

Towards Adversarial-Resilient Deep Neural Networks for False Data Injection Attack Detection in Power Grids

Jiangnan Li*, Yingyuan Yang[§], Jinyuan Stella Sun*, Kevin Tomsovic*, Hairong Qi*

*The University of Tennessee, Knoxville, jiangnanutk@gmail.com, {jysun, tomsovic, hqi}@utk.edu

[§]University of Illinois Springfield, yyang260@uis.edu

Abstract—False data injection attacks (FDIAs) pose a significant security threat to power system state estimation. To detect such attacks, recent studies have proposed machine learning (ML) techniques, particularly deep neural networks (DNNs). However, most of these methods fail to account for the risk posed by adversarial measurements, which can compromise the reliability of DNNs in various ML applications. In this paper, we present a DNN-based FDIA detection approach that is resilient to adversarial attacks. We first analyze several adversarial defense mechanisms used in computer vision and show their inherent limitations in FDIA detection. We then propose an adversarial-resilient DNN detection framework for FDIA that incorporates random input padding in both the training and inference phases. Our simulations, based on an IEEE standard power system, demonstrate that this framework significantly reduces the effectiveness of adversarial attacks while having a negligible impact on the DNNs' detection performance.

Index Terms—False Data Injection Attack, Smart Grid Communication, Deep Learning, Adversarial Attacks

I. INTRODUCTION

State estimation is a critical application in power grids. The state estimator in the control center uses analog measurements and status data from remote sensors to determine the network topology and estimate voltage, current magnitudes, and phase. The estimator also employs a built-in residue-based bad data detection mechanism to remove bad data and filter measurement errors. The refined data is then used by operators and advanced applications to compute market prices and make operational decisions.

False data injection attacks (FDIAs), first proposed by Liu *et al.* in 2009 [1], are a well-known attack vector in state estimation. They allow attackers to inject malicious false data into legitimate measurements while bypassing bad data detection, causing the estimator to output incorrect system states. These attacks can disturb critical operations, such as contingency analysis, that rely on the estimated state. Several types of FDIAs have been developed, including FDIA with incomplete information [2], blind FDIA [3], and outage masking [4].

To mitigate the threat posed by FDIAs, various detection methods have been proposed in the literature, such as measurement protection [5] and detection using phasor measurement units (PMUs) [6]. In recent years, machine learning (ML), especially deep neural networks (DNNs), has become a key technique for detecting FDIAs [7]–[15]. The basic architecture of DNN-based FDIA detection is shown in Fig. 1. DNNs

leverage the statistical properties of measurement data and can achieve state-of-the-art detection performance. Furthermore, DNNs are usually pure software systems and can be deployed in the control center without requiring additional equipment or infrastructure upgrades, making them ideal defense techniques against FDIAs.




Fig. 1: Illustration of DNN-based FDIA detection. The attacker compromises a subset of the sensors and inject false data. A trained DNN is placed in the control center to detect and remove malicious measurements.

Recent research in the field of artificial intelligence (AI) has demonstrated that even well-performing deep neural networks (DNNs) are highly vulnerable to adversarial attacks. By adding well-crafted perturbations to a DNN's input, an attacker can deceive the network into outputting incorrect decisions. Such attacks have also been shown to be effective in power system applications [16]–[19], including fault detection, isolation, and analysis (FDIA) based on DNNs [20] and state estimation [21]. Therefore, it is essential to consider the risks of adversarial attacks and investigate more robust DNNs that are resilient to such attacks for FDIA detection. However, typical approaches for mitigating adversarial attacks are primarily designed for computer vision applications and make specific assumptions that are not necessarily applicable to FDIA detection. For example, adversarial detection approaches often assume that adversarial examples follow a different distribution from normal inputs [22]. Nevertheless, such assumptions may be infeasible in FDIA detection due to physical system constraints, as demonstrated in this paper. Thus, effectively mitigating adversarial attacks in DNN-based FDIA detection represents a significant research challenge, which is the primary contribution of this paper.

In this work, we investigate defense mechanisms to mit-

igate adversarial attacks in DNN-based FDIA detection. As a baseline, we evaluate three typical defense methods in the computer vision domain and highlight their intrinsic limitations in preventing adversarial attacks in FDIA detection. Subsequently, we propose a general random input padding framework for DNN-based FDIA detection. Our framework can be easily adapted to different ML models used in existing research [7]–[15] with simple changes to the configuration during training and detection deployment. We evaluate our approach using synthetic datasets generated from the IEEE standard 118-bus system. Our results demonstrate that the proposed framework can significantly mitigate adversarial attacks. Our main contributions are as follows:

- We highlight the need for robust DNN-based methods that are resilient to adversarial attacks and summarize the adversarial defense properties and requirements in FDIA detection.
- We demonstrate that typical adversarial defense approaches (defensive distillation, adversarial training, and adversarial detection) in the computer vision domain are insufficient in mitigating adversarial attacks in FDIA detection, as shown through theoretical analysis and simulations.
- We propose a general defense framework that randomly pads the model input during both the training and testing stages. The framework reduces the effectiveness of adversarial perturbations and can be seamlessly adapted to DNNs proposed in the literature.
- We conduct simulations based on the standard IEEE 118-bus power system. The results show that our framework can effectively mitigate adversarial attacks while inducing minimal degradation in the models' detection accuracy of normal inputs.

The rest of the paper is organized as follows. The related work is presented in Section II. Section III presents the adversarial attacks in DNN-based FDIA detection. We discuss the limitations of the existing defense approaches and propose our framework in Section IV. Simulation results are presented in Section V. Section VI discusses the limitations and future work. Finally, Section VII concludes the paper.

II. RELATED WORK

A. FDIA Detection

Since FDIA is considered as a serious threat to power system security, many studies and countermeasures are proposed. [23] analyzed the impact of FDIA with static security assessment and claimed that FDIA can deceive the operator to make wrong actions, such as load shedding. Khalaf *et al.* studied and evaluated FDIA under wide area protection settings and showed that the FDIA can affect the grids' operation and stability [24]. Bi and Zhang utilized graph theory to locate the meters to be protected [25]. Some studies took advantage of new hardware devices, such as PMUs. The PMUs' measurements are synchronized to GPS signals and will increase the barrier for FDIA attacks. [26] and [27] investigated the placement strategies of PMUs.

Detection that utilized ML techniques became popular in the literature. In 2016, Ozay *et al.* first utilized ML techniques to detect FDIA [7]. They evaluated the detection performance of traditional ML algorithms, such as K-Nearest Neighbour (KNN), support vector machine (SVM), and Sparse Logistic Regression (SLR), and demonstrated that ML achieved convincing detection accuracy with fine-tuned parameters. After that, Yan *et al.* studied the performance of a supervised learning classifier on detecting both direct and stealth FDIA [8]. He *et al.* employed deep learning (DL) for FDIA and electricity theft detection [9]. In 2018, [13] evaluated the performance of plain DNN in FDIA attacks. Thereafter, different DNNs were designed. The studies in [10][11][12] used recurrent neural networks (RNNs) for FDIA detection. They evaluated the DNNs with different scale power systems and all achieves high detection results. Niu *et al.* designed a DNN that adopted both convolutional neural network (CNN) and RNN to detect FDIA in dynamic time-series measurement data [14]. False data usually follows a different distribution from legitimate data. Wang *et al.* trained an auto-encoder with pure normal measurement data and employed the reconstruction loss as the metric for FDIA detection [15].

B. Adversarial Attacks

Recent research in the AI domain has demonstrated that DNNs are vulnerable to adversarial attacks. In 2013, adversarial examples to DNNs were discovered by Szegedy *et al.* [28]. By adding a small crafted perturbation to the legitimate inputs, the DNNs will be led to output wrong results. After that, different adversarial attack algorithms were proposed. Goodfellow *et al.* presented the Fast Gradient Sign Method (FGSM) that utilized the signed gradient values to generate perturbations [29]. After that, the Fast Gradient Method by Rozsa *et al.* used the gradient values directly [30]. Other well-known adversarial attack algorithms include the DeepFool [31] and iterative attack [32].

The potential threat of adversarial attacks to critical infrastructures also draws attention in recent research. In 2018, Chen *et al.* investigated the effect of adversarial examples with both categorical and sequential applications in power systems [16]. They then designed adversarial attacks for regression models used for load forecasting [17]. Tian *et al.* extended [16] and proposed an adaptive normalized attack for power system ML applications [18]. [20] proposed constrained adversarial machine learning in CPS applications, and demonstrated an attacker could generate adversarial examples that met the intrinsic constraints defined by physical systems, such as the residual-based detection in state estimation.

The defense methods against adversarial attacks draw attention in both security and AI communities. The typical defense approaches include model distillation [33], adversarial detection [34][22][35], adversarial training [32][36][37], input reconstruction [38][39], stochastic methods [39][40], and so on.

III. ADVERSARIAL ATTACK IN FDIA DETECTION

A. Background: False Data Injection Attacks (FDIA)

Power systems utilize state estimation to estimate the state of each bus by analyzing the other sensor's measurements. FDIA enables an attacker to generate a false measurement vector \mathbf{a} to be added to legitimate \mathbf{z} , so that the polluted measurements will be $\mathbf{z}_a = \mathbf{z} + \mathbf{a}$. [1] shows that if the attacker knows the power system matrix \mathbf{H} , she/he can construct $\mathbf{a} = \mathbf{H}\mathbf{c}$ (\mathbf{c} represents the estimation error) that can bypass the fault detection in state estimation, as shown by equation (1), where $\hat{\mathbf{x}}_{bad}$ and $\hat{\mathbf{x}}$ denote the estimated \mathbf{x} using \mathbf{z}_a and \mathbf{z} respectively. The details of FDIA can be found in [1].

$$\|\mathbf{z}_a - \mathbf{H}\hat{\mathbf{x}}_{bad}\| = \|\mathbf{z} + \mathbf{a} - \mathbf{H}(\hat{\mathbf{x}} + \mathbf{c})\| \quad (1a)$$

$$= \|\mathbf{z} - \mathbf{H}\hat{\mathbf{x}} + (\mathbf{a} - \mathbf{H}\mathbf{c})\| \quad (1b)$$

$$= \|\mathbf{z} - \mathbf{H}\hat{\mathbf{x}}\| \leq \tau \quad (1c)$$

The above equation can be further represented as:

$$\mathbf{a} = \mathbf{H}\mathbf{c} \Leftrightarrow \mathbf{P}\mathbf{a} = \mathbf{P}\mathbf{H}\mathbf{c} \Leftrightarrow \mathbf{P}\mathbf{a} = \mathbf{H}\mathbf{c} \Leftrightarrow \mathbf{P}\mathbf{a} = \mathbf{a} \quad (2a)$$

$$\Leftrightarrow \mathbf{P}\mathbf{a} - \mathbf{a} = \mathbf{0} \Leftrightarrow (\mathbf{P} - \mathbf{I})\mathbf{a} = \mathbf{0} \quad (2b)$$

$$\Leftrightarrow \mathbf{B}\mathbf{a} = \mathbf{0} \quad (2c)$$

where $\mathbf{P} = \mathbf{H}(\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T$ and matrix $\mathbf{B} = \mathbf{P} - \mathbf{I}$.

B. Attack Properties

In this paper, we mainly consider supervised learning techniques for FDIA detection as they are more popular in the literature [7]–[14]. We note that our method can be easily extended to unsupervised learning approaches, such as the autoencoder used in [15], by designing corresponding loss functions.

Without loss of generality, we consider the DNN-based FDIA detection as a binary classification problem. The trained DNN $F : \mathbf{Z} \rightarrow Y$ maps the input measurement vectors \mathbf{Z} to their labels Y . An ideal DNN will map a legitimate measurement \mathbf{z} to legitimate label Y_0 and a false measurements $\mathbf{z}_a = \mathbf{z} + \mathbf{a}$ to false label Y_1 .

TABLE I: Terminologies and Notations

Terminology	Notation	Explanation
legitimate measurements	\mathbf{z}	sensors' original measurements
injected false data	\mathbf{a}	meets constraint $\mathbf{B}\mathbf{a} = \mathbf{0}$
false measurements	\mathbf{z}_a	$\mathbf{z}_a = \mathbf{z} + \mathbf{a}$
adversarial perturbation	\mathbf{v}	perturbation for \mathbf{z}_a
total injected false data	$\hat{\mathbf{a}}$	$\hat{\mathbf{a}} = \mathbf{a} + \mathbf{v}$, $\mathbf{B}\hat{\mathbf{a}} = \mathbf{0}$
adversarial measurements	\mathbf{z}_{adv}	$\mathbf{z}_{adv} = \mathbf{z}_a + \mathbf{v} = \mathbf{z} + \hat{\mathbf{a}}$

The adversarial attack in FDIA detection is a false-negative attack that deceive F to classify the false measurements as legitimate. The attacker needs to generate an adversarial perturbation vector \mathbf{v} and add it to the false measurement vector \mathbf{z}_a , so that the adversarial measurement $\mathbf{z}_{adv} = \mathbf{z}_a + \mathbf{v} =$

$\mathbf{z} + \mathbf{a} + \mathbf{v} = \mathbf{z} + \hat{\mathbf{a}}$ can be classified as legitimate measurements (with Y_0 label) by F , where $\hat{\mathbf{a}} = \mathbf{a} + \mathbf{v}$ is the total injected false data. Meanwhile, the attacker requires $\hat{\mathbf{a}}$ also meet the constraint $\mathbf{B}\hat{\mathbf{a}} = \mathbf{0}$ defined by (2c) to avoid being removed by the state estimation residual-based detection system. Therefore, the adversarial measurements in FDIA detection can also be considered as special false measurements. To be clear, Table I summarizes the terminologies and notations.

C. Defense Requirements & Threat Model

White-box adversarial attacks allow the attacker to have access to the target DNN, which are common attacks in previous literature and have been extensively studied since they help researchers to learn the weakness of DNNs more directly [41]. Robust against white-box adversarial attacks is a desired property that the DNNs should maintain [37], especially for critical infrastructure like power grids. In this paper, we expect the our defense mechanism to be resilient to white-box FDIA adversarial attacks.

From the attacker's point of view, we summarize the threat model of the adversarial attacks in FDIA detection:

- The attacker can compromise k ($k > m - n$) measurements in the power system, which is inherited from the FDIA attacks' requirement[1].
- Also as described in [1], the attacker knows the \mathbf{H} matrix to launch FDIA.
- The total injected false data $\hat{\mathbf{a}}$ should always follow $\mathbf{B}\hat{\mathbf{a}} = \mathbf{0}$ to bypass the built-in residual-based detection mechanism in state estimation.

D. Adversarial Attack Algorithm

We employ the iterative projection framework in [32] and [42] since the perturbation \mathbf{v} needs to follow the constraints defined by the power system. However, instead of projecting to the ε -neighbor ball, the attacker needs to map the adversarial perturbation \mathbf{v} to the solution space of the homogeneous equation $\mathbf{B}\mathbf{v} = \mathbf{0}$ so that the total injected false data $\mathbf{B}\hat{\mathbf{a}} = \mathbf{B}(\mathbf{a} + \mathbf{v}) = \mathbf{B}\mathbf{a} + \mathbf{B}\mathbf{v} = \mathbf{0}$ will enable the adversarial measurements \mathbf{z}_{adv} to bypass the residual-based detection.

IV. ADVERSARIAL DEFENSE IN FDIA DETECTION

The arms race between adversarial attacks and defense in the AI domain is still in progress. In recent years, different adversarial defense mechanisms are proposed in the literature, such as model distillation [33], adversarial training [32], [36], [37], adversarial detection [22], [34], [35], and input reconstruction [38][39]. However, to the best of our knowledge, no defense method was demonstrated to be effective against all adversarial attacks [41][43]. In particular, adversarial examples will always exist since none of the trained DNNs are perfect. Therefore, adversarial defense in FDIA detection mainly aims to degrade the attack performance and increase the attack cost.

In this section, we review three typical adversarial defense mechanisms proposed in the literature. We analyze these approaches and show that they have inherent limitations to mitigate adversarial attacks in FDIA detection effectively. We

verify our analysis with simulations in Section V. After that, we propose our random input padding framework to train adversarial-resilient DNNs for FDIA detection. The evaluation in Section V shows that our framework can significantly decrease the overall performance of adversarial attacks.

A. Typical Adversarial Defense Methods

1) **Defensive Distillation:** Neural network distillation was originally proposed to reduce the size of DNN architectures [44]. In 2016, Papernot *et al.* employed distillation as an adversarial defense approach in the computer vision field [33]. The basic idea behind DNN distillation is to transfer the knowledge from a trained model to a new model, and the new DNN is shown to be less sensitive to input perturbations and then becomes more robust to adversarial attacks.

Limitation Analysis: A common premise of adversarial attacks in the computer vision domain is that the size of the perturbation needs to be relatively small so that the adversarial pictures will not be noticed by human eyes. The assumption is no longer feasible in DNN-based FDIA detection. For example, a reckless attacker who aims to inject considerable false data into the state estimation only needs to consider the size of the total injected false data $\hat{\mathbf{a}}$. Although model distillation decreases the DNN's sensitivity to input perturbations, the attacker can still iteratively search for a valid adversarial measurement \mathbf{z}_{adv} that deceives the DNN model without considering the size of \mathbf{v} . Therefore, model distillation is not capable to defend against adversarial attacks in FDIA detection.

2) **Adversarial Training:** Adversarial training is one of the common methods to mitigate an adversarial attack [32][36][37]. The basic principle of adversarial training is to generate and include adversarial examples in each data batch during the training stages. As the DNN is trained to recognize adversarial examples, it becomes more robust.

Limitation Analysis: Adversarial training needs to generate adversarial examples for each batch of data during the training process, which increases the training computation overhead significantly. As demonstrated in Section III-D, to avoid being removed by the residual-based detection scheme, the adversarial perturbations must be projected to fit the constraint. The mapping process will further significantly introduce computation overhead to the adversarial training process. Therefore, adversarial training is not scalable to large systems that contain massive data resources. Meanwhile, [32] shows that adversarial training performs deficiently for iterative attacks, which makes it inappropriate for FDIA detection.

3) **Adversarial Detection:** Adversarial detection recognizes adversarial examples at the DNN inference stage [34][22][35]. In particular, an auxiliary binary classification DNN F_{adv} is trained with normal records and corresponding adversarial examples [22] to detect if an input is an adversarial example. The adversarial detection DNN F_{adv} will be employed first to recognize the input records, and only the normal records will be fed into the original functional DNN.

Limitation Analysis: Adversarial detection assumes that the adversarial examples follow a different distribution from normal inputs. The assumption is reasonable in the computer vision domain (the natural images will not contain the well-crafted perturbations) but not applicable for FDIA detection. As introduced in Section III, the manifold of the injected false data \mathbf{a} can be represented by the constraint $\mathbf{B}\mathbf{a} = \mathbf{0}$ empirically. To bypass the built-in residual-based detection of state estimation, the total injected false data $\hat{\mathbf{a}}$ is also required to meet the constraint $\mathbf{B}\hat{\mathbf{a}} = \mathbf{0}$. Since the number of the possible attack scenarios (different meters can be compromised) can be large, intuitively, the crafted adversarial measurement vector \mathbf{z}_{adv} shares a similar manifold with the false measurement vector \mathbf{z}_a in the FDIA detection. In fact, \mathbf{z}_{adv} can be considered as special \mathbf{z}_a , as analyzed in Section 4.1. Therefore, adversarial detection will not work effectively in FDIA detection tasks. This analysis can also be adapted to input reconstruction methods [38][39].

B. Random Input Padding Framework

As discussed above, the adversarial defense in DNN-based FDIA detection is non-trivial since the adversarial measurements share the same manifold as the general false measurements. Given the victim model, the attacker generates the perturbation \mathbf{v} for \mathbf{z}_a iteratively through a gradient-based optimization process. As presented in [32], the perturbation generated by multi-step attacks usually has worse transferability, which indicates that adversarial perturbation \mathbf{v} in FDIA detection is highly likely to be unique for each given \mathbf{z}_a . Therefore, there is an intuition that the perturbation will no longer work if the input to the model changes. Inspired by the stochastic-based defense mechanisms in the computer vision field [39][40], we propose a random input padding defense framework to mitigate the effect of adversarial attacks in FDIA detection.

The philosophy of our random input padding framework is straightforward, and the overall structure is shown in Fig 2. A random padding layer is added in front of the DNN in both training and inference stages. In general, the measurements of the sensors \mathbf{z} are used as the features to train the detection models. Our framework firstly requires the operator to pick a padding dimension number P ($P > m$) as the input feature numbers for the DNN. Thereafter, we pad $P - m$ zeros randomly to the plain inputs \mathbf{z} and there will be $P - m$ padding scenarios in total. The DNN is then required to learn the pattern from the plain measurements that are embedded into the padded inputs during the training process. During the inference stage, when a new measurement vector \mathbf{z} is received, the framework randomly pads \mathbf{z} to a P dimensional vector and feeds the padded vector to the DNN. Ideally, the detection rate against adversarial attacks should be $1 - \frac{1}{P-m+1}$ ($P \geq m$). The padding framework also works with possible input reshape, as shown in Fig 2.

As the padding process is random for each \mathbf{z} at the inference stage, the attacker (and even the operator) cannot know the final DNN padded input vectors even when she/he knows




Fig. 2: Illustration of random inputs padding framework. Each record in the plain inputs is a 4-d vector. The above 1-d padding scheme randomly pads four zeros to plain inputs and the padded vectors become an 8-d ($P = 8$). The below 2-d padding scheme reshapes the inputs pads five zeros, which finally generates 3×3 dimensional records ($P = 9$).

the whole framework. The attacker will be able to generate perturbations for one of the $P-m$ padding scenarios. Since the multi-steps perturbations have relatively weak transferability, the adversarial attacks should have a lower success rate under the random padding framework. Intuitively, a larger P will decrease the success rate of adversarial attacks and finally increase the robustness of the DNN used for FDIA detection.

Different from [40], our framework requires input data pre-processing (padding) during the training stage and cannot be applied to a trained model directly. This is because the measurement data of a specific power system should follow the manifold defined by the physical property of the system, which will be destroyed if the measurement vectors are reshaped, resized, or sampled directly. On the other hand, the FDIA detection performance of the legitimate data vector \mathbf{z} and false measurements \mathbf{z}_a will not be constrained by the padding/scale size if an appropriate structure of the neural network is selected. Meanwhile, our framework only increases the computation of the training process slightly and is compatible with different neural networks.

V. SIMULATION

In this section, we first evaluate the effect of adversarial attacks in DNN-based FDIA detection. After that, we evaluate the typical adversarial defense approaches discussed in Section IV-A and verify our analysis of their limitation through numeric simulations. Finally, we evaluate the performance of our random padding framework.

Synthetic Dataset Generation: The simulations are conducted based on the standard IEEE 118-bus system, which was used as a test benchmark in previous literature [1], [7], [9], [11]. We employ the MATPOWER [45] tool to derive the \mathbf{H} matrix of the system and simulate the power flow measurement data of each branch as \mathbf{Z} . In our evaluation,

each measurement vector \mathbf{z} contains $m = 186$ measurements. The FDIA is also implemented with MATLAB based on the corresponding \mathbf{H} . We simulate a dataset D_{train} that contains 30,000 legitimate measurement vectors. We pollute half of the records in D_{train} by injecting false data generated by the FDIA. To simulate practical possible attacks, we set different compromised numbers $70 < k < 100$, and the indexes of the compromised measurements are randomly selected. In addition to D_{train} , we generate four test datasets. We consider the scenarios that there were 75, 80, 85, 90 measurements being compromised by the attacker respectively, and simulate 1000 polluted (false) data for each scenario.

Target DNN: We train a feed-forward neural network F as the target DNN in our simulations, and the structure of F is shown in Table II. We randomly split D_{train} into the training part and testing part, with the test part containing 15% of records in D_{train} . We employ the categorical cross-entropy as the loss function and utilize stochastic gradient descent to optimize the loss. Through tuning parameters, F finally achieves a 98.9% detection accuracy and 98.6% recall. F is developed with the Tensorflow and Keras libraries. The simulations are conducted on a Windows 10 machine with an Intel i7 CPU and an additional NVIDIA GeForce GTX 1070 GPU to accelerate the training process.

TABLE II: Model Structure of F

Layer	1	2	3	4	5	6
Nodes	186	128	64	16	0.25 Dropout	2 Softmax

* Dense layer is used for each layer.

* The activation function is *ReLU* unless specifically noted.

Evaluation Metrics: We set three metrics to evaluate the attack performance. The first metric is the detection **Recall** of the target DNN under adversarial attacks, which represents the probability of the adversarial measurements fooling the target DNN. In addition, we consider two different attack scenarios. The attack scenario happens when the attacker aims to inject a specific false vector \mathbf{a} into state estimation, such as to gain specific profit through modifying local marginal price [46][47]. In this scenario, the perturbation \mathbf{v} is required to be small and our second metric is the L_2 -Norm of \mathbf{v} of effective adversarial measurements that successfully fool the DNN's detection. We denote this metric as **Bias L_2 -Norm**. The second scenario is that a malicious attacker aims to inject considerable false data to state estimation and the size of $\hat{\mathbf{a}}$ is expected to be large. We set the L_2 -Norm of $\hat{\mathbf{a}}$ to be the third evaluation metric, which is noted as **Valid L_2 -Norm**. In summary, from the attacker's point of view, a lower Recall, a smaller Bias L_2 -Norm, and a larger Valid L_2 -Norm indicate a more successful adversarial attack and vice versa.

A. Adversarial Attacks in FDIA Detection

This subsection studies the vulnerabilities of DNN through adversarial attacks. As described above, the size of the simulated FDIA noise in D_{train} follows a Gaussian distribution.

Therefore, it is intuitive that a larger injected false data \mathbf{a} will result in a higher probability to be detected by F . On the contrary, the false measurements \mathbf{z}_a become difficult to be distinguished by F if \mathbf{a} is small, but the FDIA performance also becomes worse.




Fig. 3: Vanilla attack recall.




Fig. 4: Vanilla attack bias.

Baseline: To demonstrate the effectiveness of adversarial attacks, we set a vanilla attack as a baseline. The vanilla attack will simply multiply the injected false data \mathbf{a} with a factor α in the four test datasets, and the newly crafted false measurements are fed into F for evaluation. The evaluation results of vanilla attack are shown in Fig. 3 and Fig. 4 respectively. As analyzed, from Fig. 3 we can learn that a larger α will result in a higher detection recall of the false measurements, and the detection performance may be various for different attack scenarios. Fig. 4 demonstrates that the trend of Bias L_2 -Norm and Valid L_2 -Norm of the vanilla attack is adverse with the α increases under vanilla attack. Overall, if the vanilla attacker aims to obtain a relatively higher probability to bypass the F detection, she/he obtains a high Bias L_2 -Norm and low Valid L_2 -Norm.

We then evaluate the performance of adversarial attacks, as summarized in Table III. The evaluation demonstrates that the adversarial attacks can decrease the detection recall significantly with a relatively low bias L_2 -Norm and high valid L_2 -Norm. The attack performance can be various for different attack scenarios and slightly affected by the attack parameters **step size** (*size*) [20], [32], [42]. By comparing Table III with Fig. 3 and Fig. 4, we can learn that the adversarial attacks significantly out-performs vanilla attacks.

TABLE III: Attack Performance of Plain DNN

Case	Size	Recall	Bias L_2 -Norm	Valid L_2 -Norm
75	0.1	3.9%	79.7	112.2
	0.5	0.9%	87.2	111.29
	1.0	1.9%	89.6	108.9
80	0.1	18.9%	64.5	114.4
	0.5	12.9%	76.7	117.5
	1.0	10.9%	80.2	114.1
85	0.1	7.0%	115.2	134.4
	0.5	3.9%	124.0	134.3
	1.0	3.9%	127.0	131.4
90	0.1	12.9%	79.7	165.5
	0.5	3.9%	109.9	178.0
	1.0	5.9%	109.9	178.0

B. Typical Adversarial Defense Methods

1) *Defensive Distillation:* Similar to [33], we select different distillation temperatures (1, 2, 5, 10, 20, 30, 50, 100) to train the corresponding distilled DNN models. Fig. 5 demonstrates the properties of the distilled models. The left figure in Fig. 5 shows that the input sensitivity decreases when the temperature increase, which coordinates with the effect of distillation, and the DNN models become less sensitive to input perturbation. From the right figure, we can learn that there is a slight decrease in both detection accuracy and recall of the distilled models with the temperature increases. However, when the temperature becomes larger, such as larger than 50, the detection accuracy decreases significantly, which indicates a high false-positive rate.




Fig. 5: Distilled DNNs performance.

We evaluate the performance of distilled DNNs under adversarial attacks, as shown in Table IV. The detection recall was low in most attack cases and the bias L_2 -Norm and valid L_2 -norm are comparable with plain DNN. Overall, model distillation presents ineffective defense performance to adversarial attacks in FDIA detection.

2) *Adversarial Training:* We follow the method described in [48] to implement the adversary training. During each training epoch, we generate the adversarial measurements of the mini-batch measurement data with the real-time trained DNN and label them as false. We then add the generated adversarial measurements to the mini-batch data and train the model. Meanwhile, we project the generated measurements to follow the linear constraints defined by Equation (2).

TABLE IV: Attack Performance of Distilled DNNs

Temp	Case	Recall	Bias L_2 -Norm	Valid L_2 -Norm
1	75	2.9%	77.2	116.8
	80	57.9%	209.1	212.5
	85	49.0%	207.0	248.7
	90	23.9%	149.1	225.1
2	75	3.9%	154.2	185.1
	80	27.0%	104.2	185.5
	85	34.9%	186.4	242.5
	90	3.9%	94.9	156.1
5	75	1.0%	90.5	108.6
	80	5.9%	84.6	126.1
	85	0.41%	195.8	148.9
	90	27.0%	158.6	252.1
10	75	5.9%	97.2	140.1
	80	10.9%	71.2	144.0
	85	43.0%	241.0	295.8
	90	7.9%	96.4	190.8
20	75	2.0%	57.6	139.5
	80	12.9%	91.57	155.6
	85	20.0%	145.1	175.4
	90	34.9%	179.2	271.6
30	75	15.9%	94.6	157.8
	80	11.9%	73.8	157.4
	85	47.9%	203.7	222.3
	90	41.9%	196.2	213.4
50	75	38.9%	172.9	233.8
	80	28.9%	123.2	193.2
	85	43.0%	215.1	212.4
	90	40.0%	202.1	268.0
100	75	20.0%	136.2	196.1
	80	18.0%	102.1	181.6
	85	36.0%	201.7	230.9
	90	10.9%	74.6	202.0

* Parameters: $size = 0.1$.

Fig. 6 shows the training loss and detection accuracy during adversarial training.




Fig. 6: Adversarial training process.

As analyzed in Section IV-A2, the projection process significantly increases the training computation overload, which makes it impractical to be employed in large-scale power systems. In our evaluation, the training process takes around 650 seconds to converge and achieves 96.5% overall detection accuracy. For comparison, the normal training process takes around 10 seconds to converge in the same computer.

We evaluate the defense performance of adversarial training, we launched our adversarial attacks to the trained model, and the results are summarized in Table V. In our simulations, overall, adversarial training is demonstrated to increase the

TABLE V: Attack Performance of Adversarial Training

Case	Size	Recall	Bias L_2 -Norm	Valid L_2 -Norm
75	0.1	15%	96.2	103.9
	0.5	7.0%	115.8	113.8
	1.0	5.9%	122.8	115.8
80	0.1	43.1%	147.1	1.53
	0.5	45.7%	359.9	354.4
	1.0	47.2%	640.8	635.1
85	0.1	63.9%	142.0	167.5
	0.5	61.0%	471.8	437.5
	1.0	62.9%	519.0	678.3
90	0.1	62.2%	187.6	231.2
	0.5	66.4%	761.8	768.4
	1.0	69.3%	1485.9	1495.7

robustness of DNN to a certain degree. The detection recall of the adversarial trained DNN is higher than plain DNN in all evaluation scenarios. Meanwhile, although the Valid L_2 -Norm values are very large in some scenarios (**Case 90**), the corresponding Bias L_2 -Norm values are also large. However, for **Case 75**, the defense performance is still limited.




Fig. 7: Training process of the auxiliary f_{adv} for adversarial detection.

3) *Adversarial Detection*: We employ the adversarial detection methods described in [22] and generate the adversarial measurements of all the false measurements in D_{train} . We use the false measurements and their corresponding adversarial measurements to train a binary auxiliary classification DNN F_{adv} . We empirically attempt different structures and parameters of the F_{adv} and observe that its performance is not reliable and the training process does not converge. As analyzed in Section IV-A3, we explain that this result is caused by the similar manifolds shared between the false measurements \mathbf{z}_a and the adversarial measurements \mathbf{z}_{adv} . Fig. 7 provides the training process of an example F_{adv} in our simulation.

To verify our analysis, we utilize the t-Distributed Stochastic Neighbor Embedding (t-SNE) to visualize the manifolds in 2 dimensions, as shown in Fig. 8. From the left figure, we can learn that the false measurements follow different manifolds with the legitimate measurements. This explains the effectiveness of DNN-based FDIA detection and the high detection recall of F . From the right figure, however, we can learn that the manifolds between the false measurements and corresponding adversarial measurements are very similar




Fig. 8: The manifolds of legitimate measurements, false measurements, and adversarial measurements with t-SNE dimensionality reduction.

(overlapping markers). This phenomenon is caused by the physical property of the power system and the constraints defined by Equation (2). The adversarial measurements \mathbf{z}_{adv} can be regarded as special false measurements. Therefore, adversarial detection can not distinguish the adversarial measurements from the DNN's inputs effectively in FDIA detection.

C. Random Input Padding Framework




Fig. 9: Detection performance of random input padding DNNs.

We evaluate the defense performance of our framework based on F and modify the number of the input neurons according to the padding number P for F . The overall detection performance of padded DNNs are shown in Fig. 9. We can learn that with the padding number P increases, the detection accuracy of DNN decreases gradually and becomes stable. The detection accuracy and recall of the padded DNNs reach around 96% and 95% respectively. Fig. 9 shows that the random padding framework only slightly decreases the FDIA detection performance, compared with the plain models ($P = 186$). Our framework pads zeros in front of and/or after the plain measurements, which will not destroy the pattern of normal inputs.

Figure. 10 shows the defense performance of the random input padding framework under different adversarial attacks. The top four figures in Fig. 10 present the detection recall of the padded DNNs. The detection recalls of all attack scenarios increase remarkably when the padding number is relatively small ($P = 187, 188$). After that, the detection recalls increase gradually and will converge to a specific range with the padding number increases. We can observe that the detection recalls under different attack scenarios do not strictly follow the expected $1 - \frac{1}{(P-186)+1}$ ($P \geq 186$) regulation. We note

this is due to the transferability of adversarial measurements under different padding cases. Overall, we can observe that the padding framework significantly increases the resiliency of DNNs against adversarial attacks.

The bottom eight figures in Fig. 10 demonstrate the bias L_2 -Norm and valid L_2 -Norm of the adversarial attacks. In general, we can observe that the bias L_2 -Norm follows a similar trend with valid L_2 -Norm. Compared with the plain DNN, our random padding framework can increase the bias L_2 -Norm, which can decrease the performance of FDIA that targets local marginal price. On the contrary, similar to adversarial training, the valid L_2 -Norm can also be large in some test scenarios. This phenomenon indicates that the attacker may inject considerable noise into the state estimation. However, our framework maintains a high detection recall under all test scenarios, which makes it outperform the other typical defensive methods discussed in this paper.

VI. DISCUSSION AND FUTURE WORK

As we discussed in Section IV, currently, there is no defense mechanism that is robust to all adversarial attacks. In general, adversary examples/measurements will always exist since the DNN is imperfect and an attacker can always modify the input to force the DNN to change the prediction output. For example, if the vanilla attacker in Section V-A sets the factor α to a small value, the resulted false measurements will have a high probability to bypass the DNN's detection. In this paper, we note the attacker can generate his/her adversarial perturbations based on multiple padding cases to increase their transferability and finally bypass the random padding system, as used in [49] and [20]. However, generating transferable perturbations will affect the attack performance, such as the valid L_2 -Norm, and increase labor and resources of the attacker [20].

We proposed the input padding framework for FDIA detection in this paper and the framework is compatible with different models. Inspired by the randomly-selected auto-encoder scheme in [39], we believe our framework can be used together with their system to further increase the DNNs' robustness in FDIA detection. In the future, we will study the implementation and performance of this joint system.

In DNN-based FDIA detection, in addition to decreasing the detection recall, we expect the defense mechanism to increase the bias L_2 -Norm and decrease the valid L_2 -Norm. In our simulations, the random padding framework increases both the bias L_2 -Norm and the valid L_2 -Norm (similar to adversarial training). We consider this as an open problem in our future work and will study defense methods that achieve satisfying performance in all three metrics.

VII. CONCLUSION

The adversarial attacks present a serious threat to DNN-based FDIA detection. In this paper, we study the defense mechanism of adversarial attacks in FDIA and methods to increase the robustness of corresponding DNN models. We first analyze the unique properties of adversarial attack and defense

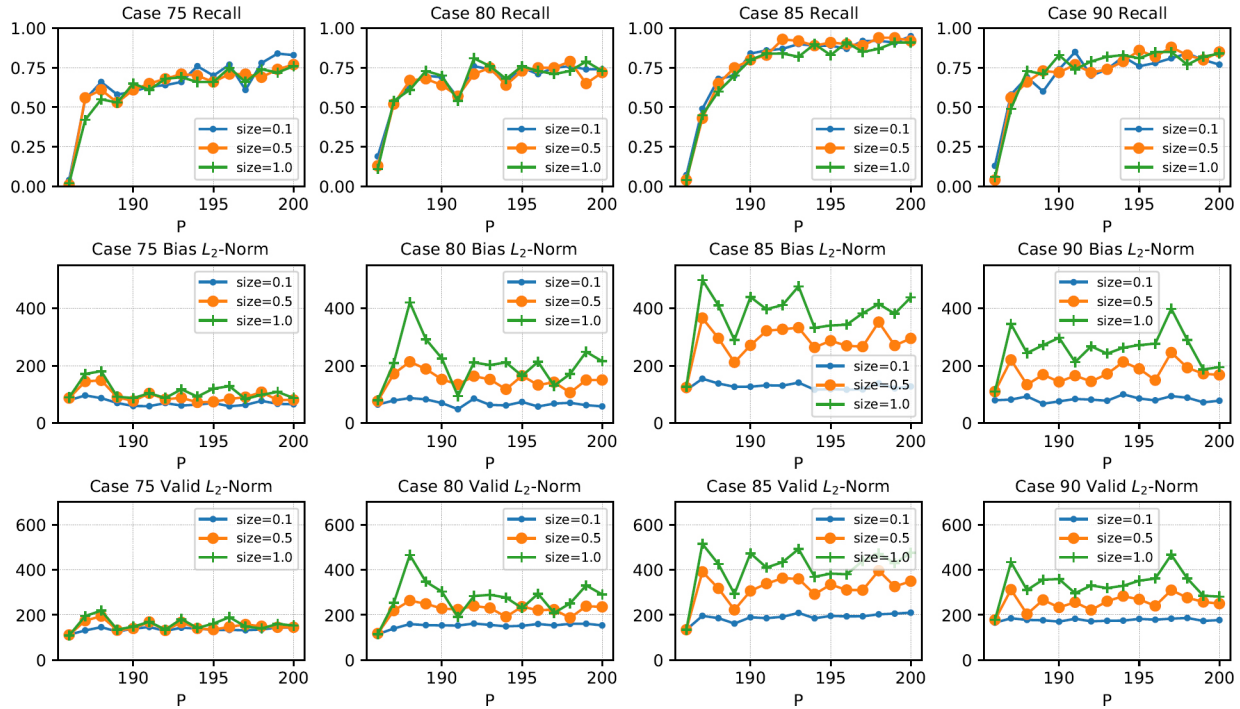


Fig. 10: Adversarial evaluation of random input padding DNNs.

in FDIA detection and summarize the defense requirements. We evaluate several typical defense methods and demonstrate that they have intrinsic limitations for this task through analysis and numeric simulations. After that, we propose a random padding framework for employing DNNs in FDIA detection. The framework is compatible with different DNNs and introduces little extra computation. Evaluation results present that our framework outperforms the typical defense mechanisms and can significantly increase the robustness of DNNs against adversarial measurements in FDIA detection. Meanwhile, the framework will hardly decrease DNNs' detection performance on the normal inputs.

ACKNOWLEDGMENT

This work was partially supported by the US National Science Foundation (NSF) under grant CNS-2038922. Meanwhile, this research was supported in part by the Engineering Research Center Program of the National Science Foundation and the Department of Energy under NSF Award Number EEC-1041877 and the CURENT Industry Partnership Program. We also thank Mr. Eric Reinsmidt and Mr. Jin Young Lee for their suggestions for improving the quality of this paper.

REFERENCES

- [1] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," in *Proceedings of the 16th ACM conference on Computer and communications security*, pp. 21–32, 2009.
- [2] M. A. Rahman and H. Mohsenian-Rad, "False data injection attacks with incomplete information against smart power grids," in *2012 IEEE GLOBECOM*, pp. 3153–3158, IEEE, 2012.
- [3] Z.-H. Yu and W.-L. Chin, "Blind false data injection attack using pca approximation method in smart grid," *IEEE Transactions on Smart Grid*, vol. 6, no. 3, pp. 1219–1226, 2015.
- [4] X. Liu, Z. Li, X. Liu, and Z. Li, "Masking transmission line outages via false data injection attacks," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 7, pp. 1592–1602, 2016.
- [5] S. Bi and Y. J. Zhang, "Defending mechanisms against false-data injection attacks in the power system state estimation," in *2011 IEEE GLOBECOM Workshops (GC Wkshps)*, pp. 1162–1167, IEEE, 2011.
- [6] Q. Yang, D. An, R. Min, W. Yu, X. Yang, and W. Zhao, "On optimal pmu placement-based defense against data integrity attacks in smart grid," *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 7, pp. 1735–1750, 2017.
- [7] M. Ozay, I. Esnaola, F. T. Y. Vural, S. R. Kulkarni, and H. V. Poor, "Machine learning methods for attack detection in the smart grid," *IEEE transactions on neural networks and learning systems*, vol. 27, no. 8, pp. 1773–1786, 2015.
- [8] J. Yan, B. Tang, and H. He, "Detection of false data attacks in smart grid with supervised learning," in *2016 International Joint Conference on Neural Networks (IJCNN)*, pp. 1395–1402, IEEE, 2016.
- [9] Y. He, G. J. Mendis, and J. Wei, "Real-time detection of false data injection attacks in smart grid: A deep learning-based intelligent mechanism," *IEEE Transactions on Smart Grid*, vol. 8, no. 5, pp. 2505–2516, 2017.
- [10] Q. Deng and J. Sun, "False data injection attack detection in a power grid using rnn," in *IECON 2018-44th Annual Conference of the IEEE Industrial Electronics Society*, pp. 5983–5988, IEEE, 2018.
- [11] J. James, Y. Hou, and V. O. Li, "Online false data injection attack detection with wavelet transform and deep neural networks," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 7, pp. 3271–3280, 2018.
- [12] A. Ayad, H. E. Farag, A. Youssef, and E. F. El-Saadany, "Detection of false data injection attacks in smart grids using recurrent neural networks," in *2018 IEEE Conference on Smart Grid Communications (SmartGridComm)*, pp. 1–6, IEEE, 2018.

- networks,” in *2018 IEEE ISGT*, pp. 1–5, IEEE, 2018.
- [13] M. Ashrafuzzaman, Y. Chakhchoukh, A. A. Jillepalli, P. T. Tomic, D. C. de Leon, F. T. Sheldon, and B. K. Johnson, “Detecting stealthy false data injection attacks in power grids using deep learning,” in *2018 14th IWCNC*, pp. 219–225, IEEE, 2018.
- [14] X. Niu, J. Li, J. Sun, and K. Tomsovic, “Dynamic detection of false data injection attack in smart grid using deep learning,” in *2019 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*, pp. 1–6, IEEE, 2019.
- [15] C. Wang, S. Tindemans, K. Pan, and P. Palensky, “Detection of false data injection attacks using the autoencoder approach,” *arXiv preprint arXiv:2003.02229*, 2020.
- [16] Y. Chen, Y. Tan, and D. Deka, “Is machine learning in power systems vulnerable?,” in *2018 IEEE SmartGridComm*, pp. 1–6, IEEE, 2018.
- [17] Y. Chen, Y. Tan, and B. Zhang, “Exploiting vulnerabilities of load forecasting through adversarial attacks,” in *Proceedings of the Tenth ACM International Conference on Future Energy Systems*, pp. 1–11, 2019.
- [18] J. Tian, T. Li, F. Shang, K. Cao, J. Li, and M. Ozay, “Adaptive normalized attacks for learning adversarial attacks and defenses in power systems,” in *2019 IEEE SmartGridComm*, pp. 1–6, IEEE, 2019.
- [19] J. Li, Y. Yang, and J. S. Sun, “Searchfromfree: Adversarial measurements for machine learning-based energy theft detection,” in *2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, pp. 1–6, IEEE, 2020.
- [20] J. Li, Y. Yang, J. S. Sun, K. Tomsovic, and H. Qi, “Conaml: Constrained adversarial machine learning for cyber-physical systems,” in *Proceedings of the 2021 ACM Asia Conference on Computer and Communications Security*, pp. 52–66, 2021.
- [21] T. Liu and T. Shu, “Adversarial false data injection attack against nonlinear ac state estimation with ann in smart grid,” in *International Conference on Security and Privacy in Communication Systems*, pp. 365–379, Springer, 2019.
- [22] J. H. Metzen, T. Genewein, V. Fischer, and B. Bischoff, “On detecting adversarial perturbations,” *arXiv preprint arXiv:1702.04267*, 2017.
- [23] C. Jiongcong, G. Liang, C. Zexiang, H. Chunchao, X. Yan, L. Fengji, and Z. Junhua, “Impact analysis of false data injection attacks on power system static security assessment,” *Journal of Modern Power Systems and Clean Energy*, vol. 4, no. 3, pp. 496–505, 2016.
- [24] M. Khalaf, A. Hooshyar, and E. El-Saadany, “On false data injection in wide area protection schemes,” in *2018 IEEE PESGM*, pp. 1–5, 2018.
- [25] S. Bi and Y. J. Zhang, “Graphical methods for defense against false-data injection attacks on power system state estimation,” *IEEE Transactions on Smart Grid*, vol. 5, no. 3, pp. 1216–1227, 2014.
- [26] T. T. Kim and H. V. Poor, “Strategic protection against data injection attacks on power grids,” *IEEE Transactions on Smart Grid*, vol. 2, no. 2, pp. 326–333, 2011.
- [27] J. Chen and A. Abur, “Placement of pmus to enable bad data detection in state estimation,” *IEEE Transactions on Power Systems*, vol. 21, no. 4, pp. 1608–1615, 2006.
- [28] 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.
- [29] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” *arXiv preprint arXiv:1412.6572*, 2014.
- [30] A. Rozsa, E. M. Rudd, and T. E. Bault, “Adversarial diversity and
- [36] A. Shafahi, M. Najibi, M. A. Ghiasi, Z. Xu, J. Dickerson, C. Studer, L. S. Davis, G. Taylor, and T. Goldstein, “Adversarial training for free!” hard positive generation,” in *Proceedings of the IEEE CVPR Workshops*, pp. 25–32, 2016.
- [31] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, “Deepfool: a simple and accurate method to fool deep neural networks,” in *Proceedings of the IEEE CVPR*, pp. 2574–2582, 2016.
- [32] A. Kurakin, I. Goodfellow, and S. Bengio, “Adversarial machine learning at scale,” *arXiv preprint arXiv:1611.01236*, 2016.
- [33] N. Papernot, P. McDaniel, X. Wu, S. Jha, and A. Swami, “Distillation as a defense to adversarial perturbations against deep neural networks,” in *2016 IEEE Symposium on Security and Privacy (SP)*, pp. 582–597, IEEE, 2016.
- [34] S. Ma and Y. Liu, “Nic: Detecting adversarial samples with neural network invariant checking,” in *Proceedings of the 26th NDSS*, 2019.
- [35] W. Xu, D. Evans, and Y. Qi, “Feature squeezing: Detecting adversarial examples in deep neural networks,” *arXiv preprint arXiv:1704.01155*, 2017.
- in *Advances in Neural Information Processing Systems*, pp. 3353–3364, 2019.
- [37] F. Tramèr, A. Kurakin, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel, “Ensemble adversarial training: Attacks and defenses,” *arXiv preprint arXiv:1705.07204*, 2017.
- [38] S. Gu and L. Rigazio, “Towards deep neural network architectures robust to adversarial examples,” *arXiv preprint arXiv:1412.5068*, 2014.
- [39] D. Meng and H. Chen, “Magnet: a two-pronged defense against adversarial examples,” in *Proceedings of the 2017 ACM CCS*, pp. 135–147, 2017.
- [40] C. Xie, J. Wang, Z. Zhang, Z. Ren, and A. Yuille, “Mitigating adversarial effects through randomization,” in *International Conference on Learning Representations*, 2018.
- [41] H. Xu, Y. Ma, H.-C. Liu, D. Deb, H. Liu, J.-L. Tang, and A. K. Jain, “Adversarial attacks and defenses in images, graphs and text: A review,” *International Journal of Automation and Computing*, vol. 17, no. 2, pp. 151–178, 2020.
- [42] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” *arXiv preprint arXiv:1706.06083*, 2017.
- [43] A. Athalye, N. Carlini, and D. Wagner, “Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples,” in *International Conference on Machine Learning*, pp. 274–283, PMLR, 2018.
- [44] G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *arXiv preprint arXiv:1503.02531*, 2015.
- [45] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas, “Matpower: Steady-state operations, planning, and analysis tools for power systems research and education,” *IEEE Transactions on power systems*, vol. 26, no. 1, pp. 12–19, 2010.
- [46] S. Bi and Y. J. Zhang, “False-data injection attack to control real-time price in electricity market,” in *2013 IEEE Global Communications Conference (GLOBECOM)*, pp. 772–777, IEEE, 2013.
- [47] L. Xie, Y. Mo, and B. Sinopoli, “False data injection attacks in electricity markets,” in *2010 First IEEE International Conference on Smart Grid Communications*, pp. 226–231, IEEE, 2010.
- [48] A. Kurakin, I. Goodfellow, and S. Bengio, “Adversarial examples in the physical world,” *arXiv preprint arXiv:1607.02533*, 2016.
- [49] S.-M. Moosavi-Dezfooli, A. Fawzi, O. Fawzi, and P. Frossard, “Universal adversarial perturbations,” in *Proceedings of the IEEE CVPR*, pp. 1765–1773, 2017.