

Towards a Safety Case for Hardware Fault Tolerance in Convolutional Neural Networks Using Activation Range Supervision

Florian Geissler^{1*}, Syed Qutub¹, Sayanta Roychowdhury¹, Ali Asgari², Yang Peng¹, Akash Dhamasia¹, Ralf Graefe¹, Karthik Pattabiraman² and Michael Paulitsch¹

¹Intel, Germany

²University of British Columbia, Canada

Abstract

Convolutional neural networks (CNNs) have become an established part of numerous safety-critical computer vision applications, including human robot interactions and automated driving. Real-world implementations will need to guarantee their robustness against hardware soft errors corrupting the underlying platform memory. Based on the previously observed efficacy of activation clipping techniques, we build a prototypical safety case for classifier CNNs by demonstrating that range supervision represents a highly reliable fault detector and mitigator with respect to relevant bit flips, adopting an eight-exponent floating point data representation. We further explore novel, non-uniform range restriction methods that effectively suppress the probability of silent data corruptions and uncorrectable errors. As a safety-relevant end-to-end use case, we showcase the benefit of our approach in a vehicle classification scenario, using ResNet-50 and the traffic camera data set MIOVision. The quantitative evidence provided in this work can be leveraged to inspire further and possibly more complex CNN safety arguments.

1 Motivation

With the widespread use of convolutional neural networks (CNN) across many safety-critical domains such as automated robots and cars, one of the most prevailing challenges is the establishment of a safety certification for such artificial intelligence (AI) components, e.g., with respect to the ISO 26262 [1] or ISO/PAS 21448 (SOTIF) [2]. This certification requires not only a high fault-tolerance of the trained network against unknown or adversarial input, but also efficient protection against hardware faults of the underlying platform [3, 4]. Importantly this includes transient *soft errors*, meaning disturbances originating from events such as cosmic neutron radiation, isotopes emitting alpha particles, or electromagnetic leakage on the computer circuitry itself.

*corresponding author, Email: florian.geissler@intel.com. Copyright © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

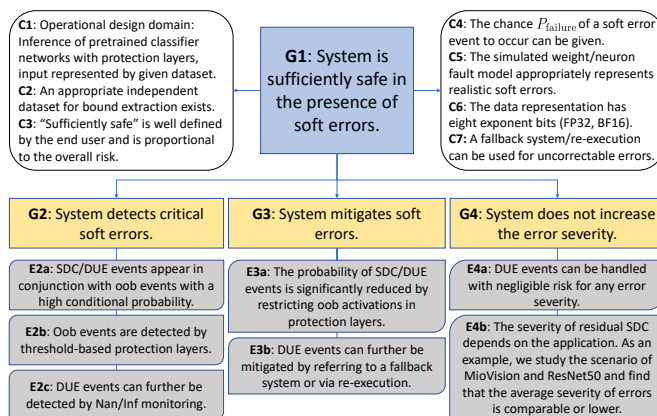


Figure 1: Structured safety argument for the fault tolerance of a CNN in the presence of soft errors, using range restrictions. The notation follows [10] including goals (G), context (C), and evidence (E). "Oob" denotes "out-of-bounds".

Soft errors typically manifest as single or multiple bit upsets in the platform's memory elements [5]. As a consequence, network parameters (*weight faults*) or local computational states (*neuron faults*) can be altered during inference time, and invalidate the network prediction in a safety-critical way, for example, by misclassifying a person as a background image in an automated driving context [6–8]. This has led to a search for strategies to verify CNN-based systems against hardware faults at the inference stage [9]. With chip technology nodes scaling to smaller sizes and larger memory density per area, future platforms are expected to be even more susceptible to soft errors [5].

In this paper, we evaluate range restriction techniques in CNNs exposed to platform soft errors with respect to the key elements of a prototypical safety case. This means that we formulate arguments (in the form of "goals") that constitute essential parts of a complete safety case, and provide quantitative evidence to support these goals in the studied context (see Fig. 1). Individual safety arguments can be reused as building blocks of more complex safety cases. The structure of our goals is based on the probabilistic, high-level safety




Figure 2: Visualization example of the impact of a weight fault using LeNet-5 and the MNIST data set. Range restriction layers (“Ranger”) are inserted following [8] (top left). The rows represent the feature maps of the individual network layers after range restriction was applied, where linear layers (FC1-FC3) were reshaped to a 2D feature map as well for visualization purposes. In (b) - (g), a large weight fault value is injected in the second filter of the first convolutional layer. For the unprotected model (b), this leads to a SDC event (“0” gets changed to “7”). The columns (c) - (g) then illustrate the effect of the different investigated range restriction methods.

objective of minimizing the overall risk [10], expressed as:

$$P_{\text{loss}}(i) = P_{\text{failure}}(i) [(1 - P_{\text{detection}}(i)) + (1 - P_{\text{mitigation}}(i))],$$

$$\text{Risk} = \sum_i P_{\text{loss}}(i) \cdot \text{Severity}(i).$$
(1)

Explicitly, for a fault type i , this includes the sub-goals of efficient error *detection* and *mitigation*, as well as a consideration of the fault *severity* in a given use case. On the other hand, the probability of occurrence of a soft error (i.e., P_{failure} in Eq. 1) is assumed to be a constant system property that cannot be controlled by run-time monitoring methods such as activation range supervision.

In a nutshell, range restriction builds on the observation that silent data corruption (SDC) and detected uncorrectable errors (DUE), e.g., *NaN* and *Inf* occurrences), stem primarily from those bit flips that cause very large values, for example in high exponential bits [6]. Those events result in large activation peaks that typically grow even more during forward

propagation due to the monotonicity of most neural network operations [11]. To suppress the propagation of such corrupted values, additional range restriction layers are inserted in the network at strategic positions following the approach of Chen et. al. [8] (see Fig. 2 for an example). At inference time, the protection layers then compare the intermediate activations against previously extracted interval thresholds in order to detect and reset anomalously large values. Derivative approaches have been shown to be efficient in recovering network performance [6–8, 12] and, advantageously, do not require the retraining of CNN parameters nor computationally expensive functional duplications.

The focus of this paper is to examine alternative restriction schemes for optimized soft error mitigation. In a CNN, the output of every kernel is represented as a two-dimensional (2D) *feature map*, where the activation magnitudes encode specific features, on which the network bases its prediction. Soft errors will manifest as distortions of feature maps in all subsequent layers that make use of the corrupted value, as shown in Fig. 2(a)-(b). The problem of mitigating soft errors in a CNN can therefore be rephrased as restoring the fault-free topology of feature maps.

Previous analyses have adopted uniform range restriction schemes that truncate out-of-bound values to a finite threshold [7, 8], e.g., Fig. 2(c)-(d). We instead follow the intuition that optimized, non-uniform range restriction methods that attempt to reconstruct feature maps (see Fig. 2(e)-(g), and details in Sec. 5) can not only reduce SDC to a comparable or even lower level, but may also lead to less critical misclassifications in the case of an SDC. This is because classes with more similar attributes will display more similar high-level features (e.g., pedestrian and biker will both exhibit upright silhouette, in contrast to car and truck classes).

Finally, a safety analysis has to consider that not all SDC events pose an equal risk to the user. We study a safety-critical use case evaluating cluster-wise class confusions in a vehicle classification scenario (Sec. 6). The example shows that range supervision reduces the severe confusions proportionally with the overall number of confusions, meaning that the total risk is indeed mitigated.

In summary, this paper make the following contributions:

- **Fault detection:** We quantify the correlation between SDC events and the occurrence of out-of-bound activations to demonstrate the high efficiency of fault detection by monitoring intermediate activations,
- **Fault mitigation:** We explore three novel range restriction methods that build on the preservation of the feature map topologies instead of mere value truncation,
- **Fault severity:** We demonstrate the benefit of range supervision in an end-to-end use case of vehicle classification where high and low severities are estimated by the generic safety-criticality of class confusions.

The article is structured as follows: Section 2 reviews relevant previous work while section 3 describes the setup used in this paper. Subsequently, the sections 4, 5, and 6 discuss error detection, mitigation, and an exemplary risk analysis, respectively, before section 7 concludes the paper.

2 Related work

Parity or error-correcting code (ECC) can protect memory elements against single soft errors [5, 13]. However, due to the high compute and area overhead, this is typically done only for selected critical memory blocks. Component replication techniques such as triple modular redundancy can be used for the full CNN execution at the cost of a large overhead. Selective hardening of hardware elements with the most salient parameters can improve the robustness of program execution in the presence of underlying faults [6, 14]. On a software level, the estimation of the CNN’s vulnerable feature maps (fmaps) and the selective protection by duplicated computations [15], or the assertive re-execution with stored, healthy reference values [16] has been investigated. Approaches using algorithm-based fault tolerance (ABFT) [17] seek to protect networks against soft errors by checking invariants that are characteristic for a specific operation (e.g., matrix multiplication). Symptom-based error detection may for example include the interpretation of feature map traces by a secondary companion network [18]. The restriction of intermediate ranges was explored [6, 12] in the form of modified (layer-insensitive) activation functions such as *tanh* or *ReLU6*. This concept was extended to find specific uniform protection thresholds for neuron faults [8] or clipping bounds for weight faults [7]. An alternative line of research is centered around fault-aware retraining [19].

3 Experimental setup

3.1 Models, data sets, and system

CNNs are the most commonly used network variant for computer vision tasks such as object classification and detection. We compare the three standard classifier CNNs ResNet-50 [20], VGG-16 [21], and AlexNet [22] together with the test data set ImageNet [23] and MIOVision [24] for the investigation of a safety-critical example use case. Since fault injection is compute-intensive, we rescale our test data set for ImageNet to a subset of 1000 images representing 20 randomly selected classes. For MIOVision, a subset of 1100 images (100 per class) that were correctly classified in the absence of faults was chosen. All experiments adopt a single-precision floating point format (FP32) according to the IEEE754 standard [25]. Our conclusions apply as well to other floating point formats with the same number of exponent bits, such as BF16 [26], since no relevant effect was observed from fault injections in mantissa bits (Sec. 4).

Experiments were performed in PyTorch (version 1.8.0) deploying torchvision models (version 0.9.0). For MIOVision, the ResNet-50 model was retrained [27]. We used Intel® Core™ i9 CPUs, with inferences running on GeForce RTX 2080, Titan RTX, and RTX 3090 GPUs.

3.2 Protection layers and bound extraction

We insert protection layers at strategic positions in the network such as after activation, pooling, reshape or concatenate layers, according to the model of Chen et al. [8]. Each protection layer requires specific bound values for the expected activation ranges as a parameter. We extract those by monitoring the minimal and maximal activations from a sepa-




Figure 3: Illustration of SDC and DUE events. Errors are detected or missed in the case of out-of-bound (oob) or in-bound (ib) events, respectively. (Green) Samples of the data set that form the subset of a given filtering stage, (Yellow) samples of the data set that are discarded at the given stage, (White) samples that were filtered out at a previous stage.

rate test input, which is taken from the training data sets of ImageNet (143K images used) and MIOVision (83K images used), respectively. This step has to be performed only once. Bound extraction depends on the data set and will in general impact the safety argument (see Fig. 1). To check the suitability of the bounds, we verify that no out-of-bound events were detected during the test phase in the absence of faults, so the baseline accuracy is the same with and without protection. While all minimum bounds are zero in the studied setup, the maximum activation values for ImageNet vary by layer in a range of (see also Sec. 5) $1 < T_{up} < 45$ for ResNet-50, $20 < T_{up} < 360$ for VGG-16, and $65 < T_{up} < 170$ for AlexNet. For MIOVision and ResNet-50, we find maximum bounds between $1 < T_{up} < 19$.

3.3 Fault model and injection

In line with previous investigations, we distinguish two different manifestations of memory bit flips referred to here as *weight faults* and *neuron faults*. The former represent soft errors affecting memory elements that store the learned network parameters, while the latter refer to errors in memory that holds temporary states such as intermediate network layer outputs. While neuron faults may also impact states used for logical instructions, it was demonstrated that bit flip injections in the output of the affected layer are generally a good model approximation [28]. Memory elements can be protected against single bit flips by mechanisms such as parity and ECC [5, 13]. However, this kind of protection is not always available due to the associated compute and area overhead. Further, ECC typically cannot correct multi-bit flips.

We inject faults either directly in the weights of CNN layers (*weight faults*) or in the output of the latter (*neuron faults*), using a customized fault injection framework based on PytorchFI [29]. To speed up the experiments we focus on bit flips in the most relevant bit positions 0 – 8 (sign bit and ex-

ponential bits, neglecting mantissa) unless stated otherwise. Fault locations (i.e., layer index, kernel index, channel etc.) in the network are randomly chosen with an equal weight, so without further constraints on the selection process to reflect the arbitrary occurrence of soft errors. As weights are typically stored in the main memory and loaded only once for a given application, we keep the same weight faults for one entire epoch, running all tested input images. In total, we run 500 epochs, i.e., fault configurations, each one applied to 1K images. Neuron faults, on the other hand, apply to memory representing temporary states that are overwritten for each new input. Therefore, we inject new neuron faults for each new input and run 100 epochs resulting in 100K fault configurations, each one applied to a single image.

3.4 Evaluation

To quantify the impact of faults on the system safety, we measure the rate of SDC events. Throughout, we consider the Top-1 prediction to determine SDC. In line with previous work [6, 8], SDC is defined as the ratio of images that are misclassified in the presence of faults (without exceptions) but correctly classified in the absence of faults and the overall number of images, $p(\text{sdc}) = N_{\text{incorrect}}/N_{\text{test, correct}}$ (Fig. 3).

During the forward pass, non-numerical exceptions in the form of *Inf* and *NaN* values can be encountered, due to the following reasons: **i)** *Inf* values occur if large activation values accumulate (for example during conv2d, linear, avgpool2d operations) until they exceed the maximum of the data representation. This effect becomes particularly apparent when flips of the most significant bit (MSB, position index 1) are injected. **ii)** *NaN* values are found when denominators are undetermined or multiple *Inf* values get added, e.g., in BatchNorm2d layers, **iii)** *NaN* values can be generated directly via bit flips in conv2d layers due to the fact that FP32 encodes *NaN* as all eight exponent bits being in state "1". In the studied classifier networks, the latter effect is very rare for single bit flips in weights (see Sec. 4) but not necessarily for single neuron bit flips or multiple flips of either type.

The creation of the above exceptions is found to differ slightly between CPU and GPU executions, as well as between experiments with different batch sizes on the accelerator. We attribute this observation to algorithmic optimizations on the GPU that are not necessarily IEEE754-compliant and thus affect the floating point precision [30]. To mitigate the effect of exception handling we monitor the occurrences of *Inf* and *NaN* in the output of any network layer. All forward passes with an exception are separated and define the detected uncorrectable error (DUE) rate, $p(\text{due}) = N_{\text{exceptions}}/N_{\text{test, correct}}$, see Fig. 3.

In a real system, DUE events can be readily monitored and the execution is typically halted on detection. However, due to the non-numerical nature of these errors we cannot apply the same mitigation strategy that is adopted for SDC events. We therefore make the assumption that either a fallback system (e.g., alternative classifier, emergency stop of vehicle, etc.) can be leveraged or a timely re-execution is possible to recover from transient DUE events. This in turn assumes that DUEs do not impact the system safety but may compromise the system availability when occurring frequently.




Figure 4: Bit-distribution across all weight parameters in conv2d layers. Values are represented in FP32, where only the sign bit (0) and the exponent bits (1 – 8) are shown.

4 Error detection coverage

To effectively protect the network against faults, we first verify the error detection coverage for silent errors. Those errors are detected by a given protection layer if the activation values exceed (fall short of) the upper (lower) bound. If at least one protection layer is triggered per inference run, we register an out-of-bound (oob) event. Otherwise, we have an in-bound (ib) event. In addition, we quantify the probabilities of SDC and regular correct classification (cl) events, as well as the respective conditional probabilities that correct and incorrect classifications occur *given that oob or ib events were detected*. This allows us to define true positive (Tp), false positive (Fp), and false negative (Fn) SDC detection rates as

$$\begin{aligned}
 \text{Tp} &= p(\text{sdc}|\text{oob}) \cdot p(\text{oob}), \\
 \text{Fp} &= p(\text{cl}|\text{oob}) \cdot p(\text{oob}), \\
 \text{Fn} &= p(\text{sdc}|\text{ib}) \cdot p(\text{ib}).
 \end{aligned} \tag{2}$$

The fault detector then is characterized by precision, $P = \text{Tp}/(\text{Tp} + \text{Fp})$, and recall, $R = \text{Tp}/(\text{Tp} + \text{Fn})$.

The Tab. 1 displays the chances of oob and sdc events resulting from a single fault per image in the absence of range protection. For weight faults, we find that all three CNNs showcase a high correlation between oob situations and either SDC or DUE events ($p(\text{sdc}|\text{oob}) + p(\text{due}|\text{oob}) > 0.99$), which can be associated with the chance of a successful error detection, $P_{\text{detection}}$ (see Eq. 1). The chance of finding SDC after ib events is very small ($\ll 1e^{-3}$), leading to a very high precision and recall performance (> 0.99). For neuron faults, while the recall remains very high, the precision is reduced (in particular VGG-16 and AlexNet) due to additional Fp events where non-MSB oob events still get classified correctly.

We further verify that SDC events from single weight faults are attributed almost exclusively to flips of the MSB. This can be explained with the distribution of parameters in the studied networks (Fig. 4). The weight values are closely centered around zero, and thus exhibit characteristic properties when

represented in an eight-exponent data format. In the fault-free case, the MSB always has state “0”, while the exponent bits 2 to 4 are almost always in state “1”. This means that among the relevant exponential bits all single bit flips of the MSB will produce large values, while those of the other exponential bits will either be from “1” → “0” or will be too small to have a significant effect.

For neuron faults, on the other hand, the distribution of fault-free values is input-dependent and broader, leading in general to a smaller quota of MSB flips to SDC, in favor of flips of other exponential bits and the sign bit. No SDC due to mantissa bit flips were observed in either weight or neuron faults. DUE events are unlikely (< 0.01) for a single bit flip as there are not multiple large values to add up. Further, network weights are usually < 1 , meaning that at least two exponent bits are in state “0”, and hence at least two bit flips are needed to directly generate a *NaN* value.

5 Range restriction methods for error mitigation

5.1 Model

We refer to a subset of the tensor given by a specific index in the batch and channel dimensions as a 2D feature map, denoted by f . Let x be an activation value from a given feature map tensor $f \in \{f_1, f_2, \dots, f_{C_{out}}\}$. Further, T_{up} and T_{low} denote the upper and lower activation bounds assigned to the protection layer, respectively.

Ranger: For a given set of (f, T_{up}, T_{low}) , Ranger [8] maps out-of-bound values to the expected interval (see Fig. 2c),

$$r_{\text{ranger}}(x) = \begin{cases} T_{up} & \text{if } x > T_{up}, \\ T_{low} & \text{if } x < T_{low}, \\ x & \text{otherwise.} \end{cases} \quad (3)$$

Clipper: In a similar way, clipping truncates activations that are out of bound to zero [7],

$$r_{\text{clipping}}(x) = \begin{cases} 0 & \text{if } x > T_{up} \text{ or } x < T_{low}, \\ x & \text{otherwise.} \end{cases} \quad (4)$$

The intuition is that it can be favorable to eliminate corrupted elements rather than to re-establish finite activations.

FmapRescale: While uniform restriction methods help in eliminating large out-of-bound values, the information encoded in relative differences of activation magnitudes is lost when all out-of-bound values are flattened to the same value. The idea of rescaling is to linearly map all large out-of-bound values back onto the interval $[T_{low}, T_{up}]$, implying that smaller out-of-bound values are reduced more. This follows the intuition that the out-of-bound values can originate from the entire spectrum of in-bound values.

$$r_{\text{rescale}}(x) = \begin{cases} \frac{(x - \min(f))(T_{up} - T_{low})}{\max(f) - \min(f)} + T_{low} & \text{if } x > T_{up}, \\ T_{low} & \text{if } x < T_{low}, \\ x & \text{otherwise.} \end{cases} \quad (5)$$

Backflip: We analyze the underlying bit flips that may have caused out-of-bound values. This reasoning holds for

	Weight faults	Neuron faults
ResNet-50:		
$p(\text{sd})$	0.018 ± 0.001	$0.013 \pm 6e^{-4}$
$p(\text{oob})$	0.019 ± 0.001	$0.013 \pm 6e^{-4}$
$p(\text{sd} \text{oob})$	0.981 ± 0.008	0.974 ± 0.008
$p(\text{sd} \text{ib})$	$5e^{-5} \pm 4e^{-5}$	0.0 ± 0.0
$p(\text{MSB} \text{sd})$	0.998 ± 0.002	0.961 ± 0.012
P	0.997 ± 0.002	0.980 ± 0.006
R	0.997 ± 0.002	1.0 ± 0.0
$p(\text{due})$	$3e^{-4} \pm 1e^{-4}$	$5e^{-4} \pm 1e^{-4}$
$p(\text{due} \text{oob})$	0.016 ± 0.008	0.006 ± 0.005
$p(\text{MSB} \text{due})$	1.0 ± 0.0	1.0 ± 0.0
VGG-16:		
$p(\text{sd})$	0.024 ± 0.001	$0.016 \pm 9e^{-4}$
$p(\text{oob})$	0.027 ± 0.001	0.020 ± 0.001
$p(\text{sd} \text{oob})$	0.893 ± 0.010	0.778 ± 0.016
$p(\text{sd} \text{ib})$	$7e^{-5} \pm 7e^{-5}$	0.0 ± 0.0
$p(\text{MSB} \text{sd})$	0.997 ± 0.003	0.397 ± 0.017
P	0.999 ± 0.001	0.820 ± 0.014
R	0.997 ± 0.003	1.0 ± 0.0
$p(\text{due})$	$0.003 \pm 4e^{-4}$	$0.006 \pm 4e^{-4}$
$p(\text{due} \text{oob})$	0.106 ± 0.011	0.051 ± 0.012
$p(\text{MSB} \text{due})$	1.0 ± 0.0	1.0 ± 0.0
AlexNet:		
$p(\text{sd})$	0.022 ± 0.001	0.013 ± 0.001
$p(\text{oob})$	0.024 ± 0.001	0.015 ± 0.001
$p(\text{sd} \text{oob})$	0.907 ± 0.012	0.877 ± 0.023
$p(\text{sd} \text{ib})$	$2e^{-4} \pm 1e^{-4}$	$9e^{-5} \pm 5e^{-5}$
$p(\text{MSB} \text{sd})$	0.995 ± 0.003	0.245 ± 0.031
P	1.0 ± 0.0	0.913 ± 0.025
R	0.989 ± 0.005	0.994 ± 0.004
$p(\text{due})$	$0.003 \pm 3e^{-4}$	$0.005 \pm 3e^{-4}$
$p(\text{due} \text{oob})$	0.093 ± 0.012	0.040 ± 0.011
$p(\text{MSB} \text{due})$	1.0 ± 0.0	1.0 ± 0.0

Table 1: Statistical absolute and conditional probabilities of SDC, DUE events and the related precision and recall of the fault detector. Experiments of 10K fault injections were repeated 10 times, where a single fault per image was injected in any of the 32 bits for each image (from ImageNet, using a batch size of one). We further list what proportion of SDC or DUE events were caused by MSB flips.

neuronal faults, where we may assume that a specific activation value is bit-flipped directly. For weight faults, on the other hand, the observed out-of-bound output activation is the result of a multiply-and-accumulate operation of an input tensor with a bit-flipped weight value. However, we argue that the presented back-flip operation will recover a representative product, given that the input component is of the order of magnitude of one. To restore a flipped value, we distinguish the following cases:

$$r_{\text{backflip}}(x) = \begin{cases} 0 & \text{if } x > T_{up} \cdot 2^{64}, \\ 2 & \text{if } T_{up} \cdot 2^{64} > x > T_{up} \cdot 2, \\ T_{up} & \text{if } T_{up} \cdot 2 > x > T_{up}, \\ T_{low} & \text{if } x < T_{low}, \\ x & \text{otherwise.} \end{cases} \quad (6)$$




Figure 5: SDC rates for weight (a) and neuron (b) faults using different range supervision techniques. Note that compared to Tab. 1 rates are around $4\times$ higher since we inject only in the bits 0 – 8 here.

The above thresholds are motivated by the following logic: Given appropriate bounds, an activation is $< T_{up}$ before a bit flip. Any flip of an exponential bit $i \in \{1 \dots 8\}$ effectively multiplies a factor of $\text{pow}(2, 2^{8-i})$. Hence, any value beyond $T_{up} \cdot 2^{64}$ must have originated from a flip "0" \rightarrow "1" of the MSB, meaning that the original value was between 0 and 2. We then set back all out-of-bound values in this regime to zero, assuming that lower reset values represent a more conservative choice in eliminating faults. Next, flipped values that are between $T_{up} \cdot 2^{64} > x > T_{up} \cdot 2$ can possibly originate from a flip of any exponential bit. Given that T_{up} is typically > 1 , a bit flip has to produce a corrupted absolute value > 2 in this regime. This is possible only if either the MSB is flipped from "0" \rightarrow "1", or the MSB is already at "1" and another exponential bit is flipped "0" \rightarrow "1". In all variants of the latter case, the original value had to be already > 2 itself, and hence we conservatively reset out-of-bound values to 2. Finally, corrupted values of $T_{up} \cdot 2 > x > T_{up}$ may originate from any non-sign bit flip. Lower exponential or even fraction bit flips result from already large values close to T_{up} in this case, which is why we set back those values to the upper bound. As in Ranger, values that are too small are reset to T_{low} .

FmapAvg: The last proposed range restriction technique uses the remaining, healthy fmaps of a convolutional layer to

reconstruct a corrupted fmap. The intuition behind this approach is as follows: Every filter in a given conv2d layer tries to establish characteristic features of the input image. Typically, there is a certain redundancy in the topology of fmaps, since not all features the network was trained to recognize may be strongly pronounced for a given image (instead mixtures of potential features may form), or because multiple features resemble each other at the given processing stage. Therefore, replacing a corrupted fmap with a non-corrupted fmap from a *different* kernel can help to obtain an estimate of the original topology. We average all healthy (i.e., not containing out-of-bound activations) fmaps by

$$\text{ind} = \{i = 1 \dots C_{out} \mid \max(f_i) \leq T_{up}, \min(f_i) \geq T_{low}\},$$

$$f_{avg} = \frac{1}{|\text{ind}|} \sum_{j \in \text{ind}} f_j. \quad (7)$$

If there are no healthy feature maps, f_{avg} will be the zero-tensor. Subsequently, we replace oob values in a corrupted fmap with their counterparts from the estimate of Eq. (7),

$$r_{favg}(x) = \begin{cases} f_{avg}(x) & \text{if } x > T_{up} \text{ or } x < T_{low}, \\ x & \text{otherwise.} \end{cases} \quad (8)$$

5.2 Results

In Fig. 5 we present results for the SDC mitigation experiments with different range supervision methods. Comparing 1 and 10 fault injections per input image, we note that the unprotected models are dramatically corrupted with an increasing fault rate (SDC rate becomes ≥ 0.50 for weights, ≥ 0.32 for neurons in the presence of 10 faults). We can associate the SDC rate with the chance of unsuccessful mitigation, $1 - P_{mitigation}$, in Eq. 1. Weight faults have a higher impact than neuron faults since they directly corrupt a multitude of activations in a layer's fmap output (in contrast to individual activations for neuron faults) and thus propagate faster than neuron faults.

All the studied range restriction methods reduce the SDC rate by a significant margin, but perform differently for weight and neuron fault types: For weight faults, we observe that *Clipper*, *Backflip*, and *FmapAvg* are highly efficient in all three networks, with SDC rates suppressed to values of $\lesssim 0.01$ (SDC reduction of $> 50\times$). *Ranger* provides a much weaker protection, in particular in the more shallow networks VGG-16 and AlexNet. *FmapRescale* performs better than *Ranger* but worse than the aforementioned methods. The deepest studied network, ResNet-50, benefits the most from any type of range restriction in the presence of weight faults.

When it comes to neuron faults (Fig. 5b), we see that *Clipper* and *Backflip* provide the best protection (SDC rate is suppressed to < 0.005 , reduction of $> 38\times$), followed by the also very effective *Ranger* (except for AlexNet). *FmapAvg* appears to be less efficient for higher fault rates in this scenario, while *FmapRescale* again falls behind all the above.

Overall, we conclude that the pruning-inspired mitigation techniques *Clipper* and *Backflip* represent the best choices among the investigated ranger supervision methods, as they succeed in mitigating both weight and neuron faults to very small residual SDC rates.




Figure 6: Formation of class clusters in MIOVision (VRU denotes vulnerable road user). We make the assumption here that confusions towards less vulnerable clusters are the most safety-critical ones.


In the experiments of Fig. 5, the encountered DUE rates for 1 weight or neuron fault (0.003 for ResNet, 0.03 for VGG-16 or AlexNet) are only slightly reduced by range restrictions. However, for a fault rate of 10 we find the following trends: **i)** For weights, the DUE is significantly reduced in ResNet (from 0.15 to 0.002), while rates in VGG (0.22) and AlexNet (0.26) remain. **ii)** For neurons, *Ranger*, *Clipper* and *Backflip* suppress the DUE rate by a factor of up to $2\times$ in all networks.

The studied range restriction techniques require different compute costs due to the different number of additional graph operations. In PyTorch, not all needed functions can be implemented with the same efficiency though. For example, *Ranger* is executed with a single *clamp* operation, while no equivalent formulation is available for *Clipper* and instead three operations are necessary (two masks to select oob values greater and smaller than the threshold, and a *masked-fill* operation to clip to zero). As a consequence, measured latencies are framework-dependent and a fair comparison cannot be made at this point. Given the complexity of the protection operations, we may instead give a qualitative performance ranking of the described methods: *FmapRescale* appears to be the most expensive restriction method due to the needed number of operations, followed by *FmapAvg* and *Backflip*. *Clipper* and *Ranger* are the least complex, with the latter outperforming the former in the used framework, due to its more efficient use of optimized built-in operations.


6 Analysis of traffic camera use case

As a selected safety-critical use case, we study object classification in the presence of soft errors with a retrained ResNet-50 and the MIOVision data set [24]. The data contains images of 11 classes including for example pedestrian, bike, car, or background, that were taken by traffic cameras. The correct identification of an object type or category can be safety-critical for example to an automated vehicle that uses the support of infrastructure sensors for augmented perception [31].

However, not every class confusion is equally harmful. To estimate the severity of an error-induced misclassification we establish three clusters of vulnerable, as well as non-vulnerable road users (VRU or non-VRU), and background,



(a) Weight faults



(b) Neuron faults

Figure 7: SDC rates for ResNet-50 and MIOvision. We inject 1 and 10 faults targeting bits 0 – 8 in the network weights (a) and neurons (b). The portion of safety-critical SDC events according to Fig. 6 is displayed as a dark-color overlay.

see Fig. 6. Misclassifications that lead to the prediction of a class in a *less vulnerable* cluster are assumed to be safety-critical (Severity ≈ 1 in Eq. 1, e.g., a pedestrian is misclassified as background), while confusions within the same cluster or towards a *more vulnerable* cluster are considered non-critical (Severity ≈ 0) as they typically lead only to similar or a more cautious behavior. This binary estimation allows us quantify the overall risk as the portion of SDC events associated with the respective critical class confusions.

From our results in Fig. 7 we make the following observations: **i)** The relative proportion of critical confusions is lower for weight than for neuron faults in the unprotected and most protected models. For weight faults, the most frequent confusions are from other classes to the class "car" (the most robust class of MIOVision, with the most images in the training set), which are statistically mostly non-critical. Neuron faults, on the other hand, distort feature maps in a way that induces with the highest frequency misclassifications towards

the class "background". Those events are all safety-critical (see Fig. 6), leading to a high critical-to-total SDC ratio. **ii)** Range supervision is not only effective in reducing the overall SDC count, but also suppresses the critical SDC count proportionally. For example, we observe that the most frequent critical class confusion caused by 1 or 10 weight faults is from the class "pedestrian" to "car" (≈ 0.2 of all critical SDC cases), where > 0.99 of those cases can be mitigated by *Clipper* or *Backflip*. For neuron faults, the largest critical SDC contribution is from "pedestrian" to "background" (1 fault) or "car" to "background" (10 faults), both in about 0.1 of all critical SDC cases. *Clipper* or *Backflip* are able to suppress > 0.91 of those events.

As a consequence, all studied range-restricted models exhibit a critical-to-total SDC ratio that is similar to or lower than one of the unprotected network (< 0.41 for weight, < 0.78 for neuron faults), meaning that faults in the presence of range supervision have on average a similar or lower severity than faults that do not face range restrictions. A lower ratio can be interpreted as a better preservation of the feature map topology: If the reconstructed features are more similar to the original features there is a higher chance of the incorrect class being similar to the original class and thus to stay within the same cluster. The total probability of critical SDC events – and therefore the relative risk according to Eq. 1 – is negligible in the studied setup in the presence of *Clipper* or *Backflip* range protection.

The mean DUE rates in the unprotected model are 0.0 (0.02) for 1 weight (neuron) fault and 0.11 (0.17) for 10 faults. Using any of the protection methods, the system's availability increases as DUE rates are negligible for 1 fault, and reduce to < 0.03 (< 0.05) for 10 weight (neuron) faults.

7 Conclusion

In this paper, we investigated the efficacy of range supervision techniques for constructing a safety case for computer vision AI applications that use Convolutional Neural Networks (CNNs) in the presence of platform soft errors. In the given experimental setup, we demonstrated that the implementation of activation bounds allows for a highly efficient detection of SDC-inducing faults, most importantly featuring a recall of > 0.99 . Furthermore, we found that the range restriction layers can mitigate the once-detected faults effectively by mapping out-of-bound values back to the expected intervals. Exploring distinct restriction methods, we observed that *Clipper* and *Backflip* perform best for both weight and neuron faults and can reduce the residual SDC rate to $\lesssim 0.01$ (reduction by a factor of $> 38\times$). Finally, we studied the selected use case of vehicle classification to quantify the impact of range restriction on the severity of SDC events (represented by cluster-wise class confusions). All discussed techniques reduce critical and non-critical events proportionally, meaning that the average severity of SDC is not increased. Therefore, we conclude that the presented approach reduces the overall risk and thus enhances the safety of the user in the presence of platform soft errors.

Acknowledgment

Our research was partially funded by the Federal Ministry of Transport and Digital Infrastructure of Germany in the project Providentia++ (01MM19008). Further, this research was partially supported by a grant from the Natural Sciences and Engineering Research Council of Canada (NSERC), and a research gift from Intel to UBC.

References

- [1] International Organization for Standardization, "ISO 26262," Tech. Rep., 2018. [Online]. Available: <https://www.iso.org/standard/68383.html>
- [2] —, "Road vehicles - Safety of the intended functionality," Tech. Rep., 2019. [Online]. Available: <https://www.iso.org/standard/70939.html>
- [3] J. Athavale, A. Baldovin, R. Graefe, M. Paulitsch, and R. Rosales, "AI and Reliability Trends in Safety-Critical Autonomous Systems on Ground and Air," *Proceedings - 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN-W 2020*, pp. 74–77, 2020.
- [4] H. D. Dixit, S. Pendharkar, M. Beadon, C. Mason, T. Chakravarthy, B. Muthiah, and S. Sankar, *Silent Data Corruptions at Scale*. Association for Computing Machinery, 2021, vol. 1, no. 1. [Online]. Available: arxiv.org/abs/2102.11245
- [5] A. Neale and M. Sachdev, "Neutron Radiation Induced Soft Error Rates for an Adjacent-ECC Protected SRAM in 28 nm CMOS," *IEEE Transactions on Nuclear Science*, vol. 63, no. 3, pp. 1912–1917, 2016.
- [6] G. Li, S. K. S. Hari, M. Sullivan, T. Tsai, K. Pattabiraman, J. Emer, and S. W. Keckler, "Understanding error propagation in Deep Learning Neural Network (DNN) accelerators and applications," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2017*, 2017.
- [7] L.-H. Hoang, M. A. Hanif, and M. Shafique, "FT-ClipAct: Resilience Analysis of Deep Neural Networks and Improving their Fault Tolerance using Clipped Activation," 2019. [Online]. Available: <https://arxiv.org/abs/1912.00941>
- [8] Z. Chen, G. Li, and K. Pattabiraman, "Ranger: Boosting Error Resilience of Deep Neural Networks through Range Restriction," 2020. [Online]. Available: <https://arxiv.org/abs/2003.13874>
- [9] J. M. Cluzeau, X. Henriquel, G. Rebender, G. Soudain, L. van Dijk, A. Gronskiy, D. Haber, C. Perret-Gentil, and R. Polak, "Concepts of Design Assurance for Neural Networks (CoDANN)," *Public Report Extract Version 1.0*, pp. 1–104, 2020. [Online]. Available: <https://www.easa.europa.eu/document-library/general-publications/concepts-design-assurance-neural-networks-codann>
- [10] P. Koopman and B. Osyk, "Safety argument considerations for public road testing of autonomous vehicles," *SAE Technical Papers*, vol. 2019-April, no. April, 2019.
- [11] Z. Chen, G. Li, K. Pattabiraman, and N. Debardeleben, "BinFI: An efficient fault injector for safety-critical machine learning systems," *International Conference for High Performance Computing, Networking, Storage and Analysis, SC*, 2019.

- [12] S. Hong, P. Frigo, Y. Kaya, C. Giuffrida, and T. Dumitras, "Terminal brain damage: Exposing the graceless degradation in deep neural networks under hardware fault attacks," in *Proceedings of the 28th USENIX Security Symposium*, 2019.
- [13] A. Lotfi, S. Hukerikar, K. Balasubramanian, P. Racunas, N. Saxena, R. Bramley, and Y. Huang, "Resiliency of automotive object detection networks on GPU architectures," *Proceedings - International Test Conference*, vol. 2019-Novem, pp. 1–9, 2019.
- [14] M. A. Hanif and M. Shafique, "SalvagedNn: Salvaging deep neural network accelerators with permanent faults through saliency-driven fault-aware mapping," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 2020. [Online]. Available: <https://royalsocietypublishing.org/doi/10.1098/rsta.2019.0164>
- [15] A. Mahmoud, S. K. Sastry Hari, C. W. Fletcher, S. V. Adve, C. Sakr, N. Shanbhag, P. Molchanov, M. B. Sullivan, T. Tsai, and S. W. Keckler, "Hardnn: Feature map vulnerability evaluation in CNNs," 2020.
- [16] J. Ponader, S. Kundu, and Y. Solihin, "MILR: Mathematically Induced Layer Recovery for Plaintext Space Error Correction of CNNs," 2020. [Online]. Available: <http://arxiv.org/abs/2010.14687>
- [17] K. Zhao, S. Di, S. Li, X. Liang, Y. Zhai, J. Chen, K. Ouyang, F. Cappello, and Z. Chen, "FT-CNN: Algorithm-Based Fault Tolerance for Convolutional Neural Networks," *IEEE Transactions on Parallel and Distributed Systems*, vol. 32, no. 7, pp. 1677–1689, 2021.
- [18] C. Schorn, A. Guntoro, and G. Ascheid, "Efficient On-Line Error Detection and Mitigation for Deep Neural Network Accelerators," in *Safecomp 2018*, vol. 11093 LNCS, 2018.
- [19] L. Yang and B. Murmann, "SRAM voltage scaling for energy-efficient convolutional neural networks," in *Proceedings - International Symposium on Quality Electronic Design, ISQED*. IEEE Computer Society, may 2017, pp. 7–12.
- [20] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2016-Decem, pp. 770–778, 2016.
- [21] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, 2015.
- [22] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Advances in Neural Information Processing Systems*, vol. 2, pp. 1097–1105, 2012.
- [23] Jia Deng, Wei Dong, R. Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [24] Z. Luo, F. B. Charron, C. Lemaire, J. Konrad, S. Li, A. Mishra, A. Achkar, J. Eichel, and P.-M. Jodoin, "MIO-TCD: A new benchmark dataset for vehicle classification and localization," *IEEE Transactions on Image Processing*, 2018.
- [25] IEEE, "754-2019 - IEEE Standard for Floating-Point Arithmetic," Tech. Rep., 2019.
- [26] Intel Corporation, "bfloat16 - Hardware Numerics Definition," Tech. Rep., 2018. [Online]. Available: <https://software.intel.com/content/www/us/en/develop/download/bfloat16-hardware-numerics-definition.html>
- [27] R. Theagarajan, F. Pala, and B. Bhanu, "EDeN: Ensemble of Deep Networks for Vehicle Classification," *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, 2017.
- [28] C. K. Chang, S. Lym, N. Kelly, M. B. Sullivan, and M. Erez, "Evaluating and accelerating high-fidelity error injection for HPC," *Proceedings - International Conference for High Performance Computing, Networking, Storage, and Analysis, SC 2018*, pp. 577–589, 2019.
- [29] A. Mahmoud, N. Aggarwal, A. Nobbe, J. R. Sanchez Vicarte, S. V. Adve, C. W. Fletcher, I. Frosio, and S. K. S. Hari, "Py-TorchFI: A Runtime Perturbation Tool for DNNs," in *DSN-DSML*, 2020.
- [30] Nvidia, "Cuda toolkit documentation," 2021. [Online]. Available: <https://docs.nvidia.com/cuda/floating-point/index.html>
- [31] A. Krämmer, C. Schöller, D. Gulati, and A. Knoll, "Providentia - A large scale sensing system for the assistance of autonomous vehicles," *arXiv*, 2019. [Online]. Available: [arxiv:1906.06789](https://arxiv.org/abs/1906.06789)