

Sponge Examples: Energy-Latency Attacks on Neural Networks

Ilia Shumailov
University of Cambridge, UK
ilia.shumailov@cl.cam.ac.uk

Yiren Zhao
University of Cambridge, UK
yiren.zhao@cl.cam.ac.uk

Daniel Bates
University of Cambridge, UK
daniel.bates@cl.cam.ac.uk

Nicolas Papernot
University of Toronto
and Vector Institute, Canada
nicolas.papernot@utoronto.ca

Robert Mullins
University of Cambridge, UK
robert.mullins@cl.cam.ac.uk

Ross Anderson
University of Cambridge, UK
ross.anderson@cl.cam.ac.uk

Abstract—The high energy costs of neural network training and inference led to the use of acceleration hardware such as GPUs and TPUs. While such devices enable us to train large-scale neural networks in datacenters and deploy them on edge devices, their designers’ focus so far is on average-case performance. In this work, we introduce a novel threat vector against neural networks whose energy consumption or decision latency are critical. We show how adversaries can exploit carefully-crafted sponge examples, which are inputs designed to maximise energy consumption and latency, to drive machine learning (ML) systems towards their worst-case performance. Sponge examples are, to our knowledge, the first denial-of-service attack against the ML components of such systems.

We mount two variants of our sponge attack on a wide range of state-of-the-art neural network models, and find that language models are surprisingly vulnerable. Sponge examples frequently increase both latency and energy consumption of these models by a factor of $30\times$. Extensive experiments show that our new attack is effective across different hardware platforms (CPU, GPU and an ASIC simulator) on a wide range of different language tasks. On vision tasks, we show that sponge examples can be produced and a latency degradation observed, but the effect is less pronounced. To demonstrate the effectiveness of sponge examples in the real world, we mount an attack against Microsoft Azure’s translator and show an increase of response time from 1ms to 6s ($6000\times$). We conclude by proposing a defense strategy: shifting the analysis of energy consumption in hardware from an average-case to a worst-case perspective.

Index Terms—availability attacks, adversarial machine learning, adversarial examples, sponge examples, latency attacks, denial of service

1. Introduction

The wide adoption of machine learning has led to serious study of its security vulnerabilities. Threat vectors such as adversarial examples [1], [2], data poisoning [3], [4], membership inference [5]–[7] and fault injection attacks [8] have been extensively explored. These attacks either target the *confidentiality* or *integrity* of machine learning systems [9], [10]. So what about the third leg

of the security triad: their *availability*? In this paper, we introduce an attack that increases the power drawn by neural networks and the time they take to make decisions. An adversary may mount our attack on a datacenter providing ML-as-a-Service to cause disruption, *i.e.* denial-of-service [11]. Increasing the energy consumption of edge devices such as smartphones can drain their batteries and make them unavailable [12]. Perhaps even more seriously, an attack that slows down decisions can subvert safety-critical or mission-critical systems.

Our key observation is that different inputs of the same size can cause a deep neural network (DNN) to use very different amounts of time and energy: this *energy-latency gap* is the vulnerability we exploit. The gap exists because of specific optimisations in hardware (*e.g.* leveraging input sparsity) and algorithms (*e.g.* a variable number of passes through the network for an input of the same size).

Our attack can be even more effective against the growing number of systems that use GPUs or custom hardware. Machine learning in general, and neural networks in particular, command workloads heavy in matrix algebra. GPUs were fundamental to the AlexNet breakthrough in 2012 [13]; in response to increasing demand, Google introduced TPUs to facilitate inference – and training – in its datacenters [14], while Apple introduced the Neural Engine to make its smartphones more energy-efficient for on-device deep learning [15]. Hardware engineers explicitly target the Operations per Watt (OPs/W) performance of DNN processing. But by increasing complexity, optimisations tend to increase the attack surface. There is ample precedent elsewhere in computer engineering for optimisations widening the gap between average-case and worst-case performance in ways that a capable attacker can exploit: a recent example is Spectre [16], which exploits hardware speculation to launch a powerful timing side-channel attack. Security engineers therefore need to pay close attention to worst-case performance. In this paper, we start this process for the optimisations used to speed up modern machine learning in both hardware and algorithms.

Sponge examples are designed to soak up energy consumed by a given neural network, forcing the underlying hardware system running DNN inference towards its worst-case performance. We present two ways of generating sponge examples, one gradient-based and one using genetic algorithms. The gradient-based approach

requires access to DNN model parameters, while the genetic algorithm only sends queries to the model and evolves inputs based on energy or latency measurements. These two attack methods cover both White-box and Black-box attacks, and our extensive experiments demonstrate the effectiveness of sponge examples under both scenarios. When considering modern ML-as-a-Service in a Black-box setting, our genetic algorithm successfully produces sponge examples that consistently increase the service response time and thus the energy consumption of the remote server.

In this paper we make the following contributions:

- We introduce a novel threat against the availability of ML systems based on energy and latency. Our *sponge examples* are designed to cause inference to take as long as possible and consume as much energy as possible.
- We show that sponge examples cause increased energy consumption and longer run-time for a wide range of vision and language models. Sponge examples are particularly powerful against language models.
- We demonstrate the portability of sponge examples, by showing they are not only transferable across hardware platforms (CPUs, GPUs, and an ASIC simulator) but also across model architectures.
- We show that modern ML-as-a-Service is vulnerable to sponge attacks. With a 50-character input to Microsoft Azure translator, we can increase the latency from 1ms up to 6s: a 6000× degradation.
- We present a simple defense against sponge examples in the form of a worst-case performance bound for some models. This can also prevent unexpected increases in energy consumption in the absence of adversaries, potentially reducing the carbon footprint of models deployed for inference at scale.

2. Motivation

Artificial Intelligence and Machine Learning have enabled real progress in automation, unleashing greater productivity and potential economic growth. Yet modern machine learning has become extremely power-hungry. It is estimated that the energy consumed in training a single transformer model is equivalent to 60% of a car’s lifetime carbon emissions [17]. And the energy cost doesn’t stop there – each inference consumes significant energy and happens ever more often¹. Energy is the lifeblood of modern machine learning, as energy-efficient inference and training are needed to scale machine learning to more use cases. In this paper we explore the adversarial manipulation of energy and latency in ML models, services and applications.

Modern hardware exploits many different optimisation techniques to maintain a high ratio of useful work to energy consumed. This often involves predicting future workloads and scheduling resources according to dynamic needs. Prediction and speculation occur at a number of levels in the stack, widening the gap between average-case and worst-case scenarios. This is particularly an

1. For example, OpenAI greatly limits number of queries one can do to the GPT3 model.

issue in time or energy sensitive tasks, such as time series forecasting for automatic trading [18] and activity recognition on wearable devices [19]. In such applications, hitting worst-case performance could cause failures in decision making or deplete the batteries of user devices. In safety-critical and real-time systems, such as autonomous vehicles which depend on scene understanding with tight latency constraints, service-denial attacks can pose a threat to life.

In this paper, we show that a capable attacker can exploit the performance dependency on hardware and model optimisations to launch a number of different attacks. We find that they can negate the effects of hardware optimisations, increase computation latency, increase hardware temperature and massively increase the amount of energy consumed. To make things worse, they can do this with very few assumptions, making attacks scalable in the real world. We further highlight the realism of sponge examples in a case study with Microsoft Azure translator, where we degraded latency up to a factor of 6000×.

On a number of occasions, despite the hardware protection provided by GPU engineers, we were able to increase temperature so that it passed the throttling point and sometimes even crashed the GPU drivers. The energy consumed by sponge examples on a machine learning model can therefore affect the underlying hardware if its power management software or hardware is not designed with adversaries in mind.

3. Background

3.1. Hardware Acceleration for Deep Learning

Deep Neural Network (DNN) inference is both compute-intensive and memory-intensive. Common hardware products such as CPUs and GPUs are now being adapted to this workload, and provide features for accelerating it. Intel’s Knights Mill CPU provides a set of SIMD instructions [20], while NVIDIA’s Volta GPU introduces Tensor Cores to facilitate the low-precision multiplications that underpin much of deep learning [21].

Hardware dedicated to deep learning is now pervasive in data centers, with examples including Big Basin at Facebook [22], BrainWave at Microsoft [23], and racks of TPUs at Google [14], [24]; the underlying hardware on these systems are either commodity hardware (Big Basin), re-configurable hardware (FPGAs for BrainWave), or custom silicon (TPUs). The latter two are specifically designed to improve the number of Operations per Watt of DNN inference. Careful modeling of average hardware efficiency allows ML-as-a-Service providers to price their services per query, rather than per energy used. As we discuss later, custom and semi-custom hardware will typically exploit sparsity in data and the adequacy of low-precision computations for DNN inference, reducing both arithmetic complexity and the amount of DRAM traffic, to achieve significantly better power efficiency [25]–[27]. Our attack targets these optimisations among others.

3.2. Attacks on Energy

Operations per Watt are an important indicator of the efficiency of cloud infrastructure [28]. Power oversubscrip-

tion is a popular method for cloud services to handle provisioning, but it leaves datacenters vulnerable to power attacks [29]–[32]. If malicious users can remotely generate power spikes on multiple hosts in the data center at the same time, they might overload the system and cause disruption of service [11], [29]. Energy attacks against mobile devices usually aim to drain the battery more quickly [12], [33], although energy management in mobile devices can also be used to perform deterministic fault injection [34]. The possible victims of energy attacks on mobile systems range from autonomous vehicles to sensors with constrained computing abilities [35]. Higher energy consumption also increases hardware temperature, which in turn increases the failure rate. For example, Anderson *et al.* note that an increase of 15°C causes component failure rates to go up by 2× [36]. Modern hardware throttles to avoid overheating; while short-term power savings may be possible through such voltage scaling, the overall energy consumption increases [37]. This creates nonlinear dependencies between energy and latency.

3.3. Security of Machine Learning

Machine learning has been shown to be vulnerable to a number of different attack vectors [38]. Adversarial examples can cause a system to classify inputs incorrectly [1], [2]. Adversarial examples can be found in the White-box setting through gradient-based optimization [1], [2] while in the Black-box setting, the adversary can transfer adversarial examples from another model [39] or approximate gradients with finite differences [40] when they can observe the model’s confidence as well as the output label. Data can be poisoned to manipulate future performance [3], [4]. Run-time bit-errors can be introduced to greatly reduce performance [8]. These attacks either target the *confidentiality* or *integrity* of machine learning systems [9], [10].

Here, we explore availability, *i.e.* timely and reliable access to information [41], and introduce a new form of service denial with samples that act as a sponge for time or energy. Service-denial attacks are well known in the context of computer networking [42], [43], but have been overlooked so far in ML. The current NIST draft on adversarial machine learning touches upon availability, but does not provide any examples of attacks [38].

Poisoning can perhaps be seen as an availability attack. If an attacker can poison data so that the machine learning model stops training or does so with reduced accuracy, this may be seen in some contexts as reducing availability. For example, Erba *et al.* presented such an attack against Industrial Control Systems [44]. However, the attacks presented in this paper do not poison data, but target either the hardware or the algorithmic complexity of the model.

4. Methodology

4.1. Threat Model

In this paper we assume an adversary with the ability to supply an input sample to a target system, which then processes the sample using a single CPU, GPU or

ASIC. We assume no rate limiting, apart from on-device dynamic power control or thermal throttling.² We assume no physical access to the systems *i.e.* an attacker cannot reprogram the hardware or change the configuration.

We consider three threat models. The first is a **White-box** setup: we assume the attackers know the model architecture and parameters. The second considers an **interactive Black-box** threat: we assume attackers have no knowledge of the architecture and parameters, but are able to query the target as many times as they want and to time operations or measure energy consumption remotely. The third is the **blind adversary**: we assume no knowledge of the target architecture and parameters, and assume no ability to take direct measurements. In this setting, the adversary has to transfer previously-discovered sponge examples directly to a new target – without prior interaction.

Our adversary models loosely capture ML-as-a-Service deployments and on-device data processing. A simple example could be a dialogue or a translation system. Users interact continuously by sending queries and can measure energy consumption, or when that is not possible by the response time (see Section 5). Indeed, in Section 5.5 we show on an example of a Microsoft Azure translator that modern ML-as-a-Service is vulnerable to sponge attacks which only rely on the adversary’s ability to observe response latency—even in presence of networking delay.

4.2. The Energy Gap

The *Energy Gap* is the performance gap between average-case and worst-case performance, and is the target for our sponge attacks. To better understand the cause of this gap, we tested three hardware platforms: a CPU, a GPU and an ASIC simulator. The amount of energy consumed by one inference pass (*i.e.* a forward pass in a neural network) depends primarily on [45]:

- the overall number of arithmetic operations required to process the inputs; and
- the number of memory accesses *e.g.* to the GPU DRAM.

The intriguing question now is:

is there a significant gap in energy consumption for different model inputs of the same dimension?

As well as fixing the dimension of inputs, *i.e.* not increasing the number of characters in a text sample or the pixel dimension of an image, we also do not consider inputs that would exceed the pre-defined numerical range of each input dimension. If models do have a large energy gap between different inputs, we describe two hypotheses that we think attackers can exploit to create *sponge examples*, that is, inputs trigger the worst-case performance and have abnormally high energy consumption.

4.2.1. Hypothesis 1: Computation Dimensions. Aside from data sparsity, modern neural networks also have a computational dimension. Along with variable input

2. Thermal throttling refers here to the deliberate slow-down of device performance when cooling is no longer able to dissipate the heat generated by a workload.

and output shapes, the internal representation size often changes as well – for example, in the Transformer-based architectures for machine translation [46]. The model is autoregressive in this case; both the input and output are sequences of words and internal computation depends on both of them. Before text gets to the model it has to go through a number of stages. First, individual components are separated within the sentence, removing punctuation and keeping useful words. Next, each word is represented as a number of tokens whose shape depends on the richness of input and output dictionaries. Because we cannot represent words mathematically, we need to map them to some numerical form. Yet we cannot build a mapping with all possible words, because that greatly increases model complexity, so in practice dictionaries with most-popular sub-words are used. Once tokenized, individual tokens are then projected into the embedding space (e.g. word2vec [47]), a high-dimensional space where knowledge about individual tokens is encoded. As computation progresses, each inference step depends on the embeddings of all of input tokens and output tokens produced so far. For example, imagine encoding the word ‘Athazagoraphobia’. With commonly used English dictionaries, it will get assigned 4 tokens for its input size of 16: ‘ath’, ‘az’, ‘agor’, ‘aphobia’. If a user makes a typing mistake, say ‘Athazagoraphbia’, then suddenly its representation turns into 7 tokens for the same size of 16: ‘ath’, ‘az’, ‘agor’, ‘aph’, ‘p’, ‘bi’, ‘a’. An adversary can exploit it and construct large token representations. For example, ‘A/h/z/g/r/p/p/i/’ will be 16 separate tokens. Ultimately, unknown words both in the input and output spaces will lead to a much larger sentence representation and many more inference runs.

Consider an input sequence x and an output sequence y . We denote the input and output token sizes (i.e. the number of individual tokens extracted from an input sentence and produced for the output sentence) with l_{tin} and l_{tout} . Each of the words in a sequence is embedded in a space of dimensionality l_{ein} , for the input, and l_{eout} , for the output. Algorithm 1 contains the pseudocode for a Transformer’s principal steps. In red, we annotate the computational complexity of the following instruction. As can be seen, several quantities can be manipulated by an adversary to increase the algorithm’s run time: 1) token size of the input sentence l_{tin} ; 2) token size of the output sentence l_{tout} ; and 3) size of the input and output embedding spaces (l_{ein} and l_{eout}). All of the above can cause a non-linear increase in algorithmic complexity and thus heavily increase the amount of energy consumed. Note that perturbing these quantities does not require that the adversary modify the dimension of input sequence x ; that is, with no changes to the input length, the adversary can increase energy consumption non-linearly.

4.2.2. Hypothesis 2: Data Sparsity. The rectified linear unit (ReLU), which computes $x \mapsto \max(0, x)$, is the de facto choice of activation function in neural network architectures. This design introduces sparsity in the activations of hidden layers when the weighted sum of inputs to a neuron is negative. A large number of ASIC neural network accelerators consequently exploit runtime data sparsity to increase efficiency [48]–[50]. For instance, ASIC accelerators may employ zero-skipping multipli-

Algorithm 1: Translation Transformer NLP pipeline

Input: Text sentence x
Result: y
 $\downarrow O(l_{\text{tin}})$
1 $x_{\text{tin}} = \text{Tokenize}(x)$;
2 $y_{\text{touts}} = \emptyset$;
 $\downarrow O(l_{\text{ein}})$
3 $x_{\text{ein}} = \text{Encode}(x_{\text{tin}})$;
 $\downarrow O(l_{\text{tin}} \times l_{\text{ein}} \times l_{\text{tout}} \times l_{\text{eout}})$
4 **while** y_{tout} has no end of sentence token **do**
 $\downarrow O(l_{\text{eout}})$
5 $y_{\text{eout}} = \text{Encode}(y_{\text{tout}})$;
 $\downarrow O(l_{\text{ein}} \times l_{\text{eout}})$
6 $y_{\text{eout}} = \text{model.Inference}(x_{\text{ein}}, y_{\text{eout}}, y_{\text{touts}})$;
 $\downarrow O(l_{\text{eout}})$;
7 $y_{\text{tout}} = \text{Decode}(y_{\text{eout}})$;
8 $y_{\text{touts}} \text{-add}(y_{\text{tout}})$;
9 **end**
 $\downarrow O(l_{\text{tout}})$;
10 $y = \text{Detokenize}(y_{\text{touts}})$

cations or encode DRAM traffic to reduce the off-chip bandwidth requirement. The latest Xilinx AI compiler provides optimisations [51] for automatically deploying sparse models to their FPGA devices, promoting the use of model sparsity in production systems. On the algorithmic level, there is a recent surge of interest in using dynamic coarse-grained sparsity for accelerating GPU inference [52], [53]. Hence, inputs that lead to less sparse activations will increase the number of operations and the number of memory accesses, and thus energy consumption.

4.3. The Laws of Physics

Before getting to the details of the attack, we need to understand what affects energy consumption and what the attacker can reliably influence. Energy E is the total consumed static power P_{static} and dynamic power P_{dynamic} for an interval of time t . This energy formulation can be analysed in more detail; we show how to do this in Section B of the Appendix.

$$\begin{aligned}
E &= (P_{\text{static}} + P_{\text{dynamic}}) \times t \\
&\quad \text{overheat or increase overall consumption} \\
&= \left(\left[\sum I_s \times \left(e^{\frac{qV_d}{kT}} - 1 \right) \times V_{\text{core}} \right] \right. \\
&\quad \left. + \left[\alpha \times C \times V_{\text{core}}^2 \times f \right] \right) \times t. \\
&\quad \text{throttle or exploit load predictor} \qquad \text{run for longer or exploit the predictor} \\
&\quad \text{more activity of the board} \qquad \text{run for longer or exploit the predictor}
\end{aligned} \tag{1}$$

The salient elements from Equation (1) are that an attacker can affect energy use through four parameters: T (temperature), α (activity ratio), f (frequency) and t (time). Our sponge examples directly exploit the activity ratio α and execution time t , since these two parameters are tightly linked to the number of operations and memory accesses performed by model inference. Although frequency f and temperature T will be influenced indirectly through optimisations performed by the underlying hardware, these are not our direct targets. We hypothesise these parameters

(f and T) can also be exploited to create hardware level availability attacks on ML systems, *e.g.* forced throttling or heating of devices, but they are beyond the scope of this paper.

4.4. Attack Methods and Setups

Having presented the intuition behind our attacks, we now introduce strategies for finding sponge examples corresponding to the threat models described in Section 4.1.

4.4.1. Genetic Algorithms in White-box and Black-box Settings. Genetic algorithms (GA) are a powerful tool for adversaries [54]. They can optimise a diverse set of objectives, and require no local gradient information. They are a particularly good fit for adversaries who only have access to the model’s prediction in a Black-box setting. The general pipeline of a GA is presented in Algorithm 2 in Appendix. We start with a pool of randomly generated samples S . These are images for computer vision models, or sentences for NLP tasks. We then iteratively evolve the population pool as is depicted in Figure 1.

- For computer vision tasks, we sample two parents A and B from the population pool, and crossover the inputs using a random mask $A * \text{mask} + (1 - \text{mask}) * B$.
- For NLP tasks, we sample two parents A and B , and crossover by concatenating the left part of parent A with the right part of parent B . We then probabilistically invert the two parts.

We explain the reasons for these choices in Section C of the Appendix. Next, we randomly mutate (*i.e.* perturb) a proportion of the input features (*i.e.* pixels in vision, words in NLP) of the children. To maintain enough diversity in the pool, where applicable we preserve the best per-class samples in the pool. We obtain a fitness score P for all pool members, namely their energy consumption. We then select the winning top 10% of samples S^3 , and use them as parents for the next iteration. This genetic algorithm is simple but effective in finding sponge examples. Parameter choice is explained in Section A in the Appendix. In Section C, we further explain the domain-specific optimisations of the GA algorithm on NLP and CV tasks for achieving a better attack performance.

Although following the same algorithm described above, we form two variants of GA for Black-box and White-box attacks respectively, each differing the way we measure fitness:

- White-box GA: We access the parameters of the neural networks and provide an estimated energy cost based on the run-time sparsity, *i.e.* number of operations based on the structure and parameters of the neural networks.
- Black-box GA: We do not access any of the neural network internals, and use purely the measured hardware cost as the fitness, *i.e.* latency or energy consumption.

3. As the sample pool is large, selecting the top 10% makes the process more tractable.

4.4.2. L-BFGS in the White-box Setting. We now consider an adversary with access to the model’s parameters. Rather than a genetic algorithm, we use L-BFGS [55] to optimise the following objective:

$$- \sum_{a_l \in A} \|a_l\|_2 \quad (2)$$

where A is the set of all activation values and a_l the activations of layer l . This generates inputs that increase activation values of the model across all of the layers simultaneously. Following Objective 1 outlined above, the increase in density prevents hardware from skipping some of the operations, which in turn increases energy consumption. We only evaluate the performance of sponge examples found by L-BFGS on computer vision tasks because of the discrete nature of the NLP tasks, which prevents differentiating the objective in Equation (2)⁴.

4.4.3. Cross-model and Cross-hardware Transferability for Blind Adversaries. When adversaries are unable to query the model, they cannot directly solve an optimisation problem to find sponge examples, even using the interactive Black-box approach, *i.e.* the GA. In this blind-adversary setting, we exploit transferability across both models and hardware. Indeed, in Section 5.4 and Section 6.3 in the Appendix, we show that sponge examples transfer across models. We examine three hardware platforms in our evaluation:

- **CPU:** The platform is an Intel(R) Xeon(R) CPU E5-2620 v4 with 2.10GHz clock frequency. We use the Running Average Power Limit (RAPL) to measure energy consumption of the CPU. RAPL has been thoroughly evaluated and found to reflect actual energy consumption, as long as the counters are not sampled too quickly [56], [57].
- **GPU:** We use a GeForce 1080 Ti GPU with a 250.0 Watts power limit, a 96°C slowdown temperature and a 84°C throttling temperature. We use the NVIDIA Management Library (NVML) to measure energy consumption. NVML was previously found to capture energy quite accurately, with occasional instability for high-low patterns and high sampling rates [58].
- **ASIC:** We also developed a deterministic ASIC simulator, which monitors and records the runtime operations and number of DRAM accesses assuming a conservative memory flushing strategy. We then use measurements by Horowitz to approximate energy consumption [45]: at 45nm technology and 0.9V, we assume 1950 pJ to access a 32 bit value in DRAM and 3.7 pJ for a floating-point multiplication.

We show in Section 5.4 that sponge examples transfer across these types of hardware.

4. It is worth noting that for NLP tasks given knowledge of the dictionary an attacker can design the worst possible input and output token sequences. In this paper we make no assumptions about the dictionary or the model deployed, instead we optimise directly over energy or time.

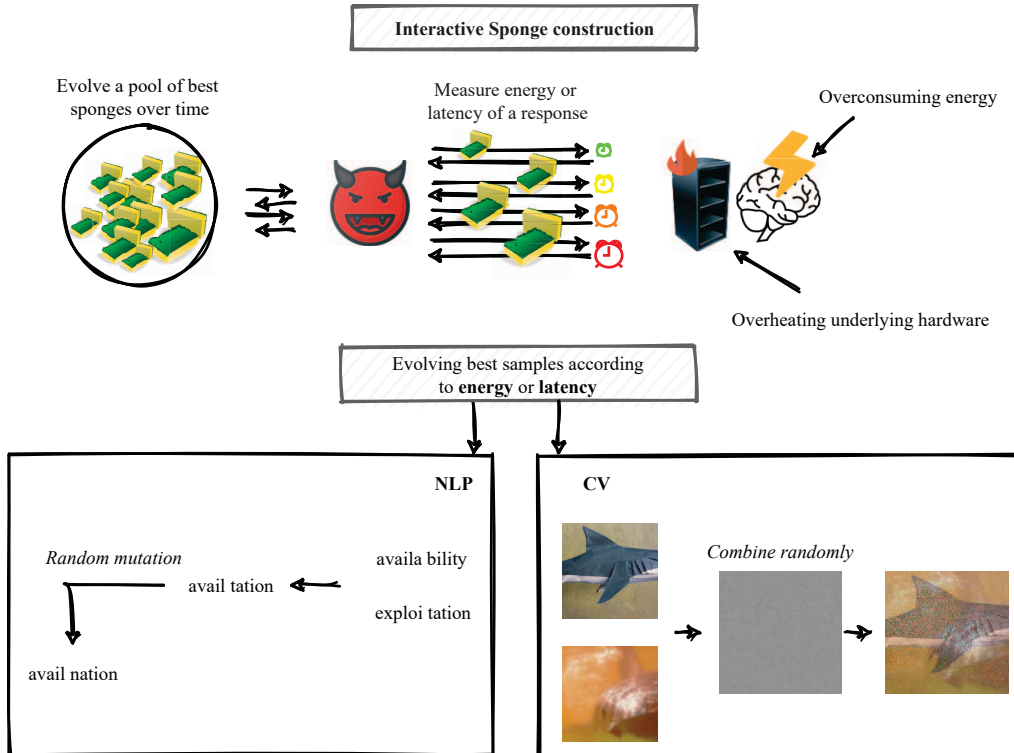


Figure 1: Availability adversary constructs sponge examples using a genetic algorithm. The adversary tries samples against the model and measures either latency or energy consumed, mixing the best performing samples in the pool. Eventually, the attacker identifies potent sponge examples.

5. Sponge Examples on Language Models

5.1. Models and Datasets

We first evaluate our sponge example attack on a range of NLP models provided by the FairSeq framework [59]. The models we consider have achieved top performance at their respective tasks and are used heavily in the real world to analyse data at scale. We report the performance of the RoBERTa [60] model, an optimised BERT [61], on three GLUE benchmarks designed to assess language understanding [62]. The datasets we considered include tasks in the SuperGLUE benchmark plus a number of machine-translation tasks. The SuperGLUE benchmark follows the style of GLUE but includes a wider range of language-understanding tasks including question answering and conference resolution [62], [63]. Further, we evaluate the attack on a number of translation tasks (WMT) using Transformer-based models [64]–[66]. Both translation and language comprehension are fundamental to human society and form a bridge between computers and humans. They are built into virtual assistants and many other applications in the real world that are used on a day-to-day basis.

Consider the pipeline for handling text. Before getting to the models, the text goes through several preprocessing steps. First, words get tokenized in a manner meaningful for the language. We used the tokenizer from the Moses toolkit [67], which separates punctuation from words and normalises characters. Next, tokenized blocks get encoded. Until recently, unknown words were simply replaced with an unknown token. Modern encoders improve performance by exploiting the idea that many words are a combination

of other words. BPE is a popular approach that breaks unknown words into subwords it knows and uses those as individual tokens [68]. In that way, known sentences get encoded very efficiently, mapping every word to a single token, and the number of computations is greatly reduced.

5.2. White-box Sponge Examples

In this section, we look at the White-box GA attack (as explained in Section 4.4.1) for generating sponge examples. In this setup, we access to the parameters and run-time information of the neural networks. The GA optimisation relies on the estimated number of operations of the neural network inference.

Table 1 shows the energy consumption of different models in the presence of our generated White-box sponge examples. For different input sequence sizes and a wide range of NLP tasks, we show the energy costs of sponge examples on both GPUs (GPU Energy) and the ASIC simulator (ASIC Energy). We use natural, random and sponge to represent the energy measured on data from the evaluation dataset, randomly formed strings and sponge examples. In addition, we also report the latency of running these samples on GPUs. Due to the limitation of the ASIC simulator, we cannot have faithful time measurements and these numbers are not reported.

We have made several important observations:

- The energy cost of sponge examples is always the highest on both GPUs and ASICs. In the best-case scenario for the attacker, sponge examples increase energy consumption by $26\times$.

Input size		GPU Energy [mJ]			ASIC Energy [mJ]			GPU Time [mS]		
		Natural	Random	Sponge	Natural	Random	Sponge	Natural	Random	Sponge
<i>SuperGLUE Benchmark with [60]</i>										
CoLA	15	2865.68	3023.705	3170.38	504.93	566.58	583.56	0.02	0.02	0.02
		1.00×	1.06×	1.11 ×	1.00×	1.12×	1.16 ×	1.00 ×	0.92×	0.92×
	30	3299.07	4204.121	4228.22	508.73	634.24	669.20	0.03	0.03	0.02
		1.00×	1.27×	1.28 ×	1.00×	1.25×	1.32 ×	1.00 ×	0.93×	0.82×
	50	3384.62	6310.504	6988.57	511.43	724.48	780.57	0.03	0.04	0.04
		1.00×	1.86×	2.06 ×	1.00×	1.42×	1.53 ×	1.00×	1.23×	1.27 ×
MNLI	15	3203.01	3573.93	3597.3	509.19	570.10	586.43	0.03	0.03	0.03
		1.00×	1.12×	1.12 ×	1.00×	1.12×	1.15 ×	1.00×	1.01 ×	0.95×
	30	3330.22	4752.84	5045.25	514.00	638.78	672.07	0.03	0.03	0.03
		1.00×	1.43×	1.51 ×	1.00×	1.24×	1.31 ×	1.00×	1.06 ×	1.03×
	50	3269.34	6373.507	7051.68	519.51	728.82	783.18	0.03	0.04	0.04
		1.00×	1.95×	2.16 ×	1.00×	1.40×	1.51 ×	1.00×	1.28×	1.30 ×
WSC	15	4287.24	13485.49	38106.98	510.84	1008.59	2454.89	0.04	0.07	0.20
		1.00×	3.15×	8.89 ×	1.00×	1.97×	4.81 ×	1.00×	2.02×	5.51 ×
	30	4945.47	36984.44	79786.57	573.78	2319.05	5012.75	0.04	0.20	0.46
		1.00×	7.48×	16.13 ×	1.00×	4.04×	8.74 ×	1.00×	4.89×	11.04 ×
	50	6002.68	81017.01	159925.23	716.96	5093.42	10192.41	0.05	0.46	0.93
		1.00×	13.50×	26.64 ×	1.00×	7.10×	14.22 ×	1.00×	10.16×	20.56 ×
<i>WMT14/16 with [64]</i>										
En→Fr	15	9492.30	25772.89	40975.78	1793.84	4961.56	8494.36	0.10	0.24	0.37
		1.00×	2.72×	4.32 ×	1.00×	2.77×	4.74 ×	1.00×	2.51×	3.89 ×
En→De	15	8573.59	13293.51	238677.16	1571.59	2476.18	48446.29	0.09	0.13	2.09
		1.00×	1.55×	27.84 ×	1.00×	1.58×	30.83 ×	1.00×	1.46×	24.18 ×
<i>WMT18 with [65]</i>										
En→De	15	28393.97	38493.96	874862.97	1624.05	2318.50	49617.68	0.27	0.33	7.25
		1.00×	1.36×	30.81 ×	1.00×	1.43×	30.55 ×	1.00×	1.20×	26.49 ×
<i>WMT19 with [69]</i>										
En→Ru	15	33181.43	91513.13	876941.24	1897.19	5380.20	47931.11	0.31	0.77	7.19
		1.00×	2.76×	26.43 ×	1.00×	2.84×	25.26 ×	1.00×	2.46×	22.85 ×

TABLE 1: Energy is reported in milli joules. We use the White-box GA attack to produce sponge examples and measure the performance on different platforms. The GPU readings are from NVML. GA was run for 1000 epochs with a pool size of 1000. A detailed explanation of the results is in Section 5.2. Standard deviation for ASIC measurements are shown in Table 5.

- Randomly generated samples are more energy-consuming than natural samples.
- When the task is quick to execute, sponge examples do not show big performance degradation in terms of GPU Time, but they increase latency significantly when the task takes more time (up to 30×).

The main reason for performance degradation appears to be the increased dimension of the computation, as described in Algorithm 1. First, for a given input sequence size, the attack maximises the size of the post-tokenisation representation (x_{tin}), exploiting the tokeniser and sub-word processing. Words with which the model is less familiar are represented inefficiently, forcing the network to do more work. Imagine holding an email conversation with an academic from another field, who uses specific technical terms. Every time an unfamiliar term appears you have to look it up in a dictionary or search engine. Second, the attack learns to maximise output sequence length, since this links directly to the computation cost. Third, internal computation coupled with output sequence length and post-tokenisation length give a quadratic increase in energy consumption. These reasons explain why sponge examples can significantly increase both energy and latency

of language models. Do note that in this paper we use relatively small input sizes and in practice the effect will be a lot more pronounced for larger texts. Indeed as we later show in Section 5.5, in a Black-box setup with 50-character long text inputs an attack on Azure Language Translator caused 6000× degradation.

Interestingly, we observe that randomly generated samples significantly reduce the performance of NLP tasks. This can be attributed to the fact that natural samples are efficiently encoded, whereas random and attack samples produce an unnecessarily long representation, meaning that random noise can be used as a scalable Black-box latency and energy attack tool. Otherwise put, many ML systems are vulnerable to simple barrage jamming.

It is also worth noting that the short execution of inference on GPUs makes it hard to provide an accurate measurement even with iterative runs. We further explain how this measurement is difficult due to a variety of hardware problems in Section 6.2.

In the upcoming sections, we turn to Black-box variants of the attack based on energy and latency measurements of the individual samples. We mentioned previously that modern hardware optimises Operations per Watt (OPs/W), making sure that energy is only actively consumed when

useful work is being done. In our experiments we see that the relationship between degradation factors of energy and time ranges between 1.15 and 1.62 (see Table 4 in Appendix), with energy scaling faster⁵.

5.3. Interactive Black-box Sponge Examples

In this section, we show the performance of the attacks running in an interactive Black-box manner against NLP tasks. In this setup, we launch the Black-box GA attack as described in Section 4.4.1. This interactive Black-box setup assumes that attackers cannot access to the neural network parameters but have the abilities to measure the energy or latency remotely. In addition, they can query the service as many times as they like, so there is no rate limiting. We evaluate two Black-box attacks, and they use GPU Time and GPU Energy as the optimisation targets for the GA. We also present results for a White-box GA attack in the third setup as a baseline, which is the same attack used in Section 5.2.

Figure 2 shows sponge example performance against a WMT14 English-to-French Transformer-based translator with an input of size 15 and pool size of 1000. In Figure 2, we use the name GPU Energy Attack, GPU Time Attacker and White-box Attacker to represent these different attacks. In addition, we report measurements of every iteration of the GA for these different attackers on GPU Energy, GPU Time and ASIC Energy respectively. In Figure 2, the legends represent attackers with different measurement proxies; and we show that these interactive Black-box attacks are transferable across hardware platforms and measurement proxies. For instance, an attack targeting GPU Time transfers well when used to increase the energy cost of the ASIC simulator.

It can be seen that although the attackers have no knowledge of any neural network internals or datasets, they are able to successfully increase the energy and time costs. The experiment further highlights the difference between using time and energy as fitness functions. While time is noisy and depends on the current state of the hardware, energy remains relatively stable during the attack. As was explained previously, that can be attributed to the hardware switching its performance modes to keep the ratio of useful work to energy constant.

5.4. Blind Black-box Sponge Examples and Hardware Transferability

In this section, we turn to the question of transferability across hardware and different models in a blind Black-box manner. As we’ve described in Section 4.1, in the blind Black-box setup the attacker blindly transfers the previously discovered sponge examples to the target. Table 2 shows the results across different models, tasks and hardware platforms. The first column is the source task that we used to produce sponge examples. We then later launch these sponge examples to the target tasks shown in the second column. We report the performance

5. Interestingly, we observe a net energy increase for the task if throttling happens. Although throttling decreases the running frequency and the voltage, it significantly increases the execution time so that the overall energy consumption has increased.

of both sponge and natural examples on the targeting task in Table 2. Since the ASIC simulator is coarse-grained, it does not produce faithful execution time estimation, we thus only report the estimated energy cost. In general, in Table 2, we observe a significant increase in energy and time in comparison to natural samples on all hardware platforms. However, the blind Black-box attacks fail to achieve the same level of energy or latency degradation when taking the White-box case as a baseline.

5.5. A Case Study: Microsoft Azure Translation

We evaluated sponge attack performance against an actually deployed service that is available on demand. We present a Black-box attack against this production system without any assumptions about its internals. As with the experiment setup in Section 5.3, we interactively evolve the pool of samples in a Black-box setting using latency as a fitness function. Note that, in this case, observed sample fitness is noisy as it includes communication latency. That in turn makes it harder to perform the attack.

We used the Microsoft Azure Translation system located on the same continent as the requesting server. We fixed the input to be 50 characters long, with a pool size of 500, and ran the attack for 50 epochs. We report four different attack runs, each running immediately when the previous attack finishes. The attacks are run sequentially, so Azure only translates a single sample at a time. It should be noted that we hold no assumptions about the actual system and do not possess any information about what architecture or dataset is used. Furthermore, we possess no information on whether Azure employs query-caching strategies or other optimisation techniques.

Figure 3 shows the performance of the attack as observed by the requesting server and reported by Azure. The server’s reported numbers are larger as they include additional noise from communication latency. It can be clearly seen that all four separate runs of GA were capable of converging to samples that were consuming considerably more time to process – up to a maximum degradation factor of 6000×. Although we have no way of telling the amount of energy that was consumed by the computation, results from Section 5 suggest that the energy consumption increase should also be in the range of thousands. Interestingly, we find that performance varies greatly within the pool during the attack. We suspect this is due to Azure’s caching mechanism, where previously performing samples get almost constant time computation and the pool has to adapt quickly. For all individual runs, we see an up–down pattern that we do not observe with experiments on our own hardware. Interestingly, Azure translator assigns high confidence scores > 0.9 to the sponge example predictions.

Sponge examples against translators strongly resemble Denial-of-Service (DoS) attacks observed in the wild. DoS attacks aim to make computer systems unresponsive and unavailable via excess connections or data requests. Instead of overwhelming the victim’s bandwidth as in the vast majority of DoS attacks [70], we target the application layer. In particular, we target the most expensive parts of the translator to get an extraordinary amplification factor by sending specifically crafted requests. For example with Azure and an input of length 50, we were getting translated

From	To		ASIC		GPU		CPU	
			Energy [mJ]	Time [S]	Energy [mJ]	Time [S]	Energy [mJ]	Time [S]
<i>Black-box</i>								
WMT16 _{en→de} [64]	WMT14 _{en→fr} [64]	Sponge	3648.219	0.174	17251.000	1.048	51512.966	
		Natural	1450.403	0.053	6146.550	0.537	23610.145	
			2.52×	3.27 ×	2.81 ×	1.95×	2.18×	
	WMT18 _{en→de} [65]	Sponge	2909.245	0.414	47723.500	3.199	181936.595	
		Natural	1507.364	0.253	27265.250	1.344	71714.201	
			1.93×	1.64×	1.75×	2.38 ×	2.54 ×	
WMT19 _{en→ru} [66]	Sponge	3875.365	0.652	67183.100	4.409	247585.091		
	Natural	1654.965	0.215	25033.620	2.193	121210.376		
		2.34×	3.03 ×	2.68 ×	2.01×	2.04×		
<i>White-box</i>								
WMT16 _{en→de} [64]	WMT16 _{en→de} [64]	Sponge	48447.093	2.414	260187.900	13.615	781758.680	
		Natural	1360.118	0.056	6355.620	0.520	23262.311	
		35.62×	42.98 ×	40.94 ×	26.20×	33.61×		

TABLE 2: Energy values are reported in milli Joules and time is reported in seconds. GA was run for 100 epochs with a pool size of 1000. More results are available in Appendix. The first column shows source task that we generate sponge examples, and the second column shows the target task to launch these sponge examples. The performance of the sponge examples are evaluated on three hardware platforms (ASIC, GPU and CPU).




Figure 2: Black-box attack performance of sponge examples on different hardware metrics against English-to-French translation model [65]. We show two Black-box attackers (GPU Energy and GPU Time attacker) and one White-box attacker, all using GA as the optimisation for finding sponge examples.

responses spanning thousands of characters. This finding bridges ML to the field of classic computer security and suggests that decades of experience with managing service-denial attacks can be applied here.

Ethics: Having established the similarity of Sponge examples to DoS attacks, it is appropriate to discuss the ethics of the experiments. First, we paid for the translation service and used only legitimate and well-formed API calls. For experiments and testing, we performed around 200k queries. Second, to minimise the impact of sponge examples on Azure and CO₂ production, we chose relatively small input and pool sizes. Although the maximum input size that Azure accepts is 10000 characters per request, we used only 50. We expect that the impact of sponges can be further increased by running GA with a larger input size [71]. Third, we ran the experiment at night in the data-center timezone, when it is easier to cool the servers and energy costs are lower. Fourth, to minimise the interaction of sponges between each other we executed a single sample at a time. Finally, we followed our standard

responsible disclosure process: we notified Microsoft of the vulnerability of their Translator to sponge examples and shared a draft of the paper with the Microsoft Azure team. We want to stress that sponges are not specific to Microsoft Azure, and other ML-as-a-Service providers will be affected by the same attack. Microsoft Azure was chosen because of our experience with the Azure platform. Since the discovery, in a joint effort with Microsoft, sponge examples have been added to the MITRE attack framework for AI security⁶.

5.6. Section summary

In this section, we demonstrate the effectiveness of sponge examples in different attack setups (White-box, Interactive Black-box and Blind Adversary). We consider a set of state-of-the-art sequence learning models, such as BERT [61] and RoBERTa [60]; and a wide variety

6. <https://github.com/mitre/advmthreatmatrix>




Figure 3: Maximum latency of the Microsoft Azure Translator model as is observed on the requesting server (a); reported by Azure servers (b). Azure servers were located on the same continent as the requesting server. Natural data mean baseline is at 1ms. We report multiple attack runs to show that the attack performs consistently with multiple restarts and the performance is not specific to the throttling of the user account.

of tasks. The performance of sponge examples are task-dependent and also model-dependent, however, all of the evaluated models and tasks show significant latency and energy increase when they are under attack. In addition, we demonstrate the transferability of sponge examples not only across hardware platforms but also across model architectures. Finally, we demonstrated how sponge examples can be generated in a Black-box manner on existing ML-as-a-service platforms, greatly increasing the response latency.

6. Sponge Examples on Vision Models

6.1. Models and Datasets

We evaluate the sponge example attack on a range of vision models provided in the TorchVision library. We show the performance of ResNet-18, ResNet-50 and ResNet-101 [72], DenseNet-121, DenseNet-161, DenseNet-201 [73], and MobileNet-V2 [74]. Networks span a range of sizes from 3.4M parameters (MobileNet-V2) to 49M (ResNet-101). The considered networks also have a relatively large architectural diversity, where MobileNet-V2 is designed to run on modern battery-powered mobile devices. All of the networks classify a canonical computer vision task – ImageNet-2017, since the ImageNet challenge serves as a golden baseline in the computer vision community.

6.2. White-Box Sponge Examples

Following objectives in Section 4.2.1 and Section 4.2.2, we can increase energy consumption by increasing either computation dimension or data density. Although theoretically we can provide larger images to increase the computation dimension for computer vision networks, very few modern networks currently deal with dynamic input or output. Usually preprocessing normalizes variable-sized images to a pre-defined size by either cropping or

scaling. Therefore, for computer vision models, we focus on increasing energy and latency via data density.

Table 3 shows the performance of sponge examples on CV models, and we focused on using White-box attacks to maximise energy consumption. We use both the White-box GA and White-box L-BFGS to generate sponge examples (named sponge and sponge L-BFGS in Table 3). Since the energy consumption is lower per inference, it is challenging to get a true measurement of energy given the interference of the GPU’s hardware temperature control, and that energy inspection tools lack the resolution. We then show the ASIC Energy readings and the Energy Ratio in the first two columns. The Energy Ratio term refers to the cost on an ASIC with data sparsity optimisations compared to the cost on an ASIC without any optimisations. We considered data sparsity optimisations including compressed DRAM accesses, zero-skipping multiplications. These optimisation techniques are widely adopted in many proposed ASIC accelerators [26], [49], [75], and there are now real implementations of these techniques in hardware. We then further look at the internals of neural networks and show how their data density is changing with different types of samples. We calculate the theoretical upper bounds of data density using Interval Bound Propagation (IBP) [76]. Although originally developed for certifiable robustness, we adopt the technique to look at internal network bounds that only take value 0 (*i.e.* lower bound = upper bound = 0) for the whole natural image range⁷. We also look at data densities after the ReLU function (Post-ReLU) and the overall densities. The results for density and energy suggest that both attacks can successfully generate sponge examples that are marginally more expensive in terms of energy. To be precise, we were able to get a 1 – 3% increase in energy consumption when compared to natural samples. Interestingly we observe that more of the density impact comes in the first few layers. To better understand

7. Note that we assume full floating point precision here. In practice, emerging hardware often uses much lower quantization which will result in a lower maximum data density.

		Energy			Density	
		ASIC Energy [mJ]	Energy ratio	Post-ReLU	Overall	Maximum
<i>ImageNet</i>						
ResNet-18	Sponge LBFGS	53.359 ± 0.004	0.899	0.685	0.896	0.981
	Sponge	51.816 ± 0.271	0.873	0.599	0.869	
	Natural	51.745 ± 0.506	0.871	0.596	0.869	
	Random	49.685 ± 0.008	0.837	0.480	0.834	
ResNet-50	Sponge LBFGS	164.727 ± 0.062	0.863	0.619	0.885	0.998
	Sponge	160.887 ± 0.609	0.843	0.562	0.868	
	Natural	160.573 ± 1.399	0.842	0.572	0.867	
	Random	155.819 ± 0.016	0.817	0.483	0.845	
ResNet-101	Sponge LBFGS	258.526 ± 0.028	0.857	0.597	0.873	0.994
	Sponge	254.182 ± 0.561	0.842	0.556	0.861	
	Natural	253.004 ± 1.345	0.839	0.545	0.857	
	Random	249.026 ± 0.036	0.825	0.507	0.846	
DenseNet-121	Sponge LBFGS	152.595 ± 0.050	0.783	0.571	0.826	0.829
	Sponge	149.564 ± 0.502	0.767	0.540	0.814	
	Natural	147.247 ± 1.199	0.755	0.523	0.804	
	Random	144.366 ± 0.036	0.741	0.487	0.792	
DenseNet-161	Sponge LBFGS	288.427 ± 0.087	0.726	0.435	0.764	0.811
	Sponge	287.153 ± 0.575	0.723	0.429	0.761	
	Natural	282.296 ± 2.237	0.711	0.404	0.751	
	Random	279.270 ± 0.065	0.703	0.387	0.744	
DenseNet-201	Sponge LBFGS	237.745 ± 0.156	0.756	0.505	0.788	0.863
	Sponge	239.845 ± 0.522	0.763	0.519	0.794	
	Natural	234.886 ± 1.708	0.747	0.487	0.781	
	Random	233.699 ± 0.098	0.743	0.479	0.777	
MobileNet v2	Sponge LBFGS	87.511 ± 0.011	0.844	0.692	0.890	0.996
	Sponge	84.513 ± 0.386	0.815	0.645	0.868	
	Natural	85.077 ± 0.683	0.821	0.646	0.873	
	Random	80.807 ± 0.022	0.779	0.567	0.844	

TABLE 3: We report the performance of two White-box attacks, Sponge and Sponge LBFGS against a number of computer vision benchmarks. They are optimised using the the GA and LBFGS respectively for finding sponge examples. We show the energy readings from the ASIC simulator and the Energy ratio. The Energy ratio is a ratio between the estimated energy of an ASIC optimised for sparse matrix multiplication and an ASIC without such optimisations. To further illustrate the internals of neural networks, we show data densities that are post-ReLU, across the entire neural network, and also the maximum possible density calculated using interval bound propagation (IBP). Details are described in Section 6.2.

the difference in performance please refer to Section D in Appendix. We show a statistical analysis across a wide range of CV models and describe the difficulties of precisely showing performance on CPUs and GPUs in Section D.

For computer vision models, we also find that different architectures will have similar class-wise computation densities and sponge examples can increase densities across model architectures.

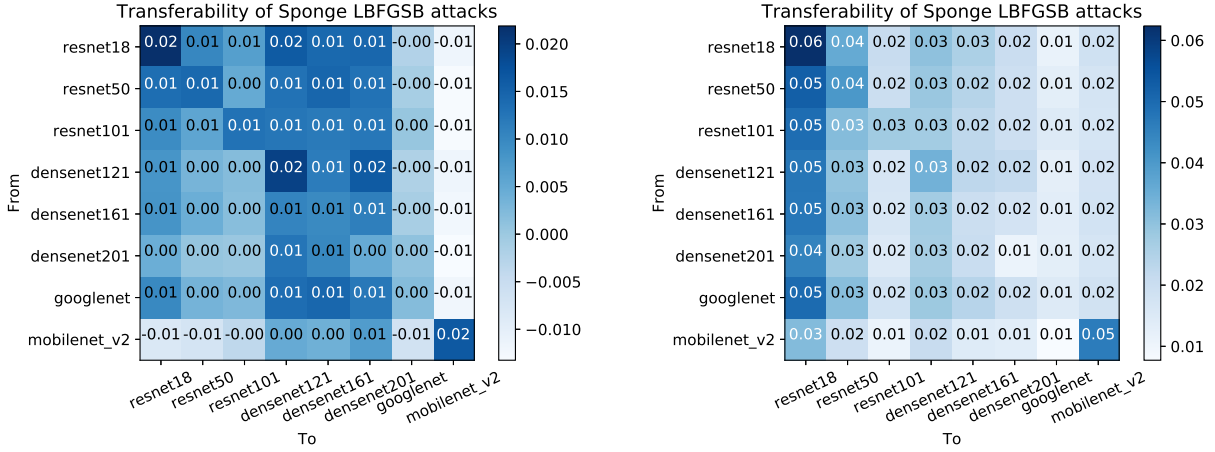
6.3. Transferability of Attacks

We observe that sponge examples are transferable and can be used to launch a blind Black-box attack. Figure 4 shows the density difference of transferred sponge samples. For all networks but one (MobileNet), the sponge samples increased the internal data density despite not having any knowledge of what natural samples look like or any

architectural and parameter information of the targeted neural networks. All of the sponge samples outperformed random noise, suggesting that sponge samples target specific features of the data set and can be applied in a blind Black-box fashion.

6.4. Class-Wise Natural Data Density

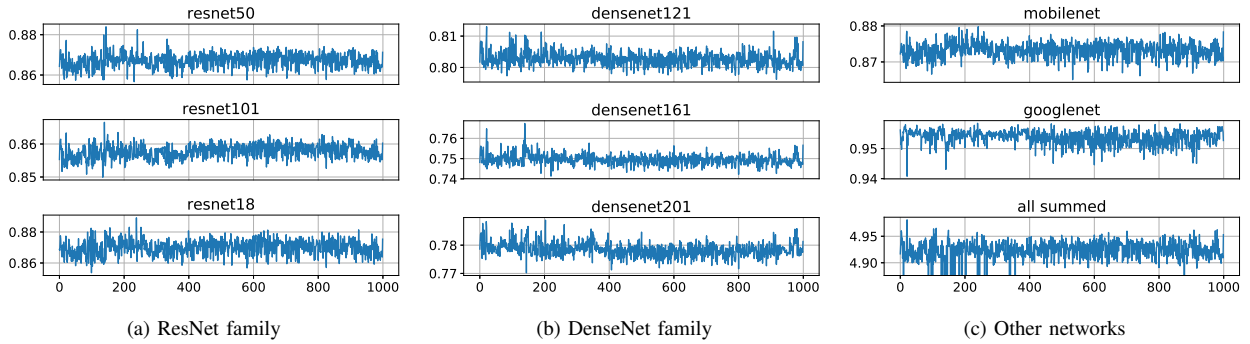
Figure 5 shows the densities of natural samples from the ImageNet dataset. On the horizontal axis, we show the 1000 classes of ImageNet and the vertical axis displays the run-time data densities for samples in that class. It can be clearly seen that there are per-class similarities between data densities of natural samples. These are particularly pronounced within ResNet and DenseNet architectures hinting that similar architectures will learn similar features so that samples of the same class have similar run-time densities across architectures. Finally, Figure 5.c shows



(a) Sponge density - Normal density

(b) Sponge density - Random density

Figure 4: Transferability of sponge examples across different computer vision benchmarks.



(a) ResNet family

(b) DenseNet family

(c) Other networks

Figure 5: Class-wise average densities of natural samples from the ImageNet validation dataset. Some classes are a lot more densely represented internally than others. X-axis shows the class numbers, whereas Y-axis shows densities.

the summed per-class densities across all of the tested networks. There are classes that are consistently more dense than others. This test is helping us to summarise that, in computer vision tasks, there exist natural samples that are producing more computation because of the increased data densities. This intrinsic property suggests that an adversary may send *natural* samples resulting in higher activation density to drain the energy of targeted devices.

6.5. Section summary

In this section, we report the results of sponge examples on computer vision models. We observe that sponge examples can successfully decrease the run-time sparsity of CV models, thus generating marginally more energy-consuming samples for ASICs that utilise data sparsity. The generated hardware differences are too small to be reliably observed on GPUs, however, we show that the GPU energy readings are statistically different between normal and sponge samples in Section D.1. The computation of CNN inference is more structured and normally only handles fixed-sized inputs. This structured computation flow provides fewer opportunities for sponge examples, and only hardware devices utilising fine-grained sparsity are vulnerable to sponge attacks.

7. Discussion

7.1. Lessons from Sponge Examples

In this paper, we demonstrated novel service-denial attacks on ML components. In their current form, our attacks work best against NLP models, whose internal complexity makes domain-specific optimisations necessary. We showed that our attacks can also target hardware optimisations, suggesting that the capable attacker will always be capable of exploiting different optimisations across the stack.

Our attacks will have a significant impact on how future ML pipelines will be deployed. The integration of different ML components can only lead to higher complexity, which will in turn be even more vulnerable to the sponge attacks described here. They may lead to ML deadlocks or livelocks, of which another precursor may be semi-trained RL agents that forever walk in circles.

The attacks presented in this paper assume that a single sample is processed at a time. This enables simple demonstrations but these are only a starting point. More complex attacks could involve samples that interact with each other. Indeed, in the world of Federated Learning, aggregators that experience delays in presence of network

failure are likely to find this leads to significant increases in overall latency. Coordinated attacks on federated systems are a natural frontier for future research.

Our attacks also used two main optimisation strategies, but others can be tried. In our attack on the Microsoft Azure Translator, it appears that a caching mechanism was making previously potent samples perform poorly in subsequent runs, but we still managed, using a genetic algorithm, to find powerful sponge examples.

Our attacks were used offensively in this paper, but such examples should also be used routinely in safety testing. Worst-case examples may also turn up by happenstance, so it is prudent to use adversarial techniques to find them in advance. Furthermore, our methodology can be used to automatically discover timing side-channels and other latent dependencies between interacting components in both ML and traditional processing pipelines.

Finally, sponge examples show that commonly deployed API rate limiting defences are not enough to protect the availability of the underlying machine learning system. Indeed, the attacker can use sponges to increase consumption of the overall system per sample without increasing the rate at which the system is queried.

7.2. Defending against Sponge Examples

Sponge examples can be found by adversaries with limited knowledge and capabilities, making the threat realistic. We further showed that sponge examples work against a deployed system that is available on demand. We now propose a simple defence to preserve the availability of hardware accelerators in the presence of sponge examples.

In Table 1, we observe a large energy gap between natural examples and random or sponge examples. We propose that before deploying the model, natural examples get profiled to measure the time or energy cost of inference. The defender can then fix a *cut-off* threshold. This way, the maximum consumption of energy per inference run is controlled and sponge examples will simply result in an error message. In this way, their impact on availability can be bounded.

This will often be sufficient to deal with the case where the threat model is battery drainage. Where the threat is a jamming attack on real-time performance, as with the vision system of an autonomous vehicle, the system will need to be designed for worst-case performance, and if need be a fallback driving mechanism should be provided. Current draft safety standards call for a self-driving car that becomes confused to slow to a stop in the same lane while alerting the human driver. This may be a good example of the appropriate response to an attack.

No single solution can tackle all of the possible abuse cases where an attacker can introduce errors into a machine-learning system. Depending on the setup both defence [77], [78] and detection [79], [80] mechanisms may be required. That problem space may be as large as the human conflict itself. At the level of technical research, serious work is needed to assess what impact different hardware platforms (*e.g.* TPUs that do not exploit sparsity) have on susceptibility to sponge examples. Above all, it is vital to take a whole system view when engineering for security or safety; to consider what threats and hazards are realistic; and to provide appropriate defences or mitigation.

In the case of attacks that cannot be prevented, the optimal strategy may often be attack detection.

7.3. Energy and Machine Learning

Most of the prior research on the carbon footprint of machine learning focuses on the energy required to train large neural network models and its contribution to carbon emissions [17], [81], [82]. This work shows that we need to study energy use at small scales as well as large. As with side-channel attacks on cryptographic systems, the fine-grained energy consumption of neural networks is a function of the inputs. In this case, the main consequence is not leakage of confidential information but a denial-of-service attack.

First, sponge examples can aim to drain a device's batteries; the operations and memory access in inference account for around a third of the work done during a complete backpropagation step, but inference happens at a much higher frequency and scale compared to training once a model is deployed. Our research characterizes the worst-case energy consumption of inference. This is particularly pronounced with natural-language processing tasks, where the worst case can take dozens of times more energy than the average case.

Second, the sponge examples found by our attacks can be used in a targeted way to cause an embedded system to fall short of its performance goal. In the case of a machine-vision system in an autonomous vehicle, this might enable an attacker to confuse the scene understanding mechanisms and crash the vehicle; in the case of a missile guided by a neural network target tracker, a sponge example might break the tracking lock. The lesson is that system engineers must think about adversarial worst-case performance and test it carefully.

8. Reproducibility

It should be noted that the performance of our attacks will vary greatly across different hardware platforms and even weather outside. When running experiments in a Black-box setup on two servers with similar configurations in some cases we found the energy and latency varied by up to a factor of 10. To help reproducibility, we release the sponge examples we found, the attack code-base we used and the ASIC simulator⁸.

9. Conclusion

We introduced energy-latency attacks, which enable an adversary to increase the latency and energy consumption of ML systems to deny service. Our attacks use specially-crafted sponge examples and are effective against deep neural networks in a spectrum of threat models that realistically capture current deployments of ML – whether as a service or on edge devices. They can be mounted by adversaries whose access varies from total to none at all. As proof of concept, we showed that we can slow down translation in Microsoft Azure by a factor of several thousand. Our work demonstrates the need for careful worst-case analysis of the latency and energy consumption of computational systems that use deep learning mechanisms.

8. https://github.com/iliaishacked/sponge_examples

Acknowledgment

We thank the reviewers for their insightful feedback. We want to explicitly thank Nicholas Carlini, Florian Tramèr, Adelin Travers, Varun Chandrasekaran and Nicholas Boucher for their help and comments. This work was supported by CIFAR (through a Canada CIFAR AI Chair), by EPSRC, by Apple, by Bosch Forschungsstiftung im Stifterverband, by NSERC, and by a gift from Microsoft. We also thank the Vector Institute’s sponsors.

References

- [1] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrndić, P. Laskov, G. Giacinto, and F. Roli, “Evasion attacks against machine learning at test time,” in *Joint European conference on machine learning and knowledge discovery in databases*. Springer, 2013, pp. 387–402.
- [2] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, “Intriguing properties of neural networks,” *arXiv preprint arXiv:1312.6199*, 2013.
- [3] B. Nelson, M. Barreno, F. J. Chi, A. D. Joseph, B. I. Rubinstein, U. Saini, C. A. Sutton, J. D. Tygar, and K. Xia, “Exploiting machine learning to subvert your spam filter.” *LEET*, vol. 8, pp. 1–9, 2008.
- [4] M. Jagielski, A. Oprea, B. Biggio, C. Liu, C. Nita-Rotaru, and B. Li, “Manipulating machine learning: Poisoning attacks and countermeasures for regression learning,” in *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2018, pp. 19–35.
- [5] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, “Membership inference attacks against machine learning models,” in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 3–18.
- [6] A. Salem, Y. Zhang, M. Humbert, P. Berrang, M. Fritz, and M. Backes, “MI-leaks: Model and data independent membership inference attacks and defenses on machine learning models,” *arXiv preprint arXiv:1806.01246*, 2018.
- [7] C. A. C. Choo, F. Tramer, N. Carlini, and N. Papernot, “Label-only membership inference attacks,” *arXiv preprint arXiv:2007.14321*, 2020.
- [8] 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 *28th USENIX Security Symposium (USENIX Security 19)*. Santa Clara, CA: USENIX Association, Aug. 2019, pp. 497–514. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity19/presentation/hong>
- [9] B. Biggio and F. Roli, “Wild patterns: Ten years after the rise of adversarial machine learning,” *Pattern Recognition*, vol. 84, pp. 317–331, 2018.
- [10] N. Papernot, P. McDaniel, A. Sinha, and M. Wellman, “Towards the science of security and privacy in machine learning,” *arXiv preprint arXiv:1611.03814*, 2016.
- [11] F. Palmieri, S. Ricciardi, U. Fiore, M. Ficco, and A. Castiglione, “Energy-oriented denial of service attacks: an emerging menace for large cloud infrastructures,” *The Journal of Supercomputing*, vol. 71, no. 5, pp. 1620–1641, 2015.
- [12] T. Martin, M. Hsiao, D. Ha, and J. Krishnaswami, “Denial-of-service attacks on battery-powered mobile computers,” in *Second IEEE Annual Conference on Pervasive Computing and Communications, 2004. Proceedings of the*. IEEE, 2004, pp. 309–318.
- [13] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [14] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers *et al.*, “In-datacenter performance analysis of a tensor processing unit,” in *Proceedings of the 44th Annual International Symposium on Computer Architecture*, 2017, pp. 1–12.
- [15] C. V. M. L. Team, “An on-device deep neural network for face detection,” in *Apple Machine Learning Journal*, 2017.
- [16] P. Kocher, J. Horn, A. Fogh, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher *et al.*, “Spectre attacks: Exploiting speculative execution,” in *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 1–19.
- [17] E. Strubell, A. Ganesh, and A. McCallum, “Energy and policy considerations for deep learning in nlp,” *arXiv preprint arXiv:1906.02243*, 2019.
- [18] A. Bernal, S. Fok, and R. Pidaparthi, “Financial market time series prediction with recurrent neural networks,” *State College: Citeseer.[Google Scholar]*, 2012.
- [19] M. Edel and E. Köppe, “Binarized-blstm-rnn based human activity recognition,” in *2016 International conference on indoor positioning and indoor navigation (IPIN)*. IEEE, 2016, pp. 1–7.
- [20] I. Intel, “Intel architecture instruction set extensions programming reference,” *Intel Corp., Mountain View, CA, USA, Tech. Rep.*, pp. 319433–030, 2016.
- [21] S. Markidis, S. W. Der Chien, E. Laure, I. B. Peng, and J. S. Vetter, “Nvidia tensor core programmability, performance & precision,” in *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. IEEE, 2018, pp. 522–531.
- [22] K. Hazelwood, S. Bird, D. Brooks, S. Chintala, U. Diril, D. Dzhulgakov, M. Fawzy, B. Jia, Y. Jia, A. Kalro *et al.*, “Applied machine learning at facebook: A datacenter infrastructure perspective,” in *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2018, pp. 620–629.
- [23] E. Chung, J. Fowers, K. Ovtcharov, M. Papamichael, A. Caulfield, T. Massengill, M. Liu, D. Lo, S. Alkalay, M. Haselman *et al.*, “Serving dnns in real time at datacenter scale with project brainwave,” *IEEE Micro*, vol. 38, no. 2, pp. 8–20, 2018.
- [24] N. Jouppi, C. Young, N. Patil, and D. Patterson, “Motivation for and evaluation of the first tensor processing unit,” *IEEE Micro*, vol. 38, no. 3, pp. 10–19, 2018.
- [25] Y.-H. Chen, T.-J. Yang, J. Emer, and V. Sze, “Eyeriss v2: A flexible accelerator for emerging deep neural networks on mobile devices,” *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 9, no. 2, pp. 292–308, 2019.
- [26] S. Han, X. Liu, H. Mao, J. Pu, A. Pedram, M. A. Horowitz, and W. J. Dally, “Eie: efficient inference engine on compressed deep neural network,” *ACM SIGARCH Computer Architecture News*, vol. 44, no. 3, pp. 243–254, 2016.
- [27] Y. Zhao, X. Gao, X. Guo, J. Liu, E. Wang, R. Mullins, P. Y. Cheung, G. Constantinides, and C.-Z. Xu, “Automatic generation of multi-precision multi-arithmetic cnn accelerators for fpgas,” in *2019 International Conference on Field-Programmable Technology (ICFPT)*. IEEE, 2019, pp. 45–53.
- [28] L. A. Barroso, “The price of performance,” *Queue*, vol. 3, no. 7, pp. 48–53, 2005.
- [29] C. Li, Z. Wang, X. Hou, H. Chen, X. Liang, and M. Guo, “Power attack defense: Securing battery-backed data centers,” *ACM SIGARCH Computer Architecture News*, vol. 44, no. 3, pp. 493–505, 2016.
- [30] G. Somani, M. S. Gaur, D. Sanghi, and M. Conti, “Ddos attacks in cloud computing: Collateral damage to non-targets,” *Computer Networks*, vol. 109, pp. 157–171, 2016.
- [31] Z. Xu, H. Wang, Z. Xu, and X. Wang, “Power attack: An increasing threat to data centers,” in *NDSS*, 2014.
- [32] Z. Xu, H. Wang, and Z. Wu, “A measurement study on co-residence threat inside the cloud,” in *24th {USENIX} Security Symposium ({USENIX} Security 15)*, 2015, pp. 929–944.
- [33] U. Fiore, F. Palmieri, A. Castiglione, V. Loia, and A. De Santis, “Multimedia-based battery drain attacks for android devices,” in *2014 IEEE 11th Consumer Communications and Networking Conference (CCNC)*. IEEE, 2014, pp. 145–150.
- [34] A. Tang, S. Sethumadhavan, and S. Stolfo, “{CLKSCREW}: exposing the perils of security-oblivious energy management,” in *26th {USENIX} Security Symposium ({USENIX} Security 17)*, 2017, pp. 1057–1074.
- [35] X. Chen, K. Makki, K. Yen, and N. Pissinou, “Sensor network security: a survey,” *IEEE Communications Surveys & Tutorials*, vol. 11, no. 2, pp. 52–73, 2009.

- [36] D. Anderson, J. Dykes, and E. Riedel, "More than an interface-scsi vs. ata."
- [37] R. Efraim, R. Ginosar, C. Weiser, and A. Mendelson, "Energy aware race to halt: A down to earth approach for platform energy management," *IEEE Computer Architecture Letters*, vol. 13, no. 1, pp. 25–28, 2014.
- [38] E. Tabassi, K. J. Burns, M. Hadjimichael, A. D. Molina-Markham, and J. T. Sexton, "A taxonomy and terminology of adversarial machine learning."
- [39] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, "Practical black-box attacks against machine learning," in *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, 2017, pp. 506–519.
- [40] P.-Y. Chen, H. Zhang, Y. Sharma, J. Yi, and C.-J. Hsieh, "Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models," in *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, 2017, pp. 15–26.
- [41] M. Nieves, K. Dempsey, and V. Pillitteri, "An introduction to information security," National Institute of Standards and Technology, Tech. Rep., 2017.
- [42] P. Ferguson and D. Senie, "rfc2827: network ingress filtering: defeating denial of service attacks which employ ip source address spoofing," 2000.
- [43] J. Bellardo and S. Savage, "802.11 denial-of-service attacks: Real vulnerabilities and practical solutions," in *USENIX security symposium*, vol. 12. Washington DC, 2003, pp. 2–2.
- [44] A. Erba, R. Taormina, S. Galelli, M. Pogliani, M. Carminati, S. Zanero, and N. O. Tippenhauer, "Real-time evasion attacks with physical constraints on deep learning-based anomaly detectors in industrial control systems," *CoRR*, vol. abs/1907.07487, 2019. [Online]. Available: <http://arxiv.org/abs/1907.07487>
- [45] M. Horowitz, "1.1 computing's energy problem (and what we can do about it)," in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, 2014, pp. 10–14.
- [46] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [47] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Advances in neural information processing systems*, 2013, pp. 3111–3119.
- [48] H. Sharma, J. Park, N. Suda, L. Lai, B. Chau, V. Chandra, and H. Esmaeilzadeh, "Bit fusion: Bit-level dynamically composable architecture for accelerating deep neural network," in *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2018, pp. 764–775.
- [49] A. Parashar, M. Rhu, A. Mukkara, A. Puglielli, R. Venkatesan, B. Khailany, J. Emer, S. W. Keckler, and W. J. Dally, "Scnn: An accelerator for compressed-sparse convolutional neural networks," *ACM SIGARCH Computer Architecture News*, vol. 45, no. 2, pp. 27–40, 2017.
- [50] M. Nikolić, M. Mahmoud, A. Moshovos, Y. Zhao, and R. Mullins, "Characterizing sources of ineffectual computations in deep learning networks," in *2019 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. IEEE, 2019, pp. 165–176.
- [51] V. Kathail, "Xilinx vitis unified software platform," in *The 2020 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, 2020, pp. 173–174.
- [52] X. Gao, Y. Zhao, Ł. Dudziak, R. Mullins, and C.-z. Xu, "Dynamic channel pruning: Feature boosting and suppression," *arXiv preprint arXiv:1810.05331*, 2018.
- [53] W. Hua, Y. Zhou, C. M. De Sa, Z. Zhang, and G. E. Suh, "Channel gating neural networks," in *Advances in Neural Information Processing Systems*, 2019, pp. 1886–1896.
- [54] W. Xu, Y. Qi, and D. Evans, "Automatically evading classifiers," in *Proceedings of the 2016 network and distributed systems symposium*, vol. 10, 2016.
- [55] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, "A limited memory algorithm for bound constrained optimization," *SIAM Journal on scientific computing*, vol. 16, no. 5, pp. 1190–1208, 1995.
- [56] M. Hähnel, B. Döbel, M. Völp, and H. Härtig, "Measuring energy consumption for short code paths using rapl," *ACM SIGMETRICS Performance Evaluation Review*, vol. 40, no. 3, pp. 13–17, 2012.
- [57] K. N. Khan, M. Hirki, T. Niemi, J. K. Nurminen, and Z. Ou, "Rapl in action: Experiences in using rapl for power measurements," *ACM Trans. Model. Perform. Eval. Comput. Syst.*, vol. 3, no. 2, Mar. 2018. [Online]. Available: <https://doi.org/10.1145/3177754>
- [58] S. Sen, N. Imam, and C. Hsu, "Quality assessment of gpu power profiling mechanisms," in *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, 2018, pp. 702–711.
- [59] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, "fairseq: A fast, extensible toolkit for sequence modeling," in *Proceedings of NAACL-HLT 2019: Demonstrations*, 2019.
- [60] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [61] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [62] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman, "Glue: A multi-task benchmark and analysis platform for natural language understanding," *arXiv preprint arXiv:1804.07461*, 2018.
- [63] A. Wang, Y. Pruksachatkun, N. Nangia, A. Singh, J. Michael, F. Hill, O. Levy, and S. Bowman, "Superglue: A stickier benchmark for general-purpose language understanding systems," in *Advances in Neural Information Processing Systems*, 2019, pp. 3261–3275.
- [64] M. Ott, S. Edunov, D. Grangier, and M. Auli, "Scaling neural machine translation," *arXiv preprint arXiv:1806.00187*, 2018.
- [65] S. Edunov, M. Ott, M. Auli, and D. Grangier, "Understanding back-translation at scale," *arXiv preprint arXiv:1808.09381*, 2018.
- [66] N. Ng, K. Yee, A. Baevski, M. Ott, M. Auli, and S. Edunov, "Facebook fair's WMT19 news translation task submission," *CoRR*, vol. abs/1907.06616, 2019. [Online]. Available: <http://arxiv.org/abs/1907.06616>
- [67] P. Koehn, H. Hoang, A. Birch, C. Callison-Burch, M. Federico, N. Bertoldi, B. Cowan, W. Shen, C. Moran, R. Zens *et al.*, "Moses: Open source toolkit for statistical machine translation," in *Proceedings of the 45th annual meeting of the association for computational linguistics companion volume proceedings of the demo and poster sessions*, 2007, pp. 177–180.
- [68] R. Sennrich, B. Haddow, and A. Birch, "Neural machine translation of rare words with subword units," *CoRR*, vol. abs/1508.07909, 2015. [Online]. Available: <http://arxiv.org/abs/1508.07909>
- [69] N. Ng, K. Yee, A. Baevski, M. Ott, M. Auli, and S. Edunov, "Facebook fair's wmt19 news translation task submission," *arXiv preprint arXiv:1907.06616*, 2019.
- [70] R. Barnett, "The dark side of apis: Denial of service attacks." [Online]. Available: <https://blogs.akamai.com/sitr/2018/08/the-dark-side-of-apis-denial-of-service-attacks.html>
- [71] M. Azure, "Request limits for translator." [Online]. Available: <https://docs.microsoft.com/en-us/azure/cognitive-services/translator/request-limits>
- [72] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [73] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4700–4708.
- [74] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4510–4520.




Figure 6: GA performance with WSC task from GLUE Benchmark running on GPUs. Words of size 29 are evaluated with pool sizes of 100, 300, 500, 700 and 900.

- [75] D. Kim, J. Ahn, and S. Yoo, “A novel zero weight/activation-aware hardware architecture of convolutional neural network,” in *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2017.* IEEE, 2017, pp. 1462–1467.
- [76] S. Gowal, K. Dvijotham, R. Stanforth, R. Bunel, C. Qin, J. Uesato, R. Arandjelovic, T. Mann, and P. Kohli, “On the effectiveness of interval bound propagation for training verifiably robust models,” *arXiv preprint arXiv:1810.12715*, 2018.
- [77] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” *International Conference on Learning Representations (ICLR)*, 2015.
- [78] E. Wong, L. Rice, and J. Z. Kolter, “Fast is better than free: Revisiting adversarial training,” 2020.
- [79] S. Chen, N. Carlini, and D. Wagner, “Stateful detection of black-box adversarial attacks,” in *Proceedings of the 1st ACM Workshop on Security and Privacy on Artificial Intelligence*, ser. SPAI ’20. New York, NY, USA: Association for Computing Machinery, 2020, p. 30–39. [Online]. Available: <https://doi.org/10.1145/3385003.3410925>
- [80] I. Shumailov, Y. Zhao, R. Mullins, and R. Anderson, “Towards certifiable adversarial sample detection,” in *Proceedings of the 13th ACM Workshop on Artificial Intelligence and Security*, ser. AISec’20. New York, NY, USA: Association for Computing Machinery, 2020, p. 13–24. [Online]. Available: <https://doi.org/10.1145/3411508.3421381>
- [81] P. Henderson, J. Hu, J. Romoff, E. Brunskill, D. Jurafsky, and J. Pineau, “Towards the systematic reporting of the energy and carbon footprints of machine learning,” *arXiv preprint arXiv:2002.05651*, 2020.
- [82] A. Lacoste, A. Luccioni, V. Schmidt, and T. Dandres, “Quantifying the carbon emissions of machine learning,” *arXiv preprint arXiv:1910.09700*, 2019.
- [83] B. Goel, S. A. McKee, and M. Sjalander, “Techniques to measure, model, and manage power,” in *Advances in Computers*. Elsevier, 2012, vol. 87, pp. 7–54.
- [84] E. García-Martín, C. F. Rodrigues, G. Riley, and H. Grahm, “Estimation of energy consumption in machine learning,” *Journal of Parallel and Distributed Computing*, vol. 134, pp. 75–88, 2019.
- [85] J. A. Butts and G. S. Sohi, “A static power model for architects,” in *Proceedings 33rd Annual IEEE/ACM International Symposium on Microarchitecture. MICRO-33 2000.* IEEE, 2000, pp. 191–201.

Algorithm 2: Sponge samples through a Genetic Algorithm

Result: S

initialise a random pool of inputs;

```

1  $S = \{S_0, S_1, \dots, S_n\}$ ;
2 while  $i < K$  do
   | Profile the inputs to get fitness scores;  $\Rightarrow$ 
   | latency or energy
   |  $P = \text{Fitness}(S)$ ;
   | Pick top performing samples;
   |  $\hat{S} = \text{Select}(P, S)$ ;
   | if NLP then
   | |  $S = \text{MutateNLP}(\hat{S})$ ;
   | | Concatenate samples A, B;
   | |  $\Rightarrow S = \text{LeftHalf}(A) + \text{RightHalf}(B)$ ;
   | |  $\Rightarrow S = \text{RandomlyMutate}(S)$ ;
   | end
   | if CV then
   | |  $S = \text{MutateCV}(\hat{S})$ ;
   | | Concatenate samples A, B, and a random
   | | mask;
   | |  $\Rightarrow A * \text{mask} + (1 - \text{mask}) * B$ ;
   | end
3 end
4 ;

```

Appendix A. Parameter Choices

We have thoroughly evaluated different parameter choices for the sponge attack and found that a small pool size and a relatively short number of GA iterations should be sufficient for a large number of tasks.

Figure 6 shows the performance of sponge samples on the RoBERTa model for the Winograd Schema Challenge (WSC) with different pool sizes and varying input sequence length. The horizontal axis shows the number of GA iterations. In terms of pool size of the GA, although there is an increase in performance for larger pool sizes, the increase is marginal. Also, smaller pool sizes significantly reduce the runtime for the attack. From the hardware perspective, using a large pool size might trigger GPUs to throttle, so that the runtime will be further increased. We observed that the convergence is consistently faster for smaller input sequences. This is mainly because the complexity of the search is less. In practice, we found almost all input sequence lengths we tested plateau within 100 GA iterations; even going to over 1000 iterations gives only a small increase in performance. For these reasons, for the experiments presented below, we report the results of the attack with a pool size of 1000 for GLUE and Computer Vision benchmarks and 1000 for translation tasks. We use 1000 GA iterations for all benchmarks tested. When displaying results, we normally use sponge for sponge examples produced using the GA and use sponge L-BFGS to identify sponge examples generated using L-BFGS.

Appendix B. Energy Cost Factors

Energy cost is a combination of static and dynamic energy.

$$E = (P_{\text{static}} + P_{\text{dynamic}}) \times t$$

Static power refers to the consumption of the circuitry in an idle state [83] there are multiple models to estimate this depending on the technology [83]–[85]. In this paper, we follow a coarse-grained approach. Cycle-accurate hardware simulation incurs a large run-time, but a coarse-grained energy simulator provides enough resolution to indicate the energy-consuming samples while using significantly less time per round of simulation.

$$P_{\text{static}} = \sum I_{\text{leakage}} \times V_{\text{core}} = \sum I_s \times (e^{\frac{qV_d}{kT}} - 1) \times V_{\text{core}} \quad (3)$$

where I_s is the reverse saturation current; V_d is the diode voltage; k is Boltzmann’s constant; q is the electronic charge; T is temperature and V_{core} is the supply voltage.

Dynamic power refers to consumption from charging and discharging the circuitry [84].

$$P_{\text{dynamic}} = \alpha \times C \times V_{\text{core}}^2 \times f \quad (4)$$

Here, α refers to the activity factor *i.e.* components that are currently consuming power; V_d is the source voltage; C is the capacitance; f is the clock frequency. Ultimately an attacker attempts to solve an optimisation problem

$$\begin{aligned} \max E, \text{ where } E = & \left(\left[\sum I_s \times \left(e^{\frac{qV_d}{kT}} - 1 \right) \times V_{\text{core}} \right] \right. \\ & \left. + \left[\alpha \times C \times V_{\text{core}}^2 \times f \right] \right) \times t. \end{aligned} \quad (5)$$

overheat or increase overall consumption
throttle or exploit load predictor
more activity of the board
run for longer or exploit the predictor

For all parameters considered in the equation, only four can be manipulated by the adversary described in Section 4.1: T , α , f and t . Of these, frequency and temperature cannot be controlled directly, but are affected through optimisations performed by the computing hardware. As we assume a single GPU, CPU or ASIC, we focus on the activity ratio α , the time t and the switching power from flipping the state of transistors. The execution time t and activity ratio α link tightly to the number of operations and memory accesses performed. In the temporal dimension, attackers might trigger unnecessary passes of a compute-intensive block; in the spatial domain, attackers can turn sparse operations into dense ones. These temporal and spatial attack opportunities can significantly increase the number of memory and arithmetic operations and thus create an increase in α and t to maximise energy usage.

Appendix C. Domain Specific Optimisations

In Section 4.4 we outlined the genetic algorithm we used to find sponge samples. That approach is generic. Here

we describe how we can improve the effectiveness of the genetic algorithm through domain-specific optimisations.

First, for NLP tasks, the greatest impact on performance was acquired from exploiting the encoding schemes used. While the genetic algorithm was fast to pick up this vulnerability, it struggled with efficiency around the mid-point, where the parents were concatenated. For example, when trying to break down individual sub-words to more tokens, we observed the GA inserting backslashes into the samples. When concatenated, we saw cases where two non-backslashes followed each other, meaning the GA was losing on a couple of characters. As a solution, we probabilistically flipped the halves and saw a slight improvement.

For CV tasks, we observed that random samples were always classified as belonging to the same class. Furthermore, random samples had very low internal density. We hypothesize that this has to do with the fact that on random samples there are very few class features, as opposed to what is observed in natural samples. As the GA improvement largely depends on randomness, that meant that we often observed that after merging two highly dense parents, uniform randomness across all pixels was decreasing sparsity to the level of random samples. In other words, uniform randomness was diluting class features. To counter this phenomenon, instead of applying uniform randomness across all pixel values, we resorted to diluting only 1% of them. That led to a bigger improvement in the whole population pool. Furthermore, after observing that the density is class-dependent, it became apparent that to preserve diversity in the pool it was important to keep samples from multiple classes. For this, we tried to ensure that at least 20 different classes were preserved in the pool.

We attempted to use domain knowledge and tried adding operations like rotation, transposition and re-scaling into the mutation process, yet we found that these did not lead to significant improvements.

Appendix D. Understanding Sponges and Their Performance

To better understand the results, we present Figure 7 which shows per-class density distributions of natural, random and sponge samples. There are 50,000 random and natural samples respectively and 1,000 sponge samples, with the bars normalised to form a probability density.

The first thing that becomes apparent is that randomly generated samples on CV models cost significantly less energy because many activations are not on. On average, random samples result in a sparser computation – around 4% more sparse for ResNet18 – and our simulator costs for natural samples are around 4–7% higher than the costs of random samples. Second, a surprising finding is that the most and least sparse samples are clustered in a handful of classes. In other words, certain classes have inputs that are more expensive than others in terms of energy. For ResNet-18, the most sparse classes are ‘wing’ and ‘spotlight’ while the least sparse are ‘greenhouse’ and ‘howler monkey’. We observe a similar effect for larger ResNet variants and also DenseNets, although the energy gap is smaller on DenseNets. Interestingly, we see that energy-expensive




Figure 7: Per-class mean density of samples evaluated on ResNet18 and DenseNet161. The natural samples are from the validation set and are compared to 50 000 randomly generated samples and 1000 Sponge GA samples. The scales are normalised to form a probability density.

classes are consistent across different architectures, and we further demonstrate this class-wise transferability in Section 6.4. Ultimately, this phenomenon implies that it is possible to burn energy or slow a system without much preparation, by simply bombarding the model with natural samples from energy-consuming classes. Finally, we see that the sponge samples are improving the population performance and tend to outperform natural sample. We observe that it is easier for sponge to outperform all natural samples for DensNets of different size, yet it struggles to outperform all of the ResNets. We further measure the energy performance statistically in Section D.1.

D.1. Measuring Difficulties and Statistical Analysis

Although we have presented in Section 6.2 that sponge attacks cause ASIC energy consumption to rise for computer vision tasks, it is still unclear what this translates to real life.

If one were to directly measure the CPU or GPU load per adversarial sample, interpreting it would be hard, especially when one talks about the energy cost improvements in the order of around 5% for ResNet18 and 3% as for DenseNet101. As is mentioned in Section 4.3 the main energy costs include the frequency of switching activities, voltage and clock frequency. Due to the heat impact from voltage and clock frequency, a large number of different optimisations are deployed by the hardware. Here, the optimisations try to balance multiple objectives – they try to be as performant as they can, whilst being as energy efficient as possible and also maintain reliability. Modern CPUs and GPUs have several performance modes between which the hardware can switch. For example, official Nvidia documentation lists 15 different performance modes.

Figure 8 shows measurements taken during the sponge GA attack running against ResNet-18. The x-axis shows the number of epochs, with each epoch the internal density is increasing from 0.75% to 0.8%. In (a), the right y-axis shows mean energy readings per sample, whereas left y-

axis shows mean power readings per-sample. In (b) the left y-axis shows mean latency values per-sample.

The amount of power consumed is strongly correlated to the amount of time taken by each sample. When the GPU speeds up, it consumes more energy but requires less time, but the rise in temperature causes the hardware then to go to a more conservative mode to cool down. We observe this heating and cooling cycle with all tasks running on GPUs, making it hard to measure the absolute performance and the attack impact. We can however measure the performance statistically. First, we turn to a question of

Can we detect energy differences between Natural, Random and Sponge samples?

To investigate the relationship between the samples we use Mann-Whitney-Wilcoxon U test (U-test), a nonparametric test for the difference between distributions. With three classes of samples, we need three pairwise comparisons. For each one, the null hypothesis that the distributions of energy consumed by the samples are identical. The complement hypothesis is that of a difference between distributions.

The U-test is based on three main assumptions:

- Independence between samples;
- The dependent variable is at least ordinal;
- The samples are random.

The first assumption is fulfilled since no sample belongs to more than one category i.e. natural, random and sponge. The second assumption is satisfied by the fact that both time and energy are cardinal variables. The third assumption, however, is harder to satisfy.

The cause of this lies in the closed nature of hardware optimisations: although some of the techniques are known, the exact parameters are unknown. Furthermore, it is hard to achieve the same state of the hardware even through power cycling. As was mentioned in Section 4.3 temperature affects energy directly, and it is hard to make sure that the hardware always comes back to the same state.

To minimise temperature effects we apply the load of natural, attack and random samples and wait until the




Figure 8: ResNet-18 solving ImageNet-2017 without any rate limiting with increasing internal density.

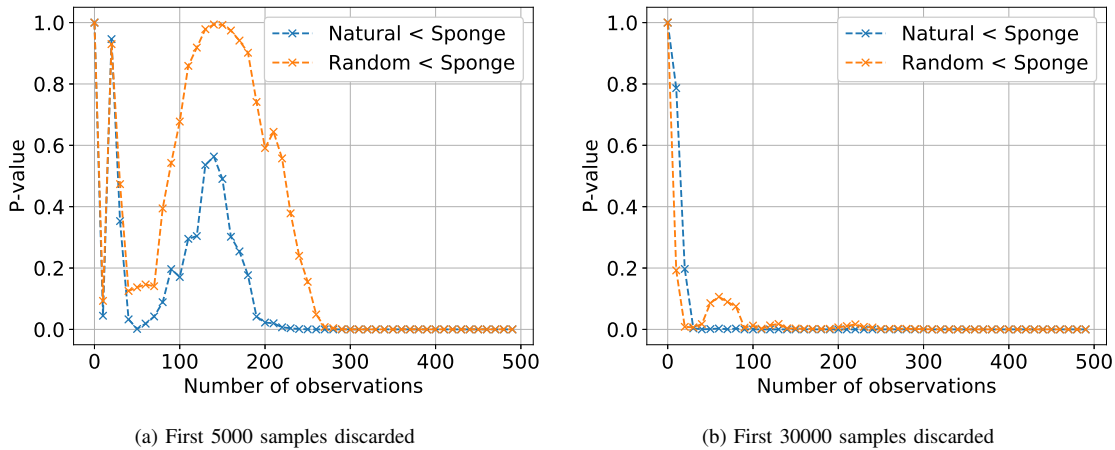


Figure 9: Mann-Whitney test on CPU measured Mobilenet execution. Number of observations is shown on x-axis and p-value on the y-axis.

temperature stabilises. That takes approximately 30000 samples. The order of the samples is random, and at this point, it can be assumed that all of the data and instruction caches are filled. Finally, because the samples are randomly shuffled, all of the predictive optimisations will work with the same probability for each of the classes.

For these reasons, we believe it is safe to assume that the samples themselves are random in that the effect of hardware optimisations is random so that the last assumption of the Mann-Whitney test is fulfilled.

Using this test we can do a pairwise comparison of the natural, random and sponge samples. The test indicates that the three types of samples generate energy consumption distributions that are statistically different (one-sided test, $p\text{-value}=0.000$) for mobilenet executed on a CPU. On a practical level, the amount of energy consumed by sponge samples is 1.5% higher on a CPU and >7% on ASIC. We could not evaluate the energy recordings on a GPU, as the standard deviation was over 15% which becomes worse as temperature increases. Figure 9 shows the confidence

of the Mann-Whitney test with mobilenet measured on the CPU as a function of the number of observations. The number of observations is on the x-axis, and the p-value on the y-axis. As can be seen, in a stable environment i.e. the temperature has stabilised, after about 100 observations per class, the differences become statistically significant at any reasonable confidence level. A similar trend is observed for unstable temperature environment, but around three times more data is required. That means that in practice, about 100–300 observations per class are sufficient to differentiate between classes with high confidence.

Appendix E. Energy–Time Relationship of Sponge Exam- ples

	Input size	GPU Energy [mJ]		GPU Time [mS]		Time \times / Energy \times
		Natural	Sponge	Natural	Sponge	
<i>SuperGLUE Benchmark with [60]</i>						
CoLA	15	1.00 \times	1.11\times	1.00\times	0.92 \times	1.21
	30	1.00 \times	1.28\times	1.00\times	0.82 \times	1.56
	50	1.00 \times	2.06\times	1.00 \times	1.27\times	1.62
MNLI	15	1.00 \times	1.12\times	1.00 \times	0.95 \times	1.26
	30	1.00 \times	1.51\times	1.00 \times	1.03 \times	1.46
	50	1.00 \times	2.16\times	1.00 \times	1.30\times	1.66
WSC	15	1.00 \times	8.89\times	1.00 \times	5.51\times	1.61
	30	1.00 \times	16.13\times	1.00 \times	11.04\times	1.46
	50	1.00 \times	26.64\times	1.00 \times	20.56\times	1.29
<i>WMT14/16 with [64]</i>						
En \rightarrow Fr	15	1.00 \times	4.32\times	1.00 \times	3.89\times	1.11
En \rightarrow De	15	1.00 \times	27.84\times	1.00 \times	24.18\times	1.15
<i>WMT18 with [65]</i>						
En \rightarrow De	15	1.00 \times	30.81\times	1.00 \times	26.49\times	1.16
<i>WMT19 with [69]</i>						
En \rightarrow Ru	15	1.00 \times	26.43\times	1.00 \times	22.85\times	1.15

TABLE 4: We use the White-box GA attack to produce sponge examples and measure the performance on different platforms and calculate how energy improvement factor relates to time improvement factor. The GPU readings are from NVML. GA was ran for 1000 epochs with a pool size of 1000. A detailed explanation of the results is in Section 5.2.

	Input size	Natural	ASIC Energy [mJ]	
			Random	Sponge
<i>SuperGLUE Benchmark with [60]</i>				
CoLA	15	504.93 \pm 1.07	566.58 \pm 2.74	583.56 \pm 0.00
	30	508.73 \pm 1.87	634.24 \pm 4.06	669.20 \pm 0.00
	50	511.43 \pm 3.64	724.48 \pm 5.12	780.57 \pm 0.59
MNLI	15	509.19 \pm 1.45	570.10 \pm 2.82	586.43 \pm 0.00
	30	514.00 \pm 2.07	638.78 \pm 3.89	672.07 \pm 0.00
	50	519.51 \pm 2.79	728.82 \pm 5.26	783.18 \pm 0.75
WSC	15	510.84 \pm 8.84	1008.59 \pm 192.22	2454.89 \pm 68.06
	30	573.78 \pm 140.12	2319.05 \pm 502.31	5012.75 \pm 154.24
	50	716.96 \pm 223.75	5093.42 \pm 1020.34	10192.41 \pm 347.32
<i>WMT14/16 with [64]</i>				
En \rightarrow Fr	15	1793.84 \pm 356.29	4961.56 \pm 1320.84	8494.36 \pm 166.22
En \rightarrow De	15	1571.59 \pm 301.69	2476.18 \pm 1586.95	48446.29 \pm 0.06
<i>WMT18 with [65]</i>				
En \rightarrow De	15	1624.05 \pm 352.99	2318.50 \pm 296.09	49617.68 \pm 0.02
<i>WMT19 with [69]</i>				
En \rightarrow Ru	15	1897.19 \pm 607.30	5380.20 \pm 2219.24	47931.11 \pm 0.00

TABLE 5: We use the White-box GA attack to produce sponge examples and measure the consistency of ASIC results. GA was ran for 1000 epochs with a pool size of 1000. A detailed explanation of the results is in Section 5.2.