Backdoor Attacks Against Deep Learning-based Massive MIMO Localization

Tianya Zhao[†], Xuyu Wang[†], Shiwen Mao[‡]
[†]Knight Foundation School of Computing and Information Sciences, Florida International University, Miami, FL 33199, USA
[‡]Department of Electrical and Computer Engineering, Auburn University, Auburn, AL 36849, USA
Emails: tzhao010@fiu.edu, xuywang@fiu.edu, smao@ieee.org

GLOBECOM 2023 - 2023 IEEE Global Communications Conference | 979-8-3503-1090-0/23/$31.00 ©2023 IEEE | DOI: 10.1109/GLOBECOM54140.2023.10437534

***Abstract*—Millimeter wave (mmWave) communications and massive MIMO play crucial roles in the development of future wireless systems. In addition to offering high data rates, these technologies enable the realization of high-precision localization systems, especially in complicated indoor rich multi-path environments without GPS coverage. While deep neural networks (DNNs) enable high accuracy in fingerprint-based indoor localization, their implementations also introduce security problems. In the field of computer vision, backdoor attacks have proven to be able to effectively deceive models using specific or imperceptible triggers. In this paper, we study the impact of backdoor attacks on 5G massive MIMO localization systems in both indoor and outdoor environments. Two different triggers are investigated: the one-pixel trigger (visible) and the random noise trigger (invisible). We evaluate the localization systems using a public dataset and demonstrate that DNN-based localization systems are vulnerable to backdoor attacks.**

***Index Terms*—Backdoor Attack, Deep Learning, Massive MIMO, Wireless Localization.**

I. Introduction

With the growing popularity of smart cities, smart homes, and autonomous driving, there is a rising emphasis on location-based applications. Precise tracking of individuals and equipment plays a vital role in these contexts. With the help of the powerful modeling capabilities of deep neural networks (DNNs), current data-driven localization systems can achieve high-precision location predictions by collecting and analyzing position-related wireless data.

Global Positioning System (GPS) is a widely recognized technology that enables people to navigate themselves. However, GPS encounters challenges when it comes to indoor environments, primarily due to its sensitivity to occlusion. To enable accurate indoor localization, alternative wireless technologies such as WiFi have gained significant attention. WiFi signals are pervasive in indoor environments, making them suitable for wireless positioning. In addition to WiFi, millimeter wave (mmWave) and massive multiple input multiple output (MIMO) technologies have been recognized as key technologies for the next generation wireless networks [1]. In comparison to WiFi-based localization systems, mmWave/massive MIMO-based positioning systems offer the potential for higher precision due to the higher data rates and multiplexing gains.

979-8-3503-1090-0/23/$31.00 © 2023 IEEE

In measurement-based wireless positioning systems, time of arrival (TOA), time difference of arrival (TDOA), angle of arrival (AOA), and received signal strength indicator (RSSI) are widely used to capture location information [2], [3]. In deep learning-driven approaches, RSSI and channel state information (CSI) data have been commonly employed to train DNNs for wireless localization. The measured CSI contains valuable information such as scattering, delays, and fading at a fine-grained level, making it capable of accurate localization. DeepFi is the first work that deploys deep learning models in indoor localization systems, which leverages WiFi CSI to achieve a desirable performance [4]. Moreover, convolutional neural networks (CNNs) are leveraged to achieve high localization accuracy as well. For example, DyLoc proposed a CNN-based mmWave/massive MIMO localization system in dynamic scenarios using angle delay profile (ADP) as input, which is a linear transformation of the CSI [5].

Indeed, deep learning-based localization systems exhibit better performance owing to the powerful capabilities of deep learning models. Nevertheless, the robustness of such systems is a critical issue due to the black-box feature of DNNs. In the field of computer vision, even minor perturbations can significantly degrade a model's performance on specific images [6]. Even worse, universal perturbations can deceive a DNN on any image and make the model completely ineffective [7]. In addition to the above evasion attacks, there exists another type of attack known as backdoor attacks. Backdoor attacks insert a trigger into the input data. The model performs normally on benign samples but will behave maliciously as intended when the trigger is activated [8].

The open nature of wireless signals makes the research on the resilience of deep learning-based wireless systems more important. Bahramali et al. demonstrate that DNN-based wireless communication systems are vulnerable to adversarial examples and remain susceptible to attacks even when well-designed defense mechanisms are deployed [9]. In the domain of wireless localization, our previous work has verified that WiFi-based and 5G-based wireless localization systems are susceptible to adversarial attacks [10], [11]. In fact, these studies underscore the importance of addressing the robustness of deep learning-based wireless localization systems against potential adversaries.

To the best of our knowledge, there is a lack of research in the existing literature concerning *backdoor attacks on deep learning-based wireless positioning systems*. Given that mmWave and massive MIMO technologies are crucial for present and future wireless systems, and the robustness of these highly accurate positioning systems is significant, this paper aims to evaluate the impact of backdoor attacks on deep learning-based mmWave/massive MIMO indoor and outdoor localization systems. We use the ADP in massive MIMO systems as our input data and model the localization problem as a regression task in deep learning. We then introduce backdoor attacks with different triggers on the DNN-based positioning systems and show that the DNN models are susceptible to backdoor attacks both in indoor and outdoor environments.

The major contributions made in this paper are summarized as follows.

- To the best of our knowledge, this is the first work to study backdoor attacks on mmWave/massive MIMO indoor and outdoor localization systems.
- We design two different triggers for backdoor attacks. The one-pixel trigger is fixed at a single point, while the random noise trigger remains invisible and has the same shape as the input. Both triggers have been shown to be capable of launching backdoor attacks.
- We experimentally demonstrate that backdoor attacks can greatly degrade the performance of positioning systems in both indoor and outdoor environments when the triggers are activated.

The remainder of this paper is organized as follows. Related work is discussed in Section II. Section III introduces the system design. Comprehensive experiments are presented in Section IV. Section V concludes this paper.

II. Related Work

Extensive research has been conducted on various attacks in wireless systems, aiming to enhance the security and resilience of wireless applications. For wireless human activity recognition (HAR) systems, jamming is a common method to attack wireless systems by interfering with the transmission and reception processes [12]. IS-WARS injected interference signals by coexisting protocols in HAR systems, which was more stealthy and harder to detect [13].

With the rapid development of deep learning, there is a growing focus on addressing security concerns associated with the models employed in these systems. For instance, WiCAM generated adversarial examples with limited perturbations by learning the temporal and spatial information. It can drastically degrade DNN-based WiFi sensing systems [14]. Ambalkar et al. studied three white-box attacks on CSI-based HAR systems [15]. Adar presents optimization-driven adversarial examples and explores adversarial training on sensor-based HAR [16]. Yang et al. investigated white-box attacks in Doppler-based HAR systems [17]. In addition to studying how to attack the system, SecureSense proposed a defense framework for device-free HAR systems [18].

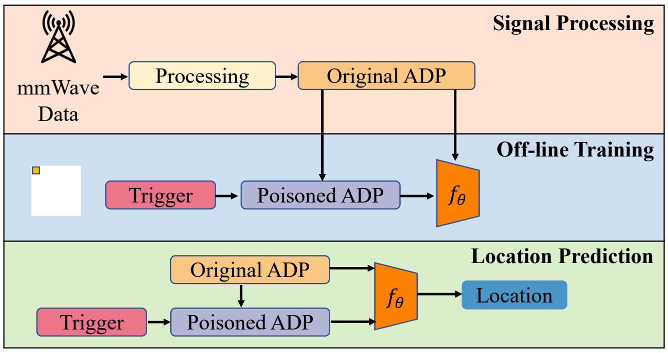

Fig. 1: Backdoor attacks on DNN-based mmWave/massive MIMO localization systems.

Three classic white-box attacks were studied in the context of WiFi-based floor classification and indoor localization in [10]. The authors comprehensively evaluated black-box and white-box attacks in indoor localization system, and examined the efficacy of adversarial training as a defense method [19].

Our work in this paper differs from the previous works in two key aspects. First, we extend the scope to both indoor and outdoor 5G localization scenarios. Second, we investigate backdoor attacks on mmWave/massive MIMO positioning systems, which have not been well studied before.

III. System Design

Data-driven localization systems are mostly based on fingerprinting, which consists of two stages: the off-line phase and the on-line phase. In the off-line phase, which is also known as the training stage, a large training dataset is constructed, and the deep learning models is trained on the dataset. During the on-line phase, the trained models are deployed to predict the location of target devices using newly collected data.

A. System Overview.

Fig. 1 presents an overview of backdoor attacks against the DNN-based mmWave/massive MIMO localization system. Poisoned ADP data is generated by directly adding a trigger to the original ADP, while the only difference between the original inputs and the poisoned inputs lies in the presence of the trigger. The number of poisoned inputs constitutes only a small fraction of the original inputs, which is a variable that can be adjusted accordingly by the attacker. During the training process, the original ADP data serve as benign inputs to optimize the model's performance, while the poisoned inputs are designed to deceive the model into making inaccurate location predictions. Once the training process is completed, the model gains the ability to predict locations. *With benign inputs, the model can accurately predict the locations of target devices. However, when triggers are injected into the inputs, the model's performance degrades significantly, leading to erroneous predictions.*

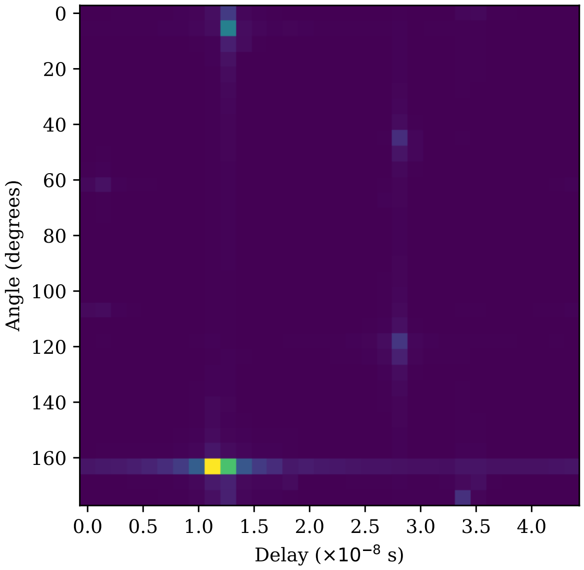

Fig. 2: An ADP image used in the massive MIMO localization system. Each pixel depicts the gain of the path with the corresponding AOA and delay.

B. Design Goals

This paper aims to achieve the following design goals in our study of backdoor attacks:

- *Effectiveness*. When the trigger is activated, the model should generate position predictions that deviate as significantly as possible from the correct locations.
- *Undetectable*. While maintaining the success rate of backdoor attacks and predictions on benign samples, minimizing the trigger value and reducing the number of poisoned samples contribute to reducing the probability of being detected.

C. Signal Processing

In this paper, we use ADP data for the massive MIMO positioning system, which is a linear transformation of the CSI. The ADP matrix $\mathbf{A}$ is computed by multiplying the channel CSI matrix with two discrete Fourier transform (DFT) matrices and taking the absolute values, as

$$\mathbf{A}_{z,q} = |\mathbf{V}_{N_t,N_t}\mathbf{H}\mathbf{F}_{N_c,N_c}|, \tag{1}$$

where V_{N_t,N_t} and F_{N_c,N_c} denote the DFT matrices, respectively, N_t is the number of antenna elements, N_c is the number of sub-carriers, and $\mathbf{H}$ is the CSI matrix.

As shown in Fig. 2, the ADP elements are presented as $[A]_{z,q}$ in a two-dimensional space, denoting the absolute gain of the zth AOA and the qth delay. More processing details can be found in [5], [11]. The ADPs obtained in the indoor and outdoor environments have varying shapes of size 32×32 and 64×64, respectively, due to differences in antennas and sub-carrier numbers.

TABLE I: Configuration of the CNN Models

Layer	*Kernel Size (I)*	*Stride (I)*	*Kernel Size (O)*	*Stride (O)*
1	16×16×1	2	32×32×1	2
2	8×8×4	2	16×16×4	2
3	7×7×16	1	8×8×8	2
4	5×5×32	1	7×7×16	1
5	3×3×64	1	5×5×32	1
6			3×3×64	1

D. CNN-based Localization

By using ADP as input, we can transform the localization problem into a regression problem. Therefore, we employ CNNs as our massive MIMO localization model. Table I presents the convolutional part of our CNN model. In this study, we have developed two distinct CNNs with slight modifications to cater to the specific requirements of indoor and outdoor localization tasks, respectively. For the indoor case, denoted as “I,” and the outdoor case, denoted as “O,” we customize the convolutional kernel sizes accordingly. The CNN designed for the outdoor scenario includes an additional convolutional layer to maintain consistency in the output dimensions with the indoor case. The architectures of our CNNs closely resemble those presented in [5]. Following the convolutional layers, three linear layers are employed to produce a two-dimensional location output.

E. Backdoor Attack

In the context of today’s deep learning landscape, the integration of various cloud platforms, pre-trained models, and public datasets has become indispensable. However, ensuring the security of such content poses significant challenges. Malicious attackers can introduce problematic datasets and pre-trained models, impairing the performance of inference tasks. Furthermore, attackers can invade the cloud infrastructure and manipulate gradients during the training process to disrupt model performance. Given these circumstances, backdoor attacks can be classified into three main types: poisoning-based backdoor attacks, weights-oriented backdoor attacks, and structure-modified backdoor attacks [20]. This paper concentrates on employing poisoning-based backdoor attacks on datasets, which is a commonly used and straightforward approach.

Let $D = (x_i, l_i)^N$ denote the benign dataset, where x_i represents the input data and l_i is the corresponding location. A poisoned dataset $P = (x_i + t, l_t)^M$, where $M \ll N$, is generated from the benign dataset. The poisoning rate

$$p \doteq \frac{M}{N} \tag{2}$$

is defined as the ratio of the number of poisoned samples to the number of benign samples.

As shown in Section III-B, we define the objective of backdoor attacks as

$$t^* = \arg\max_{t} \ d(F_\theta(x_i + t), F_\theta(x_i)), \quad s.t. \ \mid t \mid \leq \varepsilon, \tag{3}$$

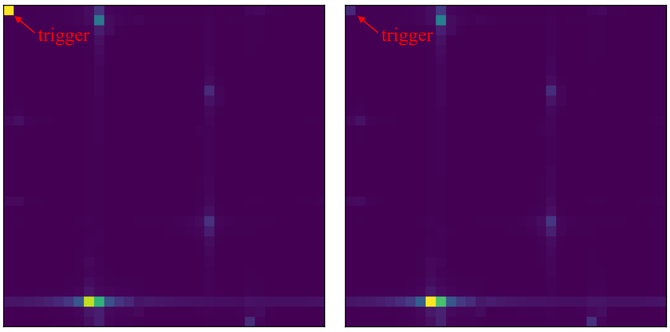

(a) Trigger value = 0.01

(b) Trigger value = 0.001

Fig. 3: Indoor ADP under the one-pixel attack. The trigger is always introduced at position (0, 0).

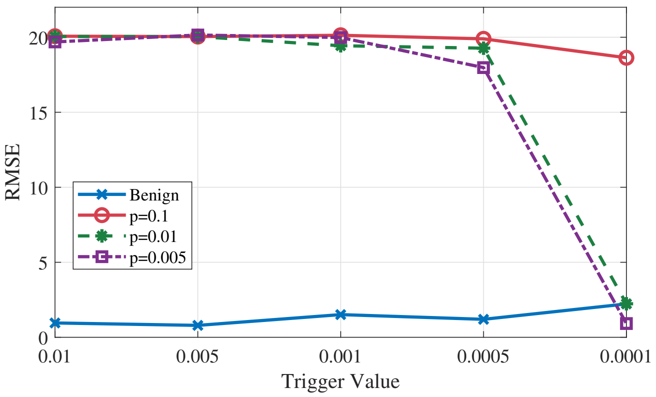

Fig. 5: RMSE of one-pixel attack on indoor localization.

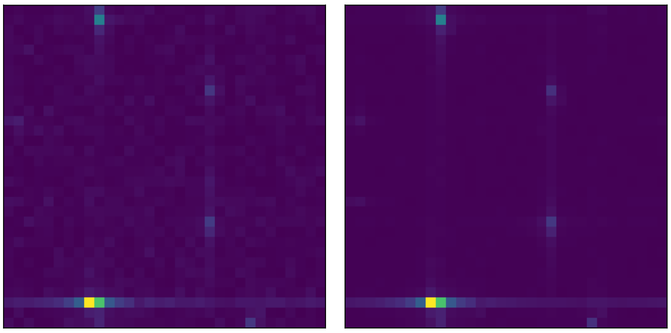

(a) $\mu = 10^{-4}, \sigma = 10^{-4}$

(b) $\mu = 10^{-5}, \sigma = 10^{-5}$

Fig. 4: Indoor ADP under the random noise attack. The right trigger is random and not visible.

where F_θ represents the CNN model, and $d(\cdot)$ denotes a distance function. In this paper, we employ the Euclidean distance as the distance function. The primary goal is to maximize the spatial distances between the benign input x_i and the poisoned input $x_i + t$, while simultaneously restricting the trigger value to be upper bounded by a specific value ε. By pursuing this objective, we intend to emphasize the divergence in the spatial distances when the trigger is present.

In this paper, we explore two types of triggers: the one-pixel trigger and the random noise trigger. The one-pixel trigger involves adding a specific value to the position $(0, 0)$ of inputs, as depicted in Fig. 3. To make the trigger more invisible, we also consider adding random noise with different mean values and standard deviations as trigger to the input, as shown in Fig. 4. By manipulating the mean value and standard deviation, random noise triggers can vary in visibility, ranging from imperceptible to noticeable.

IV. Experiment Study

A. Experimental Configuration

We perform our experiments utilizing the DeepMIMO dataset [21] and employ ADP as the input of positioning systems [5]. The DeepMIMO outdoor scenario number 1 (O1) at 3.5 GHz band and indoor scenario number 3 (I3) at 60 GHz are deployed as outdoor and indoor environments, respectively. For the outdoor environment, a single base station is equipped with a uniform linear array (ULA) with 64 antennas. The data generation process involves generating 199,100 data points by varying the locations from row 1 to row 1,100. The position coordinates range from $(242.4, 297.2)$ to $(278.4, 517.0)$. The indoor environment simulates a 10m × 11m conference room. The position coordinates range from $(26.34, 6.18)$ to $(27.54, 11.67)$. Treating localization as a regression problem, we set the target positions for backdoor attacks to be coordinates $(200, 200)$ and $(0, 0)$ for the outdoor and indoor cases, respectively.

For all experiments, we set the learning rate to 0.0001 and epochs to 10. The poisoning rate p and trigger value t are the variables to control the backdoor attack, allowing us to adjust their values to assess the effectiveness of backdoor attacks. All experiments are performed on a server with an Intel Xeon E5-2650L v4 CPU and 8 NVIDIA GeForce GTX 1080Ti GPU.

B. One-pixel Attack

We first evaluate the one-pixel attack on localization systems, where the root mean square error (RMSE) is chosen as the evaluation metric. A lower RMSE value indicates a more accurate prediction of position. In our evaluation, we consider three different poisoning rates p. Such variation in poisoning rates allows us to assess the influence of different levels of poisoning in the training data and positioning performance.

For the indoor backdoor attack case, the ADP values vary between 0.0162 and 0.0001. Therefore, we investigate a range of trigger values for the one-pixel attack, specifically, the set of trigger values $[0.01, 0.005, 0.001, 0.0005, 0.0001]$. Such selection roughly encompasses both the maximum and minimum values to examine the attack's effectiveness.

Fig. 5 shows the results of one-pixel attack on indoor localization. Without backdoor attacks, the CNN model demonstrates a satisfactory performance, yielding an RMSE value of approximately 0.95. However, when the trigger value becomes to be larger than 0.0005, the positioning system can be completely deceived, generating inaccurate locations with an

RMSE of about 20. When considering different poisoning rates, the effectiveness of the one-pixel attack diminishes significantly when the trigger value is set to 0.0001. It is noted that 0.0001 is the minimum value in the ADP matrix. In these cases, the system is robust to backdoor attacks and achieves a similar accuracy as that of an unattacked system. When the poisoning rate is increased to 0.1, the one-pixel attack recovers some abilities, resulting in a high RMSE of 18.6.

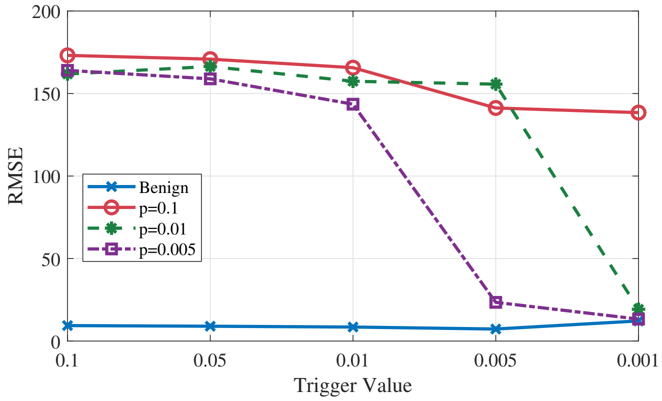

Fig. 6: RMSE of one-pixel attack on outdoor localization.

In the outdoor localization scenario, the ADP values span a range from 0.0003 to 0.457. Due to the larger values compared to the indoor case, we modify the trigger values set to $[0.1, 0.05, 0.01, 0.005, 0.001]$ while keeping the poisoning rate unchanged. Fig. 6 depicts the results of the one-pixel attack on outdoor localization and shows a similar trend to indoor localization. When the trigger value is set to 0.005, the small poisoning set with a 0.005 poisoning rate is unable to effectively deceive the localization system. Likewise, when the trigger value is set to 0.001, only the largest poisoning set with a poisoning rate of 0.1 can deceive the system, albeit causing a partial loss of effectiveness.

By comparing the above results, we have the following observations:

- In general, introducing a trigger does not impact the effectiveness of CNN when predicting locations on the benign dataset. However, there is a minor decrease in accuracy as the trigger gets smaller. This may suggest that a smaller trigger value may potentially confuse the model to recognize the trigger.
- If the trigger value is sufficiently large, the one-pixel attack can effectively impair the performance of the model without using a significant amount of poisoned samples. However, when the trigger value is too small to effectively attack the model, increasing the poisoning rate can help to enhance the attack capability.

To achieve the goals mentioned in Section III-B, the recommended settings for the one-pixel attack on the indoor localization system is to set the trigger value to 0.005 and the poisoning rate to 0.01. For the outdoor localization system, it is suggested to set the trigger value to 0.0005, accompanied by a poisoning rate of either 0.1 or 0.01.

TABLE II: Results of Random Noise Attack

	Indoor		Outdoor	
	Poisoned	Benign	Poisoned	Benign
$\mu = 10^{-4}$, $\sigma = 10^{-4}$	19.98	0.41	164.22	17.55
$\mu = 10^{-4}$, $\sigma = 10^{-5}$	19.13	0.83	169.20	20.18
$\mu = 10^{-5}$, $\sigma = 10^{-4}$	19.74	6.98	162.57	6.16
$\mu = 10^{-5}$, $\sigma = 10^{-5}$	19.67	19.67	15.21	23.07

C. Random Noise Attack

While the one-pixel attack is proved effective in deceiving localization systems, it may be easily detected and removed since the trigger is always located at a specific place with a fixed value. We next investigate the application of an invisible random noise attack on localization systems. For this purpose, we generate a matrix of normally distributed noise with the same shape as the inputs. To ensure non-negative inputs, we take the absolute value of the noise matrix elements. Once the noise matrix is generated, we establish it as the trigger. In the following, we assess the impact of varying mean values μ and standard deviations σ of the normal distribution underlying the trigger. In all the cases, we set the poisoning rate to 0.01.

The results presented in Table II demonstrate that the random noise attack can effectively deceive the model and achieve similar RMSE values compared to the one-pixel attack. In the case of indoor localization, when the mean value is set to 10^{-4}, the random noise attack successfully degrades the system's performance without affecting the prediction of benign sample locations. However, when the mean value is reduced to 10^{-5}, the noise becomes more imperceptible considering that the smallest value of the original input is 0.0001. In this scenario, the effectiveness of the random noise attack diminishes. Although the RMSE values of the poisoned samples remain the same, the system fails to accurately predict the locations of benign samples, which is contradictory to the goal of backdoor attacks.

For outdoor localization, the attack remains effective on the poisoned samples when the mean value is set to 10^{-4}. However, the system experiences a decrease in accuracy when predicting the locations of benign samples, resulting in an increase of approximately 10 in the RMSE value. Furthermore, when the mean value is set to 10^{-5}, the random noise attacks exhibit different behaviors. When a larger standard deviation 10^{-4} is used, the attack successfully causes the system to make incorrect predictions on the poisoned samples while maintaining accurate predictions on the benign samples without losing precision. Nevertheless, when the standard deviation is set to 10^{-5}, the random noise attack becomes completely ineffective in attacking the system. The presence of the trigger even leads to smaller RMSE values.

In both cases, the RMSE values on the original datasets increase when the standard deviation decreases. This suggests that reducing the fluctuations in triggers makes it more challenging for the DNN model to differentiate between benign and poisoned inputs. Due to the distinct locations between the benign inputs and the poisoned inputs, the DNN model becomes perplexed in predicting accurate locations as it struggles to recognize the triggers. Consequently, regardless of the presence or absence of the trigger, the model fails to make accurate predictions.

To accomplish the objectives outlines in Section III-B, the recommended settings for the random noise attack on the indoor localization system is setting the trigger with a mean value of 10^{-4} and a standard deviation of 10^{-5}. As shown in Fig. 4(a), the trigger is visible when the mean value and standard deviation are all set to 10^{-4}. To effective attacking the outdoor localization system, it is suggested to set the trigger with a mean value of 10^{-5} and a standard deviation of 10^{-4}.

Overall, the random noise attack can be invisible while successfully degrading the system's localization performance. Compared to the one-pixel attack, the random noise attack requires careful adjustments of the mean and standard deviation values to trade-off between invisibility and successful attacks.

D. Discussion

By implementing the one-pixel attack and the random-noise attack, we successfully mislead the system to generate incorrect locations for poisoned data, while accurately predicting positions for benign data. All of these processes do not require any knowledge of the underlying model architecture. The one-pixel attack is straightforward to launch, which only requires to modify a single input value. In contrast, the random noise attack requires to balance invisibility and effectiveness. For future work, we may explore defense mechanisms to defend against these two attack methods and study a more general and robust attack method on DNN-based localization systems. We shall also investigate more realistic attack methods that do not involve directly injecting triggers after the input data has been processed.

V. Conclusions

In this paper, we focused on investigating backdoor attacks for mmWave/massive MIMO-based localization systems. We evaluated two different triggers: a one-pixel trigger and a random noise trigger. Through comprehensive experiments, we demonstrated the effectiveness of both triggers in launching backdoor attacks on localization systems in both indoor and outdoor environments. By choosing proper trigger paramter values, positioning systems exhibited excellent performance on original, benign inputs, but failed to accurately predict the locations of triggered inputs.

Acknowledgments

This work is also supported in part by the NSF (CNS-2321763, CNS-2319343, CNS-2317190, and CNS-2107190), as well as the Wireless Engineering Research and Education Center at Auburn University, Auburn, AL.

References

[1] S. A. Busari, K. M. S. Huq, S. Mumtaz, L. Dai, and J. Rodriguez, "Millimeter-wave massive MIMO communication for future wireless systems: A survey," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 2, pp. 836–869, Secondquarter 2017.

[2] X. Wang, X. Wang, and S. Mao, "RF sensing in the internet of things: A general deep learning framework," *IEEE Communications Magazine*, vol. 56, no. 9, pp. 62–67, Sept. 2018.

[3] C. Zhang, P. Patras, and H. Haddadi, "Deep learning in mobile and wireless networking: A survey," *IEEE Communications surveys & tutorials*, vol. 21, no. 3, pp. 2224–2287, 2019.

[4] X. Wang, L. Gao, S. Mao, and S. Pandey, "CSI-based fingerprinting for indoor localization: A deep learning approach," *IEEE transactions on vehicular technology*, vol. 66, no. 1, pp. 763–776, Jan. 2016.

[5] F. Hejazi, K. Vuckovic, and N. Rahnavard, "DyLoc: Dynamic localization for massive mimo using predictive recurrent neural networks," in *Proc. IEEE INFOCOM 2021*, Virtual Conference, May 2021, pp. 1–9.

[6] N. Akhtar and A. Mian, "Threat of adversarial attacks on deep learning in computer vision: A survey," *IEEE Access Journal*, vol. 6, pp. 14 410–14 430, Mar. 2018.

[7] S.-M. Moosavi-Dezfooli, A. Fawzi, O. Fawzi, and P. Frossard, "Universal adversarial perturbations," in *Proc. IEEE/CVF CVPR*, Honolulu, HI, July 2017, pp. 1765–1773.

[8] Y. Gao, B. G. Doan, Z. Zhang, S. Ma, J. Zhang, A. Fu, S. Nepal, and H. Kim, "Backdoor attacks and countermeasures on deep learning: A comprehensive review," *arXiv preprint arXiv:2007.10760*, July 2020. [Online]. Available: https://arxiv.org/abs/2007.10760

[9] A. Bahramali, M. Nasr, A. Houmansadr, D. Goeckel, and D. Towsley, "Robust adversarial attacks against DNN-based wireless communication systems," in *Proc. 2021 ACM SIGSAC Conference on Computer and Communications Security*, Virtual Conference, Nov. 2021, pp. 126–140.

[10] M. Patil, X. Wang, X. Wang, and S. Mao, "Adversarial attacks on deep learning-based floor classification and indoor localization," in *Proc. 2021 ACM Workshop on Wireless Security and Machine Learning (WiseML'21)*, Abu Dhabi, UAE, June-July 2021, pp. 7–12.

[11] U. Boora, X. Wang, and S. Mao, "Robust massive MIMO localization using neural ODE in adversarial environments," in *Proc. IEEE ICC 2022*, Seoul, South Korea, May 2022, pp. 4866–4871.

[12] W. Xu, W. Trappe, Y. Zhang, and T. Wood, "The feasibility of launching and detecting jamming attacks in wireless networks," in *Proc. ACM MobiHoc 2005*, Urbana-Champaign, IL, May 2005, pp. 46–57.

[13] P. Huang, X. Zhang, S. Yu, and L. Guo, "IS-WARS: Intelligent and stealthy adversarial attack to Wi-Fi-based human activity recognition systems," *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 6, pp. 3899–3912, Nov.-Dec. 2021.

[14] L. Xu, X. Zheng, X. Li, Y. Zhang, L. Liu, and H. Ma, "WiCAM: Imperceptible adversarial attack on deep learning based WiFi sensing," in *Proc. IEEE SECON 2022*, Virtual Conference, Sept. 2022, pp. 10–18.

[15] H. A. X. Wang and S. Mao, "Adversarial human activity recognition using Wi-Fi CSI," in *Proc. 2021 Annual IEEE Canadian Conference of Electrical and Computer Engineering (CCECE'21)*, Virtual Conference, Sept 2021, pp. 1–5.

[16] R. K. Sah and H. Ghasemzadeh, "Adar: Adversarial activity recognition in wearables," in *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. IEEE, 2019, pp. 1–8.

[17] Z. Yang, Y. Zhao, and W. Yan, "Adversarial vulnerability in doppler-based human activity recognition," in *Proc. 2020 International Joint Conference on Neural Networks*, Glasgow, UK, July 2020, pp. 1–7.

[18] J. Yang, H. Zou, and L. Xie, "SecureSense: Defending adversarial attack for secure device-free human activity recognition," *IEEE Transactions on Mobile Computing*, Dec. 2022, Early Access.

[19] X. Wang, X. Wang, S. Mao, J. Zhang, S. Periaswamy, and J. Patton, "Adversarial deep learning for indoor localization with channel state information tensors," *IEEE Internet of Things Journal*, vol. 9, no. 19, pp. 18 182–18 194, Oct. 2022.

[20] Y. Li, Y. Jiang, Z. Li, and S.-T. Xia, "Backdoor learning: A survey," *IEEE Transactions on Neural Networks and Learning Systems*, June 2022, Early Access.

[21] A. Alkhateeb, "DeepMIMO: A generic deep learning dataset for millimeter wave and massive MIMO applications," *arXiv preprint arXiv:1902.06435*, Feb. 2019. [Online]. Available: https://arxiv.org/abs/1902.06435