

## Investigating the significance of adversarial attacks and their relation to interpretability for radar-based human activity recognition systems

Utku Ozbulak <sup>a,e,1,\*</sup>, Baptist Vandersmissen <sup>a,1</sup>, Azarakhsh Jalalvand <sup>a,d</sup>, Ivo Couckuyt <sup>b</sup>, Arnout Van Messem <sup>c,e</sup>, Wesley De Neve <sup>a,e</sup>

<sup>a</sup> Department of Electronics and Information Systems, Ghent University, Belgium

<sup>b</sup> Department of Information Technology, Ghent University–imec, Belgium

<sup>c</sup> Department of Applied Mathematics, Computer Science and Statistics, Ghent University, Belgium

<sup>d</sup> Department of Mechanical and Aerospace Engineering, Princeton University, USA

<sup>e</sup> Center for Biotech Data Science, Ghent University Global Campus, Republic of Korea

### ARTICLE INFO

#### MSC:

41A05

41A10

65D05

65D17

#### Keywords:

Radar data

Activity recognition

Adversarial examples

Neural network interpretability

Deep convolutional neural networks

### ABSTRACT

Given their substantial success in addressing a wide range of computer vision challenges, Convolutional Neural Networks (CNNs) are increasingly being used in smart home applications, with many of these applications relying on the automatic recognition of human activities. In this context, low-power radar devices have recently gained in popularity as recording sensors, given that the usage of these devices allows mitigating a number of privacy concerns, a key issue when making use of conventional video cameras. Another concern that is often cited when designing smart home applications is the resilience of these applications against cyberattacks. It is, for instance, well-known that the combination of images and CNNs is vulnerable against adversarial examples, mischievous data points that force machine learning models to generate wrong classifications during testing time. In this paper, we investigate the vulnerability of radar-based CNNs to adversarial attacks, and where these radar-based CNNs have been designed to recognize human gestures. Through experiments with four unique threat models, we show that radar-based CNNs are susceptible to both white- and black-box adversarial attacks. We also expose the existence of an extreme adversarial attack case, where it is possible to change the prediction made by the radar-based CNNs by only perturbing the padding of the inputs, without touching the frames where the action itself occurs. Moreover, we observe that gradient-based attacks exercise perturbation not randomly, but on important features of the input data. We highlight these important features by making use of Grad-CAM, a popular neural network interpretability method, hereby showing the connection between adversarial perturbation and prediction interpretability.

### 1. Introduction

Recent advancements in the field of computer vision, natural language processing, and audio analysis enabled the deployment of intelligent systems in homes in the form of assistive technologies. These so-called *smart homes* come with a wide range of functionality such as voice- and gesture-controlled appliances, security systems, and health-related applications. Naturally, multiple sensors are needed in these smart homes to capture the actions performed by household residents and to act upon them.

**Sensors for smart home applications**—Microphones and video cameras are currently two of the most commonly used sensors in smart homes. The research in the domain of video-oriented computer vision is extensive, and the combined usage of a video camera and computer vision enables a wide range of assistive technologies, including

applications related to security (e.g., intruder detection) and applications incorporating gesture-controlled functionalities (Zhang et al., 2019). However, one of the major drawbacks of using video cameras is their privacy intrusiveness (Rajpoot and Jensen, 2015). These privacy-related concerns are, at an increasing rate, being covered by media articles (Staples, 2019). Furthermore, a largely overlooked aspect of video-assisted technologies in smart homes is that video cameras are able to capture both smart home residents and visitors. Therefore, residents of smart homes need to be aware of the statutory restrictions on privacy invasion.

Low-power radar devices, as complementary sensors, are capable of alleviating the privacy concerns raised over the usage of video cameras. In that regard, the main advantages of radar devices over video cameras are as follows: (1) better privacy preservation, (2) a higher effectiveness

\* Corresponding author at: Department of Electronics and Information Systems, Ghent University, Belgium.

E-mail address: [utku.ozbulak@ugent.be](mailto:utku.ozbulak@ugent.be) (U. Ozbulak).

<sup>1</sup> Equal contribution.

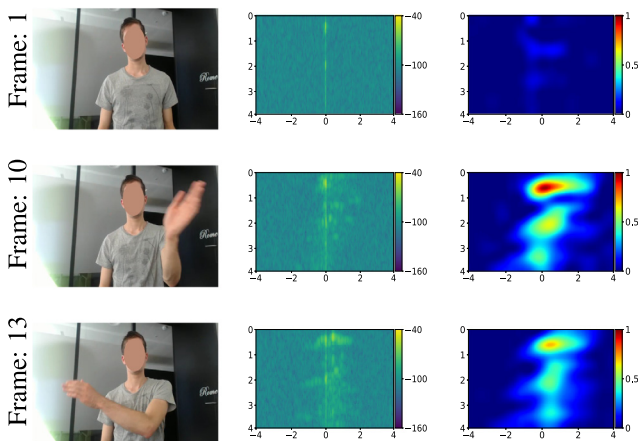


Fig. 1. (1) RD frames, (2) video frames, and (3) Grad-CAM heatmaps for the action *swipe left*. The  $X$ -axis and  $Y$ -axis of the RD frames and the Grad-CAM images, which are omitted for visual clarity, correspond to range and velocity, respectively.

in poor capturing conditions (e.g., low light, presence of smoke), and (3) through-the-wall sensing (Zhao et al.).

Frequency-modulated continuous-wave (FMCW) radars capture the environment by transmitting an electromagnetic signal over a certain line-of-sight. The reflections of this transmitted signal are then picked up by one or more receiving antennas and converted into range-Doppler (RD) and micro-Doppler (MD) frames (Chen et al., 2014). These frames contain velocity and range information about all the objects in the line-of-sight (for the duration of the recording). Recent studies show that with the help of (deep) neural networks, it is possible to leverage these RD and MD frames to recognize multiple individuals (Vandersmissen et al., 2018; Jalalvand et al., 2019) or to detect human activities with high precision (Wang et al., 2016). The aforementioned studies represent these RD and MD frames in the form of a sequence of mono-color images which are supplied as an input to deep CNNs. Three example RD frames and their corresponding video frames for the gesture *swiping left* are given in Fig. 1.

**Convolutional neural networks**—Even though CNNs were applied extensively in the work of LeCun et al. (1998) on the MNIST data set, they only became popular after some revolutionary results were obtained in the study of Krizhevsky et al. (2012) on the ImageNet data set (Russakovsky et al., 2015). That work was further improved by many researchers, showing the efficacy of deep CNNs (Simonyan and Zisserman, 2014; He et al., 2016; Szegedy et al., 2016). One of the main benefits of CNNs is their ability to automatically learn features, thus making it possible to forgo the cumbersome process of having to define hand-engineered expert features. Since CNN architectures are end-to-end differentiable, the features can be learned by an optimization method such as gradient descent. This property allows CNNs to be applied to various types of data beyond images and video sequences, such as text (i.e., natural language processing), speech, and radar data (Graves et al., 2013; Kim and Toomajian, 2016).

Nevertheless, CNNs usually also come with a number of detrimental properties, namely, (1) a high training complexity, (2) difficulty of interpretation, and (3) vulnerability to adversarial attacks.

**Adversarial examples**—The study of Szegedy et al. (2013) showed that the predictions made by CNNs may change drastically with small changes in the pixel values of the input images. Specifically, when the input data at hand are modified with a gradient-based method to maximize the likelihood of another class (i.e., targeted adversarial attacks) or to minimize the likelihood of the initially correct class (i.e., untargeted adversarial attacks), it is possible to create malicious data samples that are called *adversarial examples*. These adversarial examples are shown to exist not only in the digital domain but also in the real world (Kurakin et al., 2016), and are thus recognized as a

major security threat for models that are used in a real environment (that is, an environment where the input is not strictly administered).

Research on adversarial attacks on CNNs gained traction after the seminal studies conducted by Szegedy et al. (2013) and Goodfellow et al. (2014), with Biggio et al. (2013) providing an in-depth discussion of adversarial attacks on machine learning models deployed in security-sensitive applications. Since then, the susceptibility to adversarial attacks is recognized as one of the major drawbacks of deep learning, along with the lack of clear interpretability of these models. Some studies even go so far as to suggest that these two issues should be investigated as a single topic (Ross and Doshi-Velez, 2018; Tao et al., 2018; Etmann et al., 2019).

On the other side of this story, many defense techniques to prevent adversarial attacks have been proposed, only to be found ineffective (Carlini and Wagner, 2017; Athalye et al., 2018). As it currently stands, there are no defense mechanisms available that reliably prevent all adversarial attacks.

**Aim of this study**— Given the extensive research in the machine learning community on techniques to prevent adversarial attacks, we analyze the vulnerability of radar-based CNNs to adversarial examples, with the goal of assessing their significance as a security threat in smart homes. For this analysis, we consider the human activity recognition task presented in Vandersmissen et al. (2019), in which the proposed models were able to identify human gestures with high precision using a low-power FMCW radar. Our analysis of adversarial attacks covers a wide range of scenarios, from white-box attacks, in which the adversary is assumed to have all the knowledge about the underlying system, to localized attacks on radar frames under strict conditions, in which the adversary is assumed to have limited knowledge. Furthermore, we also attempt to analyze the connection between adversarial attacks and neural network interpretability by investigating the connection between prediction, perturbation amount, and Grad-CAM (Selvaraju et al., 2016), a popular deep neural network (DNN) interpretability technique.

This paper is organized as follows. In Section 2, we describe the mathematical notation, the data set, and the machine learning models used. In Section 3, we cover the details of the different threat models considered in this study, followed by a discussion of the experiments performed and the results obtained in Section 4. Next, we provide a number of additional experiments in Section 5, exposing the relation between adversarial examples and model interpretability. In Section 6, we conclude our paper and provide directions for future work.

## 2. Framework and notation

In this section, we outline our mathematical notation, also providing details on the data and the models used in this study.

- $\mathbf{X}$  : an arbitrary RD frame represented as a 3-D tensor (frame count  $\times$  height  $\times$  width), with values in the range of  $[0, 1]$ .
- $\mathbf{y} = g(\theta, \mathbf{X})$  : a classification function that links an input  $\mathbf{X}$  to an output vector  $\mathbf{y}$  of size  $M$ , containing predictions made by a neural network that comes with parameters  $\theta$  and that does not contain a final softmax layer.  $M$  denotes the total number of classes used. The  $k$ th element of this vector is referred to as  $g(\theta, \mathbf{X})_k$ .
- $J(g(\theta, \mathbf{X})_c) = -\log\left(\frac{e^{g(\theta, \mathbf{X})_c}}{\sum_{m=1}^M e^{g(\theta, \mathbf{X})_m}}\right)$  : the cross-entropy (CE) function, which calculates the negative logarithmic loss of the softmax prediction made by a neural network for a class  $c$ .
- $\nabla_x g(\theta, \mathbf{X})$  : the partial derivative of a neural network  $g$  with respect to an input  $\mathbf{X}$ .

**Data**—Our experiments are conducted on a data set of human gestures, containing six different hand-based actions performed in an indoor environment, published with our previous study (Vandersmissen

**Table 1**

An overview of all the activities in the selected data set.

(Class ID) Activity	Samples	Avg. duration (Std.)
(0) Drumming	390	2.92s ( $\pm 0.94$ )
(1) Shaking	360	3.03s ( $\pm 0.97$ )
(2) Swiping Left	436	1.60s ( $\pm 0.27$ )
(3) Swiping Right	384	1.71s ( $\pm 0.31$ )
(4) Thumb Up	409	1.85s ( $\pm 0.37$ )
(5) Thumb Down	368	2.06s ( $\pm 0.42$ )

et al., 2019). The different gestures, along with the number of samples per gesture and their average duration, are listed in Table 1. These activities vary from dynamic and clear movements (e.g., *swiping left*) to static (e.g., *thumbs up*) and subtle (e.g., *drumming*) motions.

The samples are recorded using nine different subjects, with each subject repeating each activity several times, and with each subject performing different activities at different speeds and with different pause intervals. This recording approach results in less generic and more diverse activities, given that the length of the activities is not predetermined, nor is their order. The gestures are performed in front of both a radar sensor and an RGB camera, with both devices recording in a synchronized manner. As shown in Table 1, the extent of time in which each activity is performed differs significantly per activity class. The data set contains 2347 activities in total, with an average duration of 2.16s per activity, thus making it one of the larger radar data sets concerning human actions (Kim and Moon, 2015; Jokanovic et al., 2016; Seyfioğlu et al., 2018).

In order to implement a number of threat scenarios, which are discussed in more detail in Section 3, and in order to work with a scenario that better reflects real-world settings, we apply a data set split different from the random stratified split used in Vandersmissen et al. (2019). Instead, we use a subject-specific split, ensuring that the data of a single subject are only present in either the training, testing, or validation set.

- ( $S_+$ ): A subject-specific split, with the training set consisting of samples obtained from subjects 2, 6, 8, and 9. Samples originating from subjects 1 and 7 are used for the validation set and samples obtained from subjects 3, 4, and 5 are used for the test set. This approach leads to a total of 1050, 572, and 725 samples for the training, testing, and validation set, respectively.
- ( $S_-$ ): This subject-specific split is the opposite of ( $S_+$ ), which means that the training set contains samples obtained from subjects 1, 3, 4, 5, and 7. Likewise, samples obtained from the subjects 2 and 9 are used for the testing set, while subjects 6 and 8 provide samples for the validation set. This approach leads to a total of 1297, 404, and 646 samples for the training, testing, and validation set, respectively.

In line with our previous study (Vandersmissen et al., 2019), we consider a fixed sample length of 50 frames, which matches the average length of the majority of the activity samples. Samples that are shorter than 50 frames are padded with the median RD frame. This median frame is calculated by using all of the samples in the data set in order to acquire a padding frame that does not disturb the prediction (i.e., that is not an out-of-distribution sample). For the samples that possess more than 50 frames, only the middle 50 frames are considered.

**Models**—In this study, we use three substantially different architectures in order to solve the multiclass classification problem of human activity recognition. The first architecture is the 3D-CNN architecture used in Vandersmissen et al. (2019), which we will refer to as  $\mathcal{A}$ . The second architecture is a variant of ResNeXt (Xie et al.), a relatively new architecture that achieved the second place in ILSVCR 2017 (Rusakovsky et al., 2015). The design of this architecture is heavily inspired by VGG (Simonyan and Zisserman, 2014) and ResNet (He et al., 2016). To handle data with a temporal dimension (e.g., radar or video data),

we use a modified version of this architecture (adopted from Hara et al. (2018)). In the remainder of this paper, we will refer to the ResNeXt architecture used as  $\mathcal{R}$ .

Furthermore, a recent trend in the field of activity recognition is the usage of CNN-LSTM architectures (Zhao et al., 2018b; Nair et al., 2018; Yao and Qian, 2018), leveraging the underlying CNN as the feature extractor and employing a Long Short-Term Memory (LSTM) (Hochreiter and Schmidhuber, 1997) layer in order to discover temporal relations. Apart from the usage of the previously explained fully convolutional architectures, we also employ a similarly capable CNN-LSTM architecture ( $\mathcal{L}$ ) in order to discover differences between fully convolutional and CNN-LSTM architectures in terms of adversarial robustness. Architectural details of all models, as well as their performance on the selected dataset, can be found in Section A of the supplementary materials.

With 8.1 million trainable parameters, the employed ResNeXt model ( $\mathcal{R}$ ) is significantly more complex than the 3D-CNN model ( $\mathcal{A}$ ), which contains approximately 647 thousand trainable parameters. Consequently, the size of the two models is also considerably different: ResNeXt occupies about 32.9MB of memory, whereas 3D-CNN only takes about 2.9MB. Although the space occupied by each of the models does not make a significant difference for many of the current commercial products, on the same hardware, the 3D-CNN model is up to 7 times faster than the ResNeXt model in terms of inference speed. When considering edge-computing and real-time applications in the context of smart homes, this means that deploying the ResNeXt model will naturally cost more than deploying the 3D-CNN model.

By employing fully-convolutional models that are significantly different in terms of both architecture and the number of trainable parameters, we are able to study the impact of adversarial examples generated by an *advanced* model on a *simpler* model, and vice versa. Moreover, by evaluating the adversarial examples generated by these fully-convolutional models on the CNN-LSTM model, we are able to analyze the effectiveness of non-LSTM adversarial examples on LSTM architectures.

The accuracies of  $\mathcal{A}$ ,  $\mathcal{R}$ , and  $\mathcal{L}$  are provided in Table A.1 in the supplementary materials, as obtained for the evaluation splits. As can be observed from this table, although the number of trainable parameters is significantly different for each architecture, they achieve comparable results. Note that the models trained on the evaluation splits  $S_{\{-,+\}}$  achieve slightly lower test and validation accuracies than the models presented in our previous work. This can be attributed to the reduced amount of training data we intentionally assigned to these splits, with the goal of covering a wide range of threat models (see Section 3). Throughout this paper, we adopt the notation  $\mathcal{A}_{\text{Dataset split}}$  in order to describe a trained model. For instance,  $\mathcal{A}_{S_+}$  means that the model is of architecture  $\mathcal{A}$  and that this model has been trained on the training set of  $S_+$ . As will be described in the next section, our approach towards selecting models and creating evaluation splits makes it possible to evaluate a wide range of white- and black-box attack scenarios.

### 3. Threat model

In this section, we discuss the threat models evaluated in this paper. To that end, recall that activities performed in smart homes cause either a global response, meaning that the assistance of a third party is required (e.g., calling the police to prevent an intruder from entering a home or calling an ambulance for a health-related emergency situation), or a local, in-house response, meaning that the request of a household resident is related to a functionality confined to the house (e.g., turning on the lights). In this study, we evaluate threat scenarios concerning an adversarial attack to the neural network that is part of the decision making mechanism, which may affect both in-house and out-house functionalities. Naturally, there are also other types of security-related topics that need to be analyzed when a smart home system requests aid from outside the house. Such topics are mainly



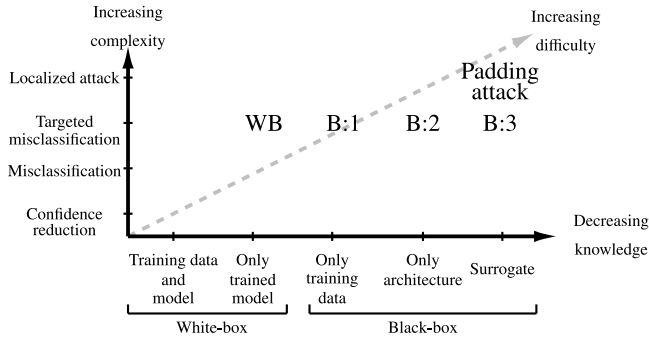


Fig. 2. Threat configurations and their taxonomy for adversarial examples. Scenarios evaluated in this study are highlighted with their abbreviation.

related to home network security and are deemed out of scope for this study.

Given the context described above, an activity and a corresponding flow of events taking place in a smart home environment are visualized in Fig. A.4 in the supplementary materials, as well as in the graphical abstract of this paper. We consider three entry points for a possible adversarial attack: (1) when the radar frames are generated, (2) when radar frames are being transferred from the detector to the on-site server, and (3) right before the inference phase, where frames are sent as an input to the underlying CNN.

In order to be consistent with past research that studied adversarial examples, we use a taxonomy similar to the one outlined in Papernot et al. (2015). In this context, the two main aspects of a threat model are (1) the knowledge of the adversary about the underlying system and (2) the complexity of the attack. Multiple levels of (1) and (2) are shown in Fig. 2, with the labels within the figure also illustrating the different scenarios assessed in this paper. Among all possible combinations of the attacks listed in Fig. 2, we only evaluate the two most restrictive cases, which are targeted misclassification and targeted misclassification with a localized attack. A detailed description of all attacks can be found in the supplementary materials.

Based on the different types of attacks and the level of knowledge of the adversary about the underlying system, we evaluate the following scenarios as threat models:

- **White-box threat model (WB):** The adversary has access to the underlying trained model (including the trained weights) that performs the classification when a radar activity is performed (e.g., adversarial examples generated by  $\mathcal{A}_{S_+}$  and tested on  $\mathcal{A}_{S_+}$ ).
- **Black-box threat models:** The adversary does not have access to the underlying trained model (including the trained weights), but the adversary does have access to the following specifications of the underlying decision-making system:
  - Scenario 1 (B:1): The adversary has access to (1) the architecture of the underlying model that performs the classification (without the trained weights) and (2) data from a similar distribution that the underlying model was trained with (e.g., adversarial examples generated by  $\mathcal{A}_{S_+}$  and tested on  $\mathcal{A}_{S_-}$ ).
  - Scenario 2 (B:2): The adversary has access to the training data used to train the underlying model, but not to the exact specifics of this model such as the weights, layers, and nodes (e.g., adversarial examples generated by  $\mathcal{A}_{S_+}$  and tested on  $\mathcal{R}_{S_+}$ ).
  - Scenario 3 (B:3): The adversary neither has access to the underlying model nor the training data used to train this model. However, the adversary has access to data obtained from a similar distribution and a model that is akin to the underlying model (e.g., adversarial examples generated by  $\mathcal{A}_{S_+}$  and tested on  $\mathcal{R}_{S_-}$ ).

In the upcoming section, we evaluate the threat models described above using various adversarial attacks.

#### 4. Threat model evaluation

In this section, we analyze the robustness of radar-based CNNs against adversarial examples using the threat models described in Section 3. To that end, we investigate the significance of commonly used adversarial attacks, as well as additional attacks that are only possible because of the existence of the temporal domain in the employed input.

##### 4.1. Evaluating common adversarial attacks for radar data

Since the inception of adversarial attacks against machine learning models, a wide range of attack methods has been proposed in the literature, with each attack method focusing on a different aspect of the adversarial optimization. Fast Gradient Sign (FGS) (Goodfellow et al., 2014) is a fast way of generating adversarial examples without any iteration. Adversarial Patch (Brown et al., 2017) produces visible patches that, when added to the input, reliably change the prediction of the model. Basic Iterative Method (BIM) (Kurakin et al., 2016) is an extension of FGS that generates adversarial examples in an iterative manner. Universal Perturbation, as proposed by Moosavi-Dezfooli et al. (2017), shows that it is possible to generate adversarial examples with a pre-selected perturbation pattern. Finally, the Carlini–Wagner Attack (CW) (Carlini and Wagner, 2016) produces *durable* adversarial examples that are resistant to defense systems. Next to these, there are also other black-box attacks that do not use a surrogate model to produce adversarial examples (Ilyas et al., 2018; Guo et al., 2019).

Recently published studies typically aim for improving the *strength* of the produced adversarial examples, making it easier to evade deployed defense systems (Madry et al., 2017). In the literature, BIM and CW are often selected as methods to evaluate newly proposed defense mechanisms against adversarial examples, given the diversity of their properties. Similarly, we use BIM and CW in order to investigate the vulnerability of radar-based CNNs.

**Basic Iterative Method**—This attack is an extension of FGS, which uses the signature of the cross-entropy loss in order to generate adversarial examples. On the one hand, BIM is seen as a method for generating *weak* adversarial examples that heavily perturb the input. On the other hand, BIM is highly efficient because it generates adversarial examples much faster than CW. It is defined as follows:

$$\mathbf{X}_{n+1} = \text{Clip}_{\mathbf{X}, \epsilon}(\mathbf{X}_n - \alpha \text{sign}(\nabla_{\mathbf{x}} J(g(\theta, \mathbf{X}_n), c))), \quad (1)$$

where the *Clip* function ensures that the adversarial example  $\mathbf{X}_{n+1}$  is a valid input (i.e., an image) and where  $\alpha$  denotes the perturbation multiplier. In this study, we use  $\alpha = 15 \times 10^{-4}$ , meaning that a single iteration of perturbation will change values by half a pixel value (i.e., 0.5/255), reporting results based on the use of this perturbation multiplier.

**Carlini–Wagner Attack**—The Carlini–Wagner attack is proposed as a method to generate *strong* adversarial examples (Carlini and Wagner, 2016). This attack uses multi-target optimization and maximizes the prediction likelihood of both the target class and second-most-likely class in order to deceive the underlying machine learning model. CW is criticized for its high computational complexity, which is primarily due to its extensive search for *strong* perturbations (Goodfellow et al., 2018). It is defined as follows:

$$\text{mimize } \|\mathbf{X} - (\mathbf{X} + \delta)\|_2^2 + \ell(\mathbf{X} + \delta), \quad (2)$$

$$\ell(\mathbf{X}') = \max(\max\{g(\theta, \mathbf{X}')_i : i \neq c\} - g(\theta, \mathbf{X}')_c, -\kappa) \quad (3)$$

where  $\delta$  is the perturbation added to the image and  $\ell$  a loss function. This equation aims at maximizing the prediction likelihood of the target class  $c$  and the second-most likely class  $i$ , with  $\kappa$  controlling the logit

**Table 2**

Median value (interquartile range) of the  $L_2$  and  $L_\infty$  distances obtained for 1000 adversarial optimizations, as well as their success rate, for the models and data sets described in Section 2. For easier comprehension, the threat models are listed from most permissive to least permissive.  $L_\infty$  values less than 0.003 are rolled up to 0.003 (this is approximately the smallest amount of perturbation required to change a pixel value by 1, so to make discretization possible).

Threat model	Source model	Target model	BIM			CW		
			$L_2$	$L_\infty$	Success %	$L_2$	$L_\infty$	Success %
WB	$\mathcal{A}_{\{s_+,s_-\}}$	$\mathcal{A}_{\{s_+,s_-\}}$	0.79 (0.72)	0.003 (0.003)	100%	0.50 (0.42)	0.02 (0.02)	97%
	$\mathcal{R}_{\{s_+,s_-\}}$	$\mathcal{R}_{\{s_+,s_-\}}$	0.19 (0.11)	0.003 (0.003)	100%	0.17 (0.11)	0.01 (0.02)	99%
BB:1	$\mathcal{A}_{\{s_+,s_-\}}$	$\mathcal{A}_{\{s_+,s_+\}}$	2.27 (3.17)	0.006 (0.011)	79%	0.72 (1.17)	0.02 (0.09)	54%
	$\mathcal{R}_{\{s_+,s_-\}}$	$\mathcal{R}_{\{s_+,s_+\}}$	0.24 (0.08)	0.006 (0.003)	35%	0.20 (0.17)	0.01 (0.03)	31%
BB:2	$\mathcal{A}_{\{s_+,s_-\}}$	$\mathcal{R}_{\{s_+,s_-\}}$	2.24 (4.77)	0.009 (0.012)	35%	0.92 (0.97)	0.03 (0.11)	30%
	$\mathcal{R}_{\{s_+,s_-\}}$	$\mathcal{A}_{\{s_+,s_-\}}$	0.24 (0.12)	0.006 (0.003)	28%	0.22 (0.18)	0.03 (0.02)	27%
BB:3	$\mathcal{A}_{\{s_+,s_-\}}$	$\mathcal{R}_{\{s_+,s_+\}}$	3.12 (5.14)	0.012 (0.07)	30%	0.87 (1.03)	0.03 (0.13)	21%
	$\mathcal{R}_{\{s_+,s_-\}}$	$\mathcal{A}_{\{s_+,s_+\}}$	0.22 (0.10)	0.006 (0.005)	19%	0.21 (0.27)	0.02 (0.04)	18%

difference between both classes. We adhere to the study of Carlini and Wagner (2016) and set  $\kappa = 20$ .

**Constraints on Adversarial Attacks**—During the generation of adversarial examples, when not considering constraints for the generated adversarial examples, (1) the optimization may result in an adversarial example that does not represent a valid input for the targeted neural network or (2) the attack may not be representative of a real-world scenario. In order to avoid such scenarios, we impose a box constraint, a time constraint, and a discretization constraint on the way adversarial examples are generated. A detailed description of these constraints can be found in the supplementary materials.

**Experiments**— In Table 2, we present the experimental results obtained in terms of white-box and black-box transferability success, for the adversarial examples created with BIM and CW. Specifically, Table 2 details the success rate obtained for 1000 adversarial examples that originate from unseen data points by their respective models during training time, as well as median  $L_2$  and  $L_\infty$  distances between adversarial examples and their initial data points, giving an idea of the minimum amount of perturbation necessary to change the prediction of a model by both attacks. Since we use the  $L_2$  and  $L_\infty$  distances between adversarial examples when they transfer successfully from a source model to a target model, we also provide the interquartile range in order to gain insight into the spread of the  $L_2$  and  $L_\infty$  distances. We use median and interquartile range over mean and standard deviation in order to mitigate the influence of outliers when the success rate of the attacks is low. Based on this experiment, we make the following observations:

- Unsurprisingly, the success rate of the adversarial attacks decreases as the knowledge of the adversary on the underlying system decreases. As opposed to this trend, the minimal required perturbation to change the prediction of the target model often increases as the knowledge of the adversary decreases.
- More often than not, adversarial attacks with BIM are more successful than the ones with CW, even though the latter is considered a more advanced attack. Using additional experiments, we observed this can be primarily attributed to the time constraint imposed on the optimization. Since CW is computationally more expensive than BIM, given its extensive search for a minimum amount of perturbation, generating an adversarial example with CW within the time limit imposed becomes challenging.
- Although BIM is more successful in generating adversarial examples than CW, the adversarial examples generated by BIM

come with much stronger perturbations in terms of  $L_2$  distance than those generated by CW. On the other hand, thanks to the  $sign(\cdot)$  function flattening the gradients to an equal level, adversarial examples created with BIM come with much less perturbation in terms of  $L_\infty$  distance. This finding for radar data is also in line with the observations we have made in the image domain (Ozbulak et al., 2020).

- Even though ResNeXt models are able to find adversarial examples with less perturbation, the time limit set on the generation of adversarial examples also affects ResNeXt models more than 3D-CNN models, since it takes them longer to perform a prediction, as well as to calculate the gradients for adversarial example generation, ultimately resulting in lower success rates.
- For all black-box cases, the ResNeXt architecture is able to find adversarial examples with much less perturbation than 3D-CNN. Our initial interpretation of this finding was that the adversarial examples generated from *stronger* models transfer with less perturbation when attacking similar or *weaker* models. However, recent results in the area of adversarial research suggest that residual models that contain skip-connections allow the generation of adversarial examples with much less perturbation (Wu et al., 2020). Our experiments also confirm this observation.
- Experimental results obtained for the CNN-LSTM architecture (presented in Table E.2 in the supplementary materials) show that adversarial examples generated by fully convolutional architectures are capable of adversarial transferability. Moreover, we observed that the CNN-LSTM architecture employed in this study provides no additional security compared to fully convolutional models. A detailed discussion of these results is presented in Section E of the supplementary materials.

Detailed visual examples of the degree of perturbation needed and the perturbation visibility can be found in Fig. F.6 in the supplementary materials.

#### 4.2. Adversarial padding for radar data

Our experiments show that, in the most restrictive case, the success rate of the adversarial attacks falls as low as 18%. The reason behind the low success rate can again be mainly attributed to the time constraint imposed on the attacks. However, an attacker may already be in possession of a pattern of adversariality that is ready to be deployed without needing any additional computation, thus nullifying the time

**Table 3**

Median value (interquartile range) of the  $L_2$  and  $L_\infty$  distances obtained for 1000 adversarial optimizations, as well as their success rate, for the models and data sets described in Section 2, hereby using the padding attack described in Section 4.2.

Source model	Target model	Padding attack		Success %
		$L_2$	$L_\infty$	
$\mathcal{A}_{\{s_r, s_s\}}$	$\mathcal{A}_{\{s_r, s_s\}}$	3.34 (0.56)	0.012 –	84%
$\mathcal{R}_{\{s_r, s_s\}}$	$\mathcal{R}_{\{s_r, s_s\}}$	3.07 (0.28)	0.012 –	74%
$\mathcal{A}_{\{s_r, s_s\}}$	$\mathcal{A}_{\{s_r, s_s\}}$	3.47 (0.82)	0.011 –	68%
$\mathcal{R}_{\{s_r, s_s\}}$	$\mathcal{R}_{\{s_r, s_s\}}$	4.19 (1.07)	0.010 –	51%
$\mathcal{A}_{\{s_r, s_s\}}$	$\mathcal{R}_{\{s_r, s_s\}}$	5.19 (0.95)	0.011 –	48%
$\mathcal{R}_{\{s_r, s_s\}}$	$\mathcal{A}_{\{s_r, s_s\}}$	5.56 (0.87)	0.012 –	28%
$\mathcal{A}_{\{s_r, s_s\}}$	$\mathcal{R}_{\{s_r, s_s\}}$	4.67 (1.21)	0.010 –	41%
$\mathcal{R}_{\{s_r, s_s\}}$	$\mathcal{A}_{\{s_r, s_s\}}$	5.47 (0.80)	0.012 –	25%

constraint set on generating an adversarial example. In the literature, this type of attacks is called *universal perturbations* (Moosavi-Dezfooli et al., 2017). The main focus of these studies is to generate a universal perturbation pattern in advance and use it during inference time.

Moosavi-Dezfooli et al. (2017) demonstrated that universal perturbations exist for DNNs and that these universal perturbations may even be diverse in nature (i.e., more than one universal perturbation may be available). Instead of evaluating the techniques proposed in Moosavi-Dezfooli et al. (2017), which allow for finding a universal perturbation eventually, we experiment with the idea of generating adversarial padding. This approach can then be employed for samples in the data set that contain less than 50 frames, facilitating prediction without requiring additional optimization. In other words, adversarial padding allows an adversary to change the prediction without even having to modify the frames in which the activity of interest takes place. Inspired by the work of Moosavi-Dezfooli et al. (2017), we use the approach described below to generate the adversarial padding.

Let  $\mathbf{P} \in [0, 1]^{80 \times 126}$  be the median padding used in this study, as described in Section 2, and let  $\mathbf{X} \in [0, 1]^{\ell \times 80 \times 126}$  be an  $\ell$ -frame-long activity for some  $\ell \in \{1, \dots, 50\}$ . We define  $\mathbf{M}_1 = [\mathbf{P} \dots \mathbf{P}] \in [0, 1]^{50 \times 80 \times 126}$  as the initial adversarial padding pattern made up of multiple  $\mathbf{P}$ s and denote by  $\mathbf{M}^i$ ,  $i \in \{1, \dots, 50\}$ , the  $i$ th frame of  $\mathbf{M}$ . In order to generate the adversarial padding, we use the following approach:

$$\tilde{\mathbf{X}} = [\underbrace{\mathbf{X} \mathbf{P} \dots \mathbf{P}}_{50-\ell}], \quad \hat{\mathbf{X}} = [\underbrace{\mathbf{X} \mathbf{M}^{\ell+1} \dots \mathbf{M}^{50}}_{50-\ell}], \quad (4)$$

$$\text{minimize } \|\tilde{\mathbf{X}} - \hat{\mathbf{X}}\|_2, \quad (5)$$

$$\text{such that } \arg \max(g(\theta, \tilde{\mathbf{X}})) = c, \quad \arg \max(g(\theta, \hat{\mathbf{X}})) \neq c \quad (6)$$

where  $c$  is the target class. We then calculate  $\mathbf{M}$  in an iterative manner as follows:

$$\mathbf{M}_{n+1} = \mathbf{M}_n + (\alpha \nabla_x (g(\theta, [\mathbf{X} \mathbf{M}_n^{\ell+1} \dots \mathbf{M}_n^{50}]))_c \odot \mathbf{K}), \quad (7)$$

$$\mathbf{K}^i = \begin{cases} 1^{80 \times 126}, & \text{if } i \in \{i_1, i_2, i_3\} \\ 0^{80 \times 126}, & \text{otherwise} \end{cases}, \quad (8)$$

$$\text{with } \|\mathbf{M}_n^{i_1, i_2, i_3} - \mathbf{P}\|_2 < \|\mathbf{M}_n^u - \mathbf{P}\|_2, \quad (9)$$

$$\forall u \in \{1, \dots, 50\} \setminus \{i_1, i_2, i_3\}, \quad (10)$$

where  $\odot$  is the element-wise tensor multiplication. The idea behind using  $\mathbf{K}$  is to force the optimization to modify the least modified padding frames (in this case, the three frames  $i_1$ ,  $i_2$ , and  $i_3$ ). This leads

to a more uniform distribution of the perturbation over the padding frames, rather than having perturbation that is concentrated in just a few frames. Indeed, we observed that the optimization focuses on just a few frames rather than all frames when  $\mathbf{K}$  is not incorporated, resulting in an adversarial example that is not able to reliably change the prediction. We also experimented with selecting more than three frames: although results were comparable, selecting three frames produced the best results in terms of adversarial transferability.

By following the aforementioned procedure, we are able to create padding sequences that convert model predictions to the targeted class, without even having to change the frames in which the activity occurs. In order to demonstrate the effectiveness of this attack, we provide Table 3, containing details about the success rate of the padding attack, obtained under the same conditions as the results presented in Table 2. Note that the  $L_\infty$  interquartile range is not presented in Table 3 because the padding attack aims at spreading the adversarial perturbation equally over all padding frames, thus keeping the  $L_\infty$  distance between genuine data points and their adversarial counterparts for multiple data points approximately the same.

Given Table 3 in the main text and Table E.3 in the supplementary materials, the first observation is that the padding attack cannot achieve a success rate upwards of 80% for white-box cases. This is because certain samples, which usually belong to activities 0 and 1, are either not padded or padded with very few frames (see Table 1 for average activity duration). Thus, it is very challenging, or downright impossible, for the padding attack to change the prediction. Trivially, the shorter activities are more affected by the padding attack. Furthermore, we can observe that the padding attack achieves higher success rates in most black-box cases than the attacks presented in Table 2, albeit by incorporating stronger perturbations.

Our experiments show that it is indeed possible to exploit the structure of data sets that contain a temporal dimension with special attacks similar to the above-described padding attack. In this case, we demonstrated the possibility of changing a model prediction by only perturbing the frames where the activity does not take place. Moreover, the adversarial padding generated by our padding attack only needs to be computed once and can then be used multiple times, thus allowing it to be incorporated in scenarios where the attacker has limited time for performing a malicious attack. In the supplementary materials, we provide a detailed illustration of adversarial padding in Fig. F.7, showing how remarkably hard it is to spot adversarial padding using the bare eye.

In the next section, we discuss a number of interesting observations related to model interpretability, as made during our analysis of adversarial attacks on radar-based CNNs.

## 5. Relation of adversarial attacks to interpretability

A major criticism regarding DNNs is their lack of interpretability; it is often challenging (if not impossible) to understand the reasoning behind the decisions made by a neural network-based model. In order to overcome this issue and to increase the trustworthiness of DNNs, several techniques have been proposed. These can broadly be divided into the following two groups: (1) perturbation-based forward propagation methods (Zeiler and Fergus, 2014; Shrikumar et al., 2017) and (2) back-propagation-based approaches (Simonyan et al., 2014; Zhou et al., 2016; Selvaraju et al., 2016). The main goal of these techniques is to highlight those parts of the input that are *important* for the prediction made by a neural network. When the input consists of a natural image, this analysis is often done subjectively, unless the evaluated data comes with, for example, weakly-supervised localization labels, which can then be used for evaluating the correctness of the selected interpretability technique. In our case, different from natural images, the input consists of a sequence of RD frames that are significantly harder to interpret for humans. However, different from prediction using a single image, radar data also bring useful features, such as allowing for a frame-by-frame analysis.

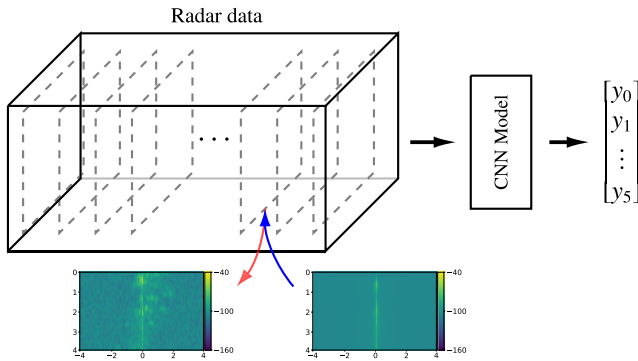


Fig. 3. Visual representation of the frame replacement operation as explained in Section 5.




Fig. 4. A boxplot representation of added perturbation, as generated by CW, displayed for individual frames of adversarial examples that transfer from  $A_S$  to  $A_S$ . The amount of added perturbation is plotted against the median frame importance, as calculated by the experiment detailed in Section 5.




Fig. 5. Mean Grad-CAM magnitude (normalized) is plotted against the mean frame importance, as calculated by the experiment detailed in Section 5.

A first peculiar observation we made during the experiments presented in Section 4 is that CW focuses on only introducing perturbation in certain frames, rather than spreading out the perturbation equally. In Fig. 4, we present the amount of perturbation added by CW to each frame in the form of boxplots. Note that padding frames at the end receive considerably less perturbation than the frames containing the action. We hypothesize that frames that are the recipient of stronger perturbations are *important* frames, making it possible to distinguish actions from one another.

In order to confirm this hypothesis, we perform an exhaustive experiment on measuring the importance of a frame. As illustrated in Fig. 3, we replace individual frames, one at a time, by the median frame we described in Section 2, subsequently performing a forward pass. Since this median frame is used throughout the training procedure to pad the data, it is not an out-of-distribution sample, thus not favoring one class over another. By doing so, for each data point, we measure the change in the prediction logit for the correct class 50 times (for each frame individually) and plot the median difference in Fig. 4, showing the relation between the perturbation amount per frame and the logit change when those frames are replaced. Specifically, the red line represents the median logit change and the shaded area represents the interquartile range. As can be observed, the frames favored by adversarial attacks in terms of added perturbation are also the ones that contribute more to the prediction, confirming our hypothesis. In the supplementary materials, an extended version of this experiment, conducted on each class individually, can be found in Fig. F.8.

Following this experiment, we investigate the applicability of CNN interpretability techniques to radar data. Among different interpretability techniques, Grad-CAM (Selvaraju et al., 2016) stands out, thanks to its superior weakly-supervised localization results obtained on ImageNet. Another reason for selecting this method is that its approach is based on backpropagation, meaning that the input is not perturbed. We especially want to avoid methods based on input perturbation because, unlike natural images, small changes in RD frames may lead to large changes in terms of correctness of the data (i.e., being a valid data point). In our setting, Grad-CAM is defined as follows:

$$\text{Grad-CAM} = \sum_k \left( \text{ReLU} \left( \sum_i \sum_j \nabla_x \mathbf{L}_{i,j}^p \right) \mathbf{L}_k^p \right), \quad (11)$$

where  $\mathbf{L}^p$  denotes the output of the forward pass after the  $p$ th layer (i.e., discriminative features) and  $\nabla_x \mathbf{L}^p$  denotes the gradient obtained with a backward pass from the same layer with respect to the input (i.e., weighted gradient). Different from adversarial attacks, as well as vanilla and guided backpropagation, Grad-CAM does not use the gradients of the first layer, thus arguably allowing for a more robust explanatory approach. Because the input is not a single image but a sequence of frames, Grad-CAM produces class activation maps for each frame individually. An example set of video frames, their corresponding radar frames, and the obtained Grad-CAM heatmaps can be found in Fig. 1. In the supplementary materials, an extended version of the same activity sequence is provided in Fig. F.10. These qualitative results show that the heatmaps usually highlight (1) those frames where the *most important* part of the activity occurs and (2) those locations where the radar activity is the largest.

Apart from the qualitative results provided in Fig. 1, which are heavily criticized in Lipton (2016) and Ghorbani et al. (2017), we now aim at performing a quantitative evaluation of the correctness of the produced Grad-CAM activation frames. We calculate the median magnitude of the produced Grad-CAM frames, which are normalized between 0 and 1, and compare it to the previously presented frame importance data in Fig. 5, where the blue line represents the median Grad-CAM magnitude and the red line the median frame importance. The bands around the lines correspond to the respective interquartile ranges. In the supplementary materials, the same type of illustration on a per-class basis is provided in Fig. F.9.

Given Fig. 5, we again observe a correlation between the importance of frames and their corresponding Grad-CAM activations. Both experiments, as presented in Figs. 4 and 5, show strong correlation between their respective data. In particular, the higher the magnitude of the positive Grad-CAM heatmap, the larger the change in the prediction will be when replacing the underlying frame with the padding frame. Consequently, our experiments confirm that the output of Grad-CAM can indeed be used to assess the relative importance of each radar frame for the prediction made. Indeed, the frames that contribute the most to a prediction are also the ones that are naturally perturbed more than the others during an adversarial optimization, pointing to a strong connection between adversarial optimization and model interpretability.

## 6. Conclusions and future work

In this study, we evaluated multiple scenarios in which adversarial attacks are performed on CNNs trained with a sequence of range-Doppler images obtained from a low-power FMCW radar sensor, with the goal of performing gesture recognition. Our analysis showed that these models are vulnerable not only to commonly used attacks, but also to unique attacks that take advantage of how the data set is crafted. In order to demonstrate a unique attack that leverages knowledge about the data set, we proposed a padding attack that creates a padding sequence that changes the predictions made by CNNs.

An often mentioned drawback of CNNs is their lack of interpretability. By taking advantage of the data selected for this study, we were



able to show the connection between the perturbation exercised by adversarial attacks and the importance of individual frames. Moreover, we were also able to demonstrate that it is possible to identify *important* frames using Grad-CAM, thus showing (1) the relation between adversarial optimization and interpretability, and (2) a quantitative method to evaluate interpretability techniques.

In future research, we aim to analyze multiple shortcomings of radar sensors against so-called real-world adversarial examples (Sun et al., 2018; Kurakin et al., 2016), as well as black-box attacks that do not use surrogate models (Tu et al., 2019; Cheng et al., 2019; Gragnaniello et al., 2019). In the case of activity detection, real-world adversarial examples may occur when radar sensors are employed in environments exhibiting poor recording conditions, such as environments that contain reflective materials (e.g., metal objects), or similarly, when the subject itself carries any reflective material. Moreover, it would be of interest to investigate the influence of multiple moving subjects in the same recording environment on adversariality.

### CRedit authorship contribution statement

**Utku Ozbulak:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing - original draft, Visualization. **Baptist Vandersmissen:** Conceptualization, Methodology, Software, Validation, Investigation, Data curation, Writing - original draft. **Azaraksh Jalalvand:** Data curation, Writing - review & editing, Supervision. **Ivo Couckuyt:** Data curation, Writing - review & editing, Supervision. **Arnout Van Messem:** Formal analysis, Writing - review & editing, Supervision, Project administration. **Wesley De Neve:** Writing - review & editing, Supervision, Project administration.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgments

We would like to thank the anonymous reviewers for their valuable and insightful comments. We believe their comments significantly improved the quality of this manuscript.

The research activities described in this paper were funded by Ghent University Global Campus, Republic of Korea, Ghent University, Belgium, imec, Belgium, Flanders Innovation & Entrepreneurship (VLAIO), Belgium, the Fund for Scientific Research-Flanders (FWO-Flanders), Belgium, and the EU.

### Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.cviu.2020.103111>.

### References

Athalye, A., Carlini, N., Wagner, D., 2018. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. *CoRR abs/1802.00420*.

Biggio, B., Corona, I., Maiorca, D., Nelson, B., Štrdić, P., Giacinto, G., Roli, F., 2013. Evasion attacks against machine learning at test time. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, pp. 387–402.

Brown, T.B., Mané, A., Abadi, M., Gilmer, J., 2017. Adversarial patch. *CoRR abs/1712.09665*.

Carlini, N., Wagner, D.A., 2016. Towards evaluating the robustness of neural networks. *CoRR 1608.04644*.

Carlini, N., Wagner, D.A., 2017. Adversarial examples are not easily detected: Bypassing ten detection methods. *CoRR abs/1705.07263*.

Chen, V.C., Tahmouh, D., Miceli, W.J., 2014. *Radar Micro-Doppler Signatures: Processing and Applications*. Radar, Sonar & Navigation, Institution of Engineering and Technology.

Cheng, S., Dong, Y., Pang, T., Su, H., Zhu, J., 2019. Improving black-box adversarial attacks with a transfer-based prior. In: *Advances in Neural Information Processing Systems*. pp. 10932–10942.

Etmann, C., Lunz, S., Maass, P., Schönlieb, C.B., 2019. On the connection between adversarial robustness and saliency map interpretability. *arXiv preprint arXiv:1905.04172*.

Ghorbani, A., Abid, A., Zou, J., 2017. Interpretation of neural networks is Fragile. *CoRR abs/1710.10547*.

Goodfellow, I., McDaniel, P., Papernot, N., 2018. Making machine learning robust against adversarial inputs. *Commun. ACM 61*, 56–66.

Goodfellow, I., Shlens, J., Szegedy, C., 2014. Explaining and harnessing adversarial examples. *CoRR abs/1412.6572*.

Gragnaniello, D., Marra, F., Poggi, G., Verdoliva, L., 2019. Perceptual quality-preserving black-box attack against deep learning image classifiers. *CoRR abs/1902.07776*.

Graves, A., Mohamed, A.R., Hinton, G., 2013. *Speech recognition with deep recurrent neural networks*. In: *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, pp. 6645–6649.

Guo, C., Gardner, J.R., You, Y., Wilson, A.G., Weinberger, K.Q., 2019. Simple black-box adversarial attacks. *CoRR abs/1905.07121*.

Hara, K., Kataoka, H., Satoh, Y., 2018. Can spatiotemporal 3D CNNs retrace the history of 2D CNNs and ImageNet?. In: *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 6546–6555.

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778.

Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural Comput. 9*, 1735–1780.

Ilyas, A., Engstrom, L., Athalye, A., Lin, J., 2018. Black-box adversarial attacks with limited queries and information. *CoRR abs/1804.08598*.

Jalalvand, A., Vandersmissen, B., D. Neve, W., Mannens, E., 2019. Radar signal processing for human identification by means of reservoir computing networks. In: *IEEE Radar Conference*, pp. 1–6.

Jokanovic, B., Amin, M., Ahmad, F., 2016. Radar fall motion detection using deep learning. In: *2016 IEEE Radar Conference (RadarConf)*. IEEE, pp. 1–6.

Kim, Y., Moon, T., 2015. Human detection and activity classification based on micro-doppler signatures using deep convolutional neural networks. *IEEE Geosci. Remote Sensing Lett. 13*, 8–12.

Kim, Y., Toomajian, B., 2016. Hand gesture recognition using micro-doppler signatures with convolutional neural network. *IEEE Access 4*, 7125–7130.

Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. *Imagenet classification with deep convolutional neural networks*. In: *Advances in Neural Information Processing Systems*. pp. 1097–1105.

Kurakin, A., Goodfellow, I., Bengio, S., 2016. Dversarial examples in the physical world. *CoRR abs/1607.02533*.

LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., 1998. Gradient-based learning applied to document recognition. *Proc. IEEE 86*, 2278–2324.

Lipton, Z.C., 2016. The mythos of model interpretability. *arXiv preprint arXiv:1606.03490*.

Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A., 2017. Towards deep learning models resistant to adversarial attacks. *CoRR abs/1706.06083*.

Moosavi-Dezfooli, S.M., Fawzi, A., Fawzi, O., Frossard, P., 2017. Universal adversarial perturbations. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1765–1773.

Nair, N., Thomas, C., Jayagopi, D.B., 2018. Human activity recognition using temporal convolutional network. In: *Proceedings of the 5th international Workshop on Sensor-based Activity Recognition and Interaction*, pp. 1–8.

Ozbulak, U., Gasparyan, M., D. Neve, W., Va. Messem, A., 2020. Perturbation analysis of gradient-based adversarial attacks. *Pattern Recognit. Lett.*

Papernot, N., McDaniel, P.D., Jha, S., Fredrikson, M., Celik, Z.B., Swami, A., 2015. The limitations of deep learning in adversarial settings. *CoRR abs/1511.07528*.

Rajpoot, Q.M., Jensen, C.D., 2015. *Video surveillance: Privacy issues and legal compliance*. In: *Promoting Social Change and Democracy Through Information Technology*. IGI global, pp. 69–92.

Ross, A.S., Doshi-Velez, F., 2018. Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients. In: *Thirty-second AAAI Conference on Artificial Intelligence*.

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L., 2015. *Imagenet large scale visual recognition challenge*. *Int. J. Comput. Vis. 115*, 211–252.

Selvaraju, R.R., Das, A., Vedantam, R., Cogswell, M., Parikh, D., Batra, D., 2016. Grad-Cam: why did you say that? Visual explanations from deep networks via gradient-based localization. In: *CVPR 2016*.

Seyfioğlu, M.S., Özbayoğlu, A.M., Gürbüz, S.Z., 2018. Deep convolutional autoencoder for radar-based classification of similar aided and unaided human activities. *IEEE Trans. Aerosp. Electron. Syst. 54*, 1709–1723.

Shrikumar, A., Greenside, P., Kundaje, A., 2017. Learning important features through propagating activation differences. *CoRR abs/1704.02685*.

Simonyan, K., Vedaldi, A., Zisserman, A., 2014. Deep inside convolutional networks: Visualising image classification models and saliency maps. In: *Workshop, Proceedings of 2th International Conference on Learning Representations (ICLR)*.



- Simonyan, K., Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition. *CoRR* abs/1409.1556.
- Staples, P., 2019. Thinking about buying a smart home device? heres what you need to know about security. <https://www.forbes.com>. (Accessed 26 July 2019).
- Sun, L., Tan, M., Zhou, Z., 2018. A survey of practical adversarial example attacks. *Cybersecurity* 1, 9.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z., 2016. Rethinking the inception architecture for computer vision. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2818–2826.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R., 2013. Intriguing properties of neural networks. *CoRR* abs/1312.6199.
- Tao, G., Ma, S., Liu, Y., Zhang, X., 2018. Attacks meet interpretability: Attribute-steered detection of adversarial samples. In: *Advances in Neural Information Processing Systems*. pp. 7717–7728.
- Tu, C.C., Ting, P., Chen, P.Y., Liu, S., Zhang, H., Yi, J., Hsieh, C.J., Cheng, S.M., 2019. Autozoom: Autoencoder-based zeroth order optimization method for attacking black-box neural networks. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 742–749.
- Vandersmissen, B., Knudde, N., Jalalvand, A., Couckuyt, I., Bourdoux, A., D. Neve, W., Dhaene, T., 2018. Indoor person identification using a low-power fmcw radar. *IEEE Trans. Geosci. Remote Sens.* 56, 3941–3952.
- Vandersmissen, B., Knudde, N., Jalalvand, A., Couckuyt, I., Dhaene, T., D. Neve, W., 2019. Indoor human activity recognition using high-dimensional sensors and deep neural networks. *Neural Comput. Appl.* 1–15.
- Wang, S., Song, J., Lien, J., Poupyrev, I., Hilliges, O., 2016. Interacting with soli: Exploring fine-grained dynamic gesture recognition in the radio-frequency spectrum. In: *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, pp. 851–860.
- Wu, D., Wang, Y., Xia, S.T., Bailey, J., Ma, X., 2020. Skip connections matter: On the transferability of adversarial examples generated with resnets. In: *International Conference on Learning Representations*.
- Xie, S., Girshick, R., Dollár, Z., He, K., 2017. Aggregated residual transformations for deep neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1492–1500.
- Yao, L., Qian, Y., 2018. DT-3DResNet-LSTM: An architecture for temporal activity recognition in videos. In: *Pacific Rim Conference on Multimedia*. Springer, pp. 622–632.
- Zeiler, M.D., Fergus, R., 2014. Visualizing and understanding convolutional networks. In: *European Conference on Computer Vision*. Springer, pp. 818–833.
- Zhang, H.B., Zhang, Y.X., Zhong, B., Lei, Q., Yang, L., Du, J.X., Chen, D.S., 2019. A comprehensive survey of vision-based human action recognition methods. *Sensors* 19 (1005).
- Zhao, M., Li, T., Ab. Alsheikh, M., Tian, Y., Zhao, H., Torralba, A., Katabi, D., 2018a. Through-Wall human pose estimation using radio signals. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7356–7365.
- Zhao, Y., Yang, R., Chevalier, G., Xu, X., Zhang, Z., 2018b. Deep residual bidir-LSTM for human activity recognition using wearable sensors. *Math. Probl. Eng.* 2018.
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A., 2016. Learning deep features for discriminative localization. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2921–2929.