

# MERCURY: An Automated Remote Side-channel Attack to Nvidia Deep Learning Accelerator

Xiaobei Yan\*, Xiaoxuan Lou\*, Guowen Xu<sup>†</sup>, Han Qiu<sup>‡</sup>, Shangwei Guo<sup>§</sup>, Chip Hong Chang\*, Tianwei Zhang\*

\* Nanyang Technological University, Singapore

<sup>†</sup> City University of Hong Kong, China

<sup>‡</sup> Tsinghua University, China

<sup>§</sup> Chongqing University, China

{xiaobei002, xiaoxuan001}@e.ntu.edu.sg, {guowen.xu, ECHChang, tianwei.zhang}@ntu.edu.sg,  
qiuhan@tsinghua.edu.cn, swguo@cqu.edu.cn

**Abstract**—DNN accelerators have been widely deployed in many scenarios to speed up the inference process and reduce the energy consumption. One big concern about the usage of the accelerators is the confidentiality of the deployed models: model inference execution on the accelerators could leak side-channel information, which enables an adversary to preciously recover the model details. Such model extraction attacks can not only compromise the intellectual property of DNN models, but also facilitate some adversarial attacks.

Although previous works have demonstrated a number of side-channel techniques to extract models from DNN accelerators, they are not practical for two reasons. (1) They only target simplified accelerator implementations, which have limited practicality in the real world. (2) They require heavy human analysis and domain knowledge. To overcome these limitations, this paper presents MERCURY, the first automated remote side-channel attack against the off-the-shelf Nvidia DNN accelerator. The key insight of MERCURY is to model the side-channel extraction process as a sequence-to-sequence problem. The adversary can leverage a time-to-digital converter (TDC) to remotely collect the power trace of the target model’s inference. Then he uses a learning model to automatically recover the architecture details of the victim model from the power trace without any prior knowledge. The adversary can further use the attention mechanism to localize the leakage points that contribute most to the attack. Evaluation results indicate that MERCURY can keep the error rate of model extraction below 1%.

**Index Terms**—profiled side-channel attacks, DNN accelerator, sequence-to-sequence learning, FPGA, model extraction

## I. INTRODUCTION

Modern deep learning technology exhibits a computationally intensive trend to perform more complex tasks. This leads to the popularity of adopting specific hardware to accelerate computation and reduce energy consumption. Field Programmable Gate Arrays (FPGAs) are a prevalent choice for implementing DNN accelerators, and have been widely deployed in large-scale datacenters by various cloud providers, such as Amazon EC2 F1 [1] and Microsoft Catapult [2].

However, DNN accelerators in the cloud pose new security challenges, and one significant threat is the model extraction attack [3]–[5]. Generally, cloud providers adopt the multi-tenancy policy that facilitates multiple users to share the same FPGA board to enhance the resource utilization [6]. Although the circuits of different users can be logically separated, they

TABLE I: Summary of existing works

Attack	Impl.	Model	Target	Aim	Remote	Automated	# of runs
[8]	VTA [16]	CNN	Layer	A	✓	×	50
[11]	Home-grown	MLP	Model	W	×	×	60K
[12]	Home-grown	CNN	Multiplication	W	×	×	40K
[13]	Home-grown	BNN	Model	A, W	×	A: ×; W: ✓	10K
[14]	Home-grown	CNN	Model	A, W	×	×	–
[15]	NVDLA	CNN	Layer	A	×	×	–
[17]	Home-grown	CNN	Layer	A	✓	×	–
[18]	FINN [19]	BNN	Layer	A	✓	×	100
Ours	NVDLA	CNN	Model	A	✓	✓	1

For Aim: A refers to architecture recovery while W refers to weight recovery.

still share the same power distribution network (PDN). A prior study [7] shows that the supply voltage at different locations of a PDN is not constant and depends on the activity of the logic. Therefore, the voltage fluctuation becomes a critical side channel to leak sensitive information across different parts of the FPGA. The adversary can abuse the multi-tenancy feature to *remotely* launch an on-chip monitor on the same board with the victim’s deep learning model and steal its information *without requiring physical access to the hardware*. It has been commonly referred to as remote side-channel attacks in prior works [8]–[10], exhibiting more flexibility and practicality than physical attacks [11]–[15].

Although various side-channel techniques have been proposed to attack FPGA-based DNN accelerators [20], there is still a huge gap to apply them in practice. (1) **Simplified implementations.** A majority of works only target their homemade DNN accelerators [11]–[14], [17], [18], which are normally simplified, and easy to break. In contrast, full-fledged architectures in the real world usually involve more complex structure designs and optimizations, which complicate the side-channel analysis and attacks. (2) **Simplified models.** A number of works aim to steal Binary Neural Networks (BNNs) with binary values of model weights and activations [13], [18]. The feasibility of their extension to non-binarized DNNs is dubious. (3) **Simplified attack goals.** Many studies only attack individual network layers, but cannot achieve end-to-end extraction of the entire model [8], [15], [17], [18]. This restricts their practical values. Table I summarizes the comparisons of prior studies.

To address the above limitations, we propose MERCURY, a novel power side-channel attack to steal the architecture of DNN models on the practical NVIDIA Deep Learning Accelerator (NVDLA). NVDLA serves as the standard way

for DNN accelerator designs, and has been the mainstream implementation in many products. It includes the complex hardware-software stack, execution pipeline and runtime environment. Such designs cause the side-channel traces to be highly redundant and noisy, and exhibit no one-to-one relationship with the victim model. This significantly increases the attack difficulty (See Section III-B for more discussions).

The core concept of MERCURY is to model the side-channel extraction process as a sequence-to-sequence learning task. We leverage powerful RNN-CTC and Transformer models to recover the architecture of the victim model from the power trace of its inference. Besides, we utilize the attention mechanism to localize the leakage point in the side-channel trace, which can shed light on the potential vulnerability of DNN accelerators and defense directions.

To our best knowledge, the only side-channel attack targeting NVDLA is [15]. However, it has the following limitations: (1) it is not a remote attack and requires physical access to the victim device; (2) it requires the attacker to manually split the execution trace for different layers, and extract each layer individually; (3) it needs to train multiple learning models for each hyper-parameter. MERCURY can overcome these limitations: it can be launched remotely in the multi-tenant cloud context; it enables the adversary to automatically steal the model without any manual analysis or prior knowledge of the victim system. Besides, MERCURY is cost-efficient: the adversary only needs to train one end-to-end model, and run one inference process for extraction of the entire model, while prior attacks need thousands of rounds [11]–[13].

Note that our attack goal is to steal the model architecture, which is the same as some previous works [21]–[24]. Stealing the architecture has high financial incentive as it is basis for building more valuable intellectual properties at incremental cost [21], [23], [25]. Besides, obtaining the architecture details can facilitate other attacks such as adversarial examples and membership inference. We provide two case studies in Section VI to show how MERCURY can enhance these attacks. Some studies proposed methods to extract the model weights [11]–[14]. However, these attacks need to physically access the victim device or inject hardware trojans to collect more informative traces (Table I). How to remotely extract model weights is challenging and we leave it as a future work.

We perform extensive experiments to validate the effectiveness of MERCURY. Evaluation results show that MERCURY can recover victim’s model information with an error rate of 1%. It is robust enough to resist certain levels of noise without a significant performance drop.

## II. BACKGROUND

### A. Nvidia Deep Learning Accelerator (NVDLA)

NVDLA is an open-source configurable architecture designed by Nvidia to accelerate deep learning inference. It is capable of computing convolution, activation, pooling, and normalization operations in the model inference. NVDLA can be configured as a large or small implementation, differing

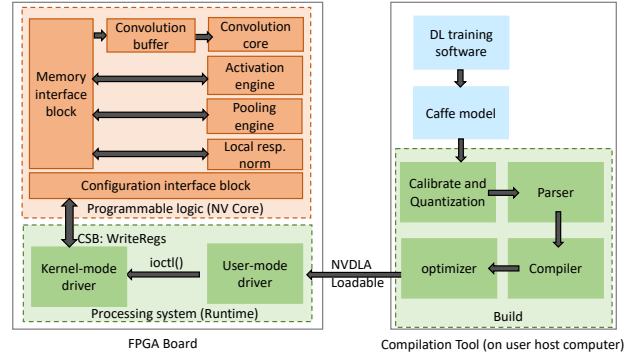


Fig. 1: Architecture overview of NVDLA

in the dimension of the cores and implementation of some specific engines (e.g., Rubik and DMA engines) [26].

Figure 1 shows the architecture overview of NVDLA. It can be divided into two parts, i.e. hardware design and software design. Hardware design is built as a series of pipeline stages containing various types of engines to regulate the behaviors of FPGA boards. Software design connects the users and hardware components, and is responsible for building and loading the DNN model for the FPGA to execute. The software design further consists of two components: (1) the compilation tools use the model pre-compiled by Caffe to generate a network of hardware layers supported by NVDLA, called loadable, which is calibrated by TensorRT. (2) The runtime environment processes the calibrated loadable and runs it directly in the NVDLA environment.

### B. Voltage Drop Sensor on FPGA

All the components on an FPGA chip share one power distribution network (PDN). Intensive switching activities may cause voltage fluctuations in the PDN. The PDN can be modeled as an RLC circuit, where a resistor (R), an inductor (L), and a capacitor (C) are connected in series or in parallel. Therefore, the transient voltage drop seen by a circuit can be modeled as [27]:  $V_{drop} = IR + L \frac{di}{dt}$ , where the transient response term  $L \frac{di}{dt}$  reflects the intensity of the switching activity on the FPGA. In typical CMOS circuits, combinational logic delays can be modeled to be inversely proportional to the voltage supplied to each gate [28]. Hence, the information of the switching activities can be inferred from the logical delay.

In this paper, we use a time-to-digital converter (TDC) to read the combinational logic delay, where a clock signal propagates through a chain of buffers as the voltage drop sensor. As there are discrepancies in switching activities for different types of calculation on other parts of the FPGA, the voltage drop value may be different, which leads to different delay measurements in TDC. These different delays can cause different propagation lengths in the delay line, resulting in different values in the latches. Therefore, the activities of other circuits can be identified through the readout of the TDC. This has been demonstrated in various studies [8]–[10].

### C. Profiled Side-Channel Attacks

Profiled side-channel attacks are one of the most powerful attacks [29]. The adversary generates a profile through a similar or the same device of the target, and computes the secret information by matching the profile with the victim's execution trace. A profiled side-channel attack can be divided into two phases. First, a profiling phase allows the adversary to characterize its physical leakages when running the target application. Suppose the adversary has an input secret set  $s = \{s_1, \dots, s_n\}$ . He obtains  $N$  side-channel traces  $\mathbf{T}_{i,n}$  for each input  $s_i$ , and builds the mapping  $f : \mathbf{T}_{i,n} \mapsto s_i$ . This mapping can be learned through template creating and machine learning methods. Second, an exploitation phase is launched to perform secret recovery by profile matching. The adversary performs the side-channel attack using the mapping  $f$ . With additional  $q$  traces  $T'_1, \dots, T'_q$  collected from the device under attack, the secret  $s'_i$  can be guessed as  $s'_i = f(T'_i)$ .

In this paper, we adopt machine learning for profile learning and matching to attack FPGA-based DNN accelerators. It has several advantages. First, the adversary is able to extract information even if there are no visible patterns in the power trace. As most calculations on FPGAs are done in parallel, manually analyzing the pattern may be futile. Second, the adversary can utilize all information in a single power trace, as machine learning can effectively handle high-dimensional data. In contrast, traditional methods require the selection of points of interest (POI) to narrow down the information for the attack. Besides, machine learning models have higher resilience against the noise in the side-channel trace than conventional statistical methods.

### D. Sequence-to-sequence (seq2seq) Learning

Seq2seq learning achieves state-of-the-art prediction accuracy in various tasks like speech recognition [30], machine translation [31], image captioning [32], question answering [33], etc. Since side-channel power traces are essentially sequential data, seq2seq learning is a natural fit for analyzing such leaking patterns. However, only very few works [21], [24] have applied seq2seq learning to side-channel analysis, and none of them realized automated attacks on FPGA-based accelerators, which have much more noise and more complicated monitoring architecture.

Commonly used architectures for seq2seq learning includes Transformer [34], Connectionist Temporal Classification (CTC), and Recurrent Neural Network (RNN). MERCURY adopts two models, i.e., Transformer and RNN-CTC, for side-channel extraction. In the RNN-CTC model, the output of the RNN is passed into the CTC decoder. Aligning the operation sequence with the variable-length power sequence presents a challenge. To address this, the CTC decoder introduces a "blank" label  $\epsilon$  that serves no specific correspondence and can be easily excluded from the output. The Transformer model consists of an encoder and a decoder. The input sequence is transformed by the encoder, generating an abstract representation that captures the learned features. Subsequently, the decoder utilizes this abstract representation to predict the

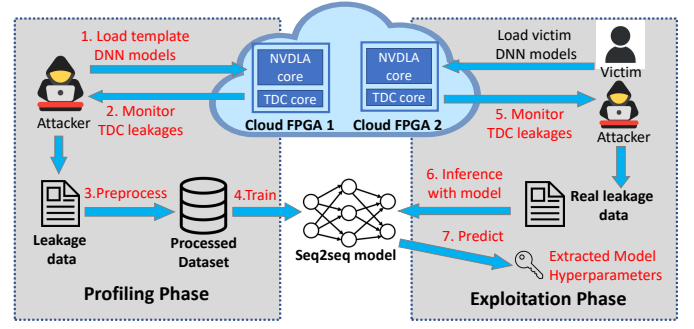


Fig. 2: Attack workflow in MERCURY

subsequent output step-by-step, building upon the previous output. To enhance feature learning from the event sequence, a convolution layer is introduced before the encoder.

## III. ATTACK OVERVIEW

### A. Threat Model

We follow the *same threat model* of remote power side-channel attacks in previous works [8], [9], [35]: the adversary and victim's accelerator share the same FPGA board (but logically separated) with the same PDN, which can be realized in the multi-tenant FPGA-based cloud. The adversary is able to deploy and fully control his own malicious circuits. However, he cannot physically access the FPGA board or control any part of the victim's circuit. He can only deduce the victim's activity from his own circuit.

Figure 2 illustrates the workflow of MERCURY, which consists of two phases. (1) In the *profiling* phase, the adversary deploys various types of DNN models on a cloud FPGA board and collects the corresponding power traces with the TDC sensor. With the profiling information he can establish the seq2seq model  $f$  that predicts the model architecture from the power trace. (2) In the *exploitation* phase, the adversary deploys the TDC circuit on the same FPGA as the victim's model. He only needs to remotely collect the readout of this sensor for *one* inference process of the target model. Then he can use  $f$  to extract the victim model's architecture details.

### B. Challenges of Attack Design

The NVDLA implementation raises the following challenges for designing remote side-channel attacks.

**(1) NVDLA has complicated hardware and software architecture.** From the software perspective, pre-compiled neural networks are first loaded by the user space runtime driver and then submitted to the kernel mode driver for the subsequent inference. The kernel mode driver schedules layer operations on NVDLA and programs the NVDLA registers to configure each functional block. From the hardware perspective, functional computation blocks are assembled as a pipeline to run submitted tasks, and the interrupt signal is asserted when the task is completed. Due to such complex hardware-software co-design, the captured power trace from NVDLA does not directly reflect the information of each layer

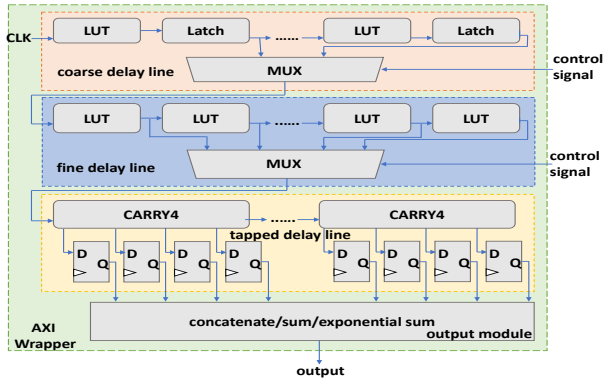


Fig. 3: Architecture details of TDC

and the data inside. It is mixed with other hardware-level power information on the FPGA, as well as the software-level power information related to CPU execution. This is much more complex than other simple accelerator designs, and significantly increases the attack difficulty.

(2) **NVDLA has a parallel design to boost the performance.** All functional blocks in NVDLA have duplicated register groups known as ping-pong buffers. Configurations of the next layer will be transmitted to the other register while running the current layer [36]. Due to this parallel design, NVDLA data flow may not have a strict one-to-one relationship with the layers of a model. These settings make the model extraction attack much more challenging.

(3) **The collected side-channel data are redundant and noisy.** A side-channel trace of an inference process can contain an extremely large amount of data points (more than 300K) due to the high monitoring frequency, which places heavy demands on the underlying compute infrastructure. Such a scale of data will induce great pressure on the data processing, making it infeasible to perform manual analysis as adopted in existing works. Moreover, the raw side-channel data may exhibit random noise originating from the hardware activities, a consequence of the intricate system optimization and runtime dynamics. This noise has the potential to substantially diminish the accuracy of the extraction process.

#### IV. DESIGN DETAILS

To address the above challenges, we introduce innovative designs for the power monitor, dataset processing and seq2seq model. The detailed mechanisms are elaborated below.

##### A. TDC-based Power Monitor

MERCURY utilizes TDC readouts to infer power information from the inference process. Figure 3 illustrates the architecture details of the TDC, where the clock signal is fed into an adjustable coarse delay line and fine delay line to obtain an initial delay, which is then sent to a tapped delay line. The initial delay can be dynamically configured by controlling the multiplexer (MUX) to modify the number of logic elements forming the coarse and fine delay lines during the calibration process, thus altering the delay duration. To ensure the TDC's

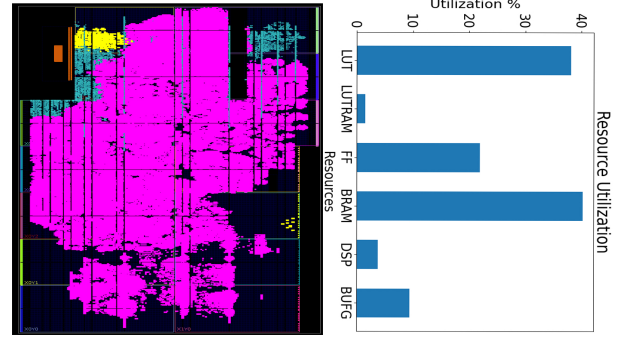


Fig. 4: (Left) The floorplan showing the location of NVDLA (in purple) and the TDC sensor (in yellow) on the FPGA. (Right) Resource utilization on FPGA

functionality is not optimized during synthesis or implementation phases, the `DONT_TOUCH` attribute is set for the coarse delay line, fine delay line, and tapped delay line.

The coarse delay line consists of replicated look-up table (LUT) and latch modules to provide a significant amount of delay. In contrast, the fine delay line is equipped with replicated LUT modules to offer smaller delay. The tapped delay line, which employs carry chains, is composed of CARRY4 primitives, with their CO outputs registered by four dedicated D flip-flops. During each readout, it counts the taps which the clock signal has reached and gives a raw value. The output can then be concatenated or converted into a sum or exponential sum, depending on the configuration set in the TDC IP settings. The TDC output is routed to the ARM processor via AXI-4 buses. We develop a C-based driver program running on the ARM processor to read the TDC output using the `mmap` system call on the addresses specified in the Vivado IP Integrator.

Note that it is important to perform TDC calibration, i.e., adjusting its initial delay, prior to output measurement. We implement the calibration process as two loops in our driver. We iteratively test all combinations of possible values for the fine and coarse delay line lengths to select the optimal initial delay value. This ensures that the Hamming weight in the tapped delay line of the TDC is half the length of the tapped delay line, providing room for power measurements to increase or decrease the Hamming weight. Although prior works [10], [37] highlight the importance of TDC placement and routing constraints for obtaining useful side-channel information, in practice, the adversary may not have the privilege of determining such configurations. In MERCURY, we do not need to manually set the TDC location, yet still successfully initiate the attack. Figure 4 shows the floorplan of our implemented design, with the resource utilization on the FPGA.

##### B. Dataset Formulation and Preprocessing

To build a generalized seq2seq model for side-channel extraction, we need to collect a dataset that covers different types of DNN models and operations in the profiling phase. To achieve this, we generate 160 different random models and deploy them on NVDLA. These models have random numbers






Fig. 5: TDC readouts for one inference process

(in the range of [2, 16]) and types of network layers. They include 12 convolution layers (with the kernel size of 2, 3, 4, 5 and output size of 10, 20, 30), 4 pooling layers (with the kernel size of 2, 3, 4, 5), 5 fully-connected layers (with the output size of 100, 200, 300, 400, 500), 1 relu layer, and 1 softmax layer. These models are pre-trained by Caffe, and then calibrated by TensorRT and compiled by the NVDLA compiler. They are generated on the host computer, and then executed by NVDLA runtime on the FPGA.

The power trace captured by TDC is a sequence in which the TDC readouts are arranged in the temporal order. They represent the switching activities on the FPGA at different moments. To better control the data collection process, we use a bash script on the board to start the TDC measurement program concurrently with the model inference on NVDLA, and terminate it when the inference task is completed.

Figure 5 shows an example of the TDC readouts for one inference. The  $x$ -axis denotes the time while the  $y$ -axis represents the corresponding power consumption value. The middle part of this trace has larger fluctuations, indicating the power-related information leakage from NVDLA. Therefore, we crop the trace and only use the data within this effective period with the time coordinate range [50,000, 200,000] for all traces. Training the model with the full trace also yields successful attack results but requires longer training time. For each model, we collect approximately 2,700 traces as the training set, and 200 traces as the test set. As every single trace has a large amount of data points which may decrease the training efficiency, we reshape it to a 2D matrix with the size of  $3 \times 50,000$ . Then we compute the average of 50 data points, normalized them by subtracting the mean and divided by the standard deviation, for efficient model training.

We label each power trace with a sequence of types for each layer. Given that convolutions take a majority in neural networks, we differentiate different kinds of convolution layers with different labels. We also set four labels for the pooling layer, fully-connected layer, relu layer, and softmax layer, respectively. Table II shows the label for each type of network layer. As required by the CTC decoder (Section IV-C), there is also one label representing a blank operation. Therefore, we have 17 different possible labels in total. The combination of these labels forms the label sequences.

### C. Seq2seq Model

MERCURY adopts two *alternative* seq2seq learning models to extract the DNN architecture from the side-channel trace,

TABLE II: Prediction labels for different types of layers. For example, the label sequence for model with one 5\*5 conv layer (10 output channels) and one fc layer is represented as [9,13].

Layer	Label
conv layer, kernel size 2*2, output channel 10, 20, 30	0-2
conv layer, kernel size 3*3, output channel 10, 20, 30	3-5
conv layer, kernel size 4*4, output channel 10, 20, 30	6-8
conv layer, kernel size 5*5, output channel 10, 20, 30	9-11
pooling layer	12
fully-connected layer	13
relu layer	14
softmax layer	15

as shown in Figure 6. The first one is RNN-CTC. This model uses some convolution layers to extract features from the input sequences, and then 2 RNN layers to propagate the information. To enhance long-term memory capabilities, the RNN module adopts the bidirectional gated recurrent unit (BiGRU). The DNN models in the profiling phase are randomly generated and may not have strong relationships between the layers as they do in the actual functional model design. Hence using RNN may not give the best results. However, we believe it can have better performance in the actual situation as the relation of layers can be utilized. The RNN layer produces a probability distribution for each input, which is subsequently passed into the CTC decoder. Despite the challenge of aligning operation sequences with varying lengths to their corresponding power sequences, the CTC decoder overcomes this hurdle by introducing a "blank" label. Ultimately, the output sequence with the highest prediction probability is determined using beam search. To obtain better training results, Adam optimization and the OneCycleLR scheduler are also used in our model.

The second one is the Transformer model, which adopts a single-layer encoder and single-layer decoder. This model utilizes weight sharing between the decoder embedding and the decoder projection to enhance generalization. It is also trained using the Adam optimizer with OneCycleLR scheduler. The performance is evaluated by the cross-entropy loss function.

The Transformer model can also localize the leakage point in the TDC trace with the attention mechanism, helping us better understand the side-channel vulnerabilities of the DNN inference. Attention is a powerful technique that allows the model to selectively focus on specific parts of the input when making predictions. It computes a weighted sum of input features with the weights representing the importance or relevance of each feature to the current prediction. We can leverage the multi-head attention in the Transformer model to identify the high attention weights, which correspond to the potential leakage points. See Section V-C for more analysis.

## V. EVALUATION

We adopt the Xilinx Zynq-7000 SoC ZC706 board (xc7z045ffg900-2) as our testbed. Due to the restriction of our hardware, we adopt the small implementation of NVDLA. MERCURY can be generalized to the large implementation as they perform the same operations [36]. The ARM processor of the board runs Ubuntu 16.04 OS, which supports NVDLA

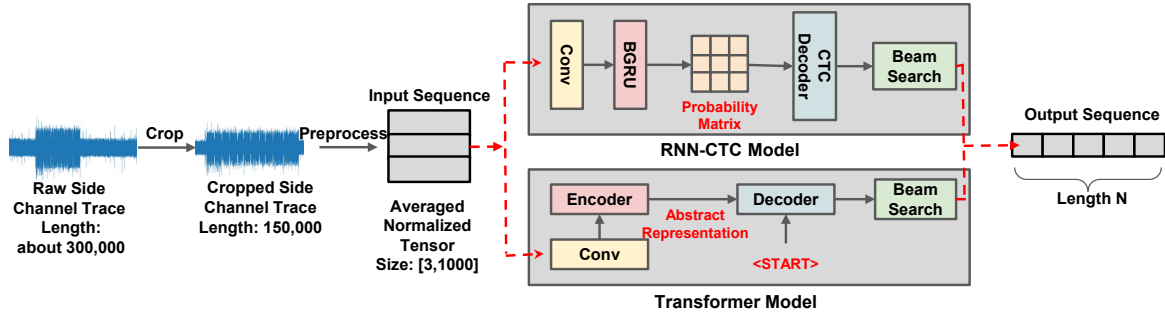


Fig. 6: Overview of two alternative attack models




Fig. 7: Loss and OER trends when training the RNN-CTC model with different numbers of CNN layers.

and the TDC driver. Vivado 2019.1 is used to design the hardware. The clock frequency is 10MHz for NVDLA. The sensor clock of TDC is 150MHz and the AXI clock of TDC operates at 10MHz. The board sends the TDC readouts to the host computer through Ethernet with the `scp` command. Pytorch (1.13) and CUDA (11.6) are adopted to train models running on the server with a Nvidia GeForce RTX 3090 GPU.

We utilize two metrics to evaluate the attack effectiveness. First, as the model architecture is represented as a sequence with each element representing the type of the layer, we adopt Operation Error Rate (OER) [21] to measure the prediction accuracy. It is calculated as  $OER = L(s', s) / \|s\|$ , where  $\|s\|$  is the sequence length of  $s$ , and  $L(s', s)$  is the edit distance (Levenshtein) between the ground-truth sequence  $s$  and predicted operation sequence  $s'$ . A smaller OER indicates higher accuracy. Second, we also use loss functions. In the RNN and Transformer models, we use the CTC loss and cross-entropy loss for evaluation, respectively.

We find it is difficult to compare MERCURY with prior attacks. As shown in Table I, most works fall into the category of physical side-channel attacks and require to set up physical instruments to obtain the power traces. For remote attacks [8], [17], [18], they only target layer extraction and cannot achieve end-to-end model stealing. Due to the distinct settings and attack goals, we mainly report the attack results of MERCURY.

#### A. Model Extraction Results

**Model performance.** We employ our collected dataset to train both the RNN-CTC and Transformer models. For the RNN-CTC model, we utilize 3 CNN layers with an RNN dimension of 128, and train the model for 120 epochs. For the Transformer model, we incorporate one encoder layer and one decoder layer. The model input and output dimension, denoted as  $d_{model}$ , is configured as 256. We incorporate 8

# of layers	Best Loss		Best OER	
	Train	Test	Train	Test
1	0.02	0.49	0.01	0.15
2	0.04	0.51	0.01	0.14
3	<b>0.04</b>	<b>0.10</b>	<b>0.01</b>	<b>0.02</b>
4	0.09	0.54	0.04	0.15
5	0.13	0.44	0.05	0.12

(a) Number of CNN layers

(b) RNN dimensions

TABLE III: Impact of RNN-CTC configurations

parallel attention layers, with projected queries, keys, and values having dimensions of  $d_k = d_v = d_{model}/h$ , where  $h$  represents the number of attention heads. The positional encoding and optimizer settings adhere to the implementation outlined in [34]. With these configurations, we train the model for 100 epochs.

We run inference for all randomly generated models in NVDLA on both MNIST and CIFAR-10 datasets. For MNIST, using RNN-CTC, we achieve OER of 0.01 and 0.02 on the training and test set, respectively; using Transformer, we get OER of 0.02 and 0.15 on the two sets. For CIFAR-10, with RNN-CTC, we achieve OER of 0.04 and 0.16 on the training and test set, respectively; with Transformer, we get OER of 0.10 and 0.15 on the two sets. We observe that both models exhibit strong performance on the training dataset. RNN-CTC outperforms Transformer on the test set, mainly because its stronger ability to handle variable-length input and output sequences. However, Transformer is more interpretable in analyzing the leakage points in the trace (Section V-C). Below we mainly use RNN-CTC on MNIST for more evaluation.

**Ablation studies.** We further investigate the impact of hyperparameters on the RNN-CTC model. We first consider different numbers of CNN layers used for feature extraction (1 ~ 5). Figure 7 shows the trends of the training and test losses and OER when training the RNN-CTC model. We observe that OER and loss in all figures drop when the training proceeds.

TABLE V: Sensitivity to TDC locations

TDC location	OER (RNN)		OER (Transformer)	
	w/o.	w/.	w/o.	w/.
Original	0.14		0.15	
Center	0.51	0.15	0.59	0.11
Top-right	0.4	0.23	0.49	0.23
Top-left	0.39	0.20	0.43	0.20
Bottom-right	0.25	0.23	0.45	0.26
Bottom-left	0.41	0.19	0.59	0.14

Table IIIa reports the best OER and loss values for different numbers of convolution layers. We find that the RNN-CTC model with 3 convolution layers gives the best performance. We also test the prediction accuracy of the RNN-CTC model with different RNN dimensions. Table IIIb reports the best loss and OER when the RNN dimension is set as 64, 96, 128, 256, and 512 respectively. We observe that RNN with the dimension of 128 has the best performance. We will adopt such optimal hyper-parameters in the following experiments.

### B. Robustness Analysis

**Side-channel noise.** We first evaluate the robustness of the prediction model against side-channel noise. We add different amounts of Gaussian noise to the side-channel trace, and measure to what extent the extraction accuracy will be affected by such noise. The scale of the injected noise is calculated as  $P_n = P_s / (10^{SNR/10})$ , where  $P_s$  is the power of the input sequence, and SNR is the signal-to-noise ratio (SNR), defined as the power ratio of the input to the noise in decibels.

Table IV reports the extraction results under various scales of noise. We observe that RNN-CTC has no accuracy drop when the SNR is above 50dB. The side-channel noise only has a minor effect on the model when the SNR is 40dB. Considering that our input sequence collected from TDC already contains a certain level of noise, the actual threshold of SNR that can preserve the extraction accuracy should be even lower than the tested value. Therefore, our model has high robustness to resist the noise in the measured trace.

**Placement location of TDC.** Next we explore the effect of the TDC locations on the board during model extraction. Prior studies [10], [37] showed that the outputs of the TDC are highly sensitive to its location. Hence, it is essential for the adversary to find the optimal place to implement the TDC. To evaluate this factor, we place the TDC in 5 different locations: top-left, bottom-left, center, top-right, and bottom-right. We set the location constraints using Pblock in Vivado. For each location, we collect 600 TDC traces and then re-evaluate the prediction of both models, as shown in Table V (the “w/o.” columns). We observe that a different location of TDC can indeed degrade the extraction accuracy, as the adversary’s side-channel power traces during training and inference is different.

TABLE IV: Accuracy with different SNRs

SNR(dB)	OER	Loss
No noise	0.05	0.18
50	0.05	0.18
40	0.09	0.35
30	0.20	0.93
25	0.34	1.34
10	0.71	2.83




Fig. 8: Attention weight (brighter color means larger weights)

In the real-world multi-tenant cloud scenario, the adversary may not have the permission to select the location for implementing his TDC, which is allocated by the cloud provider. To bridge this gap and make our attack practical, the adversary can augment his training dataset by placing the TDC in multiple locations and collecting the comprehensive power traces in the offline phase. Then the prediction model will be more general and robust against TDC placement. We train such a model with an augmented dataset which includes additional 10% of TDC trace for each of the above 5 locations. Table V (the “w/.” columns) shows that the prediction errors are significantly reduced, and close to that of the original location.

### C. Leakage Point Localization

We show how to use the Transformer model to localize the leakage point in the power trace. As described in Section IV-C, we compute the attention weight as the indicator of leakage. Given that the input trace is much longer than the output sequence, the attention weight matrix is too long to be easily illustrated. Therefore, we average every 25 attention weights. Figure 8 shows an example of the attention weights for one TDC trace. The model architecture is shown on the left side, where “fc” represents a fully connected layer and “conv\_2\*2-10” represents a convolution layer with the kernel size of  $2 \times 2$  and output feature map dimension of 10. The number at the top represents the input sequence.

From Figure 8, we observe that the highest attention weights normally appear very early in the trace, which indicates that major leakage of model information is from the beginning. Considering that NVDLA sends the layer-related information to the relevant register from ARM to FPGA before computation, we hypothesize that the major leakage may occur during the transmission or storage of this information. This suggests the importance of providing more protections over these particular locations as the countermeasure. We leave such interpretability-based defenses as future work.

## VI. CASE STUDIES

As mentioned in Section I, extracting model architectures can not only compromise model intellectual property, but also facilitate some attacks to deep learning. We present two case studies to demonstrate how MERCURY can enhance the adversarial examples and membership inference attack.

### A. Enhancing Adversarial Examples

A popular security threat to deep learning models is adversarial examples (AEs), which are created by adding human-invisible perturbations to normal samples to mislead the victim

model [38], [39]. Over the years, numerous attack methodologies have been proposed to generate effective AEs, which can be classified into two categories based on the threat model. The first one is *white-box* attacks. The adversary has knowledge of the model parameters, based on which he precisely crafts the adversarial perturbations. Typical methods include FGSM [40], C&W [41], Deepfool [42], PGD [43], etc. The second one is *black-box* attacks, where the adversary does not know any information about the target model. He can leverage the transferability property of AEs [38], which refers to the ability of AEs generated from one model to attack another different model. The adversary can train a shadow model locally, and generate the corresponding AEs using conventional white-box attack techniques. Then these AEs have a high chance to succeed in attacking the target model.

For black-box attacks, the attack success rate, i.e., AE transferability, highly depends on the similarity between the victim model and adversary’s shadow model. Therefore, our model extraction technique provides a new opportunity of improving such similarity, thus the success rate of black-box attacks. In particular, the adversary can apply MERCURY to extract the architecture of the victim model, and then train a shadow model with this architecture. The AEs generated from this model enjoy higher transferability to the victim model compared to the ones from a random shadow model.

Figure 9a validates the effectiveness of MERCURY in enhancing black-box adversarial attacks. The y-axis represents four victim models with different architectures, while x-axis represents the shadow models used for generating AEs with FGSM. The blocks with the same  $x$  and  $y$  indexes denote the success rates with the enhancement of MERCURY, while the rest blocks show the success rates with random architectures. For fair comparisons, we set the same perturbation scale (with  $\epsilon = 0.3$ ) for all cases. It is clear that AEs based on our MERCURY has the highest transferability (the diagonal blocks) for each victim model.

### B. Enhancing Membership Inference Attacks

As the second case study, we show how MERCURY can facilitate the membership inference attack (MIA) [44]. MIA is a type of privacy attack that aims to determine whether a particular data point has been used to train a machine learning model. This attack is of particular concerns in applications where the training data contain sensitive information, such as medical diagnoses, credit scoring, and fraud detection. To launch the MIA, the adversary can train multiple shadow models locally on separate datasets that are representatives of the target dataset. The adversary then queries these shadow models with a data sample and obtains the corresponding predictions. These prediction vectors, along with whether the sample is a member of the training sets as the label, are used to train a membership inference model. With this model, the adversary can infer the membership of any sample based on the prediction results from the target model.

Clearly the MIA accuracy is related to the similarity of the shadow models and target model: shadow models that are

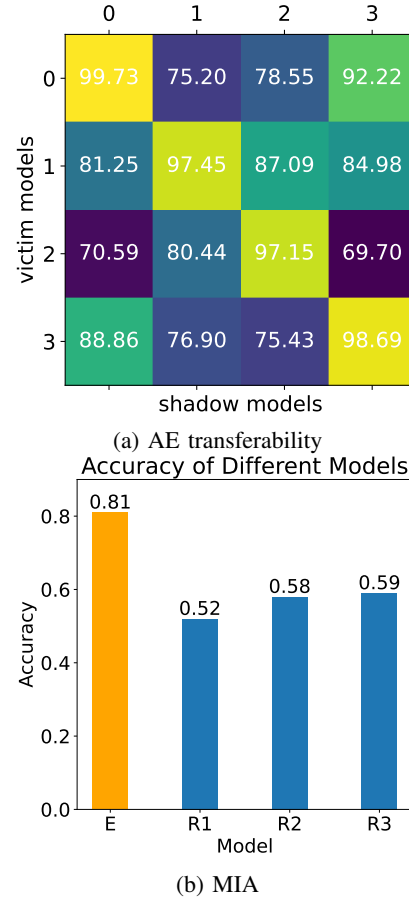


Fig. 9: MERCURY enhances two attacks

identical to the target can better reflect the attributes of the training data. In the black-box setting, our MERCURY can increase such similarity from the model architecture perspective. Figure 9b shows the MIA accuracy of four cases: shadow models with the architecture extracted by MERCURY (yellow bar,  $E$ ), and with random architectures (blue bars,  $R1 - R3$ ). All these attacks are carried out on the MNIST dataset. The target, shadow and attack inference models are trained on the datasets of sizes 2,500, 57,500 and 3,500, respectively. It is obvious that our extracted model yields the highest accuracy.

## VII. RELATED WORKS

**Model extraction attacks on DNN accelerators.** As summarized in Table I, numerous attacks have been proposed to steal the model architecture or weights from the DNN accelerators. For instance, Zhang et al. [17] used ROs to remotely steal the structure of an FPGA neural network. But it did not aim at real-world accelerator designs, and it only did layer tests instead of recovering the entire model. Meyers et al. [18] analyzed the impact of layer folding in accelerators, which makes attacks harder to perform. Subsequently, they showed how to recover the folding before the number of neurons. Gupta et al. [15] launched the EM-based side-channel attack on NVDLA using an oscilloscope. However, this attack requires physical access. Additionally, the attacker needs to manually split a full trace into trunks divided by



different layers to train the model or initiate the attack through inference. Tian et al. [8] proposed to remotely extract the architecture of the versatile tensor accelerator (VTA) using the TDC sensor. They manually distinguished different hyperparameters by observing the distinct shapes in the power trace. The attack is realized in the layer-wise setting. Yoshida et al. [11] introduced an attack to extract the model parameters from the EM leakage. Li et al. [12] initiated differential power analysis (DPA) attacks on 2D DNN accelerators to retrieve weights from the matrix multiplication accelerator. Yu et al. [13] exploited EM side-channels to recover model architecture by simple EM analysis and then recover the model weights by adversarial learning. All the above three attacks require physical access to the victim device, and only target very simple accelerator implementations.

**Other types of attacks on DNN accelerators.** In addition to model extraction, model inversion attacks are also designed to recover the input inference samples. Wei et al. [45] implemented a high-resolution oscilloscope to collect the power traces and recover the label of the input images. Moini et al. [9] designed a remote side-channel attack to recover MNIST images from a binarized neural network (BNN). Different from the above works, Luo et al. [35] demonstrated an integrity threat to the DNN accelerators, where they performed a guided fault injection attack to alter the prediction of the victim model. Those attacks are outside the scope of this paper.

## VIII. CONCLUSION

We propose MERCURY, the first automated remote side-channel attack against the Nvidia Deep Learning Accelerator (NVDLA). MERCURY leverages the readouts of TDC as the power indicator of NVDLA, and trains RNN-CTC and Transformer seq2seq models to predict the model architectures, and disclose the leakage points. Evaluation results demonstrate that MERCURY is able to extract the operation sequence of the victim model within an error rate of 1%. It can also enhance existing black-box adversarial examples and membership inference attacks.

## ACKNOWLEDGEMENT

This research is supported by National Research Foundation, Singapore, and Cyber Security Agency of Singapore under its National Cybersecurity Research & Development Programme (Cyber-Hardware Forensic & Assurance Evaluation R&D Programme <NRF2018NCRNCR009-0001>), and MoE Tier 1 RS02/19. Any opinions, findings and conclusions or recommendations expressed in this paper are those of the authors and do not reflect the view of National Research Foundation, Singapore and Cyber Security Agency of Singapore

## REFERENCES

- [1] Amazon fl web site. [Online]. Available: <https://aws.amazon.com/ec2/instance-types/fl/>
- [2] A. Putnam, A. M. Caulfield, E. S. Chung, D. Chiou, K. Constantinides, J. Demme, H. Esmailzadeh, J. Fowers, G. P. Gopal, J. Gray et al., "A reconfigurable fabric for accelerating large-scale datacenter services," in *ACM/IEEE International Symposium on Computer Architecture*, 2014.
- [3] G. Li, G. Xu, S. Guo, H. Qiu, J. Li, and T. Zhang, "Extracting robust models with uncertain examples," in *The Eleventh International Conference on Learning Representations*, 2022.
- [4] K. Chen, S. Guo, T. Zhang, X. Xie, and Y. Liu, "Stealing deep reinforcement learning models for fun and profit," in *Proceedings of the 2021 ACM Asia Conference on Computer and Communications Security*, 2021, pp. 307–319.
- [5] W. Jiang, H. Li, G. Xu, T. Zhang, and R. Lu, "A comprehensive defense framework against model extraction attacks," *IEEE Transactions on Dependable and Secure Computing*, 2023.
- [6] J. M. Mbongue, A. M.-I. Shuping, P. Bhowmik, and C. Bobda, "Architecture support for fpga multi-tenancy in the cloud," in *IEEE International Conference on Application-specific Systems, Architectures and Processors*, 2020.
- [7] D. R. Gnad, F. Oboril, S. Kiamehr, and M. B. Tahoori, "Analysis of transient voltage fluctuations in fpgas," in *International Conference on Field-Programmable Technology*, 2016.
- [8] S. Tian, S. Moini, A. Wolnikowski, D. Holcomb, R. Tessier, and J. Szefer, "Remote power attacks on the versatile tensor accelerator in multi-tenant fpgas," in *IEEE Annual International Symposium on Field-Programmable Custom Computing Machines*, 2021.
- [9] S. Moini, S. Tian, D. Holcomb, J. Szefer, and R. Tessier, "Remote power side-channel attacks on bnn accelerators in fpgas," in *Design, Automation & Test in Europe Conference & Exhibition*, 2021.
- [10] J. Gravelier, "Remote hardware attacks on connected devices," Ph.D. dissertation, Ecole des Mines de Saint-Etienne, 2021.
- [11] K. Yoshida, T. Kubota, M. Shiozaki, and T. Fujino, "Model-extraction attack against fpga-dnn accelerator utilizing correlation electromagnetic analysis," in *IEEE Annual International Symposium on Field-Programmable Custom Computing Machines*, 2019.
- [12] G. Li, M. Tiwari, and M. Orshansky, "Power-based attacks on spatial dnn accelerators," *arXiv preprint arXiv:2108.12579*, 2021.
- [13] H. Yu, H. Ma, K. Yang, Y. Zhao, and Y. Jin, "Deepem: Deep neural networks model recovery through em side-channel information leakage," in *IEEE International Symposium on Hardware Oriented Security and Trust*, 2020.
- [14] W. Hua, Z. Zhang, and G. E. Suh, "Reverse engineering convolutional neural networks through side-channel information leaks," in *ACM/ESDA/IEEE Design Automation Conference*, 2018.
- [15] N. Gupta, A. Jati, and A. Chattopadhyay, "Ai attacks ai: Recovering neural network architecture from nvla using ai-assisted side channel attack," *Cryptology ePrint Archive*, 2023.
- [16] T. Moreau, T. Chen, L. Vega, J. Roesch, E. Yan, L. Zheng, J. Fromm, Z. Jiang, L. Ceze, C. Guestrin et al., "A hardware–software blueprint for flexible deep learning specialization," *IEEE Micro*, vol. 39, no. 5, pp. 8–16, 2019.
- [17] Y. Zhang, R. Yasaei, H. Chen, Z. Li, and M. A. Al Faruque, "Stealing neural network structure through remote fpga side-channel analysis," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 4377–4388, 2021.
- [18] V. Meyers, D. Gnad, and M. Tahoori, "Reverse engineering neural network folding with remote fpga power analysis," in *2022 IEEE 30th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, 2022, pp. 1–10.
- [19] Y. Umuroglu, N. J. Fraser, G. Gambardella, M. Blott, P. Leong, M. Jahre, and K. Vissers, "Finn: A framework for fast, scalable binarized neural network inference," in *Proceedings of the 2017 ACM/SIGDA international symposium on field-programmable gate arrays*, 2017, pp. 65–74.
- [20] X. Lou, T. Zhang, J. Jiang, and Y. Zhang, "A survey of microarchitectural side-channel vulnerabilities, attacks, and defenses in cryptography," *ACM Computing Surveys (CSUR)*, vol. 54, no. 6, pp. 1–37, 2021.
- [21] X. Lou, S. Guo, J. Li, Y. Wu, and T. Zhang, "Naspy: Automated extraction of automated machine learning models," in *International Conference on Learning Representations*, 2021.
- [22] M. Yan, C. W. Fletcher, and J. Torrellas, "Cache telepathy: Leveraging shared resource attacks to learn {DNN} architectures," in *USENIX Security Symposium*, 2020.
- [23] S. Hong, M. Davinroy, Y. Kaya, D. Dachman-Soled, and T. Dumitraş, "How to Own nas in your spare time," *International Conference on Learning Representations*, 2020.
- [24] X. Hu, L. Liang, S. Li, L. Deng, P. Zuo, Y. Ji, X. Xie, Y. Ding, C. Liu, T. Sherwood et al., "DeepSniffer: A DNN model extraction framework based on learning architectural hints," in *International Conference*

on Architectural Support for Programming Languages and Operating Systems, 2020.

- [25] X. Lou, S. Guo, J. Li, and T. Zhang, "Ownership verification of dnn architectures via hardware cache side channels," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 11, pp. 8078–8093, 2022.
- [26] G. Cesarano, "Fpga implementation of a deep learning inference accelerator for autonomous vehicles."
- [27] M. Zhao and G. E. Suh, "Fpga-based remote power side-channel attacks," in *IEEE Symposium on Security and Privacy*, 2018.
- [28] S. Pant, "Design and analysis of power distribution networks in vlsi circuits." Ph.D. dissertation, 2008.
- [29] F.-X. Standaert, F. Koeune, and W. Schindler, "How to compare profiled side-channel attacks?" in *International Conference on Applied Cryptography and Network Security*, 2009.
- [30] D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen *et al.*, "Deep speech 2: End-to-end speech recognition in english and mandarin," in *International conference on machine learning*, 2016.
- [31] G. Neubig, "Neural machine translation and sequence-to-sequence models: A tutorial," *arXiv preprint arXiv:1703.01619*, 2017.
- [32] M. S. Islam, S. S. S. Mousumi, S. Abujar, and S. A. Hossain, "Sequence-to-sequence bangla sentence generation with lstm recurrent neural networks," *Procedia Computer Science*, 2019.
- [33] K. Palasundram, N. M. Sharef, K. A. Kasmiran, and A. Azman, "Enhancements to the sequence-to-sequence-based natural answer generation models," *IEEE Access*, 2020.
- [34] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [35] Y. Luo, C. Gongye, Y. Fei, and X. Xu, "Deepstrike: Remotely-guided fault injection attacks on dnn accelerator in cloud-fpga," in *ACM/IEEE Design Automation Conference (DAC)*, 2021.
- [36] Hardware architectural specification. [Online]. Available: <http://nvdla.org/hw/v1/hwarch.html>
- [37] S. Moini, X. Li, P. Stanwicks, G. Provelengios, W. Burleson, R. Tessier, and D. Holcomb, "Understanding and comparing the capabilities of on-chip voltage sensors against remote power attacks on fpgas," in *2020 IEEE 63rd International Midwest Symposium on Circuits and Systems (MWSCAS)*.
- [38] 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.
- [39] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv preprint arXiv:1412.6572*, 2014.
- [40] S. Huang, N. Papernot, I. Goodfellow, Y. Duan, and P. Abbeel, "Adversarial attacks on neural network policies," *arXiv preprint arXiv:1702.02284*, 2017.
- [41] N. Carlini and D. Wagner, "Towards evaluating the robustness of neural networks," in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 39–57.
- [42] 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 conference on computer vision and pattern recognition*, 2016, pp. 2574–2582.
- [43] 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.
- [44] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 3–18.
- [45] L. Wei, B. Luo, Y. Li, Y. Liu, and Q. Xu, "I know what you see: Power side-channel attack on convolutional neural network accelerators," in *Annual Computer Security Applications Conference*, 2018.