

Evaluating single event upsets in deep neural networks for semantic segmentation: An embedded system perspective[☆]

Jon Gutiérrez-Zaballa^{a,*}, Koldo Basterretxea^a, Javier Echanobe^b

^a Dept. of Electronics Technology, University of the Basque Country, Bilbao, 48013, Spain

^b Dept. of Electricity and Electronics, University of the Basque Country, Leioa, 48940, Spain

ARTICLE INFO

Dataset link: <https://ipaccess.ehu.es/HSI-Drive/>

Keywords:

Single bit upsets
Robustness evaluation
Model compression
Embedded artificial intelligence
Semantic segmentation

ABSTRACT

As the deployment of artificial intelligence (AI) algorithms at edge devices becomes increasingly prevalent, enhancing the robustness and reliability of autonomous AI-based perception and decision systems is becoming as relevant as precision and performance, especially in applications areas considered safety-critical such as autonomous driving and aerospace. This paper delves into the robustness assessment in embedded Deep Neural Networks (DNNs), particularly focusing on the impact of parameter perturbations produced by single event upsets (SEUs) on convolutional neural networks (CNN) for image semantic segmentation. By scrutinizing the layer-by-layer and bit-by-bit sensitivity of various encoder–decoder models to soft errors, this study thoroughly investigates the vulnerability of segmentation DNNs to SEUs and evaluates the consequences of techniques like model pruning and parameter quantization on the robustness of compressed models aimed at embedded implementations. The findings offer valuable insights into the mechanisms underlying SEU-induced failures that allow for evaluating the robustness of DNNs once trained in advance. Moreover, based on the collected data, we propose a set of practical lightweight error mitigation techniques with no memory or computational cost suitable for resource-constrained deployments. The code used to perform the fault injection (FI) campaign is available at <https://github.com/jonGuti13/TensorFI2>, while the code to implement proposed techniques is available at <https://github.com/jonGuti13/parameterProtection>.

1. Introduction

There is an increasing interest in developing domain-specific processors for the deployment of artificial intelligence (AI) algorithms at the edge and at the endpoints. This trend is driven by the necessity of freeing endpoint AI-based systems from having to send acquired data to external, more powerful computing machines to be processed and wait for a response before any action can be executed. The goal is to make embedded AI systems completely autonomous, avoid security and reliability issues associated with data transmission, and reduce response latency. Achieving embedded AI processing autonomy makes it possible to extend the applicability of complex AI algorithms, such as the increasingly widespread deep neural networks (DNN), to domains in which there may be severe communication bandwidth constraints, hard reliability and security specifications and/or real-time response requirements. Examples of such application areas are intelligent vision, autonomous navigation, advanced driving assistance systems (ADAS), autonomous driving systems (ADS), and aerospace applications such as remote sensing and airborne flight control [1].

If AI has to be pervasively integrated in systems that autonomously operate in real-world environments, AI processors must meet not only demanding performance specifications, but also strict safety and reliability standards. This is nowadays a mayor concern in the development of AI-based systems that must be integrated in safety-critical applications such as in the aerospace [2] and automotive [3] domains. For instance, the main agent that jeopardizes the performance in the aerospace applications is the exposure of aircraft and spacecraft electronics to radiation [4]. When a radiation particle interacts with electronic components, the logical value stored in a cell may be altered, resulting in a single event upset (SEU) or, more specifically, a single bit upset (SBU) when it only affects a single bit or a multiple bit upset (MBU) if multiple bits are affected. Field programmable gate arrays (FPGA) are particularly sensitive to SBUs, which can affect both the sequential elements of the implemented circuit (flip-flops or Block RAMs) and the configuration memory (LUTs). In the simplest case, a SBU will cause a bit of one of the implemented artificial neural network (ANN) model parameters to flip. Depending on the specific

[☆] This work was partially supported by the Basque Government, Spain under grants PRE_2023_2_0148 and KK-2023/00090, and by the Spanish Ministry of Science and Innovation under grant PID2020-115375RB-I00.

* Corresponding author.

E-mail address: j.gutierrez@ehu.es (J. Gutiérrez-Zaballa).

<https://doi.org/10.1016/j.sysarc.2024.103242>

Received 10 May 2024; Received in revised form 1 July 2024; Accepted 18 July 2024

Available online 20 July 2024

1383-7621/© 2024 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

position of the SBU, the network output may change and differ from the expected one (critical errors), putting the safety of the system at risk. In the context of ADS, soft errors at terrestrial altitudes occur due to the following factors: the interaction of high-energy cosmic neutrons with silicon, the interaction of low-energy cosmic neutrons with high concentrations of 10B in the device and the emission of alpha particles from trace radioactive impurities in the device materials [5]. According to [6], when exposed to terrestrial neutrons, a bit in SRAM has a probability of $1.33 \cdot 10^{-24}$ to flip in 1 ns. As explained in [7], for a typical neural network of more than 10 million parameters, the probability of at least one bit-flip occurring in a month is 10%. This problem is aggravated as a consequence of increasing miniaturization and the use of lower voltage levels in modern integrated circuits.

In addition to background radiation, embedded devices are also vulnerable to other factors such as disturbance errors in storage devices (SRAM and DRAM) and also to deliberate malicious bit-flip attacks (BFA). Moreover, alternative technologies to traditional meta-oxide-semiconductor devices for the more efficient in-memory computation, e.g. memristive devices, are particularly prone to bit instabilities [8,9]. Disturbance errors are a general class of reliability problem that affects memory and storage technologies such as SRAM, DRAM, flash and hard-disk [10]. Write failure and read disturb are two major causes of SRAM technology failures tightly related to miniaturization and low-voltage operation [11]. In the case of DRAM, the so-called RowHammer vulnerability is a well-known issue that produces data corruption and the appearance of multiple bit-flips [10]. This has attracted the attention of attackers, who can maliciously flip memory bits in DRAM without the need of any data access privileges [12,13]. If extended to safety-critical domains such as ADS and aerospace, the consequences of data-oriented attacks can be catastrophic, especially when the BFAs are targeted [14].

The outstanding performance of state-of-the-art DNNs in many applications is very often at the expense of increasing their size and complexity dramatically [15]. Although it strongly depends on the design of the digital circuit that implements their functionalities (the “application layer”, [16]), large models are generally more robust due to their overparameterized nature [17,18]. However, deploying complex AI models into resource-constrained embedded processing systems usually implies, firstly, applying one or various compression techniques to transform large models into the so-called lightweight deep learning models. These procedures aim to reduce the computational complexity and the memory footprint of original DNNs maintaining comparable performance. The most used approaches include network pruning, parameter quantization, knowledge distillation, and architecture search [19]. Secondly, it is usually needed to design custom processing units that help accelerate algorithm execution to meet speed and power consumption specifications. This is achieved by tailoring specific pipelines for data parallelism and applying arithmetic optimization techniques to make the most of both available memory and computational resources. Depending on the selected target technology (the “physical layer”, [16]), e.g., embedded GPUs, FPGAs, or application-specific integrated circuits (ASIC), there exist different design constraints to consider and specific optimization techniques to apply. In any case, our concern is focused on the effects of SBUs on the parameterization of the DNNs, i.e., how the values of trainable parameters are altered and the potential consequences of these changes on the performance of the DNN models.

In this paper, we analyze the reliability of encoder–decoder DNN models designed for image segmentation tasks against SBU disturbances. This analysis has been performed with two main objectives in mind. Firstly, the focus is on the consequences that compression techniques such as model pruning and parameter quantization may have on the robustness of these models. Secondly, a more comprehensive study has been conducted to precisely determine how such disturbances modify the model performance in the inference process. The aim was to use this knowledge to propose design techniques that help improving

model robustness, and thus system reliability, with no cost in terms of computational complexity and memory footprint. This work was originally motivated by the necessity to improve the reliability of some image segmentation DNN models designed to be applied in ADAS/ADS when implemented on the target processing devices, both embedded GPUs and FPGAs. To date, there are few published studies that analyze in detail and in a statistically significant way the impact of SBUs on the robustness of image segmentation DNN models, particularly when subjected to compression for deployments at the edge. This paper aims to fill this gap by performing a detailed layer-by-layer and bit-level analysis based on SBU emulation to better understand the mechanisms that produce failures in the behavior of the models. The analysis has been performed both for 32-bit floating-point and 8-bit quantized representations to cover final implementations on different devices. Although the study has been necessarily carried out for a specific model architecture and a particular dataset, the purpose of the work is to provide guidelines and analysis tools applicable to other models, particularly those of the encoder–decoder type. Moreover, it also has served to program a set of memory-, computation- and training-free procedures for the protection, to a certain extent, of the performance of such models against SBUs. The code has been made public at [20].

The main contributions of this paper are:

- We analyze in a statistically significant way how SBUs alter the performance of encoder–decoder DNN models in image segmentation tasks, determining the sensitivity of the output to single bit-flips in the model parameters according to the layer depth, the parameter type, and their binary representation.
- We analyze how and why model pruning affects the robustness of the models against SBU perturbations.
- We study the consequences of applying parameter quantization in terms of model robustness.
- We provide in [21] a modification of the original TensorFlow2 [22] code to allow applying fault injection (FI) campaigns on quantized TensorFlow Lite models and on big unquantized TensorFlow2 models too.
- We describe some simple rules to improve model robustness by protecting model parameters against the consequences of SBUs with no cost on computational complexity neither on the memory footprint.
- We provide in [20] the code to perform fault mitigation based on the above mentioned technique.

The paper is organized as follows: Section 2 includes the related work done in the field of DNN robustness evaluation and hardening techniques against SBUs. Then, in Section 3, we describe the architecture of the encoder–decoder model as well as the compression techniques whose effect on model robustness will be studied. All the experimental results in relation with the fault injection campaign are collected in Section 4. Section 5 describes the proposed fault mitigation technique. Finally, we provide the conclusions in Section 6.

2. Related work

2.1. Robustness of ANNs against bit-flips

Most published papers on ANN robustness against bit-flips focus on classification tasks, while very few analyze semantic segmentation models [23]. The majority of the authors have carried out experimental fault injection campaigns combined with statistical analysis, while only a few have developed a theoretical analysis based on the development of a vulnerability model.

Among the experimental works, papers such as [24,25] test the robustness of several image classification convolutional neural network (CNN) architectures against single bit-flips in their weights, but with relatively superficial analyses. In [26], the authors compare the robustness of multilayer perceptrons and CNNs of different sizes, concluding

that larger networks with more layers are more robust than smaller, shallower networks. Some authors, such as [27], investigate the impact of individual bit-flips while comparing floating-point and fixed-point representations, concluding that the latter are more robust.

Works like [28,29] evaluate the effects of both quantization and pruning on the robustness of models. The former shows that compressed models are more fault-resilient compared to uncompressed models in terms of bit error rate zero accuracy degradation, but statistical significance due to a short number of experiments leads to a high variance in the prediction results. The latter finds that integer-only quantization acts as a fault mitigation technique by reducing the overall range of the data. It also concludes that pruning enhances the resilience of deep models as a consequence of the reduction in the occupied area and execution times.

Some works focus on methods to ensure the statistical significance of fault injection campaigns. In [30], the authors present a methodology to evaluate the impact of permanent faults affecting CNNs in automotive applications. Similarly, [31] describes how to correctly specify statistical fault injections and proposes a data analysis on the parameters of a CNN for image classification tasks to reduce the number of fault injections required to achieve statistically significant results. [32] presents one of the most exhaustive and complete analysis of the vulnerability of 32-bit floating-point CNNs for image classification tasks. The author considers various factors such as bit position, bit-flip direction, parameter sign, layer width, activation function, normalization, and model architecture. The key findings are: the vulnerability is caused by drastic spikes in a parameter value, the spikes in positive parameters are more threatening, an activation function that allows negative outputs renders the negative parameters vulnerable as well, and the dropout and batch normalization (BN) layers are ineffective in preventing the massive spikes that bit-flips cause. In [33], the author presents two exhaustive tools for fault injection in models created with both TensorFlow1 and TensorFlow2 [22], and performs a thorough analysis of the consequences of fault injections in different classification models. Additionally, the article also explores techniques to identify the source of the error by the analysis of changes in the model's predictions. In [34,35], floating-point and fixed-point data type model implementations are analyzed showing that fixed-point data provide the best trade-off between memory footprint reduction and CNN resilience. Similarly, [36] performs a comprehensive layer-wise fault analysis of homogeneously and heterogeneously quantized DNNs, suggesting that quantizing the DNN model heterogeneously to fewer bits helps increase the model's resiliency.

As a consequence of the growing concern about the threat of deliberate BFAs, some authors are studying the effects of simultaneous bit disturbances across multiple model parameters. In the extensive work presented in [14], BFAs are analyzed according to their untargeted or targeted nature, and an effective mitigation methodology against targeted BFAs is proposed. In [37], the authors apply a progressive bit search algorithm to investigate the effects of bit-flip-based weight attacks and obtain some relevant observations regarding quantized DNN sensitivity: the most sensitive parameters are the ones close to zero (large parameter shift), the weights in the front-end layers are the most sensitive, and BFAs force almost all inputs to be classified into one particular output class. Similarly, in [38], the authors evaluate the accuracy degradation of an 8-bit integer quantized DNN as a consequence of untargeted random BFAs, one-shot BFAs and progressive BFAs for image classification tasks. The authors show that with the most exhaustive BFA, i.e. progressive BFA, the accuracy drops to 1% after just 5 iterations. This result aligns with the one described in [39], where the author of the progressive BFA algorithm shows that an 8-bit integer quantized ResNet-18 can malfunction after just flipping 13 weight bits out of 93 million.

One of the first works that attempted to formalize a theoretical method to evaluate the robustness of DNN models was presented in [40]. It describes a layer-wise relevance propagation model based on

the analysis of the contribution of individual neurons to the final loss. With a similar approach, gradient-based methods such as that explained in [41], analyze the sensitivity of each network layer according to the importance (relevance) of the weights during inference. In [7], the authors present a theoretical analysis of error propagation on some commonly used processing layers in image classification models when the sign bit is flipped. In [42], a vulnerability model based on some key features such as gradient and absolute value of the parameters is constructed to reduce the necessary amount of fault injections to perform robustness analysis. [43] proposes BinFI, a fault injector for finding safety-critical bits in machine learning applications that significantly outperforms random fault injection methods in terms of computational costs. Finally, in [44], the authors formulate a bit-flip-based weight fault propagation model for 32-bit floating-point CNNs to analyze the robustness of ReLU-based models and propose a hardening method based on function upper bounding. Nevertheless, the applicability of these methods to segmentation networks is not straightforward because they are primarily designed for classification networks, where the layers under study are typically simple convolutions and fully-connected layers. Additionally, gradient-based analyses can be inaccurate due to noisy gradients and challenges such as vanishing or dying gradients associated with common activation functions like Sigmoid and ReLU.

All the above-mentioned works are focused on detection and classification DNN models. Regarding the few papers that deal with segmentation networks, [45] claims to have performed the first fault injection study of DNNs performing semantic segmentation, proposing a critical/tolerable fault categorization for a 32-bit floating-point DeeplabV3+ network. [46] evaluates the reliability of neural networks for various tasks, including semantic segmentation, implemented in 32-bit floating-point representation on a GPU trained with Supervised Compression for Split computing. In [47], the author stresses the importance of statistical significance in the analyses, and proposes a fast reliability methodology exploiting statistical fault injections in a U-Net model for image segmentation. The author performs a comparison of the results obtained with this method to those obtained by random FI and improperly-defined statistical FI campaigns and shows the inability of the latter campaigns to reveal sensitive parameters of U-Net.

2.2. Hardening and protection of neural networks against fault occurrence

Most of the papers that propose methods to protect ANNs against faults are also concerned with image classification models. According to the proposed protection technique, these can be grouped as those that use redundancy, modifications of the activation functions, modifications of the parameters, and modifications of the training/inference process.

One of the most basic methods for error detection and protection is the use of a checksum together with a replication of the model, which, while highly effective, is also prohibitively costly in terms of memory consumption and computing overhead. [48,49] are two examples where full duplication has been shown to be effective. As proposed in [50–52], using a proper model sensitivity analysis makes it possible to optimize redundancy to protect only the most critical layers. In the same direction, [17] presents a software methodology based on a triple modular redundancy technique to selectively protect a reduced set of critical neurons, under a single fault assumption, by a majority voter correction technique. In [53], the authors characterize fault propagation not only by exposing the FPGA/GPU to neutron beams but also by performing a thorough fault injection campaign. Based on the observations, the authors propose a strategy to improve system reliability by adapting algorithm-based fault-tolerant solutions to CNNs, i.e., adding invariants to the code for quick error detection or correction. In [54], the authors propose protecting the matrix multiplication operation of the CNNs in GPUs based on a three-stage methodology to selectively protect CNN layers to achieve the required diagnostic coverage and performance trade-off: sensitivity analysis to misclassification per CNN

layers using a statistical fault injection campaign, layer-by-layer performance impact and diagnostic coverage analysis, and selective layer protection. Finally, [55] proposes a reduced-precision duplication with comparison technique to improve the reliability of computing devices to reduce overhead. It is suitable for mixed-precision architectures, such as NVIDIA GPUs.

The bounding of activation functions is an alternative technique to redundancy to reduce implementation overhead. In [56], the authors perform a comprehensive error resilience analysis of DNNs for image classification tasks subjected to hardware faults in the weight memory. Then, ClipAct is applied, an error mitigation technique based on squashing the high-intensity faulty activation values to alleviate the impact of faulty weights on predictions. In [57], Ranger is proposed, a low-cost fault corrector which selectively restricts the ranges of values in specific DNN layers to dampen the large deviations typically caused by transient faults leading to silent data corruptions. In [58], FitAct is proposed, a low-cost approach to enhance the error resilience of DNNs for image classification tasks by deploying fine-grained post-trainable activation functions. The main idea is to accurately bound the activation value of each individual neuron via neuron-wise bounded activation functions to prevent fault propagation in the network. In [18], the author presents a comprehensive methodology for exploring and enabling a holistic assessment of the trilateral impact of quantization on model accuracy, activation fault reliability, and hardware efficiency. The framework allows for the application of different quantization-aware techniques, fault injection, and hardware implementation and directly measure the hardware parameters. A novel lightweight protection technique integrated within the framework that ensures the dependable deployment, evaluating the maximum values of the layers' activations and replacing the out-ranged values with either lower or upper-bound to avoid fault propagation, is also proposed. Finally, [44] presents a boundary-aware ReLU to improve the reliability of DNNs by determining an upper bound of the activation function which is theoretically calculated so that the deviation between the boundary and the original output cannot affect the final result.

Some other works explore methods to enhance network robustness by directly modifying the network parameters. Based on the findings of [40] about the connection between neuron resilience and its contribution to the final prediction score, the authors of [59] propose a methodology based on architectural and feature optimization to avoid critical bottlenecks and balance the feature criticality inside each layer. In [60], the authors propose MATE, which is a low-cost CNN weight error correction technique based on the observation that, as all mantissa bits of the weights are not closely related to accuracy, some of them can be replaced with error correction codes. Therefore, MATE can provide high data protection with no memory overhead. In [61], the authors perform a comparison of the robustness among 32-bit floating-point, 16-bit floating-point, and 8-bit integer formats for image classification tasks and propose an opportunistic parity method to detect and mask errors with zero storage overhead.

There are also proposals to harden network performance by the modification of the training or inference process. The authors of [62] propose a bipolar vector classifier which can be easily integrated with any CNN structure for image classification tasks that end in a fully connected layer. The underlying idea is that as the weights of the classifier are binarized to ± 1 , the resulting final feature vector will only contain positive/negative values that will also be binarized to ± 1 , and thus, a pattern will be created. Each class has a specific reference pattern so, to assign the final feature vector to a certain class, it will be compared with all the reference patterns, and the winning class will be the one with the smallest Hamming distance. In [63], the authors compare the accuracy of quantized DNNs (QNNs) for image classification tasks during accelerated radiation testing when trained with different methodologies and implemented with a dataflow architecture in an FPGA. The authors find that QNNs trained with fault-aware training, a kind of data augmentation methodology to allow

the network to also experience errors during training, make QNNs more resilient to SEUs in FPGAs. In [64], an efficient error detection solution for object detection-oriented CNNs is proposed based on the observation that, in the absence of errors, the differences between the input frames and the inference provided by the CNN should be strictly correlated.

Finally, regarding the specific works that focus on the hardening of DNNs against BFAs, [37] proposes binarization-aware training and piecewise clustering as methods to enhance the resistance of quantized DNNs. The authors conclude that applying binary quantization, increasing network capacity, and using dropout or Batch-Norm regularization are effective techniques to build resistance, while applying adversarial weight training or pruning is shown to be ineffective. In [38], a three-step algorithm based on mean calculation, quantization and clipping is proposed to reconstruct the perturbed quantized weight matrix to tolerate the faults caused by BFAs. The overhead introduced by the process is small and the protected DNN can better cope with progressive BFA than non-protect DNN (accuracy of 60% and 1% respectively after 5 iterations). A completely different approach is presented in [14], where the authors propose a dynamic multi-exit architecture that trains extra internal classifiers for hidden layers that can tolerate the existing attacks which flip bits in one specific layer. The experiments are conducted using well-known DNN structures and image classification datasets. Apart from the above-described methods, which aim to mitigate the effect of such attacks, there are also some other methods that focus on verifying the integrity of the models. In [65], the authors propose to extract a unique signature from the original DNN prior to deployment and then verify the inference output on-the-fly while trying to add the minimum performance and resource overhead. Indeed, due to the strict temporal and memory footprint constraints that edge devices must adhere to, the protection systems of interest to us are those that introduce minimal memory/computation overhead.

3. Model development and optimization

3.1. Architectural design and training

The segmentation model used as reference for this study is a U-Net, an encoder-decoder fully convolutional network (FCN) for image segmentation, but adapted to use hyperspectral images (HSI) as inputs.




Fig. 1. A sample segmented image of HSI-Drive v2.0.

The most recent version of the model [66] was trained using the HSI-Drive v2.0 dataset [66], intended for developing ADAS/ADS systems using HSI (Fig. 1). The results on the test set can be found in Table 1.

The model, depicted in Fig. 2, features a 5-level encoder-decoder architecture comprising two sequences of 3×3 2D convolutional layers (initially with 32 filters) followed by batch normalization and ReLU activation at each level. Additionally, it includes one 2×2 max-pooling layer per encoder level and one 2×2 transposed 2D Convolutional layer per decoder level. The resulting model comprises 31.14 million parameters and requires 34.60 GFLOPS per inference to execute (Table 2). Detailed information regarding the training and testing procedures can be found in [66].




Fig. 2. Architecture of the unpruned unfolded unquantized U-Net model which has been used as reference.

Table 1
Segmentation metrics on the test set using the 32-bit floating-point unpruned and pruned models.

Class	Unpruned			Pruned		
	Metric					
	Rec.	Prec.	IoU	Rec.	Prec.	IoU
Road	99.69	98.14	97.84	99.57	98.16	97.75
Marks	91.95	95.70	88.30	92.74	95.12	88.53
Vegetation	97.66	96.32	94.15	97.30	96.11	93.61
Sky	96.45	96.71	93.38	96.04	97.52	93.75
Others	81.84	93.42	77.38	82.26	92.14	76.86
Global	97.23	97.17	94.67	97.12	97.06	94.48
Weighted	91.98	95.68	88.40	92.33	95.36	88.48

Table 2
Complexity of the different models under study after applying pruning/quantization (status).

Metric	Status			
	00	10	01	11
Params. (M)	31.14	0.31	31.14	0.31
Operat. (OPS)	34.60 ^a	8.49 ^a	34.60	8.49
Size (MB)	118.77	1.24	29.70	0.31

^a Before applying quantization, operations are FLOPS.

3.2. Model compression

The most prevalent model compression techniques among deep learning developers for later implementation on edge devices involve applying pruning and/or quantization after training the 32-bit floating-point model. To evaluate how these compression techniques affect the model's robustness against SBUs, several factors must be considered. First, the size of the model and the architecture designed for its implementation as a digital processor, which directly influences the ratio of unused to used resources and the probability for a SBU to happen. However, device occupation is not the only parameter influencing the Device Vulnerability Factor (DVF), representing the probability of a configuration bit being critical for the design [67]. Overparameterized models can absorb more SBUs without producing critical errors since there are more irrelevant weights and biases whose perturbations do not alter the output. Thus, it is one of the objectives of this paper to analyze whether large reductions in DNN model sizes can be achieved using these compression techniques without degrading model accuracy or robustness to SBUs. For this, we are assuming streaming-like architectures with independent resources allocated for all layers so the impact of a SEUs is isolated [68].

3.3. Pruning

Pruning facilitates the reduction of both the number of parameters to be stored and the number of computations to be performed. The concept involves eliminating the least significant parameters of the model. Depending on which parameters and how they are pruned, pruning can be categorized as fine-grained/sparse/unstructured, where the least important weights are rounded to 0, or coarse-grained/dense/structured, where entire filters are removed from the computational graph. While the former method generally allows for pruning more weights without harming performance, it is only justified if the processing system is able to optimize multiply and accumulate (MAC) operations with sparse matrices.

In this work, we have applied a conventional structured pruning approach that has been customized to perform model optimization in an iterative manner. The algorithm basically analyzes the computational complexity of each layer while evaluating the impact of the pruning process on the model's accuracy to guarantee negligible impact on overall performance. As shown in Table 2, applying this method a 99% reduction in the number of parameters and a 75% reduction in the number of operations was achieved.

3.4. Quantization

Quantization aims to reduce the number of bits needed to store the model parameters, thereby reducing memory footprint (see Table 2). Additionally, it can speed up both data transfer and model inference for custom processor implementations. Depending on the target device, quantization can be more or less fine-grained in terms of homogeneity, uniformity, scale factor, symmetry, and mixed-precision.

In this article, we have chosen a general and standardized heterogeneous quantization scheme, known as post-training integer quantization (PTQ), as implemented by TensorFlow Lite [69]. This procedure converts 32-bit floating-point numbers (weights and activation outputs) to the nearest 8-bit fixed-point numbers, while biases, due to the greater sensitivity of the models to perturbation on these parameters, are converted to 32-bit fixed-point numbers. This heterogeneous quantization scheme allows for model size reduction while preserving accuracy [69,70]. As shown in Table 3, comparable segmentation metrics are obtained for the quantized model to those with the unquantized model, as quantization schemes applied to already trained models, are simple to apply and produce negligible accuracy degradation for most widely used ANN models [71,72].

Table 3

Segmentation metrics on the test set using the 8-bit integer unpruned and pruned models.

Class	Unpruned			Pruned		
	Metric					
	Rec.	Prec.	IoU	Rec.	Prec.	IoU
Road	99.68	98.09	97.79	99.59	98.15	97.75
Marks	91.67	95.79	88.12	92.51	95.46	88.61
Vegetation	97.63	95.76	93.59	97.38	95.88	93.47
Sky	93.60	98.29	92.10	94.02	98.45	92.65
Others	82.49	92.56	77.37	82.81	91.54	76.92
Global	97.10	97.04	94.44	97.07	97.00	94.39
Weighted	91.44	96.06	88.24	91.99	95.81	88.55

3.5. Quantization of pruning

To fully compress the model, both techniques can be applied consecutively: first pruning and then quantization. As shown in Table 3, after the quantization of the pruned model segmentation accuracy on the test set remains mainly unaltered. Table 2 sums up model complexity figures for each version of the reference model: original (00), pruned (10), 8-bit quantized (01), and pruned and quantized (11). As can be seen, the original memory footprint is reduced from 118.77 MB to just 317.44 KB.

4. Assessment of the fault injection campaign

To test the models' robustness against SBUs, an extensive fault injection campaign was conducted on the aforementioned models using a modified version of the TensorFI2 framework [22]. This modification enables a more memory-efficient implementation by directly accessing the parameters of the model without creating additional copies or intermediate tensors, ensuring that FI on complex, large models does not result in an excessive memory overhead. The original code has also been extended to support quantized TensorFlow Lite models and is available at [21]. To evaluate the perturbation produced by injected faults, the corrupted FCNs have been assessed using ten test images that represent the diversity of driving conditions in the dataset described in Section 3.

Performed FI campaign involved injecting single bit-flips into the parameters of the FCN model under analysis. Parameter sets include the weights/kernels of the 2D convolution layers ($conv2D_k$), bias of the 2D convolution layers ($conv2D_b$), weights/kernels of the 2D transposed convolution layers ($conv2Dtr_k$), bias of the 2D transposed convolution layers ($conv2Dtr_b$), gamma of the batch normalization layers (BN_γ), and beta of the batch normalization layers (BN_β). Thus, in what follows, $p1$ set of the network refers to the weights of the first $conv2D$ layer, while $p2$ set represents the biases of that layer. In like manner, when it is mentioned that an error has been injected into parameter x , it means it has been injected into one of the m elements that comprise parameter x set.

As a general rule, for a given parameter, the higher the bit position in which the fault is injected, the greater its impact on the output, disregarding the sign bit. However, not every layer and every parameter contributes equally to the output. Previous studies [30,34,54] focused on tasks such as image classification and object detection, have shown that faults injected in the exponent bits (interval [23–30]) are more

likely to alter the output. To verify whether this is also generally the case for this encoder–decoder model aimed to image segmentation and in order to set reasonable bounds for the range of bits to be modified, a first round of 150 faults per-layer spanning the entire range [0–31] were injected. We could verify that in some cases, bit-flips in the range [20–23] and in the sign bit (31) also modified significantly the inference result. Based on these preliminary results, a comprehensive fault injection campaign was conducted on bit positions [20–31].

The quantity of fault injections was set to 1550 single bit-flip faults per layer, totaling 155 000 injections to assure statistically representative experiments according to Eq. (1) as proposed in [73].

$$n = \frac{N}{1 + e^2 * \frac{N-1}{t^2 * p * (1-p)}} \quad (1)$$

where N represents the number of possible faults per-layer (number of parameters * parameter bit-width), e denotes the error margin, which was set to 0.025 (2.5%), t is the cut-off point corresponding with the confidence level, set to 1.96 (95%), and p is the estimated probability of faults resulting in a failure, set to 0.5 since, as it is a priori unknown, a conservative approach is to use the value that maximizes the sample size [73]. This is also done in other studies such as [30,34]. Finally, n represents the minimum number of injections for the study to be statistically significant, capped at 1550 for the chosen parameters.

To assess the impact of the fault injection campaign, it is necessary to first define what constitutes an error. Contrary to the definition given in classification tasks, where the output is a single class and a prediction is deemed erroneous if the top-ranked class predicted by the original model changes, defining errors in segmentation tasks is more challenging. Changes in the predicted class of some isolated pixels may not significantly alter the overall interpretation of the image, so in order to have an unequivocally defined metric, in this analysis the error rates are defined related to changes in the predicted classes at any pixels in the test images.

4.1. Fault injections in unquantized models

As a consequence of the 32-bit floating-point data representation, the effect of a bit-flip varies with its position (Fig. 3). To maintain clarity in terminology, the most significant bit (MSB), the sign bit, is assigned position 31, while the least significant bit (LSB) is assigned position 0.

Figs. 4 and 5 depict the bit-flip error rate according to the bit position and the parameter position for both the unpruned and pruned models. As expected, the unpruned model exhibits greater robustness and can better withstand the injected faults. Once the general statistical results have been observed, let us now analyze in detail the insights of the sensitivity of the model performance to perturbations in the parameters.

4.1.1. Analysis of robustness of the unpruned model

Output $conv2D$ layer $conv2D_{22}$ ($p99$ – $p100$ sets in Fig. 2) consists of n_{class} filters of $1 * 1 * C$ dimensions and n_{class} biases, where n_{class} is the number of classes to be predicted (6 in this model) and C is the number of output channels from the previous layer, $conv2D_{21}$ (32 in this model).

A signal calibration analysis (see Fig. 6) reveals that the activations of the last convolutional layer lie within the range [−6.7656, 11.7859], while the absolute values of the bias parameters turn out to



Fig. 3. Single-precision floating-point format (IEEE 754 format).

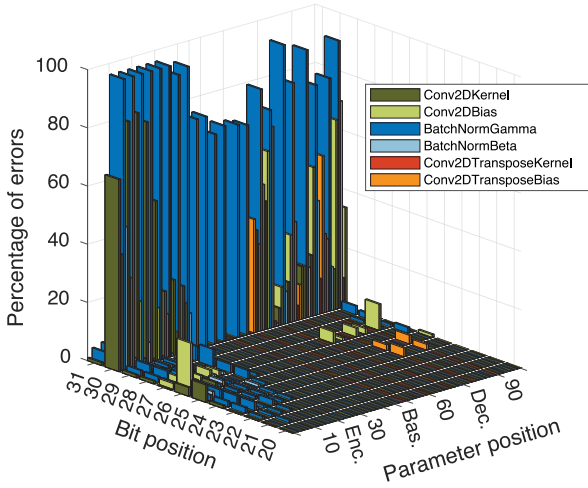


Fig. 4. Bit-flip error rate on the original unpruned model.




Fig. 5. Bit-flip error rate on the original pruned model.




Fig. 6. Graphic representation of the unpruned model's calibration study (layers conv2D_21 to conv2D_22).

be smaller than 1 (−0.8494, 0.3171, −0.0275, 0.0394, −0.1706, 0.1090). Consequently, a change in the sign (bit 31) of one of the biases is unlikely to alter the winning class. However, a bit-flip in bit 30 will significantly increase the bias' magnitude, because in 32-bit floating-point representation all six biases contain a '0' in their 30th bit. The

impact of a bit-flip on bit 30 depends on two factors: the sign of the bias and the probability of the model predicting each class. In the scenario where the bit-flip occurs in a negative bias (biases 0, 2, and 4), the bias becomes so negative that the associated class will never be predicted by the model. Consequently, if the original model predicted many pixels of that class in an image, produced error rate will be very high; conversely, if the presence of that class was low, produced error will be small. If the bit-flip happens in a positive bias (biases 1, 3, and 5), the bias becomes large enough so that the associated class is always predicted by the model. Again, depending on the frequency of that class in an image, the resulting error rate can be high or negligible. Both the sign of the bias associated to each class and the probability of the model predicting each class can be known beforehand. Therefore, the error rate can be predicted using Eq. (2).

$$\%error_{30} = \sum_{j=0}^5 P_{fi}^j * P_m^j \tag{2}$$

where j is the class index, P_{fi}^j is the probability of a bit-flip occurring in class/bias j (assumed to be $\frac{1}{6}$) and P_m^j is the probability of the faultless model predicting j as the output class. Applying this equation yields:

$$\begin{aligned} \%error_{30} &= \frac{1}{6}(0 + 55.09 + 4.41 + 73.05 + 7.47 + 83.73) \\ &= 37.29 \end{aligned}$$

The slight discrepancy between this value (37.29%) and the one obtained from the experimental data (34% in Fig. 4) is due to the statistical error of the fault injection since it may not result in the exact same number of flips for each of the 6 classes (assuming a constant probability of $\frac{1}{6}$ is thus only an approximation). Additionally, the highly unbalanced terms in the equation make the effect of the non-constant probability more noticeable.

Regarding the weights, $conv2D_{22w}$, they are multiplied by the output activation of layer $conv2D_{21}$, which, due to the use of ReLU activation functions, is always ≥ 0 . Hence, if the bit-flip occurs in a weight that is multiplied by a 0-valued activation from $conv2D_{21}$, the model will be immune to it (unless the bit-flip produces a non-desirable special value such as NaN in the weight itself). The impact of a flip in the sign bit is considered negligible and the effect of a flip in bit 30, similar to biases, will only be relevant if the weight is negative and belongs to a filter of the class predicted by the model or if it is positive but does not belong to a filter of the class predicted by the model.

Gamma parameter in batch normalization layers The gamma (γ) parameter of the BN layers is the most sensitive one in the model (located in the third, fourth, seventh and eighth positions in each of the 8-parameter blocks of Fig. 2). This high sensitivity is due to the very nature of the BN operation, as specified in Eq. (3) for a single pixel of a single channel of an activation map (the final result is a scalar value):

$$BN(Conv2D(x)) = \gamma \frac{(\sum w \odot x + b) - \mu}{\sqrt{\sigma}} + \beta \tag{3}$$

where x represents an input 3D array, w is the weight 3D array of a filter of the $conv2D$ layer, b is the bias of a filter of the $conv2D$ layer, γ is the positive gamma parameter of the BN, β is the bias parameter of the BN, and σ and μ are the variance and the mean of the training data.

It so happens that in this model, γ values are positive, smaller than 2 but higher than 0.1. Thus, in this case, three different scenarios can be considered: values higher than 1, equal to 1, and smaller than 1. In the first case, except for bit 30, the remaining 7 bits of the exponent are '1's and the mantissa is non-zero, so a bit-flip in bit 30 converts the original value to a NaN (Fig. 3). The propagation of the NaN through the network is unavoidable, resulting in a completely wrong output map. In the second case, where the mantissa is zero, the same bit-flip converts the value 1 into $+\infty$. This situation is akin to the previous one, as the network cannot digest this error except in situations where the value of

the weight of the next *conv2D* is negative so as the ReLU transforms the $-\infty$ into a 0. In the last case, considering 0.1 as an example, the bit-flip in position 30 will increase the value to something around e^{37} , which is near the maximum representable value ($\approx 3.4028e^{38}$). Moreover, as Eq. (3) shows, that result is divided by the square root of the variance of the data, which is usually much smaller than 1. Consequently, the resulting value would exceed the maximum value and would saturate to $\pm \infty$. In fact, as observed in the calibration analysis (e.g. Figs. 7, and 8), BN layers always increase the range of the processed data.




Fig. 7. Graphic representation of the unpruned model's calibration study (layers *conv2D* to *conv2D_3*).

Weight parameters in *conv2D* layers As shown in Fig. 4, faults injected in weight parameters of *conv2D* layers located in the encoder branch (p sets below 40 in Fig. 2) have more impact on the output than when injected in the decoder branch (p sets above 48). This discrepancy is due to the presence of skip-connections between encoder and decoder branches since faults injected in the first layers will directly propagate to the outputs through skip connections. Moreover, longer data paths will be also involved in the propagation of the errors through encoder/decoder branches, including BN layers which, as mentioned above, increase the range of the signals with the risk of overflows.




Fig. 8. Graphic representation of the unpruned model's calibration study (layers *conv2D_4* to *conv2D_7*).

In the worst case, if a bit-flip occurs in bit 30th of a positive weight, when multiplied by a positive input, it produces a value close to the maximum representable number (otherwise it would be absorbed by the ReLU and 0 would be propagated, which is less problematic). Thus, most likely, the value would saturate to infinity when passing through a BN layer. This observation aligns with the fact that the weights of the first 5 convolutions (Figs. 7 and 8) are significantly larger than those of the rest of convolutions in the encoder. Hence, it is more likely for a bit-flip in bit 30 to cause an output mismatch when it is produced in these initial layers, as can be verified in Fig. 4.

Table 4

Percentage of positive bias of the most sensitive *conv2D* layers in the unpruned model. Note: the tables are read in a U way (from top to bottom in the encoder layers and from bottom to top in the decoder layers).

Name	Pos.	Total	%	Name	Pos.	Total	%
conv2D	15	32	46.88	conv2D_21	30	32	93.75
conv2D_1	12	32	37.50	conv2D_20	11	32	34.38
conv2D_2	19	64	29.69	conv2D_19	31	64	48.44
conv2D_3	6	64	9.38	conv2D_18	10	64	15.63
conv2D_4	18	128	14.06	conv2D_17	34	128	26.56
conv2D_5	16	128	12.50	conv2D_16	20	128	15.63
conv2D_6	23	256	8.98	conv2D_15	174	256	67.97
conv2D_7	55	256	21.48	conv2D_14	109	256	42.58

Bias parameters in *conv2D* layers According to obtained statistics (Fig. 4), biases of *conv2D* layers in the decoder branch appear to be more sensitive than biases in the encoder branch. The explanation of this is that, for a bit-flip in position 30 of a bias in a *conv2D* layer to significantly alter the output, it must be a positive valued bias. Errors in negative biases are absorbed by ReLU layers, since BN layers between *conv2D* and the ReLU do not change the sign of that value, given that the multiplicative constants are always positive. *conv2D* bias values share the same characteristics, i.e. positive and smaller than unity, as case 3 of the γ parameters described in Section 4.1.1. Therefore, when passing through BN layers, its value will increase to infinity, causing errors in the outputs.

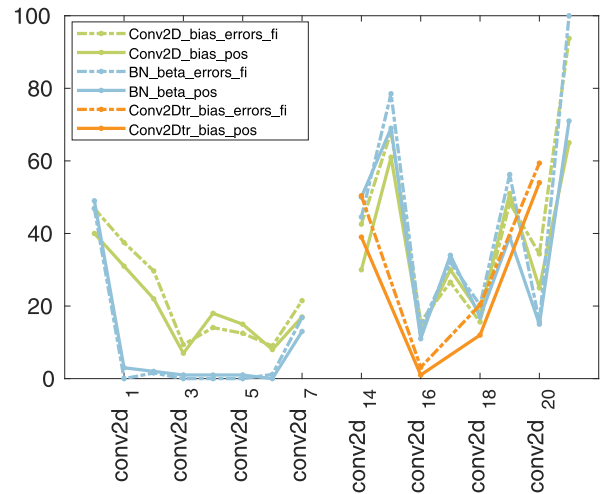


Fig. 9. Bit-flip error in the 30th bit (dashed) and ratio of positive parameters (solid) by parameter index: *conv2D_b* (green), *BN_{beta}* (blue) and *conv2Dtr_b* (orange) in the unpruned model. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

We can analyze the influence of bias parameters in more detail by inspecting the presence of positive biases in the primary *conv2D* layers (deeper layers closer to the base of the model are analyzed in a separate paragraph later on). As shown in Table 4, the percentage of positive biases in the decoder branch is much higher than in the encoder branch, which is consistent with our analysis and the obtained statistical observations. In fact, if the correlation between the presence of positive bias and produces error rates is examined, parallel trends are observed (green lines in Fig. 9).

Beta parameters in batch normalization layers Similarly to the analysis performed for biases in *conv2D* layers, comparable error sensitivity for positive β parameter values is observed. As shown in Fig. 9 (blue lines), there is also a high degree of correlation between the error rate and the percentage of positive β in each BN layer (Table 5).

Table 5

Percentage of positive betas of the most sensitive BN layers in the unpruned model. Note: the tables are read in a U way (from top to bottom in the encoder layers and from bottom to top in the decoder layers).

Name	Pos.	Total	%	Name	Pos.	Total	%
bn	15	32	46.88	bn_21	32	32	100
bn_1	0	32	0	bn_20	5	32	15.63
bn_2	1	64	1.56	bn_19	36	64	56.25
bn_3	0	64	0	bn_18	13	64	20.31
bn_4	0	128	0	bn_17	41	128	32.03
bn_5	0	128	0	bn_16	18	128	14.06
bn_6	3	256	1.17	bn_15	201	256	78.52
bn_7	43	256	16.80	bn_14	114	256	44.53

Bias parameters in conv2D_{tr} layers Finally, in Table 8, the percentage of biases that are positive for the primary conv2D_{tr} layers is grouped to see if the biases of the conv2D_{tr} also follow the above dynamics. Plotted in orange in Fig. 9, it is observed how, as in previous cases, both lines follow the same trend.

At this point, it is important to note that only bit-flips in the bit 30 of parameters have been considered so far. Indeed, bit-flip on the rest of the bits produce many fewer errors, as seen in Fig. 4. However, there are particular bits in specific parameters, such as bit 26 in the bias of the first conv2D layer, associated with a significant error rate. Explanation on this issue will be provided later on this paper. Finally, it has to be remarked that the correlation between lines depicted in Fig. 9 is not perfect. The reason is that faulty bits in negative biases can also result in a prediction error since they may turn positive activations into 0 values, and there are also corrupted positive biases that can be absorbed by a ReLU layer if they are multiplied by a null or negative weight (the error should be smaller then).

Deep layers: the base of the U-Net From the calibration and the FI analyses, it can be concluded that the deepest layers of the network (p33-p58 sets in Fig. 2) practically do not contribute to the inference process. Nevertheless, training of shallower models (encoder depths 2, 3, and 4) with fewer convolution filters (8, 16, 32 in the first convolution) produced considerably worse results. In fact, compressing trained large sparse models has been shown to be more effective than training smaller dense models [74].




Fig. 10. Graphic representation of the unpruned model's calibration study (layers conv2D_8 to conv2D_11).

Convolution weights in deep layers are so small that only the biases (Figs. 10 and 11) contribute to the activations. Theoretically, even a bit-flip in the most significant bit of the biases is quickly absorbed by the network and should not influence the outputs. However, it can be seen in Fig. 4 that error statistics for FIs in deep layers (p41-p50 sets) are not negligible. This is because the previous reasoning does not take into account events where a parameter value becomes NaN. That is what precisely occurs when the bit 30 of one of the γ parameters whose value is very close to 1 flips up.




Fig. 11. Graphic representation of the unpruned model's calibration study (layers conv2D_12 to conv2D_14).

4.1.2. Analysis of robustness of the pruned model

Output conv2D layer In the pruned model, the structure of conv2D_22 is very similar to that of the unpruned model, as it also consists of 6 filters of size 1 * 1 * 32 and 6 biases. As shown in Fig. 12, there is a slight difference in the range of the activations, which is now [-8.3144, 15.1673]. However, since the values of the bias parameters are essentially identical to those of the unpruned model (-0.8495, 0.3150, -0.02991, 0.0373, -0.1736, 0.1066), the conclusions that can be drawn are quite the same.

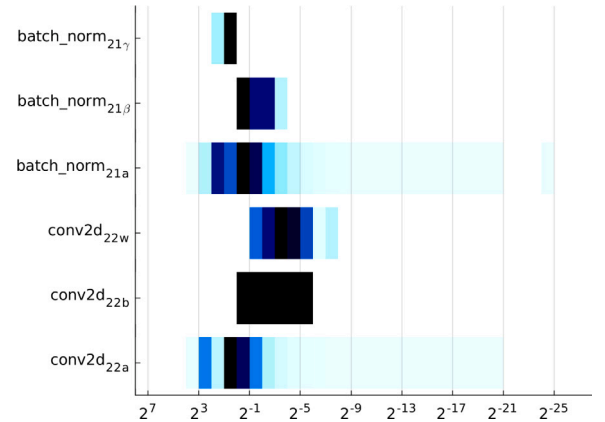


Fig. 12. Graphic representation of the pruned model's calibration studies (layers conv2D_21 to conv2D_22).

Applying Eq. (2) to this case, the following value is obtained:

$$\begin{aligned} \%error_{30} &= \frac{1}{6} (0 + 55.66 + 4.35 + 72.93 + 7.37 + 83.13) \\ &= 37.24 \end{aligned}$$

The theoretical value (37.24%) is not far from that obtained experimentally and depicted in Fig. 5 (32%), being the difference a consequence of the statistical nature of the FI campaign.

Let us now focus on lower position bits. In Fig. 5 it can also be seen that faults in bits 29, 28, 27, and 26 do not generate errors in the output, but changes in bit 25 do. If we analyze the binary exponent of the 6 biases under consideration (see Fig. 13), it is observed that all of them contain a '1' in bits [26-29]. A bit-flip in those positions will produce a reduction of the represented magnitude, and hence, the result of the bias addition will decrease too. However, since the biases are small compared to the range of activations (see Fig. 12), it is reasonable to think that this reduction will not produce a meaningful change in the output map. Same circumstance will happen if a bit-flip from 1 to 0 occurs in bit 25. But, focusing on the last bias of Fig. 13, a bit-flip

from 0 to 1 in bit 25 causes the partial exponent to fill with 1s. Thus, the value increases from +0.1 to +1.7, which is a mayor perturbation that will probably produce a prediction error.




Fig. 13. Exponent part of the 32-bit floating-point representation of the biases of the last *conv2D* of the pruned model.

Gamma parameter in batch normalization layers In the pruned model, it is also found that the γ of the BN layers is the most sensitive parameter, and the same explanation holds. However, in the seventh BN layer (*p27* set in Fig. 2), there is a clear decrease in the number of errors associated with bit-flips in bit 30 and an increase in the number of errors associated with bit-flips in bits [26–29]. The reason is that some of the γ values in this layer, unlike in the rest of the BN layers, are greater than 2 but smaller than 3. This means that the exponents in the binary representations change drastically from 01111111 to 1000000. Then, a change in bit 30 of those γ s will no longer cause a value to be a *NaN*; instead, it will be a very small number that is less likely to produce a critical error. In turn, a bit-flip in bits [26–29] will now result in a considerable increase in magnitude, so the error rates for those bit positions grow compared to the layers with γ values smaller than 2.

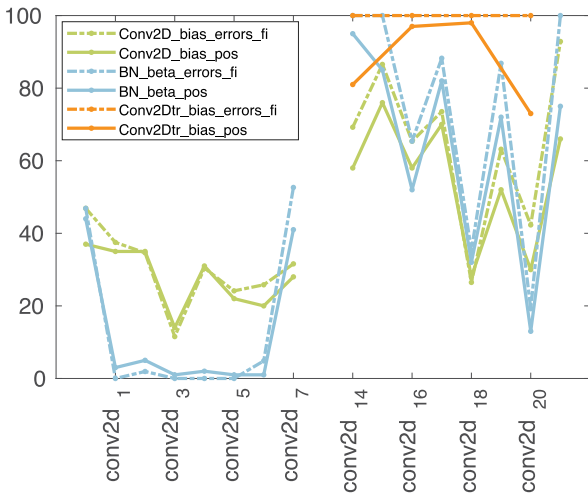


Fig. 14. Bit-flip error in the 30th bit (dashed) and ratio of positive parameters (solid) by parameter index: *conv2D_b* (green), *BN_{beta}* (blue) and *conv2Dtr_b* (orange) in the pruned model. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Bias parameters in conv2D layers As observed in the unpruned model, flips in bit 30 of the biases of *conv2D* layers located in the encoder branch do not produce as many errors as in the decoder part. Table 6 contains the total number of positive biases compared to the total number of biases in those layers. Again, as shown in Fig. 14 (green lines), there is a high correlation between the sign of the biases and the error rate.

As previously mentioned, Fig. 5 exhibits a peak for the bias parameter of the first *conv2D* layer in bit 26. Furthering the discussion from the paragraph on the last *conv2D*, here 31 out of the 32 biases have an absolute value greater than 2^{-15} (all 0s in the exponent except bits 29, 28 and 27). Hence, altering bit 26 produces the most significant increase in absolute value. While this increment may seem insignificant, considering the first layer’s activations (refer to Fig. 15)




Fig. 15. Graphic representation of the pruned model’s calibration study (layers *conv2D* to *conv2D₃*).

which are the lowest across the network ([−2.6148, 2.1645]), even small perturbations can induce notable changes.

Table 6
Percentage of positive bias of the most sensitive *conv2D* layers in the pruned model. Note: the tables are read in a U way (from top to bottom in the encoder layers and from bottom to top in the decoder layers).

Name	Pos.	Total	%	Name	Pos.	Total	%
conv2D	15	32	46.88	conv2D_21	26	28	92.86
conv2D_1	12	32	37.50	conv2D_20	11	26	42.31
conv2D_2	18	52	34.62	conv2D_19	24	38	63.16
conv2D_3	6	52	11.54	conv2D_18	9	34	26.47
conv2D_4	14	46	30.43	conv2D_17	25	34	73.53
conv2D_5	14	58	24.14	conv2D_16	17	26	65.38
conv2D_6	16	62	25.81	conv2D_15	45	52	86.54
conv2D_7	24	76	31.58	conv2D_14	36	52	69.23

There is also a noticeable peak in the error rates of *conv2D_{16b}* for changes in bit 27 (*p72* set in Fig. 2). Among the 26 biases, 4 of them have 6 of the 7 LSBs of the exponent set to 1 (all except bit 27), placing their absolute value in $[1.5259e^{-5}, 2 * 1.5259e^{-5}]$ range. There are other 2 biases that are slightly outside this range. Simultaneously, the activations of this layer are not very large (refer to Fig. 16). In consequence, a bit-flip in bit 27 notably increases the value of those 4 biases to values larger than 1 but smaller than 2 (with the other 2 biases reaching a value close to 1).




Fig. 16. Graphic representation of the pruned model’s calibration study (layers *conv2D₁₅* to *conv2D₁₇*).

According to this, we could expect small error rate peaks in certain bit positions for every layer in Fig. 5, given that all bias parameter sets contain some values with binary representations that are one bit away from completing the 7 least significant bits of the exponent (partial

exponent). However, this is not the case for every layer. The reason is that layers with higher activation ranges are more robust against such errors.

An illustrative example is provided by the biases in *conv2D_20* (p92 set in Fig. 2), where 10 out of 26 biases show an absolute value in the range $[1.5259e^{-5}, 2 * 1.5259e^{-5}]$. There are other 2 biases that are close to this range. In those cases, flipping bit 27 will cause a noticeable increase in the values, making them greater than one. However, those bit-flips do not induce more errors than in *conv2D_16* because of the broader range of activations $[-3.38, 4.93]$ for *conv2D_16* and $[-8.97, 6.94]$ for *conv2D_20*.

Beta parameter in batch normalization layers The errors resulting from injecting a flipped bit 30 on the β parameter in the batch normalization layers increase for the pruned model. The percentage of β parameters that are positive relative to the total is noted in Table 7. As shown in Fig. 14 (blue line), the correlation between parameter soundness and the error rate is almost perfect.

Table 7

Percentage of positive betas of the most sensitive BN layers in the pruned model. Note: the tables are read in a U way (from top to bottom in the encoder layers and from bottom to top in the decoder layers).

Name	Pos.	Total	%	Name	Pos.	Total	%
bn	15	32	46,88	bn_21	28	28	100
bn_1	0	32	0,00	bn_20	5	26	19,23
bn_2	1	52	1,92	bn_19	33	38	86,84
bn_3	0	52	0,00	bn_18	12	34	35,29
bn_4	0	46	0,00	bn_17	30	34	88,24
bn_5	0	58	0,00	bn_16	17	26	65,38
bn_6	3	62	4,84	bn_15	52	52	100
bn_7	40	76	52,63	bn_14	52	52	100

Bias parameter in conv2Dtr layers This parameter is one of the most pruned as Table 8 shows the difference in the number of parameters before and after applying pruning. Visualizing the two lines plotted in orange in Fig. 14, it is observed that this is the case for which the two curves are farther apart. However, it should be noted that this is the oddest case, since all the biases are positive. This would imply an error rate of 100%, but we have previously discussed that very large positive biases can still be absorbed, reducing their adverse impact on the output.

Table 8

Percentage of positive bias of the most sensitive *conv2Dtr* layers of the unpruned (left) and pruned (right) models. Note: the table is read from bottom to top.

Name	Pos.	Total	%	Name	Pos.	Total	%
conv2Dtr_4	19	32	59.38	conv2Dtr_4	8	8	50
conv2Dtr_3	13	64	20.31	conv2Dtr_3	10	10	100
conv2Dtr_2	4	128	3.13	conv2Dtr_2	4	4	100
conv2Dtr_1	129	256	50.39	conv2Dtr_1	2	2	100

Finally, *conv2Dtr_4* (p90 set in Fig. 2) is analyzed in detail as it has a peak in bits 24 and 25 (Fig. 5). Based on Fig. 17, problems arise when the bit-flip increases the magnitude of a value. Since all the biases of the parameter have bits in the range [26–29] set to ‘1’, a bit-flip in those bits will not increase their value, so an error is unlikely to happen. However, in many cases, a bit-flip in bit 24, and especially in bit 25, causes bias values to increase to even above unity (when the least significant 7 bits of the exponent are set to ‘1’).

Deep layers: the base of the U-Net The base of the pruned U-Net is also quite insensitive to single bit-flips (Fig. 5) due to small parameter values. However, it is observed that in the graph region located between the first two transposed *conv2D* layers (p51–p58 sets in Fig. 2) there is an increase in the error rate. This is not a surprise if we look at how the ranges of the weights in layer *conv2D_13* (p55 set in Fig. 2) have increased after the pruning and the subsequent fine-tuning process (there is a noticeable drift to the left when comparing Figs. 11 and 18).

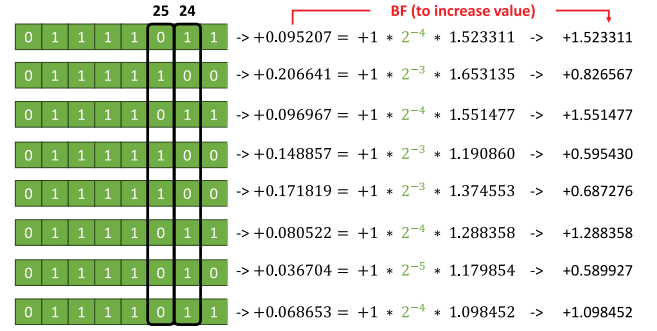


Fig. 17. Exponent part of the 32-bit floating-point representation of the *conv2Dtr_4*, p90 set in Fig. 2.




Fig. 18. Graphic representation of the pruned model's calibration study (layers *conv2D_12* to *conv2D_14*).

4.1.3. Sensitivity changes in pruned layers

The most intensely pruned *conv2D* layers of the encoder branch are *conv2D_4*, *conv2D_5*, *conv2D_6* and *conv2D_7* (p17–p18, p21–p22, p25–p26 and p29–p30 sets in Figs. 2 respectively), which are precisely the layers with the largest increase in the measured error rate. The pruning process did not result in a modification of the range of the weight values because, as expected, pruned filters contained very small weights. As can be observed comparing parameter ranges depicted in Figs. 8 and 19, the original model contains many irrelevant weights that are not present in the pruned model. In consequence, any SBU will be potentially more critical in the pruned model although, on the other hand, in most of implementations the probability of a SBU in smaller models is also lower.

4.1.4. Batch normalization folding

Since the γ parameters of the BN layers have shown to be the most sensitive ones of the network upon SBUs, fusing the BN operation with the previous *conv2D*, technique known as BN folding, may produce an improvement in the robustness of the model. Following the notation of Eq. (3), both operations can be fused into just one *conv2D* operation as follows:

$$\begin{aligned}
 BN(Conv2D(x)) &= \gamma \frac{(\sum w \odot x + b) - \mu}{\sqrt{\sigma}} + \beta \\
 &= Conv\hat{v}2D(x) = \sum (\hat{w} \odot x + \hat{b})
 \end{aligned} \tag{4}$$

where $\hat{w} = \frac{\gamma}{\sqrt{\sigma}} w$ and $\hat{b} = \gamma \frac{b - \mu}{\sqrt{\sigma}} + \beta$.

Although the original aim of this technique was to reduce the number of parameters to be stored and improve training performance, here we analyze if, indeed, it can help reduce prediction error rates




Fig. 19. Graphic representation of the pruned model's calibration study (layers *conv2D_4* to *conv2D_7*).

produced by bit-flips. Fig. 20 shows the statistics of errors for the unpruned model after having folded all BN layers. As can be seen, the layers associated with the highest error rates have been eliminated and, as a consequence of that, the base of the U-Net now does not produce any errors (there are neither *NaN*s nor $\pm \infty$). The error rates associated with the kernel of the *conv2D* layers have decreased while the error rates in the kernel and bias parameters of the *conv2D_tr* layers remains unaltered. However, the error rates to injected faults in the biases of *conv2D* layers have increased. To explain this increment, attention must be paid to the changes produced in the ranges of these values.

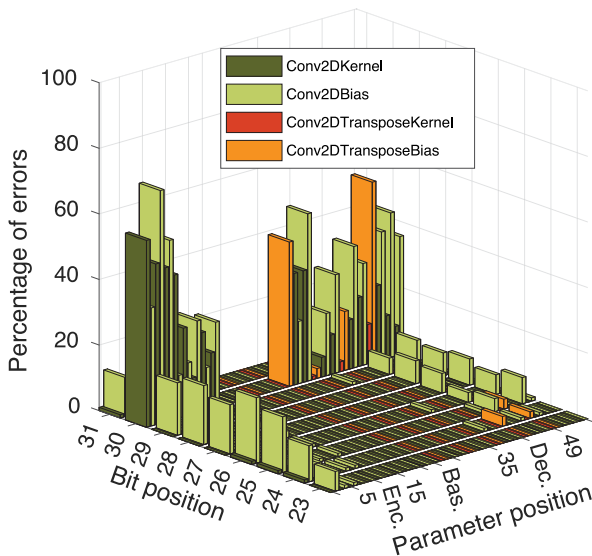


Fig. 20. Bit-flip error rate on the unpruned folded model.

Table 9 collects some bias range statistics for *conv2D* layers, both for the unfolded and the folded models. As expected, the folding operation has significantly increased the absolute value of the bias parameters. Biases with absolute values between 1 and 2 in the folded model are a bit-flip away in bit 30 from producing a *NaN* value. Similarly, bit-flips in any of the remaining bits of the exponent of biases in the range above 2 will increase their value, augmenting the probability of producing output errors. This analysis aligns with Fig. 20, showing that layers with the most values in the 'risky' range also have the highest error rates, and vice versa.

The same significant differences observed in the unpruned folded model are also evident in the pruned folded model (compare statistics in Figs. 5 and 21). Firstly, the layers associated with the highest error

Table 9

Bias values per range before and after BN folding in some sensitive *conv2D* layers of the unpruned model.

Range/ <i>conv2D</i>	Not folded				Folded				
	_3	_6	_19	_21	_3	_6	_19	_21	
$-\infty < x \leq -2$	0	0	0	0	3	0	1	0	0
$-2 < x \leq -1$	0	0	0	0	3	1	12	6	0
$-1 < x \leq 1$	32	64	256	64	32	15	63	28	26
$1 < x \leq 2$	0	0	0	0	5	0	7	27	2
$2 < x \leq +\infty$	0	0	0	0	6	0	6	3	4

rates have been removed. Secondly, the base of the U-Net does not produce any errors either. Thirdly, the error rates associated with the *conv2D* kernel decreases and the error rates in *conv2D_tr* kernel and bias parameters remains stable. Lastly, the error rates associated with the bias of *conv2D* layers have increased. Once again, some bit positions [24–29] which previously caused few or no errors, are now generating a significant number of them. Looking to Table 10, we see the same range increase of bias values as in the unpruned model after BN folding, which leads to higher error rates in the corresponding layers.




Fig. 21. Bit-flip error rate on the pruned folded model.

In conclusion, although BN folding has positive effects on the general robustness of the model by eliminating layers with the highest error rates and the most dangerous situations from areas with low information, and reducing the number of parameters to protect, hence decreasing vulnerability, it also increases the probability of errors in some layers due to the increment of the absolute value of the *conv2D* biases.

Table 10

Bias values per range before and after BN folding in some sensitive *conv2D* layers of the pruned model.

Range/ <i>conv2D</i>	Not folded				Folded				
	_3	_6	_19	_21	_3	_6	_19	_21	
$-\infty < x \leq -2$	0	0	0	0	3	0	0	0	0
$-2 < x \leq -1$	0	0	0	0	3	1	5	0	0
$-1 < x \leq 1$	32	52	62	38	28	12	51	16	22
$1 < x \leq 2$	0	0	0	0	5	0	7	19	3
$2 < x \leq +\infty$	0	0	0	0	6	0	6	3	3

4.2. 8-bit integer quantized models (QNNs)

The applied quantization scheme (Section 3.2) involves converting weights to 8-bit integers and biases to 32-bit integers. Additionally, BN layers have been folded into previous convolution layers to reduce the number of parameters and speed up computations. According to [69],




Fig. 22. Two's complement arithmetic for 8-bit and 32-bit integer/fixed-point representation.

the quantized version \hat{r} of a real number r is approximated by Eq. (5), where S is a real positive scale factor, q is an 8/32-bit integer value, and Z is the zero-point integer value (0 for symmetric quantization).

$$\hat{r} \approx r = S(q - Z) \quad (5)$$

Thus, the value q_y , resulting from $Y = Wx + b$ is:

$$q_y = \frac{S_w S_x}{S_y} (q_w q_x - Z_x q_w + q_b) \quad (6)$$

Due to the per-tensor and mainly symmetric quantization, the number of q_i values that need to be stored is considerably greater than that of Z_i and S_i values. Additionally, due to the quantization process itself, the S values have to meet certain restrictions [69] and are not of interest for the comparison being made. Consequently, bit-flips were only injected on q_w and q_b , which are represented in two's complement arithmetic, the typical computer representation for fixed-point (signed integers are just one example) binary values (Fig. 22).

In this representation, a notable drawback arises from the double functionality of the sign bit: a bit-flip in the sign bit converts a very small negative number into a very large positive number (and vice versa) and a very small positive number into a very large negative number (and vice versa). By contrast, two's complement does not define a specific representation for NaN or infinity, seemingly alleviating the issues associated with these values.




Fig. 23. Bit-flip error rate on the unpruned quantized model. Weights are quantized to 8 bits (bit 7 is the MSB). Biases are quantized to 32 bits and faults are injected across all bit positions due to disparities in their ranges.

We repeated the statistical fault injection campaign on the QNN and obtained results are graphed in Figs. 23 and 24 respectively. Based on these results, the unpruned model shows, in principle, a superior robustness.

4.2.1. Analysis of robustness of the unpruned model

Relevance of weights compared to biases The injection of a fault in the convolution weights shows minimal impact on the output (see Fig. 23). Due to skip-connections in the model, the only critical scenario is the appearance of bit-flips in the first layer. The significance of the

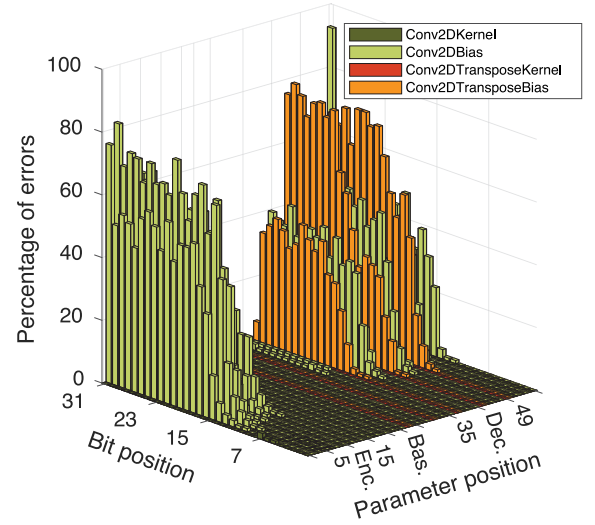


Fig. 24. Bit-flip error rate on the pruned quantized model. Weights are quantized to 8 bits (bit 7 is the MSB). Biases are quantized to 32 bits and faults are injected across all bit positions due to disparities in their ranges.

perturbations in weights is negligible compared to the biases, which is a noteworthy observation. This means that for an error protection technique based on redundancy, it could practically be limited to storing and checking the biases.




Fig. 25. Bit-flip error in [17-30] bits (dashed) and ratio of positive bias (solid) in some $conv2D_b$ encoder layers of the quantized unpruned model.

The biases Contrary to what was observed for the unquantized models, encoder biases induce more errors than decoder biases. Since bit-flips cannot result in NaNs or $\pm\infty$ in two's complement representation, the analysis differs. For QNNs the layer position is more determinant, mainly due to the presence of skip connections and longer error propagation paths. Indeed, the initial layers of the model are clearly the most sensitive ones.

However, analyzing the sign of the biases remains relevant to explain the bit-flip errors in the initial layers of the encoder. For int-8

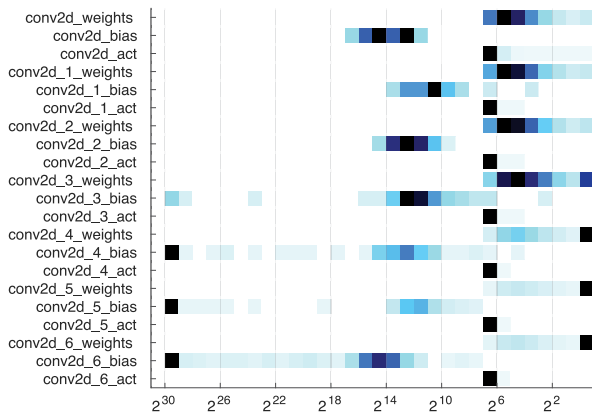


Fig. 26. Calibration study of the unpruned quantized model (q_w , q_b and q_a) from layers *conv2D* to *conv2D_6*.

representation, perturbed parameters produce output errors that depend on the position of the flipped bits, being positions [17–30] the most sensitive. Once a bit-flip in one position produces an increase in the bias value high enough to produce an error at the output, changes in higher positions become inconsequential. In order to get a general estimation of a single statistical error rate value generated by upset events in any of the considered bit positions, a linear weighting approximation based on bit significance has been applied. Fig. 25 reveals a practically identical trend between the calculated bit-flip error and the ratio of positive biases.




Fig. 27. Calibration study of the unpruned quantized model (q_w , q_b and q_a) from layers *conv2D_13* to *conv2D_17*.

Finally, we analyze why bit-flips above bit 16, on average, result in significant errors, those below bit 13 cause very few errors, and below bit 9 induce none. By referencing Table 11 and Figs. 26–29, it can be seen that the maximum bias values hovers around 2^{17} . Thus, bit-flips in higher positions lead to a magnified bias, resulting in a significant increase in the error rates. Moreover, certain layers exhibit a requirement for 30/31 bits to represent all values. To expound on this observation, we must differentiate between highly significant parameters (those in the first layers of the encoder and the last layers of the decoder) and those of low significance (found in the base zone).

If a parameter is placed in a highly relevant area of the model (e.g., *conv2D_3*, *conv2D_4*, *conv2D_5*, *conv2D_6*, *conv2D_17*, *conv2Dtr_3*, *conv2D_18*, *conv2Dtr_4*), the necessity for a high number of bits primarily stems from representing negative values (Table 11). However, there are relatively few instances of very negative values in these areas (as observed in Figs. 26–28). Notably, channels with very negative biases typically have negligible relevance to the output, as the ReLU activation function effectively nullifies their contributions.




Fig. 28. Calibration study of the unpruned quantized model (q_w , q_b and q_a) from layers *conv2D_18* to *conv2D_22*.




Fig. 29. Calibration study of the unpruned quantized model (q_w , q_b and q_a) from layers *conv2D_7* to *conv2D_12*.

Conversely, in comparatively irrelevant areas of the model (e.g., *conv2D_9*, *conv2D_10*, *conv2D_11*, *conv2Dtr*, *conv2D_12*, *conv2D_13*, *conv2Dtr_1*), the necessity for a high number of bits arises from the need to represent both very positive and very negative values (Table 11). In these regions, there are numerous instances of very positive or very negative values (as observed in Figs. 27 and 29). However, many of the channels either contain very negative biases (which are absorbed by ReLU activation) or very positive values that are subsequently multiplied by nearly zero weights (thus, contributing insignificantly).

Table 11

Necessary bits to represent all pos./neg. values of the *conv2D_b* (green) and *conv2Dtr_b* (orange) parameters in the quantized unpruned model.

Layer	Positive	Negative	Layer	Positive	Negative
0	18	17	22	13	16
1	15	15	21	17	13
2	16	16	20	15	15
3	15	31	4	19	28
4	16	31	19	16	17
5	15	31	18	15	31
6	18	31	3	20	29
7	18	20	17	17	31
8	23	23	16	16	31
9	32	31	2	18	30
10	32	31	15	19	21
11	32	31	14	19	20
0	32	31	1	32	31
12	32	31	13	32	31

In summary, although certain layers exhibit biases quantized to 30/31 bits, the high-order bits (typically 15/16 bits) are essentially

redundant as they solely extend the sign. These bits could potentially serve as a means of detecting and/or correcting sensitive bit positions.

Deep layers: the base of the U-Net Quantizing the model does not alter the significance of the base of the U-Net. Even when a bit-flip occurs at the most sensitive bit (bit 31), the deepest layers ($p16-p32$ sets in Fig. 23) remain unchanged. This is because the convolution weights are very small, and only biases (as depicted in Figs. 27 and 29) contribute to the activations.

Output conv2D layer The biases of the last convolution layer are $[-25983, 2355, -187, 300, -1494, 923]$ (as illustrated in Fig. 28), consisting of 3 positive and 3 negative biases. In the event of a bit-flip occurring in any of the bits within the $[17-30]$ range, the positive biases become more positive, while the negative biases become more negative. Applying equation from Section 4.1.1, the expected error rate for a bit-flip in these positions is:

$$\begin{aligned} \%error_{17-30} &= \frac{1}{6}(0 + 54.91 + 4.28 + 72.42 + 6.91 + 83.86) \\ &= 37.06 \end{aligned}$$

This value is similar to that obtained experimentally by linearly weighting the error rates in the $[17-30]$ range (37.64% in Fig. 23).

Considering that a bit-flip in the sign bit significantly modifies the magnitude, interpreting the results requires an inverse explanation: if the bit-flip occurs in bias 0, the value becomes very positive and will invariably be the winning class, even though it should never win (error of 100%); if it occurs in bias 1, the bias becomes very negative, and class 1 will never win (error of 45.09%). The same reasoning applies to the remaining biases. Hence, the error rate associated with bit-flips in bit 31 should be approximately the complement of the error associated with bit-flips in bit $[17-30]$, i.e., 62.94%. This value is notably smaller than the one obtained experimentally, which is 88%. The discrepancy can be attributed to the uneven distribution of errors injected into each of the six biases in the statistical analysis for bit-flip 31 together with the presence of highly unbalanced terms in the error rate equation for bit-flip 31 (complement values are 100, 45.09, 95.72, 27.58, 93.09 and 16.14).

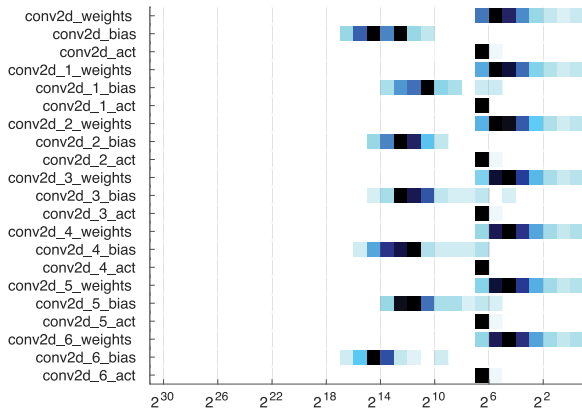


Fig. 30. Calibration study of the pruned quantized model (q_w , q_b and q_a) from layers $conv2D$ to $conv2D_6$.

4.2.2. Analysis of robustness of the pruned model

Relevance of weights compared to biases Despite the decrease in robustness of the pruned model (as illustrated in Fig. 24), the relative negligible importance of weights compared to biases persists.

The biases The pruning process removes numerous irrelevant channels containing high negative biases along the whole model (Table 12). This can be seen by comparing the same figures for the unpruned and pruned case (e.g. Figs. 26 and 30). The only region where this cleanup may not be too pronounced is the central area (as seen in the comparison of Fig. 29 with Fig. 32). However, even in this region, the




Fig. 31. Bit-flip error in $[17-30]$ bits (dashed) and percentage of positive bias (solid) in $conv2D_b$ layers of the pruned quantized model.




Fig. 32. Calibration study of the pruned quantized model (q_w , q_b and q_a) from layers $conv2D_7$ to $conv2D_{12}$.

Table 12

Necessary bits to represent all pos./neg. values of the $conv2D_b$ (green) and $conv2D_{tr}$ (orange) parameters in the quantized pruned model.

Layer	Positive	Negative	Layer	Positive	Negative
0	18	17	22	12	16
1	15	15	21	17	14
2	16	16	20	16	15
3	15	16	4	17	- ^a
4	16	17	19	17	14
5	15	15	18	15	15
6	18	18	3	19	- ^a
7	18	17	17	17	16
8	23	22	16	16	15
9	32	31	2	18	- ^a
10	32	31	15	18	12
11	32	31	14	18	18
0	32	31	1	32	- ^a
12	32	- ^a	13	32	- ^a

^a Means that there are no values of that sign.

activations show a less homogeneous distribution, which indicates that associated weights are more meaningful.

If we compare again the ratio of positive biases of the $conv2D$ layers with bit-flip error rate in the $[17-30]$ range (Fig. 31) the existing correlation between those two quantities for the first six $conv2D$ layers can be observed.

Finally, it is worth noting that, similar to the unpruned case, bit-flips occurring below position 9 do not significantly affect the output.

Similarly, errors become increasingly detrimental to the model from position 16 onwards.

Deep layers: the base of the U-Net The QNN exhibits the same characteristics at the base of the U-Net as the quantized unpruned model. Specifically, the values in $p16-p32$ sets do not generate any errors in the output, regardless of the position of bit-flips. This is a consequence of the small values of convolution weights, and only biases (as depicted in Figs. 32 and 33) contribute significantly to the activations.




Fig. 33. Calibration study of the pruned quantized model (q_w , q_b and q_a) from layers $conv2D_{13}$ to $conv2D_{17}$.

Output $conv2D$ layer The explanation given for the unpruned model remains applicable. Despite having slightly smaller magnitude bias values ($[-19945, 1730, -150, 209, -1120, 665]$), as depicted in Fig. 34, these biases retain their sign bit. The error rates experimentally obtained for bit-flips in positions [17–30] extracted from the linear weighting of the data is 37.89%, consistent with the theoretical error value 37.06%.

$$\begin{aligned} \%error_{17-30} &= \frac{1}{6}(0 + 55.68 + 4.24 + 73.03 + 6.96 + 82.48) \\ &= 37.06 \end{aligned}$$

The error associated with a bit-flip in bit 31 is 93%, exceeding the expected value due to the uneven distribution of bit-flips among the biases.

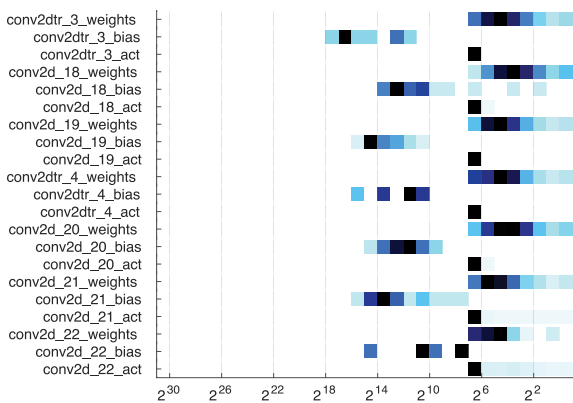


Fig. 34. Calibration study of the pruned quantized model (q_w , q_b and q_a) from layers $conv2D_{18}$ to $conv2D_{22}$.

4.2.3. Analysis of multiple random bit-flips

To estimate error rates under regular operating conditions, we have performed a statistical analysis of MBUs on both quantized unpruned and quantized pruned models. Since FIT (Failures in Time) values reported by manufacturers such as AMD-Xilinx [67] for the most up-to-date FPGA technologies are in the range of tens of FIT/Mb, and given

that targeted BFAs in embedded systems would be very improbable, this study has been performed based on a randomized fault injection campaign.

We randomly injected [1, 10, 50, 100, 250, 400, 650, 800, 950, 1250, 1550, 1750, 2000] bit-flips into the parameters of the entire model and repeated each injection 150 times to obtain average and standard deviation values. Figs. 35 and 36 show the mean multiple bit-flip error rate on the quantized unpruned and pruned DNNs, respectively. As expected, the unpruned model is quite robust since the subset of highly sensitive parameters is small compared to the whole parameter set, so the probability of failure is low compared to that of the pruned model. Nevertheless, it is worth noting that assuming a failure rate of 20 FIT/Mb (based on data from [67]), the unpruned model, with 100 times more parameters (see Table 2), experiences 5 SEUs/hour, whereas the pruned model experiences 0.05 SEUs/hour. Specifically, for models quantized to INT8, the unpruned model would accumulate 1000 upsets in just 8.3 days of operation, whereas the pruned model would require approximately 2.3 years.




Fig. 35. Mean error rate due to multiple bit-flips for the quantized unpruned model. Each purple dot represents the error rate of individual repetitions of the injections. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The results align with those reported by other authors in the context of image classification. For instance, the accuracy loss in [38] for the DNN when exposed to random BFAs is quite similar to that of the unpruned model (Fig. 35). Similarly, the accuracy loss when faults are injected into sensitive parameters is comparable to that of the pruned model (Fig. 36). In this regard, injecting random faults into an optimized and compressed model is analogous to injecting targeted faults into an uncompressed and not optimized model. In [39], it is also noted that for any given network, a very high number of random faults must be injected to significantly degrade inference quality. Finally, [37] identifies several techniques as either beneficial or detrimental to network robustness, among which pruning is detrimental, whereas increasing network capacity is beneficial.

5. Robustness enhancement techniques

As discussed in Section 2.2, neural network hardening methods against single bit-flips can be categorized into four groups: redundancy-based methods, activation function modifications, parameter modifications, and training/inference process modifications. Key comparison

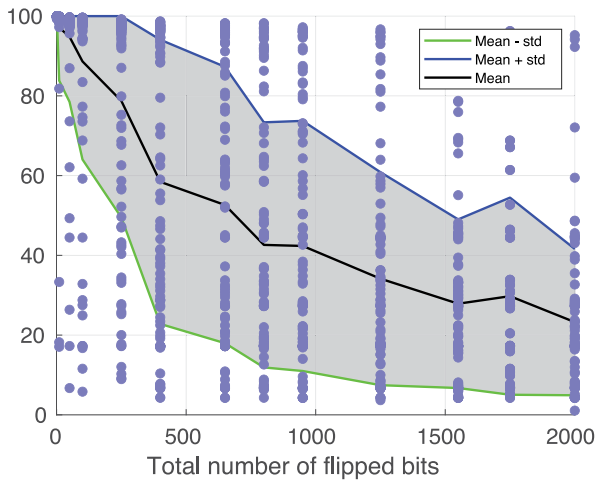


Fig. 36. Mean error rate due to multiple bit-flips for the quantized pruned model. Each purple dot represents the error rate of individual repetitions of the injections. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

features include recovered accuracy, storage overhead, time overhead, and retraining requirements.

Redundancy-based methods, such as triple-modular redundancy [17], introduce significant storage and time overhead. This is due to the need to compare the original model parameters against two additional copies. However, these methods do not require retraining.

Methods that modify activation functions also introduce storage and time overhead, which can be either fixed [56,57] or model-dependent and variable [18,44]. Some methods, such as [58], require the modified activation function to be included during training, necessitating dataset availability and introducing a possible training overhead (the authors report around 6% training overhead).

Methods based on parameter modifications generally do not introduce storage overhead. For instance, [59] does not add storage or time overhead but requires retraining to create a more homogeneous architecture with fewer critical bottlenecks. The MATE method proposed in [60] avoids retraining and memory overhead by replacing non-critical mantissa bits with error correction codes, although it does introduce a time overhead. Similarly, [61] adjusts the parity of weights by flipping the least significant bits as needed, avoiding storage overhead but increasing inference time as faults are masked with zero values.

The requirements for methods that modify the training or inference process vary. The method proposed in [63] relies on adversarial training, requiring dataset availability. In contrast, the method in [64] uses

the inference results of previous frames, adding both storage and time overhead. The method described in [62] has zero overhead, but is only applicable if the CNN ends in a fully connected layer.

Regarding methods that aim to mitigate or mask BFAs, the approach described in [14] requires dataset availability because it involves training additional internal classifiers, which increases storage overhead. However, due to the nature of this method, which promotes early exits, inference time is reported to be reduced to 46.1%–59.4% of the original model’s inference time. In a different approach, the authors of [65] extract a unique signature from the DNN offline and verify the integrity of the model on the fly. This method introduces additional time and storage overheads. While these approaches enhance resilience against faults, they often involve large memory and computation overheads, challenging their suitability for safety-critical DNN inference on edge devices. Therefore, in this work, we aim to enhance robustness without increasing memory or computational costs.

The in detail analysis of error propagation described in the previous sections allows for identifying some key factors that shape the sensitivity of segmentation FCNs to SBUs. As explained, this sensitivity depends mainly on the combination of two factors: the location of the affected parameters in the architecture of the models, and the parameter values (range, sign and their numeric representation). Based on this knowledge, it is possible to establish some design rules to protect the model from being too sensitive to SBUs by avoiding “risky” states while preserving performance. One of such states in single-precision floating-point (IEEE 754 format) is having too many parameters which are one bit-flip away from having the exponent filled, because this representation stands for two especially destructive values: NaN and $\pm\infty$. Consequently, when a SBU causes the exponent to be filled, undesirable results are produced during inference. Similarly, partial exponent filling (leaving aside the MSB) as illustrated in Figs. 13 and 17, also leads to a notable number of errors. This occurs because a number with an absolute value smaller than 1 transforms into a number with an absolute value greater than 1 (but still smaller than 2). Given that weight and bias values typically fall within the range of 0 to 1 (except for the γ values of the BN layers), such situations are common and require attention.

5.1. Modifying model’s parameterization

The proposed protection method involves identifying parameters with exponents containing either seven bits set to ‘1’ or partial exponents with six bits set to ‘1’. This is graphically illustrated on some example values in Fig. 37, where colored in green are the values with an exponent identified as “risky” by the tool. Exponents containing seven ‘1’s (with the MSB set to ‘1’) have not been included, as parameters with such high values are not likely in neural networks.

The proposed method involves incrementing or decrementing the exponent by one, with the aim of making the number of ‘0’s to be




Fig. 37. Values of the partially filled exponents (in green) after increasing (in blue) or decreasing (in orange) its value by 1. The three values which are crossed in red are not useful as either the number of ‘0’-valued positions does not increase or the value of the parameter increases above 2. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

greater than one. This prevents the partial exponent from being filled if a bit-flip event occurs. However, not all the exponents with one '0' are candidates for this method. For example, modifying exponent 01111110 by 1 would not be useful (Fig. 37), as it would either decrement the number of '0's or leave it unchanged. Similarly, incrementing 01111101 would not be useful either, as the number of '0's would remain unchanged. Furthermore, 10000000 has been omitted because bit-flips in the partial exponent would significantly increase the value of the parameter above 2.

After performing the increment/decrement of the exponent, the resultant value of the parameter is doubled/halved. Thus, it has to be compensated by modifying the mantissa. If the exponent is increased, the mantissa value needs to be reduced to the minimum value (1.0), while if the exponent is decreased, the mantissa value needs to be incremented to the maximum value (≈ 1.999). Depending on the original value of the mantissa, the protected value will differ more or less from the original. This modification of the parameter values would produce a perturbation in the model's performance, thus the tool allows for setting upper and lower thresholds for the mantissas to be identified as "full" or "empty" and, thus, candidates for protection. These thresholds define the "protection target" (PT) parameter of the tool. According to this setting, the tool can first evaluate in what extent the performance of the model is modified.

Table 13
Number of candidate parameters to protect for the pruned and unpruned models according to the PT setting.

Protection	Protected parameters			
	Not folded		Folded	
	NotPruned	Pruned	NotPruned	Pruned
PT1	2433	126	721	160
PT2	5993	1095	7497	1551
PT3	21544	5316	36534	7809
PT4	40754	10498	71141	15368

To illustrate how this technique performs on the models under study, four different PTs have been explored: PT1 (1.999, 1.001), PT2 (1.99, 1.01), PT3 (1.95, 1.05), and PT4 (1.9, 1.1). PT0 stands for unprotected models. Table 13 shows the number of candidate parameters selected for protection by the tool according to the protection target.

Table 14
Comparison of IoU values of the original model (PT0) and the protected models (PT1-PT4). BN unpruned unfolded.

Class/Target	PT0	PT1	PT2	PT3	PT4
Road	97.84	97.84	97.85	9.31	16.36
Road marks	87.99	87.99	88.00	45.77	5.50
Vegetation	94.23	94.23	94.24	1.93	12.62
Sky	92.83	92.83	92.85	5.05	2.42
Others	78.12	78.12	78.17	10.35	12.26
Global	94.71	94.71	94.72	8.76	14.20
Weighted	88.54	88.54	88.56	25.23	6.38

Tables 14 to 17 compare the IoU metric on the test set between the original unprotected model and the modified protected ones with different PT settings. Depending on the analyzed model, the degradation of the IoU metric shows different sensitivity to the PT value although, as expected, the higher the PT, the higher the model perturbation.

Table 15
Comparison of IoU values of the original model (PT0) and the protected models (PT1-PT4). BN pruned unfolded.

Class/Target	PT0	PT1	PT2	PT3	PT4
Road	97.73	97.73	97.73	97.73	97.48
Road marks	88.17	88.17	88.16	88.23	88.13
Vegetation	93.66	93.66	93.67	93.76	93.61
Sky	92.89	92.89	92.88	92.77	91.63
Others	77.34	77.34	77.36	77.24	76.46
Global	94.46	94.46	94.46	94.46	94.14
Weighted	88.48	88.48	88.48	88.47	87.97

For the unpruned unfolded model, PT2 is the highest feasible PT (Table 14). For the unpruned folded model, it has to be assessed

whether PT3 is too aggressive or not (Table 16). Regarding pruned models, PT3 is adequate so it remains to be determined whether PT4 is too demanding or not (Tables 15 and 17).

Table 16
Comparison of IoU values of the original model (PT0) and the protected models (PT1-PT4). BN unpruned folded.

Class/Target	PT0	PT1	PT2	PT3	PT4
Road	97.84	97.84	97.84	97.75	97.57
Road marks	87.99	88.00	88.00	86.80	85.96
Vegetation	94.23	94.23	94.23	94.15	93.99
Sky	92.83	92.83	92.83	92.13	90.79
Others	78.12	78.12	78.12	77.12	75.18
Global	94.71	94.71	94.71	94.47	94.05
Weighted	88.54	88.54	88.54	87.61	86.53

To evaluate the suitability of applying a certain PT setting, the following points have been taken into account. As applying this methodology does not either create new potentially dangerous situations or solve the unprotected dangerous ones, PTs are going to be evaluated by injecting a single bit-flip in the "risky" position (so that exponent or partial exponent would be filled) in the parameters selected with each PT. The assessment is based on the following three metrics: the global IoU (GloU), the weighted IoU (WIoU) and the error rate. IoU-based metrics evaluate the model just in the labeled pixels while error rate is a comparison against the prediction of the faultless model so, even in absence of bit-flips, protected models could have a certain error rate. In this section, analysis will be focused only on the results with the maximum achievable PT values.

Table 17
Comparison of IoU values of the original model (PT0) and the protected models (PT1-PT4). BN pruned folded.

Class/Target	PT0	PT1	PT2	PT3	PT4
Road	97.73	97.73	97.72	97.73	97.78
Road marks	88.17	88.17	88.16	88.08	87.89
Vegetation	93.66	93.66	93.66	93.61	93.56
Sky	92.89	92.89	92.90	92.65	92.03
Others	77.34	77.33	77.31	77.02	76.98
Global	94.46	94.46	94.45	94.40	94.38
Weighted	88.48	88.48	88.48	88.32	88.05

5.1.1. Protection of the unpruned unfolded model

Applying PT2 yields notable benefits (see Table 18), as the error rate is significantly reduced while the IoU metrics, especially for bit 26, show improvement. There is a scenario where the three indicators do not align (bit 25): despite minor improvements in the IoU metrics, the error rate considerably increases for the protected model. This could be attributed to the fact that, unintentionally, in the absence of faults, the protected model outperforms the non-protected model (see Table 14), indicating a considerable number of pixels that change their value even before bit-flip injection.

Table 18
Comparison of the non-protected and PT2 protected unpruned unfolded models on potentially dangerous bit positions.

Non-protected				Protected			
Bit	GloU	WIoU	% error	Bit	GloU	WIoU	% error
30	89.07	80.06	7.76	30	90.63	82.69	0.72
29	92.30	85.23	0.32	29	92.31	85.28	0.04
28	92.30	85.23	1.02	28	92.31	85.28	0.04
27	92.23	85.14	1.55	27	92.29	85.25	0.08
26	90.12	81.81	4.43	26	91.65	83.90	1.36
25	83.04	73.08	5.54	25	83.11	73.17	10.01
24	92.22	85.09	1.86	24	92.23	85.13	0.54

5.1.2. Protection of the pruned unfolded model

Analyzing PT4 for the pruned unfolded model (Table 19), we can see that a straightforward conclusion cannot be extracted as the benefits of the protection are not shared among all bits. As a consequence, it is necessary to assess how the protected model behaves beyond the

labeled pixels, since erroneous pixels may be situated in regions that are not critical. Otherwise, protection target should be lowered down to 3, as it has been verified that the protected model outperforms the non-protected one.

Table 19

Comparison of the non-protected and PT4 protected pruned unfolded models on potentially dangerous bit positions.

Non-protected				Protected			
Bit	GIoU	WIoU	% error	Bit	GIoU	WIoU	% error
30	0.00	0.00	100	30	26.03	24.52	79.50
29	92.48	84.50	0.00	29	91.93	84.11	1.61
28	92.48	84.50	0.00	28	91.93	84.11	1.61
27	90.83	82.33	2.35	27	91.55	83.71	2.14
26	89.82	80.76	3.51	26	90.92	82.51	2.93
25	78.25	67.49	14.58	25	79.24	68.26	13.98
24	92.32	84.35	1.03	24	91.83	83.89	2.12

5.1.3. Protection of the unpruned folded model

In the unpruned folded model, PT2 is the highest achievable one, albeit with unimpressive results, as there is neither loss nor significant improvement (see Table 20). Nevertheless, despite the lack of remarkable changes, this PT still yields benefits, as it increases the mean WIoU and reduces the mean error rates. Thus, its application keeps being beneficial for the model.

Table 20

Comparison of the non-protected and PT2 protected unpruned folded models on potentially dangerous bit positions.

Non-protected				Protected			
Bit	GIoU	WIoU	% error	Bit	GIoU	WIoU	% error
30	10.25	9.48	88.91	30	31.11	26.58	65.23
29	92.30	85.23	0.00	29	92.29	85.23	0.01
28	92.30	85.23	0.00	28	92.30	85.23	0.02
27	92.30	85.23	0.00	27	92.29	85.23	0.02
26	92.30	85.23	0.05	26	92.29	85.23	0.03
25	92.27	85.15	0.43	25	92.28	85.20	0.31
24	92.29	85.15	0.69	24	92.29	85.15	0.69

5.1.4. Protection of the pruned folded model

In the pruned folded model, the highest achievable PT is 3 (Table 21). Even though the error rate slightly increases (but not more than 0.31%), positive differences in terms of GIoU and WIoU are found for all bit positions. After applying PT4, the improvements in WIoU disappear, while the error rate still increases.

Table 21

Comparison of the non-protected and PT3 protected pruned folded models on potentially dangerous bit positions.

Non-protected				Protected			
Bit	GIoU	WIoU	% error	Bit	GIoU	WIoU	% error
30	0.00	0.00	100	30	23.84	18.86	72.09
27	92.48	84.50	0.08	27	92.54	84.54	0.39
26	92.44	84.43	0.27	26	92.53	84.52	0.43
25	92.39	84.35	0.55	25	92.52	84.49	0.61
24	92.34	84.33	0.97	24	92.40	84.34	1.07

After assessing all the models, it can be concluded that the target from which to start the protection can be deduced just by inspecting the results of Tables 14 to 17. For all the models, this target has been at least 2. Depending on the parameters of the model, a PT of 3, 4, or even higher may still be beneficial. As this method is completely memory and computation-free, it can be combined with memory-hungry methods, such as modular redundancy, which would focus on protecting just the critical parameters that have not been protected with this method.

5.2. Additional protection techniques: revisiting model's parameterization

Additional protection techniques to that proposed in Section 5.1 can be applied to enhance model protection with no memory and computational overhead. In contrast to the previous method, the following

methods exclusively focus on the redefinition of the model parameters. This adjustments ensures that, in the absence of bit-flips, the model retains its original behavior while preventing potentially critical value changes due to bit-flip events.

5.2.1. Applying sparse pruning

Applying sparse pruning can be beneficial in terms of robustness regardless of whether the processor is designed to improve inference performance in the presence of sparse matrices. In case of identifying that the model has parameters in the range $1 \leq |x| < 2$ that are not relevant for inference, applying sparse pruning (setting irrelevant parameters to 0) prevents the appearance of *NaN* values in case of bit-flips.

5.2.2. Redefining batch normalization parameters

If folding of batch normalization layers, which can be viewed as depthwise convolutions (see Eq. (7)), is not applied, γ , σ , μ and β parameters can be reconditioned as long as the weight (W) and bias (b) values do not change. Parameter reconditioning can be applied with two aims: either directly reduce the amount of parameters at risk of having the exponent filled or just modify their mantissa values so that the protection method described in the previous section can be applied.

$$BN(x) = \gamma \frac{(x - \mu)}{\sqrt{\sigma}} + \beta = \frac{\gamma}{\sqrt{\sigma}} x + (\beta - \gamma \frac{\mu}{\sqrt{\sigma}}) = Wx + b \quad (7)$$

5.2.3. Applying cross layer equalization

This method, which is similar to CLE [75] for quantization, allows for the protection of convolutional kernels, which contain most of the parameters in the model. The method is based on the positive scaling equivariance property ($f(sx) = sf(x)$) which is held for models containing Parametric Rectified Linear Unit (PReLU) functions. The activation map y of a certain layer n can be calculated as $f_n(W_n \cdot x + b_n)$ where W_n and b_n are the weights and biases of the convolutional kernel applied to layer n and x is the input cube. If we calculate the activation map of the next layer $n+1$ in terms of the previous layer n , Eq. (8) shows that model parameters can be reconditioned as long as the activation function of the n th layer fulfills the equivariance property (e.g. ReLU and Identity).

$$\begin{aligned} z &= f_{n+1}(W_{n+1} \cdot f_n(W_n \cdot x + b_n) + b_{n+1}) \\ &= f_{n+1}(W_{n+1} \cdot S \cdot \hat{f}_n(S^{-1} \cdot W_n \cdot x + S^{-1} \cdot b_n) + b_{n+1}) \\ &= f_{n+1}(\hat{W}_{n+1} \cdot f_n(\hat{W}_n \cdot x + \hat{b}_n) + b_{n+1}) \end{aligned} \quad (8)$$

where S is a diagonal matrix containing the scale factor of the parameters for each of the filters of a convolutional layer.

5.2.4. Applying bias absorption

Bias absorption [75], which is applied during post training quantization procedures, can be implemented with a different purpose in this domain: reducing the number of critical parameters in the model. The only requirement in order to apply this method is that $f(Wx + b - c) = f(Wx + b) - c$ must be fulfilled. It is straightforward to see that this condition will not be met by any non-linear activation function, will always be met by identity and will be met by ReLU activations just in certain situations ($c = \max(0, Wx + b) \forall x$) that are closely related to x data distribution and W and b parameters. The method is described in Eq. (9), where $\hat{b}_n = b_n - c$ and $\hat{b}_{n+1} = W_{n+1} * c + b_{n+1}$.

$$\begin{aligned} z &= f_{n+1}(W_{n+1} \cdot f_n(W_n \cdot x + b_n) + b_{n+1}) \\ &= f_{n+1}(W_{n+1} \cdot (f_n(W_n \cdot x + b_n) + c - c) + b_{n+1}) \\ &= f_{n+1}(W_{n+1} \cdot (f_n(W_n \cdot x + \hat{b}_n) + c) + b_{n+1}) \\ &= f_{n+1}(W_{n+1} \cdot f_n(W_n \cdot x + \hat{b}_n) + \hat{b}_{n+1}) \end{aligned} \quad (9)$$

6. Conclusions

This article analyzes the robustness of encoder–decoder type FCNs against SBUs in semantic segmentation tasks and how compression techniques applied for deployments in embedded systems alter this robustness. The analysis is based on a statistical campaign of software-injected SBUs in different network parameters to conduct a layer-by-layer and bit-by-bit in-depth study of the sources of critical errors. The reference models used in this study were designed and optimized for hyperspectral image segmentation with application to autonomous driving systems, and were trained on the HSI-Drive v2.0 dataset. We identified the potentially most problematic bit-flips according to the location of the parameters in the model architecture, their binary representation, range and sign, and relevance to the inference process. From this analysis, a tool for estimating the model’s robustness and a simple memory and computation-free error mitigation method is proposed.

This study reveals that in 32-bit floating-point FCN models with lower-bounded activation functions, such as the commonly used ReLU, critical errors are primarily associated with SBUs that result in an increment in the parameter values. This is because, as noted by other authors, commonly used machine learning computations are often monotonic. However, not all increments pose the same risk. The most sensitive bit is the MSB of the exponent, as its bit-flip significantly increases the model parameters’ values from below 2 to ∞ or *NaN*, in the worst-case scenario, propagating the error throughout the model. Given the typical values of the parameters in BN layers, this layer is the most sensitive to such bit-flips. Furthermore, we have found that this layer extends the range of output values, thus contributing negatively to the propagation of errors that may have occurred in previous layers. In this regard, it is noted that layers with many positive biases are particularly vulnerable. Another particularly “risky” situation arises from parameters with almost filled partial exponents. When a partial exponent is filled by a bit-flip, the parameter undergoes a transition from a value significantly below unity to a new value above unity, leading to a noticeable amount of errors. Due to the specific characteristics of encoder–decoder type FCNs with skip connections, SBUs in the first layers of the encoder branch and in the last layers of the decoder branch, which are closer to the model’s output, produce comparatively higher error rates. On the other hand, it is observed that perturbations in the parameters at the base of the model have little influence on the output.

Regarding compression techniques, pruned models show higher sensibility to SBUs than unpruned models. This is attributed to the over-parameterized and parallel nature of original unpruned DNN models. However, the primary advantage of pruned models lies in their smaller size, as it is less likely for an SBU to occur in a model with fewer parameters. Hence, to accurately assess the influence of pruning on the vulnerability to SBUs, model implementation details such as circuit design and selected target device should be considered.

With respect to data representation, integer-quantized models show higher robustness to bit-flips. This is due to the particularities of integer binary representation, since there is no option to extreme values such as *NaNs* or ∞ s to occur. The employed quantization method, where weights are quantized to 8 bits while biases are quantized to 32 bits, makes the latter more sensitive to SBUs. Furthermore, it is generally observed that the upper word of the bias representation simply serves to extend the sign, so it could be used to implement error detection and/or correction methods with no memory overhead. Finally, since the model sensitivity is fundamentally linked to perturbations in the biases and these constitute a tiny subset of the whole parameter-set, memory-intensive error mitigation techniques such as duplication or triple modular redundancy applied only to biases would ensure high model protection with low memory burden.

When combining pruning and quantization it was observed that the fully compressed model is more robust than the pruned unquantized

model but less robust than the unpruned quantized model. It can be concluded that integer quantization always contributes to robustness enhancement, whereas the evaluation of the effects of pruning is tied to the applied compression degree and the final implementation of the model.

Concerning the robustness analysis of quantized models, if instead of a targeted fault injection technique random multiple fault injection is performed, it is observed that the unpruned model exhibits a considerably greater robustness. This is because, in quantized models, the sensitive parameters are the biases, and not the weights, which constitute a small subset of the total parameters, reducing the likelihood of a random bit-flip occurring in them. Additionally, due to the network architecture, most parameters are concentrated in the base of the U-Net, where the least critical layers are located. As a result, after pruning, the robustness against bit-flips is significantly reduced since nearly any bit-flip can affect a sensitive parameter. Nonetheless, assuming a failure rate of 20 FiT/Mb, it is worth noting that the unpruned model experiences 5 SEUs/hour whereas the pruned model experiences 0.05 SEUs/hour. This implies that the unpruned quantized model would need just 8.3 days of operation to accumulate 1000 upsets, whereas the pruned model would require approximately 2.3 years.

Finally, based on the performed analysis, a set of complementary protection methods is proposed to mitigate some delicate situations like partial exponent completion. In particular, we describe a new memory and computation overhead-free method for the protection of certain model parameters in floating-point representation. This technique enables the selection of various protection targets, where a higher target involves attempting to protect more parameters in the model. However, setting too ambitious protection targets may result in an excessive modification of the original parameterization, thus a model performance analysis must be accomplished beforehand.

CRedit authorship contribution statement

Jon Gutiérrez-Zaballa: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Koldo Basterretxea:** Writing – review & editing, Writing – original draft, Validation, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Data curation, Conceptualization. **Javier Echanobe:** Resources, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability


I have shared the links for my code in the abstract and the dataset is available at <https://ipaccess.ehu.es/HSI-Drive/> (also included in the references)

References

- [1] Soledad Le Clairche, Esteban Ferrer, Sam Gibson, Elisabeth Cross, Alessandro Parente, Ricardo Vinuesa, Improving aircraft performance using machine learning: A review, *Aerosp. Sci. Technol.* (2023) 108354.
- [2] SAE International Aerospace, Guidelines for development of civil aircraft and systems, 2010, URL <https://www.sae.org/standards/content/arp4754a/>. Revision A.
- [3] International Organization for Standardization, Road vehicles – Functional safety, 2018, URL <https://www.iso.org/standard/68383.html>.
- [4] Corinna Martinella, Rubén G Alía, Roger Stark, Andrea Coronetti, Carlo Cazaniga, Maria Kastriotou, Yacine Kadi, Rémi Gaillard, Ulrike Grossner, Arto Javanainen, Impact of terrestrial neutrons on the reliability of SiC VD-MOSFET technologies, *IEEE Trans. Nucl. Sci.* 68 (5) (2021) 634–641.

- [5] Robert Baumann, Soft errors in advanced computer systems, *IEEE Des. Test Comput.* 22 (3) (2005) 258–266.
- [6] Soichi Hirokawa, Ryo Harada, Kenshiro Sakuta, Yukinobu Watanabe, Masanori Hashimoto, Multiple sensitive volume based soft error rate estimation with machine learning, in: 2016 16th European Conference on Radiation and Its Effects on Components and Systems, RADECS, IEEE, 2016, pp. 1–4.
- [7] Zheyu Yan, Yiyu Shi, Wang Liao, Masanori Hashimoto, Xichuan Zhou, Cheng Zhuo, When single event upset meets deep neural networks: Observations, explorations, and remedies, in: 2020 25th Asia and South Pacific Design Automation Conference, ASP-DAC, IEEE, 2020, pp. 163–168.
- [8] Wenxiao Wang, Song Gao, Yaqi Wang, Yang Li, Wenjing Yue, Hongsen Niu, Feifei Yin, Xunjian Guo, Guozhen Shen, Advances in emerging photonic memristive and memristive-like devices, *Adv. Sci.* 9 (28) (2022) 2105577.
- [9] Jiangwei Zhang, Chong Wang, Yi Cai, Zhenhua Zhu, Donald Kline, Huazhong Yang, Yu Wang, Wesco: Weight-encoded reliability and security co-design for in-memory computing systems, in: 2022 IEEE Computer Society Annual Symposium on VLSI, ISVLSI, IEEE, 2022, pp. 296–301.
- [10] Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, Onur Mutlu, Flipping bits in memory without accessing them: An experimental study of DRAM disturbance errors, *ACM SIGARCH Comput. Archit. News* 42 (3) (2014) 361–372.
- [11] Varyeon Kim, Vikas Chandra, Robert Aitken, David Blaauw, Dennis Sylvester, Variation-aware static and dynamic writability analysis for voltage-scaled bit-interleaved 8-t SRAMs, in: *IEEE/ACM International Symposium on Low Power Electronics and Design*, IEEE, 2011, pp. 145–150.
- [12] Kaveh Razavi, Ben Gras, Erik Bosman, Bart Preneel, Cristiano Giuffrida, Herbert Bos, Flip feng shui: Hammering a needle in the software stack, in: 25th USENIX Security Symposium, USENIX Security 16, 2016, pp. 1–18.
- [13] Fan Yao, Adnan Siraj Rakin, Deliang Fan, {DeepHammer}: Depleting the intelligence of deep neural networks through targeted chain of bit flips, in: 29th USENIX Security Symposium, USENIX Security 20, 2020, pp. 1463–1480.
- [14] Jialai Wang, Ziyuan Zhang, Meiqi Wang, Han Qiu, Tianwei Zhang, Qi Li, Zongpeng Li, Tao Wei, Chao Zhang, Aegis: Mitigating targeted bit-flip attacks against deep neural networks, in: 32nd USENIX Security Symposium, USENIX Security 23, 2023, pp. 2329–2346.
- [15] Wenhai Wang, Jifeng Dai, Zhe Chen, Zhenhang Huang, Zhiqi Li, Xizhou Zhu, Xiaowei Hu, Tong Lu, Lewei Lu, Hongsheng Li, Xiaogang Wang, Yu Qiao, InternImage: Exploring large-scale vision foundation models with deformable convolutions, 2023, [arXiv:2211.05778](https://arxiv.org/abs/2211.05778).
- [16] M. Violante, L. Sterpone, M. Ceschia, D. Bortolato, P. Bernardi, M.S. Reorda, A. Paccagnella, Simulation-based analysis of SEU effects in SRAM-based FPGAs, *IEEE Trans. Nucl. Sci.* 51 (6) (2004) 3354–3359, <http://dx.doi.org/10.1109/TNS.2004.839516>.
- [17] Annachiara Ruospo, Gabriele Gavarini, Ilaria Bragaglia, Marcello Traiola, Alberto Bosio, Ernesto Sanchez, Selective hardening of critical neurons in deep neural networks, in: 2022 25th International Symposium on Design and Diagnostics of Electronic Circuits and Systems, DDECS, IEEE, 2022, pp. 136–141.
- [18] Mahdi Taheri, Natalia Cherezova, Mohammad Saeed Ansari, Maksim Jenihhin, Ali Mahani, Masoud Daneshalab, Jaan Raik, Exploration of activation fault reliability in quantized systolic array-based DNN accelerators, 2024, [arXiv:2401.09509](https://arxiv.org/abs/2401.09509).
- [19] Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W Mahoney, Kurt Keutzer, A survey of quantization methods for efficient neural network inference, in: *Low-Power Computer Vision*, Chapman and Hall/CRC, 2022, pp. 291–326.
- [20] Jon Gutiérrez-Zaballa, Parameterprotection, 2024, URL <https://github.com/jonGuti13/parameterProtection>.
- [21] Jon Gutiérrez-Zaballa, TensorFI2, 2024, URL <https://github.com/jonGuti13/TensorFI2>.
- [22] Niranjhana Narayanan, Zitao Chen, Bo Fang, Guanpeng Li, Karthik Pattabiraman, Nathan DeBardeleben, Fault injection for TensorFlow applications, *IEEE Trans. Dependable Secure Comput.* 20 (4) (2023) 2677–2695, <http://dx.doi.org/10.1109/TDSC.2022.3175930>.
- [23] Mohammad Hasan AhmadiIvani, Mahdi Taheri, Jaan Raik, Masoud Daneshalab, Maksim Jenihhin, A systematic literature review on hardware reliability assessment methods for deep neural networks, *ACM Comput. Surv.* (ISSN: 0360-0300) (2023) <http://dx.doi.org/10.1145/3638242>, Just Accepted.
- [24] Austin P. Arechiga, Alan J. Michaels, The robustness of modern deep learning architectures against single event upset errors, in: 2018 IEEE High Performance Extreme Computing Conference, HPEC, IEEE, 2018, pp. 1–6.
- [25] Elaheh Malekzadeh, Nezam Rohbani, Zhonghai Lu, Masoumeh Ebrahimi, The impact of faults on DNNs: A case study, in: 2021 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems, DFT, IEEE, 2021, pp. 1–6.
- [26] Austin P. Arechiga, Alan J. Michaels, The effect of weight errors on neural networks, in: 2018 IEEE 8th Annual Computing and Communication Workshop and Conference, CCWC, IEEE, 2018, pp. 190–196.
- [27] Mohamed A Neggaz, Ihsen Alouani, Smail Niar, Fadi Kurdahi, Are cnns reliable enough for critical applications? an exploratory study, *IEEE Des. Test* 37 (2) (2019) 76–83.
- [28] Majid Sabbagh, Cheng Gongye, Yunsi Fei, Yanzhi Wang, Evaluating fault resiliency of compressed deep neural networks, in: 2019 IEEE International Conference on Embedded Software and Systems, ICCESS, IEEE, 2019, pp. 1–7.
- [29] Bruno F Goldstein, Sudarshan Srinivasan, Dipankar Das, Kunal Banerjee, Leandro Santiago, Victor C Ferreira, Alexandre S Nery, Sandip Kundu, Felipe MG França, Reliability evaluation of compressed deep learning models, in: 2020 IEEE 11th Latin American Symposium on Circuits & Systems, LASCAS, IEEE, 2020, pp. 1–5.
- [30] Alberto Bosio, Paolo Bernardi, Annachiara Ruospo, Ernesto Sanchez, A reliability analysis of a deep neural network, in: 2019 IEEE Latin American Test Symposium, LATS, IEEE, 2019, pp. 1–6.
- [31] Annachiara Ruospo, Gabriele Gavarini, Corrado De Sio, J Guerrero, Luca Sterpone, M Sonza Reorda, Ernesto Sánchez, Riccardo Mariani, J Aribido, Jyotika Athavale, Assessing convolutional neural networks reliability through statistical fault injections, in: 2023 Design, Automation & Test in Europe Conference & Exhibition, DATE, IEEE, 2023, pp. 1–6.
- [32] Sanghyun Hong, Pietro Frigo, Yiğitcan Kaya, Cristiano Giuffrida, Tudor Dumitraş, Terminal brain damage: Exposing the graceless degradation in deep neural networks under hardware fault attacks, in: 28th USENIX Security Symposium, USENIX Security 19, 2019, pp. 497–514.
- [33] Niranjhana Narayanan, Fault injection in Machine Learning applications (Ph.D. thesis), University of British Columbia, 2021.
- [34] Annachiara Ruospo, Alberto Bosio, Alessandro Ianne, Ernesto Sanchez, Evaluating convolutional neural networks reliability depending on their data representation, in: 2020 23rd Euromicro Conference on Digital System Design, DSD, IEEE, 2020, pp. 672–679.
- [35] Annachiara Ruospo, Ernesto Sanchez, Marcello Traiola, Ian O'connor, Alberto Bosio, Investigating data representation for efficient and reliable convolutional neural networks, *Microprocess. Microsyst.* 86 (2021) 104318.
- [36] Rizwan Tariq Syed, Markus Ulbricht, Krzysztof Piotrowski, Milos Krstic, Fault resilience analysis of quantized deep neural networks, in: 2021 IEEE 32nd International Conference on Microelectronics, MIEL, IEEE, 2021, pp. 275–279.
- [37] Zhezhi He, Adnan Siraj Rakin, Jingtao Li, Chaitali Chakrabarti, Deliang Fan, Defending and harnessing the bit-flip based adversarial weight attack, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 14095–14103.
- [38] Jingtao Li, Adnan Siraj Rakin, Yan Xiong, Liangliang Chang, Zhezhi He, Deliang Fan, Chaitali Chakrabarti, Defending bit-flip attack through dnn weight reconstruction, in: 2020 57th ACM/IEEE Design Automation Conference, DAC, IEEE, 2020, pp. 1–6.
- [39] Adnan Siraj Rakin, Zhezhi He, Deliang Fan, Bit-flip attack: Crushing neural network with progressive bit search, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 1211–1220.
- [40] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, Wojciech Samek, On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation, *PLoS One* 10 (7) (2015) e0130140.
- [41] Wonseok Choi, Dongyeob Shin, Jongsun Park, Swaroop Ghosh, Sensitivity based error resilient techniques for energy efficient deep neural network accelerators, in: *Proceedings of the 56th Annual Design Automation Conference* 2019, 2019, pp. 1–6.
- [42] Yangchao Zhang, Hiroaki Itsuji, Takumi Uezono, Tadanobu Toba, Masanori Hashimoto, Estimating vulnerability of all model parameters in dnn with a small number of fault injections, in: 2022 Design, Automation & Test in Europe Conference & Exhibition, DATE, IEEE, 2022, pp. 60–63.
- [43] Zitao Chen, Guanpeng Li, Karthik Pattabiraman, Nathan DeBardeleben, Binfi: An efficient fault injector for safety-critical machine learning systems, in: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2019, pp. 1–23.
- [44] Jinyu Zhan, Ruoxu Sun, Wei Jiang, Yucheng Jiang, Xunzhao Yin, Cheng Zhuo, Improving fault tolerance for reliable DNN using boundary-aware activation, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* 41 (10) (2021) 3414–3425.
- [45] Stéphane Burelt, Adrian EvansT, Lorena Anghel, Improving DNN fault tolerance in semantic segmentation applications, in: 2022 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems, DFT, IEEE, 2022, pp. 1–6.
- [46] Giuseppe Esposito, Reliability Evaluation of Split Computing Neural Networks (Ph.D. thesis), Politecnico di Torino, 2023.
- [47] G. Govarini, Annachiara Ruospo, E. Sanchez, A fast reliability analysis of image segmentation neural networks exploiting statistical fault injections, in: 2023 IEEE 24th Latin American Test Symposium, LATS, IEEE, 2023, pp. 1–6.
- [48] Davide Sabena, M Sonza Reorda, Luca Sterpone, Paolo Rech, Luigi Carro, On the evaluation of soft-errors detection techniques for GPGPUs, in: 2013 8th IEEE Design and Test Symposium, IEEE, 2013, pp. 1–6.
- [49] Daniel AG Oliveira, Paolo Rech, Heather M Quinn, Thomas O Fairbanks, Laura Monroe, Sarah E Michalak, Christine Anderson-Cook, Philippe OA Navaux, Luigi Carro, Modern GPUs radiation sensitivity evaluation and mitigation through duplication with comparison, *IEEE Trans. Nucl. Sci.* 61 (6) (2014) 3115–3122.

- [50] Lucas Weigel, Fernando Fernandes, Philippe Navaux, Paolo Rech, Kernel vulnerability factor and efficient hardening for histogram of oriented gradients, in: 2017 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems, DFT, IEEE, 2017, pp. 1–6.
- [51] Fabiano Libano, Brittany Wilson, J Anderson, Michael J Wirthlin, Carlo Cazaniga, Christopher Frost, Paolo Rech, Selective hardening for neural networks in FPGAs, IEEE Trans. Nucl. Sci. 66 (1) (2018) 216–222.
- [52] Cristiana Bolchini, Luca Cassano, Antonio Miele, Alessandro Nazzari, Selective hardening of cnns based on layer vulnerability estimation, in: 2022 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems, DFT, IEEE, 2022, pp. 1–6.
- [53] Fernando Fernandes dos Santos, Pedro Foletto Pimenta, Caio Lunardi, Lucas Draghetti, Luigi Carro, David Kaeli, Paolo Rech, Analyzing and increasing the reliability of convolutional neural networks on GPUs, IEEE Trans. Reliab. 68 (2) (2018) 663–677.
- [54] Javier Fernández, Irune Agirre, Jon Perez-Cerrolaza, Jaume Abella, Francisco J Cazorla, A methodology for selective protection of matrix multiplications: A diagnostic coverage and performance trade-off for CNNs executed on GPUs, in: 2022 6th International Conference on System Reliability and Safety, ICSRS, IEEE, 2022, pp. 9–18.
- [55] Fernando F dos Santos, Marcelo Brandalero, Michael B Sullivan, Pedro M Basso, Michael Hübner, Luigi Carro, Paolo Rech, Reduced precision dwc: An efficient hardening strategy for mixed-precision architectures, IEEE Trans. Comput. 71 (3) (2021) 573–586.
- [56] Le-Ha Hoang, Muhammad Abdullah Hanif, Muhammad Shafique, Ft-clipact: Resilience analysis of deep neural networks and improving their fault tolerance using clipped activation, in: 2020 Design, Automation & Test in Europe Conference & Exhibition, DATE, IEEE, 2020, pp. 1241–1246.
- [57] Zitao Chen, Guanpeng Li, Karthik Pattabiraman, A low-cost fault corrector for deep neural networks through range restriction, in: 2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN, IEEE, 2021, pp. 1–13.
- [58] Behnam Ghavami, Mani Sadati, Zhenman Fang, Lesley Shannon, FitAct: Error resilient deep neural networks via fine-grained post-trainable activation functions, in: 2022 Design, Automation & Test in Europe Conference & Exhibition, DATE, IEEE, 2022, pp. 1239–1244.
- [59] Christoph Schorn, Andre Guntoro, Gerd Ascheid, An efficient bit-flip resilience optimization method for deep neural networks, in: 2019 Design, Automation & Test in Europe Conference & Exhibition, DATE, IEEE, 2019, pp. 1507–1512.
- [60] Myeongjae Jang, Jeongkyu Hong, MATE: Memory-and retraining-free error correction for convolutional neural network weights, J. Inf. Commun. Convergence Eng. 19 (1) (2021).
- [61] Stéphane Burel, Adrian Evans, Lorena Anghel, Zero-overhead protection for cnn weights, in: 2021 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems, DFT, IEEE, 2021, pp. 1–6.
- [62] Suyong Lee, Insu Choi, Joon-Sung Yang, Bipolar vector classifier for fault-tolerant deep neural networks, in: Proceedings of the 59th ACM/IEEE Design Automation Conference, 2022, pp. 673–678.
- [63] Giulio Gambardella, Nicholas J Fraser, Ussama Zahid, Gianluca Furano, Michaela Blott, Accelerated radiation test on quantized neural networks trained with fault aware training, in: 2022 IEEE Aerospace Conference, AERO, IEEE, 2022, pp. 1–7.
- [64] Lucas Klein Draghetti, Fernando Fernandes dos Santos, Luigi Carro, Paolo Rech, Detecting errors in convolutional neural networks using inter frame spatio-temporal correlation, in: 2019 IEEE 25th International Symposium on on-Line Testing and Robust System Design, IOLTS, IEEE, 2019, pp. 310–315.
- [65] Mojan Javaheripi, Farinaz Koushanfar, HASHTAG: Hash signatures for online detection of fault-injection attacks on deep neural networks, in: 2021 IEEE/ACM International Conference on Computer Aided Design, ICCAD, 2021, pp. 1–9, <http://dx.doi.org/10.1109/ICCAD51958.2021.9643556>.
- [66] Jon Gutiérrez-Zaballa, Koldo Basterretxea, Javier Echanobe, M Victoria Martínez, Unai Martínez-Corral, HSI-drive v2. 0: More data for new challenges in scene understanding for autonomous driving, in: 2023 IEEE Symposium Series on Computational Intelligence, SSCI, IEEE, 2023, pp. 207–214.
- [67] AMD-Xilinx, Device reliability report. UG116 v10.17., 2023, <https://docs.amd.com/r/en-US/ug116>.
- [68] Haonan Tian, Younis Ibrahim, Rui Chen, Yixiu Wang, Chen JIN, George Belev, Li Chen, Comparative study: AutoDPR-SEM for enhancing CNN reliability in SRAM-based FPGAs through autonomous reconfiguration, Microelectron. Reliab. (ISSN: 0026-2714) 157 (2024) 115392, <http://dx.doi.org/10.1016/j.micrel.2024.115392>, URL <https://www.sciencedirect.com/science/article/pii/S0026271424000726>.
- [69] Benoit Jacob, Skirmantas Klīgys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, Dmitry Kalenichenko, Quantization and training of neural networks for efficient integer-arithmetic-only inference, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 2704–2713.
- [70] Maarten Vandersteegen, Kristof Van Beeck, Toon Goedemé, Integer-only cnns with 4 bit weights and bit-shift quantization scales at full-precision accuracy, Electronics 10 (22) (2021) 2823.
- [71] Haihao Shen, Naveen Mellempudi, Xin He, Qun Gao, Chang Wang, Mengni Wang, Efficient post-training quantization with FP8 formats, 2023, arXiv preprint [arXiv:2309.14592](https://arxiv.org/abs/2309.14592).
- [72] Hao Wu, Patrick Judd, Xiaojie Zhang, Mikhail Isaev, Paulius Micikevicius, Integer quantization for deep learning inference: Principles and empirical evaluation, 2020, arXiv preprint [arXiv:2004.09602](https://arxiv.org/abs/2004.09602).
- [73] R. Leveugle, A. Calvez, P. Maistri, P. Vanhauwaert, Statistical fault injection: Quantified error and confidence, in: 2009 Design, Automation & Test in Europe Conference & Exhibition, 2009, pp. 502–506, <http://dx.doi.org/10.1109/DATF.2009.5090716>.
- [74] Michael Zhu, Suyog Gupta, To prune, or not to prune: exploring the efficacy of pruning for model compression, 2017, arXiv preprint [arXiv:1710.01878](https://arxiv.org/abs/1710.01878).
- [75] Markus Nagel, Mart van Baalen, Tijmen Blankevoort, Max Welling, Data-free quantization through weight equalization and bias correction, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 1325–1334.



Jon Gutiérrez-Zaballa received the Bachelor's degree in Physics and Electronic Engineering from the University of the Basque Country in 2019 and the M.Sc. on Advanced Electronic Systems at the same university in 2021. He is currently working as a Ph.D. candidate in Physical Engineering at the Department of Electronics Technology of the University of the Basque Country. His current research interests mainly concern image processing, computational intelligence (data mining, Artificial Neural Networks, Convolutional Neural Networks) and software–hardware embedded co-design on PSoCs for Advanced Driver Assistance Systems.



Koldo Basterretxea received the Licenciado degree in physics (M.Sc.) with specialization in electronics and automatic control and the Ph.D. degree in physics from the University of the Basque Country (UPV/EHU), Spain, in 1994 and 2002, respectively. Currently, he is with the Department of Electronics Technology, Faculty of Engineering of Bilbao, UPV/EHU, where he has been an Associate Professor since 2009 and coordinates the Master's Program in Advanced Electronic Systems. Dr. Basterretxea has been a visiting researcher at the University of Edinburgh School of Engineering in 2009 and 2012. His research interests include design of adaptive systems on FPGAs, hardware acceleration of numerical matrix algebra and optimization algorithms, and design of application and domain specific processors for high-speed embedded control, machine learning and computational intelligence applications.



Javier Echanobe (M'06) received the Licenciado degree in physics from the University of the Basque Country (UPV/EHU), Spain, and the Ph.D. degree from the University of Navarra, Pamplona, Spain, in 1990 and 1998, respectively. He was a Predoctoral Researcher (granted by the Basque Government) from 1992 to 1996. He has been an Associate Professor in the Department of Electricity and Electronics, UPV/EHU, from 1999 to 2009. Since 2009 he is a Senior Lecturer in that department. His research interests focus on (1) digital electronics: embedded systems, reconfigurable FPGAs, DSPs, SoPC; (2) computational intelligence: artificial neural networks, fuzzy systems, and genetic algorithms; (3) ubiquitous computing: ambient intelligence, intelligent environments, intelligent vehicles. Dr. Echanobe has published many papers in international journals and conferences in most of those areas.