi.org/10.1109/IJCNN.2016.7727224>
- [16] S. Ha, J. Yun, and S. Choi. 2015. Multi-modal Convolutional Neural Networks for Activity Recognition. In *2015 IEEE International Conference on Systems, Man, and Cybernetics*. 3017–3022. <https://doi.org/10.1109/SMC.2015.525>
- [17] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. 2016. Adversarial examples in the physical world. *arXiv:1607.02533 [cs, stat]* (July 2016). <http://arxiv.org/abs/1607.02533> arXiv: 1607.02533.
- [18] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. 2016. Adversarial Machine Learning at Scale. *arXiv:1611.01236 [cs, stat]* (Nov. 2016). <http://arxiv.org/abs/1611.01236> arXiv: 1611.01236.
- [19] Oscar D. Lara and Miguel A. Labrador. 2013. A Survey on Human Activity Recognition using Wearable Sensors. *IEEE Communications Surveys & Tutorials* 15 (2013), 1192–1209. <https://doi.org/10.1109/SURV.2012.110112.00192>
- [20] Loc Tan Nguyen. 2017. Enhanced Human Activity Recognition on Smartphone by Using Linear Discrimination Analysis Recursive Feature Elimination Algorithm. In *Context-Aware Systems and Applications*. Springer International Publishing, 72–81.
- [21] Nicolas Papernot, Fartash Faghri, Nicholas Carlini, Ian Goodfellow, Reuben Feinman, Alexey Kurakin, Cihang Xie, Yash Sharma, Tom Brown, Aurko Roy, Alexander Matyasko, Vahid Behzadan, Karen Hambardzumyan, Zhishuai Zhang, Yi-Lin Juang, Zhi Li, Ryan Sheatsley, Abhibhav Garg, Jonathan Uesato, Willi Gierke, Yinpeng Dong, David Berthelot, Paul Hendricks, Jonas Rauber, and Rujun Long. 2018. Technical Report on the CleverHans v2.1.0 Adversarial Examples Library. *arXiv preprint arXiv:1610.00768* (2018).
- [22] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z. Berkay Celik, and Ananthram Swami. 2017. Practical Black-Box Attacks Against Machine Learning. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security (ASIA CCS '17)*. ACM, New York, NY, USA, 506–519. <https://doi.org/10.1145/3052973.3053009> event-place: Abu Dhabi, United Arab Emirates.
- [23] Nicolas Papernot, Patrick McDaniel, and Ian J. Goodfellow. 2016. Transferability in Machine Learning: from Phenomena to Black-Box Attacks using Adversarial Samples. *ArXiv e-prints* (2016). <http://arxiv.org/abs/1605.07277>
- [24] Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. 2015. Distillation as a Defense to Adversarial Perturbations against Deep Neural Networks. *arXiv:1511.04508 [cs, stat]* (Nov. 2015). <http://arxiv.org/abs/1511.04508> arXiv: 1511.04508.
- [25] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [26] Chawin Sitawarin and David Wagner. 2019. On the Robustness of Deep K-Nearest Neighbors. *arXiv:1903.08333 [cs, stat]* (March 2019). <http://arxiv.org/abs/1903.08333> arXiv: 1903.08333.
- [27] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. 2014. Intriguing properties of neural networks. *ICLR* (2014). <https://doi.org/abs/1312.6199>
- [28] Yizhen Wang, Somesh Jha, and Kamalika Chaudhuri. 2017. Analyzing the Robustness of Nearest Neighbors to Adversarial Examples. *arXiv:1706.03922 [cs, stat]* (June 2017). <http://arxiv.org/abs/1706.03922> arXiv: 1706.03922.