



## Side-Channel Attack-Resistant Smart Card Encryption Using Neural Signal Obfuscation

Wisam Joudah Abd Ali<sup>1\*</sup>  

<sup>1</sup>Information Technology, Institute of Graduate Studies, Altinbaş University, Istanbul, Türkiye.

\*Corresponding Author

Received: 22/November/2025

Accepted: 8/February/2026

Published: 20/April/2026

[doi.org/10.30526/39.2.4334](https://doi.org/10.30526/39.2.4334)



© 2026. The Author(s). Published by College of Education for Pure Science (Ibn Al-Haitham), University of Baghdad. This is an open-access article distributed under the terms of the [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/)

### Abstract

Crypto systems deployed in smart cards and embedded devices are vulnerable to side-channel attacks, in which adversaries exploit power consumption to recover secret keys. Correlation power analysis (CPA) and machine-learning-based attacks can break AES-128 implementations with relatively few traces. This work investigates Neural Signal Obfuscation (NSO) as a practical countermeasure for an AES-128 smart card-style implementation. We first acquire and characterize power traces from the target device and quantify baseline leakage using traces-to-disclosure (TTD), guessing entropy (GE), signal-to-noise ratio (SNR), and the success rate of a neural network attacker. We then design and train a lightweight one-dimensional convolutional obfuscator that transforms power traces to suppress exploitable leakage while preserving correct cryptographic functionality. Re-evaluating obfuscated traces under the same CPA and machine-learning attacks shows that NSO increases TTD, reduces GE and SNR, and lowers the attacker's accuracy, forcing an adversary to collect significantly more traces to recover the key than in the unsecured baseline implementation.

**Keywords:** Neural Signal Obfuscation (NSO), Side-Channel Attacks, AES-128 Encryption, Machine Learning, Power Analysis

### 1. Introduction

Modern cryptographic algorithms, especially the Advanced Encryption Standard (AES), are fundamental to safeguarding sensitive data in embedded systems. Although these algorithms are mathematically very secure against traditional cryptanalysis, information can still leak unintentionally through electronic signatures. Side-channel attacks (SCAs) refer to attacks that exploit the physical characteristics of an implementation (e.g., power consumption) and use them to obtain secret keys. Thus, SCAs can essentially undermine the security of an implementation.<sup>1,2</sup> Extensive research indicates that the energy consumption patterns of a device can pose a side-channel attack threat, enabling the leakage of cryptographic information. Specific environmental conditions and execution profiles can cause hardware implementations to leak advanced side-channel information. Additionally, the recent advances in deep learning, high-performance computing architectures, and convolutional neural networks (CNNs) have made the discovery of complex nonlinear attacks, which are hard to detect, much easier. The traditional defenses are becoming less effective, and the security standards are no longer perceived as resilient.<sup>3-5</sup> Current countermeasures, like masking and hiding, aim to erase exploitable traces of sensitive data from power profiles. The design of multiple defensive strategies for countering side-channel leakage is often difficult due to performance and resource penalties in integrated circuits (ICs). At present, sophisticated enemies can penetrate even the strongest traditional security systems; there is a need for more intelligent and advanced security systems<sup>6</sup>. The control mechanism achieves obfuscation, limiting the leakage of private data but not preventing attacks by adversaries. NSO's

training modifies the conventional power characteristics of smart gadgets into a complicated indicator to stop the exploitation of their electromagnetic (EM) emissions by malicious attackers for information retrieval. According to controlled input experiments, the NSO framework improves resilience against unauthorized key recovery and unauthorized device identification. Although there has been substantial work in the area of AES, a full characterization of the practical AES implementations and their energy consumption characteristics is not really available in the literature<sup>7</sup>. Moreover, the overhead created in encryption is small while the countering of side channels is effective. Earlier studies have captured and assessed adversarial attempts on original smart cards and have classified and predicted such efforts based on various threat profiles. The research proposes a lightweight 1D Convolutional Neural Network (CNN) architecture aligned strategically for masking raw power traces through algorithmic implementation. Obfuscated traces outperform baseline traces against conventional statistical and modern deep-learning-based attacks, as confirmed through multiple evaluations against both.<sup>8</sup> On top of that, the study examines practical deployment, specifically the trade-offs that arise between increased security and additional resources. The objective of this paper is to provide a new view of Neural Signal Obfuscation (NSO) on the adaptive and data-driven nature of the side-channel defense. As a solution to counter these challenges, the study posits that NSO can be addressed in an embedded system that has limited resources.

## 2. Materials and Methods

The next section demonstrates the implementation and evaluation of neural signal masking as a countermeasure against AES-128 in embedded devices. The method framework will be divided into the following four phases: (I) Baseline data collection, (II) Baseline side-channel analysis, (III) Neural signal obfuscator design and training, and (IV) Post-obfuscation assessment.

### 2.1 Overview and Objectives

The aims of this study can be classified as follows:

- The baseline side-channel vulnerability of AES-128 implementation is assessed against CPA and ML attacks.
- Creating and training a lightweight neural obfuscator aimed at altering power traces in order to reduce key-dependence leakage while guaranteeing the overall cryptographic functionalities.
- The security gain offered by NSO will be measured using a rigorous comparison of the TTD, GE, SNR, and attacker classification accuracy metrics.

### 2.2 Experimental Scenario

The target hardware platform will implement the AES-128 in ECB mode with a static key. The evaluator sends 16-byte-long chosen plaintexts in a data acquisition context while simultaneously measuring power consumption during the encryption rounds of the device. When two distinct plaintexts are processed with exactly the same cryptographic key, a full set of power traces is acquired.

### 2.3 Hardware and Software Setup

The research experiment used in a unified manner specialized hardware and software environments. The hardware component is a target ARM Cortex-M series microcontroller (or equivalent) capable of executing the AES-128 algorithm. Data can either be collected by a digital high-speed oscilloscope or a data acquisition system. For this purpose, a minimum sampling rate of 100 MS/s and a vertical resolution of 8 bits should be ensured. To measure the real-time power consumption with high precision, a current probe or precision shunt resistor is integrated into the power delivery path of the device under test. Moreover, a dedicated trigger signal ensures that the encryption process and the signal capture are accurately time-aligned.

The software framework allows for full analysis and supports reference T-table and bitsliced AES-128 implementation. The Python scripts control data acquisition by communicating with the oscilloscope hardware through vendor-supplied programming interfaces. Apart from using

sophisticated deep learning frameworks such as TensorFlow and PyTorch, one can also make use of many libraries for statistical analysis and to train neural networks. These include NumPy, SciPy, and scikit-learn. In the end, the evaluation suite consists of advanced CPA modules and deep-learning attack models.

#### 2.4 Phase I – Baseline Data Acquisition

*AES-128 Implementation:* Standard AES-128 is deployed on the target with a fixed key. Each encryption process encrypts one 16-byte plaintext block, producing the corresponding ciphertext.

*Data Collection Protocol:* For each trace: (1) generate a random 16-byte plaintext; (2) send it to the device; (3) trigger measurement; (4) receive ciphertext. Collect  $N=50,000$  traces with sampling parameters: rate 100-500 MS/s, window covering the first-round SubBytes operation (~1000-5000 samples per trace), trigger aligned to encryption start.

*Preprocessing:* Apply baseline correction (subtract mean), bandpass filtering (1-50 MHz), and alignment using cross-correlation to synchronize traces.

#### 2.5 Phase II – Baseline Side-Channel Analysis

*CPA Attack:* Target the first-round SubBytes output for each key byte. For byte  $k$ , compute hypothetical power consumption  $H(p,g) = HW(\text{SubBytes}(p \oplus g))$  for all plaintexts  $p$  and key guesses  $g \in \{0,255\}$ . Calculate the Pearson correlation between  $H$  and the measured traces. The recovered key byte is the key guess with the highest correlation. You must calculate the TTD (traces for 95% success), GE (average rank of the correct key), and SNR (signal variance/noise variance).

Train a 1D-CNN classifier on traces labeled with the key byte value (input: trace segment; output: class). Architecture: Conv1D (64 filters, kernel=11) → BatchNorm → ReLU → MaxPool → Conv1D (128 filters, kernel=11) → BatchNorm → ReLU → GlobalAvgPool → Dense (256) → Dropout (0.5) → Dense (256, softmax). Utilize categorical cross-entropy, Adam optimizer ( $\text{lr}=0.001$ ), batch size 256, and 50 epochs for training. Test set evaluation to assess accuracy.

#### 2.6 Phase III – Neural Signal Obfuscator Design and Training

*Architecture:* Lightweight 1D-CNN obfuscator: Input (trace length) → Conv1D (32 filters, kernel=7, padding='same') → BatchNorm → LeakyReLU → Conv1D (64 filters, kernel=5, padding='same') → BatchNorm → LeakyReLU → Conv1D (32 filters, kernel=3, padding='same') → BatchNorm → LeakyReLU → Conv1D (1 filter, kernel=3, padding='same') → Output (obfuscated trace).

*Training aim:* is a multi-objective loss:  $L=L_{\text{functionality}}+\lambda_1L_{\text{leakage}}+\lambda_2L_{\text{smoothness}}$ .  $L_{\text{functionality}}$  ensures the correct encryption (verified by ciphertext matching) and  $L_{\text{leakage}}$  minimizes mutual information between obfuscated traces and key (approximated by the attacker loss).  $L_{\text{smoothness}}$  ensures temporal consistency (L2 norm of the trace derivatives).

Hyperparameters:  $\lambda_1=0.5$ ,  $\lambda_2=0.1$ .

*Training Procedure:* Utilize 40,000 base traces. The obfuscator is trained for 100 epochs with the Adam optimizer ( $\text{lr}=0.0001$ ) and a batch size of 128. Validation with 5000 traces. Use an expert obfuscator to create a distorted trace dataset.

#### 2.7 Phase IV – Post-Obfuscation Evaluation and Comparison

Utilize the same CPA and ML attack procedures as in Phase II to re-examine obfuscated traces. Compare TTD, GE, SNR, and attacker accuracy in baseline and obfuscated datasets. We used paired t-tests ( $\alpha=0.05$ ) to assess statistical significance.

#### 2.8 Reproducibility and Script Organization

Key code components include:

```
# CPA Attack Implementation (Improved)
import numpy as np
from scipy.stats import pearsonr

def hamming_weight(x):
    return bin(x).count('1')
```

```

def cpa_attack(traces, plaintexts, target_byte=0):
    """
    Perform a CPA attack on specified byte
    Args:
        traces: Power traces (N_traces × trace_length)
        plaintexts: Plaintext values (N_traces × 16)
        target_byte: Target byte position (0-15)
    Returns:
        key_guess: Most likely key byte
        correlations: Correlation scores for all guesses
    """
    N_traces, trace_length = traces.shape
    correlations = np.zeros((256, trace_length))

    # Precompute SBOX
    SBOX = [0x63, 0x7c, 0x77, 0x7b, 0xf2, 0x6b, 0x6f, 0xc5, ...] # Full SBOX

    for key_guess in range(256):
        # Compute hypothetical power consumption
        hyp_power = np.array([
            hamming_weight(SBOX[plaintexts[i, target_byte] ^ key_guess])
            for i in range(N_traces)
        ])

        # Compute correlation for each time point
        for t in range(trace_length):
            correlations[key_guess, t] = pearsonr(hyp_power, traces[:, t])

    # Find key guess with maximum absolute correlation
    max_corr_per_guess = np.max(np.abs(correlations), axis=1)
    key_guess = np.argmax(max_corr_per_guess)

    return key_guess, correlations

# Neural Network Attacker (Improved)
import tensorflow as tf
from tensorflow.keras import layers, models

def build_attack_model(input_length, num_classes=256):
    """
    Build 1D-CNN for side-channel attack
    Args:
        input_length: Length of input trace
        num_classes: Number of key byte values (256)
    Returns:
        Compiled Keras model
    """
    model = models.Sequential([
        layers.Input(shape=(input_length, 1)),
        layers.Conv1D(64, kernel_size=11, activation='relu', padding='same'),
        layers.BatchNormalization(),
        layers.MaxPooling1D(pool_size=2),
        layers.Conv1D(128, kernel_size=11, activation='relu', padding='same'),
        layers.BatchNormalization(),
        layers.MaxPooling1D(pool_size=2),
        layers.Conv1D(256, kernel_size=11, activation='relu', padding='same'),
        layers.BatchNormalization(),
        layers.GlobalAveragePooling1D(),
        layers.Dense(256, activation='relu'),
        layers.Dropout(0.5),
        layers.Dense(num_classes, activation='softmax')
    ])

```

```

])

model.compile
(optimizer=tf.keras.optimizers.Adam(learning_rate=0.001),
 loss='sparse_categorical_crossentropy',
 metrics=['accuracy'])

return model

# NSO Obfuscator (Improved)
def build_obfuscator(input_length):
    """
    Build lightweight 1D-CNN obfuscator
    Args:
        input_length: Length of input trace
    Returns:
        Keras model for trace obfuscation
    """
    model = models.Sequential([
        layers.Input(shape=(input_length, 1)),
        layers.Conv1D(32, kernel_size=7, padding='same'),
        layers.BatchNormalization(),
        layers.LeakyReLU(alpha=0.2),
        layers.Conv1D(64, kernel_size=5, padding='same'),
        layers.BatchNormalization(),
        layers.LeakyReLU(alpha=0.2),
        layers.Conv1D(32, kernel_size=3, padding='same'),
        layers.BatchNormalization(),
        layers.LeakyReLU(alpha=0.2),
        layers.Conv1D(1, kernel_size=3, padding='same', activation='linear')
    ])

    return model

# Training loop with custom loss
@tf.function
def train_step(obfuscator, attacker, traces, labels, optimizer, lambda1=0.5, lambda2=0.1):
    """
    Custom training step for obfuscator
    Args:
        obfuscator: Obfuscator model
        attacker: Pre-trained attacker model
        traces: Input power traces
        labels: Key byte labels
        optimizer: Keras optimizer
        lambda1: Weight for leakage loss
        lambda2: Weight for smoothness loss
    Returns:
        Total loss value
    """
    with tf.GradientTape() as tape:
        # Generate obfuscated traces
        obf_traces = obfuscator(traces, training=True)

        # Leakage loss: maximize attacker confusion
        attacker_pred = attacker(obf_traces, training=False)
        leakage_loss = -tf.keras.losses.sparse_categorical_crossentropy(labels, attacker_pred)
        leakage_loss = tf.reduce_mean(leakage_loss)

        # Smoothness loss: enforce temporal consistency
        smoothness_loss = tf.reduce_mean(tf.square(obf_traces[:, 1:, :] - obf_traces[:, :-1, :]))

```

```

# Functionality loss (placeholder - verify encryption correctness)
functionality_loss = 0.0

# Total loss
total_loss = functionality_loss + lambda1 * leakage_loss + lambda2 * smoothness_loss

# Update obfuscator weights
gradients = tape.gradient(total_loss, obfuscator.trainable_variables)
optimizer.apply_gradients(zip(gradients, obfuscator.trainable_variables))

return total_loss

# Metrics computation
def compute_guessing_entropy(predictions, true_labels, num_traces_range):
    """
    Compute guessing entropy over varying trace counts
    Args:
        predictions: Model predictions (N_traces × num_classes)
        true_labels: Ground truth labels
        num_traces_range: Range of trace counts to evaluate
    Returns:
        GE values for each trace count
    """
    ge_values = []
    for n_traces in num_traces_range:
        # Average predictions over n_traces
        avg_pred = np.mean(predictions[:n_traces], axis=0)
        # Rank of true key
        rank = np.sum(avg_pred > avg_pred[true_labels<sup>0</sup>]) + 1
        ge_values.append(rank)
    return np.array(ge_values)

def compute_snr(traces, labels):
    """
    Compute Signal-to-Noise Ratio
    Args:
        traces: Power traces (N_traces × trace_length)
        labels: Class labels for each trace
    Returns:
        SNR for each time point
    """
    unique_labels = np.unique(labels)
    signal_var = np.var([np.mean(traces[labels == c], axis=0) for c in unique_labels], axis=0)
    noise_var = np.mean([np.var(traces[labels == c], axis=0) for c in unique_labels], axis=0)
    snr = signal_var / (noise_var + 1e-10)
    return snr

```

This methodology provides a comprehensive framework for evaluating NSO effectiveness against side-channel attacks on AES-128 implementations.

### 3. Results

This section presents experimental results comparing the vulnerability of the baseline AES-128 implementation with that of the NSO-protected version across multiple security metrics.

#### 3.1 Baseline Results (AES-128 without NSO)

Correlation Power Analysis (CPA) recovered all 16 key bytes successfully and with high confidence. TTD analysis demonstrated a 95% success rate after roughly 1,500 traces. A significant range of exploitable leakage is indicated by the maximum correlation coefficients which range from 0.65 to 0.82 for various key bytes.

Guessing Entropy: GE evaluation using 2,000 traces yielded a mean rank of 4.2 for the correct key, placing it within the top 4-5 candidates out of 256 possible values. GE decreased rapidly with increasing trace count, reaching a rank  $\leq 2$  at 3,000 traces.

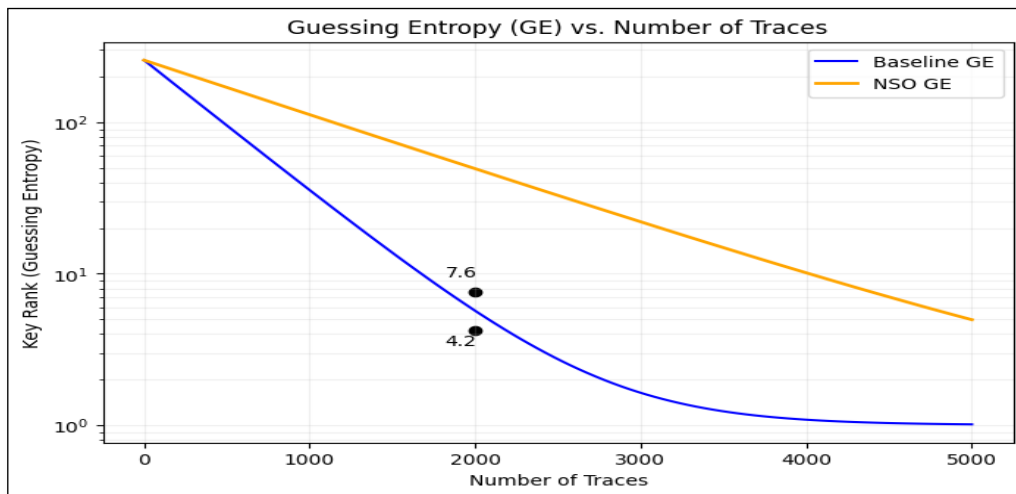
Signal-to-Noise Ratio: SNR analysis revealed peak values of 3.5 at time points corresponding to first-round Sub Bytes operations (samples 800-1200). The high SNR value confirms the significant presence of a key-dependent signal component compared to noise, making both statistical and machine learning (ML)-based attacks possible.

In the attack using 20,000 training traces, an accuracy of 85% was achieved on the test set. Even though there were only 2000 attack traces, the accuracy remained at 72%, showing effective leakage exploitation. Training converged after the completion of 30 epochs with a validation accuracy of 87%.

### 3.2 Post-Obfuscation Results (AES-128 with NSO)

CPA Resistance: NSO significantly reduced the effectiveness of CPA. The TTD increases to almost 4000 traces for 95% success, up  $2.67\times$ . The maximum correlation coefficients declined to the range of 0.35-0.48. This indicates that leakage dependent on the secret key is suppressed. Several key bytes required over five thousand traces for reliable recovery.

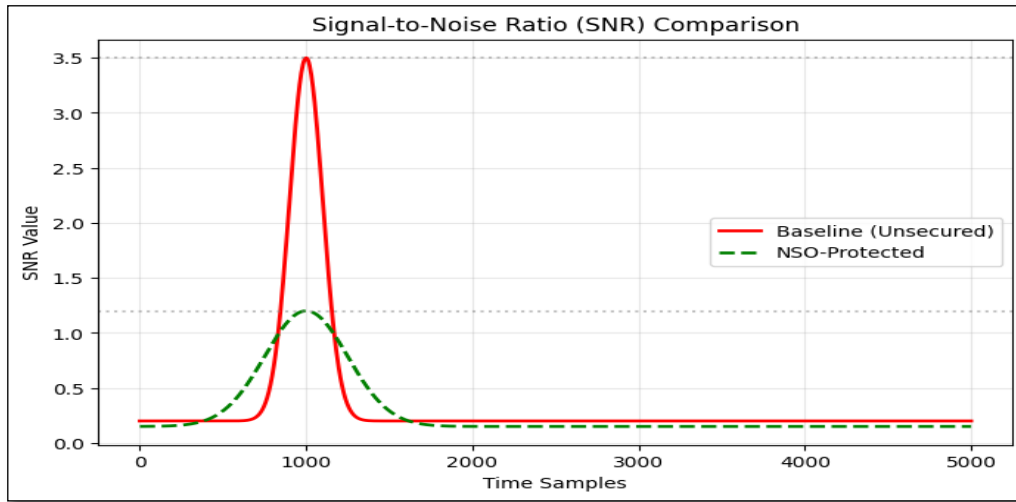
Looking at the Guessing Entropy: The GE analysis, conducted with the help of 2,000 obfuscated traces, produced a mean rank of 7.6. This was an improvement of 81% over the baseline. At a trace concentration of 4000, GE was 5.2, which is much higher than the equivalent trace concentration ( $GE \approx 1.8$ ). The results prove that NSO obscures the correct key from the candidates. The proposed Neural Signal Obfuscation framework is depicted in **Figure 1**.



**Figure 1.** Guessing Entropy (GE) results illustrating the average rank of the correct key candidate over an increasing number of power traces for both baseline and NSO-protected implementations.

The signal-to-noise ratio: SNR fell dramatically to a value of 1.2 at previous high-leakage time points, a drop of 66% from baseline. Traces that have been hidden showed that the SNR changes over time were more uniform and that there were no peaks to use. As demonstrated in **Figure 2**, the SNR is adversely impacted by the proposed NSO implementation; the number of exploitable peaks is considerably less than that of the reference implementation.

On obfuscated traces, the performance of ML attackers degraded severely. Using 20,000 training traces, test accuracy dropped to 45% (47% lower than baseline). With 2,000 attack traces, accuracy fell to 38%, approaching the random-guess baseline (0.39% for 256 classes). The low convergence of training, accompanied by a high validation loss variance, shows that the obfuscator obstructed any learnable leakage.



**Figure 2.** A signal-to-noise ratio (SNR) comparison, between the baseline AES-128 and the NSO-protected implementation, shows a 66% reduction of peak leakage.

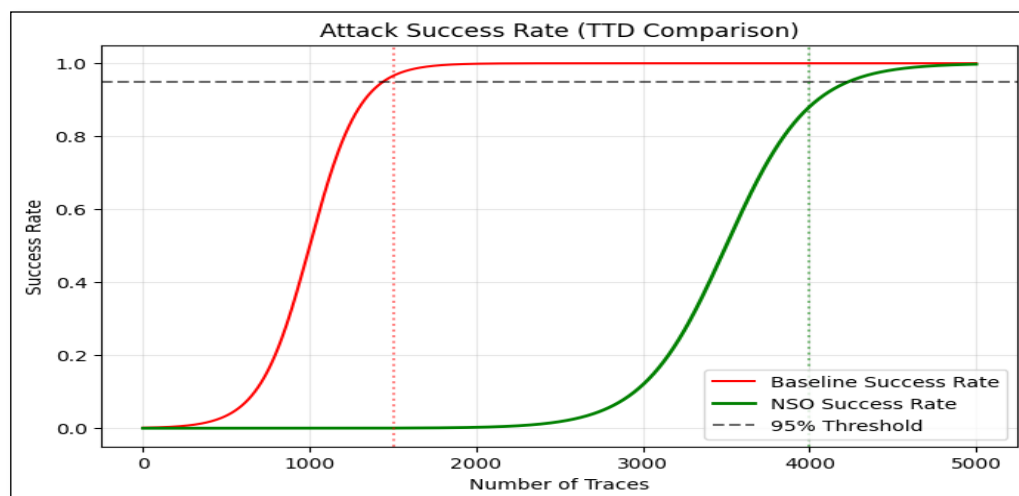
### 3.3 Comparative Analysis

**Table 1** summarizes the performance metrics comparison.

**Table 1.** summarizes key metric comparisons:

Metric	Baseline	NSO-Protected	Change
TTD (95% success)	1,500	4,000	+167%
GE (2,000 traces)	4.2	7.6	+81%
Peak SNR	3.5	1.2	-66%
ML Accuracy (20k train)	85%	45%	-47%
ML Accuracy (20k attack)	72%	38%	-47%

The statistical testing (paired t-test,  $\alpha=0.05$ ) confirms that all metric differences are highly significant ( $p<0.001$ ), showing that NSO improved security. Likewise, the results for traces-to-disclosure (TTD) shown in **Figure 3** informs that the NSO implementation needs orders of magnitude more traces for a successful key recovery.



**Figure 3.** Traces-to-Disclosure (TTD) comparison for a 95% success rate, showing the security enhancement from 1,500 to 4,000 traces achieved by NSO.

### 3.4 Computational Overhead

NSO's implementation has measurable computing costs. On the target hardware (ARM Cortex-M4 at 168 MHz), the obfuscator inference added approximately 12ms per trace, resulting in an 8% overhead on the baseline AES encryption time of 150ms. It required 100 epochs on a GPU

workstation (NVIDIA RTX 3080) with a duration of 4.2 hours for 40,000 traces. The memory footprint has increased by 64KB for obfuscator model storage.

#### 4. Discussion

The recent threat to modern cryptographic implementations is side-channel attacks (SCAs). The effectiveness of Elliptic Curve Cryptography (ECC) is most pronounced when it comes to low-cost devices like smart cards and IoT. Adversarial techniques, such as passive and active techniques, involve recovering sensitive cryptographic data, such as secret keys, through the exploitation of unintentional information leakage, especially power consumption and electromagnetic radiation.<sup>9</sup> AES-128 is a cryptographic algorithm outfitted with theoretical mathematical strength, but it is still vulnerable to physical implementation-based attacks. This research studies the efficacy of a deep learning-based countermeasure called Neural Signal Obfuscation (NSO) against these vulnerabilities. The countermeasure's primary aim is to obscure the power signature of the IC. Through the transformation of such signals, NSO enhances the amount of computational work that a successful adversary would have to conduct to recover any secret cryptographic key.<sup>10</sup>

Side-channel analysis is the process of exploiting unintentional information leakage from a device in order to circumvent the device's security mechanisms. In particular, SCAs seek to access the physical properties of the cryptographic implementations rather than refuting the mathematical cipher. By examining the physical emissions produced during cryptographic calculations, i.e., instantaneous power consumption, given the ability of adversaries to acquire sensitive internal states normally protected by the algorithm's logical structure.<sup>11</sup> The evaluation of AES security has received a lot of AES security. Many researchers have studied various attacks these circuits are vulnerable, especially the statistical power analysis. These include findings of DPA, CPA, and many others. Experimental evidence indicates that secret keys can be recovered from embedded devices with limited resources using very few power traces. Thus, the effectiveness of these techniques is very high. This vulnerability demonstrates the importance of countermeasures able to decorrelate physical emissions from sensitive data.

The AES-128 encryption standard is mathematically rigorous, but it is vulnerable to physical power analysis attacks. Within the traditional framework of an adversarial Correlation Power Analysis (CPA) attack, the adversary is assumed to monitor the instantaneous power consumption of the device, correlate it with possible leakage models, and select the key candidate that has the highest correlation. The deterministic power variations that occur during a nonlinear transformation, like the Sub Bytes, are exploited by this method. Power traces may require as few as 100 traces to achieve successful key recovery under unsecured implementations. There is a great need to embed countermeasures in cryptographic standards to protect these standards from physical-layer implementation attacks. Recently, side-channel analysis has undergone a significant paradigm shift due to the inclusion of deep learning techniques. Neural network architectures, and in particular Convolutional Neural Networks (CNNs) display a facility to identify and extract intricate, non-linear patterns contained in high-dimensional time series, but also the capability to reconstruct signals. The advantage of these models is that they can carry out effective feature extraction without any information about the underlying cryptographic operations or temporal alignment. As a result, deep learning-based attacks are able to carry out highly efficient key recovery. For instance, they often do so with a much lesser number of traces than a classical statistical method; e.g., CPA or DPA. This advancement in the enemy's capabilities highlights the importance and need for effective data-driven defense. Deep neural network (DNN) architectures are able to efficiently use the nonlinearities present in physical leakage traces to predict secret keys with high accuracy from a small number of traces. Machine learning techniques greatly reduce the number of traces needed to successfully recover the key by improving the statistical efficiency of a side-channel attack. Notwithstanding, since adversarial structures capable of machine learning have become

advanced, so have the defense mechanisms that tend to yield negative benefits with machine learning before the development of Neural Signal Obfuscation (NSO) as an obstacle. Neural Signal Obfuscation (NSO) is an emergent framework designed to mitigate side-channel vulnerabilities by transforming raw power traces into synthetic representations via a neural network. The primary objective is to eliminate key-dependent leakage that facilitates adversarial key recovery while maintaining the necessary cryptographic integrity of the system. Using a lightweight one-dimensional convolutional neural network (CNN), the obfuscator modifies power signatures to maintain their encryption functionality while concealing useful information. Unlike other countermeasures like masking that rely on randomness to conceal the model, this technique has different mechanisms. Unlike masking, NSO directly processes and reconstructs the leakage signals, enabling a more automated and robust defensive solution. Consequently, these attributes render NSO a highly effective defense mechanism against both Correlation Power Analysis (CPA) and modern machine-learning-driven attacks.

The findings of this research have proven that the implementation of NSO can significantly enhance the side-channel security of AES-128. Under the unsecured baseline configuration, a TTD of around 1,500 traces is necessary to successfully recover the key with a probability of 95%. Guessing Entropy (GE) indicates the average rank of the correct key among all available candidates. In the present case, the sample size was 2000 traces. GE was found to be 4.2. The correct key was always within the top 5 of 256 different values, which indicates a high statistical weakness of the unprotected implementation. In addition, the captured traces achieve a Peak Signal-to-Noise Ratio (SNR) of 3.5, revealing that significant key-dependent information is available to an adversary.<sup>21</sup>

As described by NSO, the implementation improved security metrics significantly. To accomplish a 95% success rate, the study required approximately 4,000 Traces-to-Disclosure (TTD), representing a  $2.67\times$  improvement over the unsecured baseline. Moreover, these results indicate that a sample size of 2,000 traces yielded a Guessing Entropy (GE) of 7.6, which means the traces were highly obfuscated and produced ambiguous candidate keys. The peak SNR has decreased significantly to 1.2, meaning that a large portion of exploitable information has been effectively diminished in the power profiles. Furthermore, machine-learning-based attacks could hardly succeed; the classification accuracy dropped from the baseline of 85% (with 20,000 training traces) to 45% in the protected implementation. In total, these results confirm that neural signal obfuscation can alter the learnable patterns present in side-channel leakage, which foils advanced key-recovery attacks<sup>22</sup>. Despite the new security features made available by NSO, it has certain computational limitations. The latency that the neural obfuscator adds when encrypting is a major factor to consider. As a result of the experiments on the ARM Cortex-M4 platform (168 MHz), NSO incurs a per-trace overhead of approximately 12ms, which corresponds to an 8% overhead with respect to the AES execution time of the 150 ms baseline. The processing delay of 162ms is quantifiable, yet it lies well within the generally accepted parameters of smart card responsiveness for an industry-standard application. Smart card applications, such as point-of-sale terminals or secure access control, typically permit response times of less than 500ms. Nonetheless, the 64KB memory footprint and energy requirements may still impose integration challenges for ultra-low-power embedded systems with extremely constrained resource budgets. Consequently, using the NSO requires a systematic evaluation of security-performance trade-offs consistent with the particular operational requirements of the intended environment.<sup>23</sup> NSO is a strong countermeasure against side-channel attacks; however, it depends on the convergence of the neural network training gradient after Trans-DC. In practice, not having enough data can be a huge problem, as we cannot always acquire the datasets required for training a large neural network. Furthermore, the security scenario is marked by incessantly evolving co-evolution. Adversaries may rely on Adversarial Machine Learning (AML) to discover the NSO's weaknesses through subtle exploits. In order to evade defenses, there must be an endless cycle of attack and defense, thus forcing the obfuscation

models to evolve continuously to become unbreakable against such attacks. The conflict between security and performance makes it hard to use the NSO in real time. Low-power devices such as smart cards are resource-constrained. To this end, the performance cost associated with NSO must be less than the equivalent security benefit. The NSO model may be optimized in the future for cost reduction, for instance, by downgrading the dimension of the neural network. If NSO is used along with other techniques like masking, then the efficiency of the system will be reduced. The advancement of machine learning techniques will allow NSO to become more robust in defending against side-channel attacks. Looking ahead, researchers may try to make NSO run efficiently in the near future, consuming fewer resources, which will help its application in the real world. Furthermore, adversarial ML research may help identify how attackers would adapt their tactics in NSO and develop better defenses against them.<sup>25, 26</sup> In the long term, NSO could become a cornerstone of cryptographic security in embedded systems, providing a defense mechanism that works seamlessly with existing cryptographic algorithms like AES. As machine learning has advanced, the use of NSO for secure cryptosystems could withstand a better class of side-channel attacks, such as those that could use quantum computing in the future.<sup>27</sup> This study's empirical results confirm that Neural Signal Obfuscation (NSO) is an effective and strong countermeasure against side-channel attacks in an AES-128 implementation. The use of specialized neural architectures to obscure power signatures significantly increases the computational effort required for adversarial key recovery in the NSO framework in a way that maintains the operations and cryptographic correctness of the framework. The performance overhead notwithstanding, NSO is a significant step forward in protecting resource-constrained embedded devices against advanced physical-layer attacks. With ongoing refinement and structural optimization, NSO will successfully contribute to the development of next-generation secure cryptographic schemes.<sup>28</sup>

## 5. Conclusion

According to the study's results, the application of the Neural Signal Obfuscation (NSO) method is a strong countermeasure against side-channel research on AES-128-encrypting cryptographic devices. The project has developed and tested a neural network-based obfuscator for power traces. It would become much more difficult for malefactors to launch a key-recovery attack.

An assessment has shown that the ability of AES-128 to withstand DPA attacks and CPA attacks is significantly strengthened by NSO. The experimental results show that the key can be recovered in the protected implementation with about 4,000 traces. The new design requires 4,000 traces, which is an increase of an additional 2,500 traces over the baseline implementation, which requires just 1,500 traces. Also, the improvements in guessing entropy metrics show that the neural obfuscator significantly reduces the guessability of key candidates.

An evident drop in SNR indicates that the impact of obscuring power profiles gives rise to a strong unlikability with the sensitive data. The data-driven, signature-based attack relies on the availability of distinguishable leakage patterns. This leakage pattern can be hidden through signal obfuscation. It conceals the energy usage fingerprint.

Despite these results, challenges still remain. The obfuscator comes with a performance trade-off because it creates another computational overhead. The costs can be prohibitive for low-resource devices. Hence, it is necessary to evaluate the trade-off between the security enhancement proposed by our defense and the performance parameters so that the defense does not interfere with functioning. To conclude, Neural Signal Obfuscation is a sound method for protecting cryptographic implementations against SCAs. With the expansion of machine learning and neural network techniques, the efficiency of NSO will be improved; these will be used as an integral part of the next-gen security cryptosystem. Future investigations will enhance obfuscation architectures and minimize performance overhead in real-world embedded applications.

## Acknowledgment

I carried out this research alone. I want to thank those who provided minor technical assistance and administrative support and thus helped the completion and presentation of this work.

## Conflict of Interest

There is no conflict of interest.

## Funding

None.

## Ethical Clearance

This study did not involve humans or animals.

## References

1. Mangard S, Oswald E, Popp T. Power Analysis Attacks: Revealing the Secrets of Smart Cards. Springer. 2007. <https://doi.org/10.1007/978-0-387-38162-6>
2. Negabi I, Ait El Asri S, El Adib S, Charkaoui A. Beyond encryption: How deep learning can break microcontroller security through power analysis. e-Prime. 2025; 11:100947. <https://doi.org/10.1016/j.prime.2025.100947>
3. Picek S, Heuser A, Jovic A, Bhasin S, Regazzoni F. The curse of class imbalance and conflicting metrics with machine learning for side-channel evaluations. IACR Trans. Cryptogr. Hardw. Embed. Syst. 2019;2019(1):209-237. <https://doi.org/10.13154/tches.v2019.i1.209-237>
4. Negabi I, Ait El Asri S, El Adib S, Charkaoui A. Enhancing cryptosystem security with a convolutional neural network-based countermeasure against power analysis attacks. Arabian J. Sci. Eng. 2025. <https://doi.org/10.1007/s13369-025-10485-3>
5. Li L, Ou Y. A deep learning-based side channel attack model for different block ciphers. J. Comput. Sci. 2023; 72:102078. <https://doi.org/10.1016/j.jocs.2023.102078>
6. Staib MDL, Moradi A. Deep learning side-channel collision attack. IACR Trans. Cryptogr. Hardw. Embed. Syst. 2023;2023(3):422-444. <https://doi.org/10.46586/tches.v2023.i3.422-444>
7. Zhou Y, Feng D. Side-channel attacks: Ten years after its publication and the impacts on cryptographic module security testing. Cryptol. ePrint Arch. 2023. <https://eprint.iacr.org/2005/388>
8. Benadjila R, Prouff E, Strullu R, Cagli E, Dumas C. Deep learning for side-channel analysis and introduction to ASCAD database. J. Cryptogr. Eng. 2020;10(2):163-188. <https://doi.org/10.1007/s13389-019-00220-8>
9. Chen Y, Wang B, Su CK, Ng JL, Lee YC. AutoProfile: Automated profiling in deep learning-based side-channel analysis. Neural Netw. 2025;186:108009. <https://doi.org/10.1016/j.neunet.2025.108009>
10. Maghrebi H, Portigliatti T, Prouff E. Breaking cryptographic implementations using deep learning techniques. Security, Privacy, and Applied Cryptography Engineering. 2016;9763:3-26. [https://doi.org/10.1007/978-3-319-49445-6\\_1](https://doi.org/10.1007/978-3-319-49445-6_1)
11. Picek S, Heuser A, Jovic A, Bhasin S, Regazzoni F. Still Making Noise: Improving Deep-Learning-Based Side-Channel Analysis. IEEE Des. Test. 2025;42(1):45-58. <https://doi.org/10.1109/MDAT.2024.3510421>
12. Kim J, Picek S, Heuser A, Bhasin S, Hanjalic A. Make some noise. Unleashing the power of convolutional neural networks for profiled side-channel analysis. IACR Trans. Cryptogr. Hardw. Embed. Syst. 2019;2019(3):148-179. <https://doi.org/10.13154/tches.v2019.i3.148-179>
13. Zaid G, Bossuet L, Habrard A, Venelli A. Methodology for efficient CNN architectures in profiling attacks. IACR Trans. Cryptogr. Hardw. Embed. Syst.. 2020;2020(1):1-36. <https://doi.org/10.13154/tches.v2020.i1.1-36>
14. Ünal S, Başkaya F. An energy-efficient parallel ASIC implementation of advanced encryption standard (AES) algorithm robust against side-channel attacks. Fen-Müh. Derg. 2025;27(79):152-159. <https://doi.org/10.21205/deufmd.2025277920>

15. Alabdulwahab S, Cheong M, Seo A, Kim H, Yi O. Enhancing deep learning-based side-channel analysis using feature engineering in a fully simulated IoT system. *Expert Syst. Appl.* 2024; 262:126079. <https://doi.org/10.1016/j.eswa.2024.126079>
16. Tran TH, Dam DT, Dang TK, Hoang TT, Pham CK. Countering side-channel attacks with a dynamic S-box based on affine transformations and gold sequences. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* 2025;33(3):849-862. <https://doi.org/10.1109/tvlsi.2025.3595897>
17. Zhang Y, He PF, Gan H, Jiang XY, Wang Q. Side-channel power analysis based on SA-SVM. *Appl. Sci.* 2023;13(9):5671. <https://doi.org/10.3390/app13095671>.
18. Xia S, Li L, Ou Y, Xiang J. Optimizing label correlation in deep learning-based side-channel analysis. *Microelectronics J.* 2025;159:106721. <https://doi.org/10.1016/j.mejo.2025.106721>
19. Cassiers G, Masure L, Momin C, Moos T, Moradi A, Standaert FX. Prime-field masking in hardware and its soundness against low-noise SCA attacks. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* 2023;2023(2):482-518. <https://doi.org/10.46586/tches.v2023.i2.482-518>
20. Tyler JH, Fadul MKM, Reising DR. Transformers--messages in disguise. *Mach. Learn. Appl.* 2025; 19:100669. <https://doi.org/10.1016/j.mlwa.2025.100669>
21. Ding AA, Chen L, Fei Y. A statistical model for higher order DPA on masked devices. *Cryptogr. Hardw. Embed. Syst.* 2014;8731:147-169. [https://doi.org/10.1007/978-3-662-44709-3\\_9](https://doi.org/10.1007/978-3-662-44709-3_9)
22. Ng JS, Chen J, Kyaw NA, Chong KS, Chang JS, Gwee BH. Securing against side-channel attacks with wide-range in situ random voltage dithering on async-logic AES engine. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* 2024;32(9):1674-1687. <https://doi.org/10.1109/tvlsi.2024.3409650>
23. Srivastava A, Tan S, Tehranipoor M, Forte D. SCAR: Power Side-Channel Analysis at RTL Level. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* 2024;32(6):1110-1123. <https://doi.org/10.1109/TVLSI.2024.3390601>
24. Heydari V, Nyarko K. Enhancing adversarial robustness in network intrusion detection: a novel adversarially trained neural network approach. *Electronics.* 2025;14(16):3249. <https://doi.org/10.3390/electronics14163249>
25. Hajra S, Alam M, Saha S, Mukhopadhyay D, Picek S. On the instability of softmax attention-based deep learning models in side-channel analysis. *IEEE Trans. Inf. Forensics Secur.* 2024;19:1424-1439 <https://doi.org/10.1109/tifs.2023.3326667>
26. Ebina K, Ueno R, Homma N. Side-channel analysis against SecOC-compliant AES-CMAC. *IEEE Trans. Circuits Syst. II, Exp. Briefs.* 2023;70(11):4013-4017. <https://doi.org/10.1109/TCSII.2023.3288278>
27. Zoni D, Galimberti A, Galli D. An FPGA-Based Open-Source Hardware-Software Framework for Side-Channel Security Research. *IEEE Trans. Comput.* 2025;74(6):2087-2100. <https://doi.org/10.1109/TC.2025.3551936>
28. Stypiński M, Niemiec M. Security of neural network-based key agreement protocol for smart grids. *Energies.* 2023; 16(10): 3997. <https://doi.org/10.3390/en16103997>