

Survey of Attacks and Defenses on Edge-Deployed Neural Networks

Mihailo Isakov*, Vijay Gadepally†, Karen M. Gettings†, Michel A. Kinsy*

* Adaptive and Secure Computing Systems (ASCS) Laboratory, Boston, MA

{mihailo, mkinsy}@bu.edu

†MIT Lincoln Laboratory, Boston, MA

{vijayg, karen.gettings}@ll.mit.edu

Abstract—Deep Neural Network (DNN) workloads are quickly moving from datacenters onto edge devices, for latency, privacy, or energy reasons. While datacenter networks can be protected using conventional cybersecurity measures, edge neural networks bring a host of new security challenges. Unlike classic IoT applications, edge neural networks are typically very compute and memory intensive, their execution is data-independent, and they are robust to noise and faults. Neural network models may be very expensive to develop, and can potentially reveal information about the private data they were trained on, requiring special care in distribution. The hidden states and outputs of the network can also be used in reconstructing user inputs, potentially violating users’ privacy. Furthermore, neural networks are vulnerable to adversarial attacks, which may cause misclassifications and violate the integrity of the output. These properties add challenges when securing edge-deployed DNNs, requiring new considerations, threat models, priorities, and approaches in securely and privately deploying DNNs to the edge. In this work, we cover the landscape of attacks on, and defenses, of neural networks deployed in edge devices and provide a taxonomy of attacks and defenses targeting edge DNNs.

Index Terms—neural networks, security, Internet of Things

I. INTRODUCTION

Since the rise of deep learning in the last decade, many different libraries and frameworks for running and training deep neural networks (DNN) have been published and open-sourced. In that time, the landscape of software tools for training neural networks has moved from difficult-to-install libraries [1], and support for static graphs only [2], to industry-ready, easy-to-deploy frameworks [3], high-development efficiency [4], and support for dynamic graphs and *just-in-time* compilation [5]. Recently, as tools have gained maturity, more businesses have started using neural networks in production and exposing services based on neural networks [6]. Deploying neural networks is non-trivial, and the research frameworks proved insufficient to handle high-bandwidth, low-latency inference, leading to the development of production-ready frameworks such as Tensorflow Serving [7] and the standardization of neural network formats with ONNX [8]. With an increasing number of mobile devices and PCs possessing GPUs and custom ASICs, networks have been pushed to smartphones [9] and even GPU-enabled JavaScript [10]. With the rise of voice assistants, wearables, and smart cameras, the need for low-power inference has led to the development of many custom DNN acceleration ASICs [11], [12].

There are many reasons why businesses or users may want to run neural networks on edge devices, as an alternative to sending the data to datacenters for processing, including:

Privacy: users may not want or may not be able to send the data to the cloud for privacy or legal reasons. For example, a hospital may want to process patient data on servers at a different location, but is not willing to risk patient privacy. Even if the patient data is encrypted, if the server is malicious, the patient data may be at risk.

Power: sending data directly to the cloud may not be the most power-efficient approach to run neural networks. For example, in [13], the authors show that in convolutional neural networks (CNN), processing a few of the first convolutional layers before sending the data to the cloud achieves higher power savings compared to processing the whole network on the device or sending the input data to the cloud. As more low-power accelerators using approaches such as quantization [14], stochastic computing [15], or sparsity [16], [17] are released, we expect the ratio between the cost of processing networks and the cost of transmitting input data to become more significant.

Latency: many applications have hard latency requirements and must process a network within a certain time limit. Furthermore, for certain mission-critical applications with hard availability guarantees, as in the case of autonomous drones or self-driving cars, being able to process data on the device is mandatory. While datacenters are able to provide virtually unlimited computing power, possibly making inference time negligible, the transmit time of inputs over the network often cannot be ignored. Hence, a device must possess the required compute power to process the inputs within the time budget.

Throughput: several industries dealing with high bandwidth data are faced with the question of whether to store data for offline processing, allowing thorough analysis at the cost of large amounts of storage, or to process the data in-flight potentially sacrificing some information, but saving on storage. Take an extreme case: the CERN Large Hadron Collider (LHC) can generate upwards of hundreds of terabytes of data per second. Storing that data is difficult, so authors of [18] propose to process the data in-flight using extremely low-latency FPGA designs.

II. TAXONOMY OF ATTACKS ON AND DEFENSES OF DEPLOYED NEURAL NETWORKS

Since DNN accelerators have only recently been deployed in commercial products, the field of attacking and defending these devices is in its infancy. In this section we aim to provide (1) a taxonomy of DNN accelerator attacks and defenses, and (2) a list of plausible attack surfaces and attacker motivations for targeting edge devices running DNNs.

A. Taxonomy of DNN Accelerator Attacks and Defenses

Of the many possible dimensions over which we could characterize attacks and defenses of edge devices, we believe that the *attacker agenda* and *level of access* to the edge device provide a useful classification. The attacker agenda represents the goal of the attacker, and ranges from local denial-of-service (DoS) to gaining full access to a network of edge devices. Level of access is a set of attack surfaces the attacker has access to and ranges from simple API accesses to probing buses or even chip internals. In Figure 1, we present an overview of attacks, defenses, and potential vulnerabilities present in the literature.

B. Attacker Agenda

The y -axis in Figure 1 represents the attacker's motivation for attacking an edge-deployed neural network. We classify attacker motivations into four categories:

Denial of Service: attackers may want to prevent a device running a neural network from properly functioning. For example, attackers may want to prevent smart cameras from properly classifying recordings in order not to raise alarms. Denial of Service (DoS) attacks prevent a device from maintaining availability and completing its function. As feedforward neural networks are data-independent and have fixed latencies, DoS attacks targeting DNNs are only applicable to accelerators running data-dependent models, e.g., recurrent neural networks [19] or neural networks with early exits like BranchyNets [20] or Tree LSTMs [21].

User Privacy Violation: smart devices are increasingly trusted with private user data such as shopping history, voice commands, or medical recordings [22]. This data is valuable for its advertising, monitoring, or polling value. User privacy violations are cases where the attacker is able to access measured or stored sensor data from the device or user data the device from the network. For example, attacks on voice assistants where the attacker can access previous voice commands constitute a local privacy violation.

Model Privacy Violation: the attacker may attempt to exfiltrate a neural network model for a number of reasons: (1) models require significant investment to develop, and, as such, may be stolen and sold, or used in ensembles as a black box [23], (2) finding adversarial examples is significantly easier if the attacker has access to a model (i.e., the white-box scenario), compared to only having access to model inputs and outputs (i.e., the black-box scenario) [24], or (3) the attacker may attempt to learn data from the dataset the model was trained on [25].

Integrity Violation: the attackers may not want to outright prevent the device from functioning, but may want to force the neural network to perform in an unacceptable way. For example, malware may craft adversarial packets in an attempt to fool a network intrusion detection system (IDS) that uses DNNs to identify packets. Local integrity violations are cases where the attacker is able to affect the correctness of a device's neural network inference.

In Figure 1, we list several attacker agendas, sorted by severity. We add two additional categories of general user and integrity violations, which consists of cases that affect not only a single device, but multiple devices, some of which are not under the attacker's physical control. An example of this is data poisoning attacks on federated learning systems [26], where attackers controlling one device can insert backdoors into all devices in the network.

C. Attacker's Level of Access

The x -axis in Figure 1 represents the attacker's level of access to an edge device. The five access categories vary by invasiveness from purely software, API-based attacks and defenses, all the way to invasive attacks, such as decapsulation and microprobing. The API attacks assume that the attacker only has access to the device through conventional channels, e.g. through the network (as in the case of machine translation systems) or the device's sensors (as in the case of voice assistants). Software side-channels additionally assume that the attacker has some ability to measure side-channels through the device's legitimate outputs, e.g., measure the latency of network responses or the amount of traffic the device is sending to the cloud. For both API and software side-channel attacks, the attacker does not need physical contact with the device. Additionally, these attacks are typically simple to automate, unlike the attacks based on physical properties of the device. In the case of physical side-channels, the attacker needs physical proximity to the device, as in the case of power or electromagnetic (EM) analysis. However, physical side-channel attacks do not require invasive sensors or access to the printed circuit board (PCB). PCB probing attacks include attacks that may measure data or timing information of any bus exposed on the PCB, but not the internals of any chip on the device. These attacks include probing RAM or non-volatile memory (NVM), as well as cold-boot attacks, etc. Finally, invasive attacks access the internals of a chip. These include approaches such as decapsulation (a procedure where the chip packaging is removed), microprobing (where the attacker can probe the internals of a chip), chemical attacks that can reveal information stored in read-only memory, and scanning electron microscope (SEM) attacks (which are able to read RAM memory) [27]. These attacks typically require specialized labs and expensive equipment. They are often destructive and may require multiple devices before a successful attack is implemented. For completeness, we also include training-time attacks and defenses that happen before devices are deployed, or during on-line training. Attacks that take place strictly before deployment are beyond the scope of this study.

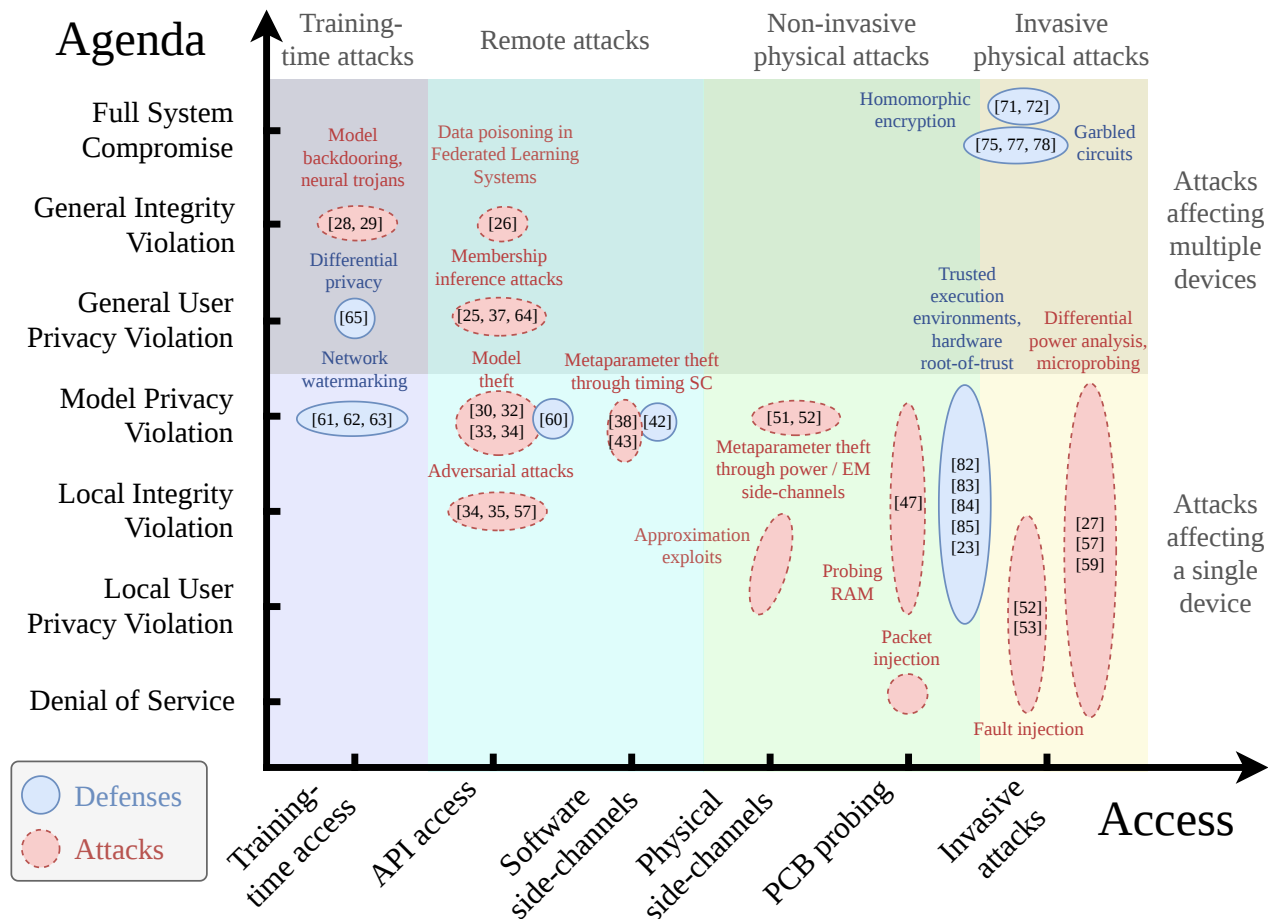


Fig. 1. Taxonomy of attacks, defenses, and potential vulnerabilities of edge devices running ML inference or on-line training.

III. ATTACKS ON DEPLOYED NEURAL NETWORKS

We present a short survey of published attacks on neural network accelerators. We focus primarily on test-time attacks (attacks on already trained models), as we assume that training-time attacks such as data poisoning [28] or Neural Trojans [29] must happen before the model is deployed.

API attacks: API attacks interact with the victim device only through the sensors, the interface, or the network. Here we assume that the attack is independent of the hardware platform running the neural network and does not rely on any side-channel information. The majority of the API attacks present in the literature either attempt to (1) exfiltrate the model or the model metaparameters, (2) find adversarial examples, or (3) infer some property of the model’s training data.

In [30], the authors show how machine learning models hosted behind APIs can be exfiltrated. Here the attacker sends crafted inputs and collects outputs from the model until the attacker is able to reconstruct the model behind the API. In the case of simpler ML models such as decision trees, the models can be perfectly reconstructed. However, for more complex models such as neural networks, the attacker cannot simply solve nonlinear equations to arrive at model weights, but must instead train a ‘student’ network on input-output pairs collected from the API [31]. A similar work [32] shows the simplicity of reverse-engineering black-box neural

network weights, architecture, optimization method and the training/data split. In [33], authors reframe the goal from model theft, to arriving at a ‘knockoff’ model exhibiting the same functionality. In [34], authors ignore model parameters and instead attempt to steal the hyperparameters of a network. Good hyperparameters, while far smaller than models, can be more difficult to arrive at, as they require many experiments and human effort to tune.

In order to violate the integrity of a machine learning model, attackers may attempt to find adversarial examples [35]. While most attacks rely on having access to the white-box model or the output gradient, several works have shown that even black-box networks [32] and networks with obfuscated gradients [36] are not resistant to determined attackers.

Lastly, attackers may attempt to infer some information about the data the neural network was trained on. Attacks which determine whether a specific input was used in training a model are called *membership inference* attacks. Though DNN models are typically smaller than the training dataset, they can nonetheless memorize potentially secret information [25], as in the example of predictive keyboards memorizing PIN codes or passwords. In [37], authors show that even when the model is behind a black-box API, and the adversary has no knowledge of the victim’s training dataset, membership inference attacks are still successful.

Software side-channel attacks: API and software side-channel (SC) attacks target a similar attack surface, but software side-channel attacks can additionally gain information through side-effects such as timing or cache side-channels. Here, the attacker abuses information about the physical device processing the attacker's request to gain an insight into the internal state of the device.

Both timing and cache side-channels typically cannot reveal anything about the data being processed on the device - timing SC reveal information about the compute intensity of a certain task, and cache SC reveal information about recently accessed addresses in the caches. As such, they are commonly employed to extract coarse-grain information such as neural network architecture running on a device. For example, in Cache Telepathy [38], attackers use the Flush+Reload [39] and Prime+Probe [40] cache SC attacks to measure the size of general matrix multiply (GEMM) operations, first counting the number of parameters in the model, and then narrowing down the model architecture. While this attack is restricted to CPUs, GPUs are no less vulnerable to cache-side channels [41]. A similar work [42] is applicable to CPUs, GPUs, and DNN accelerators, and can fingerprint a network after only a single inference operation. It leverages a priori knowledge of major DNN libraries to prime the instruction cache and learn which functions are called during inference.

Timing attacks are also used to reveal model architecture: in [43], the authors assume that the attacker knows the victim's hardware, and is able to buy the same device in order to build timing profiles of different networks. By only knowing the accuracy and the latency of the victim network, the attacker trains many candidate architectures searching for one that has the same signature. This, however, requires the attacker to first steal a part of the training dataset using a membership inference attack [37], which negates much of the need for stealing a model architecture.

Software SC attacks may be less successful in the edge domain compared to the cloud, as edge devices typically serve a single user, while SC are typically used for compromising secure multi-user systems [44]. However, as more networks are pushed to the edge, we can expect multi-network systems with different privileges, goals, and timescales to become increasingly common. An example of this may be predictive keyboards, which perform both inference (text prediction) and NN training on the same device [45].

Another potential vulnerability may be introduced with the adoption of data-dependent inference latency. For example, DARPA's N-ZERO program [46] seeks low-power edge devices that may need to stay dormant for years and have several levels of neural networks, each activating the next one once a certain pattern is sensed. These types of networks are inherently vulnerable to timing attacks, as conventional methods for defending against timing attacks, such as constant time functions negate all the benefits of variable-latency inference.

Physical side-channel attacks: Physical side-channels typically measure some physical quantity, such as power, electromagnetic radiation, vibration, etc. Several works have explored

using physical side-channels to extract the neural network architecture, weights, or user inputs to an edge device.

Memory access patterns can trivially reveal model architecture. In [47], authors attempt to steal a model and model architecture running on a secure enclave such as Intel SGX [48], by observing memory access traces. While traces allow attackers to learn the architecture, the model can only be stolen if the accelerator exploits data-dependent model properties, such as the sparsity of hidden neuron activations [49]. Power and electromagnetic (EM) side-channel attacks are explored in [50], where the authors use EM SCs to learn the activation function, simple power analysis to learn model architecture, and differential power analysis to learn network weights. While simple power analysis does not require invasive measures, differential power analysis may require chip decapsulation, and would need to be classified as an invasive attack. Finally, the authors show how user's private inputs may be extracted using power analysis. A similar attack is explored in [51], where the authors use a power side-channel to observe the processing of the first layer of a convolutional network and extract user's inputs. The authors explore both active and passive attackers, i.e., attackers that can actively input their own images to the accelerator and attackers that can only observe user inputs. Another line of attacks attempts to induce faults in order to cause misclassifications [52], [53] and relies on a microarchitectural or device-level attacks, such as RowHammer [54].

Probing attacks: Probing attacks assume that an attacker is able to access the individual components of the device, e.g., the CPU/GPU/ASIC, the RAM memory, non-volatile storage, or busses, but is not able to perform invasive attacks that access the internals of the chips. The attacker has full access to measure signals on any exposed wires or even drive wires themselves. This opens up a variety of denial-of-service, integrity, and privacy attacks. Additionally, probing attacks assume that no tamper evidence is left after the attack, unlike invasive attacks.

A simple attack the attacker can carry out is model theft - here the attacker probes the memory bus and runs an inference operation while recording the model being loaded onto the chip. This can be prevented by storing only the encrypted model in RAM and NVM, and decrypting the model on-the-fly, if power requirements permit [23]. However, even if the model is encrypted, just knowing the memory access pattern is enough to reveal the model architecture. Each layer and activation will have a different memory bandwidth, and the attacker can monitor these changes along with memory addresses to learn where layers start and end in memory. While oblivious RAM [55] can hide memory addresses, memory access timings are still sufficient to reveal the topology of the model. This forces the defender to either prefetch weights or create fake accesses in order to obfuscate memory access timings [23]. Similarly, network activations may be larger than the available on-chip memory and may be stored in RAM. These activations also need to be encrypted, because even in cases when the device manufacturer is not concerned about privacy, these activations can be used in order to infer the

model weights [56].

The attacker may also attempt to overwrite parts of RAM or feed their own inputs to the chip in order to subvert any software guards, for example in order to generate more input-output pairs used for API model theft [30]. Encrypted RAM may defend against this type of attack, but the device is still susceptible to DoS attacks, where fake accesses are inserted on busses.

Invasive attacks: Invasive attacks assume that the attacker has full control over the chip and is able to bypass any tamper-proof packaging. These attacks include freezing the device in order to extract volatile memory, probing the internals of the chip, ionizing parts of the chip in order to induce faults, feeding non-legitimate voltages and clock frequencies to the chip, etc. Mounting these attacks is typically cost-prohibitive and requires substantial expertise and equipment to execute.

Several works have explored invasive attacks on DNN accelerators, and many of the conventional (non-DNN specific) invasive attacks are still applicable to them. In DeepLaser [57], the authors decapsulate a chip and are able to induce faults by shining a laser on the chip, causing misclassifications by the neural network. This is done by causing bit-flips in the last layer's activation function, where flipping high-order bits of an output neuron's activation will cause the associated category or value to be dominant. Choosing the minimal amount of bit-flips to achieve a desired output has been studied in two works: [52] and [53]. Both these works show that, despite the robustness of neural networks to random perturbations, networks are highly susceptible to targeted bit-flips, in a manner similar to non-targeted adversarial attacks [58].

While we have not been able to find any examples of this, we expect neural network accelerators to be vulnerable to cold boot attacks [59], which may be able to steal unencrypted models stored in volatile memory, or microprobing [27], which may be able to bypass model or user data decryption.

IV. DEFENDING EDGE DEVICES RUNNING NEURAL NETWORKS

We briefly cover proposed defenses for edge devices running neural networks.

API defenses: The majority of API attacks we have mentioned attempt to steal the model or the model architecture, learn which inputs have been used to train the model, or find adversarial examples for the model running on the device. As finding adversarial examples typically involves first stealing the model [32], we focus only on defenses against model exfiltration and membership inference attacks.

In a recent work called Prada [60], the authors succeed in detecting API model-stealing attacks with a 100% detection rate and no false positives. Here, the authors do not attempt to detect if a single query is malicious (as in the case of adversarial attacks), but whether some consecutive set of them is actively trying to steal the model. The authors detect model-stealing queries as they are specifically crafted to extract the maximum amount of information out of the model. However, the authors note that attackers may introduce dummy queries

to maintain a benign query distribution, resulting in slower but more covert model-stealing attacks.

Watermarking is a method for embedding secret information into some system in order to verify the origin of that system at a later date. Watermarking has been proposed as a method of establishing ownership of neural networks [61]–[63]. Here, a watermark is applied to a neural network in such a way that it does not impact the network's accuracy, but can be used to confirm ownership from network outputs. Even if the party responsible for the theft attempts to prune or finetune the network, watermarks can be retained [62].

Defending against membership inference attacks has been explored in several works. In [64], the authors claim that overfitting is the reason why models are vulnerable to membership inference attacks and suggest that differential privacy [65] used during training can protect against these types of attacks. They propose several defenses, similar to those used in defending against adversarial attacks: (1) reducing the number of predicted classes (in the case of classification problems), (2) reducing the amount of information per class by rounding prediction probabilities, (3) increasing entropy of the prediction values and (4) using stronger regularization during training. Similarly, in [37], the authors propose two defenses: dropout [66], where authors show that randomly zeroing out neurons during training partially prevents the attackers from inferring membership, and model stacking, where multiple models are used in an ensemble to make a prediction.

Side-channel defenses: Due to the data-independent behavior of non-recurrent DNNs, all of the software side-channel attacks we have listed attempt to steal the network architecture. We have not been able to find any attacks that succeed at violating privacy of the inputs or the model parameters through software side-channels. In DeepRecon [42], where attackers prime the instruction cache in order to learn function invocations, the authors propose a defense where the defender simultaneously creates decoy function calls to similar neural network layers. These decoy layers should be small enough not to incur a performance penalty. However, this defense does not stop the attacker from using data cache-based side-channels or timing side-channels. Cache Telepathy [38] suggests less aggressive compiler optimizations, cache partitioning [67] or disallowing resource sharing as defenses against cache-based SC. However, these may not be viable solutions without hardware support for secure caches.

While cache-based defenses may help hide some of the accesses, and the defender may go so far as to remove the possibility of an attacker executing code on the same shared resources as the victim, a determined attacker may attempt to probe the memory bus. As neural networks are typically larger than the last-level cache of modern processors, caches will suffer from capacity misses and the network architecture may be exposed to memory probing attacks. In the Trusted Inference Engine (TIE) [23], the device can either create fake memory accesses in times of reduced memory bandwidth or prefetch data, given available on-chip storage. As TIE targets networks with data-independent profiles (i.e., not recurrent

neural networks), the timing of fake or prefetched accesses can be calculated at compile time.

Similar techniques can be applied to counter power and timing side-channels. As long as networks have data-independent behavior, i.e., the accelerator does not attempt to take advantage of zero values [49], or the network computation graph is static [20], [68], power and timing side-channel attacks should not be able to learn information about the network.

Defenses against invasive and semi-invasive attacks: There are two common approaches used when an organization needs to deploy software with privacy or integrity requirements. One option is to not trust the edge hardware, and assume that the hardware can be actively malicious, as in the case of untrusted CPUs/GPUs, possible hardware Trojans, broken hardware defenses [69], etc. There exist several algorithms that allow processing on private data. Homomorphic encryption [70] (HE) for neural networks has been explored in CryptoNets [71], where the authors use HE to run neural networks on encrypted data, without decrypting it at any time during the process. One of the issues with using HE is the performance reduction - inference using HE can be 100-1000 times slower than without HE. Several works have, however, been able to accelerate HE for neural networks. In Gazelle [72], authors leverage HE for linear layers and Yao's Garbled Circuits [73] for offloading calculating nonlinearities to the owner of the private data, as well as an efficient SIMD implementation and a set of homomorphic linear algebra.

While HE is very efficient for linear layers of a network, DNNs typically use nonlinear activations between the layers, requiring many rounds of computationally expensive calculations. An alternative venue for private inference is based on Yao's Garbled Circuits [73] (GC). Here, two parties want to compute the output of a function (a neural network in this case), where one party supplies the network, and the other the inputs to the network. The party that supplies the network typically creates a garbled circuit and uses a procedure such as oblivious transfer [74] (OT) to acquire the second party's inputs without learning those inputs. A naive implementation of neural networks on GC is very inefficient, and several works have presented domain-specific optimizations to them. In DeepSecure [75], authors first prune the network [76], and then convert the network to Verilog for which they can apply logic minimization. In [77], authors present a modified GC that supports free addition and constant-multiplication on a limited integer range, and a significantly cheaper activation function. As a third take on efficient DNNs using GC, XONN [78] attempts to accelerate XNOR-based networks [79] (networks where activations have only values of -1 or 1), as XNOR operations can be processed for free in GC [80]. While GC requires a linear number of rounds w.r.t. the number of network layers, both [75] and [78] are able to perform inference in a fixed amount of rounds.

The question that arises is whether it makes sense to run any of these algorithms on edge devices. In the case of inference, where both the model and user inputs should be kept private, the defender has the choice of sending encrypted inputs to the

cloud or sending the encrypted model to the edge. Since HE is computationally expensive, edge devices may not receive any latency benefits by running the models locally (unless they are not connected to the network at all).

Another option for private edge inference is *hardware root-of-trust* [81]. Here, the defender trusts some type of hardware device, which is built with certain security measures, as in the case of secure enclaves [48], [82], [83] or secure accelerators [23]. These devices are typically built to work in adversarial environments, where the threat model assumes that attacker can tamper with the device, but cannot probe chip internals. For example, using secure enclaves, such as Intel SGX [48], to perform inference can provide privacy and integrity to the user and neural network deployer, but may be very inefficient. In MLCapsule [84], authors develop a machine learning as a service (MLaaS) platform above Trusted Execution Environments (TEE) such as Intel SGX, and formally prove its security. In [85], the authors propose to use Intel SGX as a hardware root-of-trust, but leverage other hardware such as more powerful (but untrusted) CPUs cores and GPUs to perform inference. The authors are able to guarantee both the privacy of the data sent to untrusted devices, as well as the integrity of results received. An alternative venue explores building custom secure neural network accelerators [23]. Here, the design stores obfuscated or encrypted models in off-chip memory, and performs efficient decryption / deobfuscation on the device. The design leverages secure pseudo-random number generators using physical unclonable functions [86] (PUF) as a source of randomness as an alternative to the power-hungry but more secure encryption. The design also provides security against timing attacks by prefetching data or creating fake accesses to RAM memory.

Since the attacker can still probe peripherals, the device must encrypt data in RAM. However, by timing the memory accesses, the attacker can learn the model architecture. Using oblivious RAM (ORAM) does not help, as ORAM only protects the address values and not access times. Additionally, neural network weights are typically stored in ascending order, so knowing the addresses (but not timings) reveals only the complete model size. To prevent the attacker from timing the RAM, the defender, then, must either have a prefetcher and load weights in advance while maintaining a constant bandwidth, or create fake accesses in times when the bandwidth is unused [23], [87].

V. CONCLUSION

In this work, we have presented a survey of attacks and defenses on neural networks. We have created a taxonomy of attacks and defenses with regard to attackers level of access to the hardware, and attacker's agenda. We have described different types of attacks on neural networks, ranging from API-based attacks to invasive attacks such as decapsulation and microprobing. Finally, we gave an overview of the types of defenses of neural networks, with the goal of protecting the privacy of user data, the privacy of deployed neural networks, or the integrity of neural network inference.

REFERENCES

- [1] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, *et al.*, “Caffe: Convolutional architecture for fast feature embedding,” in *Proceedings of the 22nd ACM International Conference on Multimedia*, ser. MM ’14. New York, NY, USA: ACM, 2014, pp. 675–678.
- [2] R. Al-Rfou, G. Alain, A. Almahairi, C. Angermüller, D. Bahdanau, *et al.*, “Theano: A python framework for fast computation of mathematical expressions,” *CoRR*, vol. abs/1605.02688, 2016.
- [3] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, *et al.*, “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015.
- [4] F. Chollet *et al.*, “Keras,” <https://keras.io>, 2015.
- [5] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, *et al.*, “Automatic differentiation in pytorch,” 2017.
- [6] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, *et al.*, “Google’s neural machine translation system: Bridging the gap between human and machine translation,” *CoRR*, vol. abs/1609.08144, 2016.
- [7] “TensorFlow Serving,” <https://www.tensorflow.org/tfx/guide/serving>, 2019, [Online; accessed 18-April-2019].
- [8] “ONNX: Open Neural Network Exchange Format,” <https://onnx.ai/>, 2019, [Online; accessed 18-April-2019].
- [9] “Apple Core ML,” <https://developer.apple.com/documentation/coreml>, 2019, [Online; accessed 18-April-2019].
- [10] D. Smilkov, N. Thorat, Y. Assogba, A. Yuan, N. Kreeger, *et al.*, “Tensorflow.js: Machine learning for the web and beyond,” *CoRR*, vol. abs/1901.05350, 2019.
- [11] “Google Edge TPU,” <https://cloud.google.com/edge-tpu/>, 2019, [Online; accessed 18-April-2019].
- [12] “NVIDIA Jetson Nano,” <https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-nano/>, 2019, [Online; accessed 18-April-2019].
- [13] S. Teerapittayanon, B. McDanel, and H. T. Kung, “Distributed deep neural networks over the cloud, the edge and end devices,” in *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*, June 2017, pp. 328–339.
- [14] M. Courbariaux, Y. Bengio, and J.-P. David, “Binaryconnect: Training deep neural networks with binary weights during propagations,” in *Advances in Neural Information Processing Systems 28*. Curran Associates, Inc., 2015, pp. 3123–3131.
- [15] A. Ren, J. Li, Z. Li, C. Ding, X. Qian, *et al.*, “SC-DCNN: highly-scalable deep convolutional neural network using stochastic computing,” *CoRR*, vol. abs/1611.05939, 2016.
- [16] S. Han, J. Kang, H. Mao, Y. Hu, X. Li, *et al.*, “ESE: efficient speech recognition engine with compressed LSTM on FPGA,” *CoRR*, vol. abs/1612.00694, 2016.
- [17] M. Isakov, A. Ehret, and M. Kinsy, “Closnets: Batchless dnn training with on-chip a priori sparse neural topologies,” 08 2018, pp. 55–554.
- [18] J. Duarte, S. Han, P. Harris, S. Jindariani, E. Kreinar, *et al.*, “Fast inference of deep neural networks in FPGAs for particle physics,” *Journal of Instrumentation*, vol. 13, no. 07, pp. P07 027–P07 027, jul 2018.
- [19] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [20] S. Teerapittayanon, B. McDanel, and H. T. Kung, “Branchynet: Fast inference via early exiting from deep neural networks,” *CoRR*, vol. abs/1709.01686, 2017.
- [21] K. S. Tai, R. Socher, and C. D. Manning, “Improved semantic representations from tree-structured long short-term memory networks,” *CoRR*, vol. abs/1503.00075, 2015.
- [22] P. Rajpurkar, J. Irvin, K. Zhu, B. Yang, H. Mehta, *et al.*, “Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning,” *CoRR*, vol. abs/1711.05225, 2017.
- [23] M. Isakov, L. Bu, H. Cheng, and M. Kinsy, “Preventing neural network model exfiltration in machine learning hardware accelerators,” 12 2018, pp. 62–67.
- [24] N. Papernot, P. D. McDaniel, and I. J. Goodfellow, “Transferability in machine learning: from phenomena to black-box attacks using adversarial samples,” *CoRR*, vol. abs/1605.07277, 2016.
- [25] N. Carlini, C. Liu, J. Kos, Ú. Erlingsson, and D. Song, “The secret sharer: Measuring unintended neural network memorization & extracting secrets,” *CoRR*, vol. abs/1802.08232, 2018.
- [26] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, “How to backdoor federated learning,” *CoRR*, vol. abs/1807.00459, 2018.
- [27] A. Triandopoulos and H. Choukri, *Invasive Attacks*. Boston, MA: Springer US, 2011, pp. 623–629.
- [28] B. Biggio, B. Nelson, and P. Laskov, “Poisoning attacks against support vector machines,” in *ICML*, 2012.
- [29] Y. Liu, S. Ma, Y. Aafer, W.-C. Lee, J. Zhai, *et al.*, “Trojaning attack on neural networks,” in *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-22, 2018*. The Internet Society, 2018.
- [30] F. Tramèr, F. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, “Stealing machine learning models via prediction apis,” *CoRR*, vol. abs/1609.02943, 2016.
- [31] G. Hinton, O. Vinyals, and J. Dean, “Distilling the Knowledge in a Neural Network,” *arXiv e-prints*, p. arXiv:1503.02531, Mar 2015.
- [32] S. J. Oh, M. Augustin, B. Schiele, and M. Fritz, “Towards Reverse-Engineering Black-Box Neural Networks,” *arXiv e-prints*, p. arXiv:1711.01768, Nov 2017.
- [33] T. Orekondy, B. Schiele, and M. Fritz, “Knockoff Nets: Stealing Functionality of Black-Box Models,” Tech. Rep.
- [34] B. Wang and N. Z. Gong, “Stealing hyperparameters in machine learning,” *CoRR*, vol. abs/1802.05351, 2018.
- [35] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and Harnessing Adversarial Examples,” *arXiv e-prints*, p. arXiv:1412.6572, Dec 2014.
- [36] A. Athalye, N. Carlini, and D. A. Wagner, “Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples,” *CoRR*, vol. abs/1802.00420, 2018.
- [37] A. Salem, Y. Zhang, M. Humbert, P. Berrang, M. Fritz, *et al.*, “ML-Leaks: Model and Data Independent Membership Inference Attacks and Defenses on Machine Learning Models,” jun 2018.
- [38] M. Yan, C. Fletcher, and J. Torrellas, “Cache Telepathy: Leveraging Shared Resource Attacks to Learn DNN Architectures,” aug 2018.
- [39] Y. Yarom and K. Falkner, “Flush+reload: A high resolution, low noise, l3 cache side-channel attack,” in *23rd USENIX Security Symposium (USENIX Security 14)*. San Diego, CA: USENIX Association, 2014, pp. 719–732.
- [40] F. Liu, Y. Yarom, Q. Ge, G. Heiser, and R. B. Lee, “Last-level cache side-channel attacks are practical,” in *Proceedings of the 2015 IEEE Symposium on Security and Privacy*, ser. SP ’15. Washington, DC, USA: IEEE Computer Society, 2015, pp. 605–622.
- [41] H. Naghibijouybari, A. Neupane, Z. Qian, and N. Abu-Ghazaleh, “Rendered Insecure: GPU Side Channel Attacks are Practical,” *CCS*, vol. 15, 2018.
- [42] S. Hong, M. Davinroy, Y. Kaya, S. N. Locke, I. Rackow, *et al.*, “Security analysis of deep neural networks operating in the presence of cache side-channel attacks,” *CoRR*, vol. abs/1810.03487, 2018.
- [43] V. Duddu, D. Samanta, D. V. Rao, and V. E. Balas, “Stealing neural networks via timing side channels,” *CoRR*, vol. abs/1812.11720, 2018.
- [44] Y. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, “Cross-vm side channels and their use to extract private keys,” in *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, ser. CCS ’12. New York, NY, USA: ACM, 2012, pp. 305–316.
- [45] H. B. McMahan, E. Moore, D. Ramage, and B. A. y Arcas, “Federated learning of deep networks using model averaging,” *CoRR*, vol. abs/1602.05629, 2016.
- [46] “Dormant, Yet Always-Alert Sensor Awakes Only in the Presence of a Signal of Interest,” <https://www.darpa.mil/news-events/2017-09-11>, 2019, [Online; accessed 12-May-2019].
- [47] W. Hua, Z. Zhang, and G. E. Suh, “Reverse Engineering Convolutional Neural Networks Through Side-channel Information Leaks,” 2018.
- [48] V. Costan and S. Devadas, “Intel sgx explained,” *IACR Cryptology ePrint Archive*, vol. 2016, p. 86, 2016.
- [49] J. Albericio, P. Judd, T. Hetherington, T. Aamodt, N. E. Jerger, *et al.*, “Cnvlutin: Ineffectual-neuron-free deep neural network computing,” in *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*, June 2016, pp. 1–13.
- [50] L. Batina, S. Bhasin, D. Jap, and S. Picek, “CSI Neural Network: Using Side-channels to Recover Your Artificial Neural Network Information,” Tech. Rep.
- [51] L. Wei, Y. Liu, B. Luo, Y. Li, and Q. Xu, “I know what you see: Power side-channel attack on convolutional neural network accelerators,” *CoRR*, vol. abs/1803.05847, 2018.
- [52] Y. Liu, L. Wei, B. Luo, and Q. Xu, *Fault Injection Attack on Deep Neural Network*.
- [53] A. S. Rakin, Z. He, and D. Fan, “Bit-Flip Attack: Crushing Neural Network with Progressive Bit Search,” Tech. Rep.

- [54] D. Gruss, M. Lipp, M. Schwarz, D. Genkin, J. Juffinger, *et al.*, “Another Flip in the Wall of Rowhammer Defenses,” oct 2017.
- [55] E. Stefanov, M. van Dijk, E. Shi, C. Fletcher, L. Ren, *et al.*, “Path oram: An extremely simple oblivious ram protocol,” in *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security*, ser. CCS ’13. New York, NY, USA: ACM, 2013, pp. 299–310.
- [56] S. Milli, L. Schmidt, A. D. Dragan, and M. Hardt, “Model Reconstruction from Model Explanations,” Tech. Rep.
- [57] J. Breier, X. Hou, D. Jap, L. Ma, S. Bhasin, *et al.*, “DeepLaser: Practical Fault Attack on Deep Neural Networks,” Tech. Rep.
- [58] X. Yuan, Q. Zhu, and X. Li, “Adversarial Examples: Attacks and Defenses for Deep Learning,” Tech. Rep.
- [59] J. A. Halderman, S. D. Schoen, N. Heninger, W. Clarkson, W. Paul, *et al.*, “Lest we remember: Cold-boot attacks on encryption keys,” *Commun. ACM*, vol. 52, no. 5, pp. 91–98, May 2009.
- [60] M. Juuti, S. Szyller, S. Marchal, and N. Asokan, “PRADA: Protecting Against DNN Model Stealing Attacks,” Tech. Rep.
- [61] Y. Adi, C. Baum, M. Cisse Google, B. Pinkas, and J. Keshet, “Turning Your Weakness Into a Strength: Watermarking Deep Neural Networks by Backdooring,” Tech. Rep., 2018.
- [62] Y. Uchida, Y. Nagai, and S. Sakazawa, “Embedding Watermarks into Deep Neural Networks,” Tech. Rep.
- [63] B. D. Rouhani, H. Chen, and F. Koushanfar, “Deepsigns: A generic watermarking framework for IP protection of deep learning models,” *CoRR*, vol. abs/1804.00750, 2018.
- [64] R. Shokri, M. Stronati, and V. Shmatikov, “Membership Inference Attacks Against Machine Learning Models,” 2017.
- [65] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, *et al.*, “Deep Learning with Differential Privacy,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security - CCS’16*. New York, New York, USA: ACM Press, 2016, pp. 308–318.
- [66] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: A simple way to prevent neural networks from overfitting,” *Journal of Machine Learning Research*, vol. 15, pp. 1929–1958, 2014.
- [67] F. Liu, Q. Ge, Y. Yarom, F. Mckeen, C. Rozas, *et al.*, “Catalyst: Defeating last-level cache side channel attacks in cloud computing,” in *2016 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, March 2016, pp. 406–418.
- [68] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to sequence learning with neural networks,” *CoRR*, vol. abs/1409.3215, 2014.
- [69] F. Brasser, U. Müller, A. Dmitrienko, K. Kostiaainen, S. Capkun, *et al.*, “Software grand exposure: SGX cache attacks are practical,” in *11th USENIX Workshop on Offensive Technologies (WOOT 17)*. Vancouver, BC: USENIX Association, 2017.
- [70] R. L. Rivest, L. Adleman, and M. L. Dertouzos, “On data banks and privacy homomorphisms,” *Foundations of Secure Computation, Academia Press*, pp. 169–179, 1978.
- [71] N. Dowlin, R. Gilad-Bachrach, K. Laine, K. Lauter, M. Naehrig, *et al.*, “Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy,” Tech. Rep., February 2016.
- [72] C. Juvekar, V. Vaikuntanathan, and A. Chandrakasan, “GAZELLE: A Low Latency Framework for Secure Neural Network Inference,” Tech. Rep.
- [73] O. Goldreich, “Cryptography and cryptographic protocols,” *Distrib. Comput.*, vol. 16, no. 2-3, pp. 177–199, Sept. 2003.
- [74] S. Even, O. Goldreich, and A. Lempel, “A randomized protocol for signing contracts,” *Commun. ACM*, vol. 28, no. 6, pp. 637–647, June 1985.
- [75] B. Darvish, R. M. Sadegh, R. Farinaz, K. Uc, S. Diego, *et al.*, “DeepSecure: Scalable Provably-Secure Deep Learning,” 2018.
- [76] S. Han, H. Mao, and W. J. Dally, “Deep Compression - Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding,” *ICLR*, pp. 1–13, 2016.
- [77] M. Ball, B. Carmer, T. Malkin, M. Rosulek, and N. Schimanski, “Garbled Neural Networks are Practical,” Tech. Rep., 2019.
- [78] M. S. Riazi, M. Samragh, H. Chen, K. Laine, K. E. Lauter, *et al.*, “XONN: xnor-based oblivious deep neural network inference,” *CoRR*, vol. abs/1902.07342, 2019.
- [79] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi, “Xnor-net: Imagenet classification using binary convolutional neural networks,” *CoRR*, vol. abs/1603.05279, 2016.
- [80] V. Kolesnikov and T. Schneider, “Improved garbled circuit: Free xor gates and applications,” in *Automata, Languages and Programming*, L. Aceto, I. Damgård, L. A. Goldberg, M. M. Halldórsson, A. Ingólfssdóttir, *et al.*, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 486–498.
- [81] M. Tehranipoor and C. Wang, *Introduction to Hardware Security and Trust*. Springer Publishing Company, Incorporated, 2011.
- [82] G. E. Suh, D. Clarke, B. Gassend, M. van Dijk, and S. Devadas, “Aegis: Architecture for tamper-evident and tamper-resistant processing,” in *ACM International Conference on Supercomputing 25th Anniversary Volume*. New York, NY, USA: ACM, 2014, pp. 357–368.
- [83] D. Lee, D. Kohlbrenner, S. Shinde, D. Song, and K. Asanovic, “Keystone: A framework for architecting tees,” *CoRR*, vol. abs/1907.10119, 2019.
- [84] L. Hanzlik, Y. Zhang, K. Grosse, A. Salem, M. Augustin, *et al.*, “Mlcapsule: Guarded offline deployment of machine learning as a service,” *CoRR*, vol. abs/1808.00590, 2018.
- [85] F. Tramer and D. Boneh, “Slalom: Fast, Verifiable and Private Execution of Neural Networks in Trusted Hardware,” 2018.
- [86] G. E. Suh and S. Devadas, “Physical unclonable functions for device authentication and secret key generation,” in *2007 44th ACM/IEEE Design Automation Conference*, June 2007, pp. 9–14.
- [87] C. W. Fletcher, M. v. Dijk, and S. Devadas, “A secure processor architecture for encrypted computation on untrusted programs,” in *Proceedings of the Seventh ACM Workshop on Scalable Trusted Computing*, ser. STC ’12. New York, NY, USA: ACM, 2012, pp. 3–8.