

Channel-Aware Adversarial Attacks Against Deep Learning-Based Wireless Signal Classifiers

Brian Kim¹, *Graduate Student Member, IEEE*, Yalin E. Sagduyu², *Senior Member, IEEE*,
Kemal Davaslioglu, *Member, IEEE*, Tugba Erpek³, *Member, IEEE*, and Sennur Ulukus⁴, *Fellow, IEEE*

Abstract—This paper presents channel-aware adversarial attacks against deep learning-based wireless signal classifiers. There is a transmitter that transmits signals with different modulation types. A deep neural network is used at each receiver to classify its over-the-air received signals to modulation types. In the meantime, an adversary transmits an adversarial perturbation (subject to a power budget) to fool receivers into making errors in classifying signals that are received as superpositions of transmitted signals and adversarial perturbations. First, these evasion attacks are shown to fail when channels are not considered in designing adversarial perturbations. Then, realistic attacks are presented by considering channel effects from the adversary to each receiver. After showing that a channel-aware attack is selective (i.e., it affects only the receiver whose channel is considered in the perturbation design), a broadcast adversarial attack is presented by crafting a common adversarial perturbation to simultaneously fool classifiers at different receivers. The major vulnerability of modulation classifiers to over-the-air adversarial attacks is shown by accounting for different levels of information available about the channel, the transmitter input, and the classifier model. Finally, a certified defense based on randomized smoothing that augments training data with noise is introduced to make the modulation classifier robust to adversarial perturbations.

Index Terms—Modulation classification, deep learning, adversarial machine learning, adversarial attack, certified defense.

I. INTRODUCTION

BUILT upon recent advances in computational resources and algorithmic designs for deep neural networks, *deep learning* (DL) provides powerful means to learn and adapt to rich data representations [2]. Spectrum data involved in wireless communications is high dimensional due to the underlying

channel, interference and traffic effects, and reflects interactions of network protocols that need to solve complex tasks. Thus, DL has found a broad range of applications in wireless systems, e.g., waveform design [3], spectrum sensing [4], channel estimation [5]–[7], and signal classification [8].

As adversaries can manipulate the training and testing time inputs of machine learning (ML) algorithms, *adversarial ML* has emerged to study the operation of ML models in the presence of adversaries and support safe adoption of ML to the emerging applications [9]. In particular, deep neural networks (DNNs) are known to be highly susceptible to even small-scale adversarial attacks, as extensively demonstrated in the computer vision domain [10]. Shared and broadcast nature of wireless medium increases the potential for adversaries to tamper with DL-based wireless communication tasks over-the-air. However, it is not readily feasible to translate attack and defense mechanisms from other data domains such as computer vision to wireless communications that feature unique characteristics. First, a wireless adversary cannot directly manipulate input data to a DL algorithm running at a separate target and its manipulation needs to reach the target over-the-air through channel effects. Second, a wireless adversary can attack multiple targets (each with a different channel) simultaneously with a single (omnidirectional) transmission (rather than multiple transmissions that would increase the transmit energy consumed and incur more delay) by exploiting the broadcast nature of wireless communications.

In this paper, we present realistic wireless attacks built upon adversarial ML and the corresponding defense by accounting for channel and broadcast transmission effects in the algorithm design for adversarial attacks. We consider the adversarial attack that corresponds to small modifications of the original input to the DNNs that make the DL algorithms misclassify the input. These small modifications are not just random perturbations added as noise (such as in conventional jamming attacks [11]–[13]) but they are carefully designed in a way that changes the decision of the DL algorithm. This attack has been first studied for modulation classifier in [14] by directly manipulating the receiver's input data at the receiver (corresponding to the additive white Gaussian noise (AWGN) channel environment). In this setting, the perturbation has been selected by the adversary according to a power budget without considering channel effects (beyond AWGN channels) on the received perturbation; therefore, this setting corresponds to the best-case scenario for the adversary but it is not practical for realistic channel environments. However, this approach (also followed in subsequent studies [15]–[25])

Manuscript received April 21, 2021; revised September 19, 2021; accepted October 26, 2021. Date of publication November 9, 2021; date of current version June 10, 2022. This work was supported by the U.S. Army Research Office under Contract W911NF-17-C-0090. The content of the information does not necessarily reflect the position or the policy of the U.S. Government and no official endorsement should be inferred. An earlier version of this paper was presented in part at the Proceedings of Conference on Information Sciences and Systems (CISS) 2020 [1] [DOI: 10.1109/CISS48834.2020.1570617416]. The associate editor coordinating the review of this article and approving it for publication was C. Shen. (Corresponding author: Sennur Ulukus.)

Brian Kim and Sennur Ulukus are with the Department of Electrical and Computer Engineering, University of Maryland, College Park, MD 20742 USA (e-mail: bkim628@umd.edu; ulukus@umd.edu).

Yalin E. Sagduyu and Kemal Davaslioglu are with Intelligent Automation, a BlueHalo Company, Rockville, MD 20855 USA (e-mail: ysagduyu@i-a-i.com; kdavaslioglu@i-a-i.com).

Tugba Erpek is with the Hume Center, Virginia Tech, Arlington, VA 24061 USA (e-mail: terpek@vt.edu).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TWC.2021.3124855>.

Digital Object Identifier 10.1109/TWC.2021.3124855

1536-1276 © 2021 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.

See <https://www.ieee.org/publications/rights/index.html> for more information.

cannot necessarily ensure that the adversarial perturbation that reaches the receiver has the necessary power and phase characteristics to cause misclassification at the target classifier under realistic wireless channel conditions. In addition, only a single receiver has been considered in these previous studies, and broadcast attacks on multiple receivers and associated channel effects have not been studied yet. In this paper, we show the need for a channel-aware adversarial attack that can simultaneously work against multiple receivers by showing that: i) the design of adversarial perturbations without taking realistic channel effects into account cannot fool a modulation classifier, and ii) an adversarial attack crafted for the channel to a particular receiver is not effective against another receiver with different channel characteristics, i.e., the attack is receiver specific and cannot be used for a broadcast attack that is launched simultaneously against multiple receivers.

By considering a DNN at each receiver to classify wireless signals to modulation types, we design realistic adversarial attacks in the presence of channel effects from the adversary to each receiver and multiple classifiers at different receivers. For that purpose, we start with a single receiver and determine channel-aware adversarial perturbations (subject to a power constraint) to reduce the accuracy of detecting the modulation type at the receiver. We first propose two white-box attacks, a *targeted attack* with minimum power and a *non-targeted attack*, subject to channel effects known by the adversary. We show that the adversarial attack fails to fool the classifier using limited power if the channel between the adversary and the receiver is not considered when designing the adversarial perturbation. Then, we present algorithms to design adversarial perturbations by accounting for known channel effects and show that the classifier accuracy can be significantly reduced by channel-aware adversarial attacks.

By relaxing the assumption of exact channel information at the adversary, we present a white-box adversarial attack with *limited channel information* available at the adversary. We apply principal component analysis (PCA) to generate adversarial attacks using a lower-dimensional representation of channel distribution. In addition to PCA, we also apply variational autoencoder (VAE) to capture the complexity of underlying data characteristics regarding the transmitter input and the channel. We further relax assumptions regarding the information availability of transmitter input and classifier model, and introduce a *black-box universal adversarial perturbation (UAP) attack* for an adversary with limited information.

We also introduce a *broadcast* adversarial attack to fool multiple classifiers at different receivers with a single perturbation transmission by leveraging the broadcast nature of wireless communications. A practical scenario for this attack is a user authentication system with multiple signal sensors deployed at different locations. First, we show that the channel-aware adversarial perturbation is inherently selective in the sense that it can fool a target classifier at one receiver (whose channel is used in the attack design) but cannot fool a classifier at another receiver due to different channels experienced at different receivers. By considering channels from the adversary to all receivers jointly, we design the broadcast adversarial perturbation that can simultaneously fool

multiple classifiers at different receivers. Using different levels of perturbation budget available relative to noise power, our results illustrate the need to utilize the channel information in designing over-the-air adversarial attacks.

As a countermeasure, we present a *defense* method to reduce the impact of adversarial perturbations on the classifier performance. Following the *randomized smoothing* approach from [26]–[28] (previously applied in computer vision), the training data for modulation classifier is augmented with isotropic Gaussian noise to make the trained model robust to adversarial perturbations in test time. We show that this defense is effective in reducing the impact of adversarial attacks on the classifier performance. We further consider *certified defense* to guarantee classifier robustness of the classifier accuracy against adversarial attacks. The classifier is certified by augmenting the received signals with Gaussian noise samples in test time and checking statistical significance of classification results.

In summary, our contributions are given as follows:

- We propose different realistic wireless attacks against a modulation classifier by taking channel effects into account when generating the adversarial perturbations. Furthermore, we provide how to generate adversarial attacks when the adversary has limited information regarding channel, input at the classifier, and architecture of the target classifier.
- We introduce a wireless broadcast adversarial attack that manipulates multiple classifiers at different receivers with a single perturbation.
- We present a defense method to mitigate the effect of adversarial attacks on the classifier.

The rest of the paper is organized as follows. Section II discusses related work. Section III describes the system model. Sections IV and V present targeted and non-targeted white-box adversarial attacks, respectively, using channel information. Section VI considers the white-box adversarial attack with limited channel information. Section VII presents the UAP attack. Section VIII introduces the broadcast adversarial perturbation in the presence of multiple receivers. Section IX describes the defense method to make the classifier robust to adversarial perturbations. Section X concludes the paper.

II. RELATED WORK

Adversarial ML has been studied in wireless communications [29], [30] with different types of attacks. *Exploratory (inference) attacks* have been considered in [31], [32], where an adversary builds a DNN to learn the transmission pattern in the channel and jams transmissions that would otherwise be successful. Over-the-air spectrum *poisoning (causative) attacks* have been considered in [33], [34], where an adversary falsifies a transmitter's spectrum sensing data over the air by transmitting during the spectrum sensing period of the transmitter. Poisoning attacks have been extended to cooperative spectrum sensing [35], [36] and IoT data aggregation [37]. *Trojan attacks* have been studied in [38] against a modulation classifier, where an adversary slightly manipulates training data by inserting Trojans (in terms of phase shifts) that are later activated in test time. *Membership inference attacks* have been

considered in [39], [40] to learn whether a given data sample has been used during training and thereby gaining information on waveform, channel, and radio hardware characteristics. Adversarial ML has been studied in [41], [42] to launch *spoofing attacks* that aim to fool signal authentication systems based on DNNs. As a viable threat for emerging wireless systems, adversarial ML has been also studied to launch attacks against 5G and beyond communications [43]–[45].

Adversarial attacks (a.k.a *evasion attacks*) introduce small modifications to the input data that go into the DNNs such that the DL algorithm cannot reliably classify the input data. [46] has shown that the end-to-end autoencoder communication systems, proposed in [47], are vulnerable to adversarial attacks that increase the block-error-rate at the receiver. Adversarial attacks have been also considered against spectrum sensing [48], power allocation [49], [50], beam selection [51], and MIMO channel state information (CSI) feedback [52]. One particular application of DNNs that has gained recent interest in wireless communications is signal classification [8], [53]–[59]. Adversarial attacks have been studied against modulation classifiers in [14] and subsequent studies [15]–[23]. The targeted fast gradient method (FGM) attack [60] has been used in [14] by enforcing the DNNs to misclassify the input signals to a target label. Here, the target is decided by searching over all possible target labels and selecting the one with the least perturbation required to enforce misclassification. It has been shown that the modulation classifier used in [8] incurs major errors under adversarial attacks with different power levels used for perturbation signals. The Carlini-Wagner attack [61] has been considered in [18] to perturb RF data points and the corresponding defense mechanism based on pre-trained classifier using an autoencoder is considered in [20]. Furthermore, defense mechanisms against adversarial attacks for modulation classifier have been studied in [62], [63]. The setting has been extended to the use of multiple antennas for transmitting perturbations in [64] and the effects of surrogate models trained by the adversary has been studied in [65]. Adversarial attack is also used to defend against an intruder trying to intercept the modulation scheme used at the transmitter in [24], [25], where the adversarial attack is added to the transmitted signal to fool intruders while maintaining the bit error rate between the transmitter and the receiver. This setting has been used in [66] to enable covert communications by fooling deep-learning-based signal detectors.

A typical assumption in previous studies has been that the adversarial perturbations is added to the receiver input without considering channel effects from the adversary to the receiver beyond the AWGN channel environment. However, even a small channel effect (pathloss or fading) would significantly reduce the impact of adversarial perturbation by decreasing the received perturbation power just below the necessary power level or tilting the perturbation to a different direction in vector space such that the adversarial attack fails in changing classification decision over the air. Furthermore, due to the broadcast nature of wireless communications, it is possible for an adversary to attack multiple receivers (each with a different channel from the adversary) simultaneously with a common

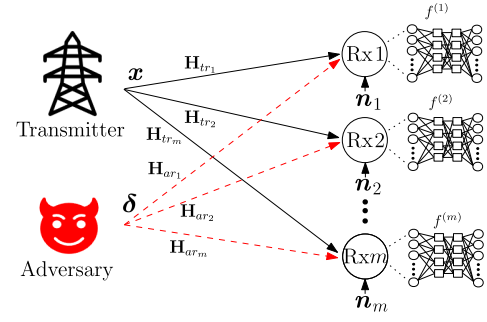


Fig. 1. System model with an adversary transmitting adversarial perturbation to fool multiple receivers.

adversarial perturbation although previous works have only considered attacks on a single receiver. This broadcast perturbation attack consumes less transmit energy and incurs less delay compared to sequentially transmitting multiple perturbations, each targeting a different receiver separately.

III. SYSTEM MODEL

We consider a wireless communications system that consists of a transmitter, m receivers, and an adversary as shown in Fig. 1. All nodes are equipped with a single antenna and operate on the same channel. Each receiver classifies its received signals with a DNN to the modulation type that is used by the transmitter. In the meantime, the adversary transmits perturbation signals over the air to fool the classifier at the receiver into making errors in modulation classification.

The DNN classifier at the i th receiver is denoted by $f^{(i)}(\cdot; \theta_i) : \mathcal{X} \rightarrow \mathbb{R}^C$, where θ_i is the parameters of the DNN at receiver i and C is the number of modulation types. Note $\mathcal{X} \subset \mathbb{C}^p$, where p is the dimension of the complex-valued (in-phase/quadrature) inputs that can be also represented by concatenation of two real-valued inputs. The classifier $f^{(i)}$ assigns a modulation type $\hat{l}^{(i)}(\mathbf{x}, \theta_i) = \arg \max_k f_k^{(i)}(\mathbf{x}, \theta_i)$ to every input $\mathbf{x} \in \mathcal{X}$ where $f_k^{(i)}(\mathbf{x}, \theta_i)$ is the output of the i th classifier corresponding to the k th modulation type.

The channel from the transmitter to the i th receiver is denoted by \mathbf{h}_{tr_i} and the channel from the adversary to the i th receiver is denoted by \mathbf{h}_{ar_i} , where $\mathbf{h}_{tr_i} = [h_{tr_i,1}, h_{tr_i,2}, \dots, h_{tr_i,p}]^T \in \mathbb{C}^{p \times 1}$ and $\mathbf{h}_{ar_i} = [h_{ar_i,1}, h_{ar_i,2}, \dots, h_{ar_i,p}]^T \in \mathbb{C}^{p \times 1}$. Note that \mathbf{h}_{tr_i} follows the channel model that is described in [67] and we assume that the channel $\mathbf{h}_{ar,j}$ is subject to Rayleigh fading with path-loss and shadowing effects, i.e., $h_{ar,j} = K(\frac{d_0}{d})^\gamma \psi h_{ray,j}$ where K is a constant factor, $h_{ray,j}$ is Rayleigh fading, γ is a path loss exponent, d_0 is the reference distance, d is the distance between the adversary and the receiver, and ψ is the shadowing effect to represent the randomness of received signal power in reality. If the transmitter transmits \mathbf{x} , the i th receiver receives $\mathbf{r}_{tr_i} = \mathbf{H}_{tr_i} \mathbf{x} + \mathbf{n}_i$, if there is no adversarial attack, or receives $\mathbf{r}_{ar_i}(\delta) = \mathbf{H}_{tr_i} \mathbf{x} + \mathbf{H}_{ar_i} \delta + \mathbf{n}_i$, if the adversary launches an adversarial attack by transmitting the perturbation signal δ , where $\mathbf{H}_{tr_i} = \text{diag}\{h_{tr_i,1}, \dots, h_{tr_i,p}\} \in \mathbb{C}^{p \times p}$, $\mathbf{H}_{ar_i} = \text{diag}\{h_{ar_i,1}, \dots, h_{ar_i,p}\} \in \mathbb{C}^{p \times p}$, $\delta \in \mathbb{C}^{p \times 1}$ and $\mathbf{n}_i \in \mathbb{C}^{p \times 1}$ is complex Gaussian noise. We assume that the adversarial perturbation δ is synchronized with the transmitter's signal,

thus superimposed with the transmitted signal \mathbf{x} at the receiver. To make the attack stealthy (i.e., hard to detect) and energy efficient, the adversarial perturbation δ is restricted as $\|\delta\|_2^2 \leq P_{\max}$ for some suitable power budget defined by P_{\max} .

The adversary determines the (common) adversarial perturbation δ for the input \mathbf{x} and all of the classifiers $f^{(i)}$ for $i = 1, 2, \dots, m$ by solving the optimization (1) problem:

$$\begin{aligned} & \arg \min_{\delta} \|\delta\|_2 \\ & \text{subject to } \hat{l}^{(i)}(\mathbf{r}_{tr_i}, \boldsymbol{\theta}_i) \neq \hat{l}^{(i)}(\mathbf{r}_{ar_i}(\delta), \boldsymbol{\theta}_i) \quad i = 1, 2, \dots, m \\ & \|\delta\|_2^2 \leq P_{\max}. \end{aligned} \quad (1)$$

In (1), the objective is to minimize the perturbation power subject to two constraints (a) each receiver misclassifies the received signal, and (b) the budget for total perturbation power is not exceeded. Note that the best solution is not necessarily realized at $\|\delta\|_2^2 = P_{\max}$ due to the complicated decision boundary of the DNN that depends on the power and phase of perturbation.

In practice, solving (1) is difficult because of the nonlinearity of the DNN. Thus, different methods have been proposed (primarily in the computer vision domain) to approximate the adversarial perturbation. For instance, FGM is a computationally efficient method for crafting adversarial attacks by linearizing the loss function of the DNN classifier. Let $L(\boldsymbol{\theta}, \mathbf{x}, \mathbf{y})$ denote the loss function of the model (at any given receiver), where $\mathbf{y} \in \{0, 1\}^C$ is the class vector. Then, FGM linearizes the loss function in a neighborhood of \mathbf{x} and uses this linearized function for optimization.

We consider two types of attacks called *targeted attacks* and *non-targeted attacks* that involve different objective functions for the adversary to optimize. In a targeted attack, the adversary aims to generate a perturbation that causes the classifier at the receiver to misclassify the input to a target class (label), e.g., a QPSK modulated signal is classified as QAM16. In a non-targeted FGM attack, the adversary searches for an attack that causes any misclassification (independent of target class). More details on these two types of attacks are provided in Section IV.

Our goal in this paper is to design adversarial perturbation attacks to fool potentially multiple classifiers while considering the channel effects. For the white-box adversarial attacks, we assume that for all i the adversary knows (a) the DNN architecture ($\boldsymbol{\theta}_i$ and $L^{(i)}(\cdot)$) of the classifier at the i th receiver, (b) the input at the i th receiver, and consequently (c) the channel \mathbf{h}_{ar_i} between the adversary and the i th receiver. We will relax these assumptions in Sections VI and VII.

We compare the performances of the attacks proposed in this paper with the benchmark attack from [14] that generates the adversarial perturbation in the AWGN channel. For performance evaluation, we use the VT-CNN2 classifier from [47] as the modulation classifier (also used in [14]), where the classifier consists of two convolution layers and two fully connected layers, and train it with GNU radio ML dataset RML2016.10a [67]. The dataset contains 220,000 samples. Each sample corresponds to a specific modulation scheme at a specific signal-to-noise ratio (SNR). There are eleven modulations in the dataset: BPSK, QPSK, 8PSK, QAM16,

QAM64, CPFSK, GFSK, PAM4, WBFM, AM-SSB, and AM-DSB. We follow the same setup of [47], using Keras with TensorFlow backend, where the input sample to the modulation classifier is 128 I/Q (in-phase/quadrature) channel symbols. Half of the samples are used for training and the other half are used in test time.

IV. TARGETED WHITE-BOX ADVERSARIAL ATTACKS USING CHANNEL INFORMATION

We start with a single receiver, i.e., $m = 1$, and omit the index i of the i th receiver for simplicity in Sections IV-VII. We will extend the setup to multiple receivers in Section VIII. For the targeted attack, the adversary minimizes $L(\boldsymbol{\theta}, \mathbf{r}_{ar}, \mathbf{y}^{target})$ with respect to δ , where \mathbf{y}^{target} is one-hot encoded desired target class (label). FGM is used to linearize the loss function as $L(\boldsymbol{\theta}, \mathbf{r}_{ar}, \mathbf{y}^{target}) \approx L(\boldsymbol{\theta}, \mathbf{r}_{tr}, \mathbf{y}^{target}) + (\mathbf{H}_{ar}\delta)^T \nabla_{\mathbf{x}} L(\boldsymbol{\theta}, \mathbf{r}_{tr}, \mathbf{y}^{target})$ that is minimized by setting $\mathbf{H}_{ar}\delta = -\alpha \nabla_{\mathbf{x}} L(\boldsymbol{\theta}, \mathbf{r}_{tr}, \mathbf{y}^{target})$, where α is a scaling factor to constrain the adversarial perturbation power to P_{\max} .

The adversary can generate different targeted attacks with respect to different \mathbf{y}^{target} that causes the classifier at the receiver to misclassify the received signals to $C - 1$ different modulation types. Thus, as in [14], the adversary can craft targeted attacks for all $C - 1$ modulation types and choose the target modulation that requires the least power. The case considered in [14] corresponds to $\mathbf{H}_{ar} = \mathbf{I}$ without realistic channel effects. We call the targeted attack perturbation in [14] as δ^{NoCh} , which is an optimal targeted attack without channel consideration (this corresponds to Algorithm 1 by setting $\mathbf{H}_{ar} = \mathbf{I}$). In the following subsections, we propose three targeted adversarial attacks to account for the effects of the channel.

A. Channel Inversion Attack

We first begin with a naive attack, where the adversary designs its attack by inverting the channel in the optimal targeted attack δ^{NoCh} , which is obtained using Algorithm 1 with $\mathbf{H}_{ar} = \mathbf{I}$. Since the adversarial attack goes through channel \mathbf{h}_{ar} , the j th element of the perturbation δ is simply designed as $\delta_j = \frac{\delta_j^{NoCh}}{h_{ar,j}}$ such that after going through the channel it has the same direction as δ_j^{NoCh} for $j = 1, \dots, p$. Furthermore, to satisfy the power constraint P_{\max} at the adversary, a scaling factor α is introduced such that $\delta^{div} = -\alpha\delta$, where $\alpha = \frac{\sqrt{P_{\max}}}{\|\delta\|_2}$. Thus, the perturbation received at the receiver becomes $\mathbf{H}_{ar}\delta^{div} = -\alpha\delta^{NoCh}$.

B. Minimum Mean Squared Error (MMSE) Attack

In the MMSE attack, the adversary designs the perturbation δ^{MMSE} so that the distance between the perturbation after going through the channel and the optimal targeted perturbation at the receiver (which corresponds to $\mathbf{H}_{ar} = \mathbf{I}$) is minimized. By designing the attack in this way, the received perturbation at the receiver is close to the optimal targeted attack as much as possible while satisfying the power constraint at the

adversary. However, since the classifier is sensitive to not only the direction but also to the power of perturbation, the squared error criterion might penalize the candidates of δ^{MMSE} that have more power with the direction of δ^{NoCh} , i.e., we set $\delta^{MMSE} = \gamma \delta^{NoCh}$ to search for all magnitudes of the δ^{NoCh} . Therefore, we formulate the optimization problem to select the perturbation δ^{MMSE} as

$$\begin{aligned} \min_{\delta^{MMSE}} \quad & \|\mathbf{H}_{ar} \delta^{MMSE} - \gamma \delta^{NoCh}\|_2^2 \\ \text{subject to} \quad & \|\delta^{MMSE}\|_2^2 \leq P_{max}, \end{aligned} \quad (2)$$

where γ is optimized by line search. We can write (2) as

$$\begin{aligned} \min_{\delta_j^{MMSE}} \quad & \sum_{j=1}^p \|h_{ar,j} \delta_j^{MMSE} - \gamma \delta_j^{NoCh}\|_2^2 \\ \text{subject to} \quad & \sum_{j=1}^p \|\delta_j^{MMSE}\|_2^2 \leq P_{max}. \end{aligned} \quad (3)$$

We solve the convex optimization problem (3) by using the Lagrangian method. The Lagrangian for (3) is given by

$$\mathcal{L} = \sum_{j=1}^p \|h_{ar,j} \delta_j^{MMSE} - \gamma \delta_j^{NoCh}\|_2^2 + \lambda \left(\sum_{j=1}^p \|\delta_j^{MMSE}\|_2^2 - P_{max} \right), \quad (4)$$

where $\lambda \geq 0$. The Karush–Kuhn–Tucker (KKT) conditions are given by

$$h_{ar,j}^* (h_{ar,j} \delta_j^{MMSE} - \gamma \delta_j^{NoCh}) + \lambda \delta_j^{MMSE} = 0, \quad (5)$$

for $j = 1, \dots, p$. From the KKT conditions, we obtain the perturbation of the MMSE attack as

$$\delta_j^{MMSE} = -\frac{\gamma h_{ar,j}^* \delta_j^{NoCh}}{h_{ar,j}^* h_{ar,j} + \lambda}, \quad (6)$$

for $j = 1, \dots, p$, where λ is determined by the power constraint at the adversary. Note that the received perturbation at the receiver is $\mathbf{H}_{ar} \delta^{MMSE} = -\alpha^T \delta^{NoCh}$, where $\alpha \in \mathbb{R}^{p \times 1}$ and the j th element of α is $\alpha_j = \frac{\gamma h_{ar,j} h_{ar,j}^*}{h_{ar,j}^* h_{ar,j} + \lambda}$.

C. Maximum Received Perturbation Power (MRPP) Attack

In the MRPP attack, the adversary selects the perturbation δ^{MRPP} to maximize the received perturbation power at the receiver and analyzes how the received perturbation power affects the decision process of the classifier. To maximize the received perturbation power and effectively fool the classifier into making a specific classification error, the adversary has to fully utilize the channel between the adversary and the receiver. Thus, if the targeted attack δ_j^{target} is multiplied by the conjugate of the channel, $h_{ar,j}^*$, to maximize the received perturbation power along the channel $h_{ar,j}$, then the received perturbation after going through the channel becomes $\|h_{ar,j} \delta_j^{target}\|_2^2$. In the MRPP attack, not only the direction of the adversarial perturbation is unaffected after going through the channel, but also the power of the adversarial perturbation is maximized by utilizing the channel. Finally, the adversary generates the targeted attack for every possible modulation

Algorithm 1 MRPP Attack

```

1: Inputs: input  $\mathbf{r}_{tr}$ , desired accuracy  $\varepsilon_{acc}$ , power constraint  $P_{max}$  and model of the classifier
2: Initialize:  $\varepsilon \leftarrow \mathbf{0}^{C \times 1}$ 
3: for class-index  $c$  in  $\text{range}(C)$  do
4:    $\varepsilon_{max} \leftarrow P_{max}, \varepsilon_{min} \leftarrow 0$ 
5:    $\delta_{norm}^c = \frac{\mathbf{H}_{ar}^* \nabla_{\mathbf{x}} L(\boldsymbol{\theta}, \mathbf{r}_{tr}, \mathbf{y}^c)}{\|\mathbf{H}_{ar}^* \nabla_{\mathbf{x}} L(\boldsymbol{\theta}, \mathbf{r}_{tr}, \mathbf{y}^c)\|_2}$ 
6:   while  $\varepsilon_{max} - \varepsilon_{min} > \varepsilon_{acc}$  do
7:      $\varepsilon_{avg} \leftarrow (\varepsilon_{max} + \varepsilon_{min})/2$ 
8:      $\mathbf{x}_{adv} \leftarrow \mathbf{x} - \varepsilon_{avg} \mathbf{H}_{ar} \delta_{norm}^c$ 
9:     if  $\hat{l}(\mathbf{x}_{adv}) == l_{true}$  then
10:        $\varepsilon_{min} \leftarrow \varepsilon_{avg}$ 
11:     else
12:        $\varepsilon_{max} \leftarrow \varepsilon_{avg}$ 
13:     end if
14:   end while
15:    $\varepsilon[c] = \varepsilon_{max}$ 
16: end for
17:  $target = \arg \min \varepsilon, \delta^{MRPP} = -\sqrt{P_{max}} \delta_{norm}^{target}$ 

```

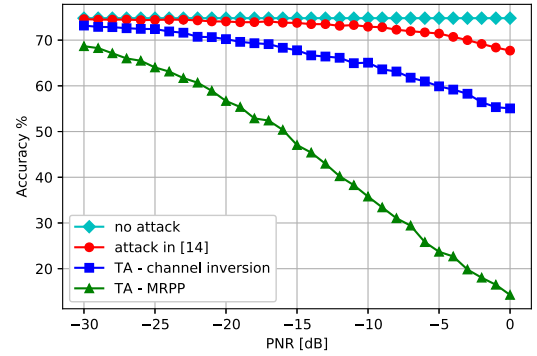


Fig. 2. Accuracy of the receiver's classifier under adversarial attacks with and without considering wireless channel effects when SNR = 10 dB.

type to decide the target class and calculate the scaling factor to satisfy the power constraint at the adversary. The details are presented in Algorithm 1.

D. Attack Performance

We set the parameters of the channel $h_{ar,j}$ as $K = 1$, $d_0 = 1$, $d = 10$, $\gamma = 2.7$, $\psi \sim \text{Lognormal}(0, 8)$ and $h_{ray,j} \sim \text{Rayleigh}(0, 1)$, where γ is selected to represent urban area cellular radio. We use the perturbation-to-noise ratio (PNR) metric from [14] that captures the relative perturbation with respect to the noise and measure how the increase in the PNR affects the accuracy of the classifier. Note that as the PNR increases, it is more likely for the attack to be detected by the receiver. In the performance evaluation figures, we denote the targeted attack by 'TA', the non-targeted attack by 'NTA' and the broadcast attack by 'BC'.

Fig. 2 presents the classifier accuracy versus the PNR under the proposed targeted white-box adversarial attack with exact channel information and compares them with the adversarial attack studied in [14]. We observe that the white-box attack from [14] without considering the necessary channel effects from the adversary to the receiver has poor performance that is

Algorithm 2 Naive Non-Targeted Attack

```

1: Inputs: input  $\mathbf{r}_{tr}$ , number of iterations  $E$ , power constraint
    $P_{max}$ , true class  $\mathbf{y}^{true}$  and model of the classifier
2: Initialize: Sum of gradient  $\Delta \leftarrow 0$ ,  $\mathbf{x} \leftarrow \mathbf{r}_{tr}$ 
3: for  $e$  in range( $E$ ) do
4:    $\delta_{norm} = \frac{\nabla_{\mathbf{x}} L(\theta, \mathbf{x}, \mathbf{y}^{true})}{\|\nabla_{\mathbf{x}} L(\theta, \mathbf{x}, \mathbf{y}^{true})\|_2}$ 
5:    $\mathbf{x} \leftarrow \mathbf{x} + \sqrt{\frac{P_{max}}{E}} \mathbf{H}_{ar} \delta_{norm}$ 
6:    $\Delta \leftarrow \Delta + \sqrt{\frac{P_{max}}{E}} \delta_{norm}$ 
7: end for
8:  $\delta^{naive} = \sqrt{P_{max}} \frac{\Delta}{\|\Delta\|_2}$ 

```

close to the no attack case in the low PNR region. The reason is that the wireless channel changes the phase and the magnitude of the perturbations perceived at the receiver. Furthermore, the targeted channel inversion attack does not perform well compared to the targeted MRPP attack, indicating the importance of the received power of perturbations for the performance of the classifier at the receiver.

V. NON-TARGETED WHITE-BOX ADVERSARIAL ATTACKS USING CHANNEL INFORMATION

The objective of the non-targeted attack is to maximize the loss function $L(\theta, \mathbf{r}_{ar}, \mathbf{y}^{true})$, where \mathbf{y}^{true} is the true class of \mathbf{x} . FGM is used to linearize the loss function as $L(\theta, \mathbf{r}_{ar}, \mathbf{y}^{true}) \approx L(\theta, \mathbf{r}_{tr}, \mathbf{y}^{true}) + (\mathbf{H}_{ar} \delta)^T \nabla_{\mathbf{x}} L(\theta, \mathbf{r}_{tr}, \mathbf{y}^{true})$ that is maximized by setting $\mathbf{H}_{ar} \delta = \alpha \nabla_{\mathbf{x}} L(\theta, \mathbf{r}_{tr}, \mathbf{y}^{true})$, where α is a scaling factor to constrain the adversarial perturbation power to P_{max} . Based on FGM, we introduce three non-targeted adversarial attacks described in the following subsections.

A. Naive Non-Targeted Attack

First, the adversary divides its power P_{max} into the number of iterations, E , in Algorithm 2 and allocates $\frac{P_{max}}{E}$ amount of power to the gradient of loss function to tilt the transmitted signal from the transmitter. Next, the adversary calculates the gradient again with respect to the transmitted signal from the transmitter and added perturbation. Then, the adversary adds another perturbation with power $\frac{P_{max}}{E}$ using the new gradient δ_{norm} as

$$\mathbf{x} \leftarrow \mathbf{x} + \sqrt{\frac{P_{max}}{E}} \mathbf{H}_{ar} \delta_{norm}. \quad (7)$$

This scheme generates the best direction to increase the loss function at that specific instance. Finally, the adversary repeats this procedure E times and sums all the gradients of the loss function that were added to the transmitted signal from the transmitter since the adversary can send only one perturbation at a time over the air. Finally, a scaling factor is introduced to satisfy the power constraint at the adversary. The details are presented in Algorithm 2.

B. Minimum Mean Squared Error (MMSE) Attack

The non-targeted MMSE attack is designed similar to the targeted MMSE attack. The adversary first obtains δ^{NoCh} from

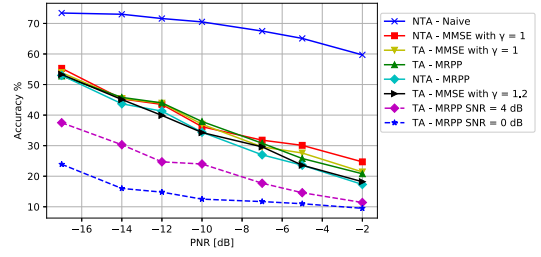


Fig. 3. Accuracy of the receiver's classifier under different white-box adversarial attacks and SNR levels.

the naive non-targeted attack with $\mathbf{H}_{ar} = \mathbf{I}$ and uses it to solve problem (2). Thus, the solution is the same as the solution to (2) except that it has the opposite direction to maximize the loss function, whereas the targeted attack minimizes the loss function. Therefore, the perturbation selected by the MMSE scheme for the non-targeted attack is $\delta^{MMSE} = \alpha^T \delta^{NoCh}$, where $\alpha \in \mathbb{R}^p$ and the j th element of α is $\alpha_j = \frac{\gamma h_{ar,j}^*}{h_{ar,i}^* h_{ar,j} + \lambda}$.

C. Maximum Received Perturbation Power (MRPP) Attack

As we have seen in the targeted MRPP attack, the attack should be in the form of $\delta^{MRPP} = \mathbf{H}_{ar}^* \delta^{target}$ to maximize the received perturbation power at the receiver. Thus, the naive non-targeted attack is changed to create the MRPP non-targeted attack by replacing δ_{norm} in Line 5 of Algorithm 2 with

$$\delta_{norm} = \frac{\mathbf{H}_{ar}^* \nabla_{\mathbf{x}} L(\theta, \mathbf{x}, \mathbf{y}^{true})}{\|\mathbf{H}_{ar}^* \nabla_{\mathbf{x}} L(\theta, \mathbf{x}, \mathbf{y}^{true})\|_2}. \quad (8)$$

D. Attack Performance

The performances of the proposed white-box attacks are compared in Fig. 3 under the same simulation environment used in Section IV-D. Note that SNR level is set to 10dB if it is not stated in the legend of the figure. As we discussed in Section IV-B, γ can be optimized with line search for the targeted MMSE attack, e.g., it performs better with $\gamma = 1.2$ compared to $\gamma = 1$. The naive non-targeted attack performs poorly compared to other attacks and the non-targeted MRPP attack outperforms all other attacks for most of the PNR values. This can be explained by the freedom of the direction that the non-targeted adversarial attack can take. For targeted attacks, there are only 10 different directions since there are 11 modulation types. However, the non-targeted attack does not have such a restriction. Thus, it is more likely that the non-targeted attack chooses a better direction to enforce a misclassification. Moreover, the computation complexity for the non-targeted attacks is lower compared to the targeted attacks which involve iterations to reach the desired accuracy. Finally, we observe that when the SNR level decreases for the MRPP attack, the accuracy of the classifier also decreases when the same PNR level is considered.

VI. WHITE-BOX ADVERSARIAL ATTACK WITH LIMITED CHANNEL INFORMATION

The adversarial attacks that are designed in the previous sections use the exact information of the channel (from the

Algorithm 3 Adversarial Attack With Limited Channel Information Using PCA

```

1: Inputs:  $N$  realizations of channels
    $\{\mathbf{H}_{ar}^{(1)}, \mathbf{H}_{ar}^{(2)}, \dots, \mathbf{H}_{ar}^{(N)}\}$ , input  $\mathbf{r}_{tr}$  and model of the
   classifier
2: Initialize:  $\Delta \leftarrow 0$ 
3: for  $n$  in range( $N$ ) do
4:   Find  $\delta^{(n)}$  from white-box attack algorithm using  $\mathbf{r}_{tr}$  and
      $\mathbf{H}_{ar}^{(n)}$ 
5:   Stack  $\delta^{(n)}$  to  $\Delta$ 
6: end for
7: Compute the first principle direction  $\mathbf{v}_1$  of  $\Delta$  using PCA
8:  $\Delta = \mathbf{U}\Sigma\mathbf{V}^T$  and  $\mathbf{v}_1 = \mathbf{V}\mathbf{e}_1$ 
9:  $\delta^{limited} = \sqrt{P_{max}}\mathbf{v}_1$ 

```

adversary to the receiver). However, this information may not always be available in practical scenarios. Therefore, in this section, we present Algorithm 3 to generate adversarial attacks with limited channel information, meaning that the adversary only knows the distribution of the channel between the adversary and the receiver and applies dimensionality reduction by accounting for limited channel information. A classical approach for dimensionality reduction is the principal component analysis (PCA) algorithm that was also used in [14] for the case when the signal is directly manipulated at the receiver. PCA is used to obtain the principal component that has the largest variance, i.e., PCA finds the principal component that provides the most information about the data in a reduced dimension by projecting the data onto it. PCA can be calculated by the eigenvalue decomposition of the data covariance matrix or the singular value decomposition of a data matrix.

To generate an adversarial attack with limited channel information, the adversary first generates N realizations of the channel from the adversary to the receiver $\{\mathbf{H}_{ar}^{(1)}, \mathbf{H}_{ar}^{(2)}, \dots, \mathbf{H}_{ar}^{(N)}\}$ which are different from the actual channel \mathbf{H}_{ar} . Then, it generates N adversarial attacks using white-box attack algorithms from the previous sections, either targeted or non-targeted, using N realizations of the observed channel and the known input \mathbf{r}_{tr} at the classifier. Finally, it stacks N generated adversarial attacks in a matrix and finds the principal component of the matrix to use it as the adversarial attack with limited channel information. The details are presented in Algorithm 3.

VII. UNIVERSAL ADVERSARIAL PERTURBATION (UAP) ATTACK

In the previous sections, the adversary designs a white-box attack with the assumptions that it knows the model of the classifier at the receiver, the channel between the adversary and the receiver, and the exact input at the receiver. However, these assumptions are not always practical in real wireless communications systems. Thus, in this section, we relax these assumptions and present the UAP attacks.

A. UAP Attack With Pre-Collected Input at the Receiver

We first relax the assumption that the adversary knows the exact input of the classifier. The adversary in the previous attacks generates an input-dependent perturbation, i.e., δ is designed given the exact input \mathbf{r}_{tr} . This requires the adversary to always know the input of the classifier, which is not a practical assumption to make and hard to achieve due to synchronization issues. Thus, it is more practical to design an input-independent UAP. We consider two different methods to design an input-independent UAP. We assume that the adversary collects some arbitrary set of inputs $\{\mathbf{r}_{tr}^{(1)}, \mathbf{r}_{tr}^{(2)}, \dots, \mathbf{r}_{tr}^{(N)}\}$ and associated classes.

1) *PCA-Based Input-Independent UAP Attack*: We present an algorithm to design the UAP attack using PCA. First, the adversary generates perturbations $\{\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(N)}\}$ with respect to the obtained arbitrary set of inputs $\{\mathbf{r}_{tr}^{(1)}, \mathbf{r}_{tr}^{(2)}, \dots, \mathbf{r}_{tr}^{(N)}\}$ and the exact channel information using algorithms from the previous sections such as the MRPP attack. To reflect the common characteristics of $\{\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(N)}\}$ in the UAP, we stack these perturbations into a matrix and perform PCA to find the first component of the matrix with the largest eigenvalue. Hence, we use the direction of the first principal component as the direction of UAP for channel \mathbf{H}_{ar} . The algorithm for the UAP with N pre-collected input data is similar to Algorithm 3 and the difference is using pre-collected input instead of different realizations of channel.

2) *VAE-Based Input-Independent UAP Attack*: We present an algorithm to design the UAP attack using a VAE that is known to effectively capture complex data characteristics. An autoencoder consists of two DNNs, an encoder to map the input latent space and a decoder to recover the input data from this latent space. VAE extends this setting by describing a probability distribution for each latent attribute instead of providing a single value to describe each latent state attribute. We first collect an arbitrary set of inputs $\{\mathbf{r}_{tr}^{(1)}, \mathbf{r}_{tr}^{(2)}, \dots, \mathbf{r}_{tr}^{(N)}\}$ to create corresponding perturbations $\{\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(N)}\}$ using Algorithm 1 with $\mathbf{H}_{ar} = \mathbf{I}$. Then, we train the VAE using these perturbations. The encoder structure of the VAE is presented in Table I and the decoder structure is presented in Table II. We use 128 number of filters and kernel size of (1,3) for Conv 1 layer and 40 number of filters and kernel size of (2,3) for Conv 2 layer at the encoder. For the deconvolution layers of the decoder, we use 40 number of filters and kernel size of (2,3) for Deconv 1 layer, 128 number of filters and kernel size of (1,3) for Deconv 2 layer, and one filter and kernel size of (3,3) for Deconv 3 layer. Since the encoder learns an efficient way to compress the data into the lower-dimensional feature space, we first use the encoder to decrease the dimension of the perturbations. Specifically, we use k number of perturbations $\{\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(k)}\}$ as inputs to the encoder and get the corresponding outputs $\{z^{(1)}, z^{(2)}, \dots, z^{(k)}\}$, where $z^{(i)} \in \mathbb{R}^{q \times 1}$, $i = 1, \dots, k$. To reflect the common characteristics of $\{\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(k)}\}$, we average the corresponding outputs $\{z^{(1)}, z^{(2)}, \dots, z^{(k)}\}$ and use $z^{avg} = \frac{1}{k} \sum_{i=1}^k z^{(i)}$ in the

TABLE I
DECODER PART OF THE VAE LAYOUT

Layer	Output dimension	Number of parameters
Input	2×128	
Conv 1	$2 \times 128 \times 128$	512
Conv 2	$2 \times 128 \times 40$	30760
Flatten	10240	
Dense 1	16	163856
Dense 2	4	68

TABLE II
DECODER PART OF THE VAE LAYOUT

Layer	Output dimension	Number of parameters
Input	4	
Dense	10240	51200
Deconv 1	$2 \times 128 \times 40$	9640
Deconv 2	$2 \times 128 \times 128$	15488
Deconv 3	2×128	1153

decoder to generate δ^{avg} . Finally, we use channel information to design an input-independent UAP as $\delta^{VAE} = \sqrt{P_{max}} \frac{\mathbf{H}_{ar}^* \delta^{avg}}{\|\mathbf{H}_{ar}^* \delta^{avg}\|}$.

B. UAP Attack With Limited Channel Information

Now, we further relax the assumption that the adversary knows the exact channel between the adversary and the receiver and assume that the adversary only knows the distribution of this channel. We consider two different methods to design the UAP attack with limited channel information.

1) *PCA-Based Channel-Independent UAP Attack*: To design the UAP attack when the adversary only knows the distribution of the channel, we first generate random realizations of the channels $\{\mathbf{H}_{ar}^{(1)}, \mathbf{H}_{ar}^{(2)}, \dots, \mathbf{H}_{ar}^{(N)}\}$ from the distribution. Then, we generate $\delta^{(n)}$ using $\mathbf{r}_{tr}^{(n)}$ and $\mathbf{H}_{ar}^{(n)}$ for $n = 1, \dots, N$ instead of using the real channel \mathbf{H}_{ar} . We use PCA to find the first component of the matrix and use it as the direction of UAP. The algorithm for PCA-based channel-independent UAP attack is analogous to Algorithm 3 except that we have random realizations of channels as opposed to exact channel used in Algorithm 3.

2) *VAE-Based Channel-Independent UAP Attack*: Similar to the PCA-based channel-independent UAP attack, we first generate random realizations of the channel $\{\mathbf{H}_{ar}^{(1)}, \mathbf{H}_{ar}^{(2)}, \dots, \mathbf{H}_{ar}^{(N)}\}$ from the distribution to train the VAE. We use the encoder to capture the common characteristics of the channel by using k number of channels as inputs to the encoder and average the outputs of the encoder to put into the decoder as we have done in the VAE-based input-independent UAP attack to get \mathbf{H}_{avg} . Note that we use the same encoder and decoder structure that is used for VAE-based input-independent UAP attack. Furthermore, we use the steps used in the VAE-based input-independent UAP attack to create δ^{avg} and design $\delta^{VAE} = \sqrt{P_{max}} \frac{\mathbf{H}_{ar}^* \delta^{avg}}{\|\mathbf{H}_{ar}^* \delta^{avg}\|}$.

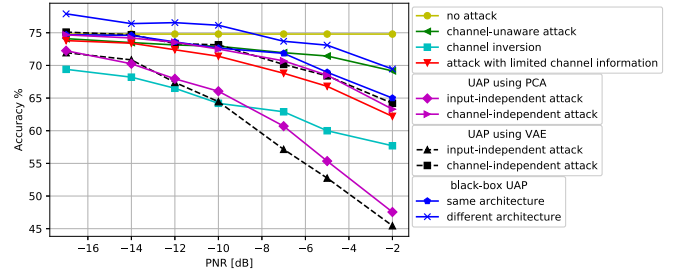


Fig. 4. Accuracy of the receiver's classifier under UAP attacks with different levels of information availability when SNR = 10 dB.

C. Black-Box UAP Attack

The last assumption that we will relax is the information about the DNN classifier model at the receiver. To relax this assumption, we use the transferability property of adversarial attacks [68]. This property states that the adversarial attack crafted to fool a specific DNN can also fool other DNNs with different models with high probability. Therefore, the adversary generates UAP attacks using a substitute DNN that is trained using an arbitrary collected dataset and applies them to fool the actual DNN at the receiver.

D. Attack Performance

Fig. 4 shows the performance of the adversarial attacks with respect to different levels of information availability. First, we observe that the input-independent UAP using VAE with 40 pre-collected data inputs, where the adversary knows the exact channel information, outperforms other attacks with limited information and even the channel inversion attack. Note that VAE is trained with 2000 samples but only 40 samples are used to actually create one UAP. This result shows the importance of the channel state information over the exact input data when crafting an adversarial attack since the input-dependent attack significantly outperforms the attack with limited channel information. Also, the input-independent UAP using VAE outperforms the input-independent UAP using PCA. However, the performance is similar when VAE and PCA are compared for the channel-independent attack suggesting that the UAP using PCA might be enough to generate an effective UAP attack. The input-independent UAP with 40 pre-collected data inputs even outperforms the targeted channel inversion attack in the high PNR region, where the adversary knows not only the exact channel but also the exact input at the receiver. Furthermore, the black-box UAP with the same architecture performs similar to the attack with limited channel information where we use the substitute DNN that has the same structure of the classifier used in [47] but is trained with a different training set so that the substitute DNN has different weights. Lastly, when the adversary uses a different DNN architecture that has one more hidden layer with 256 nodes compared to the target DNN model to create an attack, its attack performance is comparable to the channel-unaware attack in the high PNR region.

VIII. BROADCAST ADVERSARIAL ATTACK

We extend the adversarial attack to a broadcast (common) adversarial perturbation that is transmitted by the adversary to

simultaneously fool multiple classifiers at different receivers. First, we show that the channel-aware adversarial attack built for a particular channel is inherently selective, namely it is only effective against the receiver with that experienced channel and not effective against another receiver with a different experienced channel. Next, we show how to design a common adversarial perturbation by jointly accounting for multiple channels such that the adversary can fool classifiers at multiple receivers simultaneously with a single (omnidirectional) transmission. We present the design of broadcast adversarial attack only for the case of the targeted adversarial attacks, as other types of attacks can be formulated similarly. Note that again we assume the adversary knows the DNN architecture of the receiver, input at the i th receiver, and the channel between the adversary and the i th receiver in this section.

A. Individual Design of Broadcast Adversarial Attack (IDBA)

We design IDBA by building upon Algorithm 1 (that was designed against a signal receiver). By applying Algorithm 1 separately against each of m receivers, we obtain m adversarial perturbations, $\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(m)}$, where $\delta^{(i)}$ is the adversarial perturbation generated using Algorithm 1 for the i th receiver. Since the adversary aims to fool all the receivers by broadcasting a single adversarial perturbation, we combine individual perturbations as a weighted sum, $\sum_{i=1}^m w_i \delta^{(i)}$, where w_i is the weight for adversarial perturbation $\delta^{(i)}$, and then normalize it to satisfy the power constraint at the adversary (higher w_i implies more priority given to $\delta^{(i)}$). Note that the optimal weights can be found by line search subject to $\sum_{i=1}^m w_i = 1$ and $w_i \geq 0$ for each i . Numerical results suggest that the search space can be constrained by selecting weights inversely proportional to channel gains of corresponding receivers.

B. Joint Design of Broadcast Adversarial Attack (JDBA)

We generate the adversarial perturbation JDBA by modifying Algorithm 1. Specifically, we change the computation of δ_{norm}^c in Line 5 of Algorithm 1 from (8) to

$$\delta_{norm}^c = \frac{\sum_{i=1}^m w_i \mathbf{H}_{ar}^{(i)*} \nabla_{\mathbf{x}} L_i(\boldsymbol{\theta}, \mathbf{r}_{tr}^{(i)}, \mathbf{y}^c)}{\|\sum_{i=1}^m w_i \mathbf{H}_{ar}^{(i)*} \nabla_{\mathbf{x}} L_i(\boldsymbol{\theta}, \mathbf{r}_{tr}^{(i)}, \mathbf{y}^c)\|_2}, \quad (9)$$

search overall $C - 1$ classes, and choose the class that fools most of the receivers. Note that JDBA searches for a common direction of δ_{norm}^c for all receivers. On the other hand, IDBA searches separately for directions of adversarial perturbations for different receivers and then combines them to a common direction.

C. Attack Performance

For performance evaluation, we assume that there are two receivers, $m = 2$, with different classifiers which have the same architecture, but trained with different training sets and the adversary generates a broadcast adversarial perturbation to fool both receivers simultaneously. We assume Rayleigh fading for both channel $h_{ar1,j} = K(\frac{d_0}{d})^{\gamma} \psi h_{ray1,j}$ from the adversary to receiver 1, and the channel $h_{ar2,j} = K(\frac{d_0}{d})^{\gamma} \psi h_{ray2,j}$

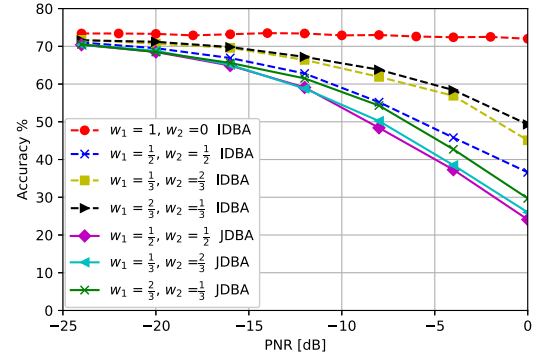


Fig. 5. Accuracy of $m = 2$ classifiers under broadcast adversarial attacks when SNR = 10 dB.

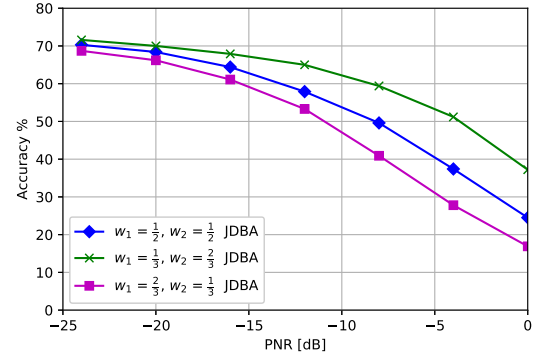


Fig. 6. Accuracy of $m = 2$ classifiers under broadcast adversarial attacks at receivers with different channel distributions when SNR = 10 dB.

from the adversary to receiver 2, where $h_{ray1,j}$ and $h_{ray2,j} \sim \text{Rayleigh}(0, 1)$.

Fig. 5 shows the accuracy of the classifiers with respect to different broadcast adversarial attacks. Note that the accuracy is the same for both classifiers since we assume that the broadcast adversarial attack is successful if it fools both receivers at the same time. JDBA outperforms IDBA for all PNRs. Furthermore, the weight selection $w_1 = \frac{1}{2}, w_2 = \frac{1}{2}$ outperforms other weight selections suggesting that the weight selection is related to the channel distribution. Also, when weights $w_1 = 1$ and $w_2 = 0$ are selected in IDBA, the attack generated only for receiver 1 has no effect on receiver 2, i.e., the attack selectively fools receiver 1 due to different channel effects. This observation validates the need to craft a broadcast adversarial perturbation.

Next, we distinguish channel distributions for the two receivers and evaluate the classifier accuracy under the JDBA attack in Fig. 6. The channel from the adversary to receiver 1, and the channel from the adversary to receiver 2 are modeled with Rayleigh fading of different variances, i.e., $h_{ray1,j} \sim \text{Rayleigh}(0, 1)$ and $h_{ray2,j} \sim \text{Rayleigh}(0, 2)$. The weight selection $w_1 = \frac{2}{3}$ and $w_2 = \frac{1}{3}$ (where the weights are selected as inversely proportional to the variance of the channel as we have seen in Fig. 5) outperforms other weight selections.

IX. DEFENSE AGAINST ADVERSARIAL ATTACKS

In this section, we introduce a defense method to increase the robustness of wireless signal classification models against adversarial perturbations. There are several methods developed

in the computer vision domain to defend against adversarial attacks. One method is *adversarial training* that calculates an adversarial perturbation and adds the perturbed samples to the training set to increase the model robustness. However, to perform adversarial training, the receiver needs to know what kind of adversarial attack is used at the adversary. In response, a line of work on *certified defense* has been considered [26]–[28], [69]. A classifier is said to be *certifiably robust* if the classifier's prediction of a sample x is constant in a small neighborhood around x , often defined by a ℓ_2 or ℓ_∞ ball. *Randomized smoothing* (also referred to as Gaussian smoothing) is a certified defense approach, which augments the training set with Gaussian noise to increase the robustness of the classifier to multiple directions. Recent work has shown that a tight robustness guarantee in ℓ_2 norm can be achieved with randomized smoothing with Gaussian noise [26]. We assume that the receiver has no information of the adversarial attack that is generated at the adversary. Hence, we apply randomized smoothing as a defense mechanism to make the modulation classifier robust against wireless adversarial perturbations over the air.

In randomized smoothing, there are two parameters σ and k that the defender can control, namely the standard deviation of the Gaussian noise σ and the number of noisy samples k added to each training sample r_i , i.e., $r_i + n_1, r_i + n_2, \dots, r_i + n_k$, where n_j is Gaussian noise with variance σ^2 . These two parameters balance the trade-off between the resulting classification accuracy and the robustness against perturbations. Note that unlike the images used as input samples in computer vision, the received signals in wireless communications are already subject to noise, and randomized smoothing slightly increases the noise level. However, as we will see in Section IX-A, as the number of data samples in the training set is less than the number of parameters in the neural network, data augmentation in fact improves the classifier performance and Gaussian smoothing does not cause any degradation.

A. Randomized Smoothing During Training

We evaluate the accuracy of the classifier for different values of σ and k selected for Gaussian noise augmentation. First, we fix $k = 10$ and change σ in Fig. 7 to study the impact of the power of the Gaussian noise added to the training set. The accuracy of the classifier trained with small values of σ such as $\sigma = 0.001$ and $\sigma = 0.005$ is larger than that of the original classifier across all PNR values. This result shows that randomized smoothing as training data augmentation improves the classifier accuracy for a small σ value, while degrading the performance when noise with a large variance is introduced.

Furthermore, there is an intersection between accuracy versus PNR curves for $\sigma = 0.001$ and $\sigma = 0.005$ due to the density of the norm ball that is created around each training set. These results suggest that when σ is small, the classifier becomes more robust to a small perturbation, but more susceptible to a large perturbation. We observe a similar crossover for large σ values (e.g., $\sigma = 0.01$) where the classifier is more robust to a large perturbation, but susceptible to a small perturbation.

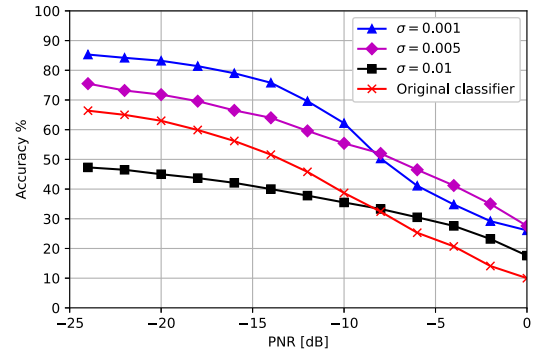


Fig. 7. Accuracy of the receiver's classifier when trained with $(10, \sigma)$ Gaussian noise augmentation with different σ and SNR = 10 dB.

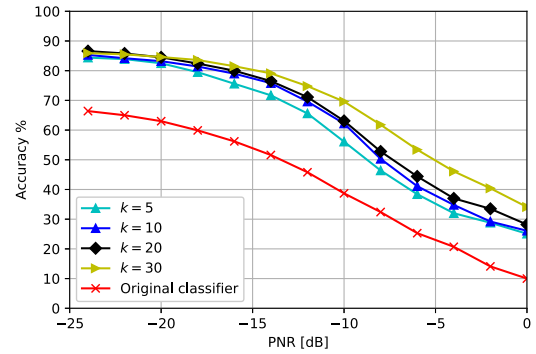


Fig. 8. Accuracy of the receiver's classifier when trained with $(k, 0.001)$ Gaussian noise augmentation with different k and SNR = 10 dB.

In Fig. 8, we keep the noise variance σ constant and vary the number of samples k added to the training set to investigate its effect on the classifier robustness. Our results demonstrate that by adding more augmentation samples to training set increases the robustness of the classifier against the adversarial perturbation when compared to the original classifier trained without any defense. However, this defense advantage comes at the cost of increased training time.

B. Certified Defense in Test Time

The defense results can be certified with a desired confidence by using randomized smoothing in test time. Consider the classifier trained with $k = 20$ and $\sigma = 0.001$. For each perturbed test sample, we draw k Gaussian noise samples with the same variance and label them with the classifier. Let \hat{c}_A and \hat{c}_B denote the classes that occurred most and second most, and n_A and n_B represent the occurrence of these two classes, respectively. We then apply a two-sided hypothesis test and check if $\text{BinomPValue}(n_A, n_A + n_B, q) \leq \alpha$ condition is satisfied, where $\text{BinomPValue}(\cdot)$ returns p -value of the two-sided hypothesis test that $n_A \sim \text{Binomial}(n_A + n_B, q)$ [26], $1 - q$ is the confidence level, and α is the probability of returning an incorrect answer. If the condition is satisfied, the classifier is very confident in its prediction. If not, then the classifier abstains and does not make a prediction. For example, when we consider 95% confidence, we observe that a

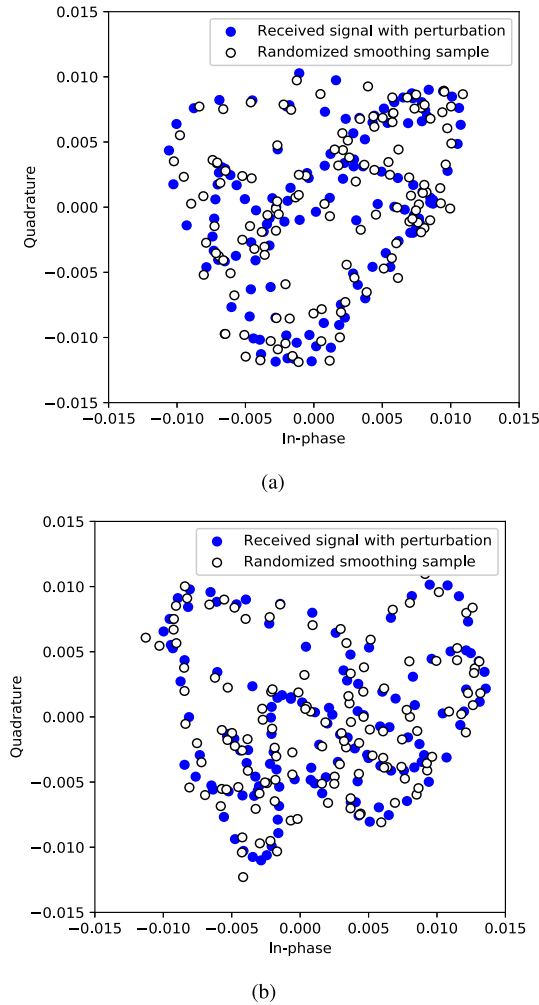


Fig. 9. Constellation points of the received signal with adversarial perturbation and randomized smoothing samples for the cases (a) when the classifier abstains and (b) when the classifier recovers the perturbed signal and correctly infers the label.

test sample that belongs to class “QAM64” is perturbed by the adversary to be classified as “QAM16” at the receiver. When randomized smoothing is applied in test time, we observe that the classifier infers the samples as class “QAM16” for 6 times and as class “QAM64” for 14 times, resulting in an $\text{BinomPValue}(\cdot) = 0.115$. Since the condition is not satisfied, the classifier abstains.

Another example is a test sample of “QAM64” that is perturbed by the adversary to be classified as “8PSK”. When randomized smoothing is applied in test time, the classifier correctly infers all k samples as “QAM64” and the defense is certified with 95% confidence. The constellation of the two examples are shown respectively in Fig. 9(a)-(b).

X. CONCLUSION

We presented over-the-air adversarial attacks against deep learning-based modulation classifiers by accounting for realistic channel and broadcast transmission effects. Specifically, we considered targeted, non-targeted and UAP attacks with different levels of uncertainty regarding channels, transmitter

inputs and DNN classifier models. We showed that these channel-aware adversarial attacks can successfully fool a modulation classifier over the air. Then, we introduced broadcast adversarial attacks to simultaneously fool multiple classifiers at different receivers with a single perturbation transmission. First, we showed that an adversarial perturbation designed for a particular receiver is not effective against another receiver due to channel differences. Therefore, we designed a common adversarial perturbation by considering all channel effects jointly and showed that this broadcast attack can fool all receivers. Finally, we presented a certified defense method using randomized smoothing, and showed that it is effective in reducing the impact of adversarial attacks on the modulation classifier performance.

REFERENCES

- [1] B. Kim, Y. E. Sagduyu, K. Davaslioglu, T. Erpek, and S. Ulukus, “Over-the-air adversarial attacks on deep learning based modulation classifier over wireless channels,” in *Proc. 54th Annu. Conf. Inf. Sci. Syst. (CISS)*, Mar. 2020, pp. 1–6.
- [2] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.
- [3] T. Erpek, T. O’Shea, Y. E. Sagduyu, Y. Shi, and T. C. Clancy, “Deep learning for wireless communications,” in *Development and Analysis of Deep Learning Architectures*. New York, NY, USA: Springer, 2020, pp. 223–266.
- [4] K. Davaslioglu and Y. E. Sagduyu, “Generative adversarial learning for spectrum sensing,” in *Proc. IEEE Int. Conf. Commun. (ICC)*, May 2018, pp. 1–6.
- [5] X. Song *et al.*, “SALDR: Joint self-attention learning and dense refine for massive MIMO CSI feedback with multiple compression ratio,” *IEEE Wireless Commun. Lett.*, vol. 10, no. 9, pp. 1899–1903, Sep. 2021.
- [6] J. Wang, G. Gui, T. Ohtsuki, B. Adebisi, H. Gacanin, and H. Sari, “Compressive sampled CSI feedback method based on deep learning for FDD massive MIMO systems,” *IEEE Trans. Commun.*, vol. 69, no. 9, pp. 5873–5885, Sep. 2021.
- [7] Y. Jin, J. Zhang, S. Jin, and B. Ai, “Channel estimation for cell-free mmWave massive MIMO through deep learning,” *IEEE Trans. Veh. Technol.*, vol. 68, no. 10, pp. 10325–10329, Nov. 2019.
- [8] T. J. O’Shea, J. Corgan, and T. C. Clancy, “Convolutional radio modulation recognition networks,” in *Proc. Int. Conf. Eng. Appl. Neural Netw.*, 2016, pp. 213–226.
- [9] Y. Vorobeychik and M. Kantarcioglu, “Adversarial machine learning,” *Synth. Lectures Artif. Intell. Mach. Learn.*, vol. 12, pp. 1–169, Dec. 2017.
- [10] C. Szegedy *et al.*, “Intriguing properties of neural networks,” 2013, *arXiv:1312.6199*.
- [11] W. Xu, W. Trappe, Y. Zhang, and T. Wood, “The feasibility of launching and detecting jamming attacks in wireless networks,” in *Proc. 6th ACM Int. Symp. Mobile ad hoc Netw. Comput. MobiHoc*, 2005, pp. 46–57.
- [12] Y. E. Sagduyu, R. A. Berry, and A. Ephremides, “Wireless jamming attacks under dynamic traffic uncertainty,” in *Proc. Int. Symp. Model. Optim. Mobile, Ad Hoc, Wireless Netw. (WiOpt)*, May 2010, pp. 303–312.
- [13] Y. Sagduyu, R. Berry, and A. Ephremides, “Jamming games in wireless networks with incomplete information,” *IEEE Commun. Mag.*, vol. 49, no. 8, pp. 112–118, Aug. 2011.
- [14] M. Sadeghi and E. G. Larsson, “Adversarial attacks on deep-learning based radio signal classification,” *IEEE Wireless Commun. Lett.*, vol. 8, no. 1, pp. 213–216, Feb. 2019.
- [15] B. Flowers, R. M. Buehrer, and W. C. Headley, “Evaluating adversarial evasion attacks in the context of wireless communications,” *IEEE Trans. Inf. Forensics Security*, vol. 15, pp. 1102–1113, 2020.
- [16] S. Bair, M. DeVecchio, B. Flowers, A. J. Michaels, and W. C. Headley, “On the limitations of targeted adversarial evasion attacks against deep learning enabled modulation recognition,” in *Proc. ACM Workshop Wireless Secur. Mach. Learn. (WiseML)*, 2019, pp. 25–30.

- [17] S. Kokalj-Filipovic and R. Miller, "Adversarial examples in RF deep learning: Detection of the attack and its physical robustness," 2019, *arXiv:1902.06044*.
- [18] S. Kokalj-Filipovic, R. Miller, and J. Morman, "Targeted adversarial examples against RF deep classifiers," in *Proc. ACM Workshop Wireless Secur. Mach. Learn. (WiseML)*, 2019, pp. 6–11.
- [19] S. Kokalj-Filipovic, R. Miller, and G. Vanhoy, "Adversarial examples in RF deep learning: Detection and physical robustness," in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, Nov. 2019, pp. 1–5.
- [20] S. Kokalj-Filipovic, R. Miller, N. Chang, and C. L. Lau, "Mitigation of adversarial examples in RF deep classifiers utilizing AutoEncoder pre-training," in *Proc. Int. Conf. Mil. Commun. Inf. Syst. (ICMCIS)*, May 2019, pp. 1–6.
- [21] D. Ke, Z. Huang, X. Wang, and L. Sun, "Application of adversarial examples in communication modulation classification," in *Proc. Int. Conf. Data Mining Workshops (ICDMW)*, Nov. 2019, pp. 877–882.
- [22] Y. Lin, H. Zhao, Y. Tu, S. Mao, and Z. Dou, "Threats of adversarial attacks in DNN-based modulation recognition," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Jul. 2020, pp. 2469–2478.
- [23] Y. Lin, H. Zhao, X. Ma, Y. Tu, and M. Wang, "Adversarial attacks in modulation recognition with convolutional neural networks," *IEEE Trans. Rel.*, vol. 70, no. 1, pp. 389–401, Mar. 2021.
- [24] M. Z. Hameed, A. Gyorgy, and D. Gunduz, "The best defense is a good offense: Adversarial attacks to avoid modulation detection," *IEEE Trans. Inf. Forensics Security*, vol. 16, pp. 1074–1087, 2021.
- [25] M. Z. Hameed, A. Gyorgy, and D. Gunduz, "Communication without interception: Defense against modulation detection," in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, Nov. 2019, pp. 1–5.
- [26] J. M. Cohen, E. Rosenfeld, and Z. Kolter, "Certified adversarial robustness via randomized smoothing," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 1310–1320.
- [27] M. Lecuyer, V. Atlidakis, R. Geambasu, D. Hsu, and S. Jana, "Certified robustness to adversarial examples with differential privacy," in *Proc. IEEE Symp. Secur. Privacy (SP)*, May 2019, pp. 656–672.
- [28] N. Carlini *et al.*, "On evaluating adversarial robustness," 2019, *arXiv:1902.06705*.
- [29] Y. E. Sagduyu *et al.*, "When wireless security meets machine learning: Motivation, challenges, and research directions," 2020, *arXiv:2001.08883*.
- [30] D. Adesina, C.-C. Hsieh, Y. E. Sagduyu, and L. Qian, "Adversarial machine learning in wireless communications using RF data: A review," 2020, *arXiv:2012.14392*.
- [31] Y. Shi, Y. E. Sagduyu, T. Erpek, K. Davaslioglu, Z. Lu, and J. H. Li, "Adversarial deep learning for cognitive radio security: Jamming attack and defense strategies," in *Proc. IEEE Int. Conf. Commun. Workshops (ICC Workshops)*, May 2018, pp. 1–6.
- [32] T. Erpek, Y. E. Sagduyu, and Y. Shi, "Deep learning for launching and mitigating wireless jamming attacks," *IEEE Trans. Cogn. Commun. Netw.*, vol. 5, no. 1, pp. 2–14, Mar. 2019.
- [33] Y. Shi, T. Erpek, Y. E. Sagduyu, and J. H. Li, "Spectrum data poisoning with adversarial deep learning," in *Proc. IEEE Mil. Commun. Conf. (MILCOM)*, Oct. 2018, pp. 407–412.
- [34] Y. E. Sagduyu, T. Erpek, and Y. Shi, "Adversarial deep learning for over-the-air spectrum poisoning attacks," *IEEE Trans. Mobile Comput.*, vol. 20, no. 2, pp. 2–14, Feb. 2021.
- [35] Z. Luo, S. Zhao, Z. Lu, J. Xu, and Y. Sagduyu, "When attackers meet AI: learning-empowered attacks in cooperative spectrum sensing," *IEEE Trans. Mobile Comput.*, early access, Oct. 12, 2020, doi: [10.1109/TMC.2020.3030061](https://doi.org/10.1109/TMC.2020.3030061).
- [36] Z. Luo, S. Zhao, R. Duan, Z. Lu, Y. E. Sagduyu, and J. Xu, "Low-cost influence-limiting defense against adversarial machine learning attacks in cooperative spectrum sensing," in *Proc. 3rd ACM Workshop Wireless Secur. Mach. Learn.*, Jun. 2022, pp. 55–60.
- [37] Z. Luo, S. Zhao, Z. Lu, Y. E. Sagduyu, and J. Xu, "Adversarial machine learning based partial-model attack in IoT," in *Proc. 2nd ACM Workshop Wireless Secur. Mach. Learn.*, Jul. 2020, pp. 13–18.
- [38] K. Davaslioglu and Y. E. Sagduyu, "Trojan attacks on wireless signal classification with adversarial machine learning," in *Proc. IEEE Int. Symp. Dyn. Spectr. Access Netw. (DySPAN)*, Nov. 2019, pp. 1–6.
- [39] Y. Shi, K. Davaslioglu, and Y. E. Sagduyu, "Over-the-air membership inference attacks as privacy threats for deep learning-based wireless signal classifiers," in *Proc. 2nd ACM Workshop Wireless Secur. Mach. Learn.*, Jul. 2020, pp. 61–66.
- [40] Y. Shi and Y. E. Sagduyu, "Membership inference attack and defense for wireless signal classifiers with deep learning," 2021, *arXiv:2107.12173*.
- [41] Y. Shi, K. Davaslioglu, and Y. E. Sagduyu, "Generative adversarial network for wireless signal spoofing," in *Proc. ACM Workshop Wireless Secur. Mach. Learn. (WiseML)*, 2019, pp. 55–60.
- [42] Y. Shi, K. Davaslioglu, and Y. E. Sagduyu, "Generative adversarial network in the air: Deep adversarial learning for wireless signal spoofing," *IEEE Trans. Cognit. Commun. Netw.*, vol. 7, no. 1, pp. 294–303, Mar. 2021.
- [43] Y. E. Sagduyu, T. Erpek, and Y. Shi, "Adversarial machine learning for 5G communications security," 2021, *arXiv:2101.02656*.
- [44] Y. Shi, Y. E. Sagduyu, T. Erpek, and M. C. Gursoy, "How to attack and defend 5G radio access network slicing with reinforcement learning," 2021, *arXiv:2101.05768*.
- [45] Y. Shi and Y. E. Sagduyu, "Adversarial machine learning for flooding attacks on 5G radio access network slicing," in *Proc. IEEE Int. Conf. Commun. Workshops (ICC Workshops)*, Jun. 2021, pp. 1–6.
- [46] M. Sadeghi and E. G. Larsson, "Physical adversarial attacks against end-to-end autoencoder communication systems," *IEEE Commun. Lett.*, vol. 23, no. 5, pp. 847–850, May 2019.
- [47] T. O'Shea and J. Hoydis, "An introduction to deep learning for the physical layer," *IEEE Trans. Cogn. Commun. Netw.*, vol. 3, no. 4, pp. 563–575, Dec. 2017.
- [48] Y. E. Sagduyu, Y. Shi, and T. Erpek, "IoT network security from the perspective of adversarial deep learning," in *Proc. 16th Annu. IEEE Int. Conf. Sens., Commun., Netw. (SECON)*, Jun. 2019, pp. 1–9.
- [49] B. R. Manoj, M. Sadeghi, and E. G. Larsson, "Adversarial attacks on deep learning based power allocation in a massive MIMO network," 2021, *arXiv:2101.12090*.
- [50] B. Kim, Y. Shi, Y. E. Sagduyu, T. Erpek, and S. Ulukus, "Adversarial attacks against deep learning based power control in wireless communications," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2021.
- [51] B. Kim, Y. Sagduyu, T. Erpek, and S. Ulukus, "Adversarial attacks on deep learning based mmWave beam prediction in 5G and beyond," in *Proc. IEEE Stat. Signal Process. Workshop (SSP)*, Jul. 2021, pp. 1–5.
- [52] Q. Liu, J. Guo, C.-K. Wen, and S. Jin, "Adversarial attack on DL-based massive MIMO CSI feedback," 2020, *arXiv:2002.09896*.
- [53] T. J. O'Shea, T. Roy, and T. C. Clancy, "Over-the-air deep learning based radio signal classification," *IEEE J. Sel. Topics Signal Process.*, vol. 12, no. 1, pp. 168–179, Feb. 2018.
- [54] Y. Shi, K. Davaslioglu, Y. E. Sagduyu, W. C. Headley, M. Fowler, and G. Green, "Deep learning for RF signal classification in unknown and dynamic spectrum environments," in *Proc. IEEE Int. Symp. Dyn. Spectr. Access Netw. (DySPAN)*, Nov. 2019, pp. 1–10.
- [55] S. Soltani, Y. E. Sagduyu, R. Hasan, K. Davaslioglu, H. Deng, and T. Erpek, "Real-time and embedded deep learning on FPGA for RF signal classification," in *Proc. IEEE Mil. Commun. Conf. (MILCOM)*, Nov. 2019, pp. 1–6.
- [56] K. Davaslioglu, S. Soltani, T. Erpek, and Y. E. Sagduyu, "DeepWiFi: Cognitive WiFi with deep learning," *IEEE Trans. Mobile Comput.*, vol. 20, no. 2, pp. 429–444, Feb. 2021.
- [57] Y. Wang, G. Gui, T. Ohtsuki, and F. Adachi, "Multi-task learning for generalized automatic modulation classification under non-Gaussian noise with varying SNR conditions," *IEEE Trans. Wireless Commun.*, vol. 20, no. 6, pp. 3587–3596, Jun. 2021.
- [58] Y. Wang *et al.*, "Automatic modulation classification for MIMO systems via deep learning and zero-forcing equalization," *IEEE Trans. Veh. Technol.*, vol. 69, no. 5, pp. 5688–5692, May 2020.
- [59] Y. Lin, Y. Tu, and Z. Dou, "An improved neural network pruning technology for automatic modulation classification in edge devices," *IEEE Trans. Veh. Technol.*, vol. 69, no. 5, pp. 5703–5706, May 2020.
- [60] A. Kurakin, I. Goodfellow, and S. Bengio, "Adversarial examples in the physical world," 2016, *arXiv:1607.02533*.
- [61] N. Carlini and D. Wagner, "Towards evaluating the robustness of neural networks," in *Proc. IEEE Symp. Secur. Privacy (SP)*, May 2017, pp. 39–57.
- [62] J. Maroto, G. Bovet, and P. Frossard, "SafeAMC: Adversarial training for robust modulation recognition models," 2021, *arXiv:2105.13746*.
- [63] L. Zhang, S. Lambotharan, G. Zheng, B. AsSadhan, and F. Roli, "Countermeasures against adversarial examples in radio signal classification," *IEEE Wireless Commun. Lett.*, vol. 10, no. 8, pp. 1830–1834, Aug. 2021.
- [64] B. Kim, Y. E. Sagduyu, T. Erpek, K. Davaslioglu, and S. Ulukus, "Adversarial attacks with multiple antennas against deep learning-based modulation classifiers," in *Proc. IEEE Globecom Workshops (GC Wkshps)*, Dec. 2020, pp. 1–6.

- [65] B. Kim, Y. E. Sagduyu, T. Erpek, K. Davaslioglu, and S. Ulukus, "Channel effects on surrogate models of adversarial attacks against wireless signal classifiers," in *Proc. IEEE Int. Conf. Commun.*, Jun. 2021, pp. 1–6.
- [66] B. Kim, Y. E. Sagduyu, K. Davaslioglu, T. Erpek, and S. Ulukus, "How to make 5G communications 'invisible': Adversarial machine learning for wireless privacy," in *Proc. Asilomar Conf.*, 2020, pp. 763–767.
- [67] T. J. O'Shea and N. West, "Radio machine learning dataset generation with GNU radio," in *Proc. 6th GNU Radio Conf.*, 2016, pp. 1–6.
- [68] N. Papernot, P. McDaniel, and I. Goodfellow, "Transferability in machine learning: From phenomena to black-box attacks using adversarial samples," 2016, *arXiv:1605.07277*.
- [69] J.-Y. Franceschi, A. Fawzi, and O. Fawzi, "Robustness of classifiers to uniform ℓ_p and Gaussian noise," in *Proc. 21st Int. Conf. Artif. Intell. Statist. (AISTATS)*, Feb. 2018, pp. 1280–1288.



Brian Kim (Graduate Student Member, IEEE) received the B.S. and M.S. degrees in electrical engineering from the Korea Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea, in 2014 and 2016, respectively. He is currently pursuing the Ph.D. degree with the Department of Electrical and Computer Engineering, University of Maryland, College Park. His research interests include wireless communications and machine learning.



Yalin E. Sagduyu (Senior Member, IEEE) received the Ph.D. degree in electrical and computer engineering from the University of Maryland, College Park, in 2007. He is currently the Director of networks and security at Intelligent Automation, a BlueHalo Company. He has been a Visiting Research Professor with the Department of Electrical and Computer Engineering, University of Maryland. His research interests include wireless communications, networks, security, machine learning, and 5G and 6G. He received the IEEE HST 2018 Best Paper Award. He chaired workshops at ACM MobiCom, ACM WiSec, IEEE CNS, and IEEE ICNP, and served as the Track Chair for IEEE PIMRC, IEEE GlobalSIP, and IEEE MILCOM.



Kemal Davaslioglu (Member, IEEE) received the Ph.D. degree in electrical and computer engineering from the University of California, Irvine. From 2012 to 2015, he held long-term internships at Broadcom Inc., Irvine, CA, USA, where he worked on 10-Gigabit Ethernet systems and later on the beamforming algorithms in the millimeter wave and 60 GHz channel characterization. At Intelligent Automation, a BlueHalo Company, he has developed several deep learning solutions for cognitive radio, adversarial learning, computer vision, synthetic data generation, and cybersecurity. He is currently a Senior Research Scientist at Intelligent Automation, a BlueHalo Company, where he is the Network and Information Systems Technical Area Lead. His research interests include deep learning for wireless networks and computer vision. He has extensive hands-on experience with deep learning software, its hardware implementation, and software defined radios. He was a recipient of the Best Paper Award at the IEEE Symposium on Technologies for Homeland Security (HST) in 2018.



wireless communications and networks, resource allocation, 5G and 6G, physical-layer security, cognitive radio, machine learning, and adversarial machine learning.

Tugba Erpek (Member, IEEE) received the Ph.D. degree in electrical engineering from Virginia Tech. She is currently a Lead Scientist at the Networks and Security Division, Intelligent Automation, a BlueHalo Company, where she is the Network Communications Technical Area Lead. She is also an Adjunct Research Professor at the Hume Center, Virginia Tech. She has been developing machine/deep learning algorithms to improve the performance, situational awareness, and security of wireless communications systems. Her research interests include



Sennur Ulukus (Fellow, IEEE) received the B.S. and M.S. degrees in electrical and electronics engineering from Bilkent University and the Ph.D. degree in electrical and computer engineering from the Wireless Information Network Laboratory (WINLAB), Rutgers University.

She is the Anthony Ephremides Professor in Information Sciences and Systems with the Department of Electrical and Computer Engineering, University of Maryland, College Park, where she also holds a joint appointment with the Institute for Systems

Research (ISR). Prior to joining UMD, she was a Senior Technical Staff Member at AT&T Labs-Research. Her research interests include information theory, wireless communications, machine learning, signal processing, and networks, with recent focus on private information retrieval, age of information, machine learning for wireless, group testing, distributed coded computation, energy harvesting communications, physical layer security, and wireless energy and information transfer.

Dr. Ulukus is a Distinguished Scholar-Teacher of the University of Maryland. She received the 2003 IEEE Marconi Prize Paper Award in Wireless Communications, the 2019 IEEE Communications Society Best Tutorial Paper Award, the 2020 IEEE Communications Society Women in Communications Engineering (WICE) Outstanding Achievement Award, the 2020 IEEE Communications Society Technical Committee on Green Communications and Computing (TCGCC) Distinguished Technical Achievement Recognition Award, the 2005 NSF CAREER Award, the 2011 ISR Outstanding Systems Engineering Faculty Award, and the 2012 ECE George Corcoran Outstanding Teaching Award. She was a Distinguished Lecturer of the IEEE Information Theory Society from 2018 to 2019. She is the TPC Chair of the 2021 IEEE GLOBECOM, and the TPC Co-Chair of 2019 IEEE ITW, 2017 IEEE ISIT, 2016 IEEE GLOBECOM, 2014 IEEE PIMRC, and 2011 IEEE CTW. She has been an Area Editor of the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS since 2019 and a Senior Editor of the IEEE TRANSACTIONS ON GREEN COMMUNICATIONS AND NETWORKING since 2020. She was an Area Editor of the IEEE TRANSACTIONS ON GREEN COMMUNICATIONS AND NETWORKING from 2016 to 2020, an Editor for the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS Series on Green Communications and Networking from 2015 to 2016, an Associate Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY from 2007 to 2010, and an Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS from 2003 to 2007. She was a Guest Editor for the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS in 2008, 2015, and 2021; *Journal of Communications and Networks* in 2012; and the IEEE TRANSACTIONS ON INFORMATION THEORY in 2011.