No Data, No Optimization: A Lightweight Method To Disrupt Neural Networks With Sign-Flips

Abstract

Deep Neural Networks (DNNs) can be catastrophically disrupted by flipping only a handful of sign bits in their parameters. We introduce Deep Neural Lesion (DNL), a data-free, lightweight method that locates these critical parameters and triggers massive accuracy drops. We validate its efficacy on a wide variety of computer vision models and datasets. The method requires no training data or optimization and can be carried out via common exploits software, firmware or hardware based attack vectors. An enhanced variant that uses a single forward and backward pass further amplifies the damage beyond DNL's zero-pass approach. Flipping just two sign bits in ResNet50 on ImageNet reduces accuracy by 99.8%. We also show that selectively protecting a small fraction of vulnerable sign bits provides a practical defense against such attacks.

Figure 1: DNL applied to RegNetY-400MF's Radosavovic et al. (2020a) first convolution layer. The original (Sobel-like) kernel, used for horizontal edge detection, is shown above the flipped version obtained by changing just one high-magnitude weight's sign bit. Even this minimal alteration leads to a drastically different output feature map. This corrupted feature propagates through the model, undermining downstream representations and severely impairing the network's overall ability to recognize the Dalmatian.

1 Introduction

Deep neural networks (DNNs) power a wide range of applications, including safety-critical tasks such as autonomous driving, unmanned aerial vehicle (UAV) navigation, medical diagnostics, and robotics, where real-time decision-making is essential. However, the increasing reliance on DNNs also raises concerns about

their resilience to malicious attacks. Ensuring the robustness of DNNs is crucial to maintaining their reliability in such critical applications.

In this paper we expose a critical vulnerability in DNNs that allows for severe disruption by flipping as few as one to ten sign bits, a tiny fraction of the model's parameters. Our method demonstrates how a small number of bit flips, within models containing up to hundred millions of parameters, can cause catastrophic degradation in performance.

We systematically analyze and identify the parameters most susceptible to sign flips, which we term "critical parameters." Our approach disrupts models by flipping the sign bits in their weights, achieving significant performance degradation with minimal computational effort across a wide range of model architectures and datasets. Crucially, our approach is data-agnostic: it requires only direct access to model weights, bypassing any need for training or validation data.

Our method is an extremely lightweight, forward-pass free (*Pass-free*) Attack, which needs no additional computational passes and is driven by a magnitude-based heuristic, considering inductive bias on learnable features and the flow of information within the networks. We also propose an enhanced *1-Pass* Attack that employs a single forward and backward pass (on random inputs) to refine the selection of critical parameters. Despite this small increase in computation, the attack remains highly efficient, preserves its data independence, and inflicts even greater damage on the model.

Malicious actors can exploit the identified parameter vulnerability through multiple system layers, including file-system intrusions, firmware compromises, direct memory access (DMA) from compromised peripherals, or memory-level exploits. In each case, once attackers gain access to the model's parameters, they can flip a small number of high-magnitude sign bits and trigger severe model failures. This lightweight approach requires no iterative optimization, thereby reducing the attacker's overhead while also increasing stealth.

For example, consider the implications for autonomous driving systems. Traditional adversarial attacks Madry et al. (2018); Goodfellow et al. (2015); Carlini & Wagner (2016) on such systems would involve manipulating the input pixels in real-time, requiring continuous communication with the vehicle and performing intensive gradient calculations to mislead the model. Physical adversarial attacks, such as placing adversarial stickers on street signs (e.g., as demonstrated in Wei et al. (2022)), demand direct access to the environment and are vulnerable to countermeasures like sensors on street signs or validation through additional traffic inputs. Other attacks targeting weights, such as Rakin et al. (2019); Park et al. (2021), rely either on the model's real data or on costly synthetic data generated from it, and demand an exhaustive search involving numerous forward and backward passes through the victim's model. Such attacks are impractical in real time because of their costly optimization process, and require the adversary to have full access to run the model as well as its data beforehand.

With our proposed method, even with limited access, an attacker could exploit any of the aforementioned hardware or software vulnerabilities and discreetly flip a small number of sign bits. By altering only a handful of critical parameters, often just one or two, the attacker can critically degrade the model's perception and decision-making, posing a far more potent threat to the reliability of autonomous driving systems. The minimal computational footprint and high impact of this attack make it exceptionally challenging to detect and mitigate in real-world deployments. All code will be made publicly available upon acceptance.

Our main contributions are summarized as follows:

The DNL Attack: We introduce the DNL attack, which exposes a severe vulnerability in DNNs by showing that heuristically flipping a small number of specific sign bits can catastrophically degrade model performance. This attack is entirely data-agnostic, requiring no knowledge of training data, domain-specific data, or synthetic inputs. Our method includes two lightweight variants: the "Pass-free" Attack, which operates without any additional computational passes, and the "1-Pass" Attack, which uses a single forward and backward pass with random inputs to enhance impact.

Characterization of Critical Parameters: We characterize the attributes that make certain parameters disproportionately vulnerable to bit flips—such as having high magnitude or being located in early layers—while showing that the vast majority of parameters remain robust to random perturbations.

Extensive Evaluation: We validate our approach on 60 classifiers across diverse tasks and datasets, including 48 ImageNet models from the publicly available timm Wightman (2019) and Torchvision Marcel & Rodriguez (2010) repositories. Flipping just a handful of bits— fewer than ten—is sufficient to significantly reduce accuracy, demonstrating the broad applicability and impact of this vulnerability.

Defense Evasion: We demonstrate that DNL circumvents various defenses, including binarization, redundancy-coding, and weight-scaling, underscoring the need for new protective strategies.

Defense Mechanisms: We leverage the insight gained from identifying critical parameters to propose efficient defenses. By selectively protecting only these most vulnerable parameters, models can become substantially more resilient to sign-flip attacks.

2 Problem Setup

Modern deep learning frameworks typically store parameters in the IEEE 754 32-bit floating-point format: one sign bit, eight exponent bits, and 23 mantissa bits: $(-1)^s \times 2^{(e-127)} \times \left(1 + \frac{m}{2^{23}}\right)$.

We focus on a standard supervised learning scenario where a model f_θ, is trained on a dataset with distribution $\mathcal{D}$. Let $\mathcal{X}, \mathcal{Y}$ be the input and label spaces, $(X, Y) \sim \mathcal{D}$, where $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$. A trained model f_θ seeks to minimize the expected risk $\min_\theta \mathbb{E}_{(X,Y)\sim\mathcal{D}}\Big[\mathcal{L}\big(f_\theta(X), Y\big)\Big]$, where $\mathcal{L}$ is a loss function.

Threat model. The attacker has write access to the stored parameters, but no access to any kind of data, nor passing anything through the model. Concretely: (i) the attacker has no samples from $\mathcal{D}$ and thus no access to $P(X)$ or $P(Y)$; and (ii) the attacker cannot evaluate the model on any input (no forward or backward passes). Equivalently, for any input random variable Z on $\mathcal{X}$ (including $Z = X$ or any synthetic/random data), the attacker has no access to $f_\theta(Z)$ or to $P\big(f_\theta(Z)\big)$ (for 1-Pass DNL we consider a slightly relaxed version where $f_\theta(z)$ could be observed for a single random input z, and allow a single forward pass and a single backward pass).

Despite this, the attacker can modify θ by flipping a small number of bits in its stored representation. Let bits$(\theta) \in \{0,1\}^B$ be the B memory bits encoding all entries of θ (e.g., IEEE-754). A *k-bit flip* chooses k distinct bit indices $j_1, \ldots, j_k \in \{1, \ldots, B\}$ and produces parameters $\theta'_{(k)}$ with

$$\text{bits}\big(\theta'_{(k)}\big)_j = \begin{cases} 1 - \text{bits}(\theta)_j, & \text{if } j \in \{j_1, \ldots, j_k\}, \\ \text{bits}(\theta)_j, & \text{otherwise.} \end{cases}$$

We refer to $\theta'_{(k)}$ as the result of a k-bit-flip attack on θ. To our knowledge, no prior method satisfies this restrictive threat model.

Mechanisms enabling parameter bit flips. The attacker can gain access to θ directly through software, firmware, or hardware-level exploits, namely Bit flip attacks Oriyano (2014). Below, we outline several exploits that adversaries can leverage to execute malicious bit-flipping operations on model parameters.
A *rootkit* (Hoglund & Butler, 2006; Sparks & Butler, 2005; Rutkowska, 2007) is malicious software running with high-level (kernel or ring-0) privileges, allowing it to intercept or modify operations. Once installed, a rootkit can scan the system's memory or storage for the model's parameter files, then surgically flip bits in place. By concealing its processes and hooking system APIs, the rootkit can evade detection from common antivirus tools and monitoring systems, enabling stealthy, ongoing tampering with model parameters without triggering suspicious activity logs.
Firmware exploits (Hudson & Rudolph, 2015) (e.g., SSD/HDD controllers, GPU firmware, BIOS, or microcode patches) can give attackers privileged memory access or the ability to inject custom commands that flip bits in system memory or on storage media. By compromising firmware updates or exploiting known bugs, attackers can precisely manipulate parameter bits.
DMA from untrustworthy peripherals (Markettos et al., 2019) can read and write system memory without involving the CPU or the operating system's normal access controls. If attackers gain low-level access to a DMA device (e.g., via Thunderbolt or FireWire interfaces), they can directly overwrite targeted bits in

protected memory regions.

Rowhammer (Kim et al., 2014b;a; Seaborn & Dullien, 2015) exploits the electrical interference between neighboring rows in modern DRAM modules. By rapidly accessing ("hammering") one row, an attacker causes bits in adjacent rows to flip, even without direct write permissions. Rowhammer attacks typically rely on high-frequency memory accesses that defeat standard refresh mechanisms; once carefully controlled, these flips can be directed at specific bit positions.

GPU cache tampering (Lipp et al., 2020; Tatar et al., 2018), which exploits a compromised kernel driver or malicious GPU code, can manipulate cache management routines to induce bit flips in stored parameters. Similar to Rowhammer's repeated DRAM accesses, continuously evicting and reloading specific cache lines may corrupt targeted parameters. Because GPU caches are often less scrutinized than CPU caches, this tampering can remain undetected, leading to stealthy yet severe degradation of model performance.

Voltage/frequency glitching (Murdock et al., 2020; Tang et al., 2017; Van der Veen et al., 2020; Frigo et al., 2018) manipulates the operating voltage or clock frequencies to induce computational errors. Certain voltage ranges can systematically cause specific bits to flip in registers or memory segments.

Adversarial objective. In all cases, the attacker's objective is to significantly degrade performance with minimal bit flips for stealth and practicality. We define the objective as: $\min_k \max \mathbb{E}_{(X,Y)\sim\mathcal{D}}\left[\mathcal{L}\big(f_{\theta'_{(k)}}(X), Y\big)\right]$, where both finding minimal k and flipping k bits to produce $\theta'_{(k)}$ are discrete optimization problems.

In other words, the attacker's goal is to induce a significant performance drop while flipping only a handful of bits, both for stealth and practical reasons, as fewer corruptions are less likely to be detected and can be exploited by the mentioned hardware attacks. For instance, Rowhammer-based exploits (Kim et al., 2014b) typically induce only sporadic bit upsets in adjacent cells, making massive coordinated flips infeasible. By stealth, we mean that the modifications to the model weights (or inputs) are minimal—while the victim may observe a performance drop and even suspect an attack, the lack of an identifiable source makes it difficult to attribute the degradation or take effective countermeasures.

As mentioned above, the attacker does *not* have access to any training or validation data, nor do they conduct extensive inference passes or iterative gradient-based searches. Such lightweight attacks are realistic in settings where the attacker's computational resources on the victim device are minimal, or where repeated forward/backward passes might raise suspicion. We therefore distinguish two scenarios: a *pass-free* attack, which uses no extra computation beyond reading and writing into the model weights (fitting our restrictive threat model), and a *1-pass* attack, which uses only a single forward (and backward) pass on a single random input, aiming for improved efficacy at a slight relaxation of the threat model. Both settings stand in contrast to existing approaches that require data samples and multiple optimization steps (see Section 5 for more details).

Although the attacker's objective can be framed as a discrete optimization problem—finding the smallest set of bits whose flips induce the greatest performance drop—exhaustive searches over millions of parameters are computationally infeasible in real-time and cannot scale with larger models. Instead, lightweight heuristics can pinpoint "critical" parameters while incurring minimal overhead. By leveraging inductive insights into how information flows through the network, an attacker can disrupt the most influential parameters without iterative optimization or a large dataset. This stands in contrast to data-driven or gradient-based methods, which demand multiple inference passes and raise both computational requirements and the risk of detection.

2.1 Accuracy Reduction Metrics

To measure the effect of bit flips, let $\theta'_{(k)}$ be the set of parameters obtained by flipping exactly k bits in θ. If $\text{Acc}(\theta)$ is the model's original accuracy, we define: $\text{AR}(k) = \frac{\text{Acc}(\theta) - \text{Acc}(\theta'_{(k)})}{\text{Acc}(\theta)}$, which captures the drop in accuracy induced by k flips. For a broader view, we also define:

$$\text{mAR}(N) = \frac{1}{N}\sum_{k=1}^{N}\text{AR}(k), \tag{1}$$

(a) AR of randomly flipping sign bits on model performance.

(b) Comparison of mAR_{10} across different strategies.

(c) Comparison of $AR(10)$ under different strategies.

Figure 2: Evaluation of model degradation under different sign-flip strategies across 48 ImageNet models.

so that a single number can represent the model's overall vulnerability across different flip counts. Because practical hardware attacks often manage only a handful of flips, we mainly focus on small k (e.g., $k \leq 10$).

3 Locating Models' Most Critical Parameters

Targeting sign bits: Considering the FP32 representation, while exponent flips can alter a weight's magnitude, flipping the most significant *sign* bit instantly switches a parameter from positive to negative (or vice versa), our experiments show sign-bit flips remain the most devastating (see Appendix C). This produces a drastic effect on the learned features, as shown in Figure 1, motivating us to focus on those bits. Moreover, localizing the sign bit in memory is straightforward (e.g., always the MSB), making it appealing as a simple target for adversaries. Various hardware-based studies show that repeated access patterns more reliably flip the *same* bit position across different addresses than arbitrarily chosen bits Kim et al. (2014b); Seaborn & Dullien (2015). Hence, focusing on sign bits aligns with how hardware attacks often achieve consistent flips in a specific bit offset across multiple weights, increasing the chance of our targeted attack success rate.

Flipping random sign bits in a network's parameters typically has a negligible impact on performance. Indeed, our experiments (visualized in Figure 2a) show that for many architectures, flipping even up to $100,000$ bits does not reduce the accuracy consistently—indicating that most parameters are not "critical." These findings motivate a more targeted strategy to identify and flip only the most sensitive parameters.

Magnitude-Based Strategy: Drawing inspiration from the pruning literature, we first examine magnitude-based strategies. Just as magnitude pruning removes low-magnitude weights to minimize the impact on final predictions Frankle & Carbin (2018), we hypothesize that flipping the sign of *high-magnitude* parameters causes significant disruption. Formally, the parameter score function is defined as follows

$$\mathcal{S}(\theta_i) \;=\; |\theta_i| \tag{2}$$

As far as we are aware, this work is the first to evaluate the efficacy of a magnitude-based attack, a surprisingly simple yet powerful strategy that disrupts neural networks without data, optimization or prior knowledge. In Figure 2b, the second boxplot from the left, shows that focusing on the top-k largest weights (in absolute value) significantly disrupts most evaluated models.

One-Flip-Per-Kernel Constraint: Empirical analyses of CNN filters (Krizhevsky et al., 2012; Zeiler & Fergus, 2014) highlight the importance of early-stage kernels (e.g., Gabor-like or Sobel-like) in extracting fundamental visual features. These studies reveal that flipping a single sign bit in a kernel can completely disrupt its feature extraction capability, altering the information the model relies on (see Figure 1 for the effect of sign flips on a real kernel). However, flipping multiple bits within the same kernel often merely changes its orientation or slightly modifies its functionality, rather than fully destroying the feature, as demonstrated in Figure 3. We observe this phenomenon consistently across multiple architectures. See Appendix B for examples.

To maximize damage, we constrain the attack to flip exactly *one* bit per kernel, ensuring the disruption does not offset itself and affects a broader range of features. Given our focus on a small number of flips, distributing them across more kernels also helps amplify the overall impact.

Figure 3: Horizontal edge detection filter (based on the Sobel Y filter) with one or two sign flips and their corresponding extracted features. With a single sign flip, the filter is severely disrupted, rendering it unable to detect edges effectively. However, with two bit flips, the resulting errors may partially offset each other, allowing the filter to retain some edge-detection capability and produce features similar to the original.

Targeted	1	2	3	5
All Layers	.13	.15	.24	.39
First 100	.19	.24	.42	.48
First 10	**93.9**	**99.6**	99.6	**99.8**
First 5	93.9	99.6	**99.7**	99.7
First 2	93.9	99.6	**99.7**	99.7
First 1	59.5	53.7	82.4	88.5
Last 10	.01	.04	.06	.17
Last 5	.03	.19	.46	1.14

(a) First l layers vs. mAR_{10}.

(b) ShuffleNetV2 layer targets ($AR\%$).

Figure 4: Layer-specific vulnerability: global trend (left) and detailed case (right).

Layer Selection   Beyond which parameters to flip, we also investigate *where* in the network to apply the attack. One might intuitively expect that targeting *final* layers—being closer to the classifier—would cause greater damage. However, our experiments reveal that in many architectures, early-layer manipulations are disproportionately damaging. Drawing on an analogy from neuroscience, early lesions (e.g., in the retina or optic nerve) can cause severe or total blindness (Kandel et al., 2000; Stewart et al., 2020; Essen et al., 1992). Similarly, flipping a single parameter in a fundamental feature detector (e.g., Sobel and Gabor filters) sends erroneous signals throughout subsequent layers, often leading to compounding error. Figure 1 illustrates this: a sign flip in a low-level "edge-detection" filter causes the network to misinterpret critical structural cues, compounding errors to later layers and severely degrading performance —more so than flips occurring in higher-level layers. Moreover, early convolutional filters encode generic edge and texture features; disrupting them degrades all downstream representations. This aligns with pruning evidence that early layers are disproportionately salient Frankle & Carbin (2018); Liu et al. (2019).

Interestingly, for most models evaluated, the largest parameters (*in absolute value*) tend to concentrate in these early layers. However, many models such as ShuffleNetV2 (Ma et al., 2018) exhibit a different pattern: their largest parameters are concentrated in later layers. As a result, naive attacks that always target the largest parameters—often located in the late layers of ShuffleNetV2—are less effective. Redirecting the attack to early layers, however, significantly amplifies the damage (see Table 4b for quantitative details).

Theoretical motivation. We can interpret our sign-flip attack through the same quadratic loss model long used to justify pruning criteria such as Optimal Brain Damage (OBD) LeCun et al. (1989b) and Optimal Brain Surgeon (OBS) Hassibi et al. (1992). For a trained network with parameters θ and loss $\mathcal{R}(\theta)$, a second-order Taylor expansion gives

$$\Delta\mathcal{R} \approx g^\top \Delta\theta + \tfrac{1}{2}\Delta\theta^\top H \Delta\theta,$$

where $g = \nabla_\theta \mathcal{R}(\theta)$ and H is the Hessian. At convergence, $g \approx 0$, so curvature dominates. Flipping the sign of weight θ_i yields $\Delta\theta_i = -2\theta_i$, with all other coordinates unchanged, giving under a diagonal-Hessian approximation:

$$\Delta\mathcal{R}_i \approx \tfrac{1}{2}(-2\theta_i)^2 H_{ii} = 2\,\theta_i^2 H_{ii}.$$

Thus, with a budget of k flips, the greedy maximizer is the set of k indices with largest $\theta_i^2 H_{ii}$. If H_{ii} is approximately constant within a layer (empirically common in early convolutional layers), this reduces to choosing the k largest $|\theta_i|$, precisely our *zero-pass* criterion. Equivalently, if $H \succeq \mu I$, then

$$\Delta\mathcal{R} \geq 2\mu \sum_{i \in S} \theta_i^2,$$

so picking the largest magnitudes maximizes a certified lower bound on loss damage.

Our *one-pass* variant (1P-DNL) refines this by replacing H_{ii} with Gauss–Newton style estimates derived from gradients, recovering the classical Taylor saliency $\propto |\theta_i g_i|$ used in pruning (Molchanov et al., 2017; Lee et al., 2019; Wang et al., 2020; Tanaka et al., 2020). Hence, magnitude- and gradient-based flips correspond to adversarial analogues of the very same criteria known to predict weight importance.

Early-layer targeting. Perturbations in early layers propagate through all subsequent transformations with multiplicative gain bounded by $\prod_{\ell>1} L_\ell$, where L_ℓ is the Lipschitz constant of layer ℓ, since the Lipschitz constant of a composition is at most the product of the layer-wise constants (Burago et al., 2001; Gouk et al., 2021; Virmaux & Scaman, 2018, Prop. 1.4.3). This amplifies their effect relative to late-layer flips.

The Deep Neural Lesion Pass-free algorithm: Based on these observations, we explore a simple heuristic algorithm that flips bits only in the first l layers of a network (with $1 \leq l \leq 10$). Algorithm 1 summarizes our Pass-free attack for a given model, number of sign flips k, and layers l. We find that any l in this range consistently degrades accuracy more than random or purely magnitude-based strategies (Figure 4a). We select $l = 10$ for simplicity and only consider the parameters of those layers as candidates for sign flips.

Algorithm 1 Deep Neural Lesion (DNL) – Pass-free Attack

1: **Inputs:** Model parameters θ, number of bits to flip k, number of layers L
2: $\theta_L \leftarrow$ parameters in the first L layers of θ
3: Sort θ_L in descending order by $|\theta_i|$
4: $\mathcal{K} \leftarrow$ top-k entries of θ_L (one per kernel)
5: **for** each θ_i in $\mathcal{K}$ **do**
6: $\quad \theta_i \leftarrow -\theta_i$ // flip sign bit
7: **end for**
8: **Output:** Modified parameters θ

3.1 Enhanced Attack Using a Single Forward Pass

When a single forward (and backward) pass is within the attacker's budget, we propose an enhanced attack, called 1P-DNL, inspired by *gradient-based pruning* methods (LeCun et al., 1989a; Mozer & Smolensky, 1988; Lee et al., 2019; Wang et al., 2020; Tanaka et al., 2020). These methods typically assign a saliency or importance score to each parameter θ_i by measuring how altering that parameter (e.g., pruning or modifying it) would affect the network's loss or outputs. Although pruning and sign-flip attacks differ in goal, the underlying idea of identifying the "most critical" weights is similar.

Hybrid Importance Score

We define a hybrid importance scoring function that combines magnitude-based saliency with second-order information. Specifically, we start from a second-order Taylor expansion of a scalar loss function $\mathcal{R}(\theta)$ around θ_i (LeCun et al., 1989a; Hassibi et al., 1992). Let α and β be tunable coefficients controlling the relative weight of magnitude- and gradient-based terms. For a given parameter θ_i,

$$\mathcal{S}(\theta_i) = \alpha \, |\theta_i| + \beta \left| \frac{\partial \mathcal{R}}{\partial \theta_i} \, \theta_i + \frac{1}{2} H_{ii} \, \theta_i^2 + \sum_{j \neq i} H_{ij} \, \theta_i \, \theta_j \right|, \qquad (3)$$

where H is the Hessian of $\mathcal{R}$ with respect to θ. In our case, we let $\alpha = \beta = 1$, $\mathcal{R}(\theta) = \sum_i c_i$, where c_i is the output logit (or class score) for a Gaussian input. Although the summation over $j \neq i$ captures inter-weight coupling, we approximate $H_{ij} = 0$ for $j \neq i$ (a common diagonal approximation in second-order pruning (LeCun et al., 1989a)), significantly reducing computation. Similarly, we replace H_{ii} by $(\frac{\partial \mathcal{R}}{\partial \theta_i})^2$ (i.e., a Gauss-Newton like approximation), which further simplifies Hessian-based estimation.

- If $\frac{\partial \mathcal{R}}{\partial \theta_i} = 0$ and $H_{ii} = 0$, Eq. equation 3 reduces to: $\mathcal{S}(\theta_i) = \alpha \, |\theta_i|$, mirroring a simple magnitude-based saliency score (identical to Equation 2).

- If $\alpha = 0$, we recover a purely second-order (Optimal Brain Damage-like) approach:

$\mathcal{S}(\theta_i) = \beta \left| \frac{\partial \mathcal{R}}{\partial \theta_i} \, \theta_i + \frac{1}{2} H_{ii} \, \theta_i^2 \right|$, which focuses on changes in $\mathcal{R}$ under small parameter perturbations.

Although one forward/backward pass on random (Gaussian) data might be required to estimate $\mathcal{S}$, it remains significantly simpler than full data-driven optimization-based attacks (e.g., iterative gradient-based bit-flips). Figure 2b shows that incorporating second-order signals consistently amplifies the attack's damage compared to purely magnitude-based methods. Consequently, this hybrid scoring approach yields a more powerful single-pass sign-flip attack in scenarios where the attacker can run a forward and backward pass on the architecture, yet does not have access to the original training set. To summarize 1P-DNL, we refer the reader to Algorithm 2. Figure 2c shows the impact of all previously suggested methods with 10 sign flips. Both DNL and 1P-DNL cause most models to collapse, with 43 out of 48 models exhibiting an accuracy reduction above 60%. Finally, in Appendix D we compare 1P-DNL with various other 1-pass methods from the weight pruning literature to find critical parameters and find 1P-DNL the most potent.

4 Additional Analysis

We further assess our sign-flip attack on various datasets to confirm its broad applicability beyond ImageNet. In particular, we evaluate DNL and 1P-DNL on DTD (Cimpoi et al., 2014), FGVC-Aircraft (Maji et al., 2013), Food101 (Bossard et al., 2014), and Stanford Cars (Krause et al., 2015). In Figure 5 and Figure 6, we plot the average accuracy reduction across EfficientNetB0 (Tan & Le, 2019), MobileNetV3-Large, and ResNet-50 He et al. (2015) after applying DNL and 1P-DNL respectively. In all cases, flipping as few as one or two sign bits leads to a sudden collapse in model performance, reaffirming that this vulnerability is not tied to a specific dataset or image distribution. Most notably, applying DNL disrupts all models across all datasets with $AR(5) \geq 85\%$, and 1P-DNL achieve $AR(4) \geq 90\%$ across all data and models. For additional evaluations on each dataset, see Appendix E.

Taken together, these results underscore both the potency and the generality of our approach: across diverse vision tasks and architectures, zero or one-pass sign-flip attacks consistently induce dramatic failures with minimal computational overhead or data requirements.

Impact of Model Size on Attack Success: To assess whether model size influences the effectiveness of our attacks, we evaluate both DNL and 1P-DNL, across five families of architectures with varying parameter counts: ResNet, RegNet, EfficientNet, ConvNeXt (Liu et al., 2022), and ViT (Dosovitskiy et al., 2020). The results, summarized in Figure 12 and Figure 13, reveal that model size does not exhibit a clear correlation with attack susceptibility. Most models collapse at similar levels regardless of their scale, demonstrating that our attack effectively scales to model size and disrupts networks of all sizes.

Figure 5: Averaged AR (%) of DNL over Efficient-NetB0, MobileNetV3-Large, and ResNet-50 vs number of sign flips. Each color represents a different dataset, confirming the impact of our pass-free attack on DTD, FGVC-Aircraft, Food101, and Stanford Cars.

Figure 6: Averaged AR (%) of 1P-DNL over EfficientNetB0, MobileNetV3-Large, and ResNet-50 vs number of sign flips. Each color represents a different dataset, confirming the fatality of our single-pass attack on DTD, FGVC-Aircraft, Food101, and Stanford Cars.

5 Comparison to Other Weight Attacks

All prior methods rely on considerable optimization efforts, typically requiring multiple forward and backward passes through the network, and in most cases, access to data samples to compute gradients. Consequently, these attacks fall outside the scope of our restrictive threat model. Early works such as *Terminal Brain Damage* (TBD) Hong et al. (2019) illustrated how manipulating exponent bits could severely harm floating-point networks. However, TBD excludes sign bits, which we find can be far more devastating to overall accuracy with fewer flips (see Appendix C). Other methods, including (Rakin et al., 2019; Yao et al., 2020), perform iterative gradient-based flips. For example, Rakin et al. (2019) requires multiple samples to compute gradients and can disrupt ResNet-50's accuracy by $\sim 99.7\%$ using 11 bit flips. Yao et al. (2020) similarly needs iterative optimization, reaching significant disruption at the cost of 23 flips.

Recent variants have attempted to relax data requirements. For instance, (Ghavami et al., 2021; Park et al., 2021) generate pseudo-samples or use partial data statistics to guide which bits to flip. Although they lessen the need for a large labeled dataset, they still rely on model feedback or approximate gradients. In contrast, our sign bit flipping approach is lightweight, data-agnostic, and can degrade a large variety of networks by over 99.8% with few flips.

Table 1 compares bit-flip attacks on ImageNet-1K INT8 quantized models, and highlights how our approach differs from prior methods. BFA Rakin et al. (2019) and DeepHammer Yao et al. (2020) require iterative gradient-based searches and at least a few validation images. ZeBRA Park et al. (2021) drops the data requirement but still performs an optimization loop (and invests further compute into generating synthetic data). In stark contrast, our attacks can be carried out without data or optimization, and are straightforward to locate in memory, making them both feasible and devastating in real-world scenarios, yet equal or exceed prior art in accuracy reduction while using fewer bit flips. For example, 1P-DNL collapses ResNet-50 by 99.4% with a *single* sign flip.

6 Defenses and Counter-Measures

Selective Defense Against Sign-Flips. A straightforward mitigation is to keep several full copies of every *sign* bit and take a majority vote at inference. Because an attacker would then have to corrupt most replicas *simultaneously*, the method is robust—but it multiplies memory and bandwidth.

A leaner alternative uses *error-correcting codes* (ECC) such as Hamming codes (Peterson & Weldon, 1972). Single-bit ECC can detect and fix isolated flips automatically, yet scaling to many parameters demands stronger (and costlier) codes. The key observation from Section 3 is that *only a tiny subset of sign bits is*

Table 1: Bit-flip attacks on ImageNet-1K. Each cell lists *# flips → AR (%)*. OF = optimization-free, DA = data-agnostic. * best of 5 trials.

Model	Method	OF	DA	# Flips → AR (%)
VGG-11	BFA Rakin et al. (2019)	✗	✗	17 → 99.7
	ZeBRA Park et al. (2021)	✗	✓	8 → 99.8
	DNL	✓	✓	3 → **99.9**
	1P-DNL	✓	✓	2 → 99.8
ResNet-50	DeepHammer Yao et al. (2020)	✗	✗	23* → 75.4
	BFA Rakin et al. (2019)	✗	✗	1 → 5.5, 5 → 99.7
	ZeBRA Park et al. (2021)	✗	✓	1 → 7.7, 5 → 99.7
	DNL	✓	✓	1 → 6.6, 8 → 99.7
	1P-DNL	✓	✓	1 → **99.4**
MobileNet-V2	DeepHammer Yao et al. (2020)	✗	✗	2* → **99.8**
	BFA Rakin et al. (2019)	✗	✗	3 → 99.8
	ZeBRA Park et al. (2021)	✗	✓	2 → 99.7
	DNL	✓	✓	2 → 99.8
	1P-DNL	✓	✓	2 → **99.9**
ViT-B/16@224	BFA Rakin et al. (2019)	✗	✗	5 → 30.1, 10 → 90.9
	ZeBRA Park et al. (2021)	✗	✓	5 → 5.1, 10 → 45.8
	DNL	✓	✓	5 → **99.3**
	1P-DNL	✓	✓	4 → 99.1

truly catastrophic. Hence we can protect just those high-scoring weights—identified by DNL—either with bit replication or ECC, while leaving the vast majority of bits unguarded.

Selective DNL Defense. We tested DNL weight selection to defend against the iterative Bit-Flip Attack (BFA) (Rakin et al., 2019).

Table 2 reports the mean accuracy reduction after the attacker is allowed up to ten flips ($k \in [1, 10]$, three runs per k). Shielding only **0.001**% of parameters already halves BFA's impact on ResNet-18 and ResNet-50; guarding **1**% of parameters nullifies the attack on every model we tried.

While these experiments illustrate a simple defense, their greater significance lies in showing that DNL reliably identifies the most critical parameters—the very ones exhaustive BFA seeks to corrupt.

An expanded evaluation on different attack strategies with results on the sixteen most vulnerable architectures is provided in Appendix H.

Table 2: Effect of (DNL Defense) against BFA.

Model	# Defended params	BFA AR(10)	Model	# Defended params	BFA AR(10)
ResNet-18	No Defense	88.87	ResNet-50	No Defense	93.87
	~ 0.001% (100 params)	58.83		~ 0.001% (250 params)	39.08
	~ 1% (100 K params)	**0.00**		~ 1% (250 K params)	**1.30**
MobileNet-V2	No Defense	99.90	ViT-B/16@224	No Defense	82.30
	~ 0.001% (30 params)	99.80		~ 0.001% (900 params)	40.51
	~ 1% (30 K params)	**44.30**		~ 1% (900 K params)	**0.21**

Evaluating DNL against existing defenses: We evaluated representative defense strategies that have been proposed to prevent bit-flip attacks and found that DNL (and its 1-pass variant) either fully or largely bypass them.

Encoding defenses.

DeepNcode (Velcický et al., 2024) defends against parameter corruption by encoding each floating-point weight into a longer binary codeword with redundancy, typically using a codebook with Hamming distance > 1 between valid codewords. During inference, the stored codeword is decoded back into a floating-point

value; if a bit flip occurs, the decoder maps the corrupted codeword to the nearest valid one, thereby correcting small errors.

However, this protection assumes the attacker cannot deliberately steer the corrupted codeword toward a *different valid codeword.* In a realistic *gray-box* setting—the attacker does not know the codebook, but can observe the resulting decoded values—we can exploit this decoding step. Specifically, by selectively flipping bits in the encoded representation, we search for the closest alternative codeword whose decoded value has the *opposite sign.* This effectively performs a sign flip *through* the encoding, bypassing the correction capability of the decoder.

Weight-scaling defenses. Weight-scaling (Fuengfusin & Tamukoh, 2024) multiplies all stored parameters by a constant c and divides by c at inference, damping additive perturbations by the factor c. Sign flips, however, are multiplicative ($\theta \mapsto -\theta$); after the rescaling they remain $-\theta$, exactly as in the undefended model: flip: $\theta \mapsto -\theta \implies$ defended: $\frac{-c\theta}{c} = -\theta$ Hence, the scaling defense has *no effect* on DNL, which is empirically confirmed by unchanged AR.

7 Concluding Remarks

This work exposes a fundamental vulnerability in deep neural networks: their susceptibility to severe disruption through minimal, computation-free interventions. We introduced a lightweight, data-free method for locating and flipping critical parameters, and validated its effectiveness across a wide range of architectures and datasets. Building on these insights, we also proposed a targeted defense that selectively protects vulnerable parameters and substantially improves model robustness to these attacks. Future work could explore architectures and training procedures that increase resistance to lightweight parameter manipulations.

References

Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – mining discriminative components with random forests. In *European Conference on Computer Vision*, 2014.

Dmitri Burago, Yuri Burago, and Sergei Ivanov. *A Course in Metric Geometry*, volume 33 of *Graduate Studies in Mathematics.* American Mathematical Society, 2001.

Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks, 2016.

M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, , and A. Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2020.

David C. Van Essen, Charles H. Anderson, and Daniel J. Felleman. Information processing in the primate visual system: An integrated systems perspective. *Science*, 255(5043):419–423, 1992. doi: 10.1126/science.1734518. URL https://www.science.org/doi/abs/10.1126/science.1734518.

Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks, 2018.

Pietro Frigo, Cristiano Giuffrida, Herbert Bos, and Kaveh Razavi. Grand pwning unit: Accelerating microarchitectural attacks with the gpu. In *2018 IEEE Symposium on Security and Privacy (SP)*, pp. 195–210, 2018. doi: 10.1109/SP.2018.00022.

Ninnart Fuengfusin and Hakaru Tamukoh. Harden deep neural networks against fault injections through weight scaling. *CoRR*, abs/2411.18993, 2024. doi: 10.48550/ARXIV.2411.18993. URL https://doi.org/10.48550/arXiv.2411.18993.

Behnam Ghavami, Mani Sadati, Mohammad Shahidzadeh, Zhenman Fang, and Lesley Shannon. Bdfa: A blind data adversarial bit-flip attack on deep neural networks, 2021.

Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL `http://arxiv.org/abs/1412.6572`.

Henry Gouk, Eibe Frank, Bernhard Pfahringer, and Michael J. Cree. Regularisation of neural networks by enforcing lipschitz continuity. *Machine Learning*, 110(2):393–416, 2021. doi: 10.1007/s10994-020-05929-w.

Hassibi, Babak, Stork, and David. Second order derivatives for network pruning: Optimal brain surgeon. In *Advances in Neural Information Processing Systems*, 1992. URL `https://proceedings.neurips.cc/paper_files/paper/1992/file/303ed4c69846ab36c2904d3ba8573050-Paper.pdf`.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.

Zhezhi He, Adnan Siraj Rakin, Jingtao Li, Chaitali Chakrabarti, and Deliang Fan. Defending and harnessing the bit-flip based adversarial weight attack. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14083–14091, 2020. doi: 10.1109/CVPR42600.2020.01410.

Greg Hoglund and Jamie Butler. *Rootkits: Subverting the Windows Kernel*. Addison-Wesley Professional, 2006. ISBN 978-0-321-29431-0.

Sanghyun Hong, Pietro Frigo, Yiğitcan Kaya, Cristiano Giuffrida, and Tudor Dumitraş. Terminal brain damage: exposing the graceless degradation in deep neural networks under hardware fault attacks. In *Proceedings of the 28th USENIX Conference on Security Symposium*, SEC'19, pp. 497–514, USA, 2019. USENIX Association. ISBN 9781939133069.

Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, Quoc V. Le, and Hartwig Adam. Searching for mobilenetv3, 2019.

Trammell Hudson and Larry Rudolph. Thunderstrike: Efi firmware bootkits for apple macbooks. In *Proceedings of the 8th ACM International Systems and Storage Conference*, SYSTOR '15, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450336079. doi: 10.1145/2757667.2757673. URL `https://doi.org/10.1145/2757667.2757673`.

E.R. Kandel, J.H. Schwartz, and T. Jessell. *Principles of Neural Science, Fourth Edition*. McGraw-Hill Companies,Incorporated, 2000. ISBN 9780838577011. URL `https://books.google.co.il/books?id=yzEFK7Xc87YC`.

Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Cheng-Yao Wilkerson, Konrad Lai, and Onur Mutlu. Flipping bits in memory without accessing them: An experimental study of DRAM disturbance errors. In *Proceedings of the 41st Annual International Symposium on Computer Architecture (ISCA)*, pp. 361–372. IEEE/ACM, 2014a. doi: 10.1109/ISCA.2014.6853210.

Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu. Flipping bits in memory without accessing them: An experimental study of dram disturbance errors. In *2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA)*, pp. 361–372, 2014b. doi: 10.1109/ISCA.2014.6853210.

Jonathan Krause, Hailin Jin, Jianchao Yang, and Li Fei-Fei. Fine-grained recognition without part annotations. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5546–5555, 2015. doi: 10.1109/CVPR.2015.7299194.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'12, pp. 1097–1105, Red Hook, NY, USA, 2012. Curran Associates Inc.

Pat Langley. Placeholder title. *Placeholder Journal*, 1(1):1–10, 2000.

Yann LeCun, John Denker, and Sara Solla. Optimal brain damage. In D. Touretzky (ed.), *Advances in Neural Information Processing Systems*, volume 2. Morgan-Kaufmann, 1989a. URL `https://proceedings.neurips.cc/paper_files/paper/1989/file/6c9882bbac1c7093bd25041881277658-Paper.pdf`.

Yann LeCun, John S. Denker, and Sara A. Solla. Optimal brain damage. In David S. Touretzky (ed.), *Advances in Neural Information Processing Systems 2, [NIPS Conference, Denver, Colorado, USA, November 27-30, 1989]*, pp. 598–605. Morgan Kaufmann, 1989b. URL `http://papers.nips.cc/paper/250-optimal-brain-damage`.

Namhoon Lee, Thalaiyasingam Ajanthan, and Philip Torr. SNIP: SINGLE-SHOT NETWORK PRUNING BASED ON CONNECTION SENSITIVITY. In *International Conference on Learning Representations*, 2019.

Moritz Lipp, Michael Schwarz, Lukas Raab, Lukas Lamster, Misiker Tadesse Aga, Clementine Maurice, and Daniel Gruss. Nethammer: Inducing rowhammer faults through network requests. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroSSamp;PW)*. IEEE, September 2020. doi: 10.1109/eurospw51379.2020. 00102. URL http://dx.doi.org/10.1109/EuroSPW51379.2020.00102.

Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning, 2019. URL https://arxiv.org/abs/1810.05270.

Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s, 2022.

Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design, 2018.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.

S. Maji, J. Kannala, E. Rahtu, M. Blaschko, and A. Vedaldi. Fine-grained visual classification of aircraft, 2013.

Sébastien Marcel and Yann Rodriguez. Torchvision the machine-vision package of torch. In *Proceedings of the 18th ACM International Conference on Multimedia*, MM '10, pp. 1485–1488, New York, NY, USA, 2010. Association for Computing Machinery. ISBN 9781605589336. doi: 10.1145/1873951.1874254. URL https://doi.org/10.1145/1873951.1874254.

A. Theodore Markettos, Colin Rothwell, Brett F. Gutstein, Allison Pearce, Peter G. Neumann, Simon W. Moore, and Robert N. M. Watson. Thunderclap: Exploring vulnerabilities in operating system iommu protection via dma from untrustworthy peripherals. *Proceedings 2019 Network and Distributed System Security Symposium*, 2019.

Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional neural networks for resource efficient inference. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL https://openreview.net/forum?id=SJGCiw5gl.

Michael C Mozer and Paul Smolensky. Skeletonization: A technique for trimming the fat from a network via relevance assessment. In *Advances in Neural Information Processing Systems*. Morgan-Kaufmann, 1988. URL https://proceedings.neurips.cc/paper_files/paper/1988/file/07e1cd7dca89a1678042477183b7ac3f-Paper.pdf.

Kit Murdock, David Oswald, Flavio D. Garcia, Jo Van Bulck, Daniel Gruss, and Frank Piessens. Plundervolt: Software-based fault injection attacks against intel sgx. In *2020 IEEE Symposium on Security and Privacy (SP)*, pp. 1466–1482, 2020. doi: 10.1109/SP40000.2020.00057.

Sean-Philip Oriyano. *CEH: Certified Ethical Hacker Version 8 Study Guide*. SYBEX Inc., USA, 1st edition, 2014. ISBN 111864767X.

Dahoon Park, Kon-Woo Kwon, Sunghoon Im, and Jaeha Kung. Zebra: Precisely destroying neural networks with zero-data based repeated bit flip attack, 2021.

W. Wesley Peterson and E. J. Weldon. *Error-Correcting Codes*. MIT Press, Cambridge, MA, 2nd edition, 1972.

Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollar. Designing network design spaces. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2020a. doi: 10.1109/cvpr42600.2020.01044. URL http://dx.doi.org/10.1109/cvpr42600.2020.01044.

Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces, 2020b.

Adnan Siraj Rakin, Zhezhi He, and Deliang Fan. Bit-flip attack: Crushing neural network with progressive bit search. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 1211–1220. IEEE, October 2019. doi: 10.1109/iccv.2019.00130. URL http://dx.doi.org/10.1109/ICCV.2019.00130.

Adnan Siraj Rakin, Li Yang, Jingtao Li, Fan Yao, Chaitali Chakrabarti, Yu Cao, Jae sun Seo, and Deliang Fan. Rabnn: Constructing robust & accurate binary neural network to simultaneously defend adversarial bit-flip attack and improve accuracy, 2021.

Joanna Rutkowska. Beyond the CPU: Defeating hardware-based RAM acquisition. Black Hat USA, 2007.

Mark Seaborn and Thomas Dullien. Exploiting the dram rowhammer bug to gain kernel privileges. Black Hat USA, August 2015. URL https://googleprojectzero.blogspot.com/2015/03/exploiting-dram-rowhammer-bug-to-gain.html.

Sherri Sparks and Jamie Butler. Shadow walker: Raising the bar for rootkit detection. In *Black Hat Federal*, 2005.

Emma E. M. Stewart, Matteo Valsecchi, and Alexander C. Schütz. A review of interactions between peripheral and foveal vision. *Journal of Vision*, 20, 2020.

Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks, 2019.

Hidenori Tanaka, Daniel Kunin, Daniel L Yamins, and Surya Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 6377–6389. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/46a4378f835dc8040c8057beb6a2da52-Paper.pdf.

Adrian Tang, Simha Sethumadhavan, and Salvatore Stolfo. CLKSCREW: Exposing the perils of Security-Oblivious energy management. In *26th USENIX Security Symposium (USENIX Security 17)*, pp. 1057–1074, Vancouver, BC, August 2017. USENIX Association. ISBN 978-1-931971-40-9. URL https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/tang.

Andrei Tatar, Radhesh Krishnan Konoth, Elias Athanasopoulos, Cristiano Giuffrida, Herbert Bos, and Kaveh Razavi. Throwhammer: Rowhammer attacks over the network and defenses. In *2018 USENIX Annual Technical Conference (USENIX ATC 18)*, pp. 213–226, Boston, MA, July 2018. USENIX Association. ISBN ISBN 978-1-939133-01-4. URL https://www.usenix.org/conference/atc18/presentation/tatar.

Victor Van der Veen, Cristiano Giuffrida, and Others. TRRespass: Exploiting the rowhammer bug in TRR-protected DRAM. In *IEEE Symposium on Security and Privacy*, 2020.

Patrik Velcický, Jakub Breier, Mladen Kovacevic, and Xiaolu Hou. Deepncode: Encoding-based protection against bit-flip attacks on neural networks. *CoRR*, abs/2405.13891, 2024. doi: 10.48550/ARXIV.2405.13891. URL https://doi.org/10.48550/arXiv.2405.13891.

Aladin Virmaux and Kevin Scaman. Lipschitz regularity of deep neural networks: Analysis and efficient estimation. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Chaoqi Wang, Guodong Zhang, and Roger Grosse. Picking winning tickets before training by preserving gradient flow. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=SkgsACVKPH.

Xingxing Wei, Ying Guo, and Jie Yu. Adversarial sticker: A stealthy attack method in the physical world. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2022. ISSN 1939-3539. doi: 10.1109/tpami.2022.3176760. URL http://dx.doi.org/10.1109/TPAMI.2022.3176760.

Ross Wightman. Pytorch image models. https://github.com/rwightman/pytorch-image-models, 2019.

Fan Yao, Adnan Siraj Rakin, and Deliang Fan. DeepHammer: Depleting the intelligence of deep neural networks through targeted chain of bit flips. In *29th USENIX Security Symposium (USENIX Security 20)*, pp. 1463–1480. USENIX Association, August 2020. ISBN 978-1-939133-17-5. URL https://www.usenix.org/conference/usenixsecurity20/presentation/yao.

Matthew D. Zeiler and Rob Fergus. *Visualizing and Understanding Convolutional Networks*, pp. 818–833. Springer International Publishing, 2014. ISBN 9783319105901. doi: 10.1007/978-3-319-10590-1_53. URL http://dx.doi.org/10.1007/978-3-319-10590-1_53.

A Full cross-model comparison

Table 3 reports *all* flip-budget measurements we collected, while Table 4 quantifies the corresponding computational savings. Both DNL and 1P-DNL remain the only attacks that are simultaneously **OF** (optimization-free) and **DA** (data agnostic), while achieving equal-or-higher accuracy reductions than optimization-based or data-requiring methods, often with fewer flipped bits. Note that Park et al. (2021) avoids requiring data by generating synthetic data from the victim model instead.

Table 3: Bit-flip attacks on ImageNet-1K. Each cell lists *# flips $\rightarrow$ AR (%)*. OF = optimization-free, DA = data-agnostic. * best of 5 trials.

Model	Method	OF	DA	Flips $\rightarrow$ AR (%)
AlexNet	BFA Rakin et al. (2019)	✗	✗	$17 \rightarrow 99.6$
	DNL	✓	✓	$10 \rightarrow 88.5,\ 17 \rightarrow 98.2$
	1P-DNL	✓	✓	$10 \rightarrow 93.2,\ 17 \rightarrow 98.2$
VGG-11	BFA Rakin et al. (2019)	✗	✗	$17 \rightarrow 99.7$
	ZeBRA Park et al. (2021)	✗	✓	$8 \rightarrow 99.8$
	DNL	✓	✓	$3 \rightarrow \mathbf{99.9}$
	1P-DNL	✓	✓	$2 \rightarrow 99.8$
ResNet-50	DeepHammer Yao et al. (2020)	✗	✗	$23^* \rightarrow 75.4$
	BFA Rakin et al. (2019)	✗	✗	$1 \rightarrow 5.5,\ 5 \rightarrow 99.7$
	ZeBRA Park et al. (2021)	✗	✓	$1 \rightarrow 7.7,\ 5 \rightarrow 99.7$
	DNL	✓	✓	$1 \rightarrow 6.6,\ 5 \rightarrow 40.4,\ 8 \rightarrow 99.7$
	1P-DNL	✓	✓	$1 \rightarrow \mathbf{99.4}$
MobileNet-V2	DeepHammer Yao et al. (2020)	✗	✗	$2^* \rightarrow \mathbf{99.8}$
	BFA Rakin et al. (2019)	✗	✗	$3 \rightarrow 99.8$
	ZeBRA Park et al. (2021)	✗	✓	$2 \rightarrow 99.7$
	DNL	✓	✓	$2 \rightarrow 99.8$
	1P-DNL	✓	✓	$2 \rightarrow \mathbf{99.9}$
Inception-V3	BFA Rakin et al. (2019)	✗	✗	$3 \rightarrow 99.8$
	ZeBRA Park et al. (2021)	✗	✓	$3 \rightarrow 99.8$
	DNL	✓	✓	$2 \rightarrow \mathbf{99.8}$
	1P-DNL	✓	✓	$3 \rightarrow 99.1$
ViT-B/16@224	BFA Rakin et al. (2019)	✗	✗	$5 \rightarrow 30.1,\ 10 \rightarrow 90.9$
	ZeBRA Park et al. (2021)	✗	✓	$5 \rightarrow 5.1,\ 10 \rightarrow 45.8$
	DNL	✓	✓	$5 \rightarrow \mathbf{99.3}$
	1P-DNL	✓	✓	$4 \rightarrow 99.1$

Table 4: Approximate computational costs for different bit-flip attacks. Here, θ is the number of parameters in the model, k is the number of flipped bits, m is the (mini-)batch size used for gradient or scoring, and B is the number of candidate bits evaluated in each iteration. All complexities assume that a forward/backward pass scales on the order of $\mathcal{O}(\theta \times m)$.

Method	Description	Complexity
BFA Rakin et al. (2019)	Iterative gradient-based search; each flip requires scoring multiple bits	$\mathcal{O}(k \times B \times \theta \times m)$
DeepHammer Yao et al. (2020)	Chain-based iterative search for each flip	$\mathcal{O}(k \times B \times \theta \times m)$
ZeBRA Park et al. (2021)	Zero real-data, but still repeated forward/backward passes per flip	$\mathcal{O}(k \times B \times \theta \times m)$
DNL (Ours)	**Pass-free**; select bits by magnitude only	$\mathcal{O}(\theta) + \mathcal{O}(k)$
1P-DNL (Ours)	**Single-pass**; one forward/backward pass	$\mathcal{O}(\theta) + \mathcal{O}(k)$

B One Flip Per Kernel Constraint Examples

MobileNetV3-Large (Howard et al., 2019): Applying our magnitude-based method for $k = 2$ selects the second highest-magnitude weight for sign flipping, which in this case belongs to the third convolutional layer, and results

in a significant accuracy drop to $AR(2) = 81.31$. Adding another magnitude-based weight results in a flip within the same kernel that reduces the degradation to $AR(3) = 46.97$, partially offsetting the attack. However, flipping the next highest-magnitude parameter from a different kernel instead raises the accuracy reduction dramatically to $AR(3) = 94.0$.

RegNet-Y 16GF (Radosavovic et al., 2020b): In the second convolutional layer, several of the top 10 highest-magnitude weights reside in the same kernel. Flipping the sixth highest weight yields $AR(6) = 74.2$, while also flipping the seventh highest weight (in the same kernel) improves accuracy to $AR(7) = 66.5$, rather than compounding the damage.

C Sign-Bit vs. Exponent-Bit Flips

While most prior work Section 5 focused on flipping exponent bits, we observe that sign-bit flips typically cause greater disruption per flip. Table 5 compares the impact of flipping the sign bit versus the most significant exponent bit on two representative architectures. Exponent-bit flips are indeed effective, as they can sharply reduce weight magnitudes. However, sign-bit flips consistently induce larger accuracy drops, especially at low flip counts. Notably, in many models, exponent-bit flips are far less impactful than sign-bit flips when ranked and applied using our method.

Table 5: Sign-bit vs. exponent-bit flips ($k = 10$). AR: accuracy reduction; mAR_{10}: mean AR over 1–10 flips on full precision models.

Model	Sign AR(10)/ mAR(10)	Exp. AR(10) / mAR(10)
VGG-11	**91.8/91.6**	53.89/ 34.48
ResNet-50	**99.7**/52.7	70.94/**65.48**
MobileNet-V2	99.85/92.17	99.86/92.0
MobileNet-V3	**99.42/88.37**	91.87/70.46
RegNetY-16GF	**82.87/61.9**	78.75 / 48.6
EfficientNet-B0	**95.80 / 77.7**	37.38 / 22.53
ViT-B/16@224	**99.84/99.52**	82.38/47.27
AlexNet	88.5 / **43.45**	**99.8**/36.03

D Weight Score Ablation

We evaluate several parameter scoring functions from the pruning literature and compare their effectiveness in identifying high-impact weights for sign-flip attacks. As shown in Figure 7, we measure the mean accuracy reduction mAR_{10} across 48 ImageNet models under the following scoring functions:

- **Magnitude-based:** $S(\theta_i) = |\theta_i|$.
- **GraSP:** $S(\theta_i) = |\theta_i \odot Hg|$, following the gradient-flow preservation principle of Wang et al. (2020) where Hg is the hessian vector product.
- **GraSP (Gauss-Newton Approx.):** Similar to GraSP but approximates the Hessian H with the square of first-order gradients.
- **SynFlow:** $S(\theta_i) = |g \odot \theta_i|$, akin to gradient *times* weight.
- **Optimal Brain Damage (OBD):** $S(\theta_i) \approx \frac{1}{2}\theta_i^t H_{ii}\theta_i$ (LeCun et al., 1989a).
- **Hybrid (Ours):** As we define in Equation (3) with and without second order term.

where $g = \frac{\partial \mathcal{R}}{\partial \theta_i}, H = \frac{\partial^2 \mathcal{R}}{\partial \theta_i^2}$.

We observe that certain models are vulnerable to second-order-based scores (e.g., OBD) even when they prove more resilient to pure magnitude-based attacks. Nevertheless, other architectures appear more robust against OBD or GraSP while showing larger drops under magnitude-based score. Motivated by these mixed results, our hybrid score combines both magnitude and gradient terms. This blend consistently identifies critical weights even in cases where either component alone fails to degrade accuracy. Overall, the hybrid approach delivers the most reliable performance drop across the tested models.

Figure 7: Comparison of mAR_{10} across different weight score functions for the model parameters applied to 48 ImageNet models.

Table 6: $AR(\cdot)$ Targeting Binary ResNet-18 with DNL

AR(1)	AR(2)	AR(3)	AR(5)	AR(10)
0.14	12.90	60.71	90.35	96.50

E   Additional Datasets Evaluation

Figures 8, 9, and 10 analyze individual dataset results on these three popular classifiers. Each shows a steep drop in accuracy with very few sign flips, highlighting the generality of the attack. Notably, although these models differ in architecture and capacity, they all exhibit severe degradation once our detected sign bits are flipped. This finding reinforces that our method targets fundamental weaknesses in DNN representations rather than exploiting quirks of a specific network or dataset.

F   1P-DNL Algorithm

Similar to Algorithm 1, Algorithm 2 shows the algorithm of 1P-DNL.

G   Defenses

More exsiting defenses

Binarization.   Binary-weight networks are often assumed to be naturally resilient to weight perturbations He et al. (2020); Rakin et al. (2021), yet flipping a sign bit still inverts the weight. We show the results on a binarized ResNet-18 in Table 6, confirming that binarization alone offers negligible protection.[1]

---

[1]Results reproduced with the RA-BNN recipe of Rakin et al. (2021).

Figure 8: *AR* (%) on DTD dataset (Cimpoi et al., 2014) with varying number of sign flips over popular image encoders.

Figure 9: *AR* (%) on FGVC Aircraft dataset (Maji et al., 2013) with varying number of sign flips over popular image encoders.

Figure 10: *AR* (%) on Food101 dataset (Bossard et al., 2014) with varying number of sign flips over popular image encoders.

Figure 11: *AR* (%) on Stanford Cars dataset (Krause et al., 2015) with varying number of sign flips over popular image encoders.

Algorithm 2 1P-DNL – Single-Pass Attack

1: **Inputs:** Model f_θ, number of bits to flip k, number of layers L
2: $X \leftarrow$ random input (e.g., Gaussian noise)
3: $\mathcal{R}(\theta) \leftarrow \sum_i f_\theta(X)[i]$ // e.g., sum of logits
4: $g \leftarrow \nabla_\theta \mathcal{R}(\theta)$ // one backward pass
5: $\theta_L \leftarrow$ parameters in the first L layers of θ
6: **for** each θ_i in θ_L **do**
7: Approx. Hessian diagonal by Gauss–Newton: $H_{ii} \approx [g_i]^2$
8: $\mathcal{S}(\theta_i) \leftarrow |\theta_i| + \left| \theta_i\, g_i + \frac{1}{2}\, \theta_i^2\, H_{ii} \right|$
9: **end for**
10: Sort θ_L in descending order by $\mathcal{S}(\theta_i)$
11: $\mathcal{K} \leftarrow$ top-k entries of θ_L (one flip per kernel)
12: **for** each θ_i in $\mathcal{K}$ **do**
13: $\theta_i \leftarrow -\theta_i$ // flip sign bit
14: **end for**
15: **Output:** Modified parameters θ

H   Selective Defense Against Sign-Flips: Additional Setups

Following Section 6, to quantify this selective defense, we tested it on 16 particularly vulnerable networks, each suffering at least a 50% accuracy reduction ($AR(100,000) \geq 50\%$) when 100K random parameters were flipped. We

Figure 12: *AR* reported across five model families of varying capacities under 1P-DNL attack. The similar vulnerability levels suggest that model size alone does not mitigate sign-flip attacks.

Figure 13: *AR* reported across five model families of varying capacities under DNL attack.

use this large-scale random flip as a strong, non-specific stress test that is not tied to our own scoring method, ensuring that the defense remains robust to other score-based sign bit flips. We then varied the fraction of protected sign bits from 1% to 20%, focusing on the largest weights in absolute value. As expected, even a modest level of

Figure 14: $AR(100,000)$ under 100k random sign flips, with selective protection on varying fractions of the most vulnerable parameters (ranked by DNL). Even partial coverage of high-scoring parameters substantially improves robustness.

protection dramatically reduced the damage inflicted by sign-flip attacks. Moreover, as illustrated in Figure 14, selectively safeguarding this small subset of sign bits mitigates the impact from sign bit flip attacks, as reflected from the stress test proposed above, demonstrating that partial protection of critical parameters offers a practical and effective defense In Appendix G, we show that naive defense mechanisms would have been unsuccessful in defending against sign flips.

In addition to selectively protecting the most impactful sign bits, we also tested a baseline defense that shields a randomly chosen subset of bits at different coverage levels. Figure 15 shows that even when 20% of the sign bits are randomly protected, the network remains highly vulnerable under 100k random sign flips. This stands in stark contrast to protecting only a small fraction of critical sign bits (e.g., the largest-magnitude weights), which can substantially preserve accuracy. The results underscore that which bits get protected is more important than how many.

Figure 15: $AR(100,000)$ under 100k random sign flips, with random subsets (1%, 5%, 10%, and 20% coverage) of sign bits protected. Unlike Figure 14, where shielding the most vulnerable bits significantly reduces damage, uniform random selection offers little resilience, as even 20% coverage barely mitigates the attack.