

# Physical Adversarial Attacks Against End-to-End Autoencoder Communication Systems

Meysam Sadeghi and Erik G Larsson

The self-archived postprint version of this journal article is available at Linköping University Institutional Repository (DiVA):

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-157535>

N.B.: When citing this work, cite the original publication.

Sadeghi, M., Larsson, E. G, (2019), Physical Adversarial Attacks Against End-to-End Autoencoder Communication Systems, *IEEE Communications Letters*, 23(5), 847-850.  
<https://doi.org/10.1109/LCOMM.2019.2901469>

Original publication available at:

<https://doi.org/10.1109/LCOMM.2019.2901469>

Copyright: Institute of Electrical and Electronics Engineers (IEEE)

<http://www.ieee.org/index.html>

©2019 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.



# Physical Adversarial Attacks Against End-to-End Autoencoder Communication Systems

Meysam Sadeghi and Erik G. Larsson,

**Abstract**—We show that end-to-end learning of communication systems through deep neural network (DNN) autoencoders can be extremely vulnerable to *physical* adversarial attacks. Specifically, we elaborate how an attacker can craft effective *physical* black-box adversarial attacks. Due to the openness (broadcast nature) of the wireless channel, an adversary transmitter can increase the block-error-rate of a communication system by orders of magnitude by transmitting a well-designed perturbation signal over the channel. We reveal that the adversarial attacks are more destructive than jamming attacks. We also show that classical coding schemes are more robust than autoencoders against both adversarial and jamming attacks. The codes are available at [1].

**Index Terms**—Adversarial attacks, Autoencoder systems, Deep learning, Wireless security, End-to-end learning, Security and Robustness of Deep Learning for Wireless Communications.

## I. INTRODUCTION

Deep neural networks (DNNs), due to their promising performance, are becoming an integral tool in many new disciplines [2]. This has created a new series of applications and concepts in wireless communications under a new paradigm, namely deep learning based wireless communications. One such concept is the end-to-end learning of communication systems using autoencoders [3]. As the name suggests, this approach enables end-to-end optimization and design of communication systems, which can potentially provide gains over the contemporary modularized design of these systems.

DNNs, despite their promising performance, are extremely susceptible to the so-called adversarial attacks [4], [5]. In adversarial attacks, the attacker adds a (small) perturbation  $\mathbf{p}$  to the input of a DNN that causes erroneous outputs [4]. In contrast to the conventional jamming attacks, this perturbation is not noise but a deliberately optimized vector in the feature space of the input domain that can fool the model. This property of DNN has raised major concerns regarding their security and robustness.

Adversarial attacks can be classified based on the nature of the attack. They can specifically be divided into white-box and black-box attacks, based on the amount of knowledge that the adversary has about the underlying NN [6]. In white-box attacks, the adversary has the full knowledge of the classifier, while in black-box attacks the adversary has no or limited knowledge of the classifier [6]. Adversarial attacks can also be classified into digital and physical attacks based on the degree of freedom that the adversary has with respect to its

access to the input of the system [7]. In digital attacks, the adversary can *precisely* design the input of the model, while in physical attacks, the adversary indirectly applies the input to the model [7].

Recently digital adversarial attacks on modulation classification was studied in [8]. Therein, digital black-box attacks were developed that can significantly reduce the performance of the DNN based classifier. However, given the inherent random nature of the wireless channel between the attacker and the receiver in a wireless system, the feasibility of a digital attack might be questioned. But physical attacks are more practical.

In this letter, we study *physical* adversarial attacks against end-to-end autoencoder communication systems, and provide the following contributions. First, we present new algorithms for crafting *effective physical* black-box adversarial attacks.<sup>1</sup> Second, we reveal that the adversarial attacks are more destructive than jamming attacks. Third, we show that classical communication schemes are more robust than the end-to-end autoencoder systems.

## II. SYSTEM MODEL

An autoencoder communication system has three main blocks, 1) transmitter, 2) channel, and 3) receiver [3]. The system receives an input message  $s \in \mathcal{M} = \{1, 2, \dots, M\}$ , where  $M = 2^k$  is the dimension of  $\mathcal{M}$  with  $k$  being the number of bits per message. The message is then passed to the transmitter, where it applies a transformation  $f : \mathcal{M} \rightarrow \mathbb{R}^{2n}$  to the message  $s$  to generate the transmitted signal  $\mathbf{x} = f(s) \in \mathbb{R}^{2n}$ . Similar to [3], the transmitter enforces an average power constraint  $\mathbb{E}[x_i^2] \leq 0.5 \forall i$ . Next, the signal  $\mathbf{x}$  is sent to the receiver using the channel  $n$  times. In this work, we consider an additive white Gaussian noise (AWGN) channel. Then, the receiver receives a signal  $\mathbf{y}$ , which contains the original signal  $\mathbf{x}$ , noise and other sources of interference. The receiver applies the transformation  $g : \mathbb{R}^{2n} \rightarrow \mathcal{M}$  to create  $\hat{s} = g(\mathbf{y})$  which is the estimate of the transmitted message  $s$ . To enable the comparability of the results developed in different sections, we set  $n = 7$  and  $k = 4$ . This also enables us to compare our results with the classical Hamming code (7,4).

In deep learning terminology the transmitter and receiver are respectively called encoder and decoder and they are implemented using neural networks [2]. The input signal  $s$  first passes through the encoder. Then the output of the encoder

<sup>1</sup>Algorithmically, the method for perturbation generation suggested here departs significantly from that in [8]. There is a superficial similarity in that both use the SVD as a building block, but the problems and solutions are different.

<sup>2</sup>Note that the output of the transmitter is an  $n$  dimensional complex vector which is transformed to a  $2n$  real vector.

The authors (m.sadeghee@gmail.com, erik.g.larsson@liu.se) are with Department of Electrical Engineering (ISY), Linköping University, Linköping, Sweden.

The codes are available at [1].

Block Name	MLP Autoencoder		CNN Autoencoder	
	Layer Name	Output Dim.	Layer Nmae	Output Dim.
Encoder	Input	M	Input	M
	Dense+eLU	M	Dense + eLU	M
	Dense+Linear	2n	conv1d+Flattening	16 × M
	Normalization	2n	Dense + Linear	2n
Channel	noise (+ perturbation)	2n	Normalization	2n
			noise (+ perturbation)	2n
Decoder	Dense + ReLU	M	conv2d	16 × 2n
	Dens+Softmax	M	conv2d+Flattening	8 × 2n
			Dense + ReLU	2M
			Dens + Softmax	M

Table I: The structure of considered autoencoders.

( $\mathbf{x}$ ) goes through the channel, where the attacker perturbation signal ( $\mathbf{p}$ ) and the noise ( $\mathbf{n}$ ) are added to the  $\mathbf{x}$ . Therefore the output of the channel will be  $\mathbf{y} = \mathbf{x} + \mathbf{p} + \mathbf{n}$ . Then  $\mathbf{y}$  is fed to the decoder. This procedure is illustrated in Fig. 1.

To enable a benchmark for comparison, we use the same autoencoder structure as [3], which is an MLP autoencoder.<sup>3</sup> Also, we use a convolutional neural network (CNN) based autoencoder to introduce the black-box adversarial attacks. The structure of these networks is given in Table I.<sup>4</sup> Detailed explanations of each element of Table I can be found in [2], [3]. We train the autoencoder in an end-to-end manner using the Adam optimizer, on the set of all possible messages  $s \in \mathcal{M}$ , using the cross-entropy loss function.

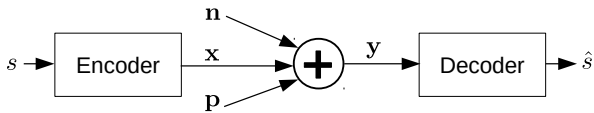


Fig. 1: An adversarial attack against an autoencoder system.

### III. CRAFTING A WHITE-BOX ADVERSARIAL ATTACK

Given the structure of an autoencoder in Fig. 1, the adversary uses the broadcast nature of the channel and attacks the decoder block of the autoencoder communication system. To model the attack mathematically, let us denote the underlying DNN classifier at the decoder by  $g: \mathbb{R}^{2n} \rightarrow \mathcal{M}$ . Also denote  $\mathbf{w} = \mathbf{x} + \mathbf{n}$ . The classifier generates an output  $\hat{s} = g(\mathbf{y})$  for every input  $\mathbf{y} = \mathbf{w} + \mathbf{p}$ . Note that when the attacker is absent  $\mathbf{p} = \mathbf{0}$  and  $\mathbf{w} = \mathbf{y}$ . Given these definitions, an adversarial perturbation  $\mathbf{p}$  for the considered autoencoder is [4], [5]

$$\begin{aligned} \min_{\mathbf{p}} \quad & \|\mathbf{p}\|_2 \\ \text{s.t.} \quad & g(\mathbf{w} + \mathbf{p}) \neq g(\mathbf{w}) \end{aligned} \quad (1)$$

where  $\mathbf{n} \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbf{I})$  with  $\sigma^2 = N_0/2RE_b$ ,  $E_b$  is energy per bit, and  $R = k/n$  is the rate in bits per channel use.

<sup>3</sup>The selected set of hyperparameters is the same as in [3], to enable a direct comparison, and we do not claim their optimality. There are infinitely many combinations of hyperparameters that could be tested, but modifying them is not likely to change the conclusions, as DNNs are known to be inherently vulnerable to adversarial attacks, see [4], [7], [9].

<sup>4</sup>We also investigated some variations of the network structure, e.g., ReLU activation, depth, different number of neurons. The results are quantitatively similar. All the codes are available at [1].

### Algorithm 1: Crafting Physical Adversarial Perturbations

Inputs:

- The full autoencoder's model, e.g.,  $g$  and the parameters of the decoder's underlying NN.
- The desired perturbation power  $p_{power}$ .
- The variance of the channel noise  $\sigma^2$ .

Output: An input-agnostic adversarial perturbation, i.e.,  $\mathbf{p}$

1: Initialize:

- Set *number-of-samples*  $\leftarrow 10$ .
- Set  $\mathbf{p} \leftarrow \mathbf{0}$ .

2: **for**  $i$  in *range(number-of-samples)* **do**

3: Choose an  $s \in \{1, 2, \dots, M\}$  uniformly at random.

4: Create a random noise  $\mathbf{n} \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbf{I})$ .

5: Set  $\hat{s} \leftarrow g(f(s) + \mathbf{p} + \mathbf{n})$ .

6: **if**  $\hat{s} == s$  **then**

7: Set  $\mathbf{w} \leftarrow \mathbf{w} + \mathbf{p}$ .

8: Solve (1) and denote its solution by  $\mathbf{p}_{update}$ .

9: **if**  $\|\mathbf{p} + \mathbf{p}_{update}\|_2^2 \leq p_{power}$  **then**

10:  $\mathbf{p} \leftarrow \mathbf{p} + \mathbf{p}_{update}$

11: **else**

12:  $\mathbf{p} \leftarrow \sqrt{p_{power}} \frac{\mathbf{p} + \mathbf{p}_{update}}{\|\mathbf{p} + \mathbf{p}_{update}\|_2}$

13: **end if**

14: **end if**

15: **end for**

Solving (1) is challenging as  $g$  does not have a convex structure [4]. Therefore, a common approach is to use the so-called fast gradient method (FGM) [4], [5], which finds an approximate solution to (1). The main idea is as follows. For the input label  $s \in \mathcal{M}$ , denote the loss function of the autoencoder by  $L(\mathbf{w} + \mathbf{p}, s)$ . Then, FGM uses a Taylor expansion of the loss function, i.e.,  $L(\mathbf{w} + \mathbf{p}, s) \approx L(\mathbf{w}, s) + \mathbf{p}^T \nabla_{\mathbf{w}} L(\mathbf{w}, s)$ , and sets  $\mathbf{p} = \alpha \nabla_{\mathbf{w}} L(\mathbf{w}, s)$ , where  $\alpha$  is a scaling coefficient. Hence, FGM effectively increases the loss function and provides an approximation of  $\mathbf{p}$ , given that the input ( $s$ ) is known.<sup>5</sup>

An adversarial attacker cannot directly apply the FGM method to attack an autoencoder communication system, as the FGM method requires knowledge of the input to the autoencoder, i.e.,  $s$ , to create a perturbation  $\mathbf{p}$ . However, the attacker does not know what symbol is being transmitted. To address this issue, Alg. 1 presents an iterative method to craft a universal (i.e., input-agnostic) perturbation  $\mathbf{p}$  that can fool the autoencoder independently of its input. The main idea comes from the literature on computer vision [9], where it has been shown that by iteratively finding image-specific perturbations one can create image-agnostic perturbations.

In the first line of the Alg. 1, we initialize the number of samples and the perturbation  $\mathbf{p}$ . The number of samples determines the number of iterations we use in order to create a universal perturbation. Then, we repeatedly apply the following steps. First, we select an input  $s$  uniformly at random and pass it through the encoder. Next, we compute the signal seen by the decoder, by adding the perturbation and the noise

<sup>5</sup>Further details and more advanced methods can be found in [4], [5], [8].

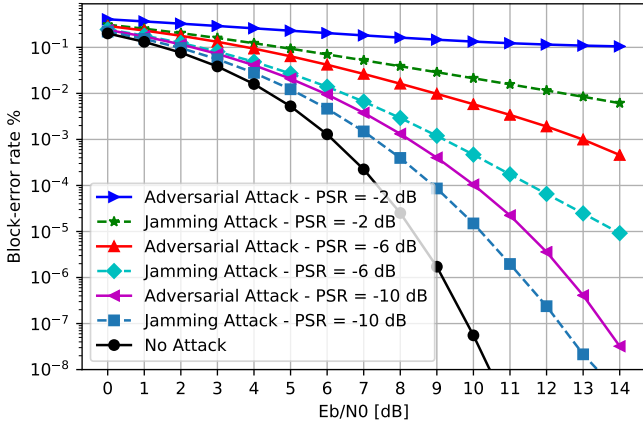


Fig. 2: BLER versus  $E_b/N_0$  under Adversarial and Jamming attacks against the MLP autoencoder in Table I.

to the output of encoder. Now we apply the decoder block and obtain an estimate of the input, i.e.,  $\hat{s}$ . If the decoder correctly classifies the input message, we set  $\mathbf{w} + \mathbf{p}$  as  $\mathbf{w}$ . Using the new value of  $\mathbf{w}$  and by solving (1), we search for an update perturbation  $\mathbf{p}_{update}$  that enforces misclassification of the new input. Then we update the existing perturbation  $\mathbf{p}$  with  $\mathbf{p}_{update}$  and check the power of the updated perturbation vector. If it is below the desired perturbation power we keep it, otherwise, we normalize it to meet the power constraint. Using this iterative procedure, we are able to craft a universal perturbation  $\mathbf{p}$ , which is independent of the input of the autoencoder.

Figure 2 presents the block-error-rate (BLER) performance of the MLP autoencoder of Table I, under adversarial attack. The adversarial attack is designed using Alg. 1. To compare the power of the adversarial perturbation at the receiver with the received signal power, we define a metric called perturbation-to-signal ratio (PSR), which is equal to the ratio of the received perturbation power to the received signal power. For the sake of comparison, we also consider a classical jamming attack, where a jammer creates Gaussian noise with the same PSR as the adversarial attack. Gaussian jamming is the most effective way to reduce the mutual information for a message encoded with a capacity-achieving Gaussian codebook (which is optimum for the AWGN channel), under a given jamming power constraint. Based on Fig. 2 it is clear that even for small PSR values, the adversarial attack can degrade the BLER by orders of magnitude. Also, the adversarial attack has a more destructive impact compared to the jamming attack.

#### IV. CRAFTING SHIFT-INVARIANT BLACK-BOX ATTACKS

In Section III, we considered two restrictive assumptions in order to craft adversarial perturbations. First, we assumed that the adversarial attacker has perfect knowledge of the autoencoder system, including the number of layers, weight and bias parameters of the decoder block. Second, the attacker is synchronous with the transmitter. More precisely, given that  $\mathbf{p}$  and  $\mathbf{x}$  are  $2n$  dimensional real vectors ( $n$ -dimensional complex vectors), then  $\mathbf{p} + \mathbf{x} + \mathbf{n}$  requires each element of  $\mathbf{p}$  to be added by its corresponding element of  $\mathbf{x}$ . We

#### Algorithm 2: Crafting Shift-Invariant Perturbations

- 1: Using the substitute network, generate  $I$  adversarial perturbations using Alg. 1.
- 2: Calculate the BLER of a randomly shifted version of each of the  $I$  perturbations on the substitute network.
- 3: Select the first  $t$  perturbations associated with the  $t$  least BLERs. Denote them as  $\{\mathbf{p}_1, \dots, \mathbf{p}_t\}$ .
- 4: Set  $\mathbf{P}_{norm} = \left[ \frac{\mathbf{p}_1}{\|\mathbf{p}_1\|_2}, \frac{\mathbf{p}_2}{\|\mathbf{p}_2\|_2}, \dots, \frac{\mathbf{p}_t}{\|\mathbf{p}_t\|_2} \right]^T$ .
- 5: Calculate the SVD of  $\mathbf{P}_{norm}$  as  $\mathbf{P}_{norm} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ .
- 6: Select the first column of  $\mathbf{V}$  as the candidate shift-invariant perturbation, i.e.,  $\mathbf{p}_{si} = \mathbf{V} \mathbf{e}_1$ .

address these problems by applying a heuristic approach and the transferability of adversarial attacks.

Consider the CNN autoencoder of Table I, and assume an attacker is interested to create an adversarial attack for it, without having any knowledge about its structure or parameters. The transferability of adversarial attacks says that attacks designed for a specific model are also effective for other models, with high probability [6]. Therefore, the attacker can use its own substitute autoencoder communication system and then design a white-box attack for it, as it has the perfect knowledge of this substitute autoencoder communication system. Then use the designed perturbation to attack the original unknown model. This attack is called a black-box adversarial attack [6].

Using this approach, we consider the MLP autoencoder in Table I as the substitute model and create an adversarial perturbation for it. Then, we use the so-created perturbation to attack the CNN autoencoder in Table I, which is unknown to the attacker. Note that this approach is general and we can use it to attack other autoencoders with different structures and parameters.

In order to remove the synchronicity requirement between the attacker and the transmitter, we present a heuristic algorithm, Alg. 2, to create attacks which are robust against random time shifts. This is a challenging task due to the considered setup:  $n=7$  defines a low-dimensional input, whereas generally, finding effective adversarial attacks is easier the larger the dimensions of the input and the model are [4].

The main idea is as follows. Using the substitute network, first we generate a pool of adversarial perturbations, e.g.,  $I$  adversarial perturbations, using Alg. 1. Then, for each of these  $I$  perturbations, we calculate the BLER of a randomly shifted version of them. Next, we rank them based on the severity of their attacks. Using this approach, the top  $t$  perturbations show a robustness against random shifts. Note that each of these  $t$  perturbations represents a direction in the feature space that even a randomly shifted version of it causes significant misclassification. Therefore, we can use a singular value decomposition (SVD) of these  $t$  vectors to find their main principal direction which hopefully would show a better shift-invariant property. More details are given in Alg. 2. Using the proposed algorithm, we are able to craft adversarial attacks which are robust against random shifts, and therefore a synchronous attack is not required.

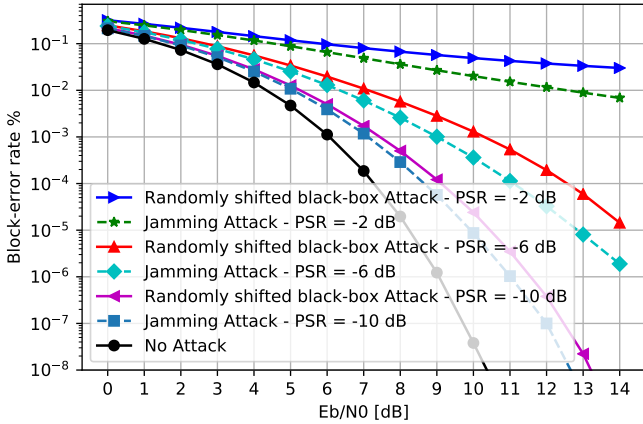


Fig. 3: BLER versus  $E_b/N_0$  for randomly shifted black-box adversarial attacks against the CNN autoencoder in Table I.

Using a black-box attack and the heuristic method proposed in Alg. 2, we attack the CNN autoencoder of Table I, while we use the parameters of the MLP autoencoder of Table I. The result is shown in Fig. 3. Fig. 3 reveals two important properties of non-synchronous black-box *physical* adversarial attacks against end-to-end autoencoder communication systems. First, they are significantly effective. Second, they can provide a more destructive effect compared to jamming attacks.<sup>6</sup>

## V. ROBUSTNESS OF AUTOENCODERS VS CLASSICAL APPROACHES

In this section, we compare the robustness of the end-to-end autoencoders with a classical modulation and coding scheme. In [3], it has been shown that the end-to-end autoencoder provides roughly the same BLER as binary phase-shift keying (BPSK) modulation combined with a Hamming (7,4) code and maximum-likelihood decoding (MLD). This also can be verified from Fig. 4. We use BPSK modulation and Hamming (7,4) coding with MLD as a benchmark and compare the robustness of the CNN autoencoder in Table I with it. For the attack, we use both the jamming attack and the shift invariant black-box attack of Section III. The results are presented in Fig. 4, where the PSR for both attacks is equal to  $-6$  dB.

From Fig. 4, we make the following observations. First, the randomly shifted black-box attack causes a bigger BLER than a jamming attack for the considered autoencoder system. Second, the BPSK modulation with Hamming coding and MLD results in roughly the same BLER for both adversarial and jamming attacks. Third, the classical BPSK modulation with Hamming coding and MLD provides a smaller BLER than the considered CNN autoencoder under a jamming attack. Therefore, Fig. 4 suggests that the performance of the classical approach is more robust, compared to autoencoders, against both adversarial attacks and jamming attacks.

<sup>6</sup>In numerical experiments not included here due to space limitations, we simulated 10 different models (with different weights and biases) and applied the same attack to them. We observed the same trend as before, suggesting that our conclusions are general and not specific to a particular model [1].

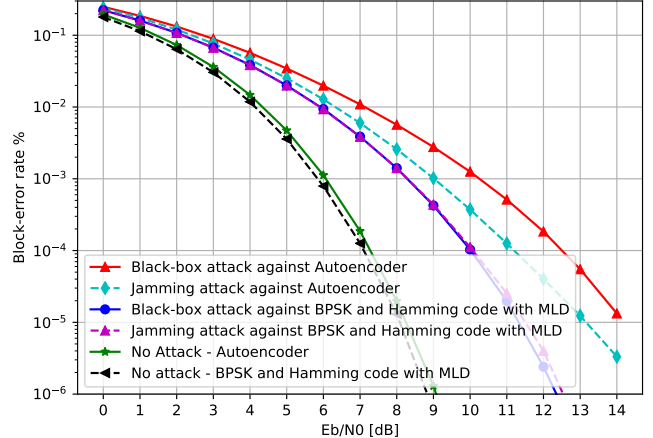


Fig. 4: Comparison of the robustness of the autoencoder (CNN, see Table I) and classical transmission (BPSK modulation with Hamming coding). The PSR is set to  $-6$  dB.

## VI. CONCLUSIONS, FUTURE WORKS AND FURTHER DISCUSSIONS

We showed end-to-end autoencoder systems are significantly vulnerable to small perturbations, which can decrease their BLER by orders of magnitude. Due to the broadcast nature of the wireless channel, creating such perturbations is easy for an adversarial attacker. We also presented algorithms to craft such *physical* adversarial attacks and verified them by simulations. These findings suggest that defense mechanisms against adversarial attacks and further research on the security and robustness of deep-learning based wireless systems is a necessity.

One possible defense mechanism is to train the autoencoder with adversarial perturbations (which is a technique to increase robustness and is known as adversarial training [4]). However, adversarial training has its own drawbacks, e.g., the attacker may use a different attack than the one used for training the network. Also, the attacker can design adversarial perturbations for an autoencoder that already has been trained with adversarial training, and craft new adversarial perturbations. Also, adversarial training can reduce the performance of the autoencoder on clean inputs.

We also compared the robustness of a CNN-based end-to-end autoencoder system with BPSK modulation and Hamming coding and MLD. We showed that the classical approach provides more robust performance against both adversarial and jamming attacks, compared to the end-to-end autoencoder.

Our results were obtained for physical adversarial attacks in AWGN channels, considering AWGN jamming. Further studies are required to investigate other channel models (e.g., Rayleigh fading), a wider set of hyperparameters and information rates (e.g., we only studied  $n = 7$  and  $k = 4$ ), and comparisons against other, more advanced, jamming strategies (see, e.g. [10]). For example, it is an open problem to determine more exactly, quantitatively, how the gap between the jamming attack and the adversarial attack behaves as function of the information rate. Also, the extension to more advanced channel models, e.g., Rayleigh fading, is an interesting direc-

tion for future work. In that case, the attacker must have an estimate of the channel between himself and the receiver. The quality of this estimate may fundamentally affect the degree of success of the attacker. We hope that our initial results will stimulate future research on this topic.

#### REFERENCES

- [1] M. Sadeghi, "Security and robustness of DL in wireless communication systems," <https://github.com/meysamsadeghi/Security-and-Robustness-of-Deep-Learning-in-Wireless-Communication-Systems>, 2019.
- [2] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>.
- [3] T. O'Shea and J. Hoydis, "An introduction to deep learning for the physical layer," *IEEE Transactions on Cognitive Communications and Networking*, vol. 3, no. 4, pp. 563–575, Dec. 2017.
- [4] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv preprint arXiv:1412.6572*, 2014.
- [5] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.
- [6] N. Papernot, P. D. McDaniel, and I. J. Goodfellow, "Transferability in machine learning: from phenomena to black-box attacks using adversarial samples," *CoRR*, vol. abs/1605.07277, 2016.
- [7] A. Kurakin, I. Goodfellow, and S. Bengio, "Adversarial examples in the physical world," *arXiv preprint arXiv:1607.02533*, 2016.
- [8] 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.
- [9] S. M. Moosavi-Dezfooli, A. Fawzi, O. Fawzi, and P. Frossard, "Universal adversarial perturbations," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jul. 2017, pp. 86–94.
- [10] X. Liu, Y. Xu, L. Jia, Q. Wu, and A. Anpalagan, "Anti-jamming communications using spectrum waterfall: A deep reinforcement learning approach," *IEEE Commun. Lett.*, vol. 22, no. 5, pp. 998–1001, May 2018.