

Exploiting Vulnerabilities in Deep Neural Networks: Adversarial and Fault-Injection Attacks

Faiq Khalid¹, Muhammad Abdullah Hanif¹, Muhammad Shafique²

¹Technische Universität Wien (TU Wien), Vienna, Austria

²Division of Engineering, New York University Abu Dhabi (NYUAD), Abu Dhabi, United Arab Emirates

Email: {faiq.khalid,muhammad.hanif}@tuwien.ac.at, muhammad.shafique@nyu.edu

Abstract—From tiny pacemaker chips to aircraft collision avoidance systems, the state-of-the-art Cyber-Physical Systems (CPS) have increasingly started to rely on Deep Neural Networks (DNNs). However, as concluded in various studies, DNNs are highly susceptible to security threats, including adversarial attacks. In this paper, we first discuss different vulnerabilities that can be exploited for generating security attacks for neural network-based systems. We then provide an overview of existing adversarial and fault-injection-based attacks on DNNs. We also present a brief analysis to highlight different challenges in the practical implementation of the adversarial attacks. Finally, we also discuss various prospective ways to develop robust DNN-based systems that are resilient to adversarial and fault-injection attacks.

Keywords—Deep Neural Networks; Adversarial Attacks; Machine Learning Security; Fault-injection Attacks.

I. INTRODUCTION

Machine Learning (ML) algorithms have become popular in many applications, especially smart Cyber-Physical Systems (CPS), because of their ability to process and classify the enormous data [1]–[3], e.g., image recognition, object detection. The state-of-the-art ML systems are mostly based on Deep Neural Networks (DNNs), which consists of many layers of neurons connected into a mesh. The input follows this mesh from the input layer, through the hidden layers, to reach the output layer. The output node/neuron with the highest value indicates the decision of the DNN. However, due to several uncertainties in feature selection, ML algorithms inherently possess several security vulnerabilities, e.g., sensitivity for small input noise. Several attacks have been proposed to exploit these security vulnerabilities, for example, adversarial attacks [4]–[6] (see Figure 1). Depending upon the extent of the access to the training process, the DNN model, or the inference, the adversarial attacks can either exploit the input sensitivity of the trained DNN during the inference or manipulate the training process, DNN model or training dataset.

Since the discovery of adversarial attacks, several studies have been performed but are mainly focused on the optimization and effectiveness of input data perturbation. In most of the practical cases, it is very challenging to manipulate the input of the DNN or corrupt the input data because of the limited access of DNN-based system or training process. However, the advancements in the communication network and computational elements of the CPS make us capable of putting the computing elements and sensors anywhere. This enables easy access to the computing elements and sensor, thereby opening a broader attack surface. Recently, this easy-to-get physical access is exploited to generate fault-injection attacks (see Figure 1). In these attacks, attackers can externally inject the faults in data stored in the memory, the control path of a DNN-based system, or the computational blocks to manipulate the DNN output. These faults can be injected using well-known techniques, e.g., variations in voltage, Electromagnetic (EM) interference, and heavy-ion radiation.

A. Novel Contributions

To encompass the complete attack surface in the DNN-based system, in this paper, we study the security vulnerabilities with respect to adversarial attack and the emerging fault-injection attacks. **In summary, the contributions of this paper are:**

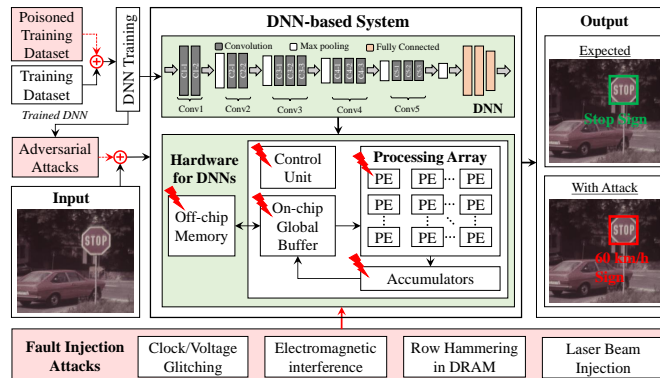


Figure 1. Security threats (i.e., adversarial and fault-injection attacks) to a DNN-based system. Fault-injection attacks can be performed by physically injecting faults in the off-chip and on-chip memories, and in the Processing Elements (PEs). The “Stop” sign image is taken from the German Traffic Sign Detection dataset [7].

- 1) We first study and discuss the different sets of assumptions and parameters of the threat model that can be used to show the effectiveness and practicability of an attack (**Section II**).
- 2) We discuss the different aspects, with respect to threat model, optimization algorithms, and computational cost, of the adversarial attacks (**Section III**) and fault-injection attacks (**Section V**).
- 3) Most of the adversarial attacks do not consider the overall pipeline in the DNN-based system and ignore the pre-processing stages. Therefore, in this paper, we provide a brief analysis to highlight the challenges in the practical implementation of the adversarial attacks (**Section IV**). Based on this analysis, we conclude that the adversarial attack must be strong enough not to be nullified by the input pre-processing (which is low-pass filtering in our analysis).
- 4) Towards the end, we highlight different research directions on the road ahead towards developing a robust DNN-based system with stronger defenses against these attacks (**Section VI**).

II. THREAT MODEL

The effectiveness of an attack depends upon different assumptions and parameters. Different configurations of these assumptions and parameters generate several scenarios that are known as threat models. Therefore, in the context of DNN security, the threat model is based on the following set of parameters (see Figure 2):

- **Attacker’s Knowledge:** Depending upon the information about the targeted ML system and attacker’s access, the attack can either be black-box or white-box (see Figure 3). In *white-box attacks*, the attacker has full access to the trained DNNs; thereby, allowing him to exploit DNN parameters to generate adversarial noise. In *black-box attacks*, attacker has access only to the input and output of the DNN.
- **Attacker’s Goal:** Depending upon the targeted payload, the attack can either be targeted or un-targeted. In the *targeted attack*, the attacker modifies the DNN, input, or other parameters to achieve a particular misclassification. On the other hand, in the *un-targeted attack*, the attacker’s aim is only to maximize the prediction error.




Figure 2. Threat model for adversarial attacks on DNNs. This model shows different assumptions and parameters that are required to generate an adversarial attack.

- **Attack Frequency:** To exploit a trade-off between resources, timing, and effectiveness, the attack can be wither one-shot or iterative. In *one-shot*, the attack is optimized only once, but *iterative* attack optimizes the payload of the attack over multiple iterations.
- **Attack Falsification:** Depending upon the payload of the attack, these attacks can either be false positive or false negative attacks. In *false-positive* attacks, a negative sample is misclassified as a positive sample. In *false-negative* attacks, a positive sample is misclassified as a negative sample.
- **Attack Type:** This parameter is related to the targeted phase of the ML design cycle. For example, depending upon the access, the attacker can target training, inference, or hardware of the DNN.
- **Dataset Access:** The attack effectiveness and strength also depends upon the attacker's access to the different datasets. For example, the strength of the evasion attacks can significantly increase if the attacker has access to training and testing datasets.

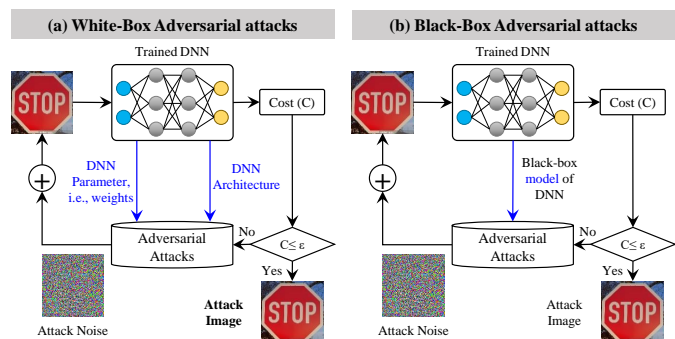


Figure 3. An overview of the adversarial attack on machine learning. (a) In a white-box setting, the attacker has access to the network architecture and network parameters, i.e., weight, number of neurons, number of layers, activation functions, and convolution filters. (b) In a black-box setting, the attacker has access to the input and output of the network.

III. ADVERSARIAL ATTACKS

The addition of imperceptible noise to the input can change the output of the DNN, and this phenomenon is known as the adversarial attack. It can also reduce the confidence classification that cause un-targeted misclassification, or it can also induce targeted misclassification. Several adversarial attacks have been proposed that manipulate the system to behave erroneously. Broadly, these attacks are categorized as causative and exploratory attacks.

A. Evasion (Exploratory) Attacks

In these attacks, an attacker introduces an imperceptible noise at the input of the trained DNN during the inference. This imperceptible noise (known as adversarial noise) can either perform targeted misclassification or maximize the prediction error. Since these attacks explore vulnerabilities of the trained DNN during inference, therefore, these attacks are also known as exploratory attacks. On the basis

of how the attacker implements the adversarial attack, the attacks can be divided into three categories: *gradient-based*, *score-based* and *decision-based* attacks, as shown in Table I.

1) *Gradient-based Attacks:* The gradient-based attacks make use of the network gradients to craft the attack noise. There are several gradient-based attacks that utilize different optimization algorithms and imperceptibility parameters to improve the effectiveness of these attacks (see the summary of gradient-based attacks in Table I). However, most of them are based basic gradient-based attacks, i.e., Fast Gradient Sign Method (FGSM) [8], iterative-FGSM (iFGSM) [9], Jacobian Saliency Map Attack (JSMA) [12], Carlini and Wagner (C&W) [14] attack and training dataset unaware attack (TrISec) [26]. In this section, we only discuss these basic gradient-based attacks.

- *Fast Gradient Sign Method (FGSM)* [8] is a gradient-based attack that utilizes use of the cost function $J(p, x, y_{True})$ for the network with parameters p , input x and the correct output class y_{True} to determine the direction in which the adversarial noise will have the greatest impact. Afterwards, a small noise ϵ is added in a single iteration to the input, in the direction of the obtained gradient.
- The *Iterative Fast Sign Method (iFGSM)* [9] is a variant of FGSM that is preferred for targeted misclassification. Here, the cost function chosen corresponds to a specific target (misclassification) class. Instead of perturbing the input in one step, the input is perturbed by α in each step.
- In *Jacobian Saliency Map Attack (JSMA)* [12], initially the forward derivative, i.e., the Jacobian, of the network F with respect to all input nodes is determined. Accordingly, saliency map is constructed, which highlights the inputs that are the most vulnerable to noise. This map can be used to perturb the minimum number of inputs to implement a successful attack. A vulnerable input node is then perturbed towards the target class. If the perturbation is found to be insufficient, then the whole process of finding Jacobian and constructing the saliency map is repeated until a successful attack is implemented or the perturbation added to the input exceeds a given threshold γ .
- The *Carlini and Wagner (C&W) attack* [14] is a white-box approach to adversarial attack. The main idea is to consider finding the appropriate perturbation for the adversarial attack as a dual optimization problem. The first objective is to minimize the noise with respect to l_p norm, added to an image with the objective of obtaining a specific target label at output. The second objective is to minimize $f(x + \epsilon)$ to a non-positive value, where the output label of the input changes to the targeted misclassification label.
- *TrISec* [26] proposes a training data “unaware” adversarial attack. The attack is again modeled as an optimization problem. Backpropagation is used to obtain the noise that causes targeted misclassification while minimizing the cost function associated with the network. The probability $P(F(x + \epsilon) = y_{Target})$ of the target class is simultaneously maximized to ensure a minimal and scattered perturbation based on two parameters, i.e., Cross-correlation Coefficient (CC) and Structural Similarity Index (SSI). CC ensures that the noise added to the input is imperceptible (to humans), while SSI ensures that noise is spread across the whole input instead of being focused in a small region, hence, improving the imperceptibility with respect to subjective analysis.

Note, these attacks require the attacker to have access to the network's internal parameters to calculate the gradients and hence are generally carried out in a white-box environment. Note, all the white-box attacks can be implemented in the black-box environment when they are combined with model stealing attacks. These attacks are known as substitute model attacks. However, it is not always possible to have the white-box access to the trained DNN.

TABLE I. A BRIEF OVERVIEW OF THE CATEGORIZATION OF THE STATE-OF-THE-ART ADVERSARIAL ATTACKS ON MACHINE LEARNING SYSTEMS. (SVM: SUPPORT VECTOR MACHINES, SSI: STRUCTURAL SIMILARITY INDEX, AND CC: CROSS-CORRELATION COEFFICIENT)

Adversarial Attacks		Type	Knowledge	Frequency	Goal	Imperceptibility	
Gradient-based Attacks	Fast Gradient Sign Method (FGSM)[8]	Evasion	White-Box	One-shot	Targeted/Un-targeted	l_1, l_2, l_{inf} norms	
	Basic Iterative Method (BIM) or Iterative FGSM[9]	Evasion	White-Box	Iterative	Targeted/Un-targeted	l_1, l_2, l_{inf} norms	
	Projected Gradient Descent (PGD)[10]	Evasion	White-Box	Iterative	Targeted/Un-targeted	ϵ norm	
	Auto-PGD[11]	Evasion	White-Box	Iterative	Targeted/Un-targeted	l_{inf} norm	
	Jacobian-based Saliency Map Attack (JSMA)[12]	Evasion	White-Box	Iterative	Targeted/Un-targeted	l_0 norm	
	Iterative Frame Saliency[13]	Evasion	White-Box	Iterative/One-shot	Un-targeted	l_0 norm	
	Carlini & Wagner l_2 attack[14]	Evasion	White-Box	Iterative	Targeted/Un-targeted	l_2 norm	
	Carlini & Wagner l_{inf} attack[14]	Evasion	White-Box	Iterative	Targeted/Un-targeted	l_{inf} norm	
	DeepFool[15]	Evasion	White-Box	Iterative	Un-targeted	l_2 norm	
	Universal Perturbations[16]	Evasion	White-Box	Iterative	Un-targeted	l_p norm	
	Newton Fool[17]	Evasion	White-Box	Iterative	Un-targeted	Tuning parameter	
	Feature Adversaries[18]	Evasion	White-Box	Iterative	Targeted	l_{inf} norm	
	Adversarial Patch[19][20]	Evasion	White-Box	Iterative	Targeted/Un-targeted	-	
	Elastic-Net (EAD)[21]	Evasion	White-Box	Iterative	Targeted/Un-targeted	l_1, l_2, l_{inf} norms	
	Dpatch[22]	Evasion	White-Box	Iterative	Targeted/Un-targeted	-	
	Score-based Attacks	High Confidence Low Uncertainty[23]	Evasion	White-Box	Iterative	Targeted	l_2 norm
Waasserstein Attack[24]		Evasion	White-Box	Iterative	Targeted/Un-targeted	l_p norm	
Shadow Attack[25]		Evasion	White-Box	Iterative	Targeted/Un-targeted	l_2, l_{inf} norms	
TriSec[26]		Evasion	White-Box	Iterative	Targeted/Un-targeted	SSI, CR	
Zeroth Order Optimization (ZOO)[27]		Evasion	Black-Box	Iterative	Targeted	l_2 norm	
Local Search[28]		Evasion	Black-Box	Iterative	Un-targeted	-	
Copy and Paste[29]		Evasion	Black-Box	Iterative	Targeted/Un-targeted	l_2 norm	
HopskipJump[30]		Evasion	Black-Box	Iterative	Targeted/Un-targeted	l_2 norm	
Decision-based Attacks		Query Efficient Attack [31]	Evasion	Black-Box	Iterative	Targeted/Un-targeted	l_2 norm
		Decision-based Attack [32]	Evasion	Black-Box	Iterative	Targeted/Un-targeted	l_2 norm
		Query Efficient Boundary Attack (QEBA)[33]	Evasion	Black-Box	Iterative	Targeted/Un-targeted	l_2 norm
		Geometry-Inspired Decision-based (qFool)[34]	Evasion	Black-Box	Iterative	Targeted/Un-targeted	l_2 norm
		Threshold Attack[35]	Evasion	Black-Box	Iterative	Targeted/Un-targeted	l_p norm
		Square Attack[36]	Evasion	Black-Box	Iterative	Targeted/Un-targeted	l_p norm
		Pixel Attacks[37][35]	Evasion	Black-Box	Iterative	Targeted/Un-targeted	l_p norm
		FaDec[38]	Evasion	Black-Box	Iterative	Targeted/Un-targeted	l_2 norm, SSI, CR
Dataset Poisoning	Poisoning Attack on SVM [39][40]	Poisoning	White-Box	Iterative	Targeted	-	
	Targeted Clean-Label Poisoning [41]	Poisoning	White-Box	One-shot	Targeted	-	
	Watermarking [42]	Poisoning	White-Box	One-shot	Targeted	-	
	Efficient Dataset Poisoning [42]	Poisoning	White-Box	Iterative	Targeted/Un-targeted	-	
	BadNets [43]	Poisoning	White-Box	Iterative	Targeted	-	
	Targeted Backdoor [44]	Poisoning	White-Box	Iterative	Targeted	-	
	Dynamic Backdoor Attacks [45]	Poisoning	White-Box	Iterative	Targeted	-	
	Feature Collision Attack [42]	Poisoning	White-Box	Iterative	Targeted/Un-targeted	-	
Model Poisoning	Weight poisoning [46]	Poisoning	White-Box	One-shot	Targeted/Un-targeted	-	
	Local Model Poisoning [47]	Poisoning	White-Box	One-shot	Targeted/Un-targeted	-	

2) *Score-based Evasion Attacks*: In these attacks, the attacker has access to output scores/probabilities [27][28]. The generation of attack is formulated as an optimization problem where the change in output scores due to the input manipulation is used to predict the direction and strength of the next input manipulation.

3) *Decision-based Evasion Attacks*: These attacks make input alterations in a reverse direction - the procedure starts with a larger input noise that causes output misclassification. The manipulations are then iteratively reduced until they are imperceptible, while still triggering the adversarial attack [30]–[37]. Since these attack aims to estimate the adversarial example at the classification boundary, therefore, these attacks are also known as boundary attacks. In these attacks, the attack starts from a seed input from the target class (in case of a targeted attack) or any other incorrect output class (for random misclassification). The algorithm progresses iteratively towards the decision boundary of the true output class for input under attack. The objective is not only to reach the decision boundary, but also to explore the different parts of the boundary to ensure that a minimum amount of noise is being added to the input. Hence, no knowledge of the DNN's gradients, parameters, or output scores is required for a successful attack. However, the cost of these attacks in terms of the number of queries is very large.

To reduce the number of queries, a *Resource Efficient Decision-based attack (FaDec-attack)* [38] finds targeted and un-targeted adversarial perturbations at a reduced computational cost. Like the boundary attack, a seed input from an incorrect class is chosen. The seed is iteratively modified to minimize its distance

to the classification boundary of the original input. To reduce the computational cost of these attacks, adaptive step sizes are used to reach the smallest perturbation in the least number of iterations.

B. Poisoning (Causative) Attacks

In these attacks, the attacker manipulates the training algorithm, un-trained model, training dataset to influence, or corrupt the ML model itself. Based on the targeted components of the training process, these attacks can be categorized as dataset poisoning and model poisoning attacks (see Table I).

- *Dataset Poisoning*: In these attacks, attacker manipulate the training dataset by adding patches (tailored noise) or randomly distributed noise [39]–[44], [48]. These poisoned images influence different parameters of the DNN model such that it performs either target misclassification or maximizes the classification error. Since these attacks introduce the backdoors in the trained network that can be exploited during inference, therefore, these attacks are also known as backdoor attacks.
- *Model Poisoning*: Another type of causative attack is to slightly modify the DNN architecture such that for a particular trigger, it performs either target misclassification or maximizes the classification error [46][47].

IV. PRACTICAL IMPLICATION OF EVASION ATTACKS

In adversarial attacks, the attack noise must be large enough to be captured by the acquisition device (for instance, camera) but




Figure 4. Effects of the preprocessing filters on some white-box adversarial attacks. It can be observed that introducing a simple low-pass filter can significantly impact the robustness of the adversarial attack. For example, in all attacks, the top-5 accuracy increases significantly, i.e., 10% to 89% with the addition of input preprocessing. However, as the smoothing factor of the filters surpasses a certain threshold, the top-5 accuracy started to decrease. The reason behind this behavior is that after this threshold, the filters started to affect the important features.

imperceptible to the subjective analysis (by a human observer). A complete pipeline of the DNN-based classifier consists of an input sensor (for instance, a camera), a preprocessing module (e.g., filters), and a classifier. The robustness of the adversarial attacks depends upon the attacker’s access to the different parts of the pipeline stages.

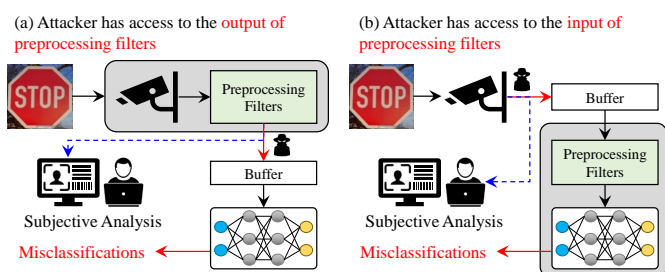


Figure 5. Practical implications of the adversarial attacks with respect to the attacker access. (a) Attack Model I: the attacker has access to the output of preprocessing filters. (b) Attack Model II: the attacker has access to the input of preprocessing filters. Note, for a successful attack, the adversarial noise should be robust to environmental changes.

Traditionally, the adversarial attacks assumed that the attacker has access to the pre-processed input, as shown in Figure 5(a). However, in real-world scenarios, it is very difficult to get access to the output of the preprocessing. The more realistic attack model considers the attacker to have access to the input of the preprocessing module (see Figure 5(b)), for example, the camera is compromised, and it is generating and adding the adversarial noise. To analyze the impact of attack model II, we perform an analysis with low pass filters as a preprocessing module for basic white-box adversarial attacks, i.e., FGSM, JSMA, iFGSM, and TrISec. In this analysis, we choose the two commonly used noise filters: *Local Average with neighborhood Pixels* (LAP) and *Local Average with Radius* (LAR). Figure 4 shows the impact of preprocessing filters on the adversarial attack and the key observations from the analysis are as follows:

- 1) In the case of attack model II, the filters significantly reduces the effectiveness of all the implemented adversarial attacks.
- 2) The increase in the number of neighboring pixels in the LAP filter worsens the performance of the DNN because it affects the key features of the input. Similarly, an increase in the LAR filter radius debilitates the performance of the DNN.

Based on the above-mentioned observations, in the context of adversarial attacks, we identify the following research directions:

- 1) To increase the robustness of adversarial attacks, it is imperative to incorporate the effects of preprocessing modules. For example, recently, researchers have presented that by incorporating the effects of preprocessing filters in optimization algorithms of existing adversarial attacks, the robustness of these attacks can be

increased [49]. Although this analysis considers only white-box attacks, it can effectively be extended to black-box attacks.

- 2) On the other hand, under a particular attack setting, preprocessing modules can also be used to nullify the adversarial attacks. For example, quantization [50], Sobel filters [51] and transformations [52] are used to defend against adversarial attacks. However, the scope of these defenses is very limited, and most of them are breakable using black-box attacks. Therefore, it is a dire need to explore the practical implications of the adversarial attack to develop more powerful defenses.

V. FAULT-INJECTION ATTACKS

Similar to adversarial attacks where the input to a DNN is modified to achieve misclassification, network parameters or computations can also be modified to achieve the same goal. Fault-injection attacks on DNNs refer to the attacks where an attacker tries to manipulate the output of a DNN by injecting faults in its parameters or in the data or control path of the hardware. There are several techniques that can be used for injecting faults, e.g., variations in voltage/clock signal, EM interference and heavy-ion radiation.

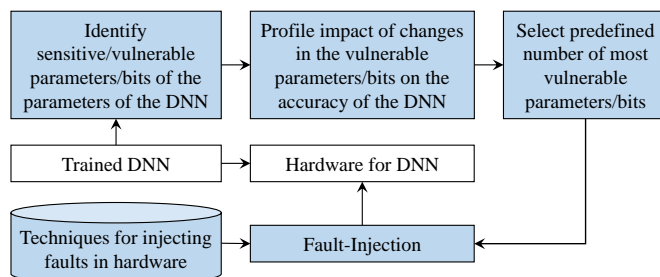


Figure 6. An overview of the design methodology for fault-injection attacks on ML-based systems.

Several studies have been conducted towards modifying the network parameters to attack DNNs. Liu et al. in [53] proposed two fault injection attacks, i.e., Single Bias Attack (SBA) and Gradient Descent Attack (GDA). Since the output of the DNNs is highly dependent on the biases in the output layer, SBA is realized by increasing only a single bias value corresponding to the neuron designated for the adversarial class. SBA is designed for the cases where the stealthiness of the attack is not necessary. For cases where stealthiness is important, GDA has been proposed that uses gradient descent to find the set of parameters to be modified and applies modifications to only some selected ones to minimize the impact of the injected faults on input patterns other than the specified one. Along the same direction, Zhao et al. in [54] proposed fault sneaking attack where they apply Alternating Direction Method of Multipliers (ADMM) [55] to optimize the attack while ensuring that the classification of images other than the ones specified is unaffected

and the modification in the parameters is minimum. A generic flow of the fault-injection attacks on DNNs is shown in Figure 6.

To increase the stealthiness of the attack, Rakin et al. in [56] proposed a methodology, Bit-Flip Attack (BFA), for attacking DNNs by flipping a small number of bits in the weights of the DNN. The bit-flips can be performed through Row-Hammer attack [57][58] or laser injection [59][60] when the weights are in the DRAM or the SRAM of the system, respectively. BFA focuses on identifying the most vulnerable bits in a given DNN that can maximize the accuracy degradation while requiring a very small number of bit-flips in the binary representation of the parameters of the DNN. It employs gradient ranking and progressive search to locate the most vulnerable bits. It is designed for quantized neural networks, i.e., where the weight magnitude is constrained based on the fixed-point representation. For floating-point representation, even a single bit-flip at the most significant location of the exponent of one of the weights of the DNN can result in the network generating completely random output [5] (see Figure 7). The results showed that BFA causes the ResNet-18 network to generate completely random output with only 13 bit-flips on the ImageNet dataset. Figure 9 presents the results for the AlexNet and the ResNet-50 networks under BFA.

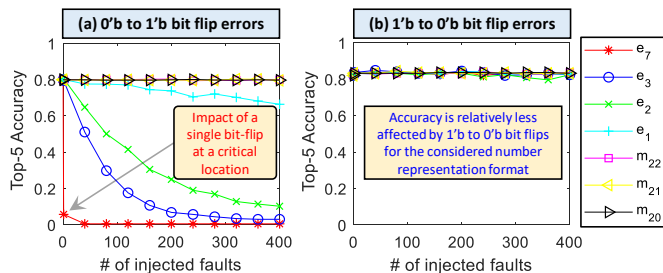


Figure 7. Impact of bit-flips in the weights of the VGG-f network on its classification accuracy on ImageNet dataset [5]. Figure 8 shows the number representation format used for weights and activations for the analysis.

Studies have also been conducted for injecting faults in the computations during the execution of the DNNs. Towards this, Breier et al. in [61] performed an analysis of using a laser to inject faults during the execution of activation functions in a DNN to achieve misclassification. They focused on the instruction skip/change attack model, as it is one of the most basic (and repeatable) attacks for microcontrollers [62], to target four different activation functions, i.e., ReLU, sigmoid, Tanh and softmax. Batina et al. in [63] showed that it is possible to reverse engineer a neural network using a side-channel analysis, e.g., by measuring the power consumption of the device during the execution of a DNN. Therefore, the attack proposed by Breier et al. can be employed even when the DNN is unknown.

VI. RESEARCH DIRECTIONS

Although DNNs are rapidly evolving and becoming an integral part of the decision-making process in CPS. However, the security vulnerabilities of DNNs, e.g., adversarial and fault-injection attacks, raises several concerns regarding their use in CPS. Therefore, stronger defenses against these attacks are required. Towards this, we identify some of the critical research directions:

- 1) Several defenses have been proposed to counter the adversarial attacks [49]–[51]. However, all the available countermeasures are effective only against a particular type of adversarial attacks. This calls for a deeper understanding of the existing attacks, to enable the design of an optimal defense.
- 2) The security measures proposed for adversarial attacks are mainly effective against only a subclass of these attacks. Therefore,




Figure 8. Single-precision floating-point representation




Figure 9. Accuracy vs. the number of bit-flips (N_{flip}) under BFA, for the AlexNet and the ResNet-50 on the ImageNet dataset (data source: [56]).

formal verification of DNNs is emerging as a promising approach to ensure adversarial robustness of DNNs [6][64], which can also help in developing verifiable security measures. The sound mathematical reasoning of formal verification techniques can provide complete and reliable security guarantees to protect DNNs against adversarial attacks.

- 3) To improve the resilience of DNNs against hardware-induced reliability threats, several low-cost fault-mitigation techniques have been proposed, e.g., range restriction-based fault mitigation [65][66]. These techniques have the potential of acting as a strong countermeasure against fault-injection attacks. However, their effectiveness as a countermeasure has not been studied so far. Alongside security, such studies can also help in further improving the reliability of DNN-based systems.

ACKNOWLEDGMENT

This work is supported in parts by the Austrian Research Promotion Agency (FFG) and the Austrian Federal Ministry for Transport, Innovation, and Technology (BMVIT) under the “ICT of the Future” project, IoT4CPS: Trustworthy IoT for Cyber-Physical Systems.

REFERENCES

- [1] M. Shafique et al., “An overview of next-generation architectures for machine learning: Roadmap, opportunities and challenges in the IoT era,” in IEEE DATE, 2018, pp. 827–832.
- [2] F. Kriebel, S. Rehman, M. A. Hanif, F. Khalid, and M. Shafique, “Robustness for smart cyber physical systems and internet-of-things: From adaptive robustness methods to reliability and security for machine learning,” in IEEE ISVLSI, 2018, pp. 581–586.
- [3] A. Marchisio et al., “Deep learning for edge computing: Current trends, cross-layer optimizations, and open research challenges,” in IEEE ISVLSI, 2019, pp. 553–559.
- [4] J. J. Zhang et al., “Building robust machine learning systems: Current progress, research challenges, and opportunities,” in IEEE/ACM DAC, 2019, pp. 1–4.
- [5] M. A. Hanif, F. Khalid, R. V. W. Putra, S. Rehman, and M. Shafique, “Robust machine learning systems: Reliability and security for deep neural networks,” in IEEE IOLTS, 2018, pp. 257–260.
- [6] M. Shafique et al., “Robust machine learning systems: Challenges, current trends, perspectives, and the road ahead,” IEEE Design & Test, vol. 37, no. 2, 2020, pp. 30–57.
- [7] S. Houben, J. Stallkamp, J. Salmen, M. Schlipsing, and C. Igel, “Detection of Traffic Signs in Real-World Images: The German Traffic Sign Detection Benchmark,” in IEEE IJCNN, no. 1288, 2013.
- [8] A. Kurakin, I. Goodfellow, and S. Bengio, “Explaining and harnessing adversarial examples,” arXiv:1412.6572, 2014.
- [9] —, “Adversarial examples in the physical world,” arXiv:1607.02533, 2016.

- [10] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," arXiv:1706.06083, 2017.
- [11] F. Croce and M. Hein, "Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks," arXiv:2003.01690, 2020.
- [12] N. Papernot, P. McDaniel, S. Jha, M. Fredrikson, Z. B. Celik, and A. Swami, "The limitations of deep learning in adversarial settings," in IEEE Euro S&P, 2016, pp. 372–387.
- [13] N. Inkawhich, M. Inkawhich, Y. Chen, and H. Li, "Adversarial attacks for optical flow-based action recognition classifiers," arXiv:1811.11875, 2018.
- [14] N. Carlini and D. Wagner, "Towards evaluating the robustness of neural networks," in IEEE S&P, 2017, pp. 39–57.
- [15] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, "Deepfool: a simple and accurate method to fool deep neural networks," in IEEE/CVF CVPR, 2016, pp. 2574–2582.
- [16] S.-M. Moosavi-Dezfooli et al., "Universal adversarial perturbations," in IEEE/CVF CVPR, 2017, pp. 1765–1773.
- [17] U. Jang, X. Wu, and S. Jha, "Objective metrics and gradient descent algorithms for adversarial examples in machine learning," in ACM ACSAC, 2017, pp. 262–277.
- [18] S. Sabour, Y. Cao, F. Faghri, and D. J. Fleet, "Adversarial manipulation of deep representations," arXiv preprint arXiv:1511.05122, 2015.
- [19] T. B. Brown, D. Mané, A. Roy, M. Abadi, and J. Gilmer, "Adversarial patch," arXiv:1712.09665, 2017.
- [20] A. Liu et al., "Perceptual-sensitive gan for generating adversarial patches," in AAAI, vol. 33, 2019, pp. 1028–1035.
- [21] P.-Y. Chen, Y. Sharma, H. Zhang, J. Yi, and C.-J. Hsieh, "Ead: elastic-net attacks to deep neural networks via adversarial examples," in AAAI, 2018.
- [22] X. Liu, H. Yang, Z. Liu, L. Song, H. Li, and Y. Chen, "Dpatch: An adversarial patch attack on object detectors," arXiv:1806.02299, 2018.
- [23] K. Grosse, D. Pfaff, M. T. Smith, and M. Backes, "The limitations of model uncertainty in adversarial settings," arXiv:1812.02606, 2018.
- [24] E. Wong, F. R. Schmidt, and J. Z. Kolter, "Wasserstein adversarial examples via projected sinkhorn iterations," arXiv:1902.07906, 2019.
- [25] A. Ghiasi, A. Shafahi, and T. Goldstein, "Breaking certified defenses: Semantic adversarial examples with spoofed robustness certificates," arXiv:2003.08937, 2020.
- [26] F. Khalid, M. A. Hanif, S. Rehman, R. Ahmed, and M. Shafique, "Trisec: training data-unaware imperceptible security attacks on deep neural networks," in IEEE IOLTS, 2019, pp. 188–193.
- [27] P.-Y. Chen, H. Zhang, Y. Sharma, J. Yi, and C.-J. Hsieh, "Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models," in ACM WAIS, 2017, pp. 15–26.
- [28] N. Narodytska and S. P. Kasiviswanathan, "Simple black-box adversarial perturbations for deep networks," arXiv:1612.06299, 2016.
- [29] T. Brunner, F. Diehl, and A. Knoll, "Copy and paste: A simple but effective initialization method for black-box adversarial attacks," arXiv:1906.06086, 2019.
- [30] J. Chen, M. I. Jordan, and M. J. Wainwright, "Hopscjumpattack: A query-efficient decision-based attack," in IEEE S&P, 2020, pp. 1277–1294.
- [31] M. Cheng, T. Le, P.-Y. Chen, J. Yi, H. Zhang, and C.-J. Hsieh, "Query-efficient hard-label black-box attack: An optimization-based approach," arXiv:1807.04457, 2018.
- [32] W. Brendel, J. Rauber, and M. Bethge, "Decision-based adversarial attacks: Reliable attacks against black-box machine learning models," arXiv:1712.04248, 2017.
- [33] H. Li, X. Xu, X. Zhang, S. Yang, and B. Li, "Qeba: Query-efficient boundary-based blackbox attack," in IEEE/CVF CVPR, 2020, pp. 1221–1230.
- [34] Y. Liu, S.-M. Moosavi-Dezfooli, and P. Frossard, "A geometry-inspired decision-based attack," in IEEE/CVF CVPR, 2019, pp. 4890–4898.
- [35] D. V. Vargas and S. Kotyan, "Robustness assessment for adversarial machine learning: Problems, solutions and a survey of current neural networks and defenses," arXiv:1906.06026, 2019.
- [36] M. Andriushchenko, F. Croce, N. Flammarion, and M. Hein, "Square attack: a query-efficient black-box adversarial attack via random search," arXiv:1912.00049, 2019.
- [37] J. Su, D. V. Vargas, and K. Sakurai, "One pixel attack for fooling deep neural networks," IEEE TEC, vol. 23, no. 5, 2019, pp. 828–841.
- [38] F. Khalid, H. Ali, M. A. Hanif, S. Rehman, R. Ahmed, and M. Shafique, "Fadec: A fast decision-based attack for adversarial machine learning," 2020, pp. 1–8.
- [39] H. Xiao, B. Biggio, B. Nelson, H. Xiao, C. Eckert, and F. Roli, "Support vector machines under adversarial label contamination," Neurocomputing, vol. 160, 2015, pp. 53–62.
- [40] B. Biggio, B. Nelson, and P. Laskov, "Poisoning attacks against support vector machines," arXiv:1206.6389, 2012.
- [41] C. Zhu et al., "Transferable clean-label poisoning attacks on deep neural nets," arXiv:1905.05897, 2019.
- [42] A. Shafahi et al., "Poison frogs! targeted clean-label poisoning attacks on neural networks," in NIPS, 2018, pp. 6103–6113.
- [43] T. Gu, K. Liu, B. Dolan-Gavitt, and S. Garg, "Badnets: Evaluating backdooring attacks on deep neural networks," IEEE Access, vol. 7, 2019, pp. 47 230–47 244.
- [44] X. Chen, C. Liu, B. Li, K. Lu, and D. Song, "Targeted backdoor attacks on deep learning systems using data poisoning," arXiv:1712.05526, 2017.
- [45] A. Salem, R. Wen, M. Backes, S. Ma, and Y. Zhang, "Dynamic backdoor attacks against machine learning models," arXiv:2003.03675, 2020.
- [46] K. Kurita, P. Michel, and G. Neubig, "Weight poisoning attacks on pre-trained models," arXiv:2004.06660, 2020.
- [47] M. Fang, X. Cao, J. Jia, and N. Gong, "Local model poisoning attacks to byzantine-robust federated learning," arXiv:1911.11815, 2019.
- [48] M. Jagielski, A. Oprea, B. Biggio, C. Liu, C. Nita-Rotaru, and B. Li, "Manipulating machine learning: Poisoning attacks and countermeasures for regression learning," in IEEE S&P, 2018, pp. 19–35.
- [49] F. Khalid, M. A. Hanif, S. Rehman, J. Qadir, and M. Shafique, "Fademl: understanding the impact of pre-processing noise filtering on adversarial machine learning," in IEEE DATE, 2019, pp. 902–907.
- [50] F. Khalid et al., "Qusecnets: Quantization-based defense mechanism for securing deep neural network against adversarial attacks," in IEEE IOLTS, 2019, pp. 182–187.
- [51] H. Ali et al., "Sscnets: Robustifying dnns using secure selective convolutional filters," IEEE Design & Test, vol. 37, no. 2, 2019, pp. 58–65.
- [52] E. Raff, J. Sylvester, S. Forsyth, and M. McLean, "Barrage of random transforms for adversarially robust defense," in IEEE CVPR, 2019.
- [53] Y. Liu, L. Wei, B. Luo, and Q. Xu, "Fault injection attack on deep neural network," in 2017 IEEE/ACM ICCAD, 2017, pp. 131–138.
- [54] P. Zhao, S. Wang, C. Gongye, Y. Wang, Y. Fei, and X. Lin, "Fault sneaking attack: A stealthy framework for misleading deep neural networks," in ACM/IEEE DAC, 2019, pp. 1–6.
- [55] Q. Liu, X. Shen, and Y. Gu, "Linearized admm for nonconvex nonsmooth optimization with convergence analysis," IEEE Access, vol. 7, 2019, pp. 76 131–76 144.
- [56] A. S. Rakin, Z. He, and D. Fan, "Bit-flip attack: Crushing neural network with progressive bit search," in IEEE ICCV, 2019, pp. 1211–1220.
- [57] Y. Kim et al., "Flipping bits in memory without accessing them: An experimental study of dram disturbance errors," ACM SIGARCH Computer Architecture News, vol. 42, no. 3, 2014, pp. 361–372.
- [58] K. Razavi, B. Gras, E. Bosman, B. Preneel, C. Giuffrida, and H. Bos, "Flip feng shui: Hammering a needle in the software stack," in USENIX security, 2016, pp. 1–18.
- [59] B. Selmke, S. Brummer, J. Heyszl, and G. Sigl, "Precise laser fault injections into 90 nm and 45 nm sram-cells," in Springer ICSCRA, 2015, pp. 193–205.
- [60] M. Agoyan, J.-M. Dutertre, A.-P. Mirbaha, D. Naccache, A.-L. Ribotta, and A. Tria, "How to flip a bit?" in IEEE IOLTS, 2010, pp. 235–239.
- [61] J. Breier, X. Hou, D. Jap, L. Ma, S. Bhasin, and Y. Liu, "Practical fault attack on deep neural networks," in ACM CCS, 2018, pp. 2204–2206.
- [62] J. Breier et al., D. Jap, and C.-N. Chen, "Laser profiling for the back-side fault attacks: with a practical laser skip instruction attack on aes," in ACM WCPS, 2015, pp. 99–103.
- [63] L. Batina, S. Bhasin, D. Jap, and S. Picek, "Csi neural network: Using side-channels to recover your artificial neural network information," arXiv:1810.09076, 2018.
- [64] M. Naseer, M. F. Minhas, F. Khalid, M. A. Hanif, O. Hasan, and M. Shafique, "Fannet: formal analysis of noise tolerance, training bias and input sensitivity in neural networks," in IEEE DATE, 2020, pp. 666–669.
- [65] Z. Chen, G. Li, and K. Pattabiraman, "Ranger: Boosting error resilience of deep neural networks through range restriction," arXiv:2003.13874, 2020.
- [66] L.-H. Hoang, M. A. Hanif, and M. Shafique, "Ft-clipact: Resilience analysis of deep neural networks and improving their fault tolerance using clipped activation," in IEEE DATE, 2020, pp. 1241–1246.