

HO-FMN: Hyperparameter optimization for fast minimum-norm attacks

Raffaele Mura^{a,1}, Giuseppe Floris^{a,1}, Luca Scionis^{a,b,1}, Giorgio Piras^{a,b}, Maura Pintor^{a,*}, Ambra Demontis^a, Giorgio Giacinto^a, Battista Biggio^a, Fabio Roli^{c,a}

^a University of Cagliari, Department of Electrical and Electronic Engineering, Cagliari, Italy

^b Sapienza University of Rome, Department of Computer Engineering, Rome, Italy

^c University of Genoa, Department of Computer Science, Bioengineering, Robotics and Systems Engineering, Genoa, Italy

ARTICLE INFO

Keywords:

Machine learning security
Adversarial examples

ABSTRACT

Gradient-based attacks are a primary tool to evaluate robustness of machine-learning models. However, many attacks tend to provide overly-optimistic evaluations as they use fixed loss functions, optimizers, step-size schedulers, and default hyperparameters. In this work, we tackle these limitations by proposing a parametric variation of the well-known fast minimum-norm attack algorithm, whose loss, optimizer, step-size scheduler, and hyperparameters can be dynamically adjusted. We re-evaluate 12 robust models, showing that our attack finds smaller adversarial perturbations without requiring any additional tuning. This also enables reporting adversarial robustness as a function of the perturbation budget, providing a more complete evaluation than that offered by fixed-budget attacks, while remaining efficient. We release our open-source code at <https://github.com/pralab/HO-FMN>.

1. Introduction

Machine learning (ML) models are susceptible to adversarial examples [1,2], i.e., input samples that are intentionally perturbed to mislead the model. Such samples are optimized using gradient-based attacks, which allow one to efficiently find adversarial perturbations close enough to the original unperturbed samples. However, using gradient-based attacks can only provide an *empirical* estimation of adversarial robustness. In particular, if an attack fails to find an adversarial example, we cannot prove that the given input is robust (i.e. there are adversarial manipulations that make the input adversarial, but the attack was not able to find it), similarly to what happens when searching for bugs in software. Ideally, through an exhaustive search or formal verification procedure, it would be possible to provide a guaranteed robustness assessment, which is however practically infeasible due to the high computational requirements and dimensionality entailing the problem [3]. This means that gradient-based attacks are most likely to provide an overly-optimistic estimation of adversarial robustness, and obtaining more reliable evaluations is not trivial [4]. It has been indeed shown that several defenses proposed to improve robustness to adversarial examples were wrongly evaluated, and rather than improving adversarial robustness they were simply *obfuscating gradients*, thereby invalidating the optimization of gradient-based attacks [4–7]. This problem is also exacerbated by the fact that each

attack presents different hyperparameters which require careful tuning to be executed correctly, i.e., to find a better optimum, along with fixed choices for the loss function, optimizer, and step-size scheduling algorithm. Nevertheless, in many of the reported evaluations, such attacks are run with their default settings, even if it has been shown that this may result in overestimating adversarial robustness [4,6,7]. Finally, many robustness evaluations are obtained by running fixed-budget attacks that only provide the adversarial robustness estimate at a fixed perturbation budget ϵ , without providing any insight on the robustness of the model when adversarial perturbations have a different size. To summarize, the main problems hindering the diffusion of more reliable adversarial robustness evaluations are: (i) the use of fixed loss, optimizer, and step-size scheduler, along with default attack hyperparameters; and (ii) the use of fixed-budget attacks, providing only limited insights on how models withstand adversarial attacks.

In this work, we aim to overcome these limitations by improving the current version of the Fast Minimum-Norm (FMN) attack originally proposed in [8]. To this end, we first propose a modular reformulation of the FMN attack that enables the use of different loss functions, optimizers, and step-size schedulers (Section 2). This facilitates the task of finding the strongest FMN configuration against each given model. We then leverage Bayesian optimization to perform a hyperparameter-optimization step that, for any given FMN configuration, automatically

* Corresponding author.

E-mail address: maura.pintor@unica.it (M. Pintor).

¹ These authors contributed equally to this work.

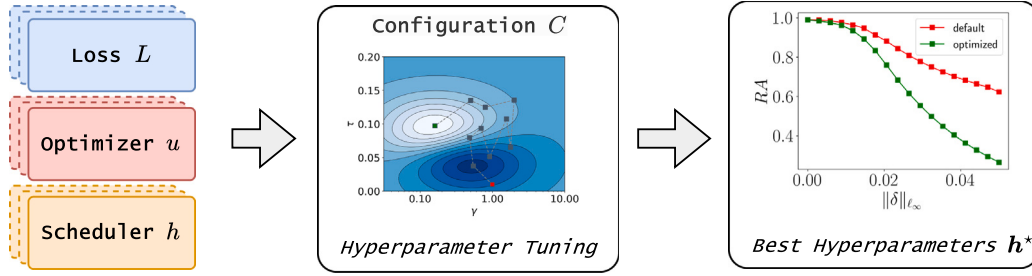


Fig. 1. Overview of our HO-FMN approach.

finds the best hyperparameters for the optimizer and the scheduler of choice (Section 3). An overview of our method, referred to as Hyperparameter Optimization for Fast Minimum-Norm (HO-FMN) attacks, is presented in Fig. 1. We extensively evaluate HO-FMN on 12 robust models against competing baseline attacks, supporting the validity of our method on efficiently obtaining complete robustness evaluation curves of ML models (Section 4). With respect to our preliminary work in [9], we extend the current approach by revisiting the FMN attack algorithm and rethinking the hyperparameter optimization framework with a Bayesian approach. In addition, we expand our experimental setup to get a more accurate evaluation. We conclude by discussing related work (Section 5), along with the main contributions, limitations, and future research directions (Section 6).

2. Revisiting fast minimum-norm attacks

We first present the FMN attack as originally proposed in [8], highlighting the changes applied to obtain the modular version of the attack algorithm in which we parameterize the loss function, the optimizer, and the step-size scheduler, along with their hyperparameters. The proposed reformulation of the algorithm enables the selection of each component independently, creating multiple parametric variations of the original attack.

Notation. Our goal is to discover minimum-norm adversarial perturbations that cause a model to misclassify an input. Let $x \in \mathcal{X} = [0, 1]^d$ represent a d -dimensional input data point with true label $y \in \mathcal{Y} = \{1, \dots, Y\}$. We denote the target function as $f : \mathcal{X} \times \Theta \mapsto \mathcal{Y}$, where $\theta \in \Theta$ is its set of parameters. We will utilize f for the label prediction function and f_y to refer to the continuous output (logit) corresponding to each class $y \in \mathcal{Y}$.

Attack Formulation. Minimum-norm attacks aim to find the smallest perturbation possible δ^* for which a sample x labeled as y gets misclassified by a model with parameters θ , i.e. $f(x + \delta^*; \theta) = \arg \max_{y \in \{1, \dots, Y\}} f_y(x + \delta^*; \theta) \neq y$. Their goal is thus to solve the following optimization problem:

$$\delta^* \in \arg \min_{\delta} \|\delta\|_p, \quad (1)$$

$$\text{s.t. } L_{\text{LL}}(x + \delta, y, \theta) < 0, \quad (2)$$

$$x + \delta \in [0, 1]^d, \quad (3)$$

where $\|\cdot\|_p$ indicates the chosen ℓ_p norm ($p = 1, 2, \infty$). The constraint in Eq. (2) is the difference-of-logits (LL) loss, defined as:

$$L_{\text{LL}}(x, y, \theta) = f_y(x, \theta) - \max_{j \neq y} f_j(x, \theta). \quad (4)$$

This loss is negative when the input is misclassified. The box constraint in Eq. (3) ensures that the sample $x + \delta$ remains in the feasible input space.

The Fast Minimum-Norm (FMN) attack [8] proposes a reformulation of the minimization problem to find the smallest ϵ for which the constraint is satisfied. The problem is reformulated as follows: as follows:

$$\min_{\epsilon, \delta} \epsilon, \quad \text{s.t. } \|\delta\|_p \leq \epsilon, \quad (5)$$

Algorithm 1: Fast Minimum-Norm (FMN) Attack Algorithm.

Input : x , the input sample; y , the target (true) class label; α_0 , the initial δ -step size; K , the total number of iterations; L , the attack loss; s , the step size scheduler; u , the optimizer.

Output: The minimum-norm adversarial example x^* .

```

1  $\epsilon_0 = \infty, \delta_0 \leftarrow \mathbf{0}, \delta^* \leftarrow \infty, \gamma_0 = 0.05$   $\triangleright$  initialization
2 for  $k = 1, \dots, K$  do
3    $\gamma_k \leftarrow s_\gamma(\gamma_0, k, K)$   $\triangleright \epsilon$ -step size decay
4    $\epsilon_k = u_\epsilon(\epsilon_{k-1}, \gamma_k, \|\delta_k\|_p)$   $\triangleright \epsilon$ -step
5    $g \leftarrow \nabla_{\delta} L(x + \delta_{k-1}, y, \theta)$   $\triangleright$  loss gradient
6    $\alpha_k \leftarrow s(\alpha_0, k, K)$   $\triangleright$  scheduler step
7    $\delta_k \leftarrow u(\delta_{k-1}, \text{proj}(g), \alpha_k)$   $\triangleright$  optimizer  $\delta$ -step
8    $\delta_k \leftarrow \Pi(x, \delta_k, \epsilon_k)$   $\triangleright$  proj. onto feasible domain
9 return  $x^* \leftarrow x + \text{best}(\delta_0, \dots, \delta_K)$   $\triangleright$  return best solution

```

plus the constraints in Eqs. (2)–(3), where ϵ is an upper bound on the perturbation size $\|\delta\|_p$. We now discuss the algorithm used by FMN.

FMN Attack Algorithm. We report in Algorithm 1 a revisited formulation of the FMN attack, in which we emphasize our specific contributions to make it parametric to its components. First, the attack is initialized (line 1), where the initial perturbation is set to $\delta = \mathbf{0}$ and the initial constraint is set to $\epsilon_0 = \infty$ to encourage the initial exploration of the loss landscape without encountering constraints in this phase.² Then, the original FMN algorithm develops as a two-step process: the ϵ -step minimizes the upper bound constraint on the maximum perturbation by reducing ϵ as long as the sample is adversarial, and the δ -step updates the perturbation towards the adversarial region trying to find it within the constraint defined by ϵ . The ϵ -step is controlled by a parameter γ_k that modifies the multiplicative factor for the current ϵ . The parameter γ_k , in turn, is reduced with cosine annealing decay (line 3). At each iteration, ϵ is reduced (increased) by a factor $1 - \gamma_k$ (by a factor $1 + \gamma_k$) if $x + \delta$ is (not) adversarial (line 4). Subsequently, the δ -step updates the perturbation with the gradient of the loss function $L(x, y, \theta)$ (line 5). While FMN uses the Logit Loss (LL) of Eq. (4), we modify the algorithm to work with any differentiable loss L .

The FMN attack normalizes the gradients $\nabla_x L(x, y, \theta)$ in the ℓ_2 norm, i.e. $g' = g / \|g\|_2$, and multiplies it by a step size α_k . In our formulation, we generalize this step with a linear projection (proj) onto a unitary-sized ℓ_p -ball. The projection maximizes a linear approximation of the gradient within a unitary ℓ_p -ball, as $g' = \arg \max_{\|v\|_p \leq 1} v^\top g$. This is accomplished, in ℓ_∞ , by taking the sign of the gradient $\text{sign}(g)$, and produces a dense update of all components of δ_k . Without loss of generality, the projection can be achieved in ℓ_1 and ℓ_2 by changing the norm used in the maximization.

² We simplify the algorithm by removing a refined estimate of ϵ_0 that approximates the distance to the boundary using the gradient of the loss used by FMN, as it might require computing an additional gradient when using our general algorithm with different losses.

In FMN, the step size is decayed with a decay schedule rule (line 6). In the original formulation, the decay was regulated with a Cosine Annealing Learning Rate scheduler (CALR). Our algorithm makes the scheduler s parametric, unlocking new scheduler rules for tuning the step size.

The perturbation is then updated with the computed δ -step (line 7), where we modify the original gradient descent (GD), replacing it with a generic optimizer u that can use different algorithms, e.g. momentum strategies. The perturbation is then projected onto the ϵ -ball and clipped to maintain the modified sample within the input domain (line 8). Finally, the best perturbation is returned (line 9).

Summary of Changes from FMN. While the overall algorithm remains conceptually unchanged, we adjust the attack loss L , the optimizer u , and the step-size scheduler s used in the δ -step to make them interchangeable. These elements were fixed in the original attack implementation, while in our work, we make them general and allow multiple choices for each component. The generalization of the loss L required an additional modification to the original algorithm, where the ϵ -step size was computed in the initial steps by estimating the distance to the boundary to speed up convergence. This estimation required computing the gradient of the LL loss, thus would require an additional backward pass for each iteration. With preliminary experiments, we found that such estimation improves the query efficiency of the initial steps, but it does not change substantially the final outcome of the attack. In addition, we use the linear projection of the gradient g instead of the normalization. Most significantly, our changes to the FMN algorithm required a thorough reevaluation of its implementation. Specifically, we enabled the choice of losses, optimizers, and schedulers that were already available in widely-adopted deep learning frameworks, which are commonly used (and efficiently implemented) for training deep neural networks.

Why FMN. Contrary to fixed-budget attacks (such as PGD [10] and AA [11]), FMN finds minimum-norm adversarial examples, solving the optimization problem in Eq. (1). It follows that, instead of having a scalar robustness evaluation associated with a predefined perturbation budget ϵ from a single run, we can obtain an entire curve, which we denote as a *robustness evaluation curve* [12], plotting how the robust accuracy of a model decreases as the perturbation budget ϵ is increased. The curve can be computed efficiently from the minimum distances $\|\delta^*\|_p$ returned by FMN, by computing:

$$\frac{1}{|D|} \sum_{(x,y) \in D} \mathbb{I}(\|\delta^*\|_p \leq \epsilon \wedge f(x + \delta^*) \neq y), \quad (6)$$

for increasing values of ϵ , where \mathbb{I} is the indicator function, which returns 1 (0) if the argument is true (false). These curves enable a more complete and informative robustness evaluation than that provided for a fixed ϵ value [6].

3. Hyperparameter optimization for fast minimum-norm attacks

A graphical representation of our HO-FMN method is depicted in Fig. 1. By leveraging the modular version of FMN presented in Section 2, and selecting a pool of losses, optimizers, and step-size schedulers, we (i) create multiple configurations of the FMN attack, (ii) optimize the hyperparameters of each configuration through Bayesian optimization, and rank them based on their effectiveness against the model under test, and (iii) run the attack with the best configurations found to estimate the robustness of the model. We discuss the choice of the configurations in Section 3.1, and, in Section 3.2, the subsequent optimization framework to get the best hyperparameters, compactly represented also in Algorithm 2.

3.1. HO-FMN: Configurations and hyperparameters

In our modular FMN re-implementation, we define the loss L , the optimizer u , and the step-size scheduler s as parametric. Accordingly, by defining a pool of losses \mathbb{L} , optimizers \mathbb{U} , and schedulers \mathbb{S} , we create diverse FMN configurations, each represented as a tuple $C = (L, u, s)$. We then define the hyperparameter search space by associating to each configuration C the set of hyperparameters $h = h_u \cup h_s$ of the corresponding optimizer u and scheduler s composing the final attack (Section 3.2). Next, given an input model, we optimize N configurations $C_1, C_2, \dots, C_N \in \mathbb{C}$ to find the best set of hyperparameters for each (Section 3.2).

Configuration Set. We denote the set of attack configurations as

$\mathbb{C} := \{C_1, C_2, \dots, C_N\}$, each represented as (L, u, s) . The total number N of possible configurations is thus obtained as the Cartesian product of each set, i.e., $N = |\mathbb{C}| = |\mathbb{L}| \times |\mathbb{U}| \times |\mathbb{S}|$. Each configuration C_i corresponds to the modular version of the FMN attack using a specific loss $L \in \mathbb{L}$, optimizer $u \in \mathbb{U}$, and scheduler $s \in \mathbb{S}$ in its attack algorithm. All-in-all, given a model with parameters $\theta_m \in \Theta$, our framework starts by considering all N (or less, since not every optimizer is necessarily associated with a scheduler) configurations.

Hyperparameter Search Space. Upon defining the configurations, the entire set \mathbb{C} undergoes a hyperparameter optimization routine. The goal of this routine is to find, for a given target model, the best set of hyperparameters h_i^* to be associated with each configuration C_i . Note that the best set might change from one model to another. Thus the hyperparameter optimization step has to be performed anew when a different model is selected. The search space and dimensionality of h_i vary depending on the configuration C_i , as the chosen optimizer and scheduler may take different (and a different number of) arguments as inputs. In this regard, given a configuration $C_i \in \mathbb{C}$, identified by (L, u, s) , its set of hyperparameters is given as $h_i = h_{i_u} \cup h_{i_s}$ (where h_{i_u} and h_{i_s} represent, respectively, the optimizer and scheduler hyperparameters, as described above). Hence, after creating the set \mathbb{C} , the optimization procedure aims to find the best set h_i^* for each C_i .

3.2. HO-FMN: Optimization procedure

Algorithm 2: HO-FMN.

Input : $D = (x, y)$, the validation dataset; C , the configuration with loss L , optimizer u , and scheduler s ; T , the number of trials; P , the number of initial samples to fit the regressor.

Output: The set of best hyperparameters h_i^* .

```

1  $S = \emptyset$ , best_median =  $\infty$            ▷ observation history
2 for  $j = 1, \dots, P$  do
3    $h_j = \text{gen\_h}()$                        ▷ sample first hyps
4    $x_j^* = \text{FMN}_{C, h_j}(x, y)$            ▷ initial observations
5    $\|\delta_j\| = \text{med}(\|x - x_j^*\|)$        ▷ compute median
6    $S \leftarrow S \cup (h_j, \|\delta_j\|)$      ▷ update observations
7 for  $j = P + 1, \dots, T$  do
8    $\text{gpr.fit}(S)$                            ▷ fit GP regressor
9    $h_j = \text{a}(\text{gpr.mean}, \text{gpr.std})$      ▷ acquire new  $h$ 
10   $x_j^* = \text{FMN}_{C, h_j}(x, y)$            ▷ new observations
11   $\|\delta_j\| = \text{med}(\|x - x_j^*\|)$        ▷ compute median
12   $S \leftarrow S \cup (h_j, \|\delta_j\|)$      ▷ update observations
13  if  $\|\delta_j\| < \text{best\_median}$  then
14     $h_j^* = h_j$                            ▷ store best  $h$ 
15     $\text{best\_median} = \|\delta_j\|$            ▷ update best median
16 return  $h^*$                                ▷ return best solution
```

We show in Algorithm 2 the complete HO-FMN procedure, which we use to find h_i^* for each C_i . This amounts to finding the set h^* that

minimizes the median perturbation $\|\tilde{\delta}\| = \text{med}(\|\mathbf{x} - \mathbf{x}^*\|)$, where med is the median function and $\mathbf{x}^* = \text{FMN}_{C,h}(\mathbf{x}, y)$, i.e., the output of Algorithm 1 when using configuration C with hyperparameters h . We select the median perturbation size as the objective to minimize for obtaining the best hyperparameters, following [8], as it reduces the impact of potential outliers that may substantially affect other metrics (e.g., the mean), and as it also represents the distance for which 50% of the samples become adversarial.

To avoid a computationally-demanding grid search on the hyperparameter space, *Bayesian Optimization* (BO) [13] can be leveraged to build a differentiable approximation of how the objective (i.e., the median perturbation size, in our case) changes as a function of the input hyperparameters. Accordingly, gradient descent can be used to efficiently optimize the choice of the best hyperparameters, while improving the approximation of the objective function after each evaluation. **Bayesian Optimization.** In our case, BO enables estimating the median perturbation that would be achieved at the end of the FMN algorithm, but without running it for all values. It requires setting a number of trials T , which refers to the number of times a new input will be sampled and a new output computed. Within T , an initial number of P trials are used as a “preliminary” stage to get a first set of observations to fit an initial model (which approximates the objective). Specifically, we first sample a set of P initial hyperparameters (line 3) with an external pseudo-random generation process, run the FMN algorithm with them, and collect a set of pairs $(h, \|\tilde{\delta}\|)$ (line 4–6). Then, we fit a Gaussian Process Regression (GPR) model on the observed median over the collected trials (line 8). GPR is a probabilistic model in which multiple regression functions are fitted and averaged, thus reporting, as a function of the value of h , a mean and uncertainty value of the metric $\|\tilde{\delta}\|$. When fitted, the GPR model predicts $\|\tilde{\delta}\|$ from a set of hyperparameters h .

For the remaining $T - P$ trials, the BO process involves the definition of an *acquisition function* $a()$, which defines the exploration strategy for navigating the possible hyperparameters. This function will produce new samplings of h to improve the approximation provided by the GPR model (line 9). Accordingly, the new values of the hyperparameters are chosen where the acquisition function is maximized. As the acquisition function, we use Noisy Expected Improvement (NEI), which aims to balance exploration (i.e. testing new *unexplored* values of the hyperparameter space) and exploitation (i.e. refining solutions closer to already seen values) in the search space. NEI extends the Expected Improvement (EI) criterion. The EI criterion is defined as the expectation on a candidate of its improvement over the function being estimated:

$$\text{EI}(x) = \mathbb{E}[(f(x) - f_{\min}) \cdot \mathbb{I}(f(x) > f_{\min})], \quad (7)$$

where $f(x)$ is the objective function to be optimized, f_{\min} is the current best observed value, and $\mathbb{I}(\cdot)$ is the indicator function. In the presence of noise, directly evaluating $f(x)$ can be unreliable. Therefore, NEI incorporates this uncertainty by modeling the noise in the objective function [14].

As new hyperparameters are evaluated, new $(h, \|\tilde{\delta}\|)$ pairs are collected, thus the GPR model is updated for an improved estimate (line 10–12).

At each iteration, the algorithm tracks the best median found (and the corresponding hyperparameters) in order to return the best solution (line 14).

The process is repeated over each of the N configurations to iteratively improve the approximation and find the best set of hyperparameters (h_1^*, \dots, h_N^*) , which are returned at the end of the algorithm (line 16).

As an example of the BO process, we show in Fig. 2 a GP regressor, isolated on the learning rate (γ) and momentum (μ) hyperparameters of a GD optimizer.³ The plot shows the mean (left) and standard deviation

(right) of the estimated $\|\tilde{\delta}\|$ for each pair of hyperparameters, obtained with $T = 32$ trials. Based on the uncertainty estimate, we notice how the acquisition function has focused on exploitation when sampling small, close-to-each-other learning rates, as the uncertainty grows for growing learning rates (indicating sporadic sampling of higher values).

4. Experimental analysis

In this section, we present the experimental details and results of our proposed optimization framework. HO-FMN, given a set of configurations C and a model M_m parameterized by θ_m , finds the best configuration C (Section 3.1) for which it identifies the best set of hyperparameters h^* (Section 3.2) on which to run the attack. Therefore, we first detail the general experimental setting details (Section 4.1). Then, we describe the creation of the configurations, the hyperparameters associated with each configuration, and the results (Section 4.2). To validate the benefits of our approach, after finding the best hyperparameters for each configuration-model pair, we run the FMN attacks and compare them with other competing attacks, as well as with the FMN baseline (Section 4.3). Additionally, we conduct a study on the computational overhead of our attack, showing that our proposed method is the best trade-off between runtime and completeness of the evaluation.

4.1. Experimental settings

We list here the main experimental details employed throughout both hyperparameter optimization and attack runs. We implement HO-FMN in PyTorch, and we run all experiments in a workstation equipped with an NVIDIA RTX A6000 GPU 48 GB. We implement our Bayesian Optimization (BO) search with the Ax framework.⁴

Datasets. We take a subset of 4096 samples (32 batches of 128 samples each) from the CIFAR-10 test set. Instead, for the ImageNet dataset, we selected 1000 samples to show how the framework can be extended and scaled. These samples serve to optimize HO-FMN, thus it can be seen as a training set. Then, for testing the capabilities of the optimized HO-FMN attack, we run the attacks with the best configurations on a separate set of 1000 samples for both the CIFAR-10 and ImageNet test set.

Perturbation Model. We restrict our analysis to the ℓ_∞ -norm perturbation model, as it is widely used in SoA evasion attacks and benchmarks [15]. We also point out that, in the original paper, the ℓ_∞ perturbation model is observed to be the most challenging for FMN [8].

Models. We consider 12 state-of-the-art robust models from the Robust-Bench repository [15], denoted as M_1 - M_{12} . We aim to verify the effectiveness of our method, with a wide range of robust models. The first 9 models are trained for robustness on the CIFAR-10 dataset on a perturbation budget of $\epsilon = 8/255$; the remaining 3 models are trained for robustness on ImageNet, within a perturbation budget of $\epsilon = 4/255$. M_1 , M_2 [16], a WideResNet-70-16 and a WideResNet-28-10 respectively, leverage an improved denoising diffusion probabilistic model (DDPM) to enhance adversarial training. M_3 , M_5 , M_9 [17], a WideResNet-70-16, a WideResNet-28-10 and a ResNet-18, use generative models to synthetically expand the original dataset and improve model resilience against ℓ_p norm attacks. M_4 [18], a WideResNet-106-16, use a combination of heuristics-based data augmentations and model weight averaging to improve the model’s robustness. M_6 , M_8 [19], respectively a WideResNet-70-16 and a WideResNet-28-10, use a self-consistent robust error measure to balance robustness and accuracy. M_7 [20], a ResNet-152, uses proxy distributions from diffusion models to enhance adversarial training. Finally, M_{10} , M_{11} , and M_{12} , are transformer architectures adversarially trained using ℓ_∞ norm perturbation bounded

³ Although the sampling of hyperparameters involves an entire set, we isolate it to create a 2-D visualization.

⁴ <https://ax.dev/versions/0.1.2/index.html>.

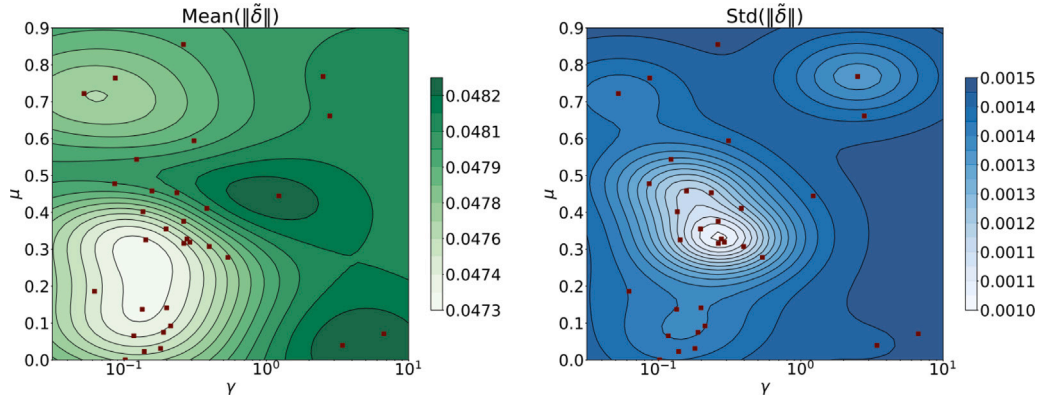


Fig. 2. Mean and standard deviation of the median perturbation size $\|\tilde{\delta}\|$ estimated by the GPR model, for a specific test configuration, as a function of the learning rate (γ) and momentum (μ) hyperparameters. The pairs (γ, μ) sampled during the process to iteratively refine the GPR model are shown as red points.

Table 1

Loss functions used in this work. We use $z_y(\mathbf{x}; \theta)$ to denote the softmax-scaled outputs of the model, and the indices π_1, \dots, π_Y to sort the logits as $f_{\pi_1} \geq \dots \geq f_{\pi_Y}$.

Loss function	Symbol	Equation
Cross-entropy	CE	$L_{\text{CE}}(\mathbf{x}, y; \theta) = \log(z_y(\mathbf{x}; \theta))$
Logits difference	LL	$L_{\text{LL}}(\mathbf{x}, y; \theta) = f_y(\mathbf{x}; \theta) - \max_{y \neq \pi} f_y(\mathbf{x}; \theta)$
Difference of logit ratio	DLR	$L_{\text{DLR}}(\mathbf{x}, y; \theta) = \frac{f_{\pi_1}(\mathbf{x}; \theta) - \max_{y \neq \pi} f_y(\mathbf{x}; \theta)}{f_{\pi_1}(\mathbf{x}; \theta) - f_{\pi_3}(\mathbf{x}; \theta)}$

at $\epsilon = 4/255$. The models M_{10} , M_{11} are two Swin-L and ConvNeXt-L models [21], while M_{12} is a ConvNeXt-L + ConvStem [22].

Performance Metrics. Within the optimization framework, we employ the smallest median perturbation $\|\tilde{\delta}\|$ as a criterion to find, for each configuration $C_1, \dots, C_N \in \mathbb{C}$, the best set of hyperparameters \mathbf{h}^* . The choice of the median follows the approach employed in the original FMN paper [8].

We then rank, based on the resulting $\|\tilde{\delta}\|$, the configurations $C_1, \dots, C_N \in \mathbb{C}$. We take the top-3 configurations, that we name C_1, C_2 , and C_3 (ordered in terms of performance, lowest median first) to evaluate the robust accuracy at ϵ (RA_ϵ) of the models at $\epsilon = 8/255$ for CIFAR-10, and at $\epsilon = 4/255$ for ImageNet (we denote it directly as RA in the rest of the section), following the RobustBench benchmark [15]. Moreover, as explained in Section 2, the benefit of FMN is that we can obtain RA by counting the successful attack samples that achieve misclassification with a perturbation size $\|\delta\|_\infty \leq \epsilon$, but we also get, with the same computational cost, the robustness evaluation curve [12].

Search Space for the Configurations. We first present the experimental settings for the hyperparameter optimization step. We list the configurations created from each loss L , optimizer u , and scheduler s , detailing the sets \mathbb{L} , \mathbb{U} , and \mathbb{S} respectively. As introduced in Section 3, we generalize the algorithm to use a selection of: (i) the loss function L , selecting between the logit loss (LL) [5], the cross-entropy loss (CE), and the difference of logits ratio (DLR) [11]; (ii) the optimizer u , selecting between Gradient Descent (GD), Adam (Adam), and AdaMax (AdaMax); and (iii) the step-size scheduler, selecting among Cosine Annealing (CALR) and Reduced On Plateau (RLRoP). We report the details of the loss used in Table 1.

As explained in Section 2, our reformulation of the FMN algorithm enables the use of components already implemented in existing libraries. Accordingly, we leverage the existing implementations of the aforementioned losses, optimizers, and schedulers as implemented in the PyTorch library, with a few exceptions. First, the LL and DLR losses are not implemented in PyTorch. Thus we took the implementations from the original FMN algorithm (LL) and from the AutoAttack repository (DLR).⁵ In addition, despite being the modular FMN version

adaptable to each kind of third-party scheduler compatible with the optimizers, we opt for a modified RLRoP implementation, as the original one adjusts a single learning rate (γ) for the entire batch. Namely, the original implementation of RLRoP decreases γ when there is no improvement (i.e., on plateau) on the average loss for a batch. However, we are interested in having a specific adaptation of the learning rate for each sample separately (as for each sample we are in a different region of the loss landscape, and we consider these optimization processes as independent from each other), thus we require a sample-wise RLRoP algorithm that tracks the improvement over the value of the loss on each sample rather than on the average loss of the batch. Therefore, we re-implemented the scheduler to have sample-wise control over the learning rates. Specifically, for a batch, we are seeking a vector of learning rates γ with one value for each sample in the batch. We configure a weighting vector initialized as $\mathbf{w} = \mathbf{1}$ containing one weight for each sample of the batch, and we obtain the learning rate by multiplying the weighting vector \mathbf{w} for the initial learning rate γ_0 , i.e., $\gamma = \mathbf{w}\gamma_0$. Subsequently, we track the individual loss for each sample, and we multiply by a reducing factor (< 1) the weight of the weighting vector w_i if the metric stops improving for sample i of the batch over a given number of iterations (as the *patience* parameter of RLRoP).

While GD is associated with a scheduler, Adam and AdaMax present an inner scheduling procedure, so we fix the scheduler to Fixed, i.e. no scheduler, in this case. Therefore, for each model, instead of having $|\mathbb{L}| = 3$, $|\mathbb{U}| = 3$, and $|\mathbb{S}| = 2$, for a total of $N = 18$ configurations, we reduce to $N = 12$ configurations.

Hyperparameters. We now list the set of hyperparameters \mathbb{H} associated to each configuration. As explained in Section 3.2, this induces a second level on the search space that is different for each C . Specifically, each optimizer u and each scheduler s comes with their own hyperparameters (respectively, \mathbf{h}_u and \mathbf{h}_s). To reduce the search space, we fix some of the hyperparameters that we denote as “fixed”. For the others, we define either the search ranges (Range) and, when different than linear, the scale used for the uniform sampling, or the possible choices (Choice). We list the hyperparameters search space \mathbb{H} , along with their sampling options, in Table 2.

Ax Framework Configuration. The Ax framework works by instantiating multiple trials sequentially. Specifically, we employed $T = 32$ trials, of which the first initialization set, i.e. $P = 8$, are quasi-randomly generated (using the SOBOL [23] approach), and the remaining 24 are sampled from the regression models, as implemented by the BOTORCH algorithm [24].

4.2. Hyperparameter optimization results

In this section, we present the results of the hyperparameter optimization. We first considered each pair of configurations and models (C_i, M_m) . Then, we tuned each configuration, finding the set of best

⁵ <https://github.com/fra31/auto-attack>.

Table 2

List of the chosen hyperparameters for each Optimizer and Scheduler selected for HO-FMN. As Optimizers, we chose GD, the one from the original FMN implementation, and Adam/AdaMax; the first requires a Scheduler while the others have an auto-scheduling mechanism. As Schedulers, we selected CALR and RLROp (our sample-wise implementation). The hyperparameters can be range, choice, or fixed; the sampling distribution can be uniform (default) or logarithmic for better exploring higher ranges. The Optimizers have the most configurable setting, resulting in a larger search space.

Optimizer	Hyperparameter	Search space
GD	Learning rate (γ)	range: [8/255, 10] logarithmic
	Momentum (μ)	range: [0.0, 0.9]
	Weight decay (λ)	range: [0.01, 1.0]
	Dampening (τ)	range: [0.0, 0.2]
Adam/AdaMax	Learning rate (γ)	range: [8/255, 10] logarithmic
	Weight decay (λ)	range: [0.01, 1.0]
	eps	fixed: $1e-8$
	betas (β_1, β_2)	range: [0.0, 0.999]
Scheduler	Hyperparameter	Search space
CALR	T_max	fixed: K
	eta_min	fixed: 0
	last_epoch	fixed: -1
RLROp	Batch size	fixed: $-^a$
	Factor	range: [0.1, 0.5]
	Patience	choice: [2, 5, 10]
	Threshold	fixed: $1e-4$
	eps	fixed: $1e-8$

^a The RLROp scheduler implements our sample-wise version, so the batch size parameter is removed.

hyperparameters h^* that achieve the smallest median perturbation. We ranked the configurations by $\|\tilde{\delta}\|$, and we selected the top-3 for each model.

Tuning Results. In Table 3, we show the resulting top-3 configurations that achieve the smallest $\|\tilde{\delta}\|$ for each model on CIFAR-10, and in Table 4, we present the corresponding top-3 configurations for ImageNet. We highlight how the DLR loss consistently finds better perturbations than CE and LL. In addition, we found that the GD-CALR-DLR configuration, ranked in the top-3 for each model, is also the best one in 6 over 9 models. It is worth noting that this configuration is also very close to the one used by the original FMN, though changing the loss from LL to DLR and having an optimized set of h . This loss works well across models as the normalization at the denominator makes the loss (and gradient) more scale invariant (compared to LL that can change of orders of magnitude from one model to the other), easing the tuning of the step size and of the other parameters for the optimization.

4.3. Best attacks

Given the top-3 configurations for each model, on the CIFAR-10 test set Table 3 and on the ImageNet test set Table 4, we run the final attack evaluation on our test sets, respectively 1000 samples from CIFAR-10/ImageNet, and compare with the baseline FMN version [8] to clearly show the benefits of HO-FMN. In addition, to rigorously validate with SoA, parameter-free approaches, we compare our attack configurations achieved through HO-FMN with the APGD attack in its CE and DLR loss versions [11]. We specify that comparing head-to-head the entire AA ensemble against a single attack from HO-FMN, would result in a four-versus-one evaluation, ultimately indicating a disproportionate analysis. Finally, to let the comparison be as fair as possible, we ensure that all algorithms are initialized in the same way. Specifically, we ensure that APGD does not perform Expectation over Transformations (EoT) steps, i.e., the computation of a smoothed gradient before the actual attack loop and the restarts. We removed this step as we want to avoid the variability given by the randomness in the EoT procedure. Thus, by avoiding this particular initialization, we can ensure that all attacks start from the same x and have no advantage

(or disadvantage) given by random initializations of δ_0 . We remark that the same method is not computed in the default configuration of APGD. Additionally, this initial EoT can also be seamlessly added later to FMN and HO-FMN.

Attack Results. We show the resulting robustness evaluation curves in Fig. 3 for the CIFAR-10 experiments; while in Fig. 4 we report the curves for the attacks against ImageNet models. We compare the curves of the baseline FMN attack against HO-FMN. The FMN baseline is defined as the original formulation in [8], thus configured with GD-CALR LL, and optimizer with $\gamma = 1.0$ and $\mu = 0.0$. In addition, we highlight, for the single perturbation norm of $\|\delta\| = 8/255$ (CIFAR-10) and $\|\delta\| = 4/255$ (ImageNet), the Robust Accuracy (RA) found by APGD_{CE/DLR}. We selected the attack that performed better, in terms of RA, between the two versions of APGD. We show the empirical results in Table 5 for CIFAR-10 and in Table 6 for ImageNet. In the first case, except for 2/9 models, HO-FMN outperforms the APGD attack (i.e., the blue line lies below the red cross). In the second case, we are able to beat all the 3 selected models with our HO-FMN version. Furthermore, our attack computes the robustness evaluation curve with one single run. Achieving the same result with APGD is only possible by executing APGD multiple times, as we discuss in the next paragraph.

Computational Overhead. We perform a set of additional experiments to have a clear understanding of the overhead added by running hyperparameter optimization. The total time of the HO process is mainly dictated by the time required for a single attack multiplied by the number of trials. Our tuning setting, as described in Section 4.1, consists of running 32 trials of HO, each running FMN on a batch of 128 samples over 200 steps. To analyze the HO overhead added to FMN, we measure the average execution time of a single FMN attack on the same setup and relate it to the number of trials. We refer to this time as T_{FMN} , for which we find $T_{FMN} = 7.479$ seconds. We then compute an estimate \tilde{T}_{HO} of the HO process, thus ignoring the Gaussian Processes (GP) overhead, as $\tilde{T}_{HO} = T_{FMN} \cdot 32 = 239.328$ seconds. Then, we run the actual HO-FMN under the same setup and measure the total execution time. On average, we find $T_{HO} = 262.612$ seconds, which indicates that the difference between our estimate and the measured time equals $\Delta_{HO} = \tilde{T}_{HO} - T_{HO} = 23.284$ seconds. Therefore, for a single trial, the required time amounts to $\Delta_{HO}/32 = 0.727$. Compared to T_{FMN} , we can assert that the overhead added by the optimization is acceptable in practice.

Comparing FMN against APGD. In Section 2, we show how FMN allows us to compute the robustness evaluation curves [12], which is instead practically unfeasible for fixed-budget attacks, such as APGD [11]. Being one of the main advantages of our HO-FMN approach the possibility to create robustness evaluation curves, we conduct additional experiments, summarized in Table 7, where we measure the required time for HO-FMN (GD-CALR-DLR/CE averaged) to compute the curves compared to APGD. In fact, through APGD, it is possible to have only a scalar robustness evaluation: for a given value of ϵ (i.e., the maximum perturbation that constrains the attack) APGD provides a scalar robust accuracy value associated to the given value ϵ . Therefore, to be compared with HO-FMN, we adapted APGD to find a minimum-norm solution using a binary search approach, that we applied sample-wise. Within this approach, we define a number of search steps that we set to 5, in addition to the search interval $[\epsilon_{low}, \epsilon_{high}]$. In particular, ϵ_{low} is the lowest value the perturbation budget can take, while ϵ_{high} is the highest. The binary search algorithm works by selecting a value for the perturbation budget which is always set as $(\epsilon_{high} - \epsilon_{low})/2$ (in the middle of the interval), and the interval is updated according to the successfulness of the attack. Hence, in the initialization phase, the search interval is set as $[\epsilon_{low}, \epsilon_{high}] = [0, 32/255]$ and at each step, APGD is run with a perturbation budget of $\epsilon_i = (\epsilon_{high} - \epsilon_{low})/2$ ($\epsilon_0 = 16/255$). If the attack finds a successful adversarial perturbation, we narrow the search to the lower half of the interval (i.e., $\epsilon_{high} = \epsilon_i$); otherwise, we search on the upper half (i.e., $\epsilon_{low} = \epsilon_i$). This process is repeated within the selected half-interval, progressively refining the

Table 3

Top-3 configurations after the hyperparameter optimization on each model (M_1 - M_9), along with the resulting median perturbation, i.e. $\|\tilde{\delta}\|$, on samples from the CIFAR-10 dataset. Then, in order, we show the learning rate (γ) and weight decay (λ), the beta coefficients ($\beta_{1,2}$) for Adam/AdaMax, and the momentum (μ) and dampening (τ) for GD. Finally, the last columns indicate the factor (fac.) and patience (pat.) for RLROp.

Model	C	u + s	L	$\ \tilde{\delta}\ $	OPTIM. (h_u)				SCHED. (h_s)
					γ	λ	β_1, β_2	μ, τ	fac./pat.
M_1	C ₁	GD + RLROp	DLR	0.048066	3.1373e-02	0.374	-	(0.5842, 0.148)	(0.345, 2)
	C ₂	AdaMax	DLR	0.048176	3.3608e-02	0.113	(0.491,0.868)	-	-
	C ₃	GD + CALR	DLR	0.048403	7.3636e-02	0.667	-	(0.3720,0.029)	-
M_2	C ₁	GD + CALR	DLR	0.048021	8.1149e-02	0.683	-	(0.0744,0.045)	-
	C ₂	Adam	DLR	0.048787	6.6351e-02	0.433	(0.412,0.000)	-	-
	C ₃	GD + RLROp	DLR	0.048801	4.2894e-02	0.111	-	(0.2158,0.100)	(0.160,2)
M_3	C ₁	Adam	DLR	0.050735	8.8306e-02	0.577	(0.688,0.713)	-	-
	C ₂	GD + CALR	DLR	0.050961	1.9881e-01	0.170	-	(0.1298, 0.173)	-
	C ₃	AdaMax	DLR	0.052110	3.1373e-02	0.982	(0.362,0.751)	-	-
M_4	C ₁	Adam	DLR	0.051502	3.1373e-02	0.435	(0.221,0.816)	-	-
	C ₂	GD + CALR	LL	0.051720	6.2728e-02	0.676	-	(0.4512,0.149)	-
	C ₃	GD + CALR	DLR	0.051725	3.1373e-02	0.924	-	(0.4195,0.130)	-
M_5	C ₁	GD + CALR	DLR	0.051760	2.9909e-01	0.010	-	(0.2493,0.105)	-
	C ₂	GD + CALR	CE	0.051958	3.8703e-02	0.511	-	(0.3857,0.074)	-
	C ₃	Adam	DLR	0.052237	3.1373e-02	0.697	(0.275,0.137)	-	-
M_6	C ₁	GD + CALR	DLR	0.047542	7.6798e-02	0.747	-	(0.4471,0.091)	-
	C ₂	Adam	DLR	0.047696	9.2127e-02	0.596	(0.687,0.286)	-	-
	C ₃	AdaMax	DLR	0.047820	8.8301e-02	0.279	(0.264,0.999)	-	-
M_7	C ₁	GD + CALR	DLR	0.049981	6.1492e-02	0.061	-	(0.3780,0.041)	-
	C ₂	GD + CALR	LL	0.050134	6.8632e-02	1.000	-	(0.1610,0.200)	-
	C ₃	Adam	DLR	0.050165	4.9743e-02	0.818	(0.622,0.255)	-	-
M_8	C ₁	GD + CALR	DLR	0.045454	2.9001e-01	0.010	-	(0.3191,0.097)	-
	C ₂	AdaMax	DLR	0.046265	5.4455e-02	0.248	(0.446,0.568)	-	-
	C ₃	GD + RLROp	DLR	0.046554	5.1549e-02	0.775	-	(0.8750,0.078)	(0.324,5)
M_9	C ₁	GD + CALR	DLR	0.043584	5.3307e-02	0.613	-	(0.7285,0.165)	-
	C ₂	GD + CALR	LL	0.043850	1.8606e-01	1.000	-	(0.0,0.200)	-
	C ₃	Adam	CE	0.044207	4.5904e-02	0.456	(0.104,0.496)	-	-

Table 4

Top-3 configurations for M_{10} - M_{12} , along with the resulting median perturbation, on samples from the ImageNet dataset. For further details please refer to Table 3.

Model	C	u + s	L	$\ \tilde{\delta}\ $	OPTIM. (h_u)				SCHED. (h_s)
					γ	λ	β_1, β_2	μ, τ	fac./pat.
M_{10}	C ₁	GD + CALR	LL	0.016130	8.2874e-02	0.266	(0.4962, 0.037)	-	-
	C ₂	AdaMax	DLR	0.017020	8.1110e-01	0.404	(0.566,0.098)	-	-
	C ₃	GD + CALR	DLR	0.017138	3.1036e-01	0.755	-	(0.8011,0.055)	-
M_{11}	C ₁	AdaMax	DLR	0.012425	1.9174e-01	0.304	(0.387,0.367)	-	-
	C ₂	Adam	DLR	0.013496	9.4361e-02	0.244	(0.354,0.231)	-	-
	C ₃	GD + RLROp	DLR	0.015658	2.1192e-02	0.109	-	(0.1497, 0.155)	(0.207,5)
M_{12}	C ₁	GD + CALR	LL	0.014309	1.0000e+01	0.582	(0.9000,0.200)	-	-
	C ₂	AdaMax	DLR	0.014430	5.8146e-01	0.401	-	(0.256,0.332)	-
	C ₃	Adam	DLR	0.014899	8.0030e+00	0.568	-	(0.806,0.222)	-

search until the maximum number of search steps is reached. As shown in Table 7, the first column is the total average time (in seconds) the attack took to complete. We show that our HO-FMN finds the best minimum-norm solution in a single run (first row). The next rows, indicated by APGD (i), represent the binary search step i performed by the two APGD versions. Table 7 shows that the best solution for APGD is found at step 4 (APGD (4)), while the binary search continues to step 5 with no improvement. Our FMN version, finds its best solution in approximately 5 s, while it takes about 20 s for each APGD version to complete the binary search, therefore showing the efficacy of our FMN approach in finding the robustness evaluation curve compared to APGD.

5. Related work

Adversarial attacks are recognized as an important tool to empirically evaluate the robustness of ML models. Many gradient-based attacks have been proposed as an effective tool to assess the models' robustness, and have evolved over time seeking for better efficiency. Among the most used attacks, the Projected Gradient Descent (PGD)

attack [10] has been extensively used as a bare essential evaluation tool. However, attacks like PGD, which solve an optimization problem, require a proper hyperparameter configuration (e.g., learning rate, step decay etc.) to avoid suboptimal solutions and, consequently, providing an overestimated adversarial robustness evaluation [6]. To mitigate this issue, parameter-free approaches that combine multiple attacks [11] have also been proposed.

AutoAttack (AA). This attack consists of ensembling 4 parameter-free attacks, including Auto-PGD (APGD), i.e., an attack that directly improves the basic PGD optimization by dynamically updating the step size. Together with APGD with both CE and DLR losses, AA also uses a gradient-based (Fast Adaptive Boundary [25]) and a black-box (SquareAttack [26]) attack, and ensembles them by retaining the first useful result found by any of them (within the fixed budget), in a sample-wise manner.

Adaptive Auto-Attack (AAA). This attack [27] provides a further evolved approach by conceiving the attacks as building blocks, thus having multiple interchangeable parts, and performing an extensive search, virtually constructing a huge ensemble of attacks. However, the implemented algorithm does not efficiently filter the searched

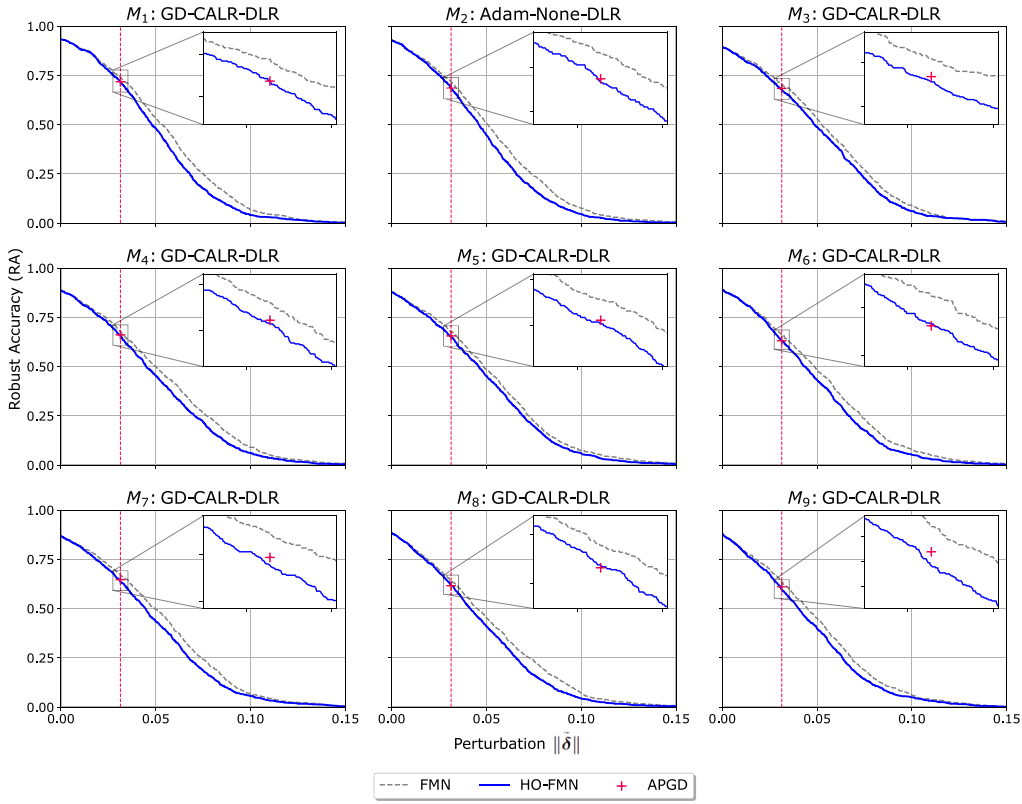


Fig. 3. Robustness evaluation curves for M_1 - M_9 . The dashed-gray and solid-blue lines represent FMN and HO-FMN. The robust accuracy (RA) value at $\epsilon = 8/255$ computed with $\text{APGD}_{\text{CE/DLR}}$ (the best value between the two) is also shown as a red cross.




Fig. 4. Robustness evaluation curves for M_{10} - M_{12} , and APGD robust accuracy at $\epsilon = 4/255$. Please refer to Fig. 3 for further details.

trials, thus potentially wasting computing resources, and does not optimize the hyperparameters of these attacks, potentially disrupting the evaluation results.

Limitations of Existing Methods. Just like standard fixed-epsilon attacks, the robustness evaluation for both AA and AAA are constrained to a single perturbation budget (e.g., $\epsilon = 8/255$), resulting in a scalar robustness estimate. Such characteristic of both AA and AAA inhibits them from constructing a full-scale robustness evaluation curve on multiple perturbation values, which would inevitably require multiple attack runs. As also noted in our work, constructing the full curve with these attacks requires running them multiple times to find the smallest ϵ that satisfies the attack's success.

In this direction, our proposed HO-FMN approach collectively takes advantage of the reliability of the parameter-free paradigm, as well as enabling a thorough robustness evaluation, contrary to the competing parameter-free approaches. Our results show the efficacy of the proposed hyperparameter optimization strategy when compared to the baseline FMN attack.

6. Conclusions and future work

In this work, we investigated the use of hyperparameter optimization to improve the performance of the FMN attack. Specifically, we reimplemented the FMN attack into a modular version that enables changing the loss, the optimizer, and the step-size scheduler to create multiple configurations of the same attack. We used Bayesian optimization to find the best attack hyperparameters for each configuration selected. Our findings highlight that hyperparameter optimization can improve FMN to reach competitive performance with existing attacks while providing a more thorough adversarial robustness evaluation (i.e., computing the whole robustness evaluation curve).

We argue that the same approach can be combined with other attacks and perturbation models. To this end, we plan to extend our analysis beyond the ℓ_∞ -norm FMN attack, considering ℓ_0 , ℓ_1 , and ℓ_2 norms. We remark that adding more hyperparameters to tune would make the search space bigger, resulting in a longer optimization time.

Table 5

Robust Accuracy (RA) with fixed perturbation $\epsilon = 8/255$ computed for each model $M_1 - M_9$ with, respectively, the **Baseline** FMN attack, the two **APGD**_{CE/DLR} versions and the top-3 HO-FMN configurations of each model (C_1, C_2, C_3). Except for two models, we beat both baseline and APGD attacks.

Model	Attack	RA	Model	Attack	RA	Model	Attack	RA
M_1	Baseline	0.744	M_2	Baseline	0.716	M_3	Baseline	0.704
	APGD _{DLR}	0.718		APGD _{DLR}	0.687		APGD _{DLR}	0.684
	APGD _{CE}	0.741		APGD _{CE}	0.716		APGD _{CE}	0.687
	C_1	0.724		C_1	0.688		C_1	0.683
	C_2	0.718		C_2	0.683		C_2	0.678
	C_3	0.717		C_3	0.693		C_3	0.681
M_4	Baseline	0.680	M_5	Baseline	0.679	M_6	Baseline	0.664
	APGD _{DLR}	0.661		APGD _{DLR}	0.659		APGD _{DLR}	0.631
	APGD _{CE}	0.678		APGD _{CE}	0.656		APGD _{CE}	0.658
	C_1	0.661		C_1	0.652		C_1	0.633
	C_2	0.661		C_2	0.658		C_2	0.637
	C_3	0.657		C_3	0.652		C_3	0.638
M_7	Baseline	0.672	M_8	Baseline	0.639	M_9	Baseline	0.635
	APGD _{DLR}	0.647		APGD _{DLR}	0.616		APGD _{DLR}	0.616
	APGD _{CE}	0.654		APGD _{CE}	0.651		APGD _{CE}	0.610
	C_1	0.638		C_1	0.618		C_1	0.596
	C_2	0.641		C_2	0.621		C_2	0.609
	C_3	0.638		C_3	0.624		C_3	0.616

Table 6

Robust Accuracy (RA) with fixed perturbation $\epsilon = 4/255$ computed for each model $M_{10} - M_{12}$ on the ImageNet dataset. We report the numerical results for, respectively, the **Baseline** FMN attack, the two **APGD**_{CE/DLR} versions and the top-1 HO-FMN configuration C_1 of each model. For all the models, we beat both baseline and APGD attacks.

Model	Attack	RA	Model	Attack	RA	Model	Attack	RA
M_{10}	Baseline	0.619	M_{11}	Baseline	0.619	M_{12}	Baseline	0.614
	APGD _{DLR}	0.611		APGD _{DLR}	0.614		APGD _{DLR}	0.609
	APGD _{CE}	0.608		APGD _{CE}	0.605		APGD _{CE}	0.594
	C_1	0.597		C_1	0.600		C_1	0.588

Table 7

Runtime comparison between HO-FMN (GD-CALR-DLR/CE) and APGD_{CE/DLR} adapted to find a minimum-norm solution (each row represents a binary search iteration). We show the total time and best median $\|\delta\|$ found by the attack on a batch of 128 samples from CIFAR-10 on model M_9 .

	Total (avg) time [s]	Best (median) $\ \delta\ $
HO-FMN _{CE (DLR)}	4.753 (5.257)	0.053 (0.053)
APGD _{CE (DLR)} 1	3.635 (4.064)	0.062 (0.062)
APGD _{CE (DLR)} 2	7.241 (8.078)	0.062 (0.062)
APGD _{CE (DLR)} 3	10.856 (12.094)	0.062 (0.062)
APGD _{CE (DLR)} 4	14.508 (16.105)	0.054 (0.054)
APGD _{CE (DLR)} 5	18.170 (20.141)	0.054 (0.054)

To this end, we will also develop sound heuristics to filter the suboptimal configurations without running the full attacks, making hyperparameter tuning more efficient. Additionally, we will design faster exploration phases in the initial steps of the FMN optimization process to enable further exploration of the loss landscape.

CRedit authorship contribution statement

Raffaele Mura: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Investigation, Data curation. **Giuseppe Floris:** Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Investigation, Data curation. **Luca Scionis:** Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Investigation, Data curation. **Giorgio Piras:** Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Maura Pintor:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Methodology, Investigation, Formal analysis, Conceptualization. **Ambr Demontis:** Writing – review & editing, Validation, Methodology,

Investigation, Conceptualization. **Giorgio Giacinto:** Writing – review & editing, Supervision, Project administration, Funding acquisition. **Battista Biggio:** Writing – review & editing, Writing – original draft, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization. **Fabio Roli:** Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work has been carried out while L. Scionis and G. Piras were enrolled in the Italian National Doctorate on AI run by the Sapienza University of Rome in collaboration with the University of Cagliari; and supported by project by project SERICS (PE00000014) and project FAIR (PE0000013) under the NRRP MUR program funded by the EU - NGEU; the European Union's Horizon Europe Research and Innovation Programme under the project Sec4AI4Sec, grant agreement No 101120393; and Fondazione di Sardegna under the project "TrustML: Towards Machine Learning that Humans Can Trust", CUP: F73C22001320007.

Data availability

We release our code publicly; it is available as a link on the abstract.

References

- [1] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrđić, P. Laskov, G. Giacinto, F. Roli, Evasion attacks against machine learning at test time, in: ECML PKDD, Part III, 8190 of LNCS, Springer Berlin Heidelberg, 2013, pp. 387–402.
- [2] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I.J. Goodfellow, R. Fergus, Intriguing properties of neural networks, in: ICLR, 2014.
- [3] A. Kumar, A. Levine, T. Goldstein, S. Feizi, Curse of dimensionality on randomized smoothing for certifiable robustness, in: International Conference on Machine Learning, PMLR, 2020, pp. 5458–5467.
- [4] M. Pintor, L. Demetrio, A. Sotgiu, A. Demontis, N. Carlini, B. Biggio, F. Roli, Indicators of attack failure: Debugging and improving optimization of adversarial examples, in: NeurIPS, 2022.
- [5] N. Carlini, D. Wagner, Towards evaluating the robustness of neural networks, in: Symp. Security and Privacy, SP, IEEE, 2017, pp. 39–57.
- [6] N. Carlini, A. Athalye, N. Papernot, W. Brendel, J. Rauber, D. Tsipras, I. Goodfellow, A. Madry, A. Kurakin, On evaluating adversarial robustness, 2019, arXiv preprint arXiv:1902.06705.
- [7] F. Tramer, N. Carlini, W. Brendel, A. Madry, On adaptive attacks to adversarial example defenses, in: NeurIPS, 2020.
- [8] M. Pintor, F. Roli, W. Brendel, B. Biggio, Fast minimum-norm adversarial attacks through adaptive norm constraints, in: NeurIPS, 2021.
- [9] G. Piras, G. Floris, R. Mura, L. Scionis, M. Pintor, B. Biggio, A. Demontis, Improving fast minimum-norm attacks with hyperparameter optimization, in: ESANN 2023, ESANN 2023, Ciaco - I6doc.Com, 2023.
- [10] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, A. Vladu, Towards deep learning models resistant to adversarial attacks, 2017, arXiv preprint arXiv:1706.06083.
- [11] F. Croce, M. Hein, Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks, in: ICML, 2020.
- [12] B. Biggio, F. Roli, Wild patterns: Ten years after the rise of adversarial machine learning, Pattern Recognit. 84 (2018) 317–331.
- [13] J. Snoek, H. Larochelle, R.P. Adams, Practical bayesian optimization of machine learning algorithms, in: Neural Information Processing Systems, 2 of NIPS'12, Curran Associates Inc., 2012, pp. 2951–2959.
- [14] B. Letham, B. Karrer, G. Ottoni, E. Bakshy, Constrained bayesian optimization with noisy experiments, Bayesian Anal. 14 (2) (2019) 495.
- [15] F. Croce, M. Andriushchenko, V. Sehwag, E. DeBenedetti, N. Flammarion, M. Chiang, P. Mittal, M. Hein, Robustbench: A standardized adversarial robustness benchmark, in: NeurIPS Datasets and Benchmarks, 2021.
- [16] Z. Wang, T. Pang, C. Du, M. Lin, W. Liu, S. Yan, Better diffusion models further improve adversarial training, in: ICML, 202 of PMLR, 2023, pp. 36246–36263.

- [17] S. Goyal, S. Rebuffi, O. Wiles, F. Stimberg, D.A. Calian, T.A. Mann, Improving robustness using generated data, in: *NeurIPS*, 2021.
- [18] S. Rebuffi, S. Goyal, D.A. Calian, F. Stimberg, O. Wiles, T.A. Mann, Fixing data augmentation to improve adversarial robustness, 2021, CoRR [abs/2103.01946](https://arxiv.org/abs/2103.01946) [arXiv:2103.01946](https://arxiv.org/abs/2103.01946).
- [19] T. Pang, M. Lin, X. Yang, J. Zhu, S. Yan, Robustness and accuracy could be reconcilable by (proper) definition, in: *ICML*, 2022.
- [20] V. Sehwal, S. Mahloujifar, T. Handina, S. Dai, C. Xiang, M. Chiang, P. Mittal, Robust learning meets generative models: Can proxy distributions improve adversarial robustness? in: *ICLR*, 2022.
- [21] C. Liu, Y. Dong, W. Xiang, X. Yang, H. Su, J. Zhu, Y. Chen, Y. He, H. Xue, S. Zheng, A comprehensive study on robustness of image classification models: Benchmarking and rethinking, 2023, [arXiv:2302.14301](https://arxiv.org/abs/2302.14301).
- [22] N.D. Singh, F. Croce, M. Hein, Revisiting adversarial training for imagenet: Architectures, training and generalization across threat models, 2023, [arXiv:2303.01870](https://arxiv.org/abs/2303.01870).
- [23] J. Bossek, C. Doerr, P. Kerschke, Initial design strategies and their effects on sequential model-based optimization: an exploratory case study based on bbob, in: *Proceedings of the 2020 Genetic and Evolutionary Computation Conference, GECCO '20*, ACM, 2020.
- [24] M. Balandat, B. Karrer, D.R. Jiang, S. Daulton, B. Letham, A.G. Wilson, E. Bakshy, Botorch: A framework for efficient monte-carlo bayesian optimization, 2020, [arXiv:1910.06403](https://arxiv.org/abs/1910.06403).
- [25] F. Croce, M. Hein, Minimally distorted adversarial examples with a fast adaptive boundary attack, in: *Int'l Conf. on Machine Learning*, PMLR, 2020, pp. 2196–2205.
- [26] M. Andriushchenko, F. Croce, N. Flammarion, M. Hein, Square attack: A query-efficient black-box adversarial attack via random search, 2020.
- [27] C. Yao, P. Bielik, P. Tsankov, M.T. Vechev, Automated discovery of adaptive attacks on adversarial defenses, in: *NeurIPS*, 2021, pp. 26858–26870.




Raffaele Mura is a Ph.D. student at the PRA Lab, in the Department of Electrical and Electronic Engineering of the University of Cagliari, Italy. He received his Master's degree in Computer Engineering, Cybersecurity, and Artificial Intelligence, from the University of Cagliari in September 2023. His research focuses on machine learning security.



Giuseppe Floris is a Ph.D. student at the PRA Lab, in the Department of Electrical and Electronic Engineering of the University of Cagliari, Italy. He received his Master's degree in Computer Engineering, Cybersecurity, and Artificial Intelligence, from the University of Cagliari in September 2023. His research focuses on machine learning security.




Luca Scionis (B.Sc. 2021, M.Sc. 2023) is a Ph.D. student at the PRA Lab, Department of Electrical and Electronic Engineering, University of Cagliari, Italy. He is currently enrolled in the national Ph.D-AI program of the Sapienza University of Rome, Italy. His research focuses on machine learning security.



Giorgio Piras (B.Sc. 2019, M.Sc. 2021) is a Ph.D. student at the PRA Lab, in the Department of Electrical and Electronic Engineering of the University of Cagliari, Italy. He is currently enrolled with the national Ph.D-AI program from the Sapienza University, Rome, Italy. His research focuses on machine learning security.



Maura Pintor is an Assistant Professor at the PRA Lab, in the Department of Electrical and Electronic Engineering of the University of Cagliari, Italy. She received her Ph.D. in Electronic and Computer Engineering from the University of Cagliari in 2022. Her research focuses on machine learning security.



Ambra Demontis is an Assistant Professor at the University of Cagliari, Italy. She received her M.Sc. degree (Hons.) in Computer Science and her Ph.D. degree in Electronic Engineering and Computer Science from the University of Cagliari, Italy. Her research interests include secure machine learning, kernel methods, biometrics, and computer security.



Giorgio Giacinto is Professor of Computer Engineering at the University of Cagliari, Italy, where he serves as the coordinator of the M.Sc. degree in Computer Engineering, Cybersecurity and Artificial Intelligence.

His research interests are in machine learning for malware analysis and detection, and he published around 200 papers in international conferences and journals. He is the Editor in Chief of the "Security Engineering & Applications" section of the *Journal of Cybersecurity and Privacy*. He is a member of the managing committees of the Cybersecurity National Lab and the Artificial Intelligence & Intelligent Systems Lab within the CINI consortium, Italy. He also represents the Cybersecurity National Lab within the European Cybersecurity Organization (ECSO). He is a Fellow of the IAPR (International Association for Pattern Recognition) and a Senior Member of the IEEE Computer Society and ACM.



Battista Biggio (M.Sc. 2006, Ph.D. 2010) is Assistant Professor at the University of Cagliari, Italy, and co-founder of the company Pluribus One. His research interests include adversarial machine learning and cybersecurity. He is Senior Member of the IEEE and of the ACM, and Member of the IAPR and ELLIS.



Fabio Roli is a Full Professor of Computer Engineering at the University of Genova and the University of Cagliari, Italy, and Founding Director of the Pattern Recognition and Applications laboratory at the University of Cagliari. He has been doing research on the design of pattern recognition and machine learning systems for thirty years. He has provided seminal contributions to the fields of multiple classifier systems and adversarial machine learning, and he has played a leading role in the establishment and advancement of these research themes. He has been appointed Fellow of the IEEE and Fellow of the International Association for Pattern Recognition. He is a recipient of the Pierre Devijver Award for his contributions to statistical pattern recognition.