



HAL
open science

Evaluation of parameter-based attacks against embedded neural networks with laser injection

Mathieu Dumont, Kévin Hector, Pierre-Alain Moellic, Jean-Max Dutertre,
Simon Pontie

► **To cite this version:**

Mathieu Dumont, Kévin Hector, Pierre-Alain Moellic, Jean-Max Dutertre, Simon Pontie. Evaluation of parameter-based attacks against embedded neural networks with laser injection. SAFECOMP 2023 - 21st Smart Card Research and Advanced Application Conference, Université de Toulouse, Sep 2023, Toulouse, France. pp.259-272, 10.1007/978-3-031-40923-3_19 . cea-04607994

HAL Id: cea-04607994

<https://cea.hal.science/cea-04607994>

Submitted on 11 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Evaluation of Parameter-based Attacks against Embedded Neural Networks with Laser Injection

Mathieu Dumont^{1,2}, Kevin Hector^{1,2}, Pierre-Alain Moëllic^{1,2}, Jean-Max Dutertre³, and Simon Pontié^{1,2}

¹ CEA Tech, Centre CMP, Equipe Commune CEA Tech - Mines Saint-Etienne, F-13541 Gardanne, France

{name.surname}@cea.fr

² Univ. Grenoble Alpes, CEA, Leti, F-38000 Grenoble, France

³ Mines Saint-Etienne, CEA, Leti, Centre CMP, F-13541 Gardanne France
dutertre@emse.fr

Abstract. Upcoming certification actions related to the security of machine learning (ML) based systems raise major evaluation challenges that are amplified by the large-scale deployment of models in many hardware platforms. Until recently, most of research works focused on API-based attacks that consider a ML model as a pure algorithmic abstraction, but new implementation-based threats have been revealed, emphasizing the urgency to propose both practical *and* simulation-based methods to properly evaluate the robustness of models. A major concern is parameter-based attacks (such as the Bit-Flip Attack – BFA) that highlight the lack of robustness of typical deep neural network models when confronted by accurate and optimal alterations of their internal parameters stored in memory. Setting in a security testing purpose, this work practically reports, for the first time, a successful variant of the BFA on a 32-bit Cortex-M microcontroller using laser fault injection. It is a standard fault injection means for security evaluation, that enables to inject spatially and temporally accurate faults. To avoid unrealistic brute-force strategies, we show how simulations help selecting the most sensitive set of bits from the parameters taking into account the laser fault model.

Keywords: Hardware Security, Fault Injection, Evaluation and certification, Machine Learning, Neural Network

1 Introduction

The massive deployment of Machine Learning (ML) models in a large spectrum of domains raises several security concerns related to their integrity, confidentiality and availability. For now, most of the research efforts are essentially focused on models seen as abstractions, i.e. the attack surface is focused on so-called *API-based attacks*, excluding threats related to their implementation in devices that may be physically accessible by an adversary, as it the case for embedded

ML models. The flaws intrinsically related to the models and the *physical* ones may be used jointly to attack an embedded model or exploit data leakages.

A typical API-based attack and one of the major threat against models are the *adversarial examples* [1] that aims at fooling the prediction of a model with input-based alteration at inference time. However, recent works demonstrated that physical attacks are realist threats by targeting critical elements of a model such as activation functions [2] or its parameters (parameter-based attacks) [3] as studied in this work.

In such a security context, with the demonstration of worrying attack vectors, there is an urgent need of certification for AI-based systems more especially for *critical* ones (see the European AI Act [4]). Therefore, alongside the demonstration of new attacks and the development of defenses, an important challenge relies on the availability of proper robustness evaluation and accurate characterization methods in addition to both simulation and experimental tools and protocols. Model robustness evaluation is one of the most important challenge of modern artificial intelligence and several remarkable works from the Adversarial Machine Learning community have already raised major issues for adversarial examples [5], with many defenses relying on weak evaluations. Dealing with embedded neural network models and weight-based adversarial attacks, this challenge encompasses both *safety* and *security* concerns: parameters stored in memory may be altered by random faults because of hostile environments or strong energy consumption limitations [6] and also be the target of an adversary that aims at optimally threatening a model.

2 Related works and objectives

2.1 Parameter-based attacks

Implementation-based threats

An important part of the ML security literature concerns *algorithmic* or so-called *API-based attacks* that exploit input/output pairs and additional knowledge from the model in case of white-box attacks. A large body of work shows that these threats concern every stage of the ML pipeline [7], both at training and inference steps, and threaten the confidentiality (model and data), integrity and availability of the models. However, these attacks do not consider that the adversary may have a direct interaction with the algorithm as it can be the case for an embedded AI system. *Implementation-based attacks* precisely exploit software or hardware flaws as well as the specific features of the device. For example, *side-channel analysis* [8] have been demonstrated for model extraction as an efficient way to extract information from the model architecture or the values of the parameters [9]. Alongside safety-related efforts that evaluate the robustness of ML models against random faults, some works demonstrate that models, and more especially deep neural networks, are highly sensitive to *fault injection analysis* [10] that alter the data, the parameters as well as the instructions flow [2, 11].

Weight-based adversarial attacks

New attack vectors have been highlighted and more essentially parameter-based attacks (also named weight-based adversarial attacks). Let’s consider a supervised neural network model $M_W(x)$, with parameters W (also referred as *weights*), trained to optimally map an input space $\mathcal{X} = \mathbb{R}^d$ (e.g., images) to a set of labels \mathcal{Y} . M is trained by minimizing a loss function $\mathcal{L}(M_W(x), y)$ (typically the cross entropy for classification task) that quantifies the error between the prediction $\hat{y} = M_W(x)$ and the correct label y . As formalized in [3] or [6] with Eq. 1, a parameter-based attack aims at maximizing the loss (i.e., increase mispredictions) on a small set of N test inputs. As for the imperceptibility criterion of adversarial examples, the attacker may add a constrain over the perturbation by bounding the bit-level Hamming distance (HD) between the initial (W) and faulted parameters (W'), corresponding to an *adversarial budget* S .

$$\max_{W'} \underbrace{\sum_{i=0}^{N-1} \mathcal{L}(M(x_i; W'), y_i)}_{\text{mispredictions}} \text{ s.t. } \overbrace{HD(W', W) \leq S}^{\text{adv budget}} \quad (1)$$

A state-of-the-art parameter-based attack is the Bit-Flip Attack (hereafter BFA) [3] that aims to decrease the performance of a model by selecting the most sensitive bits of the stored parameters and progressively flip these bits until reaching an adversarial goal (in [3] or [12], the objective is to ultimately degrade the model to a random-guess level). The selection of the bits is based on the ranking of the gradients of the loss w.r.t. to each bit $\nabla_b \mathcal{L}$, computed thanks to a small set of inputs. First, each selected bit is flipped (and restored) to measure the impact on the model accuracy. Then, the most sensitive bit is permanently flipped according to the gradient ascendant as defined in [3].

Adversarial goals

Parameter-based attacks are not limited to the alteration of the target model integrity. BFA has been recently demonstrated for powerful model extraction in [13] with an Intel i5 CPU platform: RowHammer is used to perform a BFA on the parameters of a model stored in DRAM (DDR3). The threat model in [13] follows a typical model extraction setting: the adversary knows the model’s architecture but not the internal parameters and has only access to a limited portion of the training dataset ($< 10\%$). His goal is to build a *substitute* model as close as possible as the target model (performance theft). Interestingly, this joint use of RowHammer and BFA is performed in a side-channel analysis fashion: the observation of the induced faults enables to make assumptions on the value of some bits of the parameters. Then, the knowledge of these bits enables to efficiently trained a substitute model, by constraining the value range of the parameters, with high *fidelity* compared to the target model.

2.2 Scope and objectives

Because parameter-based attacks are the basis of both powerful integrity and confidentiality threats, their practical evaluation on the different platforms where fault injection may occur is becoming a critical need for present and future standardization and certification actions of critical AI-systems. We position our work on a different set of platforms than the main works related to the BFA (CPU platforms with DRAM), that is MCU platforms (Cortex-M with Flash memory), yet a very important family of embedded AI systems regarding the massive deployment of ML models on MCU-based devices for a large variety of domains. Our main positioning is as follows:

- We set this work in a security evaluation and characterization context. Therefore we do not position ourselves through an *adversary* but an *evaluator* point of view.
- Our scope is parameter-based threats for neural network embedded in 32-bit microcontroller (hereafter MCUs).
- For that purpose, we use Laser Fault Injection (hereafter, LFI) as an advanced and very spatially and temporally accurate injection means, a reference technique that is used in many security evaluation centers.

State-of-the-art is focused on simulation-based evaluations or on RowHammer attacks (i.e., exclusively DRAM platforms that excludes MCU). To the best of our knowledge, this work is the first to demonstrate the practicability and suitability of the characterization of a weight-based adversarial perturbation against Cortex-M MCU thanks to LFI.

2.3 Related works

Since the presentation of the BFA in [3], several works analyzed the intrinsic mechanisms of the attack as well as potential protections [12,14] and evaluated its properties according to the threat model, training parameters and model architecture [15]. The standard BFA is *untargeted* since the induced misclassifications are not chosen by the adversary. Therefore, some works also proposed *targeted* versions [16] with specific target inputs and/or labels. Other recent works propose alternative methods to efficiently select the most sensitive parameters to attack [6,17]. For our work, we use and adapt (for LFI) the standard BFA of [3] as it is the state-of-the-art baseline for weight-based adversarial attacks.

To the best of our knowledge, the only work related to laser injections for MCU-based platform against embedded neural networks has been proposed in [2]. Our work differs by the target device and the elements we target in the model. In [2], the authors used an 8-bit microcontroller (**ATMega328P**) and a 32-bit neural network implemented in C (a multilayer perceptron trained on the MNIST data set). They only focused on the activation functions by inducing instruction skips (i.e., the faulted instructions are not executed, as if they were skipped). They used a laser and only targeted the last hidden layer. We used a Cortex-M 32-bit microcontroller and embedded 8-bit quantized neural

network thanks to a state-of-the-art open source library (*Neural Network on Microcontrollers*¹, hereafter NNoM [18]) for model deployment. Our attack vector (BFA-like attack) and fault model enable to evaluate the robustness of a model against an advanced adversary that aims at significantly altering the performance of a model with a very limited number of faults or extract information about parameter values for model extraction.

Thus, our work is also closely linked to [11] that demonstrated at USENIX’20 a complete exploitation of the BFA with RowHammer on an Intel i7-3770 CPU platform. Although we target a different type of platforms with 32-bit MCU and another fault injection means (LFI), we share the same objective to go beyond simulations and propose a complete practical evaluation of a parameter-based attacks against embedded quantized DNN models.

3 Evaluator assumptions and experimental setups

3.1 Goal and evaluator assumptions

Objectives. The main objective of the evaluator is to evaluate the robustness of a model against precise fault injections by significantly decrease the average accuracy on a labelled test set. More precisely, the ion scenario corresponds to a generic *untargeted* case (i.e., the incorrect labels are not controlled by the evaluator). Note that a targeted scenario (for specific test samples or target labels) is possible [19] but out of the scope of our experiments. A secondary objective is to minimize the evaluation cost with a strategy that reduces the number of faults to be injected (i.e., avoid an exhaustive search that may be unrealistic according to the complexity of the target model).

Evaluator hypothesis. Classically for security testing, the evaluator simulates a worst-case adversary that has a perfect knowledge of the model (white-box attack) and is able to query the model without limitation. The evaluator has a full access to the device (or clones of the device) and can perform elementary characterizations to adapt and optimize the fault injection set-up.

3.2 Single bit-set fault model on Flash memory

We consider an accurate fault model relevant for laser injections previously explained and demonstrated for NOR-Flash memory of Cortex-M MCU by Colom-bier and Menu [20, 21]: the bit-set fault model. As its name suggests, it consists in setting a targeted bit to a logical 1. When the bit was already at 1 the fault has no effect. When targeting a Flash memory at read time with a laser pulse, the induced bit-set is transient: it affects the data being read at that time while the stored value is left unmodified. Authors from [20] explained the underlying mechanism of the bit-set fault injection with the creation of a *photoelectric current* induced by the laser in a floating-gate (FG) transistor that flows from its

¹ <https://github.com/majianjia/nnom/>

drain to the ground. This current is added to the legitimate one so that the total current is above the reference that makes the bit read as a logical 1.

LFI is a local fault injection means, its effect is restricted to the bit line connected to the FG transistor inside the laser spot area. More precisely several current components are induced in the affected transistors, depending on the laser spot diameter: up to two adjacent bits can be faulted simultaneously [21].

3.3 Target and laser bench setup

Device under test (DUT). Our target board embeds an ARM Cortex-M3 running at 8 MHz. It includes 128 kB of Flash memory and is manufactured in the 90 nm CMOS technology node. The dimension of the chip is 3 x 2.5 mm. Since LFI requires the surface of the die being visible, the microcontroller packaging was milled away with engraving tools to provide an access to its laser-sensitive parts. The chip was then mounted into a test board compatible with the Chip-Whisperer CW308 platform².

Laser platform. Our laser fault analysis platform integrates two independent laser spots with a near infrared (IR) wavelength of 1,064 nm, focused through the same lens. Each laser spot has a diameter ranging from 1.5 to 15 μm depending on the lens magnification. Both spots can move inside the whole field of view of the lens with minimum distortion. The laser source can reach a maximum given power of 1,700 mW. The delay between the trigger and the laser shot can be adjusted with a step of a few nanoseconds. An infrared camera is used to observe the laser spot location on the target and a *XY* stage enables to move the objective above the entire DUT surface.

3.4 Datasets and models

Although theoretical and simulation-based works exclusively used complex deep neural networks trained for vision tasks (e.g., ResNet on ImageNet), it does not represent a large part of real-world applications that take benefit from classical fully-connected architecture (or multilayer perceptron, hereafter MLP) that fit and perform well on the overwhelmingly widespread constrained platforms we studied in this paper.

We considered two classical ML data sets for our fault injection experiments: The **IRIS dataset** consists of 150 samples, each containing 4 real-value inputs (width and length of the petal and sepal of iris flowers) and labelled according to 3 classes (3 different iris species). A shallow fully-connected model with few neurons and one layer is sufficient for an efficient classification of this dataset. We trained two models both with one hidden layer made of one and four neurons respectively. These simple models provide an easy insight on the bit-set fault model effect before characterizing deeper models.

The **MNIST dataset** is composed of gray-scale handwritten digits images (28x28 pixels) from 0 to 9. We reduced the complexity of our model by compressing the input data with a classical principal component analysis (PCA).

² https://rtfm.newae.com/Targets/CW308_UF0/

Our model is a MLP with a single intermediate layer of 10 neurons and ReLU as activation function. The resulting model has 620 trainable parameters (including bias). After training, the model reaches 92% of accuracy on the test set.

3.5 Model implementation on MCU

These models were trained with TensorFlow. Few tools are available to embed previously trained models in microcontroller boards such as TensorFlow Lite, X-Cube-AI or NNoM. We chose NNoM as it is an efficient and convenient open source platform that fits our security testing objectives. NNoM offers 8-bit model quantization (weights and activations with a standard uniform symmetric powers-of-two quantization scheme) and a complete white-box access to the inference code that enables to draw a timing profile of the sensitive operations we target (reading values from the Flash memory).

<pre> 1 while (rowCnt){ 2 //pA : address, stored input value 3 //pB : address, stored weight value 4 for (int j = 0; j < dim_vec; j++) { //loop on all neuron parameters 5 q7_t inA = *pA++; //load input to inA, address increment 6 q7_t inB = *pB++; //load weight to inB, address increment 7 ip_out += inA * inB; //neuron weighted sum 8 } 9 *p0++ = (q7_t)__NNOM_SSAT((ip_out >> out_shift), 8); //store 10 output rowCnt--;} </pre>	<pre> 1 ;q7_t inB = *pB++ ;Weight n+1 initialization 2 ldr r3, [r7, #80] ;Loading the address of the weight n 3 adds r2, r3, #1 ;Next weight address 4 str r2, [r7, #80] ;Input value loading into r2 reg 5 ldrsb.w r3, [r3] ;Weight value loading. LASER SHOT 6 strb r3, [r7,#23] ;Store of the weight in SRAM reg </pre>
---	---

(a) C code of the weighted-sum computation in a fully-connected layer.

(b) Assembler code of line 6 of listing 1a. Our target is the load instruction, line 5.

Fig. 1: C and Assembler codes from the NNoM inference.

Listing 1a is an extract of the C code source of the core calculation of an inference from NNoM, that is the weighted sum between the inputs (i.e. the input data or the outputs of the previous layer) and the model parameters before the non-linear activation is applied. It consists in loading the neuron input and weight values (`inA` and `inB`, line 5 and 6), then process the multiplication and accumulation in an intermediate output value (`ip_out`, line 7). Line 9 corresponds to the NNoM quantization procedure. The assembler code in listing 1b corresponds to the weight initialization of Listing 1a, line 6. The weight value, stored into the Flash memory, is loaded into register `r3` (line 5 in red). Based on our single bit-set fault model, if a laser beam is applied during the execution of the load instruction, a bit-set could be induced directly on the loaded value.

To characterize the impact of the laser pulse, we synchronize the laser thanks to a trigger signal in the C code before line 6 and monitor the parameter value before and after the trigger using UART communication.

4 A parameter-based attack with LFI

4.1 Initial characterization on a neuron

To analyze the efficiency and practicality of LFI we first demonstrated the accuracy of the induced faults on a single neuron composed of four inputs. Then, we scaled up to functional models trained on IRIS and MNIST to analyze the impact at a model-level.

An important preliminary experience is to set up our laser bench on our DUT. For conciseness purpose, we do not detail all the Flash memory mapping procedure neither the selection of the laser parameters. For that purpose, we thoroughly followed the experimental protocol from [20] and fixed the laser power to 170 mW and the pulse width to 200 ns. By selecting the lens magnification $\times 5$, we chose a spot diameter of 15 μm to have a wide laser effective area.

Then, we implemented with NNoM a 4-weights neuron. We set the Y-position to 100 μm to only focus on the X-axis motion. Fig. 2a shows that, when moving along the X-axis of the Flash memory, bit-sets are induced one after another on the whole 32-bit word line and the four quantized weights are precisely faulted. We repeated this experience with different weight values and noticed a perfect reproducibility (as also reported in [20] or [21]). By noticing positions of weights and their most significant bit (MSB), we easily conclude the little endian configuration of the Flash memory. Fig. 2b illustrates how the weights are stored according to the laser bench X-axis.

4.2 Target multilayer perceptron models

Since we can change the value of parameters related to a neuron and therefore alter the elementary neuron calculations, the next step consisted in analyzing the impact of such faults on the integrity of a target model and in measuring the potential drop of accuracy.

For that purpose, we used a MLP model trained on IRIS (accuracy of 96% on the test set) and composed of a 4-input layer, 10-neurons hidden layer and 3-output layer. For each X-step, we evaluated the robustness of the model by feeding it with 50 test samples (i.e., 50 inferences). Classically, we computed the adversarial accuracy by comparing the output predictions to the correct labels and measure the accuracy drop in comparison to the nominal performance without faults.

During one inference, all the weight loading instructions triggered a laser shot (i.e. in total 40 shots for the hidden layer). By targeting one bit line at a given X-position of the laser, only weights on the same address column are faulted with a bit-set. For illustration, as pictured in Fig. 2b, LFI actually induced bit-sets in

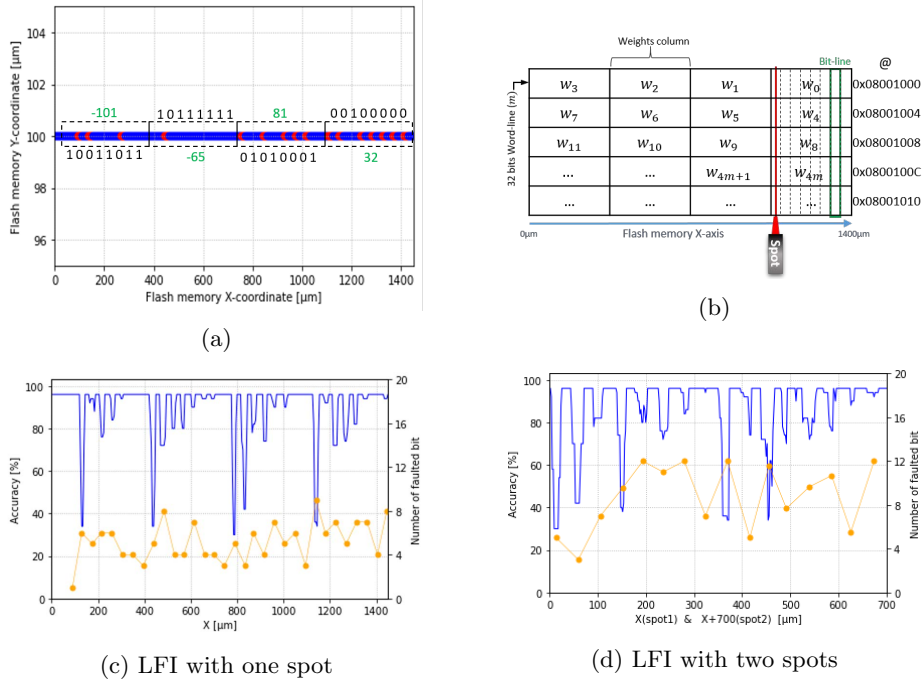


Fig. 2: (a) Bit-set faults (red dots) induced on the 4 weights of a neuron (the original decimal and binary values are given). (b) Flash memory schematic section filled with quantified weights parameters model (with m the m^{th} word-line). (c, d) Impact on the accuracy of Laser Fault Injection on a MLP model (IRIS dataset). Average accuracy over 50 inferences (blue) and number of faulted bits (orange).

refaire plus lisible

the MSBs of weights w_0, w_4, w_8, w_{4m} (with m the m^{th} word-line), which belong to the targeted bit line.

Fig. 2c shows the impact of the laser shots on the model accuracy (blue curve) on the test set. The X-axis corresponds to the position of the laser in the Flash memory. Four regular patterns of accuracy drop corresponding to the four stored weights are clearly visible. The observed accuracy drops (5 to 6 per pattern) match the coordinates of the MSBs with a decreasing correlation. For MSB bits, the model is close to a random-guess level (i.e. around 30%). The orange curve provides the number of faulted bits. On average, around 6 bit-sets are induced and the most harmful configuration corresponds to only 5 bit-sets (at 790 μm) leading to an adversarial accuracy of 30%.

One limitation of the fault model, presented in this experiment, is that only 1/4 of all weights could be faulted during one inference. However, since our laser platform integrates a second spot, we experimented a configuration where one

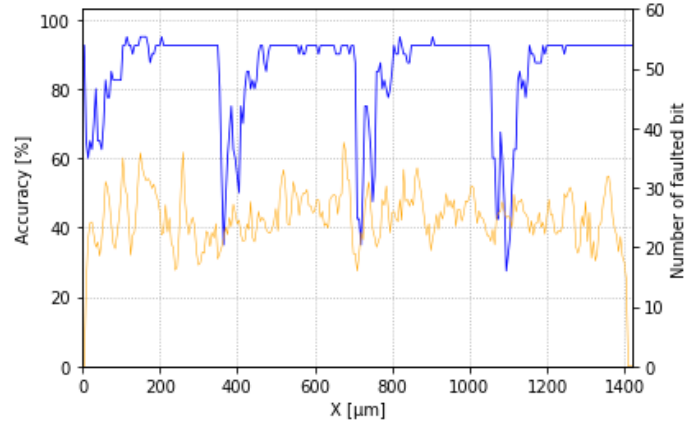
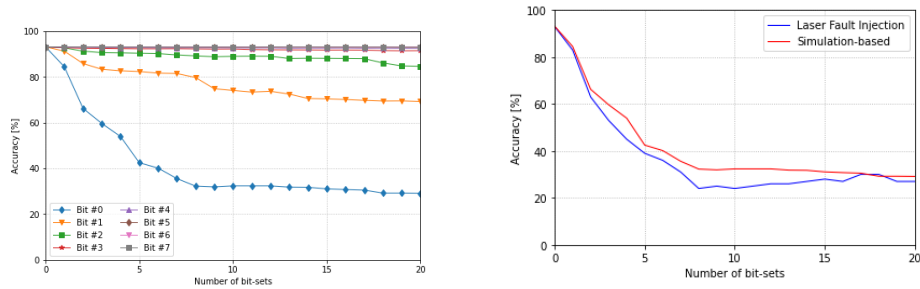


Fig. 3: Impact on the accuracy of LFI on a MLP model trained on MNIST. Average accuracy over 100 inferences (blue) and number of faulted bits (orange).

spot targets the first bit line (i.e the MSB of the weight column w_3 in Fig. 2b) and the second spot targets the bit line #16 (i.e the MSB of the weight column w_1). Then, the two spots move jointly along the X-axis, from 0 to 700 μm and 700 to 1,400 μm respectively. As a result, two weight columns can be targeted during one inference, and half of the weights are likely to be faulted. As reported in Fig. 2d more bit-set faults were injected (up to 12), and the accuracy drop for non-MSB bits is more important. Note that the whole parameters model can be targeted at the same time with a four-spot laser as in [22].

To extend the previous experiments, we embedded the 8-bit MLP model trained on MNIST (see section 3.4). The model gathers more parameters with 600 weights (i.e. 4,800 bits) and 20 biases. It reaches an average accuracy of 92.5% on the test set. We strictly kept the same experimental set-up. Results are presented in Fig. 3. On average, we observe more induced bit-sets with the most important drop of accuracy (65.5%) reached with only 25 bit-sets at $X = 1,100 \mu\text{m}$. We also point out that, in some cases, a very slight improvement of the accuracy appeared.

In comparison with the previous results (Fig. 2c) for IRIS, we note that bit lines are less distinctive. Indeed, even if the position of the most significant bits can be identified, other bits are difficult to distinguish. Actually, in Fig. 2b the bit line (green box) represents all single bit lines connected to the same bit index for different 32-bit words addresses. Since the MNIST model has more parameters, almost all bit lines are linked to a weight parameter value stored in Flash memory. The laser spot, with a size of 15 μm , is larger than a single bit line, explaining why bit line indexes are hardly discernible. Moreover, depending on the laser position, the effective area of the spot can encompass two different bit line indexes.



(a) Influence of the faulty bit number on the accuracy for the weight column 2 (simulation results).

(b) Experimental and simulation results of a LFI BSCA attack targeting the 20 most sensitive MSB of the 2^{nd} weight column.

Fig. 4: Guided-LFI on a MLP model trained on MNIST

4.3 Advanced guided-LFI

Simulation-based constraints. So far, we exhaustively attack all the parameters stored in memory. This *brute-force* strategy may be impractical with deeper models. Therefore, an important adversarial objective is to optimally select the most sensitive bits to fault. For that purpose, as in [11], we adapted the bit selection principle of the standard BFA [3] to our laser fault model. We refer to this adapted attack as BSCA (*Bit-Set Constrained Attack*).

As recommended in [15] and [6], we also adapted the adversarial objective by introducing an *adversarial budget* representing a maximum of faults the evaluator is able to process. Indeed, the random-guess level objective of [3] overestimates the number of faults and raises variability issues because of the last necessary bit-flips needed to ultimately reach the objective [15]. Therefore, we set the adversarial budget to 20 bit-sets.

As inputs, the BSCA has the target model M_W , its parameters W , the weight column index m in the Flash memory, the bit line index b and an adversarial budget S . The output of the BSCA is a new model $M_{W'}^{m,b}$ with W' and W that differ only by at most S bit-sets. Then, we can compute an adversarial accuracy on $M_{W'}^{m,b}$ and measure the accuracy drop compared to the nominal performance of $M_W^{m,b}$. The BSCA proceeds through the following steps:

1. The BFA³ ranks the most sensitive bits of W according to $\nabla_b \mathcal{L}$.
2. Exclude the bits already set to 1 and not related to m and b .
3. Pick the best bit-set and perform the fault permanently in M .
4. Repeat the process until reaching S .
5. The output is the faulted model $M_{W'}^{m,b}$.
6. Perform BSCA over all the weight columns and bit line and keep the faulted model, simply referred as $M_{W'}$, with the worst accuracy (Acc) evaluated on a test set (X^{test}, Y^{test}) : $M_{W'} = \arg \min_{m,b} Acc(M_{W'}^{m,b}, X^{test}, Y^{test})$

³ [3]: <https://github.com/elliothe/BFA/>

The same MLP MNIST model as in section 4.2 is used for the BSCA. First, we simulated the BSCA to find the 20 most significant weights to fault for each weight column. For illustration purpose, Fig 4a shows the effect of bit-set induced on the 8 bits of a weight only for the second weight column (with the MSB referred as Bit #0). The accuracy is given according to the number of bit-sets induced during the inference. Among the most sensitive weights from the second column, only few bit-sets on the most significant bits efficiently alter the model performance. Bit-sets on other bits have less or even no influence on the accuracy.

Experiments and results. The basic idea is to use the BFA to guide the LFI. For that purpose, we need to put to the test that LFI reach (near) identical performance than what expected by simulation with the BSCA.

We ran a BSCA simulation (in Python) over all the weight columns and bit lines that pointed out the MSB of the second column weight as the most sensitive. Therefore, contrary to the previous experiments, the laser source was triggered only when the 20 most sensitive weights were read from the Flash. The laser location was set accordingly to $X = 760 \mu\text{m}$. The laser power was increased to 360 mW to ensure a high fault injection success rate on weights stored in distant addresses.

The model accuracy was evaluated over 100 inferences. The blue curve in Fig.4b represents our experimental results while the red one is the BSCA simulations for the MSB. First, we can notice that experimental and simulation results are almost similar, meaning that we can guide our LFI with high reliability and confidence. The fact that experimental results are slightly more powerful than simulations may be explained by the impact of the width of the laser spot on nearby memory cells. For an adversarial budget of only 5 bit-sets (0.1% faulted bits) the embedded model accuracy drops to 39% which represents a significant loss and a strong integrity impact compared to the nominal performance of 92.5%. Moreover, after 10 bit-sets (accuracy to 25%), the most effective faults have been injected and the accuracy did not decrease anymore. This is an important observation in a security evaluation context since it positions the level of robustness of the model according to the adversarial budget.

5 Discussions

This work presents a complete LFI characterization of the robustness of a DNN model embedded on a Cortex-M platform when exposed to a state-of-the-art weight-based adversarial attack. Being the first work covering this scope, it paves the way to further analysis on other model architectures and platforms (e.g., system-on-chip). We highlight some research tracks that may contribute to the development of reliable evaluation protocols and tools.

Use other weight-based adversarial attacks. Parameter-based attack still lacks of maturity (contrary to adversarial examples with recent attacks specifically designed for robustness evaluation purposes [23]) and recent works highlight limitations of the BFA [6, 15, 17] and propose improvements or alterna-

tives. Thus, considering or combining more attack methods would improve the evaluation by simulating a more powerful adversary.

bof bofWiden the underlying threat model. Being in an evaluation and characterization context, we set in a *worst-case scenario* by simulating an adversary that has access to clone devices and able to perform spatial and temporal profiling works in order to precisely trigger the laser shots on the DUT. An interesting further experience would be to combine side-channel analysis as in [24] in which the injection timing is controlled by monitoring and matching power waveform patterns.

Evaluation of protections. Our work evaluates the intrinsic robustness of a model against a well-defined type of threat for which some protections have been proposed. Additionally to traditional countermeasures against fault injection (e.g., redundancy techniques) [10], specific defense schemes against BFA encompass weight clipping, model pruning [25], clustering-based quantization [2, 26], code-based detectors [27] or adversarial training [6]. To the best of our knowledge, none of these defenses have been practically evaluated against fault attack means such as RowHammer, glitching or LFI. As for adversarial examples (with so many defenses regularly *broken* afterwards) the definition of proper and sound evaluations of defenses against parameter-based attacks against embedded ML models is a research action of the highest importance.

6 Conclusion

This paper addresses the major issue of the robustness evaluation of embedded DNN that are increasingly deployed in large variety of devices. We focused our work on weight-based adversarial attacks that target the parameters of a model stored in memory. We present the first characterization of that threat with a laser fault injection technique and a bit-set fault model against a 8-bit quantized deep neural network model, embedded in a Cortex-M 32-bit microcontroller, a typical platform in many IoT applications. After analyzing the practicality of our fault model on a shallow model trained on the IRIS dataset, we scaled-up our experiment on a wider multilayer perceptron model trained on the MNIST dataset. We demonstrate that we can efficiently guide the laser fault injection attack by adapting a state-of-the-art parameter-based attack (the Bit-Flip Attack): only a few bit-sets were necessary to significantly decrease the model’s accuracy. Our results reveal a high accordance between simulation and practical results. This work paves the way to further experiments and analysis in order to build reliable evaluation methodologies for upcoming standardization and certification schemes of embedded AI-systems.

Acknowledgment

This work is supported by (CEA-Leti) the European project InSecTT (ECSEL JU 876038)⁴ and by the French ANR in the framework of the *Investissements*

⁴ www.insectt.eu

d'avenir program (ANR-10-AIRT-05, irtnanoelec); and (Mines Saint-Etienne) by the French program ANR PICTURE (AAPG2020). This work benefited from the French Jean Zay supercomputer with the AI dynamic access program.

References

1. C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," in *International Conference on Learning Representations*, 2014.
2. X. Hou, J. Breier, D. Jap, L. Ma, S. Bhasin, and Y. Liu, "Security Evaluation of Deep Neural Network Resistance against Laser Fault Injection," *Proceedings of the International Symposium on the Physical and Failure Analysis of Integrated Circuits, IPFA*, pp. 1–6, 2020.
3. A. S. Rakin, Z. He, and D. Fan, "Bit-flip attack: Crushing neural network with progressive bit search," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
4. "Regulation of the european parliament and of the council laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts," 2021.
5. N. Carlini, A. Athalye, N. Papernot, W. Brendel, J. Rauber, D. Tsipras, I. Goodfellow, and A. Madry, "On evaluating adversarial robustness," *arXiv preprint arXiv:1902.06705*, 2019.
6. D. Stutz, N. Chandramoorthy, M. Hein, and B. Schiele, "Random and adversarial bit error robustness: Energy-efficient and secure dnn accelerators," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
7. N. Papernot, P. McDaniel, A. Sinha, and M. P. Wellman, "Sok: Security and privacy in machine learning," in *IEEE European Symposium on Security and Privacy (EuroS&P)*, pp. 399–414, 2018.
8. S. Mangard, E. Oswald, and T. Popp, *Power analysis attacks: Revealing the secrets of smart cards*. Springer Science & Business Media, 2007.
9. R. Joud, P.-A. Moëllic, S. Pontié, and J.-B. Rigaud, "A practical introduction to side-channel extraction of deep neural network parameters," in *Smart Card Research and Advanced Applications: 21st International Conference, CARDIS 2022, Birmingham, UK, November 7–9, 2022, Revised Selected Papers*, pp. 45–65, Springer, 2023.
10. A. Barengi, L. Breveglieri, I. Koren, and D. Naccache, "Fault injection attacks on cryptographic devices: Theory, practice, and countermeasures," *Proceedings of the IEEE*, vol. 100, no. 11, pp. 3056–3076, 2012.
11. F. Yao, A. S. Rakin, and D. Fan, "DeepHammer: Depleting the intelligence of deep neural networks through targeted chain of bit flips," in *29th USENIX Security Symposium*, pp. 1463–1480, 2020.
12. Z. He, A. S. Rakin, J. Li, C. Chakrabarti, and D. Fan, "Defending and harnessing the bit-flip based adversarial weight attack," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14095–14103, 2020.
13. A. S. Rakin, M. H. I. Chowdhury, F. Yao, and D. Fan, "Deepsteal: Advanced model extractions leveraging efficient weight stealing in memories," in *2022 IEEE Symposium on Security and Privacy (SP)*, pp. 1157–1174, IEEE, 2022.
14. L. Liu, Y. Guo, Y. Cheng, Y. Zhang, and J. Yang, "Generating robust dnn with resistance to bit-flip based adversarial weight attack," *IEEE Transactions on Computers*, 2022.

15. K. Hector, P.-A. Moëllic, M. Dumont, and J.-M. Dutertre, "A closer look at evaluating the bit-flip attack against deep neural networks," in *IEEE 28th International Symposium on On-Line Testing and Robust System Design (IOLTS)*, pp. 1–5, 2022.
16. A. S. Rakin, Z. He, J. Li, F. Yao, C. Chakrabarti, and D. Fan, "T-bfa: Targeted bit-flip adversarial weight attack," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2021.
17. K. Lee and A. P. Chandrakasan, "Sparsebfa: Attacking sparse deep neural networks with the worst-case bit flips on coordinates," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4208–4212, 2022.
18. J. Ma, "A higher-level Neural Network library on Microcontrollers (NNoM)," Oct. 2020.
19. A. S. Rakin, Z. He, J. Li, F. Yao, C. Chakrabarti, and D. Fan, "T-bfa: Targeted bit-flip adversarial weight attack," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
20. B. Colombier, A. Menu, J. M. Dutertre, P. A. Moëllic, J. B. Rigaud, and J. L. Danger, "Laser-induced Single-bit Faults in Flash Memory: Instructions Corruption on a 32-bit Microcontroller," *IEEE International Symposium on Hardware Oriented Security and Trust, HOST*, 2019.
21. A. Menu, J. M. Dutertre, J. B. Rigaud, B. Colombier, P. A. Moëllic, and J. L. Danger, "Single-bit Laser Fault Model in NOR Flash Memories: Analysis and Exploitation," *Workshop on Fault Detection and Tolerance in Cryptography, FDTC*, pp. 41–48, 2020.
22. B. Colombier, P. Grandamme, J. Vernay, E. Chanavat, L. Bossuet, L. de Laulanié, and B. Chassagne, "Multi-spot laser fault injection setup: New possibilities for fault injection attacks," in *Smart Card Research and Advanced Applications: 20th International Conference, CARDIS 2021*, p. 151–166, 2022.
23. Y. Liu, Y. Cheng, L. Gao, X. Liu, Q. Zhang, and J. Song, "Practical evaluation of adversarial robustness via adaptive auto attack," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15105–15114, 2022.
24. Y. Fukuda, K. Yoshida, and T. Fujino, "Fault injection attacks utilizing waveform pattern matching against neural networks processing on microcontroller," *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, vol. 105, no. 3, pp. 300–310, 2022.
25. Z. Gao, X. Wei, H. Zhang, W. Li, G. Ge, Y. Wang, and P. Reviriego, "Reliability evaluation of pruned neural networks against errors on parameters," in *2020 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, pp. 1–6, 2020.
26. F. Libano, B. Wilson, M. Wirthlin, P. Rech, and J. Brunhaver, "Understanding the impact of quantization, accuracy, and radiation on the reliability of convolutional neural networks on fpgas," *IEEE Transactions on Nuclear Science*, vol. 67, no. 7, pp. 1478–1484, 2020.
27. M. Javaheripi and F. Koushanfar, "Hashtag: Hash signatures for online detection of fault-injection attacks on deep neural networks," in *IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, pp. 1–9, 2021.