

Swarm in EM Hay: Particle Swarm-guided Probe Placement for EM SCA

Dev Mehta*, Seyedmohammad Nouraniboosjin*, Maryam S. Safa, Shahin Tajik, Fatemeh Ganji
Worcester Polytechnic Institute
Email: {dmmehta2, snouraniboosjin, msafa, stajik, fganji}@wpi.edu

Abstract

Despite decades of research in electromagnetic (EM) side-channel analysis (SCA), practical attacks still require manual effort and domain expertise to identify informative probe locations on target devices. Existing approaches rely heavily on exhaustive grid scanning or handcrafted alignment, limiting attack scalability and realism. In this work, we present the first automated and adaptive EM SCA framework that uses particle swarm optimization (PSO) to navigate the probe. Particles are guided by mutual information (MI) leakage maps, enabling efficient recovery of secret-dependent emissions. We introduce a novel application of the Nyström approximation to accelerate MI estimation across EM trace windows, allowing real-time swarm guidance without full kernel computations. Unlike prior work, our method requires no leakage templates, manual tuning, or alignment assistance, facilitating automated attacks with minimal assumptions. We validate our framework on both microcontroller and FPGA platforms running AES-128. PSO-guided scanning identifies high-leakage points faster than grid search and reduces the number of traces required for successful key recovery by up to a factor of 16, i.e., saving tens of thousands of traces; see artifacts.

1 Introduction

Side-channel analysis (SCA) has significantly advanced the security evaluation of real-world systems, including YubiKey 5Series [70], KeeLoq keyless entry systems [26, 45], Xilinx FPGA bit-stream encryption [60], or Atmel CryptoMemory authentication [7]. Such devices are often electronically and/or architecturally simple, at least in terms of their cryptographic functionality. Besides such traditional targets, e.g., contact-based/contactless payment cards, the Internet of Things (IoT) expands, where ensuring the security of edge nodes has become increasingly important. Given that many of these nodes rely on simple micro-controllers, they are particularly vulnerable to powerful side-channel attacks [40, 41]. In addition to targets embedded in micro-controllers, increasingly complex devices with advanced functionality are now integrated in systems whose vulnerability to side-channel attacks should be assessed. For instance, modern smartphones typically integrate multi-core System-on-Chip (SoC) components running at multi-gigahertz clock speeds; however, such design complexity cannot rule out the fact that they might be vulnerable to SCA [30, 31, 53, 80]. In this context, exploiting EM side-channel leakage is particularly appealing.

Introduced and contextualized in [27, 68, 71], EM SCA has proven itself effective. Compared to its close counterpart, power side-channel leakage, EM leakage benefits from a contactless and less invasive nature to avoid complications such as on-chip voltage regulation and, crucially, enables targeting specific regions or components of the device-under-test (DUT). Unlike currents in wires, which can be represented as scalar values, electric and magnetic fields are vector quantities in three-dimensional space. As such, electromagnetic emanation offers a richer source of information than current flows, with specific properties that can be selectively exploited for attacks [34]. This, of course, comes at a higher cost than power side-channel leakage. Selecting the appropriate setup is part of this cost. Electromagnetic fields or emanation can be measured using either custom-built or commercially available magnetic coils and probes. The literature has investigated several aspects of such measurements, e.g., coil diameters, probe-to-die distances, bandwidths, and spatial measurement resolutions [4, 5, 74, 80]. Another factor contributing to the cost of EM SCA is selecting the probe location in the system to capture the EM measurements.

State-of-the-art of EM trace acquisition. Unlike power-based attacks, EM leakage depends strongly on probe placement, and the difference between a good and poor location can alter the number of required traces by more than an order of magnitude. Because EM amplitude is attenuated by several physical effects, attackers must often rely on exhaustive spatial scanning or arbitrary probe placement, both of which substantially increase acquisition effort. Prior work has, therefore, explored ways to improve the efficiency or quality of the acquisition process itself. For instance, [38] studies acquisition parameters such as probe distance, angle,

*These authors contributed equally to this work.

and bandwidth to obtain higher-quality traces with fewer repetitions. [74] further quantifies how probe-to-die distance and bandwidth influence SNR, showing degradation proportional to $1/r^2$ in amplitude and $1/r$ in correlation, with bandwidths below 300MHz reducing signal quality significantly. As systems grow in size and complexity, locating informative EM hotspots becomes increasingly impractical without a principled method for guiding the probe toward high-leakage regions. Several works [35, 36, 72] emphasize that EM leakage is highly location-dependent and that effective attacks require precise probe placement, which is typically obtained through exhaustive spatial scanning. They further show that spatial leakage patterns form coherent regions tied to device layout, supporting the use of statistical models for capturing such structure. However, these approaches do not seek to reduce acquisition cost or automate the search for informative positions. The same observation applies to [39, 74, 78], which confirm that EM leakage is spatially structured and that attack success depends critically on probe location, yet still rely on exhaustive or heuristically adaptive scans.

More recent efforts aim to reduce the number of measurements required to identify informative probe locations. [40, 41] apply multistage analysis of variance to pre-characterize and identify high-leakage regions, thereby lowering the number of traces required across the grid. [64] proposes a single-trace-per-location EM heatmap to quickly highlight high-SNR regions and restrict subsequent acquisition to a small subset of promising positions. Similarly, [20] employs simulation-driven leakage estimation to predict informative probe locations before physical measurement, reducing the need for real traces. Complementary to these approaches, [69] trains a CNN to infer probe position from EM traces and enable accurate automatic repositioning, while [65] simplifies this process using image-processing and navigation-inspired techniques for increased speed and interpretability. [6] introduces ProbeShooter, a two-mode probe-aiming technique combining passive AM sideband detection and active intermodulation analysis to accelerate trace acquisition under both controllable and uncontrollable conditions. [18] proposes a low-cost platform that applies a greedy gradient-based search to locate high-leakage points in $O(N)$ time on an $N \times N$ grid, substantially reducing the cost relative to exhaustive methods. This platform integrates a precise automated stage, a high-resolution probe, and open-source control software, demonstrating efficient key recovery on AES implementations. Yet, while such methods reduce the scanning burden, they often remain vulnerable to local optima [18] or depend on restrictive assumptions [39].

Particle Swarm Optimization (PSO) has emerged as a compelling alternative by offering global search with improved robustness. Unlike greedy or local heuristics, PSO models each probe position as a particle guided by both personal experience and the swarm's global knowledge, enabling exploration of noisy leakage landscapes and effective convergence to high-leakage points.

Multivariate Gaussian leakage. Leakage can often be reasonably modeled as a multivariate Gaussian distribution [8, 50, 66, 67, 77]. Each captured EM trace is the sum of contributions from switching gates, supply-rail fluctuations, jitter, coupling from neighboring circuitry, thermal noise, probe coupling variations, and environmental effects [14]. By the Central Limit Theorem, the aggregate effect of many weak, independent sources tends toward a Gaussian distribution, even if the individual contributions are not Gaussian themselves. This makes the Gaussian assumption both physically meaningful and empirically accurate for EM side-channel modeling. Additionally, EM measurements form correlated features across time samples, giving rise to a multivariate structure whose covariance matrix reflects the temporal and spatial correlation of the underlying hardware activity [37, 54].

The spatial distribution of leakage often exhibits smooth variations and localized regions of similarity, meaning that neighboring positions do not necessarily provide independent information. This smoothness is a result of EM emissions governed by physical coupling mechanisms that decay gradually with distance and are shaped by the local geometry of current-carrying structures [24, 51, 75, 81]. Adjacent probe locations intercept overlapping portions of the same radiating sources, and the dominant field components vary continuously across the surface rather than changing abruptly from one point to the next [59]. As a result, measurements taken at nearby positions capture largely similar combinations of circuit activity, producing spatial correlation in the observed leakage [32, 76]. This inherent continuity explains why probing more points beyond a certain resolution does not provide much new information. This raises the following problem: *can one identify informative EM leakage regions accurately enough to guide probe placement without exhaustively measuring every grid point?*

Method at a glance and contributions. Our framework proceeds in three stages. First, we collect a small number of traces across a coarse spatial grid and use a probabilistic model (Gaussian Mixture Model) to identify promising regions for probe placement. Second, we estimate leakage over the grid using mutual information (MI), computed via matrix-based Rényi entropy and accelerated with Nyström sampling so that MI can be evaluated efficiently during the search. Third, we run PSO on this leakage landscape, using the GMM-defined points as initial particle locations and MI as the fitness objective, to navigate toward high-leakage regions and then validate those locations with correlation power analysis (CPA). Our contribution goes beyond the combination of these techniques to propose an automated EM probe-localization workflow. Below, we summarize the key technical contributions of this work.

• A mutual-information estimator based on matrix-based Rényi entropy that naturally extends to multivariate leakage. We employ matrix-based Rényi entropy to compute mutual in-

formation directly from high-dimensional EM traces without assuming explicit probability densities. This approach seamlessly handles multivariate leakage samples and provides a robust, nonparametric measure of information content across probe positions. *For the first time in the SCA domain,* we reduce the cost of MI computation by applying Nyström sampling to the matrix-based Rényi entropy estimator, thereby avoiding full kernel-matrix construction over all traces and replacing the quadratic cost of pairwise kernel evaluation with a low-rank approximation. By approximating the eigenspectrum using a small, carefully selected subset of samples, the method preserves the accuracy of the MI estimate while reducing computational and memory requirements from quadratic to near-linear in the number of traces, making multivariate MI evaluation practical within the PSO search loop.

• A measurement-driven Gaussian Mixture Model (GMM) [57] for spatial leakage modeling and acquisition cost reduction. We introduce a GMM-based representation of the spatial distribution of EM leakage, obtained from extremely low-dimensional PCA features. This model captures the multivariate structure of the leakage landscape and enables the selection of informative probe positions without requiring exhaustive measurements, thereby reducing the overall trace-collection cost.

• A particle swarm optimization (PSO) [46] framework that navigates the MI landscape to localize high-leakage regions. We develop a PSO-based search strategy that uses MI values as the fitness objective, enabling adaptive exploration of the spatial grid. By combining model-guided initialization with MI-driven navigation, the method efficiently converges to regions exhibiting strong side-channel leakage. We evaluate the method on both microcontroller and FPGA platforms, outperforming standard Test Vector Leakage Assessment (TVLA) and grid-based strategies. This leads to orders-of-magnitude attack cost reduction in terms of the number of traces needed (e.g., from more than 30,000 traces to 7,000 traces).

Our method enables practical black-box EM attacks by automating probe placement, significantly reducing the time, effort, and expertise required for key recovery. As attack workflows become more autonomous, certification bodies could incorporate such techniques into scalable evaluation pipelines.

2 Background

2.1 Notations

We use calligraphic letters such as $\mathcal{X}$ to denote sets. Random variables are written in uppercase letters (e.g., X), and random vectors in bold uppercase (e.g., $\mathbf{X}$). Realizations are denoted in lowercase (e.g., x for scalars, $\mathbf{x}$ for vectors), and matrices in bold capital letters (e.g., $\mathbf{K}$). For a vector $\mathbf{x}$, its t-th component is denoted by $\mathbf{x}[t]$. The expected value is denoted $\mathbb{E}[\cdot]$, and subscripts are added for clarity, e.g., $\mathbb{E}_X[\cdot]$.

Entropy is denoted by H, and mutual information by I. For discrete random variables, entropy is given by $H(X) = -\sum_x p(x)\log p(x)$, and mutual information by $I(X;Y) = H(X)+H(Y)-H(X,Y)$.

2.2 Particle Swarm Optimization

Particle Swarm Optimization (PSO) [46] is a population-based metaheuristic designed for efficient navigation of high-dimensional search spaces. In our context, each particle represents a candidate probe position in the spatial leakage landscape, and its objective is to maximize a leakage-based fitness score, such as mutual information (MI). Unlike exhaustive grid scans, PSO adaptively explores the space by updating particles based on local and global knowledge, reducing the number of measurements required to locate high-leakage regions.

Each particle maintains a current position $\mathbf{x}_i$, a velocity vector $\mathbf{v}_i$, and its best position $\hat{\mathbf{x}}_i$. The swarm also tracks the best position $\hat{\mathbf{x}}$ observed across all particles. At iteration t, particle i updates its velocity and position according to:

$$\begin{aligned}\mathbf{v}_i(t+1) &= \mathbf{v}_i(t) + \phi_1 r_1 (\hat{\mathbf{x}}_i - \mathbf{x}_i(t)) + \phi_2 r_2 (\hat{\mathbf{x}} - \mathbf{x}_i(t)),\\ \mathbf{x}_i(t+1) &= \mathbf{x}_i(t) + \mathbf{v}_i(t+1),\end{aligned} \tag{1}$$

where ϕ_1 and ϕ_2 are acceleration coefficients and $r_1, r_2 \sim \mathcal{U}(0,1)$ are sampled independently per iteration. The algorithm balances exploration and exploitation: early iterations favor distributed search, while convergence is driven by the swarm's collective experience. PSO may fail to converge when the objective is noisy, the exploration-exploitation trade-off is poorly balanced, the initialization is weak, or the update rules mismatch the search space.

For metaheuristic algorithms like PSO, the initial state of swarm particles has a very strong influence on the final outcome [52]. In many swarm-based optimization algorithms, particle positions are initialized using either random sampling or Latin Hypercube Sampling (LHS) [55]. Random initialization places particles uniformly at random over the search space, which is simple and unbiased but can lead to clustering, uneven coverage, and poor exploration in high-dimensional or structured domains. LHS improves upon this by stratifying each dimension of the search space into equal intervals and ensuring that each interval is sampled exactly once. As a result, LHS provides more uniform coverage of the domain with the same number of particles, reducing the risk of large unexplored regions [17]. Both approaches are agnostic to the underlying objective function and do not incorporate prior knowledge about where promising regions may lie. While this generality makes them widely applicable, it also means that many particles may be initialized in low-quality or irrelevant regions, potentially increasing convergence time and computational cost.

Compared with alternative population-based optimizers such as genetic algorithms, PSO is particularly attractive here because it updates probe locations directly through continuous position and velocity dynamics, requiring little parameterization while remaining effective for expensive black-box search.

2.3 Side-channel Leakage

SCA uses physical observations $\mathbf{L}$ collected during the execution of a cryptographic algorithm on hardware or software to infer secret information [11]. In a non-profiling attack such as correlation power analysis (CPA), the adversary does not require access to a clone of the target device [11]. Instead, the adversary collects a set of leakage traces and associated plaintexts $\mathcal{D}_{\text{attack}} \doteq (\ell i, p_i) 1 \leq i \leq N_a$, where k^* denotes the fixed but unknown key. Each trace ℓ_i is a sequence of sampled physical measurements—e.g., power or EM emissions—recorded during a single cryptographic execution. We denote each leakage trace ℓ as a vector $x = [x_1, x_2, \ldots, x_T] \in \mathcal{X}^T$, where each entry x_i corresponds to a discrete-time sample and $\mathcal{X}$ denotes the space of possible sample values. A leakage model $h(p,k)$ is computed for each key hypothesis k using a power model such as Hamming weight (HW) or Hamming distance (HD), and the correlation coefficient $\rho(x, h(p,k))$ is evaluated to determine the key hypothesis that maximizes alignment with the measured traces.

The distribution of leakage traces, denoted $g_{\mathbf{L}}$, can often be modeled as a finite mixture of multivariate Gaussians [66, 67]. Let g_θ denote such a Gaussian mixture model (GMM), where for each trace $x \in \mathcal{X}^T$, $g_\theta(x) = \sum_{t=1}^{T} a_t \cdot g_{\mu_t, \Sigma_t}(x)$. Here, $\theta = (a_t, \mu_t, \Sigma_t)$ with $(1 \leq t \leq T)$ includes the mixing weights a_t (summing to 1), the component means μ_t, and covariance matrices Σ_t. This generative model captures the structure and variability in leakage traces across multiple executions and probe positions, enabling downstream information-theoretic and clustering analyses in high-dimensional spaces.

3 EM-SCA

EM SCA leverages radiation emitted by a chip during computation to extract sensitive information such as cryptographic keys. The EM radiation primarily originates from switching activity in on-chip metal layers, which carry transient signals and generate localized fields. To reduce environmental noise and improve spatial resolution, it is common to use a near-field EM probe placed close to the chip surface [36]. By systematically scanning the probe across the device, an adversary can capture spatially resolved leakage that reflects activity from specific internal components [2, 18, 72]. However, exhaustive grid-based scans are time-consuming and typically require hundreds or thousands of measurements to locate regions of interest. This limits the practicality of EM localization in real-world attack scenarios. Efficient navigation strategies that accelerate the search for high-leakage regions, without relying on golden-device fingerprinting or invasive techniques, are therefore essential.

EM probes. The choice of EM probe directly impacts the spatial resolution and signal quality in near-field EM side-channel analysis [21]. Probes vary in terms of measured field component—magnetic (H- or B-field) or electric (E-field)—as well as sensitivity, bandwidth, and footprint. Magnetic probes are commonly used to capture high-frequency emissions from switching currents in integrated circuits (ICs) and printed circuit boards (PCBs). They are typically divided into two types: B-field probes, which measure magnetic flux density via voltage induction, and H-field probes, which capture magnetic field intensity through direct coupling.

B-field probes often use multi-turn coils and are sensitive to the time derivative of the field, making them effective for detecting short, high-frequency transients. However, they are prone to integration artifacts and require precise orientation and amplifier matching. In contrast, H-field probes consist of shielded single-turn loops and produce an output voltage linearly proportional to the local field strength. They offer flatter frequency response, reduced environmental sensitivity, and are better suited for broadband EM side-channel analysis.

In practice, H-field probes are generally preferred for near-field measurements due to their robustness to ambient interference and compatibility with non-contact scanning. Their linear response and spatial precision make them particularly effective for detecting fine-grained leakage patterns from localized switching activity. This spatial dependence, however, creates a large search space for attackers. Manually locating high-leakage regions through exhaustive grid scanning is time-consuming, which motivates the need for automated probe localization strategies. In this work, we adopt H-field probes in combination with mutual information–guided particle swarm optimization to accelerate leakage discovery.

Electric field probes. Electric field (E-field) probes detect near-field voltage gradients caused by high dv/dt transitions through capacitive coupling and are typically placed in direct contact with the device under test. While they can capture useful side-channel signals, they are more susceptible to ambient noise and surface interference, and their contact-based operation limits spatial maneuverability. As a result, they are less suited than H-field probes for high-resolution, spatially resolved leakage scans. In this work, we therefore focus exclusively on H-field probes because of their non-contact operation, spatial precision, and compatibility with automated localization strategies.

Effect of probe diameter. The diameter of a magnetic probe directly influences its spatial resolution and sensitivity. Larger-diameter probes integrate emissions over a broader area, making them suitable for capturing switching activity from large circuit regions. In contrast, smaller probes provide finer spatial resolution, allowing the attacker to isolate emissions from localized sources, including passive components like resistors or capacitors. Although these components do not switch actively, they can emit unintentionally due to transient current flow or capacitive/inductive coupling from nearby activity. To accurately localize such emissions, magnetic probes with diameters in the range of 1–3mm are commonly used. Smaller probes also tend to support higher bandwidths, which is advantageous for capturing short-duration switching events and high-frequency leakage [22, 73].

4 Metrics in SCA

Side-channel evaluation often relies on multiple metrics to assess both the presence and exploitability of leakage across different stages of an attack. For EM leakage at each probe position, we employ three metrics: TVLA, CPA, and MI. While MI serves as a general-purpose indicator of statistical dependence, TVLA and CPA provide concrete metrics aligned with the capabilities and goals of a real adversary.

TVLA [9] is a standard leakage detection technique based on the Welch t-test. It determines whether two sets of side-channel traces, collected under different inputs but identical conditions, exhibit statistically significant differences. This non-specific test requires no knowledge of the intermediate values or keys and is widely used for leakage certification.

In our experiments, we collect two trace sets: one with fixed plaintext inputs and one with uniformly random inputs. Let μ_1, μ_2 denote the sample means and s_1, s_2 the standard deviations of the two populations, with n_1, n_2 traces each. The t-score is computed as:

$$t = \frac{\mu_1 - \mu_2}{\sqrt{(s_1^2/n_1) + (s_2^2/n_2)}} \tag{2}$$

A t-value exceeding ± 4.5 allows rejection of the null hypothesis with 99.999% confidence, indicating the presence of data-dependent leakage.

CPA [15,47] is used as an exploitation technique that attempts to recover the secret key. Given a leakage trace ℓ_i and a predicted leakage model $h(p_i, k)$ (typically based on Hamming weight or Hamming distance), CPA computes the Pearson correlation coefficient $\rho(k) = \text{corr}(\ell, h(p,k))$ across all key hypotheses k. The key hypothesis with the highest correlation is selected as the most likely candidate.

In our experiments, we use both Hamming weight (HW) and Hamming distance (HD) models applied to various rounds of AES-128, depending on the setup. While TVLA highlights positions with generic leakage, CPA quantifies how much of this leakage is exploitable for key recovery, providing a complementary perspective.

4.1 Mutual information

In this section, we explore MI as a fundamental metric for evaluating the information leakage in SCA. Estimating MI from observed side-channel traces allows for a model-independent assessment of how much information about the secret key is revealed by a given leakage model. We focus on two state-of-the-art MI estimation methods that are particularly relevant in the context of SCA: the Kraskov estimator (KSG) [48], which is based on k-nearest neighbors, and the Matrix-Based Rényi Entropy (MRE) estimation [1, 29, 61, 85], a kernel-based approach that leverages the spectral properties of the data. In the following, we provide the mathematical definitions of both estimators, and in the subsequent experiments, we compare their performance in terms of accuracy, noise robustness, and sample efficiency, particularly on the ASCAD dataset [12] and synthetically generated data.

Matrix-Based Rényi Entropy. When Rényi entropy is used as a generalized alternative to Shannon entropy, information-theoretic quantities can be evaluated by applying appropriate matrix functionals to normalized positive definite matrices [29, 84]. In this framework, the entropy is computed from a normalized Gram matrix $\mathbf{A}$ obtained by evaluating a kernel over all sample pairs. A kernel $\kappa : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ assigns a real-valued similarity score to two inputs drawn from the space $\mathcal{X}$. Such functions are symmetric, meaning $\kappa(x, x') = \kappa(x', x)$, and are positive definite, ensuring that for any finite set $\{x_1, \ldots, x_n\}$ the Gram matrix K satisfies $c^\top K c \geq 0$ for all $c \in \mathbb{R}^n$. A kernel is said to be infinitely divisible if $\kappa_r(x,y) = \kappa(x,y)^r$ remains positive definite for every $r > 0$, a property that guarantees closure under Hadamard powers and products [13, 84]. This property is essential when constructing multivariate entropy expressions from multiple kernel matrices. Given a Gram matrix $\mathbf{K}$, the normalized version is defined by $\mathbf{A} = \mathbf{K}/\text{tr}(\mathbf{K})$, where the trace is the sum of the diagonal entries. The eigenvalues of $\mathbf{A}$ are then inserted into a functional $\mathbf{S}_\alpha(\mathbf{A})$ to compute matrix-based Rényi entropy without performing any density estimation. Choosing $\alpha = 1.01$ produces an accurate approximation of Shannon entropy, following the analysis in [29, 84, 85]. Because the kernel implicitly defines an RKHS, the resulting matrices are positive semi-definite, ensuring non-negative eigenvalues and well-behaved entropy computations. This formulation enables information-theoretic measures to be estimated directly from samples without assuming any parametric probability model.

The joint Rényi entropy of k variables is defined as

$$J_\alpha(X_1, \cdots, X_k) = S_\alpha\left(\frac{A_1 \odot \cdots \odot A_k}{\text{tr}(A_1 \odot \cdots \odot A_k)}\right), \tag{4}$$

where each matrix $(A_z)_{ij} = \kappa_z(x_z^i, x_z^j)$ for $1 \leq z \leq k$, and $\odot$ denotes the Hadamard product. The functional used to compute Rényi entropy is

$$S_\alpha(A) = \frac{1}{1-\alpha}\log_2(\text{tr}(A^\alpha)) = \frac{1}{1-\alpha}\log_2\left(\sum_{i=1}^{n}\lambda_i(A^\alpha)\right), \tag{5}$$

where $\lambda_i(A)$ is the ith eigenvalue. The normalized matrix can equivalently be expressed as $A_{ij} = K_{ij}/\left(n\sqrt{K_{ii}K_{jj}}\right)$, as used in the original matrix-based Rényi entropy formulation [28, 29]. Setting $\alpha = 1.01$ throughout these computations yields a practical and smooth approximation to Shannon entropy [28, 29, 84].

KSG estimation of MI. Unlike correlation, which assumes linearity and a known leakage model, MI captures arbitrary statistical dependence between the measured leakage ℓ and the known plaintext p. To compute MI from empirical trace data without requiring density estimation, we adopt the Kraskov–Stögbauer–Grassberger (KSG) estimator [48], a

non-parametric method based on k-nearest neighbor statistics. KSG estimates $I(X;Y)$ directly from joint samples $\{(x_i,y_i)\}_{i=1}^n$ by counting neighbors in the marginal spaces around each sample, adapting naturally to variations in local density. This makes it robust for high-dimensional, non-Gaussian, and non-uniform leakage — all common in real EM settings. We use KSG with $k=3$ throughout, which balances bias and variance well in our experiments. Because the estimator is fast, model-free, and does not require alignment or template profiles, it is ideal for fitness evaluation in PSO-based leakage localization.

MRE vs. KSG We compare the effectiveness of KSG and matrix-based Rényi entropy (MRE) estimation in SCA contexts where fast and accurate MI evaluation is needed to guide the EM probe. As a sanity check, we first apply both estimators to leakage traces from the ASCAD dataset [12], using histogram-based MI as a reference. The ASCAD fixed dataset contains side-channel traces from an 8-bit AVR μC executing a masked AES-128 implementation. Each trace captures 700 features from electromagnetic emanations aligned to a fixed time window, with all traces collected under the same encryption key [10]. While KSG is non-parametric and easy to deploy, it shows higher variance and underestimation in sparse or noisy regions. In contrast, MRE more closely matches the histogram baseline, as shown in Figure 1, indicating superior precision in leakage quantification. To test noise robustness, we construct a synthetic dataset of 700 Gaussian samples where the ground-truth MI is analytically known as $I(X;Y) = -\frac{1}{2}\log(1-\rho^2)$, with ρ the correlation coefficient. We then add increasing Gaussian noise and evaluate both estimators. As shown in Figure 2(a), MRE remains stable and tracks the true MI across noise levels, whereas KSG degrades sharply and outputs invalid negative values under high noise.

Nyström sampling The Nyström method enables fast low-rank approximation of matrices, making it suitable for large-scale or time-critical learning tasks where full matrix evaluation is infeasible [86, 87]. A key aspect of the method is how the columns are selected from the original matrix. We also evaluate computationally efficient MRE through Nyström-based kernel approximation, which accelerates MI estimation by subsampling the trace matrix. Figure 2(b) shows that reducing the number of reference points to 100 incurs minimal accuracy loss, enabling fast fitness evaluation during particle updates without full-matrix cost. In summary, MRE offers better robustness, accuracy, and scalability, making it a more practical choice for guiding probe navigation in noisy, high-dimensional EM SCA scenarios.

5 Adaptive and Model-guided EM SCA

A critical step in side-channel attacks that leverage local EM measurements is selecting an appropriate probe position to collect traces. A naive yet common strategy is to perform a grid search over the chip surface to identify the location that yields the strongest leakage signal. Several works [18, 35, 36, 39, 41, 74, 78] have sought to improve upon this brute-force approach by introducing heuristic methods that prioritize locations with higher observed leakage. However, EM traces collected over a spatial grid are often degraded by significant noise, making scan-based techniques inefficient and susceptible to false positives—particularly in regions far from the cryptographic core (e.g., AES). To address this, our framework introduces an automated method for selecting spatial sampling points, i.e., EM probe positions, by exploiting the latent structure embedded in the collected trace data. Rather than relying on exhaustive scanning or uniform random initialization, the goal is to model the spatial distribution of informative leakage and use it to guide sampling. This framework operates in two key stages: (1) identifying an initial probe placement using unsupervised modeling of trace variability, and (2) guiding subsequent probe movements to minimize the number of traces required for a successful attack. We describe these steps in detail below.

5.1 Initializing Probe Positions

EM traces collected across the chip surface capture variations in signal intensity and shape as a function of spatial proximity to active circuit regions. These spatial differences manifest as distinct clusters when the traces are projected into a lower-dimensional feature space. Despite the high dimensionality of the raw traces, the underlying activity is often governed by correlated switching patterns, implying that the data lie on a low-dimensional manifold. Repeated hardware operations, such as clock-driven transitions and structured computation (e.g., AES rounds), further induce statistical regularities across samples. The noise in EM measurements is typically additive and approximately Gaussian, making Gaussian mixture modeling well-suited to representing both the variability within each cluster and the overall distribution of signal behavior. As a result, GMMs provide a flexible and statistically grounded framework for modeling the latent structure in spatially distributed EM data cf. [66, 67]. To build the GMM-based initialization pipeline, we proceed in four steps, summarized below and then explained in detail: **Step 1.** collect 30 traces at each grid point; **Step 2.** extract local PCA features from the traces; **Step 3.** fit candidate GMMs and select the number of components using AIC together with the elbow criterion; **Step 4.** remove low-signal or outlier-dominated regions; and sample valid probe locations from the resulting model to initialize the particles.

Step 1. Initial trace collection. At each location on a grid covering the chip (device under test), 30 traces are collected in order to obtain a statistically stable estimate of the local leakage distribution, as this sample size is commonly regarded as sufficient for Gaussian approximations to become reliable under the central limit effect [63, 79]. EM traces are typically high-dimensional, but much of the variance is concentrated in a few principal components due to correlated switching behavior across time samples. By projecting the data onto

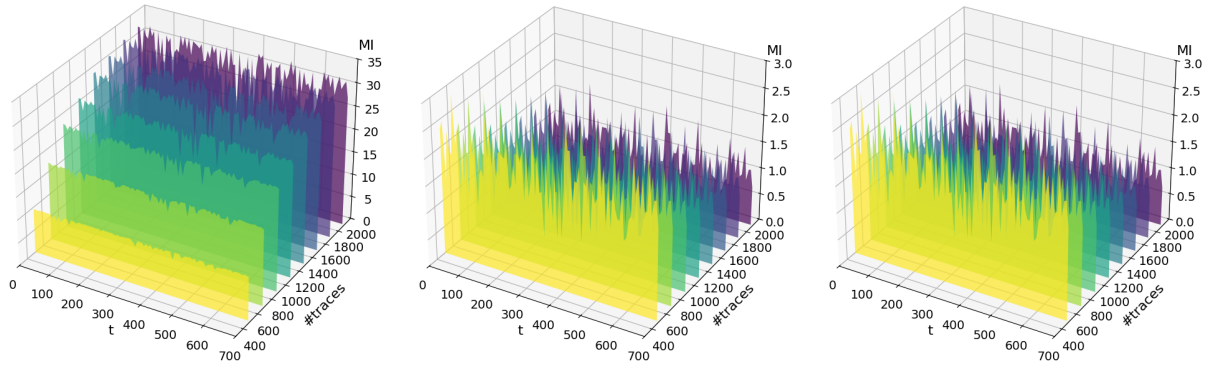

(a) KSG (b) Histogram (c) MRE

Figure 1: MI estimation on the ASCAD dataset with fixed key [10]. Each of the 50,000 traces has 700 samples. MI is computed between each sample point and the SBox output.

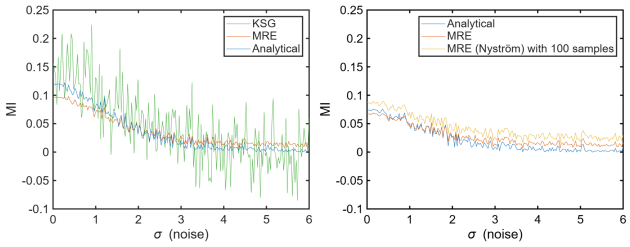

(a) (b)

Figure 2: (a) MI estimation for 700 synthetic Gaussian samples with noise variance matching ASCAD [10], illustrating real-world noise levels. (b) Impact of Nyström sampling on MRE with 100 samples.

a low-dimensional subspace, principal component analysis (PCA) [42] emphasizes structure related to spatial differences in leakage. Moreover, GMMs [56] are known to perform better in lower-dimensional spaces where clusters are more clearly separated, and covariance estimation is more reliable. Hence, applying PCA before fitting a GMM improves the robustness and stability of GMM fitting by avoiding overfitting to noisy or redundant dimensions. Local PCA at each measurement location makes sense in EM SCA because traces collected from the same position tend to exhibit consistent signal patterns driven by underlying hardware activity. By performing PCA locally, one can extract the dominant temporal variations that are characteristic of that specific probe position, effectively summarizing the most informative leakage behaviors observed at that location. This approach isolates spatially localized signal structures and filters out irrelevant or noisy components that may be present across all positions. Since each location captures a different perspective of the device's internal transitions, local PCA highlights these differences by compressing trace variability into a compact and interpretable representation. The resulting principal components serve as low-dimensional descriptors of spatial leakage characteristics, making them well-suited for modeling downstream.

Step 2. Extracting local PCA features. In EM SCA, the number of distinguishable leakage regions is not known a priori and may vary depending on circuit structure, probe sensitivity, and noise. In order to find a balance model fit against complexity, we fit GMMs with varying numbers of components and select the model with the lowest Akaike Information Criterion [3]. The AIC is a widely used model selection metric that balances the goodness of fit with model complexity. Lower AIC values indicate better models by penalizing those that overfit through excessive complexity, even if their likelihood is high. In EM SCA, the spatial structure of leakage patterns is not known in advance, and the goal is to discover meaningful clusters that correspond to regions with similar trace characteristics. To refine this further, the second derivative of the AIC curve is used to locate the point of maximum curvature, i.e., the "elbow," which serves as a robust heuristic for selecting the number of clusters before diminishing returns set in. This approach stabilizes the modeling process by avoiding underfitting (too few components) and overfitting (too many components), which is especially important when working with noisy EM data.

Step 3. Fitting GMM. In EM SCA, spatial locations dominated by noise can skew the generative model and lead to poor sampling or misinterpretation of leakage structure [49]. After fitting an initial GMM, the posterior probabilities allow identifying which traces are most likely explained by low-mean or low-weight components that do not represent a meaningful signal. By discarding samples with high posterior mass under these weak components and interpolating over the resulting gaps, the procedure effectively removes outliers and stabilizes the subsequent modeling step. This strategy ensures that the final GMM is trained only on signal-rich regions, improving the model's fidelity and the quality of the samples it draws.

Step 4. Sampling probe locations from GMM. In EM SCA, the spatial leakage distribution modeled by a GMM may include regions of low confidence or low density, especially in the tails of each component. Sampling from the

full GMM without constraints can produce unlikely or physically implausible locations that lie far from the core of the modeled leakage clusters. By rejecting samples whose Mahalanobis distance [16] exceeds 3 standard deviations from their component mean, the method restricts sampling to the high-probability mass of the distribution. This 3σ threshold aligns with standard statistical practice, covering 99.7% of the distribution under Gaussian assumptions cf. [63], and ensures that only confident, representative probe positions are retained for downstream use.

5.2 Creating MI Maps

At each spatial location drawn from the GMM, multiple EM traces, typically 1000's, can be collected to enable statistical characterization of the local leakage behavior via MI. Each trace $\ell \in \mathbb{R}^n$ represents a single observation comprising n time samples and is treated as a realization from an underlying multivariate distribution. To eliminate location-dependent DC offsets or scale differences, the traces are (optionally) preprocessed using z-score normalization. This step stabilizes downstream computations and ensures that kernel-based similarity measures are not biased by irrelevant global variations.

As mentioned earlier, MI provides a principled and model-agnostic measure of leakage by quantifying the full statistical dependence between observed traces and sensitive variables, rather than restricting the analysis to differences in means as in TVLA. MI captures all orders of statistical dependence and remains sensitive even when leakage manifests through subtle distributional changes rather than mean shifts. This makes MI particularly suitable for EM measurements, where noise, spatial averaging, and probe placement can significantly attenuate first-order effects. When coupled with Gaussian mixture modeling, MI becomes even more powerful, as GMMs provide a flexible yet structured approximation of complex leakage distributions arising from spatially heterogeneous and noisy measurements [50]. The GMM captures multimodality and variability in the leakage, while MI evaluates how much information about the sensitive variable is preserved through this distribution, yielding a robust and expressive leakage detection framework that goes beyond hypothesis testing and supports fine-grained spatial analysis.

To analyze the structure of EM leakage across spatial locations in a nonparametric, multivariate fashion, we adopt the matrix-based Rényi entropy framework [29, 84]. This approach computes entropy directly from a normalized Gram matrix $\mathbf{K}$, constructed using a positive definite kernel evaluated on all pairwise combinations of input traces. A key design choice in this framework is the selection of the kernel function and its bandwidth parameter σ, which strongly influence the resulting entropy estimate. While the Silverman rule-of-thumb is a common choice for setting the kernel bandwidth, it assumes unimodal Gaussian inputs and can severely underperform when applied to structured or high-dimensional leakage traces [33, 43]. In our setting, we observed that Silverman's formula often underestimates the optimal kernel size, leading to poor entropy resolution or over-smoothing. To address this, we explored alternative strategies, including the use of fixed percentile-based heuristics on pairwise distances or adaptive bandwidths derived from local trace geometry. These alternatives yielded more stable and meaningful entropy estimates across spatial regions, especially in the presence of multimodal or heavy-tailed trace distributions.

Moreover, when computing mutual information between leakage traces and discrete labels (e.g., intermediate values, Hamming weights), it is critical to define a kernel on the categorical label space. Unlike continuous variables, categorical labels lack a natural distance metric, making standard kernels inapplicable. To address this, we construct a label kernel $\mathbf{K}y = \mathbf{C}\mathbf{C}^\top$, where $\mathbf{C}$ is a class indicator matrix with normalized rows. This results in a symmetric, positive semi-definite matrix that reflects class membership similarity and ensures trace-normalization, as required by the entropy framework. This construction enables rigorous, kernel-based estimation of information-theoretic quantities between continuous-valued EM data and categorical side-channel labels.

Furthermore, to avoid the challenge of defining valid kernels over categorical labels—which can introduce bias or instability in kernel-based mutual information (MI) estimation—we adopt an approximation strategy instead. Let X denote the observed EM trace and Y the sensitive variable (e.g., a key-dependent intermediate). We assume a leakage model of the form $X = L(Y) + N$, where $L(Y)$ is a deterministic leakage function and N is additive noise independent of Y. To approximate the mutual information $I(X;Y)$, we estimate the entropy of the noise term $H(N)$ using traces collected at the four corner cells of the chip, which are assumed to predominantly capture background noise due to their distance from active cryptographic regions. We then subtract this estimated baseline from each cell's trace entropy $H(X)$ to obtain an approximate leakage map: $\widehat{I}(X;Y) \approx H(X) - \overline{H}_{\text{corners}}$. This approach leverages the spatial stationarity of background noise and avoids the pitfalls of kernel construction for discrete label spaces, yielding a stable and conservative estimate of spatial information leakage. In addition to sidestepping kernel-related issues, this approximation also avoids the need to estimate $H(X|Y)$, which is often unreliable in practice due to label imbalance or limited sample sizes in high-cardinality spaces. Finally, the method does not require any assumptions about the form of the leakage function, making it robust to unknown or nonlinear signal transformations—particularly relevant in real-world hardware implementations.

Nyström sampling for approximating $H(X)$. As discussed earlier, this method offers an efficient approximation strategy for kernel-based entropy estimation by reducing the computational burden of large Gram matrices. Rather than computing the full $n \times n$ kernel matrix, which scales quadratically with the number of traces per location, the Nyström method selects a smaller set of landmark samples, a representative subset of

the original data, that serve as reference points for constructing a low-rank approximation [86]. These landmarks effectively summarize the structure of the full dataset, allowing the method to interpolate relationships between all samples using only their similarities to the selected subset. This is particularly useful in SCA, where thousands of EM traces may be collected per spatial location, making full kernel evaluations impractical. To ensure meaningful approximation, the number of landmarks is selected based on the effective rank of the kernel matrix, which captures how much of the variance is explained by leading components. Moreover, repeated trials with random landmark selections allow quantifying the variance of entropy estimates, and the elbow point in the entropy curve (as the number of landmarks increases) can guide stopping criteria. This approach maintains fidelity to the full entropy landscape while dramatically reducing runtime and memory requirements, enabling scalable $H(X)$ estimation across spatial grids.

5.3 Navigating Probes on MI Maps

Navigating particles on the MI map makes intuitive and practical sense in the context of EM side-channel analysis. The MI landscape reflects the spatial distribution of information leakage across the chip surface. Each cell's MI value estimates how strongly that location's traces correlate with internal secret-dependent operations, thereby guiding the search toward regions of interest. In our setting, each particle represents a probe location on a spatial grid, and its fitness is the entropy (also compared with TVLA in Section 6) at that location. Particles are initialized from GMM-sampled points (compared with randomly and LHS initialization in Section 6), after which PSO updates their positions using the standard velocity and best-position rules until convergence or the maximum number of iterations is reached. By initializing particles using samples drawn from a GMM fitted to spatial trace features, the search avoids low-leakage regions and begins in areas already known to be structurally informative. One PSO iteration corresponds to evaluating the current particle locations by collecting the traces required for MI estimation at those positions and updating the personal and global best locations accordingly. Once initialized, particles can explore the MI surface using local gradients and swarm behavior to converge on global or local maxima—corresponding to optimal probe positions. This strategy combines statistical prior knowledge with data-driven optimization, allowing efficient navigation of a noisy, high-dimensional search space without exhaustive scanning.

Post-processing for maps. Entropy maps constructed from spatially distributed EM traces often contain both informative peaks and spurious fluctuations due to noise. Without proper post-processing, optimization algorithms like PSO may converge to misleading local maxima that do not reflect true leakage but rather stem from background variability or measurement artifacts. To enhance the robustness of probe placement, we apply spatial post-processing to the entropy maps. First, we compute a Laplacian-based curvature map to emphasize sharp local variations, which often indicate transitions between leaking and non-leaking regions. However, sharp curvature alone is not sufficient: high curvature can occur in flat or low-information regions. To mitigate this, we gate the curvature map using normalized entropy values, masking out curvature responses in areas where the entropy is below the median. This ensures that only entropy-rich regions are emphasized, and noisy border effects are suppressed. This enhanced MI map becomes the objective function for PSO. To further improve stability, we repeat PSO multiple times and score all outcomes using the fused map directly, selecting the point with the highest value rather than relying on clustering or frequency. This approach avoids overfitting to noisy results and provides a consistent way to identify optimal probe positions across varying map conditions.

6 Results

6.1 Experimental Setup

The experimental setup requires five major components: the EM probe, the XYZ stage, the control PC, the DUT, and the capture setup. The EM probe is a Riscure EM probe with a 1.25mm diameter. This probe has enough bandwidth for the frequency of our devices. The XYZ stage, which holds the probe, is also from Riscure; it uses Trinamic motors for all axes, and we use an open-source Python API (pytrinamic [23]) for control. The control PC controls all the other devices. The Chipwhisperer Python interface is used to control and configure the DUT, scope, and motor.

For the microcontroller (μC) experiments, the DUT is an ATXmega128D4-AU processor on the CW-lite board, similar to [18]. The design is an unmasked AES implementation provided by the device maker, Chipwhisperer (CW) [62]. Lastly, the capture uses a CW-husky, enabling synchronous measurement with the clock to reduce jitter and eliminate the need for trace alignment. The ADC sampling rate is 30x the clock speed of the μC, which is the maximum sampling frequency supported by CW-husky for the best resolution. For the experiments, we used the SCAPEgoat library, created explicitly for SCA [58]. It supports CW devices and allows handling individual datasets for organizing experiments.

We have also performed EM SCA on the CW305 FPGA board. It uses an Artix-7 FPGA running at 5 MHz for the attack. Again, CW-husky is used to capture traces at 40x the FPGA's clock speed. An FPGA is used for a specific purpose, namely, the different placement of the AES compared to an ASIC. We use multiple bitstreams for the same design, each with a different logic placement, using the Vivado placement GUI. Additionally, experiments on an FPGA exhibit different EM characteristics compared to those on a μC, due to differences in logic type, transistor sizes, architecture, power profile, and packaging. A μC implements logic using standard cells and a specific routing design, whereas an FPGA

uses slices, which consist of LUTs, FFs, and a switch box for routing. This variation in architecture, sizes, and power profiles across devices leads to different current consumption in a given region of the chip. Especially in FPGA, changing the design's location changes the routing, which affects EM radiation [19]. Lastly, we show all attacks as completely non-invasive by keeping the packaging intact, which adds significant noise due to the distance from the die [74]. All of these parameters lead to different signal quality across devices and implementations, as shown in the results section.

In both setups, the chip is divided into an 11x11 grid for EM analysis. For each location, 1,000 traces are collected for both random and fixed plaintext, enabling the generation of TVLA maps. From these datasets, random ones are used for MI and CPA computation. The TVLA and MI results (approximated as discussed in Section 5.2) are used to run our PSO algorithm to find optimal locations. For FPGAs, noise is higher, resulting in poor CPA performance, so we have employed averaging as a noise-reduction technique. Therefore, to calculate CPA, 10 traces are collected for each plaintext in the FPGA scenario and then averaged. We want to stress that, for TVLA and MI, averaging is not necessary, as the locations with leakage remain similar with or without it. Next, the EM SCA results on the μC will be discussed.

6.2 Microcontroller Results

The μC results use the standard AES attack model, i.e., the output of the S-box for round 0. In a μC, AES computations occur in the ALU, and sensitive values are stored in registers. The computation is serialized and uses the same ALU and register bank for each byte. As all byte processing occurs at the exact location (the ALU/register bank), once the area is identified, all bytes can be extracted from it. As we will see in the next section, this is important compared to other architectures. The goal of EM analysis is to identify leaking locations on a chip, and it is typically performed using a full-grid attack, especially when the packaging is not removed, and no information about the die placement is available to the attacker [35]. The goal of our PSO-based localization is to identify the chip location with the highest leakage faster than the grid-based approach. We use two different leakage detection methods: TVLA and MI for this purpose. Figure 3-4 shows the maps for both TVLA and MI for the μC implementation, respectively. Note that this map has been generated using an exhaustive search for representation purposes and that a real attacker does not need this map to use PSO. As seen in the maps, the leakage points are clearly visible. For demonstration purposes, we apply PSO to these maps; however, in a real application, PSO governs trace collection and location selection from the outset. The results for the PSO are provided in Table 1 for swarm sizes 3 and 5. "Max traces" denotes the number of traces needed to identify a location in the worst-performing run among 10 PSO executions, and we use this worst case consistently for all comparisons. The "CPA trace" column indicates whether CPA is possible at the selected location; if so, we report the number of traces at which the threshold is crossed permanently, and otherwise report N/A. The "Training Traces" column gives the number of traces used to fit the GMM. Accordingly, the main reduction achieved by GMM-based initialization and Nyström appears in the search cost: GMM initializes PSO in promising regions instead of random locations, which reduces the number of iterations and thus the number of traces collected during localization. Nyström further reduces the trace budget required for MI-based guidance. Overall, MI requires fewer traces than TVLA for key extraction, and in the case of GMM + Nyström, the attack can be achieved with less than 10k traces, improving over the work of Daniel et al. [18].

6.3 FPGA Results

The AES used in the FPGA is the default implementation provided by CW. We use two bitstreams with different AES placements to demonstrate the flexibility of our approach across device configurations. As shown in Figure 4, the TVLA heatmaps vary, reflecting spatially correlated leakage around the AES logic. We perform CPA at the identified high-leakage locations.

For key recovery, we use locations with higher TVLA scores. Now, we will show the correlation results for the FPGA setup with the default bitfile. We consider location (4,0) for FPGA-1 heatmap shown in Figure 4(b), which has a t-score of 30.73. For this location, we plot the max CPA vs traces and CPA vs time samples for 1,000 traces as shown in Figure 5. The model used here is a Hamming-distance model comparing the outputs of rounds 6 and 7 for a single byte. We selected rounds 6 and 7 by analyzing the TVLA plot for the location with the highest TVLA score. The max CPA graph shows that the key could be recovered with as few as 80 averaged traces. The ADC sampling rate is 40 times the clock frequency for higher resolution, whereas the FPGA AES takes 11 clock cycles to run, with the 1st cycle dedicated to key scheduling and control. Therefore, the graph contains 440 samples. The spike appears around sample 320, where the round 7 output is expected to be generated and used in round 8, indicating accurate correlation and observable leakage. The result shown is for key byte 2. This distinction is important in the FPGA setting, where different bytes are processed by different gates at different locations. Accordingly, at location (4,0), byte 10 shows no leakage, as shown in Figure 5, whereas only three locations away, at (7,0), the behavior is reversed: byte 10 leaks and byte 2 does not, as shown in Figure 6. Lastly, Table 1 shows the results of the TVLA- and MI-based PSO searches. MI outperforms TVLA across all cases in terms of the number of traces, since TVLA requires two datasets, one fixed and one random. As in the μC case, the main gain of GMM-based initialization and Nyström lies in reducing the search cost by lowering the number of PSO iterations and the trace budget needed for MI-based guidance.

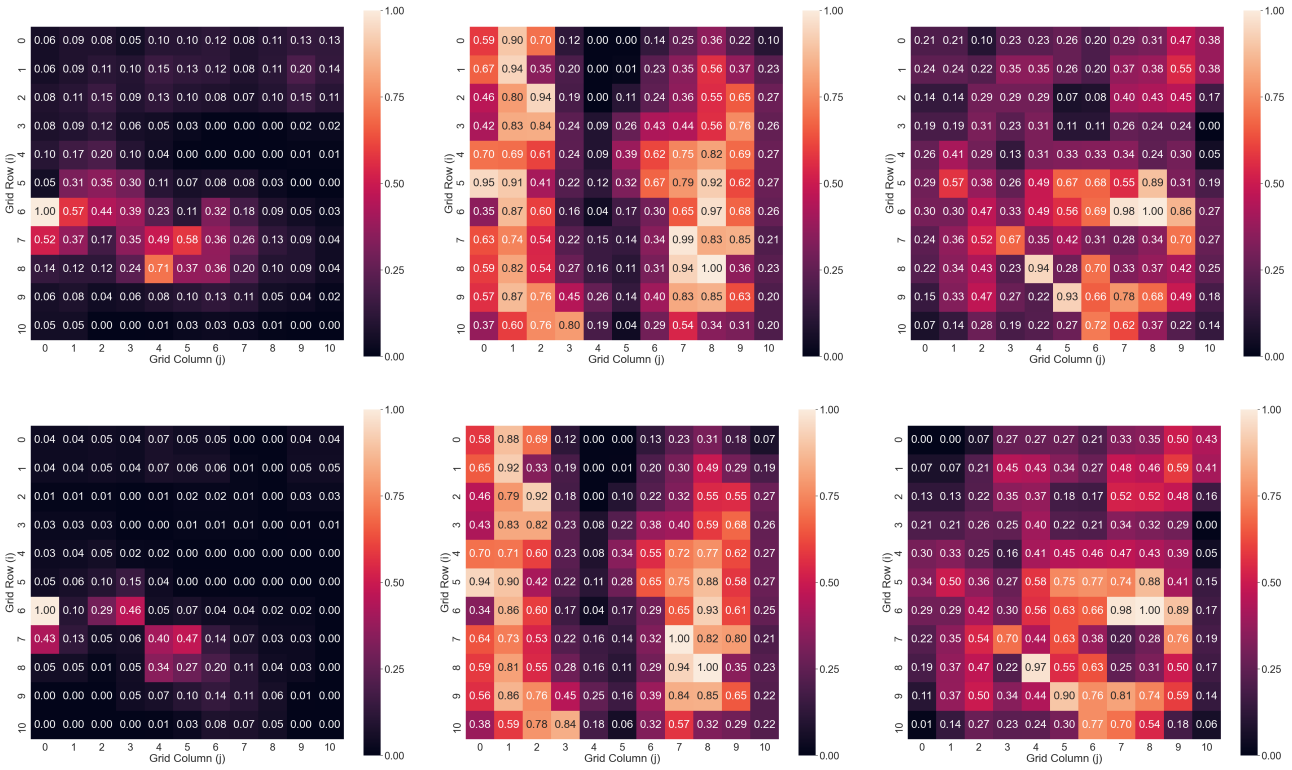

(a) μC (b) FPGA-1 (c) FPGA-2

Figure 3: Entropy map of the EM datasets. The top row shows results computed from 1,000 traces per location, while the bottom row uses Nyström sampling with 700 traces for FPGA-1, 280 for FPGA-2, and 180 for the μC dataset.

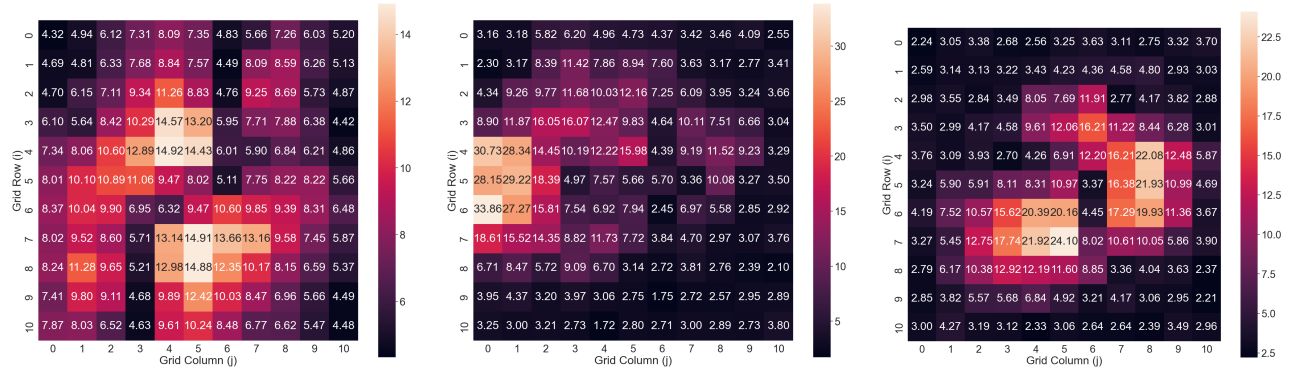

(a) μC (b) FPGA-1 (c) FPGA-2

Figure 4: TVLA heatmaps for different datasets. The peaks in the entropy and TVLA maps need not coincide exactly, since TVLA emphasizes mean-shift leakage under fixed-versus-random partitioning, whereas MI captures broader statistical dependence and can distribute strong but localized effects differently in the presence of noise.

For MI, we observe two distinct outcomes depending on the bitstream. For FPGA-1, LHS performs best, but only slightly better than GMM + Nyström, whereas GMM + Nyström performs best for FPGA-2. This difference is due to the different noise floors of the two implementations, as explained before, whereas random initialization either fails or performs worse than the other methods. Different placements and routing configurations change the sharpness of the leakage landscape; when that prior is less informative, LHS can match or slightly outperform GMM-based initialization, whereas a sharper leakage structure makes GMM + Nyström more effective. Overall, when the map contains informative leakage regions, PSO can identify them faster than the grid-based approach.

6.4 Comparison to SCNIFFER [18]

The closest state-of-the-art to our work is the work by Danial et al. [18]. They have used a gradient descent-based approach for finding the best location. Unfortunately, a direct comparison cannot be made with their results due to differences in setup. In our μC experiments, we observe a similar TVLA

Table 1: Merged comparison of TVLA and MI results across all implementations and swarm sizes.

Dataset	Swarm	Metric	Method	Min Iter	Max Iter	Location	Max Traces	CPA Traces	Training Traces	Total Traces
μC	3	TVLA	Random	3.00	5.00	5 5	30,000	700	0	30,700
			LHS	3.00	6.00	8 6	36,000	600	0	36,600
			GMM	3.00	4.00	4 6	24,000	400	3,630	28,030
		MI	Random	3.00	5.00	7 2	15,000	3,600	0	18,600
			LHS	3.00	4.00	7 1	12,000	700	0	12,700
			GMM	3.00	4.00	1 2	12,000	2,700	3,630	18,330
			GMM+Nyst.	3.00	6.00	7 2	3,240	700	3,630	**7,570**
	5	TVLA	Random	3.00	6.00	8 6	60,000	300	0	60,300
			LHS	3.00	5.00	8 6	50,000	300	0	50,300
			GMM	3.00	4.00	5 5	40,000	700	3,630	44,330
		MI	Random	3.00	6.00	7 2	30,000	700	0	30,700
			LHS	3.00	4.00	7 1	20,000	600	0	20,600
			GMM	3.00	4.00	5 2	20,000	600	3,630	24,230
			GMM+Nyst.	3.00	6.00	6 1	5,400	500	3,630	**9,530**
FPGA-1	3	TVLA	Random	3.00	4.00	7 1	24,000	3,200	0	27,200
			LHS	3.00	6.00	7 1	36,000	3,200	0	39,200
			GMM	3.00	6.00	5 1	36,000	700	3,630	43,330
		MI	Random	3.00	5.00	1 2	15,000	N/A	0	N/A
			LHS	3.00	4.00	3 3	12,000	1,600	0	**13,600**
			GMM	3.00	4.00	4 1	12,000	2,800	3,630	18,430
			GMM+Nyst.	3.00	5.00	4 1	10,500	2,800	3,630	16,930
	5	TVLA	Random	3.00	5.00	7 1	50,000	3,200	0	53,200
			LHS	3.00	5.00	7 1	50,000	3,200	0	53,200
			GMM	3.00	5.00	7 1	50,000	3,200	3,630	56,830
		MI	Random	3.00	5.00	1 2	25,000	N/A	0	N/A
			LHS	3.00	4.00	6 9	20,000	6,100	0	26,100
			GMM	3.00	4.00	4 1	20,000	2,800	3,630	26,430
			GMM+Nyst.	3.00	5.00	4 1	17,500	2,800	3,630	**23,930**
FPGA-2	3	TVLA	Random	3.00	6.00	8 6	36,000	900	0	36,900
			LHS	3.00	5.00	5 9	30,000	2,000	0	32,000
			GMM	3.00	4.00	8 3	24,000	4,200	3,630	31,830
		MI	Random	3.00	4.00	10 6	24,000	9,700	0	33,700
			LHS	3.00	5.00	7 9	30,000	2,300	0	32,300
			GMM	3.00	4.00	7 1	12,000	700	3,630	16,330
			GMM+Nyst.	2.00	4.00	8 5	3,360	4,300	3,630	**11,290**
	5	TVLA	Random	3.00	6.00	8 5	60,000	4,300	0	64,300
			LHS	3.00	4.00	8 6	40,000	900	0	40,900
			GMM	3.00	7.00	8 6	70,000	900	3,630	74,530
		MI	Random	3.00	4.00	7 9	20,000	2,300	0	22,300
			LHS	3.00	5.00	7 9	20,000	2,300	0	22,300
			GMM	3.00	5.00	7 1	25,000	700	3,630	29,330
			GMM+Nyst.	3.00	6.00	6 1	8,400	700	3,630	**12,730**

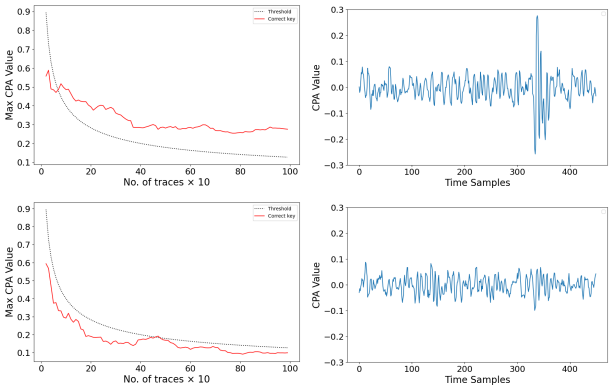

Figure 5: Correlation graphs at location (4,0) using Hamming distance between round 6 and round 7 for Byte 2 (top row) and Byte 10 (bottom row).

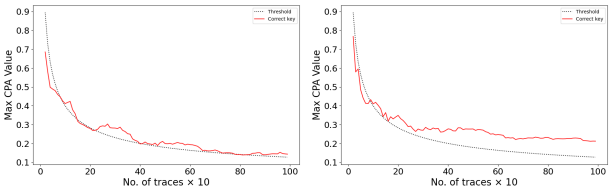

Figure 6: Correlation graphs for Bytes 2 (left) and 10 (right) at location (7,0) using Hamming distance between round 6 and round 7.

heatmap on the same processor. The significant difference is in the probe used in our setup. They have used a larger H-probe (16mm), whereas we have a smaller H-probe (1.25mm). These setup differences substantially affect the required trace counts through changes in noise and signal quality. Despite these disadvantages, we improve over the SNR-based result in their work (10,134 traces), while remaining comparable to their TVLA-based result (5,807 traces). In our TVLA setting, the GMM model with swarm size 3 requires 28,030 traces because we use 1,000 fixed and 1,000 random traces per location, whereas their work uses 400 in total. When we reduce this to 100 fixed and 100 random traces per location, the GMM model recovers the key with 5,630 traces.

In EM side-channel analysis on microcontrollers, increasing the number of traces does not necessarily improve leakage detection and can even degrade the effectiveness of statistical tests such as TVLA [25, 82]. EM leakage is spatially distributed, round-dependent, and highly sensitive to probe coupling, resulting in multiple leakage sources with different timing and polarity. As more traces are collected, non-stationarities such as clock drift, temperature variation, supply noise, or slight probe movement inflate variance and spread leakage energy across time and space, while leakage from different rounds may partially cancel when averaged [44]. Since TVLA detects mean differences normalized by variance at fixed sample points, these effects can reduce the t-statistic despite exploitable leakage, explaining why fewer, cleaner EM traces can outperform much larger acquisitions.

In this setting, TVLA can also outperform MI, even though MI is theoretically more general. TVLA is a hypothesis-driven test that concentrates statistical power on specific mean differences induced by controlled input partitioning, making it sensitive when leakage is strong, localized, and consistently aligned. By contrast, MI measures general statistical dependence without privileging specific leakage models or temporal locations, which can dilute strong, but localized effects when leakage is distributed across rounds or mixed with noise.

Now the question is: why does the certification methodology [83] recommend collecting thousands of traces to compute TVLA? That methodology is designed for certification-oriented leakage assessment rather than attack efficiency. Its recommendation reflects a conservative, worst-case approach aimed at minimizing false negatives under assumptions of stationary, power-like leakage and controlled measurement conditions. These assumptions do not fully hold for EM measurements, where spatial and temporal variability dominate.

7 Conclusion and Discussion

EM side-channel analysis is often viewed as less practical than power-based analysis due to its higher noise levels and complex spatial dependencies. In this work, we introduced a framework that makes EM-based attacks more efficient and scalable by combining GMM, matrix-based Rényi entropy, and PSO. By modeling the spatial distribution of leakage and guiding probe placement accordingly, our method significantly reduces the number of traces required for successful key recovery. We demonstrated that our approach enables faster and more effective attacks than standard methods such as TVLA, while also supporting automation and broader applicability across platforms. Below, we highlight several future directions and implications based on our findings:

- **Multi-Probe Deployment:** Our framework naturally supports extension to multi-probe setups by assigning each PSO particle to a physical probe. This parallelization can reduce acquisition time in proportion to the number of probes, particularly when combined with real-time entropy evaluation. Such configurations can accelerate testing in industrial settings where time and reproducibility are critical.
- **Adaptability Across Devices:** We validated the methodology on both an FPGA and a microcontroller, demonstrating robustness across architectures, packaging styles, and operating conditions. While device-specific parameters such as voltage, frequency, and physical footprint affect leakage characteristics, our unsupervised and information-theoretic approach remains adaptable without requiring prior profiling.
- **Countermeasure Evaluation and Moving Targets:** One potential extension is to evaluate spatially dynamic countermeasures, such as moving target defenses. Our method could help assess how such strategies affect the entropy landscape and whether adaptive probe placement still enables effective attacks. Alternatively, the framework might itself be adapted into a probing countermeasure by dynamically shifting high-leakage points out of reach.

Acknowledgment

This effort was sponsored by NSF Grants CNS-2541809 and CNS-2338069.

References

[1] Rabin Y Acharya, Fatemeh Ganji, and Domenic Forte. Information theory-based evolution of neural networks for side-channel analysis. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2022.

[2] Dakshi Agrawal, Bruce Archambeault, Josyula R Rao, and Pankaj Rohatgi. The em side—channel (s). In *International workshop on cryptographic hardware and embedded systems*, pages 29–45. Springer, 2002.

[3] Hirotogu Akaike. Information theory and an extension of the maximum likelihood principle. In *Selected papers of hirotugu akaike*, pages 199–213. Springer, 1998.

[4] Christos Andrikos, Lejla Batina, Lukasz Chmielewski, Liran Lerman, Vasilios Mavroudis, Kostas Papagiannopoulos, Guilherme Perin, Giorgos Rassias, and Alberto Sonnino. Location, location, location: Revisiting modeling and exploitation for location-based side channel leakages. In *Advances in Cryptology–ASIACRYPT 2019: 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8–12, 2019, Proceedings, Part III 25*, pages 285–314. Springer, 2019.

[5] Christos Andrikos, Giorgos Rassias, Liran Lerman, Kostas Papagiannopoulos, and Lejla Batina. Location-based leakages: New directions in modeling and exploiting. In *2017 International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS)*, pages 246–252. IEEE, 2017.

[6] Daehyeon Bae, Sujin Park, Minsig Choi, Young-Giu Jung, Changmin Jeong, Heeseok Kim, and Seokhie Hong. Probeshooter: A new practical approach for probe aiming. In *Proceedings of the 20th ACM Asia Conference on Computer and Communications Security*, pages 1158–1174, 2025.

[7] Josep Balasch, Benedikt Gierlichs, Roel Verdult, Lejla Batina, and Ingrid Verbauwhede. Power analysis of atmel cryptomemory–recovering keys from secure eeproms. In *Cryptographers' Track at the RSA Conference*, pages 19–34. Springer, 2012.

[8] Lejla Batina, Benedikt Gierlichs, Emmanuel Prouff, Matthieu Rivain, François-Xavier Standaert, and Nicolas Veyrat-Charvillon. Mutual information analysis: a comprehensive study. *Journal of Cryptology*, 24(2):269–291, 2011.

[9] George Becker, Jim Cooper, Elke DeMulder, Gilbert Goodwill, Joshua Jaffe, Gary Kenworthy, Timofei Kouzminov, Andrew Leiserson, Mark Marson, Pankaj Rohatgi, et al. Test vector leakage assessment (tvla) methodology in practice. In *International Cryptographic Module Conference*, volume 1001, page 13. sn, 2013.

[10] Ryad Benadjila, Emmanuel Prouff, Rémi Strullu, Eleonora Cagli, and Cécile Dumas. ASCAD: the ATMega8515 SCA traces databases (fixed key). [Online]https://github.com/ANSSI-FR/ASCAD/tree/master/ATMEGA_AES_v1/ATM_AES_v1_fixed_key [Accessed: Jan.8, 2024], 2017.

[11] Ryad Benadjila, Emmanuel Prouff, Rémi Strullu, Eleonora Cagli, and Cécile Dumas. Study of deep learning techniques for side-channel analysis and introduction to ascad database. 2018.

[12] Ryad Benadjila, Emmanuel Prouff, Rémi Strullu, Eleonora Cagli, and Cécile Dumas. Deep learning for side-channel analysis and introduction to ascad database. *Journal of Cryptographic Engineering*, 10(2):163–188, 2020.

[13] Rajendra Bhatia. Infinitely divisible matrices. *The American Mathematical Monthly*, 113(3):221–235, 2006.

[14] Eric Bogatin. *Signal integrity: simplified*. Prentice Hall Professional, 2004.

[15] Eric Brier, Christophe Clavier, and Francis Olivier. Correlation power analysis with a leakage model. In *International workshop on cryptographic hardware and embedded systems*, pages 16–29. Springer, 2004.

[16] Mahalanobis Prasanta Chandra et al. On the generalised distance in statistics. In *Proceedings of the National Institute of Sciences of India*, volume 2, pages 49–55, 1936.

[17] Guillaume Damblin, Mathieu Couplet, and Bertrand Iooss. Numerical studies of space-filling designs: optimization of latin hypercube samples and subprojection properties. *Journal of Simulation*, 7(4):276–289, 2013.

[18] Josef Danial, Debayan Das, Santosh Ghosh, Arijit Raychowdhury, and Shreyas Sen. Scniffer: Low-cost, automated, efficient electromagnetic side-channel sniffing. *IEEE Access*, 8:173414–173427, 2020.

[19] Debayan Das, Mayukh Nath, Baibhab Chatterjee, Santosh Ghosh, and Shreyas Sen. Stellar: A generic em side-channel attack protection through ground-up root-cause analysis. In *2019 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*, pages 11–20. IEEE, 2019.

[20] Debayan Das, Mayukh Nath, Baibhab Chatterjee, Raghavan Kumar, Xiaosen Liu, Harish Krishnamurthy, Manoj Sastry, Sanu Mathew, Santosh Ghosh, and Shreyas Sen. Em sca whitebox analysis-based reduced leakage cell design and presilicon evaluation. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 41(11):4927–4938, 2022.

[21] Fred Debeer, Marc Witteman, Bartek Gedrojc, and Yijun S Riscure. Practical electro-magnetic analysis. In *Proc. Non-Invasive Attack Testing Workshop NIAT*, 2011.

[22] Bernd Deutschmann, Harald Pitsch, and Gunter Langer. Near field measurements to predict the electromagnetic emission of integrated circuits. In *Proc. 5th Int. Workshop Electromagn. Compat. Integr. Circuits*, pages 27–32, 2005.

[23] Analog devices inc. pytrinamics: A library for controlling trinamics devices. https://github.com/analogdevicesinc/PyTrinamic, 2025. Accessed: 2025-08-07.

[24] Sonia Ben Dhia, Mohamed Ramdani, and Etienne Sicard. *Electromagnetic Compatibility of Integrated Circuits: Techniques for low emission and susceptibility*. Springer Science & Business Media, 2006.

[25] A Adam Ding, Cong Chen, and Thomas Eisenbarth. Simpler, faster, and more robust t-test based leakage detection. In *International workshop on constructive side-channel analysis and secure design*, pages 163–183. Springer, 2016.

[26] Thomas Eisenbarth, Timo Kasper, Amir Moradi, Christof Paar, Mahmoud Salmasizadeh, and Mohammad T Manzuri Shalmani. On the power of power analysis in the real world: A complete break of the keeloq code hopping scheme. In *Annual International Cryptology Conference*, pages 203–220. Springer, 2008.

[27] Karine Gandolfi, Christophe Mourtel, and Francis Olivier. Electromagnetic analysis: Concrete results. In *International workshop on cryptographic hardware and embedded systems*, pages 251–261. Springer, 2001.

[28] Luis G Sanchez Giraldo and Jose C Principe. Rate-distortion auto-encoders. *arXiv preprint arXiv:1312.7381*, 2013.

[29] Luis Gonzalo Sanchez Giraldo, Murali Rao, and Jose C Principe. Measures of entropy from data using infinitely divisible kernels. *IEEE Transactions on Information Theory*, 61(1):535–548, 2014.

[30] Cheng Gongye, Yukui Luo, Xiaolin Xu, and Yunsi Fei. Side-channel-assisted reverse-engineering of encrypted dnn hardware accelerator ip and attack surface exploration. In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 4678–4695. IEEE, 2024.

[31] Gregor Haas and Aydin Aysu. Apple vs. ema: electromagnetic side channel attacks on apple corecrypto. In *Proceedings of the 59th ACM/IEEE Design Automation Conference*, pages 247–252, 2022.

[32] Michael Haider and Johannes A Russer. Principal component analysis for efficient characterization of stochastic electromagnetic fields. *International Journal of Numerical Modelling: Electronic Networks, Devices and Fields*, 31(4):e2246, 2018.

[33] Jared K Harpole, Carol M Woods, Thomas L Rodebaugh, Cheri A Levinson, and Eric J Lenze. How bandwidth selection algorithms impact exploratory data analysis using kernel density estimation. *Psychological methods*, 19(3):428, 2014.

[34] Johann Heyszl. *Impact of localized electromagnetic field measurements on implementations of asymmetric cryptography*. PhD thesis, Technische Universität München, 2013.

[35] Johann Heyszl, Stefan Mangard, Benedikt Heinz, Frederic Stumpf, and Georg Sigl. Localized electromagnetic analysis of cryptographic implementations. In *Cryptographers' track at the RSA conference*, pages 231–244. Springer, 2012.

[36] Johann Heyszl, Dominik Merli, Benedikt Heinz, Fabrizio De Santis, and Georg Sigl. Strengths and limitations of high-resolution electromagnetic field measurements for side-channel analysis. In *International Conference on Smart Card Research and Advanced Applications*, pages 248–262. Springer, 2012.

[37] Gabriel Hospodar, Benedikt Gierlichs, Elke De Mulder, Ingrid Verbauwhede, and Joos Vandewalle. Machine learning in side-channel analysis: a first study. *Journal of Cryptographic Engineering*, 1(4):293–302, 2011.

[38] Xiaoran Huang, Yiwen Gao, Wei Cheng, Yuejun Liu, Jingdian Ming, Yongbin Zhou, and Jian Weng. Towards high-quality electromagnetic leakage acquisition in side-channel analysis. In *2024 IEEE 23rd International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pages 370–377. IEEE, 2024.

[39] Vishnuvardhan V Iyer and Ali E Yilmaz. An adaptive acquisition approach to localize electromagnetic information leakage from cryptographic modules. In *2019 IEEE Texas Symposium on Wireless and Microwave Circuits and Systems (WMCS)*, pages 1–6. IEEE, 2019.

[40] Vishnuvardhan V Iyer and Ali E Yılmaz. An anova method to rapidly assess information leakage near cryptographic modules. *IEEE Transactions on Electromagnetic Compatibility*, 64(4):915–929, 2022.

[41] Vishnuvardhan V Iyer and Ali E Yilmaz. Rapid pre-characterization of fine-grained em side-channel (in) vulnerability of aes modules. In *2022 IEEE USNC-URSI Radio Science Meeting (Joint with AP-S Symposium)*, pages 54–55. IEEE, 2022.

[42] Ian Jolliffe. Principal component analysis. In *International encyclopedia of statistical science*, pages 1094–1096. Springer, 2011.

[43] Michael C Jones, James Stephen Marron, and Simon J Sheather. Progress in data-based bandwidth selection for kernel density estimation. *Computational Statistics*, 11(3):337, 1996.

[44] Shugo Kaji, Daisuke Fujimoto, Youngwoo Kim, and Yuichi Hayashi. A fundamental evaluation of em information leakage induced by iemi for a device with differential signaling. In *2021 Asia-Pacific International Symposium on Electromagnetic Compatibility (APEMC)*, pages 1–4. IEEE, 2021.

[45] Markus Kasper, Timo Kasper, Amir Moradi, and Christof Paar. Breaking keeloq in a flash: on extracting keys at lightning speed. In *International Conference on Cryptology in Africa*, pages 403–420. Springer, 2009.

[46] James Kennedy and Russell Eberhart. Particle swarm optimization. In *Proceedings of ICNN'95-international conference on neural networks*, volume 4, pages 1942–1948. ieee, 1995.

[47] Paul Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In *Annual international cryptology conference*, pages 388–397. Springer, 1999.

[48] Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual information. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 69(6):066138, 2004.

[49] Florian Lemarchand, Cyril Marlin, Florent Montreuil, Erwan Nogues, and Maxime Pelcat. Electro-magnetic side-channel attack through learned denoising and classification. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2882–2886. IEEE, 2020.

[50] Kerstin Lemke-Rust and Christof Paar. Gaussian mixture models for higher-order side channel analysis. In *International Workshop on Cryptographic Hardware and Embedded Systems*, pages 14–27. Springer, 2007.

[51] Er-Ping Li, Xing-Chang Wei, Andreas C Cangellaris, En-Xiao Liu, Yao-Jiang Zhang, Marcello D'amore, Joungho Kim, and Toshio Sudo. Progress review of electromagnetic compatibility

analysis technologies for packages, printed circuit boards, and novel interconnects. *IEEE Transactions on Electromagnetic Compatibility*, 52(2):248–265, 2010.

[52] Qian Li, San-Yang Liu, and Xin-She Yang. Influence of initialization on the performance of metaheuristic optimizers. *Applied Soft Computing*, 91:106193, 2020.

[53] Jake Longo, Elke De Mulder, Dan Page, and Michael Tunstall. Soc it to em: electromagnetic side-channel attacks on a complex system-on-chip. In *International Workshop on Cryptographic Hardware and Embedded Systems*, pages 620–640. Springer, 2015.

[54] Haocheng Ma, Max Panoff, Jiaji He, Yiqiang Zhao, and Yier Jin. Emsim: A fast layout level electromagnetic emanation simulation framework for high accuracy pre-silicon verification. *IEEE Transactions on Information Forensics and Security*, 18:1365–1379, 2023.

[55] Michael D McKay, Richard J Beckman, and William J Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61, 2000.

[56] Geoffrey J McLachlan and David Peel. *Finite mixture models*. John Wiley & Sons, 2000.

[57] Geoffrey J McLachlan and Suren Rathnayake. On the number of components in a gaussian mixture model. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 4(5):341–355, 2014.

[58] Dev Mehta, Trey Marcantino, Mohammad Hashemi, Sam Karkache, Dillibabu Shanmugam, Patrick Schaumont, and Fatemeh Ganji. Scapegoat: Side-channel analysis library. In *2025 IEEE 43rd VLSI Test Symposium (VTS)*, pages 1–7. IEEE, 2025.

[59] Tomas Monopoli, Xinglong Wu, Sergio A Pignari, Johannes Wolf, and Flavia Grassi. Inspection tools for gaussian process regression modeling of electromagnetic fields of electronic boards and chips. In *2024 14th International Workshop on the Electromagnetic Compatibility of Integrated Circuits (EMC Compo)*, pages 56–59. IEEE, 2024.

[60] Amir Moradi, Alessandro Barenghi, Timo Kasper, and Christof Paar. On the vulnerability of fpga bitstream encryption against power analysis attacks: Extracting keys from xilinx virtex-ii fpgas. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 111–124, 2011.

[61] Seyedmohammad Nouraniboosjin and Fatemeh Ganji. Uncertainty estimation in neural network-enabled side-channel analysis and links to explainability. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2026(2):764–807, 2026.

[62] Colin O'flynn and Zhizhang Chen. Chipwhisperer: An open-source platform for hardware embedded security research. In *International Workshop on Constructive Side-Channel Analysis and Secure Design*, pages 243–260. Springer, 2014.

[63] Hossein Pishro-Nik. *Introduction to probability, statistics, and random processes*. Kappa Research, LLC Blue Bell, PA, USA, 2014.

[64] Matthias Probst, Michael Gruber, Manuel Brosch, Tim Music, and Georg Sigl. Switch-glitch: Location of fault injection sweet spots by electro-magnetic emanation. In *2024 Workshop on Fault Detection and Tolerance in Cryptography (FDTC)*, pages 22–27. IEEE, 2024.

[65] Matthias Probst, Alexander Wiesent, Michael Gruber, and Georg Sigl. Probenav-fast, precise and repeatable positioning of electromagnetic probes for local side-channel attacks. *Cryptology ePrint Archive*, 2025.

[66] Emmanuel Prouff and Matthieu Rivain. Theoretical and practical aspects of mutual information based side channel analysis. In *Applied Cryptography and Network Security: 7th International Conference, ACNS 2009, Paris-Rocquencourt, France, June 2-5, 2009. Proceedings 7*, pages 499–518. Springer, 2009.

[67] Emmanuel Prouff and Matthieu Rivain. Theoretical and practical aspects of mutual information-based side channel analysis. *International Journal of Applied Cryptography*, 2(2):121–138, 2010.

[68] Jean-Jacques Quisquater and David Samyde. Electromagnetic analysis (ema): Measures and counter-measures for smart cards. In *International Conference on Research in Smart Cards*, pages 200–210. Springer, 2001.

[69] Bastian Richter, Alexander Wild, and Amir Moradi. Automated probe repositioning for on-die em measurements. In *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 1–6. IEEE, 2019.

[70] Thomas Roche. Eucleak side-channel attack on the yubikey 5 series (revealing and breaking infineon ecdsa implementation on the way). In *2025 IEEE Symposium on Security and Privacy (SP)*, pages 4026–4043. IEEE, 2025.

[71] Pankaj Rohatgi. Electromagnetic attacks and countermeasures. In *Cryptographic Engineering*, pages 407–430. Springer, 2009.

[72] Laurent Sauvage, Sylvain Guilley, and Yves Mathieu. Electromagnetic radiations of fpgas: High spatial resolution cartography and attack on a cryptographic module. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, 2(1):1–24, 2009.

[73] Nimisha Sivaraman. *Design of magnetic probes for near field measurements and the development of algorithms for the prediction of EMC*. PhD thesis, Université Grenoble Alpes, 2017.

[74] Robert Specht, Johann Heyszl, and Georg Sigl. Investigating measurement methods for high-resolution electromagnetic field side-channel analysis. In *2014 International Symposium on Integrated Circuits (ISIC)*, pages 21–24. IEEE, 2014.

[75] Hiroshi Suenaga, Akihiro Tsukioka, Kosuke Jike, and Makoto Nagata. Compact simulation of chip-to-chip active noise coupling on a system pcb board. *IEEE Letters on Electromagnetic Compatibility Practice and Applications*, 2(1):15–20, 2020.

[76] Xin Tong, David WP Thomas, Angela Nothofer, Phillip Sewell, and Christos Christopoulos. Modeling electromagnetic emissions from printed circuit boards in closed environments using equivalent dipoles. *IEEE Transactions on Electromagnetic Compatibility*, 52(2):462–470, 2010.

[77] Thomas Unterluggauer, Thomas Korak, Stefan Mangard, Robert Schilling, Luca Benini, Frank K Gürkaynak, and

Michael Muehlberghuber. Leakage bounds for gaussian side channels. In *International Conference on Smart Card Research and Advanced Applications*, pages 88–104. Springer, 2017.

[78] Florian Unterstein, Johann Heyszl, Fabrizio De Santis, and Robert Specht. Dissecting leakage resilient prfs with multivariate localized em attacks: A practical security evaluation on fpga. In *International Workshop on Constructive Side-Channel Analysis and Secure Design*, pages 34–49. Springer, 2017.

[79] Gerald Van Belle. *Statistical rules of thumb*. John Wiley & Sons, 2011.

[80] Debao Wang, Yiwen Gao, Yongbin Zhou, and Xian Huang. Revisiting a realistic em side-channel attack on a complex modern soc. *Cryptology ePrint Archive*, 2024.

[81] Xing-Chang Wei. *Modeling and design of electromagnetic compatibility for high-speed printed circuit boards and packaging*. CRC Press, 2017.

[82] Carolyn Whitnall and Elisabeth Oswald. A cautionary note regarding the usage of leakage detection tests in security evaluation. *Cryptology ePrint Archive*, 2019.

[83] Carolyn Whitnall and Elisabeth Oswald. A critical analysis of iso 17825 ('testing methods for the mitigation of non-invasive attack classes against cryptographic modules'). In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 256–284. Springer, 2019.

[84] Shujian Yu and José C Príncipe. Simple stopping criteria for information theoretic feature selection. *Entropy*, 21(1):99, 2019.

[85] Shujian Yu, Kristoffer Wickstrøm, Robert Jenssen, and Jose C Principe. Understanding convolutional neural networks with information theory: An initial exploration. *IEEE transactions on neural networks and learning systems*, 32(1):435–442, 2020.

[86] Kai Zhang and James T Kwok. Clustered nyström method for large scale manifold learning and dimension reduction. *IEEE Transactions on Neural Networks*, 21(10):1576–1587, 2010.

[87] Kai Zhang, Ivor W Tsang, and James T Kwok. Improved nyström low-rank approximation and error analysis. In *Proceedings of the 25th international conference on Machine learning*, pages 1232–1239, 2008.