# Poster: s-MDM: Generative Virtualization of Multi-Device Hardware Variations for Portable DL-SCA

Niloufar Sayadi✉
Centrum Wiskunde & Informatica
Vrije Universiteit Amsterdam
Amsterdam, The Netherlands
n.sayadi@cwi.nl

Marten van Dijk
Centrum Wiskunde & Informatica
Vrije Universiteit Amsterdam
Amsterdam, The Netherlands
marten.van.dijk@cwi.nl

Chenglu Jin
Centrum Wiskunde & Informatica
Amsterdam, The Netherlands
chenglu.jin@cwi.nl

## Abstract

Deep Learning-based Side-Channel Analysis (DL-SCA) frequently suffers from catastrophic performance degradation across unseen hardware due to printed circuit board routing differences, silicon process variations, and measurement noise shifts. This poster presents the Synthetic Multiple Device Model (s-MDM), a zero-target-trace generative framework designed to improve cross-device portability. s-MDM combines a structured cVAE generator, a Walsh–Hadamard leakage anchor, continuous style modulation, and decoupled leakage–style–domain critics to synthesize virtual source-device profiles offline. Benchmarked on 32-bit side-channel traces (`AES_PTv2`), s-MDM maps a precise operational boundary: while physical MDM remains superior on identical electrical clones ($D_4$), s-MDM achieves consistently low key rank on the layout/acquisition-shifted Piñata target, where physical baselines are unstable or misaligned.

## CCS Concepts

• **Security and privacy** → **Hardware security implementation**; **Side-channel analysis and countermeasures**; • **Computing methodologies** → *Machine learning approaches.*

## Keywords

Side-Channel Analysis, Deep Learning, Generative Modeling, Process Variation

**ACM Reference Format:**
Niloufar Sayadi, Marten van Dijk, and Chenglu Jin. 2026. Poster: s-MDM: Generative Virtualization of Multi-Device Hardware Variations for Portable DL-SCA. In *Proceedings of the 2026 ACM SIGSAC Conference on Computer and Communications Security (CCS '26), November 15–19, 2026, The Hague, Netherlands.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3830454.3846456

## 1 Introduction

DL-SCA breaks cryptographic implementations efficiently [3]. However, classifiers suffer severe performance degradation across unseen target hardware due to the *portability problem*: process variations, PCB routing differences, and noise shifts alter physical leakage distributions [1, 8]. Existing approaches for cross-device portability have severe operational limits:

(1) **Physical Multi-Device Modeling (MDM):** MDM trains classifiers across a physical fleet of boards to learn fleet-invariant features and treat single-board traits as noise [1]. However, procuring and profiling physical fleets incurs severe financial and logistical overheads [1].
(2) **Target-Assisted Domain Adaptation:** Methods like CDPA [4] and UDA fine-tuning [9] align features across hardware, but strictly require capturing unlabeled target-device traces during profiling—a condition infeasible in non-permissive scenarios [6, 8].

***Our Contributions.*** We present the **Synthetic Multiple Device Model (s-MDM)** for zero-target-trace *device-efficient portability*. Our key contributions are:

- We introduce a generative virtualization framework that constructs a continuous synthetic device manifold from minimal source hardware.
- We design a hybrid cVAE-GAN architecture that extends basis-projected modeling [2, 11] with Gradient Reversal Layers (GRL), continuous style modulation, and decoupled critics to separate deterministic leakage from physical style variations.
- We empirically characterize the operational boundary between physical pooling and synthetic virtualization on `AES_PTv2`, establishing when direct multi-device modeling remains preferable and when s-MDM improves generalization under layout and acquisition shift.

***Background and Related Work.*** Portability methods following MDM [1], including Cross-Device Profiled Attacks (CDPA) [4], MMD minimization [9], and adversarial domain learning [5, 10], align feature distributions between devices, but unlike MDM, they require unlabeled victim-device traces during profiling [4, 8, 9]. Generative SCA methods have been used for augmentation [7], explicit noise-leakage factorization over monomial bases (cVAE-SA [11], cVAE-OSM [2]), or transferring white-box reference features to black-box targets (CGAN-SCA [6]). Existing generative SCA architectures mostly rely either on victim-device measurements for domain transfer [5, 6] or restrict noise modeling to a single physical device without spanning distinct hardware configurations [2, 11]. In contrast, s-MDM uses only source-device traces and synthesizes virtual device-style variation before deployment.

## 2 Methodology

The s-MDM framework is engineered for zero-target-trace side-channel portability. Instead of relying on target-dependent domain adaptation, which requires unlabeled traces from the target device, s-MDM learns a generative model only from the available profiling devices and uses it offline to synthesize virtual device styles. The objective is to expose the downstream attack classifier to a wider

This work is licensed under a Creative Commons Attribution 4.0 International License.
*CCS '26, The Hague, Netherlands*
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2871-6/2026/11
https://doi.org/10.1145/3830454.3846456

range of physically plausible leakage conditions before deployment, while keeping the final target device completely unseen during training and validation.

***Leakage Decomposition.*** s-MDM builds on cVAE-OSM [2] for structured leakage factorization. Given a raw trace observation vector $\mathbf{T} \in \mathbb{R}^D$, s-MDM models the trace through the following decomposition:

$$\mathbf{T} = \Psi(Y) + \mathbf{N}_{\text{content}}(\mathbf{z}, \boldsymbol{\sigma}^2) + \mathbf{N}_{\text{style}}(\mathbf{z}_{\text{style}}) \tag{1}$$

- $\Psi(Y) = b(Y)^\top \mathbf{W}_\psi$ is the deterministic AES-dependent leakage template, where $b(Y)$ is a 256-dimensional Walsh–Hadamard basis representation of the 8-bit intermediate value $Y$ [2].
- $\mathbf{N}_{\text{content}}(\mathbf{z}, \boldsymbol{\sigma}^2)$ models stochastic content variation using a latent code $\mathbf{z}$ and heteroscedastic variance $\boldsymbol{\sigma}^2$.
- $\mathbf{N}_{\text{style}}(\mathbf{z}_{\text{style}})$ models a device-specific physical style signature through a continuous 16-dimensional style code.

***The Hybrid cVAE-GAN Architecture.*** s-MDM implements the decomposition in Eq. 1 through a hybrid cVAE–GAN. It couples the structured density modeling of a cVAE with the adversarial feedback of GAN-style critics to bypass traditional VAE waveform blurring. The generator disentangles execution content from style, while three decoupled critics ( $D_{\text{style}}$, $D_{\text{leak}}$, and $D_{\text{domain}}$) enforce physical realism, cryptographic alignment, and device invariance. The proposed architecture comprises four core functional components:

(1) **Structured cVAE Generator:** Leveraging a cVAE-OSM backbone [2], the content encoder estimates $q_\phi(\mathbf{z} \mid \mathbf{T}, b(Y))$ and heteroscedastic residual variance $\boldsymbol{\sigma}^2$. The deterministic subspace $\Psi(Y) = b(Y)^\top \mathbf{W}_\psi$ remains strictly anchored to prevent feature bleaching during optimization.

(2) **Continuous Style Modulation:** A blind, secret-unaware encoder extracts a style code $\mathbf{z}_{\text{style}}$ capturing physical waveform appearance. The decoder re-renders the synthesized trace $\widehat{\mathbf{T}}$ by modulating the residual noise via Feature-wise Linear Modulation (FiLM). While Eq. 1 gives the conceptual trace decomposition, the implemented decoder realizes the style term through FiLM modulation of the residual content-noise pathway as

$$\widehat{\mathbf{T}} = \Psi(Y) + \boldsymbol{\gamma}(\mathbf{z}_{\text{style}}) \odot \mathbf{N}_{\text{content}}(\mathbf{z}, \boldsymbol{\sigma}^2) + \boldsymbol{\beta}(\mathbf{z}_{\text{style}}), \tag{2}$$

where $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ are trace-length scaling and shift vectors. Eq. 2 explicitly realizes the style component from Eq. 1 as $N_{\text{style}}(z_{\text{style}}) = (\gamma(z_{\text{style}}) - 1) \odot N_{\text{content}} + \beta(z_{\text{style}})$.

(3) **Decoupled Adversarial Critics:** $D_{\text{style}}$ evaluates secret-blind physical trace morphology. Concurrently, the leakage critic

$$D_{\text{leak}}(\mathbf{T}, Y) = \psi(\phi(\mathbf{T})) + \phi(\mathbf{T})^\top \mathbf{V} b(Y)$$

forces the CNN feature map $\phi(\mathbf{T})$ to preserve trace–label cryptographic alignment.

(4) **GRL Domain Safeguard:** A domain head $D_{\text{domain}}$ predicts profiling board identity from $\phi(\mathbf{T})$. A Gradient Reversal Layer (GRL), acting as $\mathcal{R}_\lambda(\phi(\mathbf{T})) = \phi(\mathbf{T})$ with reversed backpropagation gradients ($\frac{\partial \mathcal{R}_\lambda}{\partial \phi} = -\lambda I$), penalizes board-identifying features and encourages device-invariant representations.

***Continuous Manifold Virtualization.*** Once the hybrid cVAE–GAN is trained on the available profiling fleet, the generator is frozen and used as an offline virtual-device synthesizer. Instead of sampling only from the observed profiling styles, s-MDM samples the continuous style code from an expanded latent envelope, $\mathbf{z}_{\text{style}} \sim \mathcal{N}(0, \alpha \mathbf{I})$, $\alpha = 1.5$, and combines $\mathbf{z}_{\text{style}}$ with $Y$ and content latents via Eq. 2. This produces a virtual fleet of style-diverse traces that interpolate and mildly extrapolate beyond the observed source-board appearances without using any traces from the final target device. The purpose is to broaden the training distribution seen by the downstream attack classifier before deployment. The resulting dataset $\mathcal{D}^{\text{attack}}_{\text{train}} = \mathcal{D}^{\text{profiling}}_{\text{real}} \cup \mathcal{D}^{\text{s-MDM}}_{\text{synthetic}}$ trains a downstream attack classifier $C_\theta$ offline. During the final attack (key recovery), target traces are processed directly by the trained classifier while the generator and all adversarial critics are discarded:

$$\mathbf{T}_{\text{target}} \longrightarrow C_\theta(\mathbf{T}_{\text{target}}) \longrightarrow p(Y \mid \mathbf{T}_{\text{target}}) \longrightarrow \text{key ranking.}$$

No target traces are used to train, adapt, validate, or augment the generator. To avoid contaminating the sealed target evaluation, s-MDM uses a Leave-One-Board-Out (LOBO) validation protocol. The model is trained on the source boards and validated on a held-out board from the source-device family that acts as a pseudo-target. This held-out board provides a portability proxy for selecting training checkpoints and monitoring leakage-fidelity diagnostics such as $\text{SNR}_\Psi$, guessing entropy, and success rate. The final target remains completely unseen until the final deployment evaluation.

## 3 Experimental Evaluations

***Experimental Setup.*** The s-MDM framework was evaluated on an unprotected 32-bit software AES-128 engine using the AES_PTv2 side-channel trace infrastructure. The device fleet contains four nominally identical STM32F411E-DISCO boards ($D_1$–$D_4$) and one distinct Riscure Piñata board. Boards $D_1$–$D_4$ use the STM32F411VE Cortex-M4 platform and are measured through a noisy, less-invasive EM setup, while Piñata is a Cortex-M4F-based training board with a modified power network and cleaner direct power acquisition. We use $D_1$ and $D_2$ for profiling, $D_3$ for LOBO validation, $D_4$ as an unseen electrical clone target, and Piñata as a layout/acquisition-shifted target.

Each device is characterized by its leakage-peak sample index, peak SNR, $\rho = \text{SNR}/(1 + \text{SNR})$, and device-specific characterization window in Table 1. For model ingestion, all traces are cropped around each device's leakage peak to a common 120-sample width. We report guessing entropy GE, the mean zero-based rank of the correct key, and success rate SR, the fraction of runs reaching rank zero. While the profiling boards are strongly noise-dominated ($\rho_{D_1} = 0.072$), the Piñata target shifts from the noise-dominated profiling regime to a signal-dominated regime ($\rho_{\text{Piñata}} = 0.737$), consistent with a substantial layout/acquisition shift rather than ordinary clone-to-clone variation.

**Table 1: Physical characterization of the device fleet.**

| Device | Role | Peak idx. | Peak SNR | $\rho$ | Device-optimal window |
|---|---|---|---|---|---|
| $D_1$ | profiling | 48 | 0.077 | 0.072 | [0, 128) |
| $D_2$ | profiling | 63 | 0.135 | 0.119 | [0, 143) |
| $D_3$ | validation | 39 | 0.260 | 0.206 | [0, 119) |
| $D_4$ | sealed target | 39 | 0.224 | 0.183 | [0, 119) |
| Piñata | sealed target | 104 | **2.796** | **0.737** | [24, 184) |

***In-Manifold Evaluation (Electrical Clone).*** On the electrical clone target $D_4$ (SNR = 0.224, $\rho_{D_4} = 0.183$), physical pooling is optimal: Physical MDM reaches full recovery in 109 traces and

**Table 4: Structural ablation of generative components (Piñata target, single run per configuration). For reference, the full configuration spans GE = 0.42–3.28 across the seeds of Table 3.**

| Ablation State | Final GE | Final SR | Fidelity ($SNR_\Psi$) |
|---|---|---|---|
| s-MDM Complete | 0.42 | 0.71 | 0.057 |
| Without Style Diversity | **7.39** | **0.33** | 0.058 |
| Without Device-Invariance | 0.56 | 0.81 | 0.058 |
| Without Adversarial Critics | 0.03 | 0.97 | 0.055 |
| Without Subspace Anchor | 0.43 | 0.80 | 0.057 |

the two-device baseline in 586 traces (Table 2). In contrast, our s-MDM improves over single-device training but does not converge within the evaluated trace budget, as shown in Table 2. This result is consistent with the clone setting where real multi-device traces are more informative than synthetic re-rendering. In this regime, the generator can act as a compression bottleneck rather than a portability advantage.

**Table 2: Portability Profiles on Electrical Clone Target $D_4$.**

| Strategy | Physical Support | Final GE | Final SR | Conv. (Traces) |
|---|---|---|---|---|
| Single-Device | 1 ($D_1$) | 18.36 | 0.16 | — |
| Two-Device | 2 ($D_1, D_2$) | 0.00 | 1.00 | 586 |
| Physical MDM | 3 ($D_1$–$D_3$) | 0.00 | 1.00 | 109 |
| s-MDM (Full) | Virtualized (from 2) | 5.72 | 0.21 | — |

***Out-of-Manifold Evaluation (Layout and Acquisition Shift:*** Piñata***).*** The benefit of s-MDM appears in the out-of-support setting. Piñata has a much stronger leakage regime than profiling boards (SNR = 2.796, $\rho_{\text{Piñata}} = 0.737$).

Table 3 shows that the physical baselines are unstable under this shift. Single-device training fails to recover, two-device training is highly seed-sensitive, and Physical MDM ranks the key worse than random. In contrast, s-MDM achieves low guessing entropy, with GE $\leq$ 3.3 across the three independent realizations. The improvement suggests that continuous style virtualization exposes the classifier to a broader range of leakage appearances before deployment, reducing its dependence on the low-SNR morphology of the profiling boards.

**Table 3: Replication of final GE across three independent random seeds on the Piñata target (SNR = 2.796, $\rho_{\text{Piñata}}$ = 0.737). $GE_{\text{rand}} = 127.5$ (256 keys); higher values indicate anti-correlated (worse-than-random) ranking.**

| Strategy | Seed 0 | Seed 1 | Seed 2 | Status |
|---|---|---|---|---|
| Single-Device | 36.32 | 18.00 | 16.95 | Fails |
| Two-Device | 25.50 | 0.98 | 235.47 | Unstable |
| Physical MDM | 146.36 | 239.07 | 243.79 | Anti-Correlated |
| **s-MDM (Full)** | **0.42** | **0.69** | **3.28** | **Recovers** |

***Ablation Study.*** To identify which architectural component drives the observed transfer, we ablate the main components of s-MDM on the layout-shifted target Piñata, as detailed in Table 4. Each ablation reflects a single training run. For calibration, the full configuration spans GE = 0.42–3.28 across the three independent random seeds in Table 3. Only removal of continuous style diversity (GE = 7.39) degrades performance beyond this range, reducing the SR from 0.71 to 0.33. Thus, style diversity is the only component whose removal shows a resolved performance degradation under the observed run-to-run variability. The remaining ablations fall inside the observed seed spread and are therefore not individually resolved here. In particular, although removing the adversarial critics yields GE = 0.03 in this run, this result cannot establish a definitive improvement over the complete model without multi-seed replication.

## 4 Conclusion and Future Work

s-MDM demonstrates that zero-target-trace generative virtualization can improve DL-SCA portability bottlenecks under severe layout and acquisition shifts. On the electrical clone target $D_4$, physical MDM remains superior, showing that synthetic re-rendering acts as an information filter when the target lies inside the source manifold. On Piñata, s-MDM reaches low GE while physical baselines are unstable or anti-correlated. The ablation study identifies continuous style diversity as the only component whose removal causes a resolved performance degradation; the roles of the critics require multi-seed replication. Future work will refine critic optimization, calibrate style priors, and evaluate larger multi-board fleets.

## Acknowledgments

Marten van Dijk and Chenglu Jin are (partially) supported by project CiCS of the research programme Gravitation, which is (partly) financed by the Dutch Research Council (NWO) under the grant 024.006.037.

## References

[1] Shivam Bhasin, Anupam Chattopadhyay, Annelie Heuser, Dirmanto Jap, Stjepan Picek, and Ritu Ranjan Shrivastwa. 2020. Mind the portability: A warriors guide through realistic profiled side-channel analysis. In *NDSS 2020-Network and Distributed System Security Symposium*. Internet Society, Reston, VA, 1–14.

[2] Sana Boussam, Mathieu Carbone, Benoît Gérard, Guénaël Renault, and Gabriel Zaid. 2025. Optimal Dimensionality Reduction using Conditional Variational AutoEncoder. *IACR Transactions on Cryptographic Hardware and Embedded Systems* 2025, 3 (2025), 164–211.

[3] Eleonora Cagli, Cécile Dumas, and Emmanuel Prouff. 2017. Convolutional Neural Networks with Data Augmentation Against Jitter-Based Countermeasures. In *Cryptographic Hardware and Embedded Systems (CHES)*. Springer, Taipei, Taiwan, 45–68.

[4] Pei Cao, Chi Zhang, Xiangjun Lu, and Dawu Gu. 2021. Cross-Device Profiled Side-Channel Attack with Unsupervised Domain Adaptation. *IACR Transactions on Cryptographic Hardware and Embedded Systems (TCHES)* 2021, 4 (2021), 27–56.

[5] Pei Cao, Hongyi Zhang, Dawu Gu, Yan Lu, and Yidong Yuan. 2022. AL-PA: Cross-Device Profiled Side-Channel Attack Using Adversarial Learning. In *Proceedings of the 59th ACM/IEEE Design Automation Conference (DAC)*. ACM, San Francisco, CA, USA, 691–696.

[6] Sengim Karayalçın, Marina Krček, Lichao Wu, Stjepan Picek, and Guilherme Perin. 2024. It's a Kind of Magic: A Novel Conditional GAN Framework for Efficient Profiling Side-Channel Analysis. In *Advances in Cryptology – ASIACRYPT 2024*. Springer, Kolkata, India, 99–131.

[7] Naila Mukhtar, Lejla Batina, Stjepan Picek, and Yinan Kong. 2022. Fake It Till You Make It: Data Augmentation Using Generative Adversarial Networks for All the Crypto You Need on Small Devices. In *CT-RSA 2022*. Springer, San Francisco, CA, 297–321.

[8] Stjepan Picek, Guilherme Perin, Luca Mariot, Lichao Wu, and Lejla Batina. 2023. SoK: Deep Learning-Based Physical Side-Channel Analysis. *Comput. Surveys* 55, 11 (2023), 1–35.

[9] Ji-Eun Woo, Yongsung Jeon, Ju-Hwan Kim, and Dong-Guk Han. 2025. Novel deep learning-based side-channel attack on different-device. *Soft Computing* 29, 8 (2025), 3847–3854.

[10] Honggang Yu, Haoqi Shan, Maximillian Panoff, and Yier Jin. 2021. Cross-Device Profiled Side-Channel Attacks Using Meta-Transfer Learning. In *ACM/IEEE Design Automation Conference (DAC)*. IEEE, San Francisco, CA, 703–708.

[11] Gabriel Zaid, Lilian Bossuet, Mathieu Carbone, Amaury Habrard, and Alexandre Venelli. 2023. Conditional variational autoencoder based on stochastic attacks. *IACR TCHES* 2023, 2 (2023), 310–357.