

Learning to Localize Leakage of Cryptographic Sensitive Variables

Anonymous authors
Paper under double-blind review

Abstract

While cryptographic algorithms such as the ubiquitous Advanced Encryption Standard (AES) are secure, *physical implementations* of these algorithms in hardware inevitably ‘leak’ sensitive data such as cryptographic keys. A particularly insidious form of leakage arises from the fact that hardware consumes power and emits radiation in a manner that is statistically associated with the data it processes and the instructions it executes. Supervised deep learning has emerged as a state-of-the-art tool for carrying out *side-channel attacks*, which exploit this leakage by learning to map power/radiation measurements throughout encryption to the sensitive data operated on during that encryption. In this work we develop a principled deep learning framework for determining the relative leakage due to measurements recorded at different points in time, in order to inform *defense* against such attacks. This information is invaluable to cryptographic hardware designers for understanding *why* their hardware leaks and how they can mitigate it (e.g. by indicating the particular sections of code or electronic components which are responsible). Our framework is based on an adversarial game between a classifier trained to estimate the conditional distributions of sensitive data given subsets of measurements, and a budget-constrained noise distribution which probabilistically erases individual measurements to maximize the loss of this classifier. We demonstrate our method’s efficacy and ability to overcome limitations of prior work through extensive experimental comparison on 6 publicly-available power/EM trace datasets from AES, ECC and RSA implementations. Our PyTorch code is available here.

1 Introduction

The Advanced Encryption Standard (AES)¹ (Daemen & Rijmen, 1999; 2013) is widely used and trusted for protecting sensitive data. For example, it is approved by the United States National Security Agency for protecting top secret information (Committee on National Security Systems, 2003), it is a major component of the Transport Layer Security (TLS) protocol (Rescorla, 2000) which underlies the security of HTTPS (Rescorla, 2000), and it is used in payment card readers to secure card information before transmission to financial institutions (Bluefin Payment Systems, 2023).

AES aims to keep data secret when it is transmitted over insecure channels that are accessible to unknown and untrusted parties (e.g. via wireless transmissions which may be intercepted, or storage on hard drives accessible to untrusted individuals). Prior to transmission, the data is first encoded and partitioned into a sequence of fixed-length bitstrings called *plaintexts*. Each plaintext is then *encrypted* into a *ciphertext* by applying an invertible function from a family of functions indexed by an integer called a *cryptographic key*. This family of functions is designed so that if the key is sampled uniformly at random, then the plaintext and ciphertext are marginally independent. The key is known to the sender and intended recipients of the transmission, and is kept secret from potential eavesdroppers. Thus, the intended recipients can use the key to *decrypt* the ciphertext back into the original plaintext, while eavesdroppers who possess the ciphertext but not the key learn nothing about the plaintext.

¹For clarity of exposition we discuss AES here, but our technique also applies to other algorithms.

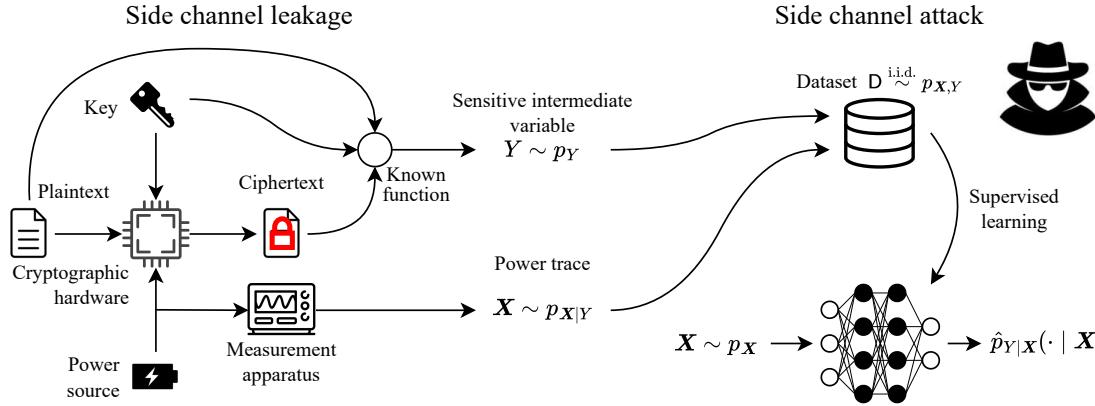


Figure 1: Diagram illustrating our probabilistic framing of side-channel leakage in the special case of power side-channel leakage from a symmetric-key (e.g. AES) cryptographic implementation. A cryptographic device consumes power over time while encrypting data. The power consumption leaks the secret key because it is statistically associated with key-dependent internal variables. Our work considers profiling side-channel attacks, a worst-case scenario where attackers can freely model the relationship between an implementation’s secret data and its power consumption over time, then use this model to attack the same device.

37 Clearly, such an algorithm is effective only if the cryptographic key remains outside of the hands of attackers.
 38 AES is believed to be ‘algorithmically secure’ in the sense that it is infeasible to determine the cryptographic
 39 key by exploiting its intended inputs and outputs: plaintexts and ciphertexts (Mouha, 2021; Tao & Wu, 2015).
 40 Despite this, *physical implementations* of AES in hardware ‘leak’ information about their cryptographic keys.
 41 This phenomenon, called *side-channel leakage*, occurs because hardware inevitably emits measurable physical
 42 signals that are statistically associated with the data it processes and the instructions it executes (Mangard
 43 et al., 2007). There are many diverse physical side-channels which can leak, such as program/operation
 44 execution time (Kocher, 1996; Lipp et al., 2018; Kocher et al., 2019), temperature (Hutter & Schmidt, 2014),
 45 and sound due to vibration of electronic components (Genkin et al., 2014). In this work the side-channels
 46 we consider are *power and electromagnetic (EM) radiation over time* (Kocher et al., 1999; Mangard et al.,
 47 2007), which are major security vulnerabilities for AES implementations (Genkin et al., 2016; Bronchain &
 48 Standaert, 2020).

49 *Side-channel attacks* exploit this leakage to break cryptographic implementations by revealing secret internal
 50 variables such as cryptographic keys. In this work we consider *profiling* side-channel attacks (Chari et al.,
 51 2003; Mangard et al., 2007), which assume a worst-case scenario where the attacker possesses a clone of the
 52 target device and can repeatedly measure its power/radiation over time while encrypting arbitrary plaintexts
 53 using arbitrary keys. These sequences of power/radiation measurements are recorded as real vectors called
 54 *traces*, where each element encodes a measurement at a fixed point in time relative to the start of encryption.
 55 The attacker can thereby gather data from the clone device to model the conditional distribution of the
 56 secret variable given a trace. They can then defeat the target device by measuring its power/radiation traces,
 57 feeding them to the model, and revealing its secret internal variables.

58 Supervised deep learning has emerged as a state-of-the-art technique for this modeling task, achieving
 59 comparable or superior performance to prior approaches while requiring far less data preprocessing and
 60 feature selection (Maghrebi et al., 2016; Benadjila et al., 2020; Zaid et al., 2020; Wouters et al., 2020; Lu et al.,
 61 2021; Bursztein et al., 2023). Older non-deep learning attacks were mostly based on parametric statistics and
 62 had major limitations such as restrictive modeling assumptions (Messerges, 2000; Chari et al., 2003; Agrawal
 63 et al., 2005; Schindler et al., 2005; Hospodar et al., 2011) and inability to scale to long traces (Chari et al.,
 64 2003; Archambeau et al., 2006). Deep learning overcomes these limitations and consequently poses a major
 65 and growing threat to a wide assortment of security measures and evaluations that were designed with the
 66 limitations of older attacks in mind.

Table 1: A summary of the important quantities and variables used throughout this work.

| Notation | Explanation |
|--|---|
| $p_A, p_{A B}$, etc. | Probability distribution of random variable A , conditional distribution of A given B , etc. |
| $\mathbb{E} f(A, B), \mathbb{E}_A f(A, B)$ | Expected value of $f(A, B)$ (left: over all randomness, right: over randomness of A) |
| $\mathbb{I}[A; B C] \in \mathbb{R}_+$ | Conditional mutual information between A and B given C |
| $T \in \mathbb{Z}_{++}$ | Number of measurements/timesteps in power trace |
| $\mathbf{X} \equiv (X_1, \dots, X_T) \in \mathbb{R}^T$ | Power trace (random variable) consisting of T measurements |
| $Y \in \mathbb{Y}$ | Secret data (random variable) from finite set \mathbb{Y} (e.g. $\mathbb{Y} = \{0, 1, \dots, 255\}$ for <code>uint8</code> Y) |
| $\gamma \equiv (\gamma_1, \dots, \gamma_T) \in [0, 1]^T$ | Occlusion probabilities, where γ_t denotes probability of occluding X_t |
| $C \in \mathbb{R}_{++}$ | Budget for occlusion probability vector, which incurs cost $\sum_{t=1}^T c(\gamma_t)$ |
| $\bar{\gamma} = C/(C + T) \in (0, 1)$ | Re-parameterized version of C which is less-sensitive to dataset |
| $\mathcal{A}_\gamma \equiv (\mathcal{A}_{\gamma,1}, \dots, \mathcal{A}_{\gamma,T}) \in \{0, 1\}^T$ | Occlusion mask (random variable), where $\text{Prob}(\mathcal{A}_{\gamma,t} = 0) = \gamma_t$ |
| $\mathbf{X} \odot \mathcal{A}_\gamma \in \mathbb{R}^T$ | Occluded trace (elementwise product of trace and occlusion mask) |
| $\Phi_\theta(\cdot \mathbf{X} \odot \mathcal{A}_\gamma, \mathcal{A}_\gamma) \in [0, 1]$ | Predicted distribution of secret data given occluded trace and mask, by classifier $w/$ weights θ |

67 In this work, we seek to leverage deep learning to *defend* against side-channel attacks by identifying specific
 68 points in time at which power/radiation measurements leak sensitive data. Our aim is to aid the designers of
 69 implementations in understanding *why* their implementations leak (e.g. by indicating the particular sections
 70 of code or electronic components which are responsible) and thereby enable targeted mitigation of the leaks.
 71 Our key contributions are:

- 72 • We propose a principled information theoretic quantity which measures the ‘leakiness’ of individual
 73 power/EM radiation measurements. Unlike prior approaches, ours is sensitive to arbitrary statistical
 74 associations between a chosen secret variable and subsets of measurements. Our quantity is implicitly
 75 defined through a constrained optimization problem.
- 76 • We propose a novel deep learning algorithm called Adversarial Leakage Localization (**ALL**) which
 77 approximately solves this optimization problem. **ALL** is based on an adversarial game played between
 78 a neural net ‘attacker’ trained to classify secret data using power/radiation traces, and a budget-
 79 constrained noise distribution trained to ‘defend’ against the attack by introducing noise to individual
 80 measurements in the traces. Due to the budget constraint, noise cannot be added everywhere and
 81 must be rationed for the leakiest measurements. After training we can thereby surmise the leakiness
 82 of each measurement from its noisiness.
- 83 • We compare **ALL** with 11 baseline methods on 6 publicly-available power and EM radiation side-
 84 channel leakage datasets from implementations of the AES, ECC, and RSA cryptographic standards.
 85 To our knowledge this is by far the most comprehensive and quantitative comparison of leakage
 86 localization algorithms which has been done, and we release our code and procedure in the hope of
 87 facilitating reproducibility and benchmarking of future work in this area.

88 2 Background and Setting

89 See Table 1 for a summary of the important notation we will use throughout the rest of the paper. Our
 90 setting is standard in the context of profiling power/EM side-channel analysis (Chari et al., 2003; Mangard
 91 et al., 2007) and is illustrated in Fig. 1. We assume to have a cryptographic device which encrypts data in a
 92 manner dependent on secret variable $y \in \mathbb{Y}$, where \mathbb{Y} is a finite set (e.g. consisting of bitstrings encoding
 93 all possible values of the variable). We assume to have some measurement apparatus that allows us to
 94 measure the power consumption or EM radiation throughout encryption. We view the resulting measurement
 95 sequences as vectors $\mathbf{x} \in \mathbb{R}^T$, called *traces*, where $T \in \mathbb{Z}_{++}$ denotes the number of measurements per trace.
 96 We view the secret variable as a random variable $Y \sim p_Y$ where p_Y is a simple (e.g. uniform) distribution
 97 under our control. The resulting trace $\mathbf{X} | Y \sim p_{\mathbf{X}|Y}$ then comes from a complicated and *a priori* unknown
 98 distribution dictated by factors such as the hardware, environment, and measurement setup. Here each
 99 element X_t of the random vector $\mathbf{X} \equiv (X_1, \dots, X_T)$ represents a power/radiation measurement at a fixed
 100 point in time relative to the start of encryption. In this work we assume that the distributions $p_{\mathbf{X},Y}(\cdot | y)$

101 exist and have full support for all $y \in \mathbb{Y}$, which is reasonable because empirically power consumption usually
 102 has a ‘random’ component which is well-described by additive Gaussian noise (Mangard et al., 2007). Most
 103 profiling side-channel attacks amount to collecting a dataset $\mathcal{D} \stackrel{\text{i.i.d.}}{\sim} p_{\mathbf{X},Y}$ and using supervised (deep or
 104 otherwise) learning to model $p_{Y|\mathbf{X}}$.

105 Given these jointly-distributed \mathbf{X}, Y , we seek to define for each X_t a scalar quantifying its ‘leakiness,’ i.e. the
 106 extent to which it can be exploited by attackers to learn Y from \mathbf{X} . Towards this end it is useful to consider
 107 the Shannon conditional mutual information (Shannon, 1948)

$$\mathbb{I}[Y; X_t | \mathcal{S}] := \mathbb{E} [\log p_{Y|X_t, \mathcal{S}}(Y | X_t, \mathcal{S}) - \log p_{Y|\mathcal{S}}(Y | \mathcal{S})], \quad \mathcal{S} \subseteq \{X_1, \dots, X_T\} \setminus \{X_t\}. \quad (1)$$

108 Intuitively, $\mathbb{I}[Y; X_t | \mathcal{S}]$ tells us the extent to which our uncertainty about Y is reduced upon observing X_t ,
 109 provided we have already observed the elements of \mathcal{S} .

110 Each individual quantity $\mathbb{I}[Y; X_t | \mathcal{S}]$ tells us something about the leakiness of X_t . However, for each t there
 111 are 2^{T-1} such quantities, and it is not obvious how they should be combined into a single scalar. Clearly,
 112 X_t should be considered leaky if $\mathbb{I}[Y; X_t] > 0$ and non-leaky if $\mathbb{I}[Y; X_t | \mathcal{S}] = 0 \forall \mathcal{S}$. More-subtly, there are
 113 many practical scenarios in which we have *second-order leakage*, where $\mathbb{I}[Y; X_t] = 0$ but $\mathbb{I}[Y; X_t | X_{t'}] > 0$ for
 114 some $t' \neq t$ (i.e. X_t *alone* reveals nothing about Y , but reveals something useful about Y *when combined*
 115 with $X_{t'}$). For example, many cryptographic implementations use a Boolean masking countermeasure (Chari
 116 et al., 1999; Benadjila et al., 2020) whereby a sensitive variable is decomposed into a pair of independent
 117 ‘random shares’ which are operated on at distant points in time. Hence, the naïve choice of $\mathbb{I}[Y; X_t]$ as our
 118 definition of the leakiness of X_t would not work.

119 Another naive choice would be to define the leakiness of X_t as $\mathbb{I}[Y; X_t | \{X_1, \dots, X_T\} \setminus \{X_t\}]$. This addresses
 120 the insensitivity of $\mathbb{I}[Y; X_t]$ to second-order leakage. However, it introduces a new shortcoming: when there
 121 are many leaky measurements with ‘redundant’ information, we may have $\mathbb{I}[Y; X_t | \{X_1, \dots, X_T\} \setminus \{X_t\}] \approx 0$
 122 even if $\mathbb{I}[Y; X_t] \gg 0$. In other words, X_t has little *new* information about Y which is not already provided by
 123 the other measurements. As we will show, this phenomenon creates issues for many prior deep learning-based
 124 leakage localization algorithms.

125 We will subsequently propose a natural notion of leakiness which is sensitive to all the quantities described
 126 by Eqn. 1. Before we do so, let us consider relevant prior work and its limitations.

127 3 Existing work and its limitations

128 We consider prior work in the side-channel analysis literature which may be leveraged for leakage localization.
 129 One prominent category of such work is *parametric statistics-based methods* which use non-deep learning
 130 techniques to look for pairwise associations between the measurements X_t and Y . The other is *neural net*
 131 *attribution-based methods*, where 1) a profiling side-channel attack is carried out with supervised deep learning,
 132 and 2) the neural net is ‘interpreted’ to determine the relative importance of its input features. Refer to
 133 Appendix C for further details.

134 3.1 First-order parametric statistics-based methods

135 First-order parametric statistics-based methods (Mangard et al., 2007; Chari et al., 2003; Brier et al., 2004)
 136 are widely used for understanding leakage due to their simplicity, interpretability, and low computational
 137 cost. However, these cannot detect leakage of order 2 or higher, and make restrictive assumptions about the
 138 relationship between \mathbf{X} and Y . Thus, such methods are ill-suited to our work’s ‘black-box’ leakage localization
 139 setting where we make minimal assumptions about the cryptographic implementation being evaluated.

140 In practice, leakage mitigation will likely be done in a ‘white-box’ setting where hardware designers understand
 141 the implementation and have access to its internal variables. In this setting, n^{th} -order leakage can often be
 142 decomposed into first-order leakage of $\geq n$ internal variables, which can then be localized individually with
 143 first-order methods. However, such analysis is error-prone and relies on careful analysis of the implementation.
 144 For example, Benadjila et al. (2020) released the second order-leaking ASCADv1-fixed dataset and analyzed
 145 its leakage by decomposing it into 2 pairs of first-order-leaking internal variables. Subsequently, Egger et al.

146 (2022) noted additional internal variables which contribute to leakage but were missed in the initial analysis.
 147 We view black-box deep learning-based methods such as ours as *complementary* with parametric white-box
 148 analysis: the latter provides an interpretable and hyperparameter-free assessment of known-leaky internal
 149 variables, while the former can detect leakage stemming from both known and unknown sources.

150 **3.2 Neural net attribution-based methods**

151 There is a great deal of prior work on localizing leakage by applying interpretability techniques to neural
 152 nets which have been trained to perform side-channel attacks (Masure et al., 2019; Hettwer et al., 2020; Jin
 153 et al., 2020; Zaid et al., 2020; Wouters et al., 2020; van der Valk et al., 2021; Golder et al., 2022; Li et al.,
 154 2022; Perin et al., 2022; Schamberger et al., 2023; Yap et al., 2023; Li et al., 2024; Yap et al., 2025). Most
 155 of these techniques can be summarized as follows: 1) use standard supervised deep learning techniques to
 156 train a model $\hat{p} \approx p_{Y|X}$ using data, and 2) use interpretability techniques to estimate the influence of each
 157 input feature on the model, on average over the dataset. For example, the Gradient Visualization (GradVis)
 158 technique of Masure et al. (2019) estimates the leakiness of X_t by $\mathbb{E}_{X,Y} \left| -\frac{\partial}{\partial x_t} \log \hat{p}(Y | X) \right|$, and the
 159 1-Occlusion technique of Hettwer et al. (2020) estimates it as $\mathbb{E}_{X,Y} |\hat{p}(Y | X) - \hat{p}(Y | (1 - I_t) \odot X)|$ where
 160 I_t denotes column t of the identity matrix. In this work we consider as baselines the recent OccPOI method
 161 (Yap et al., 2025), the m -Occlusion and 2nd-order m -Occlusion techniques (Schamberger et al., 2023), as
 162 well as GradVis (Masure et al., 2019), Saliency (Simonyan et al., 2014; Hettwer et al., 2020), 1-Occlusion
 163 (Zeiler & Fergus, 2014; Hettwer et al., 2020), LRP (Bach et al., 2015; Hettwer et al., 2020), and Input *
 164 Grad (Shrikumar et al., 2017; Wouters et al., 2020). Note that these subsume the deep learning baselines
 165 considered by Yap et al. (2025); Schamberger et al. (2023).

166 Most of these methods are prone to detecting only some leaking measurements while ignoring others.
 167 GradVis, Saliency, 1-Occlusion, LRP and Input * Grad compute the leakiness of X_t by occluding or
 168 differentiating the input x_t to $\hat{p}(Y | X_1, \dots, X_{t-1}, x_t, X_{t+1}, \dots, X_T)$ and observing its change in output.
 169 However, as discussed in Sec. 2, when many of the measurements carry ‘redundant’ information, one may
 170 have $\mathbb{I}[Y; X_t | \{X_1, \dots, X_T\} \setminus \{X_t\}] \approx 0$ even if $\mathbb{I}[Y; X_t] \gg 0$. In this case a well-fit \hat{p} becomes essentially
 171 constant with respect to x_t , causing these methods to spuriously estimate low leakiness for X_t .

172 While m -Occlusion, 2nd-order m -Occlusion, and OccPOI occlude multiple inputs simultaneously and are
 173 thus less-susceptible to this issue, they have their own shortcomings. m -Occlusion is like 1-Occlusion except
 174 that it occludes m -diameter windows rather than single points, which has an undesirable smoothing effect
 175 and is helpful only when the ‘redundant’ measurements are temporally-local. 2nd-order m -Occlusion entails
 176 occluding all pairs of windows, which is expensive because it requires $\Theta(T^2)$ passes through the dataset
 177 with the neural net. Additionally, while it provides the interesting and unique ability to discern whether
 178 leakage is first-order or second-order, we find that it gives little improvement over m -Occlusion when adapted
 179 to estimate the leakiness of a single point. Unlike the other considered methods, OccPOI does not assign
 180 a leakiness estimate to every measurement. Rather, it is a heuristic which aims to identify a non-unique
 181 minimal-cardinality set of measurements sufficient for \hat{p} to attain some classification performance when all
 182 other measurements are occluded. Additionally, it is expensive because it requires $\Omega(T)$ non-parallelizable
 183 passes through the dataset with the neural net.

184 **4 Our method: Adversarial Leakage Localization (ALL)**

185 Here we propose a novel algorithm called Adversarial Leakage Localization (ALL) for localizing leakage which
 186 addresses the shortcomings of prior work. In line with Sec. 2, we have jointly-distributed power/EM radiation
 187 traces $\mathbf{X} := (X_1, \dots, X_T)$ and secret data Y . We seek to define for each X_t a scalar quantifying its ‘leakiness,’
 188 i.e. its usefulness to attackers for learning Y from \mathbf{X} .

189 Intuitively, ALL is based on an adversarial game played between a neural net ‘attacker’ trained to predict Y
 190 from \mathbf{X} , and a budget-constrained noise distribution trained to ‘defend’ against the attack by introducing
 191 noise to individual measurements to maximize the loss of the classifier. Because of the budget constraint,
 192 increasing the noise applied to one measurement requires reducing the noise of other measurements. Thus,
 193 noise cannot simply be applied everywhere, and must be ‘triaged’ so that leakier measurements get more

194 noise. After training we can surmise the leakiness of a measurement from the amount of noise which has
 195 been applied to it.

196 In this section we first propose a constrained optimization problem which implicitly defines a notion of
 197 ‘leakiness’ which is sensitive to all the terms $\mathbb{I}[Y; X_t | S] : S \subseteq \{X_1, \dots, X_T\} \setminus \{X_t\}$. We then derive **ALL** as a
 198 practical deep learning algorithm which approximately solves this optimization problem. We conclude by
 199 explicitly contrasting **ALL** with prior work. Refer to Appendix D for an extended version of this section with
 200 proofs and derivations, and Algorithm 2 for pseudocode.

201 **4.1 Implicit definition of leakiness through a constrained optimization problem**

202 We define a vector $\gamma \in [0, 1]^T$ which we name the *occlusion probabilities*. γ parameterizes a distribution over
 203 a binary random vector with range $\{0, 1\}^T$ which we call the *occlusion mask*: $\mathcal{A}_\gamma \equiv (\mathcal{A}_{\gamma,1}, \dots, \mathcal{A}_{\gamma,T}) \sim p_{\mathcal{A}_\gamma} :$
 204 $\mathcal{A}_{\gamma,t} \sim \text{Bernoulli}(1 - \gamma_t)$. For arbitrary vectors $\mathbf{x} \in \mathbb{R}^T$, $\alpha \in \{0, 1\}^T$, let us denote $\mathbf{x}_\alpha := (x_t : t = 1, \dots, T :
 205 \alpha_t = 1)$, i.e. the sub-vector of \mathbf{x} containing its elements for which the corresponding element of α is 1. We
 206 can accordingly use \mathcal{A}_γ to obtain random sub-vectors $\mathbf{X}_{\mathcal{A}_\gamma}$ of \mathbf{X} . Note that γ_t denotes the probability that
 207 X_t will *not* be an element of $\mathbf{X}_{\mathcal{A}_\gamma}$.

208 We assign to each element of γ a strictly-increasing ‘cost’ $c : [0, 1] \rightarrow \mathbb{R}_+ : x \mapsto \frac{x}{1-x}$, with the properties
 209 $c(0) = 0$ and $\lim_{x \rightarrow 1} c(x) = \infty$. We seek to solve the constrained optimization problem

$$\min_{\gamma \in [0,1]^T} \quad \mathcal{L}(\gamma) := \mathbb{I}[Y; \mathbf{X}_{\mathcal{A}_\gamma} | \mathcal{A}_\gamma] \quad \text{such that} \quad \sum_{t=1}^T c(\gamma_t) = C \quad (2)$$

210 where $C \in \mathbb{R}_{++}$ is a ‘budget’ hyperparameter. Intuitively, the mutual information term tells us the extent
 211 to which the ‘occluded’ trace $\mathbf{X}_{\mathcal{A}_\gamma}$ ‘leaks’ Y , and γ is optimized to distribute a fixed ‘budget of occlusion
 212 probability’ among the individual elements of \mathbf{X} to minimize this leakage.

213 As discussed in Appendix D.1, during this optimization process each γ_t is ‘pushed’ up towards 1 in proportion
 214 to a weighted sum over *all* values $\mathbb{I}[Y; X_t | S] : S \subseteq \{X_1, \dots, X_T\} \setminus \{X_t\}$. Thus, **ALL** is sensitive to all
 215 associations between Y and subsets of $\{X_1, \dots, X_T\}$. This is in contrast to parametric methods which consider
 216 only pairwise associations between each X_t and Y , methods like 1-Occlusion, GradVis, Saliency, Input * Grad
 217 and LRP which are sensitive to associations between Y and the sets $\{X_1, \dots, X_T\}$, $\{X_1, \dots, X_T\} \setminus \{X_t\}$, and
 218 OccPOI, m -Occlusion, and 2^{nd} -order m -Occlusion, which consider larger yet still tiny subsets of the power
 219 set of $\{X_1, \dots, X_T\}$.

220 Due to the budget constraint, increasing γ_t requires reducing other γ_τ , $\tau \neq t$. Let us denote by γ^* a solution
 221 to Eqn. 2. Each γ_t^* will be closer to 1 if X_t is ‘leakier’ in the sense that it has greater mutual information
 222 with Y , conditioned on other X_τ , $\tau \neq t$. Thus, we propose using γ_t^* to measure the ‘leakiness’ of X_t .

223 **4.2 Deep learning-based implementation**

224 We will now re-frame this problem in a way that is amenable to standard deep learning techniques. Refer to
 225 Fig. 2 for a diagram.

226 We first re-parameterize it into an unconstrained problem by defining the variable $\eta := \text{softmax}(\tilde{\eta})$, $\tilde{\eta} \in \mathbb{R}^T$
 227 and letting γ be a function from $\mathbb{R}^T \rightarrow [0, 1]^T$ satisfying

$$c(\gamma_t(\tilde{\eta})) = C\eta_t \iff \gamma_t(\tilde{\eta}) = \text{sigmoid}(\log C + \log(\text{softmax}(\tilde{\eta}))) \quad (3)$$

228 We can now optimize with respect to $\tilde{\eta}$ instead of γ , letting us drop the constraint because it is satisfied for
 229 any $\tilde{\eta}$. Note that it is cheap and numerically-stable to map $\tilde{\eta}$ to γ in PyTorch.

230 Next, as described in Appendix D.2, we can approximate the mutual information term of Eqn. 2 with a
 231 neural net. Note that $\mathbb{I}[Y; \mathbf{X}_{\mathcal{A}_\gamma} | \mathcal{A}_\gamma] = \sum_{\alpha \in \{0,1\}^T} p_{\mathcal{A}_\gamma}(\alpha) \mathbb{I}[Y; \mathbf{X}_\alpha]$ where each $\mathbb{I}[Y; \mathbf{X}_\alpha] = \mathbb{E} \log p_{Y|\mathbf{X}_\alpha}(Y |
 232 \mathbf{X}_\alpha) - \mathbb{E} \log p_Y(Y)$. The right terms can be dropped because their corresponding terms in the full expression
 233 do not depend on γ . The conditional distributions in the left terms can each be approximated by using
 234 supervised deep learning with cross-entropy loss to classify Y from \mathbf{X}_α . There are 2^T such distributions and

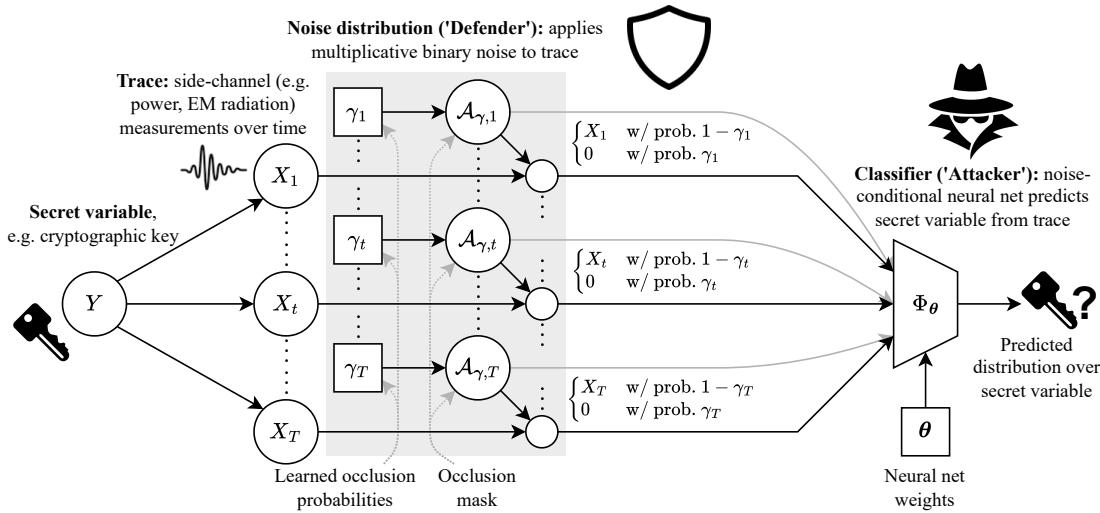


Figure 2: A diagram illustrating our Adversarial Leakage Localization (ALL) algorithm. A classifier Φ_θ is trained to ‘attack’ a cryptographic implementation by predicting its secret data Y from power/EM radiation traces $\mathbf{X} \equiv (X_1, \dots, X_T)$. Simultaneously, a noise distribution is trained to ‘defend against the attack’ by occluding the classifier’s individual input features X_t with probabilities γ_t , subject to a budget constraint which prevents trivially occluding every feature with probability 1. Because of the constraint, increasing γ_t necessarily entails decreasing γ_τ for some $\tau \neq t$, so the noise distribution must preferentially apply noise to leakier features. Thus, after training we may interpret γ_t as the ‘leakiness’ of X_t .

235 it would be infeasible to train this many neural nets independently. Instead, similarly to Lippe et al. (2022),
 236 we train a single neural net to estimate all the distributions by occluding its inputs according to \mathcal{A}_γ and
 237 feeding \mathcal{A}_γ as an auxiliary input. Thus, we can approximate Eqn. 2 with the optimization problem

$$\min_{\tilde{\eta} \in \mathbb{R}^T} \max_{\theta \in \mathbb{R}^P} \mathcal{L}_{\text{adv}}(\tilde{\eta}, \theta) := \mathbb{E} \log \Phi_\theta(Y | \mathbf{X} \odot \mathcal{A}_{\gamma(\tilde{\eta})}, \mathcal{A}_{\gamma(\tilde{\eta})}) \quad (4)$$

238 where $\Phi_\theta : \mathbb{Y} \times \mathbb{R}^T \times \{0, 1\}^T \rightarrow (0, 1)$ is a neural net with weights θ and softmax output activation, and
 239 $\Phi_\theta(y | \mathbf{x} \odot \alpha, \alpha)$ denotes its estimated probability that $Y = y$ given $\mathbf{X}_\alpha = \mathbf{x}_\alpha$. This can be approximately
 240 solved using alternating SGD-style algorithms, similarly to GANs (Goodfellow et al., 2014).

241 To use SGD-like algorithms we must estimate $\nabla_\theta \mathcal{L}_{\text{adv}}(\tilde{\eta}, \theta)$ and $\nabla_{\tilde{\eta}} \mathcal{L}_{\text{adv}}(\tilde{\eta}, \theta)$ with Monte Carlo integration.
 242 The former is routine in the context of DNN training. However, the latter is nontrivial because \mathcal{L}_{adv} has the
 243 form $\mathbb{E}_{\alpha \sim p_{\tilde{\eta}}} f(\alpha)$ where α is discrete. There is a large body of work on gradient estimation for functions of
 244 this nature, which can broadly be categorized into unbiased REINFORCE (Williams, 1992)-based estimators
 245 with variance reduction strategies, and biased estimators based on relaxing $p_{\tilde{\eta}}$ into a continuous distribution
 246 for which we can use the reparameterization trick (Rezende et al., 2014; Kingma & Welling, 2014). In our
 247 experiments we use the biased CONCRETE estimator (Maddison et al., 2017) with fixed temperature $\tau = 1$
 248 because it is cheap and simple, and we find it yields nearly the same performance as more-complicated
 249 estimators we tried. We conjecture that we get strong performance with this simple and biased estimator
 250 because our performance metrics are sensitive only to the *relative* leakiness assigned to measurements, and
 251 the bias does not significantly affect this. Note that $f(\alpha) \equiv \mathbb{E} \log \Phi_\theta(Y | \mathbf{X} \odot \alpha, \alpha)$ has been defined for
 252 $\alpha \in \{0, 1\}^T$, but after the CONCRETE relaxation we may have any $\alpha \in [0, 1]^T$. Thus, we replace it with the
 253 modified function $f(\alpha) \equiv \mathbb{E} \log \Phi_\theta(Y | \mathbf{X} \odot \alpha + \varepsilon \odot (\mathbf{1} - \alpha), \alpha)$ where $\varepsilon \sim \mathcal{N}(0, 1)^T$.

254 Our method is mainly sensitive to 3 hyperparameters: the learning rates of θ and $\tilde{\eta}$, and the budget
 255 hyperparameter C . Rather than tuning C directly, we find it easier to tune the hyperparameter $\bar{\gamma} := \frac{C}{C+T}$. $\bar{\gamma}$
 256 is equal to the occlusion probability of each measurement when $\tilde{\eta}$ is constant, and is less-sensitive to the data
 257 dimensionality T than C .

258 **4.3 Differences from prior work**

259 Whereas GradVis, Saliency, LRP, Input * Grad and 1-Occlusion effectively perturb single input features to
 260 the classifier and analyze the change in its outputs, ALL generally perturbs many inputs simultaneously. This
 261 is useful in settings where there are many ‘redundant’ leaking measurements and the impact of perturbing
 262 only one of them is small.

263 Like ALL, m -Occlusion, 2nd-order m -Occlusion and OccPOI also simultaneously occlude multiple input
 264 features. The key differences of our method are: 1) ALL samples from a distribution over all 2^T possible
 265 occlusion masks optimized to maximally hurt the performance of the classifier, whereas prior work iterates over
 266 a tiny subset of possible occlusion masks chosen heuristically. 2) ALL leverages the gradient of classifier loss
 267 with respect to relaxed occlusion masks, whereas prior work uses only zeroeth-order information from forward
 268 passes with ‘hard’ occlusion masks. 3) We simultaneously optimize the mask distribution to maximally hurt
 269 the classifier, and the classifier weights to be optimal for the current mask distribution. In contrast, prior
 270 work trains the classifier with standard supervised learning techniques, then ‘interprets’ the fixed classifier
 271 with occlusion.

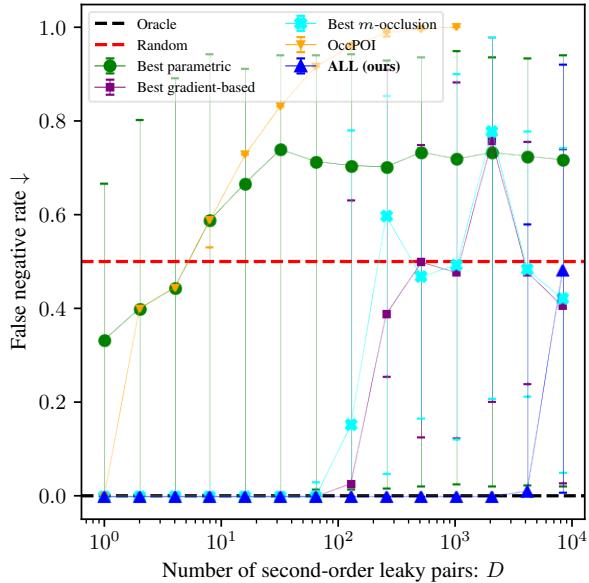
272 **5 Experimental results**273 **5.1 Synthetic datasets where we know ‘ground truth’ leakiness**

Figure 3: A toy setting described in Sec. 5.1 where ALL (ours) significantly outperforms baselines. We sample 1 non-leaky feature and D second-order leaky pairs, then plot the false negative rate, defined as the proportion of points incorrectly assigned leakiness less than or equal to that of the non-leaky point, as we increase D . ALL (ours) succeeds for D up to 64 \times higher the best prior deep learning-based approach, and the first-order parametric methods completely fail in this setting. Dots denote median and error bars denote min–max over 5 random seeds.

274 **Toy setting where ALL succeeds and prior work fails** As we will show, these differences lead to
 275 significant performance gains, as well as speedup relative to 2nd-order m -Occlusion and OccPOI.

276 We generate a sequence of binary-label $2D + 1$ -feature classification datasets consisting of ordered pairs (\mathbf{X}, Y)
 277 sampled independently as follows (see Appendix E.1 for details): $Y \sim \mathcal{U}\{0, 1\}$, $R \sim \mathcal{U}\{0, 1\}$, $M_i \sim \mathcal{U}\{0, 1\}$
 278 for $i = 1, \dots, D$, $X_R \sim \mathcal{N}(R, 1)$, $X_{M_i} \sim \mathcal{N}(M_i, 1)$, $X_{Y \oplus M_i} \sim \mathcal{N}(Y \oplus M_i, 1)$ for $i = 1, \dots, D$. Here we denote

279 by \oplus the exclusive-or operation and $\mathbf{X} := (X_R, X_{M_1}, X_{Y \oplus M_1}, \dots, X_{M_D}, X_{M_D \oplus Y})$. Intuitively, we can view
 280 the features X_R , X_{M_i} , and $X_{Y \oplus M_i}$ as noisy observations of R , M_i , and $Y \oplus M_i$, respectively. Note that R
 281 tells us nothing about Y , and while *individually* each variable M_i , $Y \oplus M_i$ is independent of Y , the *pairs* of
 282 variables $\{M_i, Y \oplus M_i\}$ allow us to determine Y through the identity $Y = Y \oplus M_i \oplus M_i$. An ideal leakage
 283 localization algorithm should indicate that each feature X_{M_i} , $X_{Y \oplus M_i}$ is leaking, while X_R is not.

284 In Fig. 3, we plot the performance of ALL and prior work as we sweep the value of D . While prior deep
 285 learning-based methods succeed for small D , they fail when D grows large because the individual contribution
 286 of each $\{X_{M_i}, X_{Y \oplus M_i}\}$ is ‘drowned out’ in the sense that $\mathbb{I}[Y; \{X_{M_i}, X_{Y \oplus M_i}\} \mid \{X_{M_j}, X_{Y \oplus M_j} : j \neq i\}]$
 287 vanishes. ALL succeeds for D up to 64 \times larger than prior work because as the classifier becomes more-
 288 reliant on particular features they are subject to a higher occlusion probability, mitigating this effect. For
 289 completeness we also include first-order parametric methods in our comparison; these fail because they are
 290 not sensitive to the second-order associations between Y and $\{X_{M_i}, X_{Y \oplus M_i}\}$.

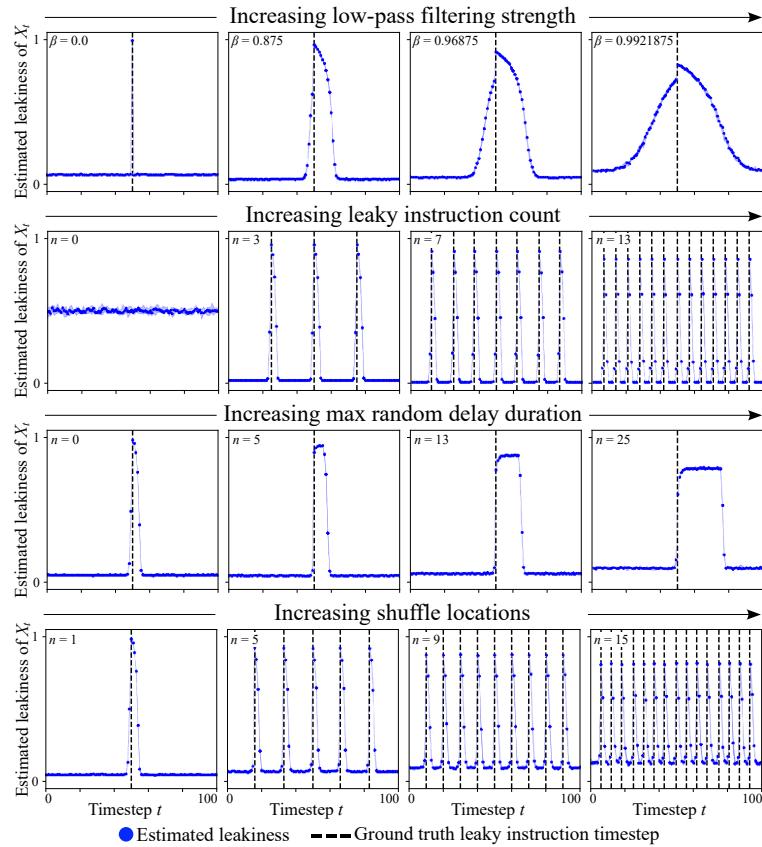


Figure 4: Output $\gamma^* \equiv \gamma(\tilde{\eta}^*)$ of our ALL algorithm when applied to simulated AES-128 power trace datasets based on the Hamming weight model of (Mangard et al., 2007, ch. 4), as described in Appendix E.2. Leakage estimated by ALL is consistent with the ground truth timestep at which leaky instruction(s) are executed across varying low-pass filtration strength, leaky instruction count, random no-op insertion, and random shuffling.

291 **Simulated AES-128 datasets** We next apply ALL to a variety of simulated AES datasets. These are
 292 a useful complement to subsequent experiments on real datasets because we can validate ALL’s outputs
 293 against ground truth knowledge about which timesteps are leaking, as well as gain insight into its behavior
 294 when individually varying particular dataset properties. Traces are simulated using the Hamming weight
 295 leakage model of (Mangard et al., 2007, ch. 4), which decomposes total power consumption as $X_t =$
 296 $X_{\text{data}, t} + X_{\text{op}, t} + X_{\text{resid}, t}$. Here $X_{\text{data}, t}$ is a function of the Hamming weight of the data currently being
 297 operated on, $X_{\text{op}, t}$ is a function of the operation currently being applied to the data, and $X_{\text{resid}, t}$ models

remaining sources of noise as a Gaussian random variable. Further details can be found in Appendix E.2. Note that while this model applies to a particular device studied by Mangard et al. (2007), the relationship between data and power consumption is device-dependent and often eludes simple characterization.

In Fig. 4 we simulate several factors of variation which may be expected to occur in realistic settings, and observe the change in behavior of ALL. *First*, we apply a varying-strength discrete low-pass filter to the simulated traces. As the strength increases, the peak of the estimated leakiness remains centered at the timestep of the leaky instruction while becoming more-diffuse. *Second*, we vary the number of leaky instructions and see that ALL consistently produces similar-height peaks at every leaky instruction. *Third*, we introduce a random delay before the leaky instruction, which causes ALL’s peak diffuses over the set of timesteps at which the instruction may occur. *Fourth*, we shuffle the location of the leaky instruction so that it may occur at various points in time. As when varying the number of instructions, ALL produces similar-height peaks at each point in time at which the instruction sometimes occurs.

5.2 Real power and EM radiation leakage datasets

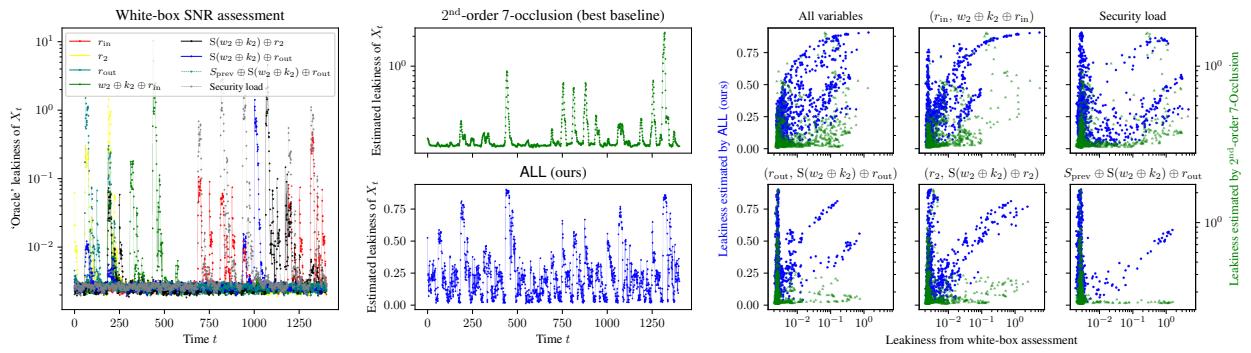


Figure 5: A qualitative comparison on ASCADv1-variable between the estimated leakiness by a white-box SNR-based assessment following Egger et al. (2022), ALL (our method), and 2nd-order 7-Occlusion (the strongest baseline). **(left)** Superimposed per-timestep leakiness of 8 internal AES variables which contribute to leakage of our targeted ‘secret variable’ $Y \equiv S(w_2 \oplus k_w)$. We consider Y to leak to the extent that at least one of these internal variables leaks. **(center)** The per-timestep leakiness of Y , as estimated by ALL and the baseline. Ideally, these estimates should align with peaks in the white-box assessment. **(right)** The per-timestep leakiness as estimated by ALL and the baseline, vs. the mean SNR of subsets of the internal AES variables (indicated in plot titles). In the plot labeled ‘All variables’, ALL exhibits a stronger positive association with the aggregate white-box assessment than the baseline. For individual variable subsets, ALL consistently produces a $\|$ -shaped structure with a diagonal trend indicating agreement between estimated leakiness and SNR for the variables of interest, and a vertical band corresponding to leakage from other variables. In contrast, the baseline often exhibits an L-shaped structure, with the horizontal band reflecting spuriously low leakiness estimates at timesteps where the internal variables are known to have significant leakage. It appears that the baseline is mainly sensitive to $(r_{in}, w_2 \oplus k_2 \oplus r_{in})$ and possibly the security load (which leaks at similar times as r_{in}), but understates or misses leakage due to the other variables.

We compare our method to prior work on 6 publicly-available datasets of real recorded side-channel emissions and metadata covering diverse settings, as described in Appendix E.3.1. See Appendix E.3 for an extended version of this section including full implementation details, additional experiments, and ablation studies.

Experimental setup Despite their methodological differences, all considered methods may be viewed as a function which maps a dataset to a sequence of scalars encoding the estimated leakiness of X_t for each t . For all deep learning baselines apart from ALL, we first train a supervised classifier to predict secret variables Y from power traces (X_1, \dots, X_T) . We then use the ‘importance’ of X_t to the classifier’s prediction about Y , on average over the profiling set, as the estimated leakiness of X_t . The definition of ‘importance’ is prescribed

319 by the method. For **ALL**, we simply run our method as described in Sec. 4, then directly use the trained
 320 occlusion probability γ_t^* as the estimated leakiness of X_t .
 321 For the supervised classifiers and the classifier denoted Φ_θ in **ALL**, we use the same ReLU MLP architecture
 322 trained with the AdamW optimizer (see Appendix E.3.2). For both supervised classification and **ALL** training
 323 runs, we use the same minibatch size and training step count, tune the important hyperparameters for each
 324 dataset using random search with a 50-run budget, and leave the rest at reasonable defaults (see Appendix
 325 E.3.3). Supervised classifier hyperparameters are tuned to minimize correct-key rank (similar to maximizing
 326 accuracy). For **ALL** we cannot use this approach, so we instead use a criterion based on occluding the inputs
 327 to a frozen supervised classifier and observing its change in performance (see Appendix E.3.5). Baselines are
 328 implemented using Captum (Kokhlikyan et al., 2020) where possible. OccPOI and (2nd-order) m -Occlusion
 329 are adapted to the setting of our paper as described in Appendix E.3.2.

330 **Performance evaluation strategies and results** In Fig. 5, in line with previous work (Masure et al.,
 331 2019; Wouters et al., 2020; Schamberger et al., 2023; Yap et al., 2025), we compute a white-box leakage
 332 assessment on ASCADv1-variable and qualitatively compare it to the outputs of **ALL** and the strongest
 333 baseline. As is standard, we target the variable $Y \equiv S(w_2 \oplus k_2)$. Because this AES implementation employs
 334 Boolean masking, Y does not directly influence the power consumption. Nonetheless, there are 3 pairs of
 335 internal variables which *do* directly influence power consumption and can be combined to determine Y (as
 336 well as more-complicated identities involving the security load and $S_{\text{prev}} \oplus S(w_2 \oplus k_2) \oplus r_{\text{out}}$) (Egger et al.,
 337 2022). Ideally, a leakage localization algorithm should consider each X_t leaky to the extent that at least one
 338 such variable with utility for determining Y leaks. In Fig. 5 **ALL** successfully identifies the leakage of all 3
 339 pairs of variables, while the best considered baseline clearly identifies leakage from only one of the pairs.
 340 In the interest of scaling our comparison to many baselines and datasets, we next consider quantitative
 341 performance metrics. For real-world datasets quantitative performance evaluation is challenging because we
 342 lack ‘ground truth’ knowledge about leakage, and there is no consensus on the best way to do so. We thus
 343 employ 4 performance evaluation metrics which are conceptually similar to evaluation strategies in prior work.
 344 Because the specific numbers used to encode leakiness by the various baselines are not directly comparable,
 345 we use metrics which are sensitive only to the *relative* leakiness of different timesteps – i.e. under which
 346 a vector of estimated leakiness values $(\gamma_1^*, \dots, \gamma_T^*)$ has the same performance as $(f(\gamma_1^*), \dots, f(\gamma_T^*))$ for any
 347 strictly-increasing $f : \mathbb{R} \rightarrow \mathbb{R}$.
 348 Table 2 lists performance across four complementary metrics which emphasize different aspects of performance.
 349 See Appendix E.3.4 for details about each metric. The *oracle agreement* metric assesses agreement with a
 350 white-box leakiness evaluation, and is given by the Spearman rank correlation coefficient between estimated
 351 leakiness and the average SNR of known leaky internal variables (similar to Fig. 5). The *template attack*
 352 *minimum traces to disclosure (MTD)*² metric assesses the utility of the estimated-leakiest measurements
 353 for conducting a side-channel attack – higher-fidelity leakiness estimates are expected to allow leakier
 354 measurements to be fed to the attacker, leading to better attack performance. The *forward DNN occlusion*
 355 *test*² measures the extent to which the performance of a deep learning-based side-channel attack is preserved
 356 as we occlude all but the estimated-leakiest measurements, and is mainly sensitive to spurious or incorrectly-
 357 ordered high estimates. The *reverse DNN occlusion test*² measures the extent to which performance is
 358 preserved as we occlude all but the estimated-least-leaky measurements, and is mainly sensitive to spurious
 359 low estimates. As seen in Table 2, **ALL** outperforms all baselines on the majority of datasets under every
 360 metric except for the forward DNN occlusion test.

361 6 Conclusion

362 We have proposed a novel algorithm for localizing side-channel leakage from cryptographic implementations.
 363 Unlike prior work, ours is sensitive to arbitrary statistical associations responsible for leakage and operates
 364 in a ‘black box’ manner, requiring only a supervised learning-style dataset with labels denoting the secret

²The template attack MTD metric is similar to evaluations done in Masure et al. (2019); Yap et al. (2025). The DNN occlusion tests were inspired by the ZB-KGE and KRPC algorithms of Hettwer et al. (2020). As described in Appendix E.3.4, we alter implementation details to summarize these as scalars and improve reliability on second-order datasets.

Table 2: A comparison of the considered deep learning-based leakage localization algorithms on 6 datasets and under 4 quantitative performance metrics. We report mean \pm 1 standard deviation over 5 random seeds. **Bold numbers** denote the method with the best mean performance (ties broken by lower variance). \uparrow denotes that a higher number is better, and \downarrow denotes that a lower number is better. **(Oracle agreement \uparrow)** The Spearman rank correlation coefficient between the predicted leakiness and an ‘oracle’ leakiness derived from a white-box SNR-based assessment exploiting knowledge of the internal variables contributing to leakage of the target. **(Template attack MTD \downarrow)** The minimum traces to disclosure (MTD) of a template attack when using the predicted leakiness for point of interest (feature) selection. **(Forward DNN occlusion test \downarrow)** The mean single-trace correct-key rank of a DNN-based attacker when occluding all but the k predicted-leakiest timesteps, on average for $k \in \{1, \dots, T\}$. **(Reverse DNN occlusion test \uparrow)** The mean single-trace correct-key rank of a DNN-based attacker when occluding all but the k predicted-least-leaky timesteps, on average for $k \in \{1, \dots, T\}$.

| Method | 2nd-order datasets | | 1st-order datasets | | | |
|---------------------------------|---|-------------------------------------|-------------------------------------|-------------------------------------|-----------------------------------|-------------------------------------|
| | ASCADv1-f | ASCADv1-r | DPAv4 | AES-HD | OTiAiT | OTP |
| Oracle agreement \uparrow | GradVis | 0.48 \pm 0.02 | 0.27 \pm 0.01 | 0.198 \pm 0.009 | 0.07 \pm 0.01 | 0.55 \pm 0.05 |
| | Saliency | 0.47 \pm 0.02 | 0.26 \pm 0.01 | 0.198 \pm 0.008 | 0.07 \pm 0.01 | 0.67 \pm 0.06 |
| | Input * Grad | 0.47 \pm 0.02 | 0.25 \pm 0.01 | 0.202 \pm 0.009 | 0.08 \pm 0.02 | 0.71 \pm 0.05 |
| | LRP | 0.47 \pm 0.02 | 0.25 \pm 0.01 | 0.202 \pm 0.009 | 0.08 \pm 0.02 | 0.71 \pm 0.05 |
| | OccPOI | 0.07 \pm 0.01 | 0.064 \pm 0.004 | 0.030 \pm 0.008 | 0.044 \pm 0.009 | 0.07 \pm 0.02 |
| | 1-Occlusion | 0.47 \pm 0.02 | 0.25 \pm 0.01 | 0.202 \pm 0.009 | 0.08 \pm 0.01 | 0.71 \pm 0.05 |
| | m^* -Occlusion | 0.49 \pm 0.02 | 0.41 \pm 0.01 | 0.32 \pm 0.01 | 0.18 \pm 0.05 | 0.72 \pm 0.04 |
| | 2 nd -order 1-Occlusion | 0.51 \pm 0.01 | 0.27 \pm 0.01 | 0.206 \pm 0.009 | 0.08 \pm 0.01 | 0.74 \pm 0.05 |
| | 2 nd -order m^* -Occlusion | 0.52 \pm 0.01 | 0.42 \pm 0.01 | 0.330 \pm 0.009 | 0.19 \pm 0.05 | 0.75 \pm 0.04 |
| | ALL (ours) | 0.794 \pm 0.006 | 0.60 \pm 0.01 | 0.317 \pm 0.002 | 0.22 \pm 0.03 | 0.782 \pm 0.001 |
| Tmpl. attack MTD \downarrow | GradVis | 686 \pm 100 | 1162 \pm 1000 | 2.7 \pm 0.1 | 20014 \pm 6000 | 1.4 \pm 0.3 |
| | Saliency | 726 \pm 100 | 1412 \pm 2000 | 2.7 \pm 0.1 | 19438 \pm 6000 | 1.14 \pm 0.02 |
| | Input * Grad | 675 \pm 100 | 1194 \pm 2000 | 2.6 \pm 0.1 | 19893 \pm 6000 | 1.14 \pm 0.02 |
| | LRP | 675 \pm 100 | 1194 \pm 2000 | 2.6 \pm 0.1 | 19893 \pm 6000 | 1.14 \pm 0.02 |
| | OccPOI | 787 \pm 100 | 942 \pm 200 | 71 \pm 30 | 25000.000 | 1.08 \pm 0.03 |
| | 1-Occlusion | 667 \pm 100 | 1376 \pm 2000 | 2.65 \pm 0.08 | 20011 \pm 6000 | 1.14 \pm 0.02 |
| | m^* -Occlusion | 673 \pm 70 | 727 \pm 400 | 9 \pm 1 | 16283 \pm 10 | 1.17 \pm 0.02 |
| | 2 nd -order 1-Occlusion | 709 \pm 100 | 1086 \pm 1000 | 2.65 \pm 0.08 | 20222 \pm 6000 | 1.14 \pm 0.02 |
| | 2 nd -order m^* -Occlusion | 642 \pm 60 | 710 \pm 400 | 9 \pm 1 | 16033 \pm 700 | 1.16 \pm 0.03 |
| | ALL (ours) | 459 \pm 40 | 394 \pm 20 | 2.22 \pm 0.01 | 17582 \pm 5000 | 1.11 \pm 0.02 |
| Fwd. DNN occlusion \downarrow | GradVis | 108.6 \pm 0.5 | 96.8 \pm 0.3 | 9.5 \pm 0.6 | 125.6 \pm 0.3 | 1.9 \pm 0.2 |
| | Saliency | 108.5 \pm 0.4 | 96.3 \pm 0.4 | 9.5 \pm 0.7 | 125.6 \pm 0.3 | 1.8 \pm 0.1 |
| | Input * Grad | 108.5 \pm 0.4 | 96.8 \pm 0.4 | 9.4 \pm 0.7 | 125.6 \pm 0.3 | 1.7 \pm 0.2 |
| | LRP | 108.5 \pm 0.4 | 96.8 \pm 0.4 | 9.4 \pm 0.7 | 125.6 \pm 0.3 | 1.7 \pm 0.2 |
| | OccPOI | 122.3 \pm 0.8 | 120.8 \pm 0.2 | 58 \pm 2 | 127.4 \pm 0.3 | 2.6 \pm 0.2 |
| | 1-Occlusion | 108.5 \pm 0.4 | 96.7 \pm 0.4 | 9.4 \pm 0.7 | 125.6 \pm 0.3 | 1.7 \pm 0.2 |
| | m^* -Occlusion | 108.2 \pm 0.5 | 95.7 \pm 0.6 | 9.0 \pm 0.6 | 125.3 \pm 0.2 | 1.8 \pm 0.2 |
| | 2 nd -order 1-Occlusion | 108.4 \pm 0.4 | 97.0 \pm 0.4 | 9.4 \pm 0.7 | 125.5 \pm 0.3 | 1.7 \pm 0.2 |
| | 2 nd -order m^* -Occlusion | 108.2 \pm 0.5 | 95.9 \pm 0.6 | 9.0 \pm 0.5 | 125.3 \pm 0.2 | 1.7 \pm 0.2 |
| | ALL (ours) | 107.3 \pm 0.4 | 104 \pm 1 | 9.9 \pm 0.7 | 125.5 \pm 0.3 | 1.8 \pm 0.1 |
| Rev. DNN occlusion \uparrow | GradVis | 125.9 \pm 0.2 | 127.6 \pm 0.1 | 122 \pm 1 | 128.0 \pm 0.3 | 4.2 \pm 0.4 |
| | Saliency | 125.8 \pm 0.2 | 127.4 \pm 0.2 | 122 \pm 1 | 128.0 \pm 0.3 | 5.1 \pm 0.3 |
| | Input * Grad | 125.7 \pm 0.3 | 127.5 \pm 0.2 | 121.9 \pm 0.9 | 128.1 \pm 0.3 | 5.2 \pm 0.3 |
| | LRP | 125.7 \pm 0.3 | 127.5 \pm 0.2 | 121.9 \pm 0.9 | 128.1 \pm 0.3 | 5.2 \pm 0.3 |
| | OccPOI | 122.3 \pm 0.4 | 124.6 \pm 0.2 | 43 \pm 1 | 127.0 \pm 0.3 | 3.6 \pm 0.3 |
| | 1-Occlusion | 125.8 \pm 0.3 | 127.4 \pm 0.2 | 122 \pm 1 | 128.1 \pm 0.3 | 5.2 \pm 0.3 |
| | m^* -Occlusion | 126.0 \pm 0.2 | 127.4 \pm 0.2 | 121 \pm 1 | 128.5 \pm 0.2 | 5.3 \pm 0.2 |
| | 2 nd -order 1-Occlusion | 125.8 \pm 0.3 | 127.5 \pm 0.2 | 122.0 \pm 0.9 | 128.1 \pm 0.3 | 5.3 \pm 0.2 |
| | 2 nd -order m^* -Occlusion | 126.1 \pm 0.2 | 127.4 \pm 0.2 | 121.3 \pm 0.9 | 128.5 \pm 0.2 | 5.3 \pm 0.2 |
| | ALL (ours) | 126.4 \pm 0.2 | 127.96 \pm 0.06 | 125 \pm 1 | 128.3 \pm 0.2 | 5.6 \pm 0.2 |

365 variable under consideration. In light of the ever-increasing efficacy of deep side-channel attack algorithms
 366 and the failure of existing work to detect all leakage they may exploit, our work marks a critical step towards
 367 understanding and mitigating the emerging vulnerabilities of cryptographic hardware.

368 **Broader Impact Statement**

369 The goal of our work is to enhance the security of cryptographic implementations against side-channel attacks
 370 by identifying the points in time at which they reveal sensitive information, thereby facilitating targeted
 371 defenses and mitigation strategies. We foresee ALL as being a useful complement to widely-adopted parametric
 372 leakage localization techniques due to its ability to identify leakage in a ‘black box’ manner without being
 373 limited by the domain knowledge of its users.

374 Our work is primarily defensive in nature. We do not introduce strategies that directly improve the
 375 performance of profiling side-channel attacks, and we solely consider cryptographic datasets which have been
 376 made available for research purposes and are already widely studied and understood. Nonetheless, improving
 377 the ability to identify and understand the weaknesses of cryptographic systems could potentially benefit
 378 attackers as well as defenders. We believe the utility of our work for defense outweighs this risk.

379 **References**

380 Dakshi Agrawal, Josyula R Rao, Pankaj Rohatgi, and Kai Schramm. Templates as master keys. In
*381 Cryptographic Hardware and Embedded Systems—CHES 2005: 7th International Workshop, Edinburgh, UK,
 382 August 29–September 1, 2005. Proceedings 7*, pp. 15–29. Springer, 2005.

383 Cédric Archambeau, Eric Peeters, F X Standaert, and J J Quisquater. Template attacks in principal subspaces.
 384 In *Cryptographic Hardware and Embedded Systems—CHES 2006: 8th International Workshop, Yokohama,
 385 Japan, October 10–13, 2006. Proceedings 8*, pp. 1–14. Springer, 2006.

386 Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and
 387 Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance
 388 propagation. *PLoS one*, 10(7):e0130140, 2015.

389 Ryad Benadjila, Emmanuel Prouff, Rémi Strullu, Eleonora Cagli, and Cécile Dumas. Deep learning for
 390 side-channel analysis and introduction to ASCAD database. *Journal of Cryptographic Engineering*, 10(2):
 391 163–188, 2020.

392 Shivam Bhasin, Nicolas Bruneau, Jean-Luc Danger, Sylvain Guillet, and Zakaria Najm. Analysis and
 393 improvements of the dpa contest v4 implementation. In *Security, Privacy, and Applied Cryptography
 394 Engineering: 4th International Conference, SPACE 2014, Pune, India, October 18–22, 2014. Proceedings 4*,
 395 pp. 201–218. Springer, 2014.

396 Shivam Bhasin, Dirmanto Jap, and Stjepan Picek. AES HD dataset - 50 000 traces. AISyLab repository,
 397 2020. https://github.com/AISyLab/AES_HD.

398 Bluefin Payment Systems. Bluefin and ID TECH partner to deliver PCI validated Advanced Encryption Stan-
 399 dard (AES) P2PE solution. Online, November 2023. URL [https://www.bluefin.com/news/bluefin-and-id-tech-partner-to-deliver-pci-validated-advanced-encryption-standard-aes-p2pe-solution/](https://www.bluefin.com/news/bluefin-

 400-and-id-tech-partner-to-deliver-pci-validated-advanced-encryption-standard-aes-p2pe-

 401-solution/).

402 Eric Brier, Christophe Clavier, and Francis Olivier. Correlation power analysis with a leakage model. In
 403 Marc Joye and Jean-Jacques Quisquater (eds.), *Cryptographic Hardware and Embedded Systems - CHES
 404 2004*, pp. 16–29, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg. ISBN 978-3-540-28632-5.

405 Olivier Bronchain and François-Xavier Standaert. Side-channel countermeasures’ dissection and the limits of
 406 closed source security evaluations. *IACR Transactions on Cryptographic Hardware and Embedded Systems*,
 407 pp. 1–25, 2020.

408 Elie Bursztein, Luca Invernizzi, Karel Král, Daniel Moghimi, Jean-Michel Picod, and Marina Zhang. Generic
 409 attacks against cryptographic hardware through long-range deep learning. *arXiv preprint arXiv:2306.07249*,
 410 2023.

411 Suresh Chari, Charanjit S Jutla, Josyula R Rao, and Pankaj Rohatgi. Towards sound approaches to counteract
 412 power-analysis attacks. In *Advances in Cryptology—CRYPTO’99: 19th Annual International Cryptology*
 413 *Conference Santa Barbara, California, USA, August 15–19, 1999 Proceedings 19*, pp. 398–412. Springer,
 414 1999.

415 Suresh Chari, Josyula R. Rao, and Pankaj Rohatgi. Template attacks. In Burton S. Kaliski, çetin K. Koç,
 416 and Christof Paar (eds.), *Cryptographic Hardware and Embedded Systems - CHES 2002*, pp. 13–28, Berlin,
 417 Heidelberg, 2003. Springer Berlin Heidelberg. ISBN 978-3-540-36400-9.

418 Committee on National Security Systems. Committee on National Security Systems Policy No. 15, Fact
 419 Sheet No. 1, June 2003. URL <https://csrc.nist.gov/csrc/media/projects/cryptographic-module-validation-program/documents/cnss15fs.pdf>.

421 Jean-Sébastien Coron and Ilya Kizhvatov. An efficient method for random delay generation in embedded
 422 software. In *International Workshop on Cryptographic Hardware and Embedded Systems*, pp. 156–170.
 423 Springer, 2009. Dataset available at <https://github.com/ikizhvatov/randomdelays-traces>.

424 Joan Daemen and Vincent Rijmen. AES proposal: Rijndael document version 2. AES Algorithm
 425 Submission, September 1999. URL <https://csrc.nist.gov/csrc/media/projects/cryptographic-standards-and-guidelines/documents/aes-development/rijndael-ammended.pdf>.

427 Joan Daemen and Vincent Rijmen. *The Design of Rijndael*. Springer Berlin, Heidelberg, March 2013. URL
 428 <https://link.springer.com/book/10.1007/978-3-662-04722-4>.

429 Josef Danial, Debayan Das, Anupam Golder, Santosh Ghosh, Arijit Raychowdhury, and Shreyas Sen. Em-x-dl:
 430 Efficient cross-device deep learning side-channel attack with noisy em signatures. *ACM Journal on Emerging*
 431 *Technologies in Computing Systems (JETC)*, 18(1):1–17, 2021.

432 Debayan Das, Anupam Golder, Josef Danial, Santosh Ghosh, Arijit Raychowdhury, and Shreyas Sen. X-
 433 deepsca: Cross-device deep learning side channel attack. In *Proceedings of the 56th Annual Design*
 434 *Automation Conference 2019*, pp. 1–6, 2019.

435 Maximilian Egger, Thomas Schamberger, Lars Tebelmann, Florian Lippert, and Georg Sigl. A second look
 436 at the ASCAD databases. In *International Workshop on Constructive Side-Channel Analysis and Secure*
 437 *Design*, pp. 75–99. Springer, 2022.

438 Guangjun Fan, Yongbin Zhou, Hailong Zhang, and Dengguo Feng. How to choose interesting points for
 439 template attacks? *Cryptology ePrint Archive*, Paper 2014/332, 2014. URL <https://eprint.iacr.org/2014/332>.

441 Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge,
 442 and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):
 443 665–673, 2020.

444 Daniel Genkin, Adi Shamir, and Eran Tromer. RSA key extraction via low-bandwidth acoustic cryptanalysis.
 445 In Juan A. Garay and Rosario Gennaro (eds.), *Advances in Cryptology – CRYPTO 2014*, pp. 444–461,
 446 Berlin, Heidelberg, 2014. Springer Berlin Heidelberg. ISBN 978-3-662-44371-2.

447 Daniel Genkin, Lev Pachmanov, Itamar Pipman, Eran Tromer, and Yuval Yarom. ECDSA key extraction
 448 from mobile devices via nonintrusive physical side channels. In *Proceedings of the 2016 ACM SIGSAC*
 449 *Conference on Computer and Communications Security*, pp. 1626–1638, 2016.

450 Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks.
 451 In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 249–256.
 452 JMLR Workshop and Conference Proceedings, 2010.

453 Anupam Golder, Ashwin Bhat, and Arijit Raychowdhury. Exploration into the explainability of neural
 454 network models for power side-channel analysis. In *Proceedings of the Great Lakes Symposium on VLSI*
 455 *2022, GLSVLSI ’22*, pp. 59–64, New York, NY, USA, 2022. Association for Computing Machinery. ISBN
 456 9781450393225. doi: 10.1145/3526241.3530346. URL <https://doi.org/10.1145/3526241.3530346>.

457 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron
 458 Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing*
 459 *systems*, 27, 2014.

460 Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *International Conference*
 461 *on Learning Representations*, 2021. URL <https://openreview.net/forum?id=1QdXeXDoWtI>.

462 Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David
 463 Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus,
 464 Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark
 465 Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer
 466 Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):
 467 357–362, September 2020. doi: 10.1038/s41586-020-2649-2. URL <https://doi.org/10.1038/s41586-020-2649-2>.

468 Katherine Hermann and Andrew Lampinen. What shapes feature representations? exploring datasets,
 469 architectures, and training. *Advances in Neural Information Processing Systems*, 33:9995–10006, 2020.

471 Benjamin Hettwer, Stefan Gehrer, and Tim Güneysu. Deep neural network attribution methods for leakage
 472 analysis and symmetric key recovery. In Kenneth G. Paterson and Douglas Stebila (eds.), *Selected*
 473 *Areas in Cryptography – SAC 2019*, pp. 645–666, Cham, 2020. Springer International Publishing. ISBN
 474 978-3-030-38471-5.

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

477 Michael Hutter and Jörn-Marc Schmidt. The temperature side channel and heating fault attacks. In *Smart*
 478 *Card Research and Advanced Applications: 12th International Conference, CARDIS 2013, Berlin, Germany,*
 479 *November 27-29, 2013. Revised Selected Papers 12*, pp. 219–235. Springer, 2014.

480 Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparametrization with gumble-softmax. In *International*
 481 *Conference on Learning Representations (ICLR 2017)*. OpenReview. net, 2017.

482 Minhui Jin, Mengce Zheng, Honggang Hu, and Nenghai Yu. An enhanced convolutional neural network in
 483 side-channel attacks and its visualization. *arXiv preprint arXiv:2009.08898*, 2020.

484 Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun
 485 (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April*
 486 *14-16, 2014, Conference Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6114>.

487 Paul Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In Michael Wiener (ed.), *Advances*
 488 *in Cryptology — CRYPTO’ 99*, pp. 388–397, Berlin, Heidelberg, 1999. Springer Berlin Heidelberg. ISBN
 489 978-3-540-48405-9.

490 Paul Kocher, Jann Horn, Anders Fogh, , Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz
 491 Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom. Spectre attacks: Exploiting
 492 speculative execution. In *40th IEEE Symposium on Security and Privacy (S&P’19)*, 2019.

493 Paul C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In Neal
 494 Koblitz (ed.), *Advances in Cryptology — CRYPTO ’96*, pp. 104–113, Berlin, Heidelberg, 1996. Springer
 495 Berlin Heidelberg. ISBN 978-3-540-68697-2.

496 Narine Kokhlikyan, Vivek Miglani, Miguel Martin, Edward Wang, Bilal Alsallakh, Jonathan Reynolds,
 497 Alexander Melnikov, Natalia Kliushkina, Carlos Araya, Siqi Yan, et al. Captum: A unified and generic
 498 model interpretability library for pytorch. *arXiv preprint arXiv:2009.07896*, 2020.

499 Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. Numba: A llvm-based python jit compiler. In *Proceedings*
 500 *of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pp. 1–6, 2015.

501 Yanbin Li, Yuxin Huang, Fuwei Jia, Qingsong Zhao, Ming Tang, and Shougang Ren. A gradient deconvolutional
 502 network for side-channel attacks. *Computers & Electrical Engineering*, 98:107686, 2022.

503 Yanbin Li, Jiajie Zhu, Zhe Liu, Ming Tang, and Shougang Ren. Deep learning gradient visualization-based
 504 pre-silicon side-channel leakage location. *IEEE Transactions on Information Forensics and Security*, 2024.

505 Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Anders Fogh, Jann Horn,
 506 Stefan Mangard, Paul Kocher, Daniel Genkin, Yuval Yarom, and Mike Hamburg. Meltdown: Reading
 507 kernel memory from user space. In *27th USENIX Security Symposium (USENIX Security 18)*, 2018.

508 Phillip Lippe, Taco Cohen, and Efstratios Gavves. Efficient neural causal discovery without acyclicity
 509 constraints. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=eYciPrLuUhG>.

511 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on
 512 Learning Representations*, 2018.

513 Xiangjun Lu, Chi Zhang, Pei Cao, Dawu Gu, and Haining Lu. Pay attention to raw traces: A deep learning
 514 architecture for end-to-end profiling attacks. *IACR Transactions on Cryptographic Hardware and Embedded
 515 Systems*, pp. 235–274, 2021.

516 Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation
 517 of discrete random variables. In *International Conference on Learning Representations*, 2017. URL
 518 <https://openreview.net/forum?id=S1jE5L5gl>.

519 Houssem Maghrebi, Thibault Portigliatti, and Emmanuel Prouff. Breaking cryptographic implementations
 520 using deep learning techniques. In *Security, Privacy, and Applied Cryptography Engineering: 6th Inter-
 521 national Conference, SPACE 2016, Hyderabad, India, December 14-18, 2016, Proceedings 6*, pp. 3–26.
 522 Springer, 2016.

523 Stefan Mangard, Elisabeth Oswald, and Thomas Popp. *Power analysis attacks*. Springer New York, NY,
 524 1st edition, March 2007. doi: 10.1007/978-0-387-38162-6. URL <https://link.springer.com/book/10.1007/978-0-387-38162-6>.

526 Loïc Masure and Rémi Strullu. Side-channel analysis against ANSSI’s protected AES implementation on
 527 ARM: end-to-end attacks with multi-task learning. *Journal of Cryptographic Engineering*, 13(2):129–147,
 528 2023.

529 Loïc Masure, Cécile Dumas, and Emmanuel Prouff. Gradient visualization for general characterization in
 530 profiling attacks. In *Constructive Side-Channel Analysis and Secure Design: 10th International Workshop,
 531 COSADE 2019, Darmstadt, Germany, April 3–5, 2019, Proceedings 10*, pp. 145–167. Springer, 2019.

532 Thomas S Messerges. Using second-order power analysis to attack dpa resistant software. In *International
 533 Workshop on Cryptographic Hardware and Embedded Systems*, pp. 238–251. Springer, 2000.

534 Nicky Mouha. Review of the Advanced Encryption Standard. NIST Interagency/Internal Report (NISTIR)
 535 8319, National Institute of Standards and Technology, July 2021. URL <https://csrc.nist.gov/pubs/ir/8319/final>.

537 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen,
 538 Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep
 539 learning library. *Advances in neural information processing systems*, 32, 2019.

540 Guilherme Perin, Lichao Wu, and Stjepan Picek. I know what your layers did: Layer-wise explainability
 541 of deep learning side-channel analysis. Cryptology ePrint Archive, Paper 2022/1087, 2022. URL <https://eprint.iacr.org/2022/1087>.

543 Stjepan Picek, Guilherme Perin, Luca Mariot, Lichao Wu, and Lejla Batina. SoK: Deep learning-based
 544 physical side-channel analysis. *ACM Computing Surveys*, 55(11):1–35, 2023.

545 Jean-Jacques Quisquater and David Samyde. ElectroMagnetic Analysis (EMA): Measures and counter-
 546 measures for smart cards. In Isabelle Attali and Thomas Jensen (eds.), *Smart Card Programming and*
 547 *Security*, pp. 200–210, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg. ISBN 978-3-540-45418-2.

548 Christian Rechberger and Elisabeth Oswald. Practical template attacks. In Chae Hoon Lim and Moti Yung
 549 (eds.), *Information Security Applications*, pp. 440–456, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
 550 ISBN 978-3-540-31815-6.

551 Eric Rescorla. HTTP over TLS. RFC 2818, May 2000. URL <https://www.rfc-editor.org/info/rfc2818>.

552 Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate
 553 inference in deep generative models. In Eric P. Xing and Tony Jebara (eds.), *Proceedings of the 31st*
 554 *International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*,
 555 pp. 1278–1286, Bejing, China, 22–24 Jun 2014. PMLR. URL <https://proceedings.mlr.press/v32/rezende14.html>.

556 Kotaro Saito, Akira Ito, Rei Ueno, and Naofumi Homma. One truth prevails: A deep-learning based
 557 single-trace power analysis on rsa–crt with windowed exponentiation. *IACR Transactions on Cryptographic*
 558 *Hardware and Embedded Systems*, pp. 490–526, 2022.

559 Thomas Schamberger, Maximilian Egger, and Lars Tebelmann. Hide and seek: Using occlusion techniques
 560 for side-channel leakage attribution in cnns. In Jianying Zhou, Lejla Batina, Zengpeng Li, Jingqiang
 561 Lin, Eleonora Losiouk, Suryadipta Majumdar, Daisuke Mashima, Weizhi Meng, Stjepan Picek, Moham-
 562 mad Ashiqur Rahman, Jun Shao, Masaki Shimaoka, Ezekiel Soremekun, Chunhua Su, Je Sen Teh, Aleksei
 563 Udoenko, Cong Wang, Leo Zhang, and Yury Zhauniarovich (eds.), *Applied Cryptography and Network*
 564 *Security Workshops*, pp. 139–158, Cham, 2023. Springer Nature Switzerland. ISBN 978-3-031-41181-6.

565 Werner Schindler, Kerstin Lemke, and Christof Paar. A stochastic model for differential side channel
 566 cryptanalysis. In *Cryptographic Hardware and Embedded Systems–CHES 2005: 7th International Workshop*,
 567 *Edinburgh, UK, August 29–September 1, 2005. Proceedings* 7, pp. 30–46. Springer, 2005.

568 Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27
 569 (3):379–423, 1948.

570 Alexander Shekhovtsov. Bias-variance tradeoffs in single-sample binary gradient estimators. In *DAGM*
 571 *German Conference on Pattern Recognition*, pp. 127–141. Springer, 2021.

572 Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating
 573 activation differences. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International*
 574 *Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 3145–3153.
 575 PMLR, 06–11 Aug 2017. URL <https://proceedings.mlr.press/v70/shrikumar17a.html>.

576 K Simonyan, A Vedaldi, and A Zisserman. Deep inside convolutional networks: visualising image classification
 577 models and saliency maps. In *Proceedings of the International Conference on Learning Representations*
 578 (*ICLR*). ICLR, 2014.

579 Biaoshuai Tao and Hongjun Wu. Improving the biclique cryptanalysis of AES. In Ernest Foo and Douglas
 580 Stebila (eds.), *Information Security and Privacy*, pp. 39–56, Cham, 2015. Springer International Publishing.
 581 ISBN 978-3-319-19962-7.

582 George Tucker, Andriy Mnih, Chris J Maddison, John Lawson, and Jascha Sohl-Dickstein. Rebar: Low-
 583 variance, unbiased gradient estimates for discrete latent variable models. *Advances in Neural Information*
 584 *Processing Systems*, 30, 2017.

585 Daan van der Valk, Stjepan Picek, and Shivam Bhasin. Kilroy was here: The first step towards explainability
 586 of neural networks in profiled side-channel analysis. In Guido Marco Bertoni and Francesco Regazzoni (eds.),
 587 *Constructive Side-Channel Analysis and Secure Design*, pp. 175–199, Cham, 2021. Springer International
 588 Publishing. ISBN 978-3-030-68773-1.

590 Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni
591 Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett,
592 Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric
593 Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold,
594 Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro,
595 Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for
596 Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.

597 Zhiguang Wang, Weizhong Yan, and Tim Oates. Time series classification from scratch with deep neural
598 networks: A strong baseline. In *2017 International joint conference on neural networks (IJCNN)*, pp.
599 1578–1585. IEEE, 2017.

600 Leo Weissbart, Stjepan Picek, and Lejla Batina. One trace is all it takes: Machine learning-based side-channel
601 attack on eddsa. In *Security, Privacy, and Applied Cryptography Engineering: 9th International Conference,
602 SPACE 2019, Gandhinagar, India, December 3–7, 2019, Proceedings 9*, pp. 86–105. Springer, 2019.

603 Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning.
604 *Machine learning*, 8:229–256, 1992.

605 Lennert Wouters, Victor Arribas, Benedikt Gierlichs, and Bart Preneel. Revisiting a methodology for efficient
606 CNN architectures in profiling attacks. *IACR Transactions on Cryptographic Hardware and Embedded
607 Systems*, pp. 147–168, 2020.

608 Yuxin Wu and Justin Johnson. Rethinking" batch" in batchnorm. *arXiv preprint arXiv:2105.07576*, 2021.

609 Trevor Yap, Adrien Benamira, Shivam Bhasin, and Thomas Peyrin. Peek into the black-box: Interpretable
610 neural network using sat equations in side-channel analysis. *IACR Transactions on Cryptographic Hardware
611 and Embedded Systems*, pp. 24–53, 2023.

612 Trevor Yap, Stjepan Picek, and Shivam Bhasin. Occpois: Points of interest based on neural network's key
613 recovery in side-channel analysis through occlusion. In Sourav Mukhopadhyay and Pantelimon Stănică
614 (eds.), *Progress in Cryptology – INDOCRYPT 2024*, pp. 3–28, Cham, 2025. Springer Nature Switzerland.
615 ISBN 978-3-031-80311-6.

616 Gabriel Zaid, Lilian Bossuet, Amaury Habrard, and Alexandre Venelli. Methodology for efficient CNN
617 architectures in profiling attacks. *IACR Transactions on Cryptographic Hardware and Embedded Systems*,
618 pp. 1–36, 2020.

619 Matthew D. Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In David Fleet,
620 Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (eds.), *Computer Vision – ECCV 2014*, pp. 818–833,
621 Cham, 2014. Springer International Publishing. ISBN 978-3-319-10590-1.

622 **A Notation and variable names**

See table 3 below for a list of the notation we use, and table 4 for a list of the main variables we define.

Table 3: List of the notation used in our paper.

| Symbol | Description |
|---|--|
| Variable case: X, x | Upper case: random variable, lower-case: actualization |
| \mathbb{R}, \mathbb{Z} | Set of real numbers, set of integers |
| $[a .. b]$ where $a < b, a, b \in \mathbb{Z}$ | Interval of integers $\{a, \dots, b\}$ |
| Serif font: S | Other sets |
| $A \subseteq B$ | A is a non-strict subset of B |
| $A \setminus B$ | Complement of B in A , i.e. $A \setminus B := \{x \in A : x \notin B\}$ |
| S_+ | Nonnegative elements of $S \subseteq \mathbb{R}$ |
| S_{++} | Positive elements of $S \subseteq \mathbb{R}$ |
| Bold font: \mathbf{x} | Vector in \mathbb{R}^D for some $D \in \mathbb{Z}_{++}$ |
| $\mathbf{a} \odot \mathbf{b}$ | Elementwise product of \mathbf{a} and \mathbf{b} |
| $f(\mathbf{x})$ where $f : \mathbb{R} \rightarrow \mathbb{R}$ | Elementwise application of f to \mathbf{x} |
| $\nabla_{\mathbf{x}} f(\mathbf{x}, \dots)$ | Gradient with respect to \mathbf{x} |
| \mathbf{x}_α where $\alpha \in \{0, 1\}^D$ | Sub-vector of \mathbf{x} according to α : $(x_d : d = 1, \dots, D : \alpha_d = 1)$ |
| $X \sim p$ | X is a random variable with distribution p |
| p_A | The distribution of A |
| $p_{A,B}$ | The joint distribution of A and B |
| $p_{A B}$ | The conditional distribution of A given B |
| $\mathbf{X} \sim p^N$ | \mathbf{X} is a random vector with elements $X_1, X_2, \dots \stackrel{\text{i.i.d.}}{\sim} p$ |
| $\mathcal{N}(\mu; \sigma^2)$ | Normal distribution with mean μ , scale σ |
| $\mathcal{U}(S)$ | Uniform distribution over set S |
| $\mathbb{E} f(A, B, \dots)$ | Expectation w.r.t. all random variables A, B, \dots in expression |
| $\mathbb{E}_A f(A, B, \dots)$ | Expectation w.r.t. only to A |
| $\mathbb{I}[A; B]$ | Mutual information between A and B |
| $\mathbb{I}[A; B C]$ | Conditional mutual information between A and B given C |
| $A \perp\!\!\!\perp B, A \not\perp\!\!\!\perp B$ | A is marginally independent, dependent on B |
| $A \perp\!\!\!\perp B C, A \not\perp\!\!\!\perp B C$ | A is conditionally independent, dependent on B given C |

623

624 **B Extended background**

625 Here we provide a high-level overview of the AES algorithm and power side-channel attacks aimed at a
 626 machine learning audience. Since our algorithm views the cryptographic algorithm and hardware as a black
 627 box to be characterized with data, a deep understanding is not necessary to understand and appreciate our
 628 work. Thus, we omit many details and aim to impart an intuitive understanding of these topics. Interested
 629 readers may refer to Daemen & Rijmen (2013) for a detailed introduction to the AES algorithm, to Mangard
 630 et al. (2007) for a detailed introduction to power side-channel attacks, and to Picek et al. (2023) for a survey
 631 of supervised deep learning-based power side-channel attacks on AES implementations. Additionally, note
 632 that while for clarity of exposition we focus here on power side-channel attacks on AES implementations, our
 633 algorithm requires only a supervised learning-style dataset of side-channel emission traces and metadata which
 634 enables computation of the target variable, so is applicable in a more-general setting. We have demonstrated
 635 that it works with various target variables on AES, ECC and RSA implementations with both power and
 636 EM radiation measurements, and suspect it is relevant in far more contexts.

Table 4: List of the main variables defined in our paper.

| Variable | Description |
|---|--|
| $T \in \mathbb{Z}_{++}$ | Dimensionality of power trace |
| $\mathbf{X} \sim p_{\mathbf{X}}$ | Power trace (random variable) |
| $\mathbf{x} \in \mathbb{R}^T$ | Power trace (actualization) |
| X_t, x_t | Power measurement at time t , i.e. t -th element of power trace |
| $\mathbb{Y} \subseteq \mathbb{Z}$ | Set of values the targeted variable may take on |
| $Y \sim p_Y$ | Targeted variable (random variable) |
| $y \in \mathbb{Y}$ | Targeted variable (actualization) |
| $\mathbf{D} \subseteq \mathbb{R}^T \times \mathbb{Y}$ | Dataset of power traces/targeted variable pairs, sampled i.i.d. from $p_{\mathbf{X}, Y}$ |
| $\gamma \in [0, 1]^T$ | Occlusion probabilities |
| γ_t | Occlusion probability at timestep t |
| $\gamma^* \in [0, 1]^T$ | The optimal value of γ after solving Eqn. ?? |
| $\tilde{\eta} \in \mathbb{R}^T$ | Unconstrained logits which parameterize γ |
| $\mathcal{A}_\gamma \sim p_{\mathcal{A}_\gamma}$ | Multiplicative binary noise vector parameterized by occlusion probabilities γ |
| $\alpha \in \{0, 1\}^T$ | Actualization of \mathcal{A}_γ |
| $c : [0, 1] \rightarrow \mathbb{R}_+$ | Cost function for occlusion probability elements γ_t |
| $C \in \mathbb{R}_{++}$ | Budget for occlusion probabilities: they must satisfy $C = \sum_{t=1}^T c(\gamma_t)$ |
| $\bar{\gamma} \in (0, 1)$ | Reparameterized version of C which is more stable w/ data dimensionality |
| $\Phi_\theta : \mathbb{R}^T \times [0, 1]^T \rightarrow \mathbb{R}$ | Noise-conditional neural net w/ weights θ ; returns softmax logits for Y |

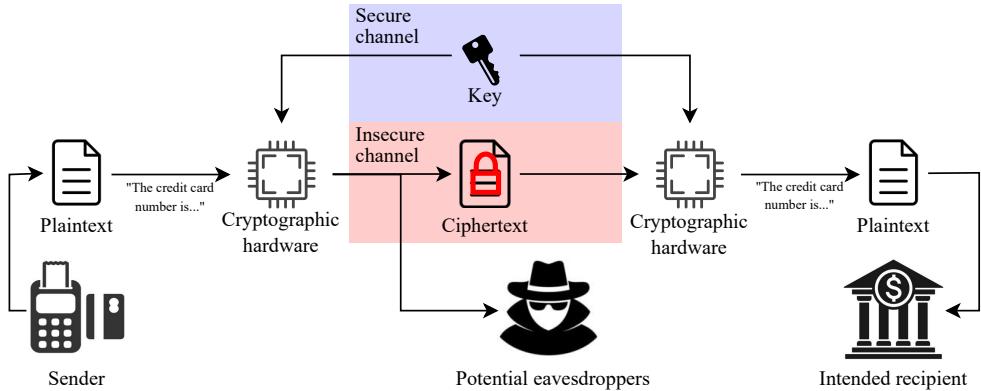


Figure 6: Diagram illustrating the main components of symmetric-key cryptographic algorithms, which enable secure transmission of data over insecure channels where it may be intercepted by eavesdroppers. The data is first partitioned and encoded as a sequence of plaintexts. Each plaintext is transformed into a ciphertext by an invertible function indexed by a cryptographic key. The key is transmitted over a secure channel to intended recipients of the data, allowing them to invert the function and recover the original plaintext. The set of functions is designed so that absent this key, the ciphertext gives no information about the plaintext. Thus, the data remains secure even if eavesdroppers have access to the ciphertext.

637 **B.1 Cryptographic algorithms**

638 Data is often transmitted over insecure channels which leave it accessible not only to intended recipients,
 639 but also to unknown and untrusted parties. For example, when a signal is wirelessly transmitted from one
 640 antenna to another, an eavesdropper could set up a third antenna between the two and intercept the signal.
 641 Alternately, data stored on a hard drive by one user of a computer may be accessed by a different user.
 642 Cryptographic algorithms aim to preserve the privacy of data under such circumstances by transforming it so
 643 that it is meaningful only in combination with additional data which is known to its intended recipients but
 644 not to the untrusted parties.

645 We focus here on the ubiquitous advanced encryption standard (AES), which is a symmetric-key cryptographic
 646 algorithm. See Fig. 6 for a diagram illustrating the important components of such algorithms. The unencrypted
 647 data to be transmitted is encoded and partitioned into a sequence of fixed-length bitstrings called *plaintexts*.
 648 The cryptographic algorithm encrypts each plaintext into a *ciphertext* by applying an invertible function
 649 from a set of functions indexed by an integer called the *cryptographic key*. This set of functions is designed so
 650 that if one were to sample a key and plaintext uniformly at random from the sets of all possible keys and
 651 plaintexts, then the plaintext and ciphertext would be marginally independent. Thus, such an algorithm may
 652 be used to securely transmit data by ensuring that the sender and recipient of the data know a shared key,³
 653 and that the key is kept secret from all potential eavesdroppers on the data.

654 **B.2 Side-channel attacks**

655 Many symmetric-key cryptographic algorithms are believed to be secure in the sense that it is not feasible to
 656 determine their cryptographic key by encrypting known plaintexts and observing the resulting ciphertexts.
 657 Any such algorithm with a finite number of possible keys is vulnerable to ‘brute-force’ attacks based on
 658 arbitrarily guessing and checking keys until success, but doing so requires checking half of all possible keys in
 659 the average case, which is unrealistic for algorithms such as AES which has either 2^{128} , 2^{192} , or 2^{256} possible
 660 keys. To our knowledge the best known such attack against AES reduces the required number of guesses by
 661 less than a factor of 8 compared to a naive brute force attack (Mouha, 2021; Tao & Wu, 2015).

662 However, while algorithms may be secure when considering only their intended inputs and outputs, *hardware*
 663 *executing these algorithms* will inevitably emit measurable physical signals which are statistically associated
 664 with their intermediate variables and operations. Examples of such signals include a device’s power con-
 665 sumption over time (Kocher et al., 1999), the amount of time it takes to execute a program or instruction
 666 (Kocher, 1996; Lipp et al., 2018; Kocher et al., 2019), electromagnetic radiation it emits (Quisquater &
 667 Samyde, 2001; Genkin et al., 2016), and sound due to vibrations of its electronic components (Genkin et al.,
 668 2014). This phenomenon is called *side-channel leakage*, and can be exploited to determine sensitive data such
 669 as a cryptographic key through *side-channel attacks*.

670 As a simple example of side-channel leakage, consider the following Python function which checks whether a
 671 password is correct:

```
672 def is_correct(provided_password: str, correct_password: str) -> bool:
 673     if len(provided_password) != len(correct_password):
 674         return False
 675     for i in range(len(provided_password)):
 676         if provided_password[i] != correct_password[i]:
 677             return False
 678     return True
```

679 Suppose the password consists of n characters, each with c possible values. Consider an attacker seeking
 680 to determine the correct password by feeding various guessed passwords until the function returns `True`.
 681 Naively, the attacker could simply guess and check all possible m -length passwords for $m = 1, \dots, n$. This
 682 would require $O(c^n)$ calls to the function, which would be extremely costly for realistically-large c and
 683 n . However, an attacker with knowledge of the function’s implementation could dramatically reduce this
 684 cost by observing that the function’s *execution time* depends on `correct_password`. Because the function
 685 exits immediately if `len(provided_password) != len(correct_password)`, the attacker can determine
 686 the length of `correct_password` in $O(n)$ time by feeding increasing-length guesses to `is_correct` until its
 687 execution time increases. Next, because `is_correct` exits the first time it detects an incorrect character, the
 688 attacker can sequentially determine each of the characters of `correct_password` by checking all c possible
 689 values of each character and noting that the correct value leads to an increase in execution time. Thus,
 690 although `is_correct` is secure against attackers which use only its intended inputs and outputs, it provides
 691 *essentially no security* against attackers which measure its execution time.

³The key is typically shared using an asymmetric-key cryptographic algorithm such as RSA or ECC. Asymmetric-key cryptography is slow and resource-intensive, so when a sufficiently-large amount of data must be transmitted, it is more-efficient to share the key with an asymmetric-key algorithm and then transmit data using a symmetric-key algorithm than to simply transmit the data with an asymmetric-key algorithm.

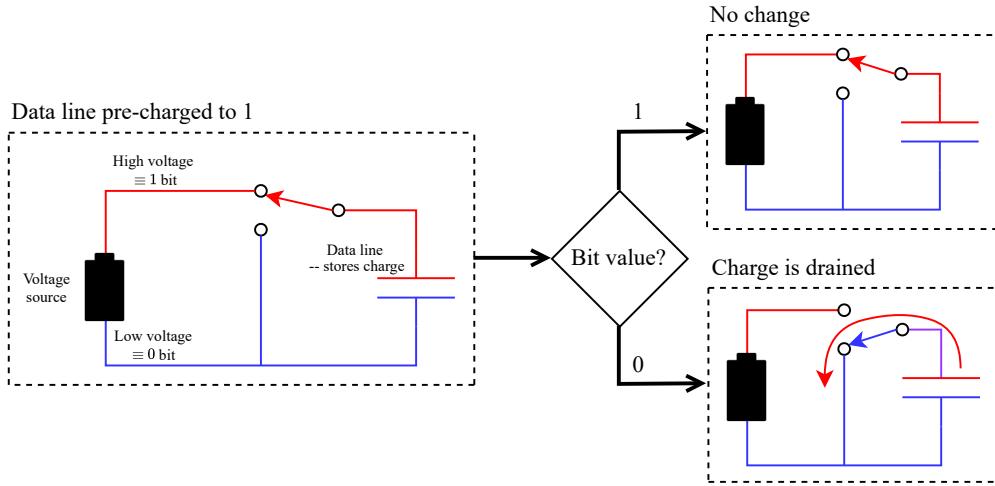


Figure 7: Diagram illustrating one reason there is power side-channel leakage in the device characterized by (Mangard et al., 2007, ch. 4). Data is transmitted over a bus consisting of multiple wires, with one wire representing each bit. Each wire represents a 0 bit as some prescribed ‘low’ voltage and a 1 bit as a ‘high’ voltage. Energy is consumed when the voltage of a wire changes from low to high because positive and negative charges, which are attracted to one-another, must be separated to create a high concentration of positive charge on the wire. When ‘writing’ data to the bus, this particular device first ‘pre-charges’ all wires to 1, then drains charge from the wires which should represent 0. Thus, because the 0’s must be changed to 1’s before the next write, energy is consumed in proportion to the number of 0’s, thereby creating a statistical association between the device’s power consumption and the data it operates on.

692 In this work we focus on side-channel leakage due to the power consumption over time of a device (as well as
 693 EM radiation, which is closely related due to being dominated by the time derivative of power consumption).
 694 A device’s power consumption is inevitably statistically-associated with the operations it executes and the
 695 data it operates on, because these dictate which components are active and the order and manner in which
 696 they operate. There are many types of components with different functionality, and components with the same
 697 intended functionality are not identical due to imperfect manufacturing processes. These differences impact
 698 power consumption. While in general the association between power consumption and data is multifactorial
 699 and difficult to describe, in Fig. 7 we illustrate a simple relationship which accounts for a significant portion
 700 of the leakage in a device characterized by Mangard et al. (2007).

701 B.3 Power side-channel attacks on AES implementations

702 *Side-channel attacks* are techniques which exploit side-channel leakage to learn sensitive information such as
 703 cryptographic keys. There are many categories of attacks, but in this work we focus on a category called
 704 *profiled* side-channel attacks on symmetric-key cryptographic algorithms. These attacks assume that the
 705 ‘attacker’ has access to a clone of the actual cryptographic device to be attacked, and the ability to encrypt
 706 arbitrary plaintexts with arbitrary cryptographic keys, observe the resulting ciphertexts, and measure the
 707 side-channel leakage during encryption. In practice, these assumptions almost certainly overestimate the
 708 capabilities of attackers – for example, while in some cases an attacker could plausibly identify the hardware
 709 and source code of a cryptographic implementation, purchase copies of this hardware, program them with the
 710 source code, and characterize these devices, the nature of the side-channel leakage of these purchased copies
 711 would differ from those of the actual device due to PVT (pressure, voltage, temperature) variations (e.g. due
 712 to imperfect manufacturing processes, environment, and measurement setup). It has been demonstrated
 713 that profiled side-channel attacks can be effective despite this, especially when numerous copies of the target
 714 hardware are used for profiling (Das et al., 2019; Danial et al., 2021). Regardless, this type of attack provides
 715 an upper bound on the vulnerability of a device to side-channel attacks, which is a useful metric for hardware
 716 designers.

717 While there are diverse types of profiled side-channel attacks, at a high level the following steps encompass
 718 the important elements of these attacks:

- 719 1. Select some ‘sensitive’ intermediate variable of the cryptographic algorithm which reveals the
 720 cryptographic key (or part of it).
- 721 2. Compile a dataset of (side-channel leakage, intermediate variable) pairs by repeatedly randomly
 722 selecting a key and plaintext, encrypting the plaintext using the key and recording the resulting
 723 ciphertext and side-channel leakage during encryption, and computing the intermediate variable
 724 based on knowledge of the cryptographic algorithm.
- 725 3. Use supervised learning to train a parametric function approximator to predict intermediate variables
 726 from recordings of side-channel leakage during encryption.
- 727 4. Measure side-channel leakage during encryptions by the actual target device. Use the trained
 728 predictor to predict sensitive variables from side-channel leakage. Potentially, these predictions can
 729 be combined to get a better estimate of the key.

730 In the case of power side-channel attacks on AES, it is generally infeasible to directly target the cryptographic
 731 key because care is taken by hardware designers to prevent it from directly influencing power consumption.
 732 Instead, it is common to target an intermediate variable which the algorithm directly operates on. A common
 733 target is the `SubBytes` output, which is computed as

$$y := \text{Sbox}(k \oplus w) \quad (5)$$

734 where $k \in \{0, 1\}^{n_{\text{bits}}}$ is the key, $w \in \{0, 1\}^{n_{\text{bits}}}$ is the plaintext, $n_{\text{bits}} \in \mathbb{Z}_{++}$ is the number of bits of the key
 735 and plaintext, \oplus is the bitwise exclusive-or operation, and $\text{Sbox} : \{0, 1\}^{n_{\text{bits}}} \rightarrow \{0, 1\}^{n_{\text{bits}}}$ is an invertible
 736 function which is widely known and the same for all AES implementations. Note that if the plaintext is
 737 known, the key can be computed as

$$k = \text{Sbox}^{-1}(y) \oplus w. \quad (6)$$

738 Additionally, it is common to independently target subsets of the bits of the cryptographic key (e.g. the
 739 individual bytes). This is reasonable because the common AES target variables are largely leaked by
 740 instructions which operate on individual bytes.

741 **B.3.1 Template attack: example of a classical profiled side-channel attack**

742 In order to underscore the advantage of deep learning over previous side-channel attack algorithms, we will
 743 here describe the template attack algorithm of Chari et al. (2003), variations of which are the state-of-the-art
 744 non-deep learning based attacks. The attack is based on modeling the joint distribution of power consumption
 745 and intermediate variable as a Gaussian mixture model, as described in algorithm 1.

746 Note that this algorithm assumes that the joint distribution is well-described by a Gaussian mixture model,
 747 which may not hold in practice. Additionally, due to the near-cubic runtime of the matrix inversion of each Σ_y
 748 required to compute the Gaussian density functions, this algorithm requires pruning power traces down to a
 749 small number of ‘high-leakage’ timesteps. Follow-up work (Rechberger & Oswald, 2005) proposed performing
 750 principle component analysis on the traces and modeling the coefficients of the top principle components
 751 rather than individual timesteps. Nonetheless, these constraints mean that the efficacy of this attack is
 752 contingent on simplifying assumptions and judgement of which points are ‘leaky’ using simple statistical
 753 techniques and implementation knowledge, limiting its usefulness as a way for hardware designers to evaluate
 754 the amount of side-channel leakage from their device.

755 **B.3.2 Practical profiled deep learning side-channel attacks on AES implementations**

756 Here we will give a common and concrete setting and method for performing profiled power side-channel
 757 attacks on AES implementations, which is used for all of our experiments.

Algorithm 1: The Gaussian template attack algorithm of Chari et al. (2003)

Input: Profiling (training) dataset $D := \{(\mathbf{x}^{(n)}, y^{(n)} : n \in [1..N]\} \subseteq \mathbb{R}^T \times \{0,1\}^{n_{\text{bits}}}$, attack (testing) dataset $D_{\text{attack}} := \{(\mathbf{x}_a^{(n)}, w_a^{(n)} : n \in [1..N_a]\} \subseteq \mathbb{R}^T \times \{0,1\}^{n_{\text{bits}}}$, ‘points of interest’ $T_{\text{poi}} := \{t_m : m = 1, \dots, \tilde{T}\} \subseteq [1..T]$

Output: Predicted key k^*

```

1 Function get_y ( $k, w$ )
2   return  $\text{Sbox}(k \oplus w)$                                 // calculate intermediate variable for given key
3 for  $n \in [1..N]$  do
4    $\tilde{\mathbf{x}}^{(n)} \leftarrow (x_{t_m}^{(n)} : m = 1, \dots, \tilde{T})$           // prune power traces to ‘points of interest’
5 for  $y \in \{0,1\}^{n_{\text{bits}}}$  do
6   // fit a multivariate Gaussian mixture model to the training dataset
7    $D_y \leftarrow \{\tilde{\mathbf{x}}^{(n)} : n \in [1..N], y^{(n)} = y\}$ 
8    $N_y \leftarrow |D_y|$ 
9    $\boldsymbol{\mu}_y \leftarrow \frac{1}{N_y} \sum_{\tilde{\mathbf{x}} \in D_y} \tilde{\mathbf{x}}$ 
10   $\boldsymbol{\Sigma}_y \leftarrow \frac{1}{N_y - 1} \sum_{\tilde{\mathbf{x}} \in D_y} (\tilde{\mathbf{x}} - \boldsymbol{\mu}_y)(\tilde{\mathbf{x}} - \boldsymbol{\mu}_y)^\top$ 
11 for  $n \in [1..N_a]$  do
12    $\tilde{\mathbf{x}}_a^{(n)} \leftarrow (x_{a,t_m}^{(n)} : m = 1, \dots, \tilde{T})$           // prune power traces of attack dataset
13   // predict key value which maximizes log-likelihood of attack dataset
14    $k^* \leftarrow \arg \max_{k \in \{0,1\}^{n_{\text{bits}}}} \sum_{n=1}^{N_a} \left[ \log \mathcal{N} \left( \tilde{\mathbf{x}}^{(n)}; \boldsymbol{\mu}_{\text{get}_y(k, w_a^{(n)})}, \boldsymbol{\Sigma}_{\text{get}_y(k, w_a^{(n)})} \right) + \log N_{\text{get}_y(k, w_a^{(n)})} \right]$ 
15 return  $k^*$ 

```

758 Consider an AES-128 implementation, which has a 128-bit cryptographic key and plaintext. Typically,
 759 attackers target each of the 16 bytes of the key independently rather than attacking the full key at once. This
 760 practice tacitly assumes that the bytes of the sensitive variable are statistically-independent given the power
 761 trace, which is reasonable because many AES operations (including those which are commonly targeted) are
 762 performed independently on the individual bytes. Thus, it is a convenient way to simplify the attack without
 763 significantly impacting performance.

764 Additionally, it is difficult and uncommon to try to directly map power traces to associated cryptographic
 765 keys, because great care is taken by hardware designers to ensure that the key does not directly impact power
 766 consumption. Instead, attackers generally target ‘sensitive’ intermediate variables which unavoidably directly
 767 impact power consumption and can be combined with the plaintext and ciphertext to learn the key. We
 768 consider one such intermediate variable which is referred to as the first **SubBytes** output, and is equal to

$$y := \text{Sbox}(k \oplus w), \quad (7)$$

769 where $k \in \{0,1\}^8$ is one byte of the cryptographic key, $w \in \{0,1\}^8$ is the corresponding byte of the plaintext,
 770 \oplus denotes the bitwise exclusive-or operation, and $\text{Sbox} : \{0,1\}^8 \rightarrow \{0,1\}^8$ is an invertible function which is
 771 publicly-available and the same for all AES implementations. Note that if w is known, as is assumed in the
 772 profiled side-channel attack setting, then k can be recovered as

$$k = w \oplus \text{Sbox}^{-1}(y). \quad (8)$$

773 In the context of profiled power side-channel analysis, one assumes to have a ‘profiling’ dataset (i.e. a training
 774 dataset) and an ‘attack’ dataset (i.e. a test dataset). Suppose we target n_{bytes} bytes of the sensitive variable.
 775 In our setting, the profiling dataset consists of ordered pairs of power traces and their associated sensitive
 776 intermediate variables:

$$D := \{(\mathbf{x}^{(n)}, y^{(n)} : n \in [1..N]\} \subseteq \mathbb{R}^T \times \{0,1\}^{n_{\text{bytes}} \times 8} \quad (9)$$

777 and the attack dataset consists of ordered pairs of power traces and their associated plaintexts:

$$D_a := \left\{ (\mathbf{x}_a^{(n)}, w_a^{(n)}) : n \in [1 \dots N_a] \right\} \subseteq \mathbb{R}^T \times \{0, 1\}^{n_{\text{bytes}} \times 8}. \quad (10)$$

778 Many works prove the concept of their approaches by targeting only a single byte of the sensitive variable.
 779 When multiple bytes are targeted, it is common to either train a separate neural network for each byte of the
 780 sensitive variable, or to amortize the cost of targeting these bytes by training a single neural network with a
 781 shared backbone and a separate head for each byte. In this work we exclusively target single bytes, though it
 782 would be straightforward to extend our approach to the multitask learning setting.

783 Consider a neural network architecture $\Phi : \mathbb{Y} \times \mathbb{R}^T \times \mathbb{R}^P \rightarrow \mathbb{R}_+ : (y, \mathbf{x}, \boldsymbol{\theta}) \mapsto \Phi(y \mid \mathbf{x}; \boldsymbol{\theta})$, where each $\Phi(\cdot \mid \mathbf{x}; \boldsymbol{\theta})$
 784 is a probability mass function over \mathbb{Y} . In the case of a multi-headed network with each head independently
 785 predicting a single byte, we compute this probability mass of $y \in \mathbb{Y}$ as the product of the mass assigned to
 786 each of its bytes. We train the network by approximately solving the optimization problem

$$\max_{\boldsymbol{\theta} \in \mathbb{R}^P} \mathcal{L}(\boldsymbol{\theta}) := \frac{1}{N} \sum_{n=1}^N \log \Phi(y^{(n)} \mid \mathbf{x}^{(n)}; \boldsymbol{\theta}). \quad (11)$$

787 Given $\hat{\boldsymbol{\theta}} \in \arg \max_{\boldsymbol{\theta} \in \mathbb{R}^P} \mathcal{L}(\boldsymbol{\theta})$, we then identify the key which maximizes our estimated likelihood of our
 788 attack dataset and key as follows:

$$\hat{k} \in \arg \max_{k \in \{0, 1\}^{n_{\text{bytes}} \times 8}} \sum_{n=1}^{N_a} \log \Phi \left(\left(\text{Sbox}(k_i \oplus w_{a,i}^{(n)}) : i = 1, \dots, n_{\text{bytes}} \right) \mid \mathbf{x}_a^{(n)}; \hat{\boldsymbol{\theta}} \right) \quad (12)$$

789 where we denote by k_i and $w_i^{(n)}$ the individual bytes of k and $w^{(n)}$.

790 C Extended related work

791 Here we consider existing work which has been applied to leakage localization in the context of power or EM
 792 radiation side-channel analysis. In line with the problem framing given in Sec. 2, we view these methods
 793 as functions which map joint emission-target variable distributions $p_{\mathbf{X}, Y}$ to vectors in \mathbb{R}^T which assign
 794 to each emission measurement variable X_t a scalar ‘leakiness’ measurement. Prior approaches to leakage
 795 localization can largely be categorized as either 1) parametric statistics-based methods which check for
 796 pairwise associations between X_t and Y , or 2) neural net attribution methods which use standard supervised
 797 deep learning techniques to train a model $\hat{p}_{Y \mid \mathbf{X}} \approx p_{Y \mid \mathbf{X}}$, then use ‘attribution’ techniques to estimate the
 798 average ‘importance’ of each feature X_t to the predictions made by the model.

799 C.1 First-order parametric statistics-based methods

800 In the side-channel attack literature it is common to use parametric first-order statistical methods to localize
 801 leakage. In this work we consider the signal to noise ratio (SNR) (Mangard et al., 2007), the sum of squared
 802 differences (SOSD) (Chari et al., 2003), and correlation power analysis with a Hamming weight leakage model
 803 (CPA) (Brier et al., 2004) due to their popularity and efficacy for ‘point of interest’ selection (Fan et al.,
 804 2014). Below we summarize and discuss these methods.

805 The SNR is a standard tool for leakage localization and is defined as

$$\text{SNR}(p_{\mathbf{X}, Y}) := \frac{\text{Var}_{Y \sim p_Y} \mathbb{E}_{\mathbf{X} \sim p_{\mathbf{X} \mid Y}} [\mathbf{X}]}{\mathbb{E}_{Y \sim p_Y} \text{Var}_{\mathbf{X} \sim p_{\mathbf{X} \mid Y}} (\mathbf{X})}. \quad (13)$$

806 The SOSD method was introduced as a point of interest (feature) selection technique for the Gaussian
 807 template attack, a parametric side-channel attack based on modeling emission measurements as a multivariate
 808 Gaussian mixture model with a component corresponding to each possible value of the target variable,

809 then using Bayes' rule to estimate the conditional distribution of the target variable given the emission
 810 measurements. The SOSD is defined as

$$\text{sosd}(p_{\mathbf{X},Y}) := \sum_{y \in \mathbf{Y}} \sum_{y' \in \mathbf{Y}} \left(\mathbb{E}_{\mathbf{X} \sim p_{\mathbf{X}|Y}(\cdot|y)}[\mathbf{X}] - \mathbb{E}_{\mathbf{X} \sim p_{\mathbf{X}|Y}(\cdot|y')}[\mathbf{X}] \right)^2. \quad (14)$$

811 CPA is based on the assumption that power consumption is a noisy linear function of the Hamming weight
 812 of the target variable, which is a useful model for certain devices. It is defined as the elementwise Pearson
 813 correlation between emission measurements and the target variable's Hamming weight:

$$\text{cpa}(p_{\mathbf{X},Y}) := \frac{\mathbb{E}[(\mathbf{X} - \mathbb{E}[\mathbf{X}])(\text{HW}(Y) - \mathbb{E}[\text{HW}(Y)])]}{\sqrt{\text{Var}(\mathbf{X}) \text{Var}(\text{HW}(Y))}} \quad (15)$$

814 where $\text{HW} : [0 \dots 2^d - 1] \rightarrow [0 \dots d]$ maps integers to the sum of their bits when writing them as an unsigned
 815 integer.

816 While techniques such as these are invaluable due to their simplicity, interpretability and low cost, they have
 817 major shortcomings. Notably, they consider they are sensitive only to pairwise associations between single
 818 emission measurements and the target variable and will consider the measurement X_t to be non-leaky when
 819 it is nonleaky in isolation but gives exploitable information *when combined* with X_τ for some $\tau \neq t$, i.e. when
 820 $\mathbb{I}[Y; X_t] = 0$ but $\exists \tau \neq t$ such that $\mathbb{I}[Y; X_t | X_\tau] > 0$.

821 Additionally, these techniques make strong assumptions about the nature of $p_{\mathbf{X},Y}$ which have generally been
 822 observed to hold in practice, but nonetheless introduce the risk of failing to detect leaking measurements.
 823 SNR and SOSD are sensitive only to the influence of Y on the mean of \mathbf{X} , and would fail to identify X_t is
 824 leaking if Y changes its distribution while leaving its mean unchanged (e.g. if the variance of X_t changes
 825 with Y). CPA is sensitive only to associations between X_t and the Hamming weight of Y , and additionally
 826 assumes a linear relationship between these variables.

827 While the present work concerns mainly 'black box' leakage localization algorithms which make minimal
 828 assumptions about the cryptographic implementation being evaluated, in practice these parametric methods
 829 are often employed as tools in white-box analyses of implementations. For example, in our work we consider
 830 the ASCADv1 datasets (Benadjila et al., 2020), which use a Boolean masking countermeasure and thus
 831 have mainly second-order leakage. This renders the above methods ineffective when directly analyzing
 832 leakage of their canonical target variable. However, Egger et al. (2022) identified 4 *pairs* of internal AES
 833 variables which individually have first-order leakage and may be combined to determine the target variable.
 834 Thus, if one is aware *a priori* that these variables leak and has access to the internal randomly-generated
 835 Boolean mask variables, they may individually analyze leakage of these variables with the above methods
 836 and accumulate the results. We use such an approach to compute 'ground truth' leakiness measurements
 837 when running experiments on the ASCADv1 datasets. We emphasize that this type of analysis is challenging
 838 and error-prone: multiple leaking variable identified by Egger et al. (2022) were unknown or overlooked by
 839 Benadjila et al. (2020), who introduced the dataset. This underscores the importance of black box techniques
 840 such as ours to supplement white box analysis.

841 C.2 Neural net attribution methods

842 There is a great deal of prior work on localizing leakage based on interpretability techniques to determine the
 843 relative importance of input features to a deep neural net which has been trained to model $\hat{p}_{Y|\mathbf{X}} \approx p_{Y|\mathbf{X}}$
 844 using standard supervised deep learning techniques (Masure et al., 2019; Hettwer et al., 2020; Jin et al., 2020;
 845 Zaid et al., 2020; Wouters et al., 2020; van der Valk et al., 2021; Wu & Johnson, 2021; Golder et al., 2022; Li
 846 et al., 2022; Perin et al., 2022; Schamberger et al., 2023; Yap et al., 2023; Li et al., 2024; Yap et al., 2025).
 847 As baselines we consider the recent works Yap et al. (2025); Schamberger et al. (2023) as well as a variety of
 848 older neural net attribution techniques which were compared in Masure et al. (2019); Hettwer et al. (2020);
 849 Wouters et al. (2020). Note that our choice of deep learning baselines subsumes those of Yap et al. (2025);
 850 Schamberger et al. (2023). Since these methods all 'interpret' a trained deep neural net, we view them as
 851 functions mapping the data distribution $p_{\mathbf{X},Y}$ as well as a model $\hat{p}_{Y|\mathbf{X}}$ s.t. $\text{softmax } \hat{p}_{Y|\mathbf{X}} \approx p_{Y|\mathbf{X}}$ to a vector

852 of leakiness estimates in \mathbb{R}^T . Here we summarize and discuss the works Masure et al. (2019); Hettwer et al.
 853 (2020); Wouters et al. (2020); Schamberger et al. (2023); Yap et al. (2025) which we consider as baselines.

854 To our knowledge Masure et al. (2019) were the first to explore neural net interpretability for leakage
 855 localization, proposing the Saliency-like GradVis leakage assessment, defined as

$$\text{GradVis}(p_{\mathbf{X},Y}; \hat{p}_{Y|\mathbf{X}}) := \mathbb{E}_{\mathbf{X},Y} \left| -\nabla_{\mathbf{x}} \log \text{softmax} \hat{p}_{Y|\mathbf{X}}(Y | \mathbf{x}) \right|_{\mathbf{x}=\mathbf{X}}. \quad (16)$$

856 Hettwer et al. (2020) subsequently compared the 1-Occlusion (Zeiler & Fergus, 2014), Saliency (Simonyan
 857 et al., 2014), and layerwise relevance propagation (LRP) (Bach et al., 2015) as leakage localization techniques.
 858 The 1-Occlusion technique is based on computing the size of change in the model's prediction as each
 859 individual input feature is 'occluded' (replaced by 0), and is defined as

$$\text{1-Occlusion}(p_{\mathbf{X},Y}; \hat{p}_{Y|\mathbf{X}}) := \mathbb{E}_{\mathbf{X},Y} \left(\left| \hat{p}_{Y|\mathbf{X}}(Y | \mathbf{X}) - \hat{p}_{Y|\mathbf{X}}(Y | (\mathbf{1} - \mathbf{I}_t) \odot \mathbf{X}) \right| : t = 1, \dots, T \right) \quad (17)$$

860 where \mathbf{I}_t denotes the vector in \mathbb{R}^T with element t equal to 1 and all other elements equal to 0. Saliency is
 861 defined as

$$\text{Saliency}(p_{\mathbf{X},Y}; \hat{p}_{Y|\mathbf{X}}) := \mathbb{E}_{\mathbf{X},Y} \left| \nabla_{\mathbf{x}} \hat{p}_{Y|\mathbf{X}}(Y | \mathbf{x}) \right|_{\mathbf{x}=\mathbf{X}}. \quad (18)$$

862 LRP (Bach et al., 2015) is a gradient-based explainability technique which is more-complicated than the
 863 above, and we refer readers to Bach et al. (2015) for an explanation. Wouters et al. (2020) applied the Input *
 864 Grad method (Shrikumar et al., 2017) to leakage localization; as its name suggests, this method is defined as

$$\text{Input * Grad}(p_{\mathbf{X},Y}; \hat{p}_{Y|\mathbf{X}}) := \mathbb{E}_{\mathbf{X},Y} \left| \mathbf{X} \odot \nabla_{\mathbf{x}} \hat{p}_{Y|\mathbf{X}}(Y | \mathbf{x}) \right|_{\mathbf{x}=\mathbf{X}}. \quad (19)$$

865 We find that all these techniques have a tendency to incorrectly assign low leakiness to certain measurements,
 866 particularly in scenarios where lots of features have significant leakage. We suspect that a primary reason
 867 for this is that these methods rely on perturbing the input x_t to the function $\hat{p}_{Y|\mathbf{X}}(y | x_1, \dots, x_t, \dots, x_T)$,
 868 but when Y is almost entirely determined by X_τ for $\tau \neq t$, it becomes nearly independent of X_t conditioned
 869 on $\{X_1, \dots, X_T\} \setminus \{X_t\}$, i.e. $\hat{p}_{Y|\mathbf{X}}(y | x_\tau : \tau = 1, \dots, T) \approx \hat{p}_{Y|\mathbf{X}}(y | x_\tau : \tau = 1, \dots, t-1, t+1, \dots, T)$. We
 870 demonstrate this phenomenon with a simple Gaussian mixture model setting in Sec. C.3.

871 We also consider the recent works Schamberger et al. (2023); Yap et al. (2025) which estimate leakiness by
 872 perturbing many inputs simultaneously to a classifier, and do not necessarily suffer from the same issue.
 873 Schamberger et al. (2023) presents the m -Occlusion technique, which is like 1-Occlusion except that it
 874 occludes m -diameter windows rather than single points. This could plausibly overcome the aforementioned
 875 issue if the 'redundant' points are temporally-local. However, it has an undesirable 'smoothing' effect where
 876 the estimated leakiness of a single point is tied to those of nearby points. Schamberger et al. (2023) also
 877 proposes to use 2nd-order m -Occlusion to analyze leakiness, where *pairs* of windows are occluded rather
 878 than only individual windows. This is computationally-expensive because it requires $\Theta(T^2)$ passes through
 879 the dataset where T is the data dimensionality. Additionally, Schamberger et al. (2023) proposes it as a
 880 means to determine whether a measurement has first-order leakage or is part of a second-order leaking pair,
 881 and does not explore its use for single-measurement leakiness estimation. In our experiments we find it
 882 only marginally-better than 1-Occlusion for this task, and not worth the significantly-higher computational
 883 cost. Yap et al. (2025) proposes the OccPOI technique, which aims to identify a non-unique minimal set
 884 of measurements sufficient for a neural net to attain some chosen classification performance when all other
 885 measurements are occluded. This differs from our other considered methods in that it does not assign a
 886 leakiness estimate to every measurement, and we find that it is ill-suited for our leakage localization task and
 887 performance metrics. Additionally, it is computationally-expensive to run, requiring $\Omega(T)$ *non-parallelizable*
 888 passes through the dataset.

889 **C.3 Numerical experiment illustrating conditional mutual information decay when many redundant
 890 leaking measurements are present**

891 Here we provide a simple numerical experiment to illustrate conditional mutual information decay. Consider
 892 random variables $Y \sim \mathcal{U}\{-1, 1\}$ and $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(Y, \sigma^2)$. In Fig. 8 we plot the quantity $\mathbb{I}[Y; X_n |$

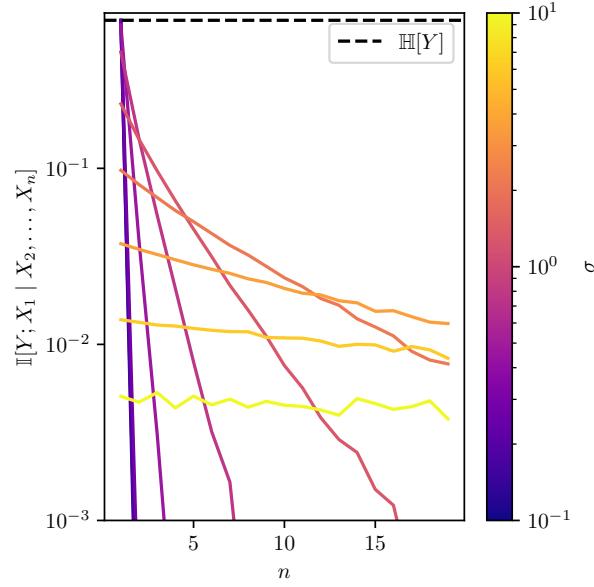


Figure 8: A numerical experiment evaluating the scaling behavior of $\mathbb{I}[Y; X_n | X_1, \dots, X_{n-1}]$ vs. n for various values of σ , where $Y \sim \mathcal{U}\{-1, 1\}$ and $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(Y, \sigma^2)$. Observe that for small σ , each X_i is approximately a point mass on Y and we have $\mathbb{I}[Y; X_1] \approx \mathbb{H}[Y]$ and $\mathbb{I}[Y; X_n | X_1, \dots, X_{n-1}] \approx 0$ for all $n > 1$. When σ is large, each X_i gives us little information about Y and we have $\mathbb{I}[Y; X_n | X_1, \dots, X_{n-1}] \ll \mathbb{H}[Y]$ approximately constant with n sufficiently small.

893 $X_1, \dots, X_{n-1}]$ vs. n . We see that for various values of σ the quantity decays rapidly with n , and appears
 894 to be well-described by the function $\mathbb{I}[Y; X_1 | X_2, \dots, X_n] \approx I k^n$ for some $I \in \mathbb{R}_+$, $k \in (0, 1)$. Informally,
 895 the fact that conditional mutual information decays with n makes sense for the following reason: each
 896 successive X_i can be viewed as an independent ‘noisy observation’ of Y which reduces uncertainty about it
 897 without fully determining it. For all n we have $0 \leq \mathbb{I}[Y; X_1, \dots, X_n] = \mathbb{I}[Y; X_1] + \mathbb{I}[Y; X_2 | X_1] + \dots + \mathbb{I}[Y; X_n |$
 898 $X_1, \dots, X_{n-1}] \leq \mathbb{H}[Y]$. If we assume that each $\mathbb{I}[Y; X_m | X_1, \dots, X_{m-1}]$ is nonnegative, then for their sum
 899 to stay bounded they must converge to 0.

900 D Extended method with derivations

901 Given $\mathbf{X}, Y \sim p_{\mathbf{X}, Y}$ as defined in section 2, where $\mathbf{X} := (X_1, \dots, X_T)$, we seek to assign to each timestep
 902 t a scalar γ_t^* indicating the ‘amount of leakage’ about Y due to X_t . All of the quantities $\mathbb{I}[X; X_t | \mathbf{S}]$ for
 903 $\mathbf{S} \subseteq \{X_1, \dots, X_T\} \setminus \{X_t\}$ are relevant to the ‘leakiness’ of X_t , but it is not clear how to weight these 2^T
 904 quantities into a single scalar measurement. Most prior work simply ignores most of these quantities: the
 905 first-order parametric methods (Mangard et al., 2007; Brier et al., 2004; Chari et al., 2003) consider only the
 906 pairwise terms $\mathbb{I}[Y; X_t]$, and GradVis (Masure et al., 2019), Saliency (Simonyan et al., 2014; Hettwer et al.,
 907 2020), 1-Occlusion (Zeiler & Fergus, 2014; Hettwer et al., 2020), LRP (Bach et al., 2015; Hettwer et al., 2020),
 908 Input * Grad (Shrikumar et al., 2017; Wouters et al., 2020) may loosely be viewed as computing proxies for
 909 $\mathbb{I}[Y; X_t | X_{\tau \neq t}]$. While m -Occlusion, 2^{nd} -order m -Occlusion (Schamberger et al., 2023), and OccPOI (Yap
 910 et al., 2025) are sensitive to more of these terms, they still ignore almost all of them in an *ad hoc* manner.

911 In this section we propose a constrained optimization problem which implicitly defines an intuitively-
 912 reasonable definition of γ_t^* which is sensitive to $\mathbb{I}[Y; X_t | \mathbf{S}]$ for all $\mathbf{S} \subseteq \{X_1, \dots, X_T\} \setminus \{X_t\}$. We then propose
 913 an adversarial deep learning algorithm which lets us approximately solve it by modeling all the conditional
 914 distributions $p_{Y|\mathbf{S}}$ in an amortized manner which emphasizes those with a large impact on the objective.
 915 While our objective is a sum over 2^T occlusion mask-like values, in practice we find we can efficiently optimize
 916 it using the reparameterization trick (Kingma & Welling, 2014; Rezende et al., 2014) with a CONCRETE-like

917 relaxation (Jang et al., 2017; Maddison et al., 2017) of our objective. This lets us exploit first-order gradient
 918 information, in contrast to ‘hard’ occlusion-based methods such as Schamberger et al. (2023); Yap et al.
 919 (2025) which leverage only zeroth-order information.

920 **D.1 Optimization problem**

921 We define a vector $\gamma \in [0, 1]^T$ which we name the *occlusion probabilities*. We use γ to parameterize a
 922 distribution over binary vectors in $\{0, 1\}^T$ as follows:

$$\mathcal{A}_\gamma \sim p_{\mathcal{A}_\gamma} \quad \text{where} \quad \mathcal{A}_{\gamma,t} = \begin{cases} 1 & \text{with probability } 1 - \gamma_t \\ 0 & \text{with probability } \gamma_t, \end{cases} \quad (20)$$

923 i.e. \mathcal{A}_γ is a vector of independent Bernoulli random variables where the t -th element has parameter $p = 1 - \gamma_t$.
 924 For arbitrary vectors $\mathbf{x} \in \mathbb{R}^T$, $\alpha \in \{0, 1\}^T$, let us denote $\mathbf{x}_\alpha := (x_t : t = 1, \dots, T : \alpha_t = 1)$, i.e. the sub-vector
 925 of \mathbf{x} containing its elements for which the corresponding element of α is 1. We can accordingly use \mathcal{A}_γ to
 926 obtain random sub-vectors $\mathbf{X}_{\mathcal{A}_\gamma}$ of \mathbf{X} . Note that γ_t denotes the probability that X_t will *not* be an element
 927 of $\mathbf{X}_{\mathcal{A}_\gamma}$ (hence, ‘occlusion probability’).

928 We assign to each element of γ a ‘cost’, defined as

$$c : [0, 1] \rightarrow \mathbb{R}_+ : x \mapsto \begin{cases} \frac{x}{1-x} & x < 1 \\ \infty & x = 1. \end{cases} \quad (21)$$

929 We seek to solve the constrained optimization problem

$$\min_{\gamma \in [0, 1]^T} \mathcal{L}_{\text{ideal}}(\gamma) := \mathbb{I}[Y; \mathbf{X}_{\mathcal{A}_\gamma} | \mathcal{A}_\gamma] \quad \text{such that} \quad \sum_{t=1}^T c(\gamma_t) = C \quad (22)$$

930 for hyperparameter $C > 0$. Note that c is strictly-increasing with $c(0) = 0$ and $\lim_{x \rightarrow 1^-} c(x) = \infty$ so that for
 931 finite C any optimal γ will be in $[0, 1]^T$, and increasing some γ_t necessarily entails reducing some other γ_τ ,
 932 $\tau \neq t$. Additionally, for each t we can re-write our objective as

$$\mathcal{L}_{\text{ideal}}(\gamma) = \sum_{\alpha \in \{0, 1\}^T} p_{\mathcal{A}_\gamma}(\alpha) \mathbb{I}[Y; \mathbf{X}_\alpha] \quad (23)$$

$$= \sum_{\substack{\alpha \in \{0, 1\}^T \\ \alpha_t = 0}} p_{\mathcal{A}_{\gamma, -t}}(\alpha_{-t}) [(1 - \gamma_t) \mathbb{I}[Y; X_t, \mathbf{X}_\alpha] + \gamma_t \mathbb{I}[Y; \mathbf{X}_\alpha]] \quad (24)$$

$$= \sum_{\substack{\alpha \in \{0, 1\}^T \\ \alpha_t = 0}} p_{\mathcal{A}_{\gamma, -t}}(\alpha_{-t}) [\mathbb{I}[Y; X_t, \mathbf{X}_\alpha] - \gamma_t \mathbb{I}[Y; X_t | \mathbf{X}_\alpha]], \quad (25)$$

933 which implies

$$\frac{\partial \mathcal{L}_{\text{ideal}}(\gamma)}{\partial \gamma_t} = - \sum_{\substack{\alpha \in \{0, 1\}^T \\ \alpha_t = 0}} p_{\mathcal{A}_{\gamma, -t}}(\alpha_{-t}) \mathbb{I}[Y; X_t | \mathbf{X}_\alpha]. \quad (26)$$

934 **D.2 Estimating mutual information with deep neural nets**

935 We cannot solve equation 22 directly because we lack an expression for $p_{\mathbf{X}, Y}$. Here we derive an equivalent
 936 optimization problem which uses deep learning to characterize $p_{\mathbf{X}, Y}$ using data.

937 Consider the family $\{\Phi_\alpha\}_{\alpha \in \{0, 1\}^T}$ with each element a deep neural net

$$\Phi_\alpha : \mathbb{Y} \times \mathbb{R}^{\sum_{t=1}^T \alpha_t} \times \mathbb{R}^P \rightarrow [0, 1] : (y, \mathbf{x}_\alpha, \theta) \mapsto \Phi_\alpha(y | \mathbf{x}_\alpha; \theta). \quad (27)$$

938 We assume each $\Phi_{\alpha}(\cdot | \mathbf{x}; \boldsymbol{\theta})$ is a probability mass function over \mathbf{Y} (e.g. the neural net has a softmax output
939 activation). We define the optimization problem

$$\min_{\boldsymbol{\gamma} \in [0, 1]^T} \max_{\boldsymbol{\theta} \in \mathbb{R}^P} \mathcal{L}_{\text{adv}}(\boldsymbol{\gamma}, \boldsymbol{\theta}) := \mathbb{E} \log \Phi_{\mathbf{A}_{\boldsymbol{\gamma}}}(Y | \mathbf{X}_{\mathbf{A}_{\boldsymbol{\gamma}}}; \boldsymbol{\theta}) \quad \text{such that} \quad \sum_{t=1}^T c(\gamma_t) = C. \quad (28)$$

940 **Proposition D.1.** Consider the objective function \mathcal{L}_{adv} of equation 28. Suppose there exists some $\boldsymbol{\theta}^* \in \mathbb{R}^P$
941 such that $\Phi_{\alpha}(y | \mathbf{x}_{\alpha}; \boldsymbol{\theta}^*) = p_{Y|\mathbf{X}_{\alpha}}(y | \mathbf{x}_{\alpha})$ for all $\alpha \in \{0, 1\}^T$, $\mathbf{x} \in \mathbb{R}^T$, $y \in \mathbf{Y}$. Then

$$\boldsymbol{\theta}^* \in \arg \max_{\boldsymbol{\theta} \in \mathbb{R}^P} \mathcal{L}_{\text{adv}}(\boldsymbol{\gamma}, \boldsymbol{\theta}) \quad \forall \boldsymbol{\gamma} \in [0, 1]^T. \quad (29)$$

942 Furthermore, for all $y \in \mathbf{Y}$ and for all $\boldsymbol{\gamma} \in [0, 1]^T$, $\alpha \in \{0, 1\}^T$ such that $p_{\mathbf{A}_{\boldsymbol{\gamma}}}(\alpha) > 0$,

$$\Phi_{\alpha}(y | \mathbf{x}_{\alpha}; \hat{\boldsymbol{\theta}}) = p_{Y|\mathbf{X}_{\alpha}}(y | \mathbf{x}_{\alpha}) \quad p_{\mathbf{X}}\text{-almost surely} \quad \forall \hat{\boldsymbol{\theta}} \in \arg \min_{\boldsymbol{\theta} \in \mathbb{R}^P} \mathcal{L}_{\text{adv}}(\boldsymbol{\gamma}, \boldsymbol{\theta}). \quad (30)$$

943 *Proof.* Note that since each $\Phi_{\alpha}(\cdot | \mathbf{x}, \boldsymbol{\theta})$ is a probability mass function over \mathbf{Y} , by Gibbs' inequality we have

$$\mathbb{E} \log \Phi_{\alpha}(Y | \mathbf{X}_{\alpha}; \boldsymbol{\theta}) \leq \mathbb{E} \log p_{Y|\mathbf{X}_{\alpha}}(Y | \mathbf{X}_{\alpha}) \quad \forall \alpha \in \{0, 1\}^T, \boldsymbol{\theta} \in \mathbb{R}^P. \quad (31)$$

944 Thus,

$$\mathcal{L}_{\text{adv}}(\boldsymbol{\gamma}, \boldsymbol{\theta}^*) \geq \mathcal{L}_{\text{adv}}(\boldsymbol{\gamma}, \boldsymbol{\theta}) \quad \forall \boldsymbol{\theta} \in \mathbb{R}^P, \boldsymbol{\gamma} \in [0, 1]^T, \quad (32)$$

945 which implies the first claim.

946 Next, consider some fixed $\boldsymbol{\gamma} \in [0, 1]^T$ and $\hat{\boldsymbol{\theta}} \in \arg \min_{\boldsymbol{\theta} \in \mathbb{R}^P} \mathcal{L}_{\text{adv}}(\boldsymbol{\gamma}, \boldsymbol{\theta})$. We must have $\mathcal{L}_{\text{adv}}(\boldsymbol{\gamma}, \hat{\boldsymbol{\theta}}) = \mathcal{L}_{\text{adv}}(\boldsymbol{\gamma}, \boldsymbol{\theta}^*)$.
947 Thus,

$$0 = \mathcal{L}_{\text{adv}}(\boldsymbol{\gamma}, \boldsymbol{\theta}^*) - \mathcal{L}_{\text{adv}}(\boldsymbol{\gamma}, \hat{\boldsymbol{\theta}}) \quad (33)$$

$$= \mathbb{E} \left[\log p_{Y|\mathbf{X}_{\alpha}}(Y | \mathbf{X}_{\alpha}) - \log \Phi_{\alpha}(Y | \mathbf{X}_{\alpha}; \hat{\boldsymbol{\theta}}) \right] \quad (34)$$

$$= \sum_{\alpha \in \{0, 1\}^T} p_{\mathbf{A}_{\boldsymbol{\gamma}}}(\alpha) \mathbb{E} \left[\log p_{Y|\mathbf{X}_{\alpha}}(Y | \mathbf{X}_{\alpha}) - \log \Phi_{\alpha}(Y | \mathbf{X}_{\alpha}; \hat{\boldsymbol{\theta}}) \right]. \quad (35)$$

948 By Gibbs' inequality, each of the expectations in the summation is nonnegative, which implies that whenever
949 $p_{\mathbf{A}_{\boldsymbol{\gamma}}}(\alpha) > 0$ we must have

$$0 = \mathbb{E} \left[\log p_{Y|\mathbf{X}_{\alpha}}(Y | \mathbf{X}_{\alpha}) - \log \Phi_{\alpha}(Y | \mathbf{X}_{\alpha}; \hat{\boldsymbol{\theta}}) \right] \quad (36)$$

$$= \int_{\mathbb{R}^T} \sum_{t=1}^T \alpha_t p_{\mathbf{X}_{\alpha}}(\mathbf{x}_{\alpha}) \mathbb{KL} \left[p_{Y|\mathbf{X}_{\alpha}}(\cdot | \mathbf{x}_{\alpha}) \| \Phi_{\alpha}(\cdot | \mathbf{x}_{\alpha}; \hat{\boldsymbol{\theta}}) \right] d\mathbf{x}_{\alpha}. \quad (37)$$

950 Since $\mathbb{KL} \left[p_{Y|\mathbf{X}_{\alpha}}(\cdot | \mathbf{x}_{\alpha}) \| \Phi_{\alpha}(\cdot | \mathbf{x}_{\alpha}; \hat{\boldsymbol{\theta}}) \right] \geq 0$ with equality if and only if $p_{Y|\mathbf{X}_{\alpha}}(y | \mathbf{x}_{\alpha}) = \Phi_{\alpha}(y | \mathbf{x}_{\alpha}; \hat{\boldsymbol{\theta}})$
951 $\forall y \in \mathbf{Y}$, this must be the case except possibly for $\mathbf{x} \in \mathbb{R}^T$ where

$$\int_{\{\mathbf{x}_{\alpha}: \mathbf{x} \in \mathbb{R}^T\}} p_{\mathbf{X}_{\alpha}}(\mathbf{x}_{\alpha}) d\mathbf{x}_{\alpha} = 0 \implies \int_{\mathbb{R}^T} p_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} = 0. \quad (38)$$

952 This implies the second claim. \square

953 **Corollary D.2.** Under the assumptions of Proposition D.1, equations 22 and 28 are equivalent.

954 *Proof.* Observe that for any $\gamma \in [0, 1]^T$,

$$\max_{\theta \in \mathbb{R}^P} \mathcal{L}_{\text{adv}}(\gamma, \theta) = \max_{\theta \in \mathbb{R}^P} \mathbb{E} \log \Phi_{\mathcal{A}_\gamma}(Y \mid \mathbf{X}_{\mathcal{A}_\gamma}(Y \mid \mathbf{X}_{\mathcal{A}_\gamma}; \theta)) \quad (39)$$

$$= \sum_{\alpha \in \{0,1\}^T} p_{\mathcal{A}_\gamma}(\alpha) \mathbb{E} \log p_{Y \mid \mathbf{X}_\alpha}(Y \mid \mathbf{X}_\alpha) \quad \text{by Prop. D.1} \quad (40)$$

$$= - \sum_{\alpha \in \{0,1\}^T} p_{\mathcal{A}_\gamma}(\alpha) \mathbb{H}[Y \mid \mathbf{X}_\alpha] \quad (41)$$

$$\equiv \sum_{\alpha \in \{0,1\}^T} p_{\mathcal{A}_\gamma}(\alpha) [\mathbb{H}[Y] - \mathbb{H}[Y \mid \mathbf{X}_\alpha]] \quad \text{because } \mathbb{H}[Y] \text{ is not a function of } \gamma \quad (42)$$

$$= \sum_{\alpha \in \{0,1\}^T} p_{\mathcal{A}_\gamma}(\alpha) \mathbb{I}[Y; \mathbf{X}_\alpha] \quad (43)$$

$$= \mathbb{I}[Y; \mathbf{X}_{\mathcal{A}_\gamma} \mid \mathcal{A}_\gamma] \quad (44)$$

$$= \mathcal{L}_{\text{ideal}}(\gamma). \quad (45)$$

955 This implies the result. \square

956 **Corollary D.3.** *Suppose the assumptions of Proposition D.1 are satisfied, and let $\hat{\theta} \in \arg \min_{\theta \in \mathbb{R}^P} \mathcal{L}_{\text{adv}}(\gamma, \theta)$*
957 *for some $\gamma \in [0, 1]^T$. Consider $\alpha := \alpha' + \alpha''$ where $\alpha', \alpha'' \in \{0, 1\}^T$ such that $\alpha'_t = 1 \implies \alpha''_t = 0$ and*
958 *$\alpha''_t = 1 \implies \alpha'_t = 0$, and $p_{\mathcal{A}_\gamma}(\alpha) > 0$. For all $y \in \mathbb{Y}$, it follows immediately from Proposition D.1 that*
959 *$p_{\mathbf{X}}$ -almost everywhere we can use our classifiers to compute the pointwise mutual information quantities*

$$\text{pmi}(y; \mathbf{x}_{\alpha'} \mid \mathbf{x}_{\alpha''}) := \log p_{Y \mid \mathbf{X}_\alpha}(y \mid \mathbf{x}_\alpha) - \log p_{Y \mid \mathbf{X}_{\alpha''}}(y \mid \mathbf{x}_{\alpha''}) \quad (46)$$

$$= \log \Phi_\alpha(y \mid \mathbf{x}_\alpha; \hat{\theta}) - \log \Phi_{\alpha''}(y \mid \mathbf{x}_{\alpha''}; \hat{\theta}). \quad (47)$$

960 This is useful because it allows us to assess leakage from *single* power traces, as opposed to merely summarizing
961 distributions of power traces. There are scenarios where a power measurement might leak for some traces
962 but not for others. For example, a common countermeasure is to randomly delay leaky instructions or swap
963 their order with another instruction so that they do not occur at a deterministic time relative to the start of
964 encryption. One could use pmi computations to determine the timetep at which the leaky instruction has
965 been run in a single trace.

966 Since it would be impractical to train 2^T deep neural networks independently, we implement the family of
967 classifiers by a single neural net with input dropout and with the dropout mask fed to the neural net as an
968 auxiliary input:

$$\Phi : \mathbb{Y} \times \mathbb{R}^T \times \{0, 1\}^T \times \mathbb{R}^P \rightarrow [0, 1] : (y, \mathbf{x}, \alpha, \theta) \mapsto \Phi(y \mid \mathbf{x} \odot \alpha, \alpha; \theta) \quad (48)$$

969 where $\Phi_\alpha(y \mid \mathbf{x}_\alpha; \theta) := \Phi(y \mid \mathbf{x} \odot \alpha, \alpha; \theta)$. This approach was inspired by Lippe et al. (2022).

970 D.3 Re-parametrization into an unconstrained optimization problem

971 We would like to approximately solve equation 28 using an alternating stochastic gradient descent-style
972 approach, similarly to GANs (Goodfellow et al., 2014). Thus, it is convenient to express it as an unconstrained
973 optimization problem. We first define a new vector $\eta \in \Delta^{T-1}$ where $\Delta^{T-1} := \{(\delta_1, \dots, \delta_T) \in \mathbb{R}_+^T : \sum_{t=1}^T \delta_t =$
974 $1\}$ denotes the T -simplex. We then define γ to be the vector satisfying the equality

$$c(\gamma_t) = C\eta_t \quad (49)$$

$$\implies \frac{\gamma_t}{1 - \gamma_t} = C\eta_t \quad (50)$$

$$\implies \log \gamma_t - \log(1 - \gamma_t) = \log C + \log \eta_t \quad (51)$$

$$\implies \gamma_t = \text{sigmoid}(\log C + \log \eta_t). \quad (52)$$

975 If we define $\boldsymbol{\eta} := \text{softmax}(\tilde{\boldsymbol{\eta}})$ for $\tilde{\boldsymbol{\eta}} \in \mathbb{R}^T$, then we can express

$$\gamma_t = \text{sigmoid}(\log C + \log \tilde{\eta}_t - \text{logsumexp}(\tilde{\boldsymbol{\eta}})), \quad (53)$$

976 which allows us to map the unconstrained vector $\tilde{\boldsymbol{\eta}}$ to $\boldsymbol{\gamma}$ or $\log \boldsymbol{\gamma}$ using numerically-stable PyTorch operations.
 977 Our constrained optimization problem 28 is thus equivalent to the following unconstrained problem:

$$\min_{\tilde{\boldsymbol{\eta}} \in \mathbb{R}^T} \max_{\boldsymbol{\theta} \in \mathbb{R}^P} \mathcal{L}(\tilde{\boldsymbol{\eta}}, \boldsymbol{\theta}) := \mathbb{E} \log \Phi(Y \mid \mathbf{X} \odot \mathcal{A}_{\boldsymbol{\gamma}(\tilde{\boldsymbol{\eta}})}, \mathcal{A}_{\boldsymbol{\gamma}(\tilde{\boldsymbol{\eta}})}; \boldsymbol{\theta}). \quad (54)$$

978 **D.4 Implementation details**

979 It is infeasible to exactly compute the expectation with respect to $\mathcal{A}_{\boldsymbol{\gamma}(\tilde{\boldsymbol{\eta}})}$ because doing so would require
 980 summing over 2^T terms. As is routine in deep learning contexts, we instead approximate the gradient
 981 with Monte Carlo integration. Note that our objective⁴ takes the form $\mathcal{L}(\tilde{\boldsymbol{\eta}}) = \mathbb{E} f(\mathcal{A}_{\boldsymbol{\gamma}(\tilde{\boldsymbol{\eta}})})$ where $f(\boldsymbol{\alpha}) :=$
 982 $\mathbb{E}_{\mathbf{X}, Y} \log \Phi(Y \mid \mathbf{X} \odot \boldsymbol{\alpha}, \boldsymbol{\alpha}; \boldsymbol{\theta})$ and the distribution of $\mathcal{A}_{\boldsymbol{\gamma}(\tilde{\boldsymbol{\eta}})}$ depends on $\tilde{\boldsymbol{\eta}}$.

983 Unbiased estimators for $\nabla_{\tilde{\boldsymbol{\eta}}} \mathcal{L}(\boldsymbol{\eta})$ of this form are usually based on the REINFORCE estimator (Williams,
 984 1992) with control variates, and tend to have complicated implementations and high variance. We tried
 985 using the vanilla REINFORCE estimator with simple control variates as well as the more-sophisticated
 986 REBAR estimator (Tucker et al., 2017), and found that the former works poorly, while the latter works well.
 987 Subsequent ablation studies revealed that the biased CONCRETE estimator (Maddison et al., 2017) works
 988 almost as well as REBAR for our application and is considerably simpler, so in this work we use CONCRETE.

989 The CONCRETE estimator lets us write $\mathcal{A}_{\boldsymbol{\gamma}(\tilde{\boldsymbol{\eta}})}$ as a deterministic function of $\tilde{\boldsymbol{\eta}}$ and $\tilde{\boldsymbol{\eta}}$ -independent noise and
 990 thereby use the reparameterization trick to estimate $\nabla_{\tilde{\boldsymbol{\eta}}} \mathcal{L}(\tilde{\boldsymbol{\eta}})$ using standard automatic differentiation tools.
 991 For binary random variables such as the elements of $\mathcal{A}_{\boldsymbol{\gamma}(\tilde{\boldsymbol{\eta}})}$, the estimator is built on the observation that

$$x \sim \text{Bernoulli}(p) \equiv x = H(\log p - \log(1 - p) + \log u - \log(1 - u)) \quad \text{for } u \sim \mathcal{U}(0, 1) \quad (55)$$

992 where $H(x) := \begin{cases} 1 & x \geq 0 \\ 0 & x < 0 \end{cases}$ denotes the unit step function. H is not amenable to gradient descent because
 993 its derivative is zero almost everywhere, but we can approximate it with the tempered sigmoid function
 994 $\text{sigmoid}_\tau(x) := \text{sigmoid}(x/\tau)$. The temperature parameter $\tau > 0$ can be tuned to control a bias-variance
 995 tradeoff for the estimator: $\lim_{\tau \rightarrow 0} \text{sigmoid}_\tau(x) = H(x) \forall x \in \mathbb{R} \setminus \{0\}$, but variance increases as $O(1/\tau)$
 996 (Shekhovtsov, 2021). We find that results are reasonable when we simply leave τ fixed at 1, and we do this
 997 throughout the present work. We conjecture that this is because all our performance evaluation metrics
 998 consider only the *relative* leakiness estimates produced by our method, and the nature of the bias is not to
 999 significantly impact the relative sizes of the elements of $\boldsymbol{\gamma}^*$.

1000 Note that while our original loss function

$$\ell(\tilde{\boldsymbol{\eta}}, \boldsymbol{\theta}, \mathbf{x}, y, \boldsymbol{\alpha}) := \log \Phi(y \mid \mathbf{x} \odot \boldsymbol{\alpha}, \boldsymbol{\alpha}; \boldsymbol{\theta}) \quad (56)$$

1001 is written for a ‘hard’ occlusion mask $\boldsymbol{\alpha} \in \{0, 1\}^T$, the relaxed occlusion masks lie inside the open ball $(0, 1)^T$.
 1002 Thus, we must relax our loss function to accept these inputs as well. While the right-hand side of Eqn. 56 is
 1003 still a valid expression for $\boldsymbol{\alpha}$ in $(0, 1)^T$, its optimal value of this loss with respect to $\boldsymbol{\theta}$ does not vary smoothly
 1004 with $\boldsymbol{\alpha}$ because rescaling the elements of \mathbf{X} by nonzero constants does not change its mutual information
 1005 with Y . Thus, we instead use the stochastic relaxed loss function

$$\ell_{\text{relaxed}}(\tilde{\boldsymbol{\eta}}, \boldsymbol{\theta}, \mathbf{x}, y, \boldsymbol{\alpha}) := \log \Phi(y \mid \mathbf{x} \odot \boldsymbol{\alpha} + \boldsymbol{\varepsilon} \odot (\mathbf{1} - \boldsymbol{\alpha}); \boldsymbol{\theta}) \quad \text{where } \boldsymbol{\varepsilon} \sim \mathcal{N}(0, 1)^T. \quad (57)$$

1006 In Alg. 2 we provide pseudocode for a practical implementation of ALL, omitting details such as minibatch use
 1007 for clarity. Also refer to this link for a minimal self-contained PyTorch (Paszke et al., 2019) implementation.

⁴ignoring its $\boldsymbol{\theta}$ -dependence because differentiating with respect to $\boldsymbol{\theta}$ is straightforward here

Algorithm 2: Simplified implementation of our Adversarial Leakage Localization (ALL) algorithm.

Input: Dataset $D := \{(\mathbf{x}^{(n)}, y^{(n)}) : n = 1, \dots, N\} \subseteq \mathbb{R}^T \times \mathcal{Y}$, mask-conditional classifier architecture $\tilde{\Phi} : \mathbb{R}^T \times [0, 1]^T \rightarrow \mathbb{R}^{|\mathcal{Y}|}$, initial classifier weights $\theta_0 \sim \mathbb{R}^P$, initial pre-constraint occlusion logits $\tilde{\eta}_0 \in \mathbb{R}^T$, occlusion budget $\bar{\gamma} \in (0, 1)$

Output: Per-timestep ‘leakiness’ estimate $\gamma^* \in [0, 1]^T$

```

1 Function getOcclProbLogits ( $\tilde{\eta} \in \mathbb{R}^T$ : pre-constraint logits of occlusion probabilities)
2   return  $\tilde{\eta} - \text{logsumexp}(\tilde{\eta}) + \log T + \log \bar{\gamma} - \log(1 - \bar{\gamma})$ 
3 Function sampleFromCONCRETE ( $\tilde{\gamma} \in (0, 1)^T$ : logits of occlusion probabilities)
4   return sigmoid(logsigmoid( $\tilde{\gamma}$ ) - logsigmoid( $-\tilde{\gamma}$ ) + log  $\mathbf{u}$  - log( $\mathbf{1} - \mathbf{u}$ )),  $\mathbf{u} \sim \mathcal{U}(0, 1)^T$ 
5 Function getMaskedCrossEntropy ( $\theta \in \mathbb{R}^P$ : classifier weights,  $(\mathbf{x}, y) \in \mathbb{R}^T \times \mathcal{Y}$ : input and label,
    $\alpha \in [0, 1]^T$ : relaxed input mask)
6   return logsoftmax  $\tilde{\Phi}_\theta(y \mid (\mathbf{1} - \alpha) \odot \mathbf{x} + \alpha \odot \varepsilon, \mathbf{1} - \alpha)$ ,  $\varepsilon \sim \mathcal{N}(0, 1)^T$ 
7 for  $t = 0, 1, \dots$  until convergence do
8    $(\mathbf{x}_t, y_t) \leftarrow \text{sampleDatapoint}(D)$ 
9    $\tilde{\gamma}_t \leftarrow \text{getOcclProbLogits}(\tilde{\eta}_t)$ 
10   $\alpha_t \leftarrow \text{sampleFromCONCRETE}(\tilde{\gamma}_t)$ 
11   $\ell_t \leftarrow \text{getMaskedCrossEntropy}(\theta_t, (\mathbf{x}_t, y_t), \alpha_t)$ 
12   $\mathbf{g}_t^\theta \leftarrow \nabla_\theta \ell_t$ ,  $\mathbf{g}_t^\eta \leftarrow -\nabla_{\tilde{\eta}} \ell_t$ ,  $\theta_{t+1} \leftarrow \text{OptStep}(\theta_t, \mathbf{g}_t^\theta)$ ,  $\tilde{\eta}_{t+1} \leftarrow \text{OptStep}(\tilde{\eta}_t, \mathbf{g}_t^\eta)$ 
13 return sigmoid getOcclProbLogits( $\tilde{\gamma}_{t+1}$ )

```

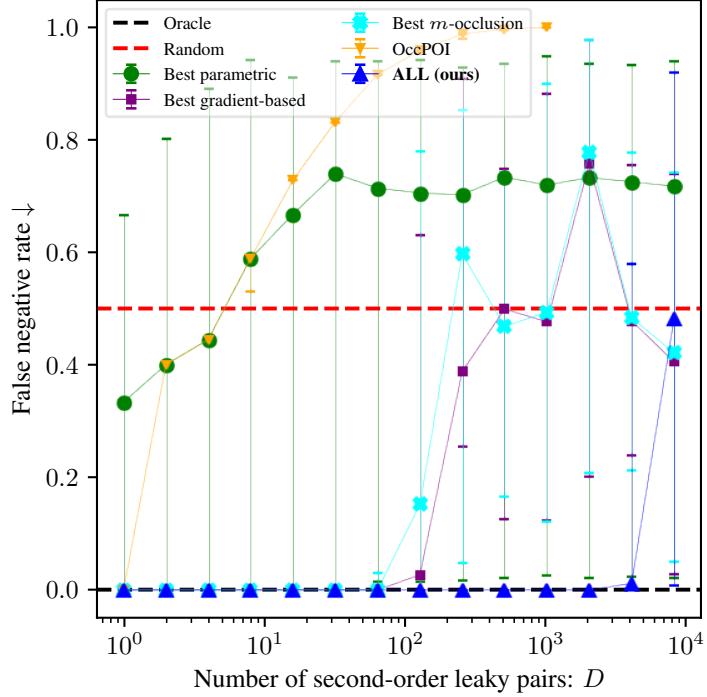
1008 **E Extended experimental details and results**1009 **E.1 Toy setting where our method succeeds and prior work fails**

Figure 9: A toy setting where ALL (ours) significantly outperforms baselines. We sample 1 non-leaky feature and D second-order leaky pairs, then plot the false negative rate, defined as the proportion of points incorrectly assigned leakiness less than or equal to that of the non-leaky point, as we increase D . ALL (ours) succeeds for D up to 64× higher the best prior deep learning-based approach, and the first-order parametric methods completely fail in this setting. Dots denote median and error bars denote min–max over 5 random seeds.

1010 As previously discussed, first-order parametric statistics-based methods are insensitive to associations of
 1011 order 2 or higher. Prior deep learning-based leakage localization algorithms tend to exploit few of the
 1012 available 2^{T-1} associations between X_t and Y given subsets of $\{X_1, \dots, X_T\} \setminus \{X_t\}$, with many of them
 1013 using only the maximal conditioning set $\{X_1, \dots, X_T\} \setminus \{X_t\}$ itself. This creates issues when there is a large
 1014 number of leaky measurements and the individual contribution of each is ‘drowned out’ in the sense that
 1015 $\mathbb{E}[Y; X_t | \{X_1, \dots, X_T\} \setminus \{X_t\}]$ vanishes. Here we construct a simple setting where both of these issues are
 1016 present, and demonstrate that ALL succeeds whereas the prior approaches face issues.

1017 We generate a sequence of binary-label $2D + 1$ -feature classification datasets consisting of ordered pairs
 1018 (\mathbf{X}, Y) sampled independently as follows:

$$Y \sim \mathcal{U}\{0, 1\}, R \sim \mathcal{U}\{0, 1\}, M_i \sim \mathcal{U}\{0, 1\} : i = 1, \dots, D, \quad (58)$$

$$X_R \sim \mathcal{N}(R, 1), X_{M_i} \sim \mathcal{N}(M_i, 1), X_{Y \oplus M_i} \sim \mathcal{N}(Y \oplus M_i, 1) : i = 1, \dots, D. \quad (59)$$

1019 Here we denote by \oplus the exclusive-or operation and $\mathbf{X} \equiv (X_R, X_{M_1}, X_{Y \oplus M_1}, \dots, X_{M_D}, X_{Y \oplus M_D})$. Intuitively,
 1020 we can view X_R as a noisy observation of R and each $X_{M_i}, X_{Y \oplus M_i}$ as noisy observations of $M_i, Y \oplus M_i$,
 1021 respectively. Here the variable Y is analogous to a targeted sensitive variable, and the values $(M_i, Y \oplus M_i)$
 1022 are analogous to the pairs of second-order leaky variables which arise in Boolean masked implementations
 1023 such as Benadjila et al. (2020).

1024 Clearly R , and thus X_R , tells us nothing about Y . Additionally, the values M_i in isolation tell us nothing
 1025 about Y . Similarly, the values $Y \oplus M_i$ in isolation tell us nothing about Y because

$$\mathbb{P}(Y = 0 \mid Y \oplus M_i = 0) = \mathbb{P}(M_i = 0) = \frac{1}{2} = \mathbb{P}(M_i = 1) = \mathbb{P}(Y = 1 \mid Y \oplus M_i = 0) \quad (60)$$

1026 and similarly

$$\mathbb{P}(Y = 0 \mid Y \oplus M_i = 1) = \mathbb{P}(M_i = 1) = \frac{1}{2} = \mathbb{P}(M_i = 0) = \mathbb{P}(Y = 1 \mid Y \oplus M_i = 1). \quad (61)$$

1027 Despite this, given the *pair* of values $\{M_i, Y \oplus M_i\}$ we can recover Y via the identity

$$M_i \oplus M_i = 0 \implies (Y \oplus M_i) \oplus M_i = Y. \quad (62)$$

1028 Thus, we would like a leakage localization algorithm to indicate that X_R is non-leaky and the values
 1029 $X_{M_i}, X_{Y \oplus M_i}$ are leaky.

All experiments use the hyperparameters shown in Table 5.

Table 5: Hyperparameters used for our toy Gaussian dataset experiments. We use the default PyTorch settings except where specified. [†]These hyperparameters apply only to **ALL**. Other hyperparameters are used both for **ALL** and for our baseline methods.

| Hyperparameter | Value |
|---|---|
| Classifier architecture | ReLU MLP with 1×500 -neuron hidden layer |
| Classifier optimizer | AdamW |
| Classifier learning rate | 10^{-4} |
| Classifier weight decay | 0.01 |
| Dataset size | 10k |
| Training steps | 5k |
| Minibatch size | 800 |
| Noise distribution learning rate [†] | 10^{-3} |
| Budget $\bar{\gamma}^{\dagger}$ | $1 - 2^{-0.1 \cdot D - 1}$ |

1030

1031 Results can be seen in Fig. 9. We measure the performance of methods by the percent of measurements
 1032 in $\{X_{M_i} : i = 1, \dots, D\} \cup \{X_{Y \oplus M_i} : i = 1, \dots, D\}$ assigned a leakiness greater than or equal to that of
 1033 X_R . For clarity we report for each D the best result out of SNR, SOSD and CPA as ‘best parametric’,
 1034 out of GradVis, Saliency, Input * Grad, and LRP as ‘best gradient-based’, and the best result out of
 1035 $\{2m + 1 : m \in [0..24], m < 2D + 1\}$ -Occlusion as ‘best m -Occlusion’. Due to their high cost we use OccPOI
 1036 only for D up to 256, and we do not use 2nd-order m -Occlusion. Note that in subsequent experiments
 1037 2nd-order m -Occlusion does not significantly outperform first-order occlusion, and OccPOI performs poorly
 1038 due to identifying only a small number of leaky measurements rather than assigning a leakiness to every
 1039 measurement.

1040 **E.2 Simulated AES datasets where we have ground truth knowledge about leakage**

1041 Here we present experiments done on synthetic AES power traces. These are a useful complement to the
 1042 experiments on real datasets because 1) here we have ground truth knowledge about which timesteps are
 1043 leaking, which we can use to validate our model’s output, 2) we can generate infinitely-large datasets to
 1044 eliminate dataset size as a confounding variable in results, and 3) we can observe the change in our technique’s
 1045 behavior as we individually vary particular dataset properties such as low-pass filtering strength and leaky
 1046 instruction count.

Algorithm 3: Pseudocode for our synthetic data generation procedure, based on the Hamming weight leakage model of Mangard et al. (2007). For clarity we omit the random delay and shuffling procedures, but these are straightforward and may be found in our code.

Input: Dataset size $N \in \mathbb{Z}_{++}$,
 Timesteps per power trace $T \in \mathbb{Z}_{++}$,
 Bit count $n_{\text{bits}} \in \mathbb{Z}_{++}$,
 Operation count $n_{\text{ops}} \in \mathbb{Z}_{++}$,
 Data-dependent noise variance $\sigma_{\text{data}}^2 \in \mathbb{R}_+$,
 Operation-dependent noise variance $\sigma_{\text{op}}^2 \in \mathbb{R}_+$,
 Residual noise variance $\sigma_{\text{resid}}^2 \in \mathbb{R}_+$,
 Low-pass filtering strength $\beta \in [0, 1)$,
 Leaking timestep count $n_{\text{lkg}} \in \mathbb{Z}_+$

Output: Synthetic dataset $\mathcal{D} \subseteq \mathbb{R}^T \times [0 \dots 2^{n_{\text{bits}}} - 1]$

```

1  $\{k^{(n)} : n \in [1 \dots N]\} \sim \mathcal{U}(\{0, 1\}^{n_{\text{bits}}})^N$  // AES keys
2  $\{w^{(n)} : n \in [1 \dots N]\} \sim \mathcal{U}(\{0, 1\}^{n_{\text{bits}}})^N$  // plaintexts
3  $\{o_t : t \in [1 \dots T]\} \sim \mathcal{U}([1 \dots n_{\text{ops}}])^N$  // operations
4  $\{\tilde{x}_{\text{op}, i} : i \in [1 \dots n_{\text{ops}}]\} \sim \mathcal{N}(0, \sigma_{\text{op}}^2)^{n_{\text{ops}}}$  // operation-dependent power consumption
5  $\mathcal{T}_{\text{lkg}} \sim \mathcal{U}\left(\frac{[1 \dots T]}{n_{\text{lkg}}}\right)$  // timesteps of leaky instructions
6 for  $n \in [1 \dots N]$  do
7    $y^{(n)} \leftarrow \text{Sbox}(k^{(n)} \oplus w^{(n)})$  // targeted variable: first SubBytes output
8    $\mathbf{x}_{\text{resid}}^{(n)} \sim \mathcal{N}(0, \sigma_{\text{resid}}^2)^T$  // residual power consumption
9   for  $t \in \mathcal{T}_{\text{lkg}}$  do
10     $d_t^{(n)} \leftarrow y^{(n)}$  // leaky timesteps: targeted variable is the data
11   for  $t \in [1 \dots T] \setminus \mathcal{T}_{\text{lkg}}$  do
12     $d_t^{(n)} \sim \mathcal{U}(\{0, 1\}^{n_{\text{bits}}})$  // rest of data treated as random
13   for  $t \in [1 \dots T]$  do
14     $x_{\text{data}, t}^{(n)} \leftarrow \sigma_{\text{data}} \left(4 - \text{HW}(d_t^{(n)})\right) / \sqrt{2}$  // data-dependent power consumption
15    $\mathbf{x}^{(n)} \leftarrow \mathbf{x}_{\text{data}}^{(n)} + \mathbf{x}_{\text{op}} + \mathbf{x}_{\text{resid}}^{(n)}$  // total power consumption
16   for  $t \in [2 \dots T]$  do
17     $x_t^{(n)} \leftarrow \beta x_{t-1}^{(n)} + (1 - \beta)x_t^{(n)}$  // Discrete low-pass filtering of  $\mathbf{x}^{(n)}$ .
    // In practice we prepend 'burn-in' timesteps to allow transient effects to decay.
18 return  $\{(\mathbf{x}^{(n)}, y^{(n)}) : n \in [1 \dots N]\}$ 

```

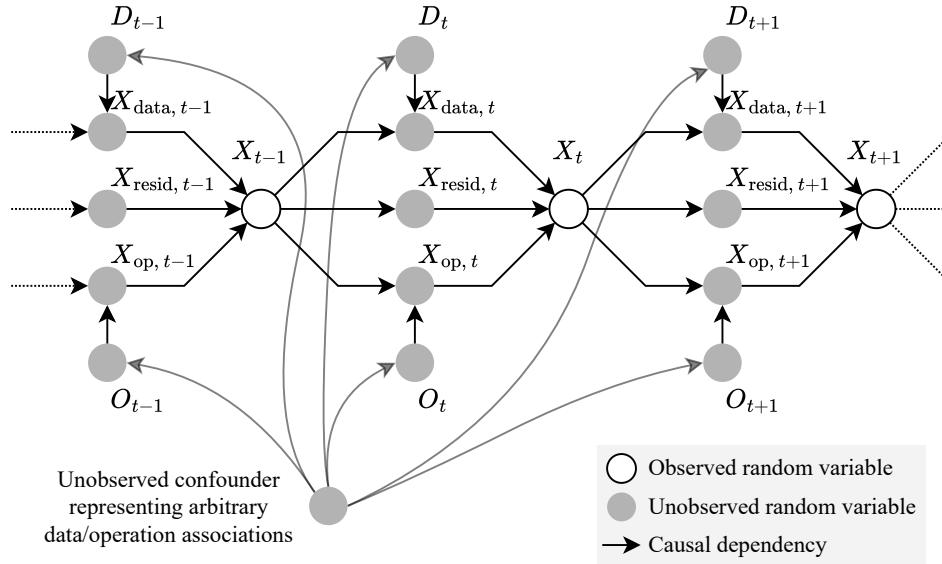


Figure 10: Causal diagram representing the assumed data-generating process for our synthetic AES datasets. We assume this is an I-map for the data-generating process – i.e. we assume independence conditions present in the diagram hold for the data-generating process, but the data-generating process might have additional independence conditions not reflected here. This is an effort to make more precise the Hamming weight leakage model of (Mangard et al., 2007, ch. 4). We represent the data of the AES algorithm by the random variables D_t , its operations O_t , and its power consumption X_t , which is further broken down into data-dependent power consumption $X_{\text{data}, t}$, operation-dependent power consumption $X_{\text{op}, t}$, and ‘residual’ power consumption $X_{\text{resid}, t}$ (e.g. due to random noise, other processes running in parallel, etc.). We assume that only the composite power consumption X_t is observed.

1047 **E.2.1 Data generation procedure**

1048 We base our synthetic data generation procedure on the Hamming weight leakage model of (Mangard et al.,
 1049 2007, ch. 4), which we will subsequently describe.⁵ As above, let us represent our power/EM traces as
 1050 a random vector $\mathbf{X} := (X_t : t = 1, \dots, T)$ with range \mathbb{R}^T . We represent the cryptographic algorithm as
 1051 sequences of data $\mathbf{D} := (D_t : t = 1, \dots, T)$ and operations $\mathbf{O} := (O_t : t = 1, \dots, T)$, where each D_t has range
 1052 $\{0, 1\}^{n_{\text{bits}}}$ and each O_t has range $[1 \dots n_{\text{ops}}]$ for some $n_{\text{bits}}, n_{\text{ops}} \in \mathbb{Z}_{++}$. For each $t \in [1 \dots T]$, we can decompose

$$X_t = X_{\text{data}, t} + X_{\text{op}, t} + X_{\text{resid}, t} \quad (63)$$

1053 with dependency structure illustrated in the causal diagram of Fig. 10. Here we have represented by $X_{\text{data}, t}$
 1054 the data-dependent component of power consumption, $X_{\text{op}, t}$ the operation-dependent component of power
 1055 consumption, and $X_{\text{resid}, t}$ the ‘residual’ power consumption due to random noise, other processes running
 1056 simultaneously on the same hardware, etc.

1057 Mangard et al. (2007) experimentally characterized the power consumption of a cryptographic device and found
 1058 that it is reasonable to approximate $X_{\text{data}, t}$ as Gaussian noise with D_t -dependent mean, $X_{\text{op}, t}$ as Gaussian
 1059 noise with O_t -dependent mean, and $X_{\text{resid}, t}$ as Gaussian noise with constant mean (which we will assume to
 1060 be zero because it contains no information and is thus irrelevant). For their device, they found that the mean
 1061 of $X_{\text{data}, t}$ was roughly proportional to $n_{\text{bits}} - \text{HW}(D_t)$ where $\text{HW} : \{0, 1\}^{n_{\text{bits}}} \rightarrow [0 \dots n_{\text{bits}}] : \mathbf{x} \mapsto \sum_{k=1}^{n_{\text{bits}}} x_k$.
 1062 We adopt these approximations for our synthetic dataset experiments. See Alg. 3 for pseudocode giving a
 1063 simplified version of our data generation procedure, and refer to our code for full details.

1064 We simulate several factors of variation which can reasonably be expected to occur in realistic settings. We
 1065 apply discrete-time low-pass filtering to the power traces, to simulate the low-pass filtering which occurs due
 1066 to measurement apparatus as well as the fact that power consumption does not change instantaneously in
 1067 real circuits. We allow for the presence of multiple leaky instructions. Additionally, we simulate random
 1068 delays to the leaky instruction which might result from countermeasures such as random no-op insertion
 1069 (Coron & Kizhvatov, 2009), and random ‘shuffling’ where the leaky instruction is randomly placed at one of
 1070 several points in time, as done in Masure & Strullu (2023).

1071 We emphasize that these approximations are specific to the device studied by Mangard et al. (2007) and may
 1072 hold to a limited extent or not at all for other devices. For example, the Hamming weight dependence of
 1073 power consumption stems from the fact that their device ‘pre-charges’ all its data bus lines to 1, then drains
 1074 the charge from the lines which should represent 0, thereby consuming power proportional to the number
 1075 of lines which represent 0. Many devices operate differently, and cryptographic hardware is often designed
 1076 with the explicit goal of obfuscating this data/power consumption dependence. Thus, while we expect all
 1077 real cryptographic devices to have some exploitable dependence between data and power consumption given
 1078 sufficient quality and quantity of data, in many cases the nature of the dependence will likely elude a simple
 1079 characterization such as this.

1080 **E.2.2 Experimental details**

1081 We run experiments on many variations of this dataset and verify that **ALL** produces outputs which align with
 1082 our expectations. For all these experiments, our classifier Φ_{θ} is a 3-layer multilayer perceptron with hidden
 1083 dimension 500, ReLU activations, an input dropout rate of 0.1, hidden dropout rate of 0.2, and pre-logits
 1084 dropout rate of 0.3. We generate data continuously so that our dataset size is effectively infinite. Additional
 1085 hyperparameters are listed in table 6, and default settings corresponding to Alg. 3 are listed in table 7, which
 1086 we use except where otherwise stated.

1087 **E.2.3 Results**

1088 We vary 4 parameters of the data generation process of Alg. 3 and observe its effect on the output of **ALL**.
 1089 We find that **ALL** consistently produces results consistent with our expectations given the timestep(s) of

⁵For clarity we alter the notation and explicitly define a causal structure for the data-generating process. The decomposition of power consumption in Eqn. 63 and the definition of $X_{\text{data}, t}$ in terms of the Hamming weight of the data are based on Mangard et al. (2007), but the additional details are our own.

Table 6: List of hyperparameters used for experiments on synthetic AES datasets. We use the default PyTorch settings unless otherwise stated. Note that in general for ALL classifiers we disable input dropout, but for these experiments the dropout rate was set to 0.1 due to an oversight.

| Hyperparameter | Value |
|--|--|
| Classifier architecture | ReLU MLP with 3×500 -neuron hidden layers |
| Input dropout | 0.1 |
| Hidden dropout | 0.2 |
| Output dropout | 0.3 |
| Optimizer for both θ and $\tilde{\eta}$ | <code>torch.optim.AdamW</code> |
| Weight decay for θ | 0.01 for weights, 0 for biases |
| Weight decay for $\tilde{\eta}$ | 0 |
| Training steps | 10^4 |
| Minibatch size | 10^3 |
| Noise budget $\bar{\gamma}$ | 0.5 |
| Weight initializer | <code>torch.nn.init.xavier_uniform_</code> |

Table 7: Default synthetic AES dataset configuration, corresponding to the inputs of Alg. 3. Subsequent experiments will use these settings unless otherwise stated.

| Setting | Value |
|---|----------|
| Dataset size N | ∞ |
| Timesteps per power trace T | 101 |
| Bit count n_{bits} | 8 |
| Operation count n_{op} | 32 |
| Data-dependent noise variance σ_{data}^2 | 1.0 |
| Operation-dependent noise variance σ_{op}^2 | 1.0 |
| Residual noise variance σ_{resid}^2 | 1.0 |
| Number of leaky instructions n_{lkg} | 1 |
| Low-pass filtering strength β | 0.5 |
| Maximum random delay size | 0 |
| Possible leaky timestep location count (i.e. shuffling) | 1 |

1090 the leaky instruction(s), and that the variance of its output is quite low in this context where we have an
 1091 infinitely-large dataset and a long training duration. Here we describe the parameters being swept and justify
 1092 the output of ALL given this.

1093 **Low pass filtering strength β** Recall that we are discrete low-pass filtering traces via the recursive
 1094 function $x_t^{\text{lpf}} := (1 - \beta)x_t + \beta x_{t-1}$. See the first row of Fig. 11, where from left to right β takes on the
 1095 values 0, 0.5, 0.75, 0.875, 0.9375, 0.96875, 0.984375, 0.9921875. Note that the peak estimated leakiness always
 1096 corresponds to the ground-truth leaky instruction timestep. As we increase β we see that measurements to
 1097 the left and right are assigned high leakiness as well, which makes sense for the following reason:

1098 Let us denote by t^* the timestep at which the leaky instruction was executed. We then have $X_{t^*} =$
 1099 $c_1 \text{HW}(Y) + c_2 U_{t^*} + c_3 X_{t^*-1}$ where U_{t^*} denotes a ‘residual’ random variable independent of Y , and c_1, c_2 are
 1100 appropriate constants. Note that $X_{t^*+1} = \beta X_{t^*} + (1 - \beta)U_{t^*+1} = \beta c_1 \text{HW}(Y) + \beta c_2 U_{t^*} + (1 - \beta)U_{t^*+1}$, so
 1101 X_{t^*+1} also leaks Y . Recursively it is clear that the same can be said for $X_{t^*+2}, X_{t^*+3}, \dots$. Less-intuitively,
 1102 X_{t^*-1} also leaks Y . This is because although X_{t^*-1} is *marginally* independent of Y , $X_{t^*} - c_3 X_{t^*-1}$ has a
 1103 higher correlation with $\text{HW}(Y)$ than X_{t^*} does – i.e. X_{t^*-1} is dependent on the Y -independent noise of X_{t^*} ,
 1104 and can be used to reduce the noise. Recursively, since $X_{t^*-2}, X_{t^*-3}, \dots$ are correlated with X_{t^*-1} , they
 1105 also leak Y by the same mechanism.

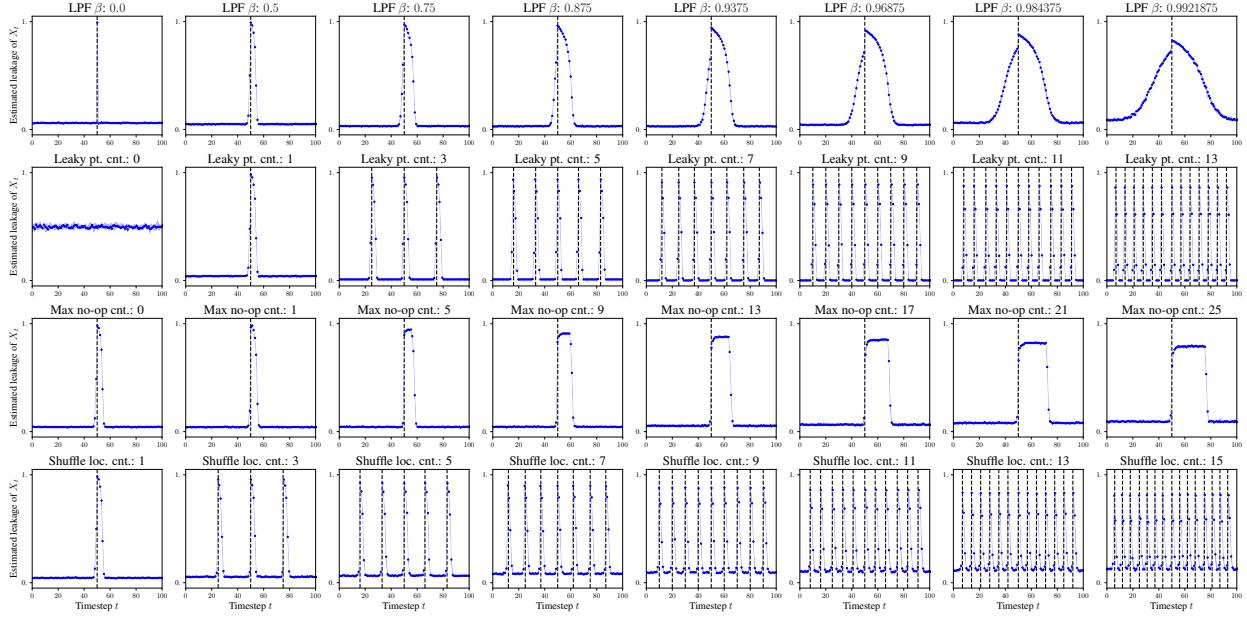


Figure 11: Results of applying ALL (ours) to synthetic AES datasets with varying parameters. The estimated leakage by ALL is denoted by blue dots and the ground truth timestep of the leaking instruction is denoted by the vertical black dotted lines. Blue dots denote mean and shading denotes median over 5 random seeds. Note that the ALL output is consistent with the ground truth leaky instruction timestep in all cases, and the variance between runs is quite low in the infinite data regime. **(first row)** Increasing low pass filtering strength β from left to right. **(second row)** Increasing number of leaky instructions from left to right. **(third row)** Increasing random delay insertion from left to right. **(fourth row)** Increasing number of possible shuffling locations for the leaky point from left to right.

1106 **Leaky point count** Here we sweep the number of leaky instructions. See the second row of Fig. 11 where
 1107 from left to right the leaky point count n_{lkg} takes on the values 0, 1, 3, 5, 7, 9, 11, 13. As expected, all leaky
 1108 instruction timesteps correspond to a peak of ALL-estimated leakiness.

1109 **Random delay size** Here we insert random delays – i.e. instead of always occurring at time t , the leaky
 1110 instruction occurs at $t + u$ where $u \sim \mathcal{U}\{0, \dots, d_{\text{max}}\}$. See the third row of Fig. 11 where from left to right
 1111 d_{max} takes on values 0, 1, 5, 9, 13, 17, 21, 25. We see that the estimated leakiness becomes ‘spread out’ over
 1112 the interval $[t \dots t + d_{\text{max}}]$.

1113 **Shuffle location count** Here we randomly ‘shuffle’ the leaky instruction – i.e. instead of always occurring
 1114 at time t , the leaky instruction occurs at $t' \sim \mathcal{U}\{t_1, \dots, t_{n_{\text{shuff}}}\}$. See the fourth row of Fig. 7 where from left to right
 1115 n_{shuff} takes on values 1, 3, 5, 7, 9, 11, 13, 15. We see that the leakiness becomes ‘spread out’ over
 1116 the set of timesteps at which the leaky instruction might occur.

1117 E.3 Experiments on real power and EM radiation leakage datasets

1118 Here we run experiments on a variety of publicly-available side-channel attack datasets, where we attempt to
 1119 localize leakage of their canonical target variable.

1120 E.3.1 Datasets

1121 We compare ALL with prior work on the 6 datasets described in Table 8, which consist of traces of power
 1122 and EM radiation measurements and associated cryptographic variables recorded from real implementations.
 1123 Note that we evaluate on AES, ECC and RSA implementations implemented on several MCUs and an FPGA

Table 8: A list of the datasets used in our paper, with a summary of their salient attributes. Note that our experiments cover a variety of settings: AES, RSA and ECC implementations on both microcontrollers (MCUs) and a field-programmable gate array (FPGA), both power and EM radiation traces, and various types of countermeasures. For all datasets we localize leakage of the canonical target variable. We denote by subscripts the targeted byte of the variable. We denote by $k_n, w_n, m_n, k_n^*, c_n$ the n -th byte (counting from 0) of the AES key, plaintext, mask, last round key, and ciphertext, respectively. [†]As described below, we deviate from the canonical profiling/attack split.

| | | | |
|--------------------------------------|--|------------------------------------|--|
| Dataset | ASCADv1 (fixed key) | ASCADv1 (variable key) | DPAv4 (Zaid version) |
| Citation | Benadjila et al. (2020) | Benadjila et al. (2020) | Bhasin et al. (2014); Zaid et al. (2020) |
| Link | (here) | (here) | (here) |
| Algorithm | AES-128 | AES-128 | AES-128 |
| Hardware | ATMega8515 (MCU) | ATMega8515 (MCU) | ATMega163 (MCU) |
| Emission measured | Power | Power | Power |
| Countermeasures | Boolean masking | Boolean masking | Rotating Sbox mask (known) |
| Targeted variable | $Sbox(k_3 \oplus w_3)$ | $Sbox(k_3 \oplus w_3)$ | $Sbox(k_0 \oplus w_0)$ |
| Dataset size (profile/attack) | 50k/10k | 200k/100k | 3k/500 [†] |
| Feature count T | 0.7k | 1.4k | 4k |
| Dataset | AES-HD | One Trace is All it Takes (OTiAiT) | One Truth Prevails (OTP) (1024-bit) |
| Citation | Bhasin et al. (2020) | Weissbart et al. (2019) | Saito et al. (2022) |
| Link | (here) | (here) | (here) |
| Algorithm | AES-128 | EdDSA w/ Curve2559 | 1024-bit RSA-CRT |
| Hardware | XiLinx Virtex-5 (FPGA) | STM32F4 (MCU) | STM32F4 (MCU) |
| Emission measured | EM radiation | Power | EM radiation |
| Countermeasures | None | None | Dummy load |
| Targeted variable | $Sbox^{-1}(k_{11}^* \oplus c_{11}) \oplus c_7$ | Ephemeral key nibble | Dummy load? |
| Dataset size (profile/attack) | 50k/25k | 5.12k/1.28k | 100k/98.304k [†] |
| Feature count T | 1.25k | 1k | 1k |

1124 with various target variables. **ALL** as well as most of our baseline algorithms are in principle agnostic to most
 1125 of these details, requiring only a supervised learning-style dataset with power traces and the associated value
 1126 of the targeted variable as labels. Through these experiments we demonstrate that this is true in practice
 1127 across a diverse array of settings.

1128 Note that the ASCADv1 datasets have primarily second-order leakage due to their Boolean masking
 1129 countermeasure, whereas the other 4 datasets have primarily first-order leakage. Our comparisons include
 1130 both deep learning methods as well as simple first-order parametric methods which are widely used due to
 1131 their low cost and interpretability. We find that the latter are competitive or superior to the deep learning
 1132 methods on the first-order datasets but perform significantly worse on the second-order datasets due to failing
 1133 to exploit second-order leakage. Our experiments do not compellingly show that deep learning methods
 1134 improve on these simpler methods for first-order datasets, but we nonetheless include them as additional
 1135 points of comparison between the deep learning methods and to show that **ALL** works in a variety of settings.

1136 In general we use the canonical target variable and profiling/attack dataset split (note that in the context
 1137 of profiling side-channel analysis the training dataset is called the profiling dataset, and the test dataset is
 1138 called the attack dataset). We deviate from the canonical dataset configuration in the following cases:

- 1139 • We find that the canonical attack dataset of DPAv4 is too small to compute useful oracle leakiness
 1140 assessments. We thus concatenate the canonical 4.5k-trace profiling dataset and 0.5k-trace attack
 1141 dataset into a single 5k-length dataset, and use the first 3k traces for profiling and the last 2k to
 1142 compute the oracle assessments. Since some of our experiments require metadata which is only
 1143 available for the attack dataset, for everything other than oracle assessment computation we use the
 1144 canonical attack dataset and leave the remaining 1.5k traces unused.
- 1145 • We use the version of DPAv4 which was preprocessed and distributed by Zaid et al. (2020) here
 1146 rather than the original version. This version has shortened traces which have been cropped around
 1147 the leaky instruction, and has the rotating Sbox mask effectively ‘disabled’ by providing the masked
 1148 SubBytes variable as the target.

1149 • The One Truth Prevails (OTP) dataset consists of approximately 64M traces and has a high label
 1150 imbalance. To save computational resources we extract a 100k-trace randomly-selected balanced
 1151 subset, which we find is more than sufficient for strong supervised classification and leakage localization
 1152 performance. See our code for details.

1153 **E.3.2 Implementation details for the leakage localization algorithms**

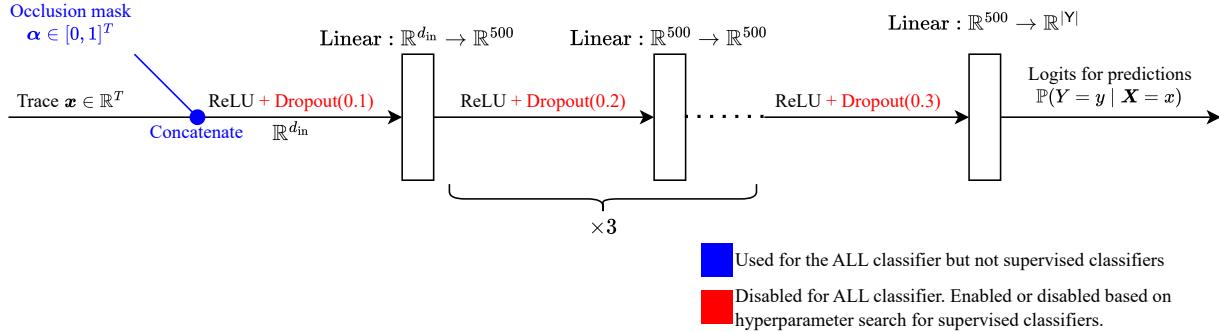


Figure 12: Diagram of the multilayer perceptron architecture used for classifiers in the deep learning methods, based on the architecture proposed in Wang et al. (2017) for time-series classification.

1154 For ALL and all considered baseline methods the classifier uses the simple ReLU + Dropout MLP architecture
 1155 of Wang et al. (2017) shown in Fig. 12, which has 3 500-neuron hidden layers, input dropout rate of 0.1,
 1156 hidden dropout rate of 0.2, and output dropout rate of 0.3. For the deep learning baselines we enable or
 1157 disable the input, hidden and output dropout based on our hyperparameter search outcome, and for ALL we
 1158 leave it disabled. The classifier for ALL takes the occlusion mask as an auxiliary input, which we implement
 1159 by concatenating it with the masked trace.

1160 We also explored convolutional architectures, but preliminary experiments indicated that these achieved
 1161 weaker classification and leakage localization performance across the board, as well as training more-slowly
 1162 than the MLP architecture due to a higher layer count. We suspect that the inductive biases of convolutional
 1163 layers are not useful for the datasets we consider. As a sanity check for this design choice, we run the deep
 1164 learning baseline methods using both our MLP architecture and a handful of open-weight classifiers which
 1165 were released with Wouters et al. (2020) on the datasets for which they are available.

1166 All deep learning methods are implemented in PyTorch (Paszke et al., 2019). Most non-deep methods are
 1167 implemented with Numpy (Harris et al., 2020), with a handful of compute-intensive methods implemented
 1168 with Numba (Lam et al., 2015). We use Scipy (Virtanen et al., 2020) implementations of statistical methods
 1169 where available.

1170 We use the AdamW optimizer (Loshchilov & Hutter, 2018; Kingma & Welling, 2014) with the default PyTorch
 1171 settings $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\lambda = 0.01$, $\epsilon = 10^{-8}$ and the learning rate chosen through hyperparameter
 1172 search. Weight decay is applied only to the weights of the linear layers, not to the biases. We use a
 1173 minibatch size of 256. Weights are initialized with the uniform Glorot initialization (Glorot & Bengio, 2010)
 1174 `torch.nn.init.xavier_uniform_` (default for Keras) rather than the default PyTorch initialization, and
 1175 we find that this is a critical detail – many supervised classifier runs on the ASCADv1 datasets completely
 1176 fail to generalize beyond the training dataset when the default PyTorch weight initialization is used. We
 1177 randomly set aside 20% of the profiling datasets for validation and use the remaining 80% for training. We
 1178 standardize the traces as $x \mapsto \frac{x - \mu}{\max(\sigma, 10^{-6})}$ where μ and σ denote the elementwise sample mean and standard
 1179 deviation computed using the profiling dataset.

1180 In general we measure the performance of supervised classifiers with their *mean rank* rather than their
 1181 accuracy, as accuracy tends to be low and too-coarse in the context of side-channel analysis. Given a label
 1182 $y \in \mathbb{Y}$ and predicted label distribution $\hat{p}_Y \in \Delta^{|\mathbb{Y}|-1}$ (e.g. the softmaxed output of a classifier neural net), we
 1183 define the rank as the number of possible labels assigned at least as much probability mass as the true label,

1184 i.e.

$$\text{Rank}(\hat{p}_Y; y) := |\{y' \in \mathcal{Y} : \hat{p}_Y(y') \geq \hat{p}_Y(y)\}|. \quad (64)$$

1185 This metric has range $[1 \dots |\mathcal{Y}|]$, with lower being better.

1186 **Implementation of the baseline methods** We use Captum (Kokhlikyan et al., 2020) implementations
 1187 of the Saliency, Input * Grad, LRP, and m -Occlusion methods. We implement SNR, SOSD, CPA, GradVis,
 1188 and 2nd-order m -Occlusion ourselves, and we implement a PyTorch version of OccPOI based on the Keras
 1189 implementation released by the authors here.

1190 Note the following choices we have made in implementing and evaluating OccPOI (Yap et al., 2025):

- 1191 • OccPOI differs from the other methods we consider in that rather than assigning a leakiness value to
 1192 every measurement, it aims to identify a non-unique subset of measurements which are sufficient
 1193 for a classifier to attain some specified performance level when all other measurements are occluded.
 1194 Since our evaluation metrics require a leakiness value for every measurement, we assign a leakiness of
 1195 0 to measurements not identified by OccPOI.
- 1196 • Their method uses the attack dataset for probing the classifier's sensitivity to input features, which
 1197 introduces data contamination in the context of our performance metrics. We cannot easily use the
 1198 validation dataset instead because the attack procedure requires having many traces corresponding to
 1199 a fixed AES key, which is generally only available for the attack dataset. We find that **ALL** significantly
 1200 outperforms OccPOI despite this contamination, so as the aim of our work is to demonstrate the
 1201 efficacy of **ALL**, we allow OccPOI to use the attack dataset.
- 1202 • Since OTiAiT and OTP do not have attack datasets which facilitate this kind of multi-trace
 1203 prediction, for these datasets we simply use the mean rank of the classifier on the attack dataset as
 1204 our performance metric.
- 1205 • Yap et al. (2025) proposes an extension of OccPOI which ranks the leakiness of the points it has
 1206 identified using a 1-occlusion-like strategy. We use this extension in our implementation.
- 1207 • Yap et al. (2025) proposes an extension of OccPOI where they apply it repeatedly on the residual
 1208 measurements not selected during the last iteration. We do not use this extension because it is
 1209 very computationally expensive, requiring $O(T)$ applications of OccPOI which each require $\Omega(T)$
 1210 *non-parallelizable* passes through the attack dataset. In preliminary experiments this extension
 1211 performed better than basic OccPOI, but still far below the performance of the other considered
 1212 methods despite requiring orders of magnitude more wall clock time.
- 1213 • Similarly to Yap et al. (2025), to save compute we use only a subset of the attack datasets for classifier
 1214 evaluation. We use 1.8k traces for ASCADv1-fixed, 2.8k for ASCADv1-variable, 140 for DPAv4, 25k
 1215 for AES-HD, 100 for OTiAiT, and 100 for OTP. These are approximately 10–100 \times the necessary
 1216 number of traces to successfully attack the AES methods.
- 1217 • Yap et al. (2025) define their performance threshold to be that the classifier correctly predicts the
 1218 AES key after accumulating all traces in the attack dataset, and we use the same threshold for the
 1219 ASCADv1 datasets, AES-HD, and DPAv4. For the non-AES dataset OTiAiT and OTP, our threshold
 1220 is that the mean rank of the classifier rises by 0.1 relative to its mean rank when no measurements
 1221 are occluded.

1222 As we will show, OccPOI attains significantly lower performance than other methods according to our
 1223 performance metrics, due to not assigning a leakiness to most measurements. This reflects that the aim of
 1224 Yap et al. (2025) is somewhat different from the present work: Yap et al. (2025) heavily emphasized the
 1225 usefulness of OccPOI as a feature selection tool with leakage localization being an auxiliary goal, whereas our
 1226 work is concerned solely with leakage localization. We consider OccPOI as a baseline because it is similar in
 1227 spirit to our work, but these results demonstrate that it is ill-suited to the task considered in our paper.

1228 We make the following choices in implementing m -Occlusion and 2nd-order m -Occlusion (Schamberger et al.,
 1229 2023):

1230 • Schamberger et al. (2023) explore different ways to occlude measurements and conclude that it
 1231 works well to replace measurements by their average value over the profiling dataset, whereas other
 1232 heuristics such as replacing them with 0 or with Gaussian noise works poorly. We thus replace
 1233 measurements by their mean. Since we are element-wise standardizing the traces, this is the same as
 1234 replacing them by 0.

1235 • Schamberger et al. (2023) propose 2nd-order m -Occlusion as a means of estimating the leakiness of
 1236 *pairs* of windows, which is useful for discerning whether a measurement has first-order leakage or is
 1237 part of a second-order leaking pair of measurements. They do not propose a means of using it to
 1238 estimate the leakiness of individual measurements. We choose to define the leakiness of X_t as the
 1239 *average* leakiness of the pairs $\{\{X_t, X_{t'}\} : t' \in [1..T]\}$.

1240 • Schamberger et al. (2023) uses a large stride for 2nd-order m -Occlusion, and thus get leakiness values
 1241 for windows of measurements rather than single measurements. We set the stride to 1 for consistency
 1242 with our other baselines.

1243 • These methods introduce the occlusion window size m as a new hyperparameter which must be tuned.
 1244 For m -Occlusion we tune m by testing successive odd-numbered window sizes starting from 1 until
 1245 oracle agreement performance starts decreasing, and using the window size which maximizes oracle
 1246 agreement. We denote this as m^* -Occlusion. Note that this introduces some data contamination,
 1247 which we accept because our aim is to demonstrate the efficacy of ALL, and ALL outperforms
 1248 m -Occlusion despite the contamination.

1249 • For 2nd-order m -Occlusion we use the optimal value of m for m -Occlusion. We denote this 2nd-order
 1250 m^* -Occlusion. We don't do another sweep because 2nd-order m -Occlusion is very computationally-
 1251 expensive, requiring $\Theta(T^2)$ passes through the dataset.

1252 • Like Yap et al. (2025), Schamberger et al. (2023) uses the attack dataset to evaluate the sensitivity
 1253 of the classifier to occluding measurements. To avoid data contamination, and for compatibility
 1254 with the OTiAiT and OTP datasets which do not have the same kind of attack dataset as the AES
 1255 implementations, we track the average change in logits over the profiling dataset as we occlude
 1256 measurements (similarly to 1-Occlusion).

1257 E.3.3 Hyperparameter tuning procedure

Table 9: Outcome of a 50-trial random hyperparameter search for the supervised classification models used by the deep learning baseline methods. All trials are early stopped at the point of lowest validation rank, and we choose the hyperparameter configuration which minimizes the lowest validation rank, with ties broken based on validation loss. Models are trained with AdamW and weight decay applied only to the weights of dense and conv1d layers. We use the default PyTorch settings everywhere unless otherwise stated.

| Hyperparameter | Search space | ASCADv1 (fixed) | ASCADv1 (variable) | DPAv4 (Zaid) | AES-HD | OTiAiT | OTP |
|---|---|-------------------|--------------------|-------------------|-------------------|-------------------|-----------------------|
| Learning rate | $\bigcup_{m=1}^9 \bigcup_{n=3}^5 \{m \cdot 10^{-n}\}$ | $3 \cdot 10^{-4}$ | $2 \cdot 10^{-4}$ | $4 \cdot 10^{-3}$ | $9 \cdot 10^{-5}$ | $8 \cdot 10^{-3}$ | $6 \cdot 10^{-3}$ |
| LR schedule | {constant, cos annealing} | cos annealing | cos annealing | constant | cos annealing | cos annealing | cos annealing |
| Input dropout | {0.0, 0.1} | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.1 |
| Hidden dropout | {0.0, 0.2} | 0.2 | 0.2 | 0.2 | 0.2 | 0.2 | 0.0 |
| Output dropout | {0.0, 0.3} | 0.3 | 0.0 | 0.0 | 0.3 | 0.3 | 0.0 |
| Training steps | n/a | 20k | 40k | 10k | 20k | 1k | 1k |
| Classification performance of chosen model | | | | | | | |
| Test rank ↓ | | 101 ± 1 | 79.1 ± 0.2 | 5.6 ± 0.7 | 124.7 ± 0.2 | 1.010 ± 0.007 | 1.00171 ± 0.00007 |
| Test loss ↓ | | 5.59 ± 0.04 | 5.380 ± 0.009 | 3.0 ± 0.2 | 5.65 ± 0.04 | 0.10 ± 0.08 | 0.0085 ± 0.0006 |
| Traces to AES key disclosure ↓ | | 179 ± 71 | 288 ± 126 | 1.4 ± 0.5 | 4355 ± 1915 | n/a | n/a |

1258 All the deep learning methods we consider require hyperparameter tuning. Work in other deep learning
 1259 subfields such as Gulrajani & Lopez-Paz (2021) has emphasized the importance of a fair hyperparameter
 1260 tuning process and realistic model selection criterion when comparing the performance of different algorithms,
 1261 and in this work we aim to follow these recommendations. Accordingly, all methods are tuned with a 50-trial
 1262 random hyperparameter search. Note that while our main performance evaluation metric is the oracle

Table 10: Outcome of a 50-trial random hyperparameter search for adversarial leakage localization. Models are trained with AdamW and weight decay applied only to the weights of the dense and conv1d layers. We use the default PyTorch settings everywhere unless otherwise stated. For ASCADv1 (fixed), ASCADv1 (variable) and AES-HD we find it helpful to ‘pretrain’ the classifier with fixed $\tilde{\eta}$ before beginning the simultaneous phase of training, but for DPAv4, OTiAiT, and OTP we find this unnecessary.

| Hyperparameter | Search space | Selected value | | | | | |
|-----------------------------------|---|-------------------|--------------------|-------------------|-----------|-----------|-------------------|
| | | ASCADv1 (fixed) | ASCADv1 (variable) | DPAv4 (Zaid) | AES-HD | OTiAiT | OTP |
| θ learning rate (pretrain) | $\bigcup_{m=1}^{10} \{m \cdot 10^{-4}\}$ | 10^{-4} | 10^{-4} | n/a | 10^{-3} | n/a | n/a |
| $\tilde{\gamma}$ (pretrain) | n/a | 0.5 | 0.5 | n/a | 0.5 | n/a | n/a |
| θ learning rate | $\bigcup_{m=1}^9 \bigcup_{n=4}^6 \{m \cdot 10^{-n}\}$ | $6 \cdot 10^{-6}$ | $7 \cdot 10^{-5}$ | $2 \cdot 10^{-5}$ | 10^{-4} | 10^{-4} | $4 \cdot 10^{-4}$ |
| $\tilde{\eta}$ learning rate | $\bigcup_{m=1}^9 \bigcup_{n=0}^2 \{m \cdot 10^n\}$ | 50 | 6 | 9 | 20 | 3 | 3 |
| $\tilde{\gamma}$ | $\bigcup_{m=1}^{19} \{0.05 \cdot m\}$ | 0.3 | 0.4 | 0.7 | 0.85 | 0.8 | 0.65 |
| Training steps (pretrain) | n/a | 10k | 20k | 0 | 10k | 0 | 0 |
| Training steps | n/a | 10k | 20k | 10k | 10k | 1k | 1k |

1263 agreement, it would be unrealistic to use this for model selection, even when computed on a validation dataset,
1264 because it relies on ‘white box’ knowledge about cryptographic implementations that we assume not to have
1265 at training time. Instead, for **ALL** we use the model selection criterion proposed in Sec. E.3.5. The prior
1266 deep learning methods are based on ‘interpreting’ a classifier trained with supervised learning, and in line
1267 with prior work we tune its associated hyperparameters to optimize classification performance via minimizing
1268 the early-stopped mean rank on a validation dataset. We also visualize the distribution of results over the
1269 hyperparameter sweep for each method.

1270 **ALL** is sensitive to the noise budget parameter $\tilde{\gamma}$ and the learning rates of the classifier weights θ and the
1271 noise distribution parameter $\tilde{\eta}$. We consistently find that $\tilde{\eta}$ should have a higher learning rate than θ , so
1272 in order to focus the search on better hyperparameter configurations we tune the *ratio* of the learning rate
1273 of $\tilde{\eta}$ to that of θ rather than the learning rate of $\tilde{\eta}$. For ASCADv1 (fixed, variable) and AES-HD we find
1274 that performance is much better if we pretrain the classifier for half of the training steps with fixed $\tilde{\eta} = 0$
1275 and noise budget $\tilde{\gamma} = 0.5$, then proceed as normal for the next half. Thus, for these trials, we first do a
1276 10-trial grid search of the learning rate for θ to minimize mean rank during this pretraining phase, then
1277 tune all hyperparameters as normal using this trained classifier as the starting point for the remaining 40
1278 trials. In preliminary experiments we explored tuning the β_1 , β_2 , ϵ and weight decay strength of the AdamW
1279 optimizer as well as the number of $\tilde{\eta}$ steps per θ step, but chose to leave these fixed for the final search
1280 because they had little impact on performance. See Table 10 for the hyperparameter search space and chosen
1281 configurations for **ALL**.

1282 The deep learning baselines are based on ‘interpreting’ a trained supervised classifier and require tuning this
1283 classifier – we tune its learning rate, dropout rates, and decide whether to use a constant or cosine decay
1284 learning rate schedule. In preliminary experiments we also explored tuning the β_1 , β_2 , ϵ and weight decay
1285 strength λ of the AdamW optimizer, but chose to leave them at fixed values for the final search because
1286 they had little effect on classification performance. See Table 9 for the hyperparameter search space, chosen
1287 configurations, and resulting classification performance for the supervised classifiers.

1288 E.3.4 Performance evaluation methods

1289 Unlike for the experiments on synthetic datasets, here we lack ground truth knowledge about the leakiness of
1290 individual measurements. It is challenging to evaluate the performance of leakage localization algorithms in
1291 this setting, and there is currently no consensus about the best way to do so. We consider 4 quantitative
1292 performance evaluation strategies which are conceptually-similar to performance evaluation strategies of prior
1293 work. To account for the varying ‘shapes’ of leakage assessments returned by the compared methods, all of
1294 our evaluation metrics are sensitive only to the *relative* leakiness assigned to measurements.

1295 **‘Oracle’ leakiness via SNR with relevant leaking first-order variables** The present work is concerned
1296 with ‘black box’ leakage localization algorithms which require only a supervised learning-style dataset of
1297 traces and associated target variable values, and minimal *a priori* knowledge about the cryptographic

1298 implementations being analyzed. However, it is also possible to analyze devices in a ‘white box’ manner
 1299 which does incorporate *a priori* knowledge. In particular, while second-order datasets such as ASCADv1
 1300 (fixed and variable) are not amenable to black box analysis with the first-order parametric methods such as
 1301 SNR, it is possible to study the implementation and ‘decompose’ the second-order leakage into first-order
 1302 leakage of *pairs* of internal variables, then use parametric methods to analyze leakage of these variables
 1303 individually. As our main performance evaluation strategy, we use such a white box analysis to compute
 1304 per-measurement leakiness predictions, which we treat as an ‘oracle’ against which to compare output. This
 1305 is a useful way to validate deep learning methods because 1) it is interpretable and hyperparameter-free, and
 1306 2) it lets us check whether an output is consistent with an analysis which a domain expert might do.

1307 For the ASCADv1 datasets, we use the white box analysis of Egger et al. (2022) to compute ‘oracle’ leakiness
 1308 estimates for each of the measurements. We use the canonical target variable for both datasets, which
 1309 is $\text{Sbox}(k_2 \oplus w_2)$ where Sbox is an invertible function which is publicly-known and shared by all AES
 1310 implementations, \oplus denotes the bitwise exclusive-or operation and k_2 and w_2 denote byte 2 of the key and
 1311 plaintext, respectively, with indexing starting from 0. The underlying AES-128 implementation of these
 1312 datasets uses Boolean masking, as shown in Alg. 1 of Benadjila et al. (2020). By design, this Boolean masking
 1313 prevents the algorithm from ever directly operating on $\text{Sbox}(k_2 \oplus w_2)$, making all power measurements X_t
 1314 nearly marginally statistically-independent of $\text{Sbox}(k_2 \oplus w_2)$.

1315 Alg. 1 does directly operate on the following pairs of variables: $(r_2, \text{Sbox}(k_2 \oplus w_2) \oplus r_2)$, $(r_{\text{out}}, \text{Sbox}(k_2 \oplus$
 1316 $w_2) \oplus r_{\text{out}})$, $(r_{\text{in}}, k_2 \oplus w_2 \oplus r_{\text{in}})$. The variables r_2 , r_{out} , r_{in} are called *masks* and are internal variables which
 1317 are randomly generated during each encryption. Thus, each of these variables *is* marginally dependent on
 1318 some measurements X_t and can be detected with first-order parametric methods. Each of these pairs is said
 1319 to leak $\text{Sbox}(k_2 \oplus w_2)$ because given both, one can calculate $\text{Sbox}(k_2 \oplus w_2)$ using the identity $a \oplus b \oplus b = a$.
 1320 In addition to these pairs of variables, Egger et al. (2022) identified that the variables $\text{Sbox}(S_{\text{prev}} \oplus r_{\text{in}}) \oplus r_{\text{out}}$
 1321 and a ‘security load’ $S_{\text{prev}} \oplus \text{Sbox}(w_2 \oplus k_2) \oplus r_{\text{out}}$ also have a strong first-order association with power
 1322 consumption and might contribute to leakage, where for byte 2 $S_{\text{prev}} = \text{Sbox}(k_{11} \oplus w_{11}) \oplus r_{11}$.

1323 All 8 of these internal variables may be computed using the metadata published with the ASCADv1 datasets.
 1324 As our oracle assessment for the ASCADv1 datasets, we compute the SNR of each of these 8 variables using
 1325 the attack dataset and average them together. We can then qualitatively assess agreement between the output
 1326 of leakage localization algorithms and these oracle assessments. For the 4 first-order datasets, we use as our
 1327 oracle assessment the SNR of the target itself from the attack dataset, which amounts to an assumption that
 1328 there is no leakage of order 2 or higher.

1329 We use the Spearman rank correlation coefficient as a scalar summary of this agreement. This quantity
 1330 is defined as the Pearson correlation coefficient between the *ranks* of a pair of sequences, and is useful for
 1331 our purposes because it tells us the extent to which leakage localization algorithms assign the same *relative*
 1332 leakiness to measurements as the oracle, while being insensitive to differences in their ‘shape’.

1333 Most prior work (Masure et al., 2019; Wouters et al., 2020; Schamberger et al., 2023; Yap et al., 2025) has
 1334 used white box assessments similar to this for qualitative evaluation of leakage localization algorithms. To our
 1335 knowledge, ours is the first work to summarize agreement with a scalar and use it for large-scale comparison
 1336 between a large number of methods.

1337 We refer to this performance evaluation strategy as *oracle agreement*. Note that we use the word ‘oracle’ for
 1338 clarity of exposition, and we believe this is the least-flawed of the evaluation metrics we consider, but it does
 1339 not give us genuine ground truth leakiness measurements. It is sensitive only to first-order leakage of variables
 1340 which can be identified *a priori* as leaky, and will ignore any other exploitable variables. Additionally, SNR
 1341 is not perfectly sensitive even to first-order leakage: it relies on changes in the expected values $\mathbb{E}[X_t | Y = y]$
 1342 with y and will not detect dependencies which do not influence the mean (e.g. if X_t is a Gaussian random
 1343 variable with Y -independent mean but Y -dependent variance).

1344 **DNN occlusion tests** Hettwer et al. (2020) proposed a variety of tests based on plotting the performance
 1345 of a trained supervised classifier as its input features are successively occluded in order of their leakiness.
 1346 The intuition is that leakier features should have a larger impact on the performance of the classifier, so the
 1347 rate at which its performance changes as we successively occlude its inputs tells about the extent to which

Algorithm 4: Pseudocode for the DNN occlusion tests.

Input: Trained supervised classifier $\Phi^* : \mathbb{R}^T \times \mathcal{Y} \rightarrow [0, 1]$, attack dataset $D_{\text{attack}} \subseteq \mathbb{R}^T \times \mathcal{Y}$, leakiness estimates $\ell \in \mathbb{R}^T$, direction $d \in \{\text{'forward'}, \text{'reverse'}\}$

Output: Area under DNN occlusion curve \bar{r}

```

1  $m_0 \leftarrow \mathbf{0}$                                      // occlusion mask
2  $\ell_{\text{idx}} \leftarrow \text{argsort}(\ell)$            // indices of sorted leakiness values, from low-high
3 if  $d = \text{'forward'}$  then
4    $\ell_{\text{idx}} \leftarrow \text{reverseOrder}(\ell_{\text{idx}})$       // sort from high-low instead
5 else if  $d = \text{'reverse'}$  then
6   pass
7 for  $t = 1, \dots, T$  do
8    $m_t \leftarrow m_{t-1} + I_{\ell_{\text{idx}, t}}$  // un-occlude  $t$ -th least (reverse) or most (forward)-leaky feature
9    $r_t \leftarrow \frac{1}{|D_{\text{attack}}|} \sum_{x, y \in D_{\text{attack}}} \text{Rank}(\Phi^*(\cdot | (1 - m_t) \odot x); y)$  // record average classifier
   performance on attack dataset under this occlusion mask
10 return  $\frac{1}{T} \sum_{t=1}^T r_t$                                 // area under the DNN occlusion curve

```

1348 these inputs were leaky. In a similar spirit, we propose 2 evaluation metrics which we name the *forward* and
 1349 *reverse DNN occlusion tests*.

1350 See Alg. 4. For the *forward DNN occlusion test* we initially occlude all the input features of a trained
 1351 classifier, then successively un-occlude one feature at a time from most- to least-leaky as predicted by the
 1352 leakage localization algorithm under test. At each occlusion level we measure the performance of the classifier
 1353 on the attack dataset in terms of mean rank. We then report the average performance across all occlusion
 1354 levels. For a ‘good’ leakage assessment we expect the average performance to be better (lower) because useful
 1355 features are un-occluded at a greater proportion of occlusion levels. Conversely, for a ‘bad’ leakage assessment
 1356 we expect the average performance to be worse because useful features stay occluded for longer. The *reverse*
 1357 *DNN occlusion test* is identical except that we un-occlude features from least- to most-leaky. For this test we
 1358 expect the average performance to be worse (higher) for a ‘good’ leakage assessment and better (lower) for
 1359 a ‘good’ leakage assessment. In general we expect the forward test to be sensitive to the extent to which
 1360 the predicted-leakiest measurements are truly among the leakiest (similar to true/false positives), and the
 1361 reverse test to be sensitive to the extent to which the predicted-nonleaky features are truly nonleaky (similar
 1362 to true/false negatives).

1363 A major limitation of the DNN occlusion tests is that they rely on an imperfect DNN classifier, and are only
 1364 sensitive to associations insofar as the classifier exploits them. Additionally, we use the same architecture,
 1365 training procedure and hyperparameters for these classifiers as for those ‘interpreted’ by the neural net
 1366 attribution baseline methods, so the test may be ‘biased’ in favor of these. Nonetheless, we consider them
 1367 a useful supplement to the oracle agreement metric because they do not suffer from the same restrictive
 1368 assumptions about the nature of associations.

1369 **Feature selection efficacy for Gaussian template attack** Similarly to Masure et al. (2019); Yap
 1370 et al. (2023), we also evaluate leakage localization assessments based on their ability to do feature selection
 1371 for Gaussian template attacks. To carry out this test, we first select the top 20 measurements with the
 1372 highest estimated leakiness. We then perform a Gaussian template attack (Chari et al., 2003) using these
 1373 measurements. The Gaussian template attack is a well-known parametric side-channel attack based on
 1374 modeling $p_{\mathbf{X}|Y}$ with a Gaussian mixture model with one component per value Y may take on, then using
 1375 Bayes’ rule to estimate $p_{Y|\mathbf{X}}$. The leakier these features are, the more-performant we expect the attack to be.

1376 For AES datasets accuracy is often low when predicting Y from a single value of \mathbf{X} . Thus, attack datasets
 1377 typically consist of many traces $\mathbf{X}_1, \dots, \mathbf{X}_M$ recorded with a *fixed* AES key but varying plaintext. The
 1378 target variable Y is typically chosen so that given the corresponding plaintext and ciphertext, Y is a known
 1379 invertible function of the key. One can thereby make many predictions about the key using these traces, then
 1380 ‘accumulate’ these predictions through the identity $\log p_{\mathbf{X}_1, \dots, \mathbf{X}_M|K} = \sum_{m=1}^M \log p_{\mathbf{X}_m|K}$ where K denotes the

1381 key. A common performance metric for attacks is the *minimum traces to disclosure (MTD)*, given by the
 1382 number of traces one must accumulate before the true key has the highest predicted probability mass (lower
 1383 is better). We use this metric to measure the performance of Gaussian template attacks on AES datasets.
 1384 For OTiAiT and OTP, which are not AES datasets, we simply use the mean rank of the target variable on
 1385 the attack dataset.

1386 Note that for the second-order ASCADv1 datasets, the algorithms we consider do not reveal which of the
 1387 leaky internal AES variables a measurement leaks. This is problematic when selecting features for a template
 1388 attack, because a successful attack must have features corresponding to both of a pair of second-order leaky
 1389 variables. If we simply used the top 20 predicted-leakiest features for an attack, this would be left to chance
 1390 and make the performance metric unreliable. To address this issue, for the second-order datasets we instead
 1391 segment the T measurements into 10 bins each containing $\lfloor \frac{T}{10} \rfloor$ consecutive measurements, then select the
 1392 top 2 leakiest measurements from each bin.

1393 The main shortcoming of this performance metric is that it is only sensitive to the 20 predicted-leakiest
 1394 measurements and ignores all others (e.g. it cannot detect that leaky measurements have been assigned
 1395 spuriously-low leakiness). Additionally, it assumes that the relationship between measurements and the target
 1396 variable is well-described by a Gaussian mixture model, which may not hold in practice. However, unlike the
 1397 oracle agreement metric it does not rely on human-identified first-order leaky variables, and unlike the DNN
 1398 occlusion tests it is hyperparameter-free and may be biased towards different associations than the DNN
 1399 classifiers. Thus, it is also a useful supplement to the aforementioned metrics.

1400 E.3.5 Model selection criterion

1401 While we consider the oracle agreement metric our most straightforward and useful performance metric, we
 1402 cannot use it for model selection (e.g. choosing hyperparameters, or early-stopping runs). This is because
 1403 it relies on ‘white box’ knowledge of the internal first-order leaky variables of second-order algorithms and
 1404 knowledge of their random masks, which we assume not to have at training time. Thus, we must devise a
 1405 model selection criterion which does not rely on this information.

1406 Note that we can freely use the forward and reverse DNN occlusion tests and the template attack feature
 1407 selection test for model selection by running them on our validation dataset rather than the attack dataset.
 1408 Additionally, we find that because ‘good’ runs typically converge to similar leakiness assessments whereas
 1409 ‘bad’ runs resemble random noise, a reasonable model selection heuristic is to 1) compute the average leakiness
 1410 value for each measurement over all hyperparameter tuning runs, then 2) use the Spearman rank correlation
 1411 coefficient between the average leakiness assessments and those for a particular run as a proxy for the run’s
 1412 performance. We refer to this model selection strategy as the *mean agreement criterion*.

1413 In Fig. 13 we visualize the relationship between the oracle agreement and the forward and reverse DNN
 1414 occlusion tests as well as the mean agreement criterion. We find that the forward and reverse DNN occlusion
 1415 criterion are weakly correlated with oracle agreement, and that the mean agreement is strongly correlated
 1416 for every dataset apart from OtiAiT. However, while we can consistently discard bad runs using these
 1417 criteria, they typically lead to selection of suboptimal models. In this work we select **ALL** models using a
 1418 ‘composite criterion’ based on ranking runs according to the forward and reverse DNN occlusion tests and
 1419 mean agreement criterion, then selecting the model with the highest mean ranking across these 3 criteria.
 1420 More research into model selection strategies for leakage localization algorithms is warranted.

1421 E.3.6 Summary of experiments

1422 We report and visualize our results in a variety of ways, which we list and summarize here.

1423 **Visualization of the best-performing leakage localization results found by **ALL**** In Fig. 14 we
 1424 visualize the oracle leakiness measurements and qualitatively compare them with the best **ALL** runs found
 1425 during our hyperparameter sweeps. For the ASCADv1 datasets we draw distinct curves for the individual
 1426 leaky variables – note the similarity between these plots and (Egger et al., 2022, Fig. 3a). We see a strong
 1427 visual resemblance between the **ALL** outputs and the oracle leakiness assessments, despite the nonlinear
 1428 relationship between them. Additionally, there are few points assigned a low oracle leakiness but high *relative*

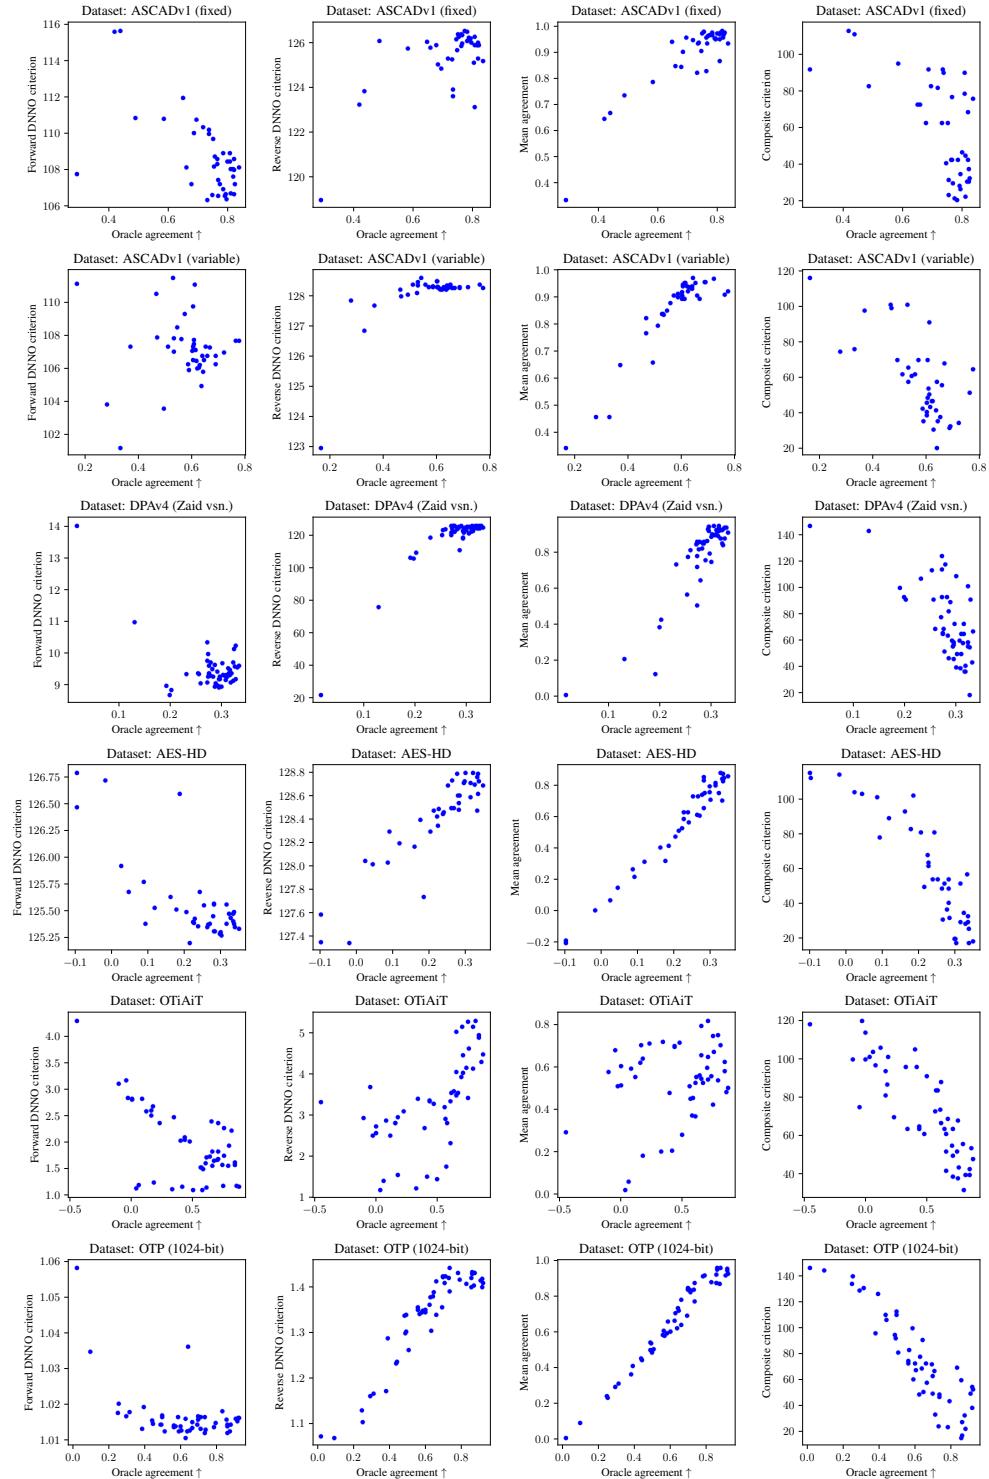


Figure 13: Visualization of the relationship between various model selection criteria from Sec. E.3.5 and the oracle agreement for **ALL** runs produced during hyperparameter search. We find consistently across datasets that the forward/reverse DNN occlusion and mean agreement criteria are consistently positively or negatively correlated with the oracle agreement, though this correlation is often weak. We achieve slightly better results using a composite criterion which uses considers the ‘votes’ according to all these criterion, and we adopt this composite criterion when selecting **ALL** models for comparison with baselines. However, this criterion often selects suboptimal models, and future research on leakage localization model selection strategies is warranted.

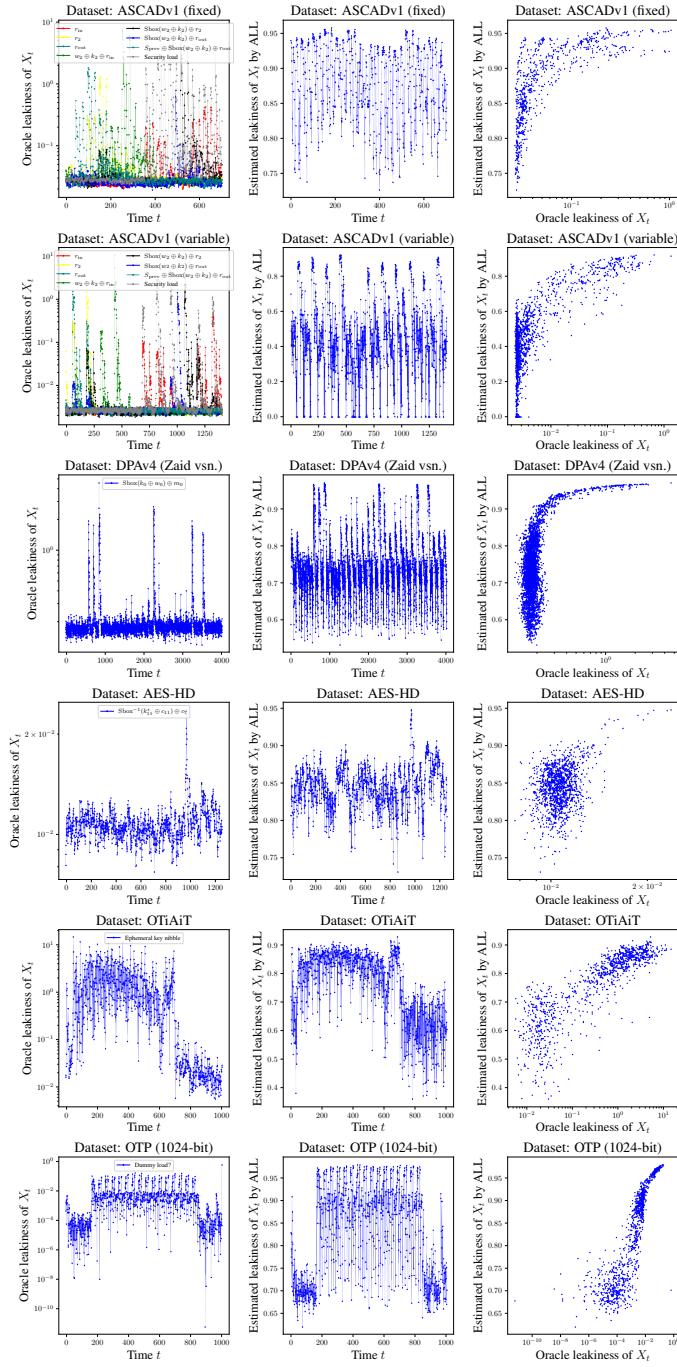


Figure 14: A plot which qualitatively compares the estimated leakiness of the best-performing ALL runs with the oracle leakiness values. **(left column)** A plot of oracle leakiness of X_t vs. timestep t . For the second-order ASCADv1 datasets (top two rows), note that our plots are similar to (Egger et al., 2022, Fig. 3a) as they are based on the same first-order variables; differences are because we measure leakiness with SNR whereas they use CPOI. **(middle column)** A plot of estimated leakiness of X_t vs. timestep t according to ALL, for the best-performing ALL run as measured by oracle agreement. We want this column to look like the left column, up to a strictly-increasing nonlinear transform. **(right column)** A plot of the ALL-estimated leakiness of X_t vs. the oracle leakiness of X_t . These curves show that the ALL estimates are good in the sense that they tend to apply similar *relative* leakiness values to the measurements as the oracle.

1429 ALL-predicted leakiness or vice-versa. Most of the disagreement appears to lie in the predicted relative
 1430 leakiness within groups of high-leakiness or low-leakiness measurements.

1431 **ALL training curves** In Fig. 15 we plot the evolution of various metrics over time during the ALL training
 1432 procedure for the runs chosen using the model selection criterion of Sec. E.3.5. In Fig. 16 we compare the
 1433 oracle agreement at different timesteps on ASCADv1-variable for ALL and selected baselines. In general
 1434 we observe that for successful ALL runs the classifier validation rank drops significantly below the random-
 1435 guessing threshold at some point in training, though it may begin to rise again as the noise distribution
 1436 trains adversarially against it. These curves are generally smooth and we do not observe significant training
 1437 instability. Unfortunately, we are not aware of a reliable way to predict oracle agreement performance from
 1438 the training curves.

1439 **Sensitivity of ALL to hyperparameters** The main hyperparameters of ALL are the noise budget $\bar{\gamma}$ and
 1440 the learning rates of the classifier weights θ and the noise distribution parameter $\tilde{\eta}$. In Fig. 17 we evaluate
 1441 the sensitivity of ALL to these hyperparameters by varying them individually using the optimal configuration
 1442 with respect to oracle agreement as a starting point. We find in general that performance generally varies
 1443 smoothly with these hyperparameters and stays significantly above random-guessing over a large search space,
 1444 which is desirable from a standpoint of hyperparameter tuning.

1445 **Attack performance and training curves of supervised classifiers** All of the deep learning-based
 1446 baseline methods are based on ‘interpreting’ a fixed classifier which has been trained using supervised learning
 1447 to predict the target variable Y from the trace \mathbf{X} . In Fig. 18 we plot the cross-entropy loss and mean rank
 1448 over time for these classifiers during training. Additionally, for the AES datasets we plot the rank of the
 1449 correct key as we accumulate traces from the attack dataset. For reference we superimpose the correct key
 1450 rank during trace accumulation for the following publicly-available open-weight classifiers: the CNN_{Best} and
 1451 MLP_{Best} model of Benadjila et al. (2020), both of which are available here, and the models of Zaid et al.
 1452 (2020) and their simplified versions from Wouters et al. (2020) distributed here. Also note that in Table 9 we
 1453 list the early-stopped validation cross-entropy loss, rank, and minimum traces to disclosure (MTD) for these
 1454 models. We find that our classifiers are able to ‘successfully’ attack all datasets (i.e. they can successfully
 1455 predict the key by accumulating all traces in the provided attack dataset), and they achieve comparable
 1456 MTD to these open-weight models on ASCADv1-fixed, ASCADv1-variable, and DPAv4, and somewhat-worse
 1457 MTD on the AES-HD dataset.

1458 **m -Occlusion window size sweep and smoothing effect** We consider as baselines m -Occlusion and
 1459 2nd-order m -Occlusion. These baselines introduce the occlusion window size as an additional hyperparameter
 1460 which must be tuned. In Fig. 19 we plot the oracle agreement performance of m -Occlusion as we sweep m , and qualitatively
 1461 show how the resulting leakiness vector is smoothed out with increasing m . In subsequent
 1462 experiments, we denote by m^* the optimal value of m found in these experiments, and report results for
 1463 1-Occlusion, m^* -Occlusion, 2nd-order 1-Occlusion and 2nd-order m^* -Occlusion (due to the high cost of
 1464 2nd-order m -Occlusion we do not separately sweep its window size). Note that this introduces some data
 1465 leakage into the results, as we are doing validation with our test metric which incorporates implementation
 1466 knowledge that we assume not to have at training time. Because the goal of this work is to demonstrate the
 1467 efficacy of ALL, and ALL generally outperforms m^* -Occlusion and 2nd-order m^* -Occlusion despite the data
 1468 leakage, we consider this acceptable.

1469 Compared to 1-Occlusion, m -Occlusion has two major differences: it occludes multiple inputs simultaneously
 1470 to increase the influence on classifier predictions, and it has a ‘smoothing’ effect which causes nearby
 1471 measurements to be assigned similar leakiness values. The latter effect can be easily simulated using average-
 1472 pooling, so we also plot the performance of ALL as we average-pool it with stride 1 and kernel size m . We
 1473 find that while m -Occlusion significantly improves performance over 1-Occlusion on the DPAv4 and AES-HD
 1474 datasets, on these same datasets we can significantly improve the performance of ALL by average-pooling,
 1475 and ALL convincingly outperforms m -Occlusion when accounting for this. We conjecture that the smoothing
 1476 effect provides a useful inductive bias for these datasets, and emphasize that it can easily be applied to any
 1477 other leakage localization technique.

Table 11: Performance comparison between various leakage localization algorithms according to the oracle agreement metric (**larger is better**) described in Sec. E.3.4. Results are reported as mean \pm standard deviation over 5 random seeds. The best result is **boxed** and the best deep learning result is underlined. We consider a result to be ‘best’ if its mean lies inside of the error bars of the result with the highest mean.

| Method | 2nd-order datasets | | 1st-order datasets | | | |
|-------------------------------|-------------------------------------|-----------------------------------|-------------------------------------|-----------------------------------|-------------------------------------|-------------------------------------|
| | ASCADv1 (fixed) | ASCADv1 (random) | DPAv4 (Zaid vsn.) | AES-HD | OTiAiT | OTP (1024-bit) |
| Random | -0.00 ± 0.04 | -0.02 ± 0.01 | 0.01 ± 0.01 | -0.01 ± 0.02 | 0.02 ± 0.03 | 0.03 ± 0.03 |
| SNR | 0.031 | <u>-0.092</u> | 0.344 | 0.185 | 0.989 | 0.944 |
| SOSD | <u>-0.253</u> | 0.272 | 0.259 | 0.063 | 0.886 | 0.803 |
| CPA | 0.521 | <u>-0.095</u> | 0.420 | 0.303 | 0.630 | 0.945 |
| GradVis | 0.48 \pm 0.02 | 0.27 \pm 0.01 | 0.198 \pm 0.009 | 0.07 \pm 0.01 | 0.55 \pm 0.05 | 0.57 \pm 0.02 |
| Saliency | 0.47 \pm 0.02 | 0.26 \pm 0.01 | 0.198 \pm 0.008 | 0.07 \pm 0.01 | 0.67 \pm 0.06 | 0.58 \pm 0.02 |
| Input * Grad | 0.47 \pm 0.02 | 0.25 \pm 0.01 | 0.202 \pm 0.009 | 0.08 \pm 0.02 | 0.71 \pm 0.05 | 0.60 \pm 0.02 |
| LRP | 0.47 \pm 0.02 | 0.25 \pm 0.01 | 0.202 \pm 0.009 | 0.08 \pm 0.02 | 0.71 \pm 0.05 | 0.60 \pm 0.02 |
| OccPOI | 0.07 \pm 0.01 | 0.064 \pm 0.004 | 0.030 \pm 0.008 | 0.044 \pm 0.009 | 0.07 \pm 0.02 | 0.01 \pm 0.02 |
| 1-Occlusion | 0.47 \pm 0.02 | 0.25 \pm 0.01 | 0.202 \pm 0.009 | 0.08 \pm 0.01 | 0.71 \pm 0.05 | 0.60 \pm 0.02 |
| m^* -Occlusion | 0.49 \pm 0.02 | 0.41 \pm 0.01 | <u>0.32 \pm 0.01</u> | 0.18 \pm 0.05 | 0.72 \pm 0.04 | 0.77 \pm 0.01 |
| 1-Occlusion ² | 0.51 \pm 0.01 | 0.27 \pm 0.01 | 0.206 \pm 0.009 | 0.08 \pm 0.01 | 0.74 \pm 0.05 | 0.60 \pm 0.02 |
| m^* -Occlusion ² | 0.52 \pm 0.01 | 0.42 \pm 0.01 | <u>0.330 \pm 0.009</u> | 0.19 \pm 0.05 | 0.75 \pm 0.04 | 0.788 \pm 0.007 |
| WoutersNet 1-Occlusion | 0.18 \pm 0.03 | n/a | 0.21 \pm 0.02 | 0.11 \pm 0.03 | n/a | n/a |
| WoutersNet m^* -Occlusion | 0.20 \pm 0.03 | n/a | 0.29 \pm 0.02 | <u>0.21 \pm 0.04</u> | n/a | n/a |
| WoutersNet GradVis | 0.19 \pm 0.03 | n/a | 0.21 \pm 0.02 | 0.11 \pm 0.03 | n/a | n/a |
| WoutersNet Input * Grad | 0.18 \pm 0.03 | n/a | 0.21 \pm 0.02 | 0.11 \pm 0.03 | n/a | n/a |
| WoutersNet Saliency | 0.19 \pm 0.03 | n/a | 0.21 \pm 0.02 | 0.11 \pm 0.03 | n/a | n/a |
| ZaidNet 1-Occlusion | 0.24 \pm 0.03 | n/a | 0.19 \pm 0.01 | 0.13 \pm 0.02 | n/a | n/a |
| ZaidNet m^* -Occlusion | 0.25 \pm 0.04 | n/a | <u>0.273 \pm 0.009</u> | <u>0.21 \pm 0.05</u> | n/a | n/a |
| ZaidNet GradVis | 0.25 \pm 0.03 | n/a | 0.19 \pm 0.01 | 0.13 \pm 0.02 | n/a | n/a |
| ZaidNet Input * Grad | 0.25 \pm 0.04 | n/a | 0.19 \pm 0.01 | 0.13 \pm 0.02 | n/a | n/a |
| ZaidNet Saliency | 0.25 \pm 0.03 | n/a | 0.19 \pm 0.01 | 0.13 \pm 0.02 | n/a | n/a |
| ALL (ours) | 0.794 \pm 0.006 | <u>0.60 \pm 0.01</u> | 0.317 \pm 0.002 | <u>0.22 \pm 0.03</u> | 0.782 \pm 0.001 | <u>0.848 \pm 0.003</u> |

1478 **Quantitative comparison between ALL and baseline methods using oracle agreement, forward**
 1479 **and reverse DNN occlusion tests, and template attack feature selection test** In Tables 11, 12,
 1480 13, 14 we compare the performance of ALL with our baseline methods according to the oracle agreement,
 1481 forward DNN occlusion, reverse DNN occlusion, and template attack feature selection tests, respectively. As
 1482 a sanity check against our supervised neural net architecture and training procedure, we also include results
 1483 for GradVis, Saliency, LRP, Input * Grad, 1-Occlusion, and m^* -Occlusion computed using the models of
 1484 Zaid et al. (2020) and their simplified versions from Wouters et al. (2020) distributed here.

1485 We find that ALL outperforms all prior deep learning-based leakage localization algorithms on all datasets
 1486 except for DPAv4 according to the oracle agreement metric. The first-order parametric methods outperform
 1487 the deep learning methods on the first-order datasets but generally do poorly on the second-order ASCADv1
 1488 datasets due to not being sensitive to second-order associations. Unsurprisingly, SNR and SOSD achieve
 1489 near-random performance on ASCADv1-fixed and SNR and CPA achieve near-random performance on
 1490 ASCADv1-variable. Surprisingly, CPA performs fairly well on ASCADv1-fixed and SOSD performs fairly well
 1491 on ASCADv1-variable. We are not sure why this is the case, but conjecture it is because these methods have
 1492 some natural proclivity to ‘rule out’ measurements which are not at useful points in time for these particular
 1493 datasets (e.g. those which are not close to a clock edge). Note that this surprisingly-strong performance
 1494 compared to the deep learning baselines does not appear to carry over to the evaluations with the DNN
 1495 occlusion tests or the template attack feature selection test.

1496 On average, ALL is the best method on the majority of datasets according to the reverse DNN occlusion test
 1497 and template attack test, but results are mixed according to the forward DNN occlusion test. However, we
 1498 find that the DNN occlusion tests have high variance, and in general there is a large overlap in error bars
 1499 (note the large number of boxed and underlined methods in Tables 12 and 13).

Table 12: Performance comparison between various leakage localization algorithms according to the forward DNN occlusion test (**smaller is better**) described in Sec. E.3.4. Results are reported as mean \pm standard deviation over 5 random seeds. The best result is **boxed** and the best deep learning result is underlined. We consider a result to be ‘best’ if its mean lies inside of the error bars of the result with the highest mean.

| Method | 2nd-order datasets | | 1st-order datasets | | | |
|-------------------------------|-----------------------------------|----------------------------------|---------------------------------|-----------------------------------|---------------------------------|-------------------------------------|
| | ASCADv1 (fixed) | ASCADv1 (random) | DPAv4 (Zaid vsn.) | AES-HD | OTiAiT | OTP (1024-bit) |
| Random | 111 \pm 1 | 111.7 \pm 0.8 | 20 \pm 2 | 126.8 \pm 0.1 | 1.30 \pm 0.05 | 1.065 \pm 0.010 |
| SNR | 117.525 | 118.448 | 11.275 | 125.053 | 1.209 | 1.015 |
| SOSD | 115.516 | 106.072 | 11.455 | 125.365 | 1.213 | 1.032 |
| CPA | 111.449 | 117.349 | 11.811 | 125.267 | 2.049 | 1.015 |
| GradVis | 108.6 \pm 0.5 | 96.8 \pm 0.3 | 9.5 \pm 0.6 | 125.6 \pm 0.3 | 1.9 \pm 0.2 | <u>1.013 \pm 0.002</u> |
| Saliency | 108.5 \pm 0.4 | 96.3 \pm 0.4 | 9.5 \pm 0.7 | 125.6 \pm 0.3 | <u>1.8 \pm 0.1</u> | <u>1.014 \pm 0.001</u> |
| Input * Grad | 108.5 \pm 0.4 | 96.8 \pm 0.4 | 9.4 \pm 0.7 | 125.6 \pm 0.3 | <u>1.7 \pm 0.2</u> | <u>1.013 \pm 0.001</u> |
| LRP | 108.5 \pm 0.4 | 96.8 \pm 0.4 | 9.4 \pm 0.7 | 125.6 \pm 0.3 | <u>1.7 \pm 0.2</u> | <u>1.013 \pm 0.001</u> |
| OccPOI | 122.3 \pm 0.8 | 120.8 \pm 0.2 | 58 \pm 2 | 127.4 \pm 0.3 | 2.6 \pm 0.2 | 1.09 \pm 0.04 |
| 1-Occlusion | 108.5 \pm 0.4 | 96.7 \pm 0.4 | 9.4 \pm 0.7 | 125.6 \pm 0.3 | <u>1.7 \pm 0.2</u> | <u>1.013 \pm 0.001</u> |
| m^* -Occlusion | 108.2 \pm 0.5 | 95.7 \pm 0.6 | 9.0 \pm 0.6 | <u>125.3 \pm 0.2</u> | <u>1.8 \pm 0.2</u> | <u>1.013 \pm 0.001</u> |
| 1-Occlusion ² | 108.4 \pm 0.4 | 97.0 \pm 0.4 | 9.4 \pm 0.7 | 125.5 \pm 0.3 | <u>1.7 \pm 0.2</u> | <u>1.013 \pm 0.001</u> |
| m^* -Occlusion ² | 108.2 \pm 0.5 | 95.9 \pm 0.6 | 9.0 \pm 0.5 | <u>125.3 \pm 0.2</u> | <u>1.7 \pm 0.2</u> | <u>1.013 \pm 0.001</u> |
| WoutersNet 1-Occlusion | 109.1 \pm 0.7 | n/a | 8.7 \pm 0.8 | <u>125.4 \pm 0.3</u> | n/a | n/a |
| WoutersNet m^* -Occlusion | 109.0 \pm 0.8 | n/a | 9.0 \pm 0.7 | <u>125.4 \pm 0.2</u> | n/a | n/a |
| WoutersNet GradVis | 109.3 \pm 0.8 | n/a | 8.8 \pm 0.8 | <u>125.4 \pm 0.3</u> | n/a | n/a |
| WoutersNet Input * Grad | 109.1 \pm 0.7 | n/a | 8.7 \pm 0.8 | <u>125.4 \pm 0.3</u> | n/a | n/a |
| WoutersNet Saliency | 109.2 \pm 0.8 | n/a | 8.7 \pm 0.8 | <u>125.4 \pm 0.3</u> | n/a | n/a |
| ZaidNet 1-Occlusion | 110.0 \pm 0.4 | n/a | 8.7 \pm 0.7 | 125.5 \pm 0.2 | n/a | n/a |
| ZaidNet m^* -Occlusion | 109.5 \pm 0.9 | n/a | 9.0 \pm 0.7 | <u>125.4 \pm 0.2</u> | n/a | n/a |
| ZaidNet GradVis | 109.7 \pm 0.6 | n/a | 8.7 \pm 0.7 | 125.5 \pm 0.2 | n/a | n/a |
| ZaidNet Input * Grad | 109.9 \pm 0.3 | n/a | 8.7 \pm 0.7 | 125.5 \pm 0.2 | n/a | n/a |
| ZaidNet Saliency | 109.7 \pm 0.6 | n/a | 8.7 \pm 0.7 | 125.5 \pm 0.2 | n/a | n/a |
| ALL (ours) | 107.3 \pm 0.4 | 104 \pm 1 | 9.9 \pm 0.7 | 125.5 \pm 0.3 | <u>1.8 \pm 0.1</u> | 1.014 \pm 0.001 |

Table 13: Performance comparison between various leakage localization algorithms according to the reverse DNN occlusion test(**larger is better**) described in Sec. E.3.4. Results are reported as mean \pm standard deviation over 5 random seeds. The best result is **boxed** and the best deep learning result is underlined. We consider a result to be ‘best’ if its mean lies inside of the error bars of the result with the highest mean.

| Method | 2nd-order datasets | | 1st-order datasets | | | |
|-------------------------------|-----------------------------------|-------------------------------------|-----------------------------------|-----------------------------------|---------------------------------|-----------------------------------|
| | ASCADv1 (fixed) | ASCADv1 (random) | DPAv4 (Zaid vsn.) | AES-HD | OTiAiT | OTP (1024-bit) |
| Random | 111.9 \pm 0.5 | 114 \pm 3 | 25 \pm 4 | 126.7 \pm 0.2 | 1.276 \pm 0.010 | 1.079 \pm 0.010 |
| SNR | 110.679 | 123.284 | 125.924 | 128.599 | 5.267 | 1.373 |
| SOSD | 112.074 | 126.765 | 108.267 | 128.216 | 4.957 | 1.369 |
| CPA | 119.875 | 117.885 | 114.857 | 128.367 | 3.792 | 1.364 |
| GradVis | 125.9 \pm 0.2 | 127.6 \pm 0.1 | 122 \pm 1 | 128.0 \pm 0.3 | 4.2 \pm 0.4 | 1.34 \pm 0.07 |
| Saliency | 125.8 \pm 0.2 | 127.4 \pm 0.2 | 122 \pm 1 | 128.0 \pm 0.3 | 5.1 \pm 0.3 | 1.33 \pm 0.06 |
| Input * Grad | 125.7 \pm 0.3 | 127.5 \pm 0.2 | 121.9 \pm 0.9 | 128.1 \pm 0.3 | 5.2 \pm 0.3 | 1.34 \pm 0.06 |
| LRP | 125.7 \pm 0.3 | 127.5 \pm 0.2 | 121.9 \pm 0.9 | 128.1 \pm 0.3 | 5.2 \pm 0.3 | 1.34 \pm 0.06 |
| OccPOI | 122.3 \pm 0.4 | 124.6 \pm 0.2 | 43 \pm 1 | 127.0 \pm 0.3 | 3.6 \pm 0.3 | 1.09 \pm 0.04 |
| 1-Occlusion | 125.8 \pm 0.3 | 127.4 \pm 0.2 | 122 \pm 1 | 128.1 \pm 0.3 | 5.2 \pm 0.3 | 1.34 \pm 0.06 |
| m^* -Occlusion | 126.0 \pm 0.2 | 127.4 \pm 0.2 | 121 \pm 1 | <u>128.5 \pm 0.2</u> | 5.3 \pm 0.2 | 1.30 \pm 0.04 |
| 1-Occlusion ² | 125.8 \pm 0.3 | 127.5 \pm 0.2 | 122.0 \pm 0.9 | 128.1 \pm 0.3 | 5.3 \pm 0.2 | 1.34 \pm 0.06 |
| m^* -Occlusion ² | 126.1 \pm 0.2 | 127.4 \pm 0.2 | 121.3 \pm 0.9 | <u>128.5 \pm 0.2</u> | 5.3 \pm 0.2 | 1.30 \pm 0.04 |
| WoutersNet 1-Occlusion | 121.9 \pm 0.7 | n/a | 116 \pm 5 | 128.2 \pm 0.1 | n/a | n/a |
| WoutersNet m^* -Occlusion | 122.4 \pm 0.8 | n/a | <u>119.3 \pm 0.7</u> | <u>128.4 \pm 0.2</u> | n/a | n/a |
| WoutersNet GradVis | 121.8 \pm 0.7 | n/a | 116 \pm 5 | 128.2 \pm 0.1 | n/a | n/a |
| WoutersNet Input * Grad | 121.9 \pm 0.7 | n/a | 116 \pm 5 | 128.2 \pm 0.1 | n/a | n/a |
| WoutersNet Saliency | 121.8 \pm 0.7 | n/a | 116 \pm 4 | 128.2 \pm 0.1 | n/a | n/a |
| ZaidNet 1-Occlusion | 122 \pm 3 | n/a | 116 \pm 3 | 128.1 \pm 0.2 | n/a | n/a |
| ZaidNet m^* -Occlusion | 123 \pm 3 | n/a | 119 \pm 2 | <u>128.4 \pm 0.3</u> | n/a | n/a |
| ZaidNet GradVis | 122 \pm 3 | n/a | 117 \pm 3 | 128.1 \pm 0.2 | n/a | n/a |
| ZaidNet Input * Grad | 122 \pm 3 | n/a | 116 \pm 3 | 128.1 \pm 0.2 | n/a | n/a |
| ZaidNet Saliency | 122 \pm 3 | n/a | 117 \pm 3 | 128.1 \pm 0.2 | n/a | n/a |
| ALL (ours) | <u>126.4 \pm 0.2</u> | <u>127.96 \pm 0.06</u> | <u>125 \pm 1</u> | 128.3 \pm 0.2 | <u>5.6 \pm 0.2</u> | <u>1.39 \pm 0.05</u> |

Table 14: Performance comparison between various leakage localization algorithms according to the template attack feature selection test (**smaller is better**) described in Sec. E.3.4. Results are reported as mean \pm standard deviation over 5 random seeds. The best result is **boxed** and the best deep learning result is underlined. We consider a result to be ‘best’ if its mean lies inside of the error bars of the result with the highest mean.

| Method | 2nd-order datasets | | 1st-order datasets | | | |
|-------------------------------|----------------------------------|----------------------------------|-------------------------------------|-------------------------------------|-----------------------------------|---------------------------------------|
| | ASCADv1 (fixed) | ASCADv1 (random) | DPAv4 (Zaid vsn.) | AES-HD | OTiAiT | OTP (1024-bit) |
| Random | 1293 \pm 1000 | 56421 \pm 40000 | 449 \pm 100 | 25000.000 | 1.2 \pm 0.1 | 1.44 \pm 0.02 |
| SNR | 5496.010 | 54834.990 | 2.590 | 17159.690 | [1.058] | 1.385 |
| SOSD | 7733.870 | 3237.250 | 91.410 | 17520.530 | <u>1.059</u> | 1.398 |
| CPA | 3826.440 | 100000.000 | 10.540 | 18974.510 | 1.101 | 1.385 |
| GradVis | 686 \pm 100 | 1162 \pm 1000 | 2.7 \pm 0.1 | 20014 \pm 6000 | 1.4 \pm 0.3 | 1.378 \pm 0.007 |
| Saliency | 726 \pm 100 | 1412 \pm 2000 | 2.7 \pm 0.1 | 19438 \pm 6000 | 1.14 \pm 0.02 | 1.379 \pm 0.005 |
| Input * Grad | 675 \pm 100 | 1194 \pm 2000 | 2.6 \pm 0.1 | 19893 \pm 6000 | 1.14 \pm 0.02 | 1.378 \pm 0.003 |
| LRP | 675 \pm 100 | 1194 \pm 2000 | 2.6 \pm 0.1 | 19893 \pm 6000 | 1.14 \pm 0.02 | 1.378 \pm 0.003 |
| OccPOI | 787 \pm 100 | 942 \pm 200 | 71 \pm 30 | 25000.000 | <u>1.08 \pm 0.03</u> | 1.47 \pm 0.05 |
| 1-Occlusion | 667 \pm 100 | 1376 \pm 2000 | 2.65 \pm 0.08 | 20011 \pm 6000 | 1.14 \pm 0.02 | 1.379 \pm 0.003 |
| m^* -Occlusion | 673 \pm 70 | 727 \pm 400 | 9 \pm 1 | [16283 \pm 10] | 1.17 \pm 0.02 | 1.382 \pm 0.007 |
| 1-Occlusion ² | 709 \pm 100 | 1086 \pm 1000 | 2.65 \pm 0.08 | 20222 \pm 6000 | 1.14 \pm 0.02 | 1.378 \pm 0.003 |
| m^* -Occlusion ² | 642 \pm 60 | 710 \pm 400 | 9 \pm 1 | [16033 \pm 700] | 1.16 \pm 0.03 | 1.381 \pm 0.008 |
| WoutersNet 1-Occlusion | 6454 \pm 4000 | n/a | 2.9 \pm 0.6 | 20278 \pm 3000 | n/a | n/a |
| WoutersNet m^* -Occlusion | 4408 \pm 4000 | n/a | 11.7 \pm 0.6 | [16124 \pm 300] | n/a | n/a |
| WoutersNet GradVis | 6230 \pm 4000 | n/a | 2.9 \pm 0.5 | 19539 \pm 5000 | n/a | n/a |
| WoutersNet Input * Grad | 4988 \pm 4000 | n/a | 3.0 \pm 0.6 | 20546 \pm 4000 | n/a | n/a |
| WoutersNet Saliency | 5878 \pm 4000 | n/a | 2.8 \pm 0.5 | 20151 \pm 4000 | n/a | n/a |
| ZaidNet 1-Occlusion | 2236 \pm 2000 | n/a | 2.6 \pm 0.3 | 20696 \pm 2000 | n/a | n/a |
| ZaidNet m^* -Occlusion | 2485 \pm 2000 | n/a | 9 \pm 2 | [16124 \pm 300] | n/a | n/a |
| ZaidNet GradVis | 2560 \pm 3000 | n/a | 3.0 \pm 0.8 | 22790 \pm 3000 | n/a | n/a |
| ZaidNet Input * Grad | 3234 \pm 3000 | n/a | 2.5 \pm 0.3 | 21600 \pm 3000 | n/a | n/a |
| ZaidNet Saliency | 2295 \pm 2000 | n/a | 2.4 \pm 0.4 | 22561 \pm 3000 | n/a | n/a |
| ALL (ours) | [459 \pm 40] | [394 \pm 20] | [2.22 \pm 0.01] | 17582 \pm 5000 | 1.11 \pm 0.02 | [1.363 \pm 0.007] |

1500 **Distribution of performance during hyperparameter sweeps** In Fig. 20 we show the distribution of
 1501 performance during the random hyperparameter searches for **ALL** and selected baselines. We see that **ALL**
 1502 convincingly outperforms all baselines besides m^* -Occlusion on all datasets. Additionally, it has a higher peak
 1503 performance than m^* -Occlusion on every dataset except for DPAv4, and a higher median performance on
 1504 ASCADv1-fixed, ASCADv1-variable, and AES-HD. As previously noted and considered again in subsequent
 1505 ablation studies, we can further improve the performance of **ALL** relative to m^* -Occlusion by mimicking its
 1506 smoothing effect with stride-1 average-pooling.

1507 **E.3.7 Ablation studies**

1508 **ALL** has many differences from prior work, so here we run ablation studies to evaluate the impact of some of
 1509 the important individual differences. For each ablated design decision, we run a new 50-trial hyperparameter
 1510 sweep and plot the distribution of performance in terms of oracle agreement. Results are shown in Fig. 21,
 1511 with the salient results from Fig. 20 copied over for reference. We ablate the following design decisions:

1512 **Heavy input dropout for the supervised classifiers used by baselines** See the distributions labeled
 1513 ' m^* -Occl + heavy dropout'. The **ALL** classifier is trained on occluded inputs with the possible 'heaviness'
 1514 of the occlusion chosen spanning a wide range and chosen as a hyperparameter to optimize our proxy for
 1515 oracle agreement. In contrast, the baseline methods are based on 'interpreting' fixed classifiers which have
 1516 been trained with input dropout chosen from $\{0.0, 0.1\}$ to optimize classification performance. A plausible
 1517 conjecture is that **ALL** has strong performance because the heavy input corruption 'encourages' the classifier
 1518 to compensate by leveraging a wider variety of input-output associations, whereas because the supervised
 1519 classifiers train on uncorrupted or lightly-corrupted inputs, they have no such 'incentive'. We test this
 1520 assumption by tuning the classifiers with input dropout chosen from $\{0.05, 0.1, \dots, 0.95\}$ (the same search
 1521 space that **ALL** uses for $\bar{\gamma}$). We plot the performance distribution for m^* -occlusion and 1-occlusion. For clarity
 1522 we omit GradVis, Saliency, LRP and Input * Grad, but these have similar trends as 1-occlusion. We did not
 1523 test OccPOI or second-order occlusion because they are costly and in our prior experiments second-order
 1524 m -occlusion performs similarly to m -occlusion and OccPOI performs poorly compared to baselines.

1525 We see that this heavy input dropout leads to significant improvements to the maximum and median
 1526 performance on all datasets, though **ALL** still convincingly outperforms 1-occlusion and m -occlusion except
 1527 on DPAv4 and OTiAiT. To our knowledge no prior work has explored the effect of regularization strategies on
 1528 leakage localization performance of methods which 'interpret' supervised classifiers, and this result suggests
 1529 that such research may be fruitful.

1530 **Adversarial → 'Cooperative' leakage localization** See the distributions labeled '**ALL** (cooperative)'.
 1531 The intuition behind **ALL** is that we train a noise distribution to distribute a fixed amount of noise to
 1532 minimize the performance of a classifier, and we then interpret noisier measurements as leakier. Along the
 1533 lines of this intuition, we could also train it to distribute noise to *maximize* the performance of a classifier,
 1534 then interpret less-noisy measurements as leakier. We try this approach and find that it typically degrades
 1535 performance relative to the adversarial version of the algorithm. We conjecture that this is because the
 1536 adversarial approach encourages the classifier to rely on a diverse assortment of input-output associations,
 1537 whereas the cooperative approach does the opposite.

1538 **Omitting the noise conditioning to the classifier** See the distributions labeled '**ALL** (unconditional)'.
 1539 We feed the occlusion mask as an auxiliary input to the **ALL** classifiers, motivated by our theory which views
 1540 it as a family of classifiers trained in an amortized manner via conditioning a single neural net. We test
 1541 omitting this auxiliary input and find mixed results. On ASCADv1-fixed and DPAv4, **ALL** achieves stronger
 1542 performance without this auxiliary input, whereas on AES-HD it achieves stronger performance with this
 1543 input and on ASCADv1-variable, OTiAiT, and OTP results are approximately the same. For future work it
 1544 is likely justifiable to simplify **ALL** by omitting this conditioning.

1545 **Average pooling **ALL** to mimic the smoothing effect of m^* -Occlusion** See the distributions labeled
 1546 '**ALL** + AvgPool(m^*)'. Recall that we have chosen the occlusion window size m^* by sweeping it over successive
 1547 odd numbers until finding a local maximum in oracle agreement. We find that this consistently improves

1548 performance. One plausible explanation for the performance improvement is that large-window occlusion has a
 1549 ‘smoothing’ effect which amounts to an assumption that temporally-close measurements have similar leakiness.
 1550 This is likely true for some datasets, and as 1-occlusion is always included in the window size search space,
 1551 m^* occlusion will never degrade performance relative to 1-occlusion. To test this explanation we average
 1552 pool **ALL** with stride 1, kernel size m^* , and zero-padding to preserve dimensionality, which creates a similar
 1553 smoothing effect. Our results seem to support this explanation, with average pooling significantly improving
 1554 the performance of **ALL** on the DPAv4 and AES-HD datasets, where m^* -occlusion also has significantly
 1555 stronger performance than 1-occlusion. Note that using oracle agreement to choose an occlusion/pooling
 1556 window size causes data contamination, so while it is fair to compare results for m^* -average-pooled **ALL** to
 1557 those of m^* -occlusion, they cannot fairly be compared to any other baseline.

1558 E.3.8 Theoretical and empirical computational cost of deep learning methods

1559 In table 15 we list the theoretical computational complexity of the considered methods, as well as the measured
 1560 wall clock time to run them. While for fairness our experiments use an equal number of training steps for **ALL**
 1561 and supervised learning, note that in general we suspect the latter will converge in fewer training steps; thus,
 1562 these measurements likely overestimate the practical wall-clock time of supervised training. Additionally,
 1563 note that **OccPOI** takes significantly more time than m -occlusion despite having the same computational
 1564 complexity; this is because it requires $\Omega(T)$ sequential forward passes through the model, whereas all forward
 1565 passes may be done in parallel for m -occlusion. There is additional variance for the runtime of this method
 1566 because we choose the attack dataset size to be as small as possible while comfortably allowing the classifier
 1567 to attain a correct-key rank of 0 – e.g. AES-HD takes the longest because we use the full attack dataset. We
 1568 omit the parametric statistics-based methods from consideration, but these are done on the CPU and take a
 negligible amount of time compared to the deep learning methods.

Table 15: Comparison of the computational cost of methods considered in our work. We denote by C_F
 and C_B the cost of a forward and backward pass through our neural net respectively, N the dataset size,
 n_{sup} and n_{all} the number of epochs required for supervised learning and **ALL** respectively, and T the data
 dimensionality. We omit the parametric statistics-based baseline methods because they are done on the CPU
 and take a negligible amount of time relative to the deep learning methods. Runtimes are reported as mean
 \pm standard deviation over 5 runs. † These methods are used to ‘interpret’ a trained supervised classifier, but
 for comparison we report the cost of running them *after* training the classifier. In practice one would also
 incur the cost of supervised training (first row).

| Method | Total FLOPS | A6000 minutes per trial | | | | | |
|--|--------------------------------------|-------------------------|------------------|-----------------|-----------------|-------------------|-----------------|
| | | ASCADv1 (fixed) | ASCADv1 (random) | DPAv4 (Zaid) | AES-HD | OTiAiT | OTP |
| Supervised training | $\Theta(Nn_{\text{sup}}(C_F + C_B))$ | 2.09 ± 0.01 | 3.68 ± 0.02 | 1.98 ± 0.04 | 1.78 ± 0.03 | 0.17 ± 0.02 | 0.448 ± 0.009 |
| GradVis † | $\Theta(N(C_F + C_B))$ | 0.0675 ± 0.0008 | 0.263 ± 0.002 | 0.0080 ± 0.0001 | 0.0384 ± 0.0003 | 0.0080 ± 0.0001 | 0.0666 ± 0.0005 |
| Saliency † | $\Theta(N(C_F + C_B))$ | 0.078 ± 0.001 | 0.300 ± 0.003 | 0.0086 ± 0.0002 | 0.048 ± 0.002 | 0.0085 ± 0.0002 | 0.0881 ± 0.0004 |
| Input * Grad † | $\Theta(N(C_F + C_B))$ | 0.080 ± 0.002 | 0.3023 ± 0.0007 | 0.0086 ± 0.0002 | 0.049 ± 0.002 | 0.0086 ± 0.0002 | 0.086 ± 0.003 |
| LRP † | $\Theta(N(C_F + C_B))$ | 0.080 ± 0.001 | 0.3045 ± 0.0008 | 0.0088 ± 0.0002 | 0.050 ± 0.002 | 0.00861 ± 0.00007 | 0.087 ± 0.005 |
| m -Occlusion † | $\Theta(NC_F T)$ | 0.1235 ± 0.0004 | 0.952 ± 0.003 | 0.0644 ± 0.0003 | 0.1781 ± 0.0007 | 0.018 ± 0.002 | 0.247 ± 0.002 |
| 2 nd -order m -Occlusion † | $\Theta(NC_F T^2)$ | 16.6 ± 0.1 | 327 ± 1 | 81.1 ± 0.5 | 68.4 ± 0.6 | 4.42 ± 0.02 | 70.6 ± 0.3 |
| OccPOI † | $\Omega(NC_F T)$ | 2.29 ± 0.09 | 7.6 ± 0.4 | 0.709 ± 0.009 | 36.5 ± 0.6 | 0.042 ± 0.002 | 0.0437 ± 0.0004 |
| ALL (Ours) | $\Theta(Nn_{\text{all}}(C_F + C_B))$ | 3.2 ± 0.2 | 4.7 ± 0.2 | 2.44 ± 0.05 | 2.28 ± 0.05 | 0.200 ± 0.005 | 0.323 ± 0.004 |

1569

1570 F Limitations

1571 For real side-channel leakage datasets we lack ground truth knowledge about the leakiness of each measurement,
 1572 and all evaluation metrics considered in our paper have limitations. In particular:

1573 • The ‘oracle’ leakage assessments used in the main paper ignore leakage of order greater than 1 except
 1574 where higher-order leakage may be decomposed into first-order leakage of multiple variables, similarly
 1575 to the analysis of Egger et al. (2022). This relies on careful analysis of implementations by humans,
 1576 and is subject to error and oversights. Additionally, the SNR is not guaranteed to detect even
 1577 first-order leakage – it is sensitive only to the influence of the secret variable Y on the *mean* of each

1578 $X_t | Y$, and will not detect cases where the distribution of $X_t | Y$ changes with the mean remaining
 1579 fixed (e.g. if X_t is Gaussian distributed with Y -dependent variance and Y -independent mean).

- 1580 • The DNN occlusion tests and similar metrics proposed by Hettwer et al. (2020) are sensitive only to
 1581 the associations between X and Y that the neural net exploits. The superior performance of **ALL**
 1582 compared to prior deep learning-based algorithms, as well as work in other domains such as Geirhos
 1583 et al. (2020); Hermann & Lampinen (2020), suggests that DNNs are prone to exploiting some but
 1584 not all of the associations at their disposal. Additionally, such metrics may be biased in favor of deep
 1585 learning methods due to both leveraging SGD-trained DNNs.
- 1586 • Evaluation via feature selection for a Gaussian template attack as done by Masure et al. (2019); Yap
 1587 et al. (2023) is sensitive only to the top ≈ 10 leakiest measurements identified by an algorithm and
 1588 ignores all others, which is a major limitation. Additionally, for second-order datasets a template
 1589 attack will be unsuccessful unless *both* of a leaking pair of variables (e.g. r_3 and $\text{Sbox}(w_3 \oplus k_3) \oplus r_3$
 1590 for ASCADv1) are leaked through the selected measurements, and because no method considered in
 1591 our work except for second-order occlusion has any ability to discern the particular variable leaking
 1592 through a measurement, this would rely on random chance.

1593 We believe it is important for work in this domain to use a variety of evaluation strategies to compensate for
 1594 these individual limitations. This is similar to image synthesis research where it is common to use precision,
 1595 recall, and FID score to avoid the individual limitations of each of these metrics.

1596 **ALL** (alongside the other deep learning-based method we consider) has major advantages over manual analysis
 1597 and simpler parametric methods. However, it also comes with its own limitations:

- 1598 • Deep learning methods such as ours have a far larger computational cost than parametric methods.
 1599 Additionally, **ALL** is more computationally-expensive than the prior deep learning algorithms apart
 1600 from OccPOI and second-order occlusion. This increased cost is not fully reflected in our reported
 1601 runtimes because while for fairness we have run both **ALL** and supervised learning-based methods for
 1602 the same number of training steps and post-hoc early-stopped the latter, in practice we find that
 1603 supervised learning usually converges to good solutions in fewer training steps than **ALL**.
- 1604 • Deep learning methods such as ours require hyperparameter tuning and will give incorrect leakiness
 1605 estimates if poorly tuned. We view **ALL** as *complementary* with traditional approaches rather than
 1606 as a replacement. Traditional approaches can cheaply exploit domain/implementation knowledge
 1607 and detect many types of leakage for a low computational cost. **ALL** can then be used to search for
 1608 additionally leaking measurements not detected by these approaches.

1609 Additionally, a limitation of the literature on deep side-channel leakage localization which we do not address
 1610 in this work is that experiments are done at a smaller scale than the state of the art in deep side-channel
 1611 attacks, and the considered algorithms would likely need to be significantly scaled up for practical use. The
 1612 largest-scale experiments in this and prior work consider the ASCADv1-variable dataset, which consists of
 1613 300k 1400-measurement power traces and can be attacked with simple MLP and CNN architectures. The
 1614 1400-length power traces are extracted from longer 250k-length traces via downsampling and cropping in
 1615 the general vicinity of known leaky instructions which were themselves located with a white box analysis.
 1616 We believe that a critical direction for future work in this area is scaling existing methods to larger-scale
 1617 architectures applied to uncropped high-dimensionality datasets, such as the transformer architecture of
 1618 Bursztein et al. (2023) and the raw ASCADv2 dataset (Masure & Strullu, 2023), which consists of 800k
 1619 1M-measurement power traces.

1620 Our experiments consider mainly temporally-synchronized power trace datasets, which we believe is reasonable
 1621 because in practice implementation designers are likely able to collect synchronized traces, as the dataset
 1622 authors have done (e.g. by using the clock line of the hardware to trigger an oscilloscope). In practice a
 1623 violation of this condition would mean that leakage appears ‘spread out’ (e.g. see row 3 of Fig. 4), making it
 1624 more-challenging for designers to identify its source.

1625 A limitation of ALL, alongside the parametric statistical methods and OccPOI, is that they produce only a
1626 single vector summarizing leakage over the entire dataset, rather than for individual traces. As observed by
1627 Wouters et al. (2020), the neural net attribution methods can also be used to assign leakiness estimates to
1628 individual traces, though to our knowledge no work has systematically studied this ability. This would be a
1629 useful direction for future work, and would likely require innovation in performance evaluation strategies
1630 beyond those introduced in this paper or prior work.

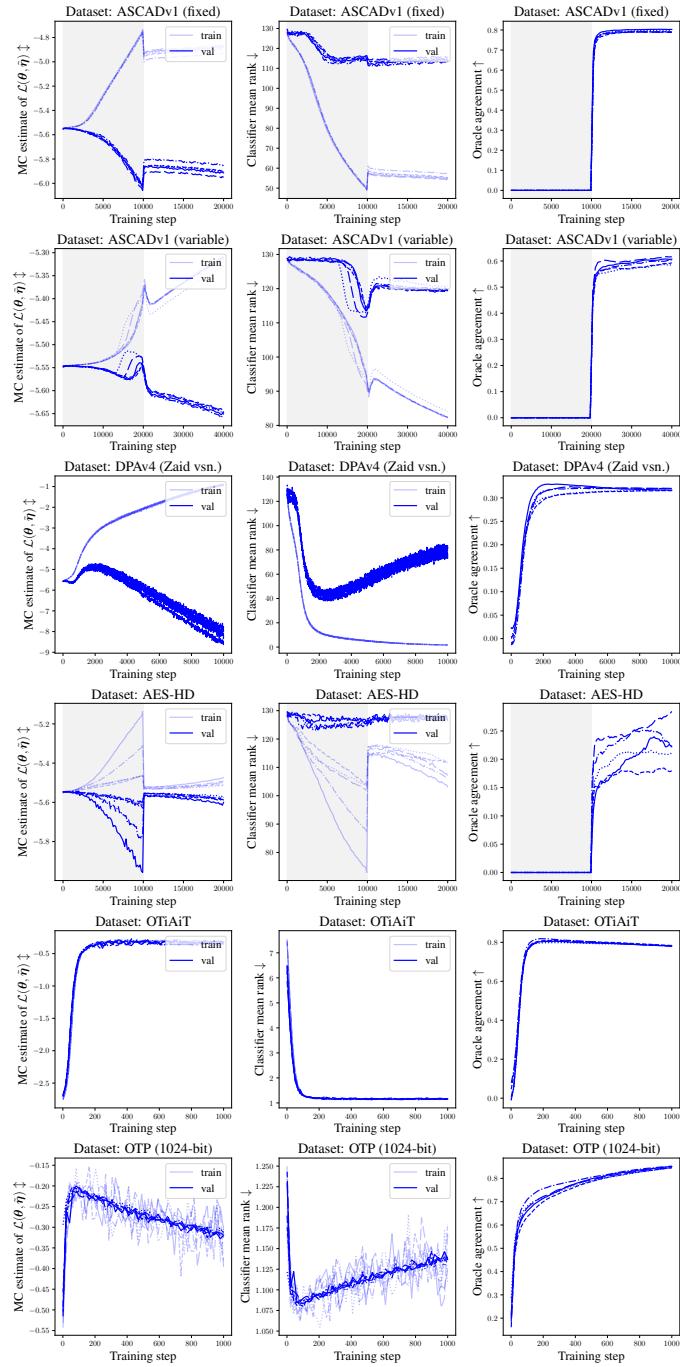


Figure 15: Training curves for ALL with the hyperparameter configuration chosen using our composite model selection criterion. Grey shaded regions denote ‘pretraining’ phases where we optimize θ but not $\tilde{\eta}$. Note that jumps in the traces happen because we set $\bar{\gamma} = 0.5$ for pretraining and change it to the setting chosen through hyperparameter search for training. **(left column)** The per-minibatch estimates of our objective function $\mathcal{L}(\theta, \tilde{\eta})$ (i.e. negative cross-entropy classification loss of the classifier) during training. We are optimizing θ to maximize this value and $\tilde{\eta}$ to minimize it. **(center column)** The mean rank of the correct label in the logits of the classifier (lower corresponds to higher classifier performance). We use this instead of accuracy for its finer granularity. **(right column)** The performance in terms of oracle agreement over the course of training; higher is better, as it indicates that the current ALL leakiness estimates are closer to being a strictly-increasing function of the oracle assessments.

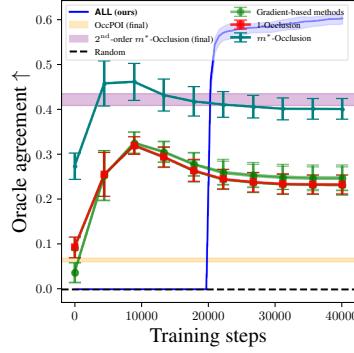


Figure 16: A comparison of the evolution of oracle agreement vs. training steps for ALL and selected baselines. Note that the oracle agreement for ALL is flat for the first 20k training steps and jumps up for the remaining steps because we leave all elements of γ fixed at 0.5 during our ‘pretraining’ phase of training, and update it only during the second half. For 2nd-order m^* -Occlusion and OccPOI, due to their high computational cost we report only the final performance after training via horizontal lines, rather than the evolution of performance during training.

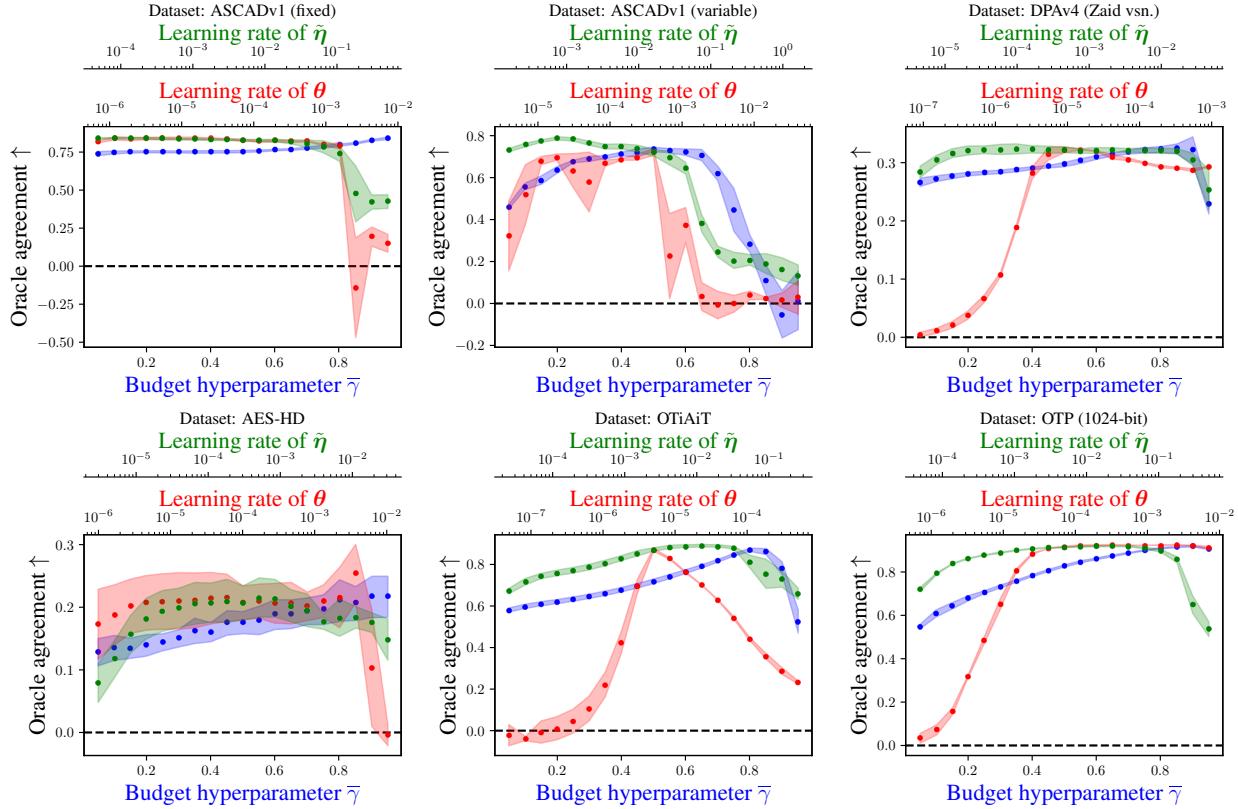


Figure 17: A plot of the performance of ALL as we perturb its 3 main hyperparameters: the noise budget $\bar{\gamma}$ and the learning rates of the classifier weights θ and the noise distribution parameter $\tilde{\eta}$. We find that performance varies smoothly with these hyperparameters and stays significantly better than random guessing over a large region of the search space, which is useful from a standpoint of hyperparameter tuning. We consider $\bar{\gamma}$ values in `np.arange(0.05, 1.0, 0.05)` and learning rates scaled by values in `np.logspace(-2, 2, 19)` relative to their optimal values according to the oracle agreement metric. All hyperparameters other than those being perturbed are left at their optimal values according to the oracle agreement metric. We repeat trials for 5 random seeds and report their mean with dots and ± 1 standard deviation with shading.

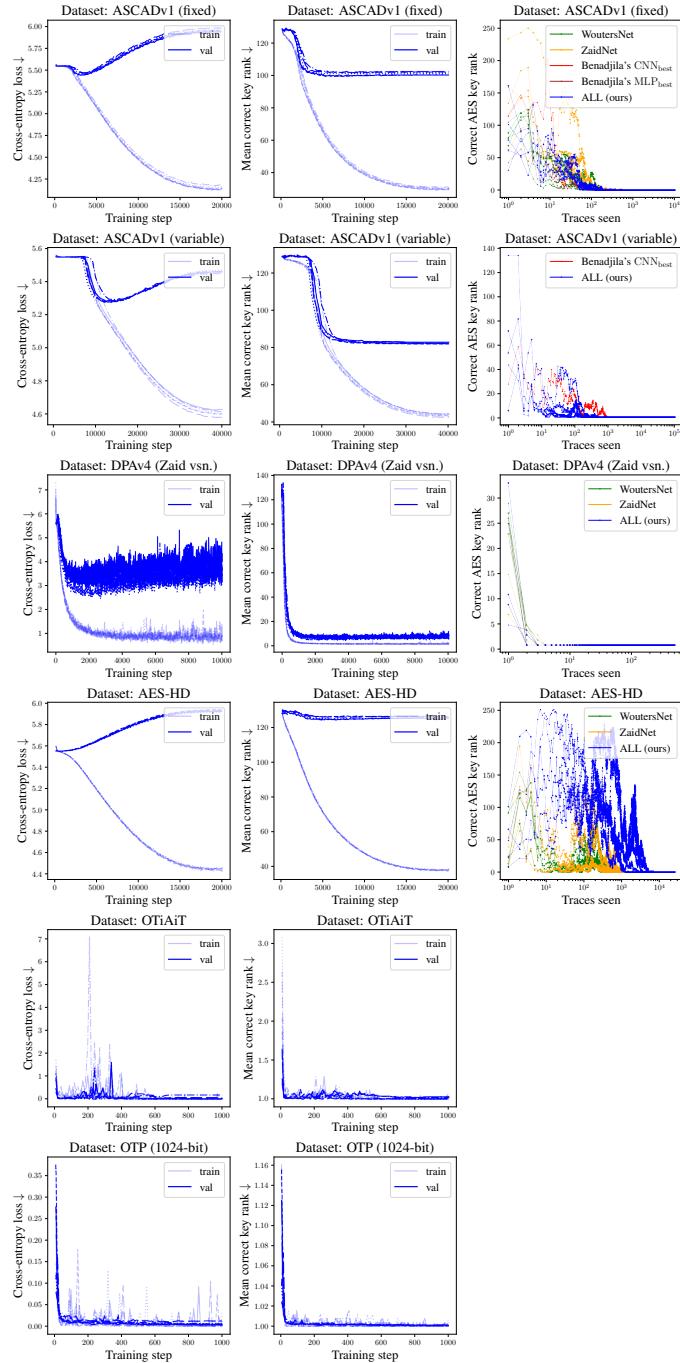


Figure 18: Training curves and attack performance of the supervised classifiers which are ‘interpreted’ by the deep learning baseline methods. **(Left column)** The training + validation cross-entropy loss vs. training steps for the supervised classifiers. **(center column)** The training + validation rank vs. training steps for the supervised classifiers. **(right column)** The rank of the correct key as we accumulate predictions on the attack dataset for the supervised classifiers, which is a common way of evaluating attack performance in the side-channel literature. For reference, we superimpose the results using open-weight classifiers provided by Benadjila et al. (2020) and Wouters et al. (2020). Note that our classifiers can successfully attack all datasets. They get comparable performance to the open-weight classifiers on ASCADv1-fixed, ASCADv1-variable and DPAv4, and somewhat worse performance on AES-HD.

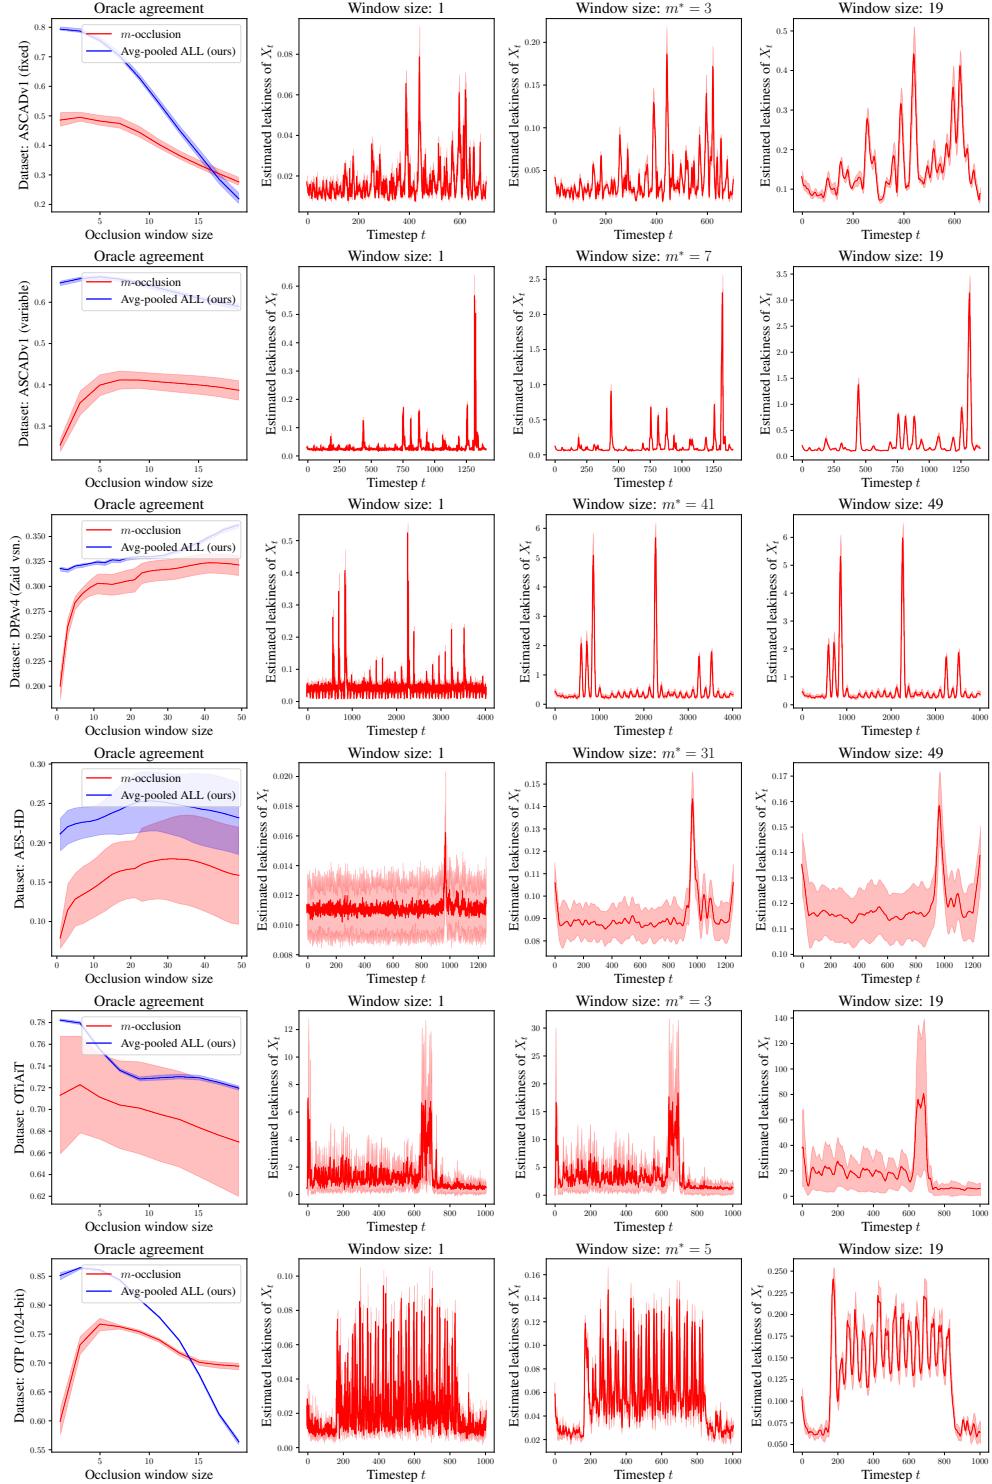


Figure 19: A sweep of the window size m for m -Occlusion. In the **first column** we plot the oracle agreement vs. window size (red curve). For reference we also plot the output of ALL (blue curve) as we average-pool it with stride 1 and kernel size m . Note that ALL consistently outperforms m -Occlusion for a wide range of window sizes. In the **second, third and fourth columns** we plot the m -Occlusion leakiness assessment for $m = 1$, optimal m , and maximum considered m . Observe that there is a smoothing effect as we increase m .

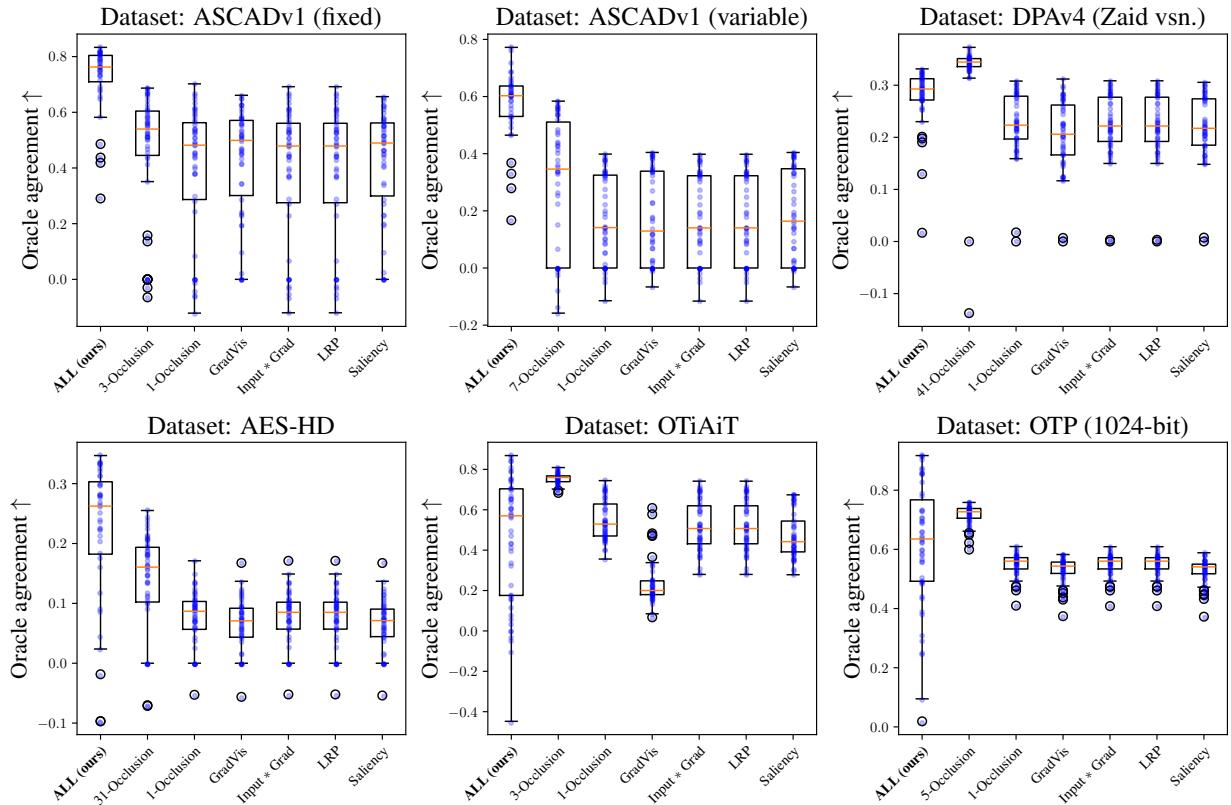


Figure 20: Distribution of performance of **ALL** and selected baseline methods during a 50-trial random hyperparameter search. Note that **ALL** generally outperforms baselines over a wide range of configurations. Blue dots denote individual samples, and boxes extend from first quartile to third quartile with a line at the median and whiskers extending to the furthest dot lying within $1.5 \times$ the interquartile range from the box.

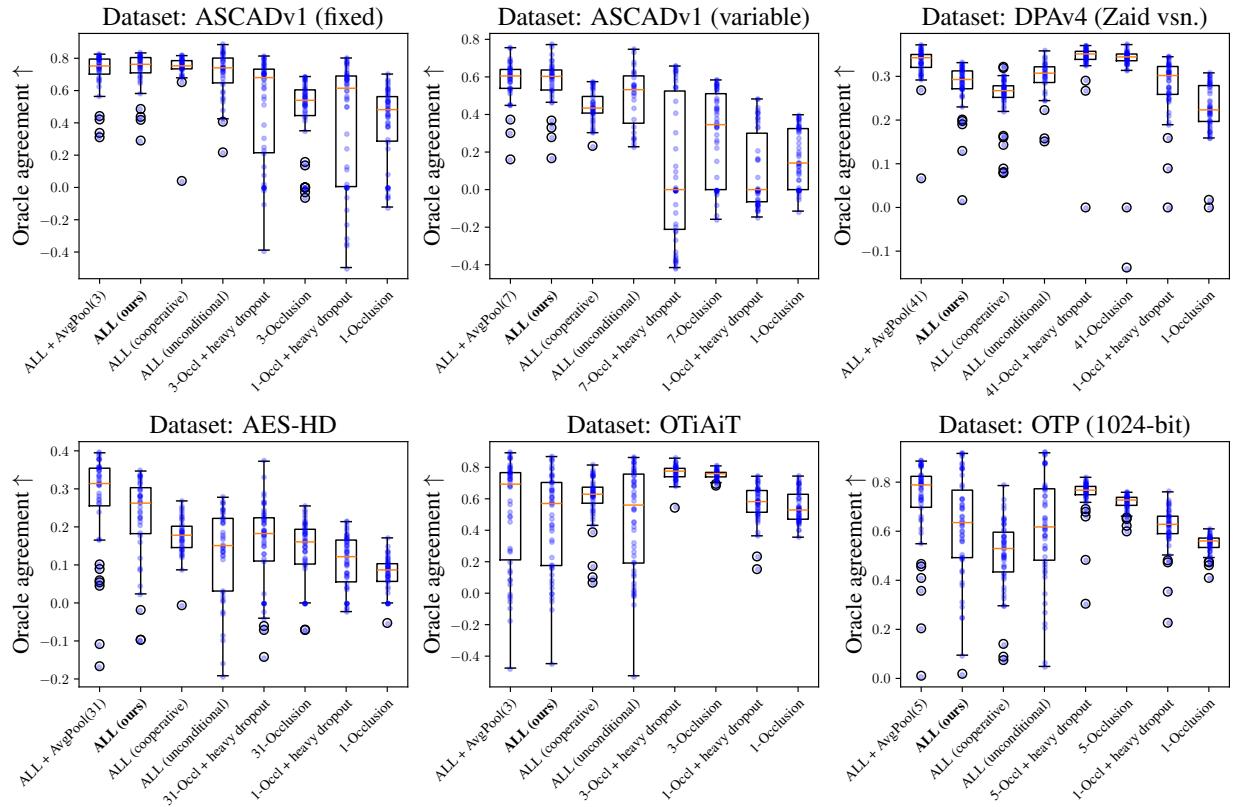


Figure 21: Ablation studies where we evaluate the influence of individual aspects of ALL on its performance gains relative to prior work, as described in Sec. E.3.7. These are plots of the distribution of oracle agreement values after a 50-trial random hyperparameter search with each ablation in place. ‘ALL (ours)’ denotes our method without modification. ‘ALL + AvgPool(m^*)’ denotes an average-pooled version of ALL to mimic the smoothing effect of m^* -Occlusion. ‘ALL (cooperative)’ denotes ALL with the adversarial objective replaced by a ‘cooperative’ objective where both the classifier and noise distribution are trained to maximize the performance of the classifier. ‘ALL (unconditional)’ denotes ALL without the occlusion masks being fed as an auxiliary input to the classifier. ‘ m -Occl + heavy dropout’ denotes m -Occlusion with the input dropout to the classifier chosen from $\{0.05, 0.1, \dots, 0.95\}$ rather than $\{0.0, 0.1\}$ as in our other experiments. ‘ m -Occlusion’ denote the results with the unmodified m -Occlusion techniques.