
Sparse-RS: a versatile framework for query-efficient sparse black-box adversarial attacks

Francesco Croce
University of Tübingen

Maksym Andriushchenko
EPFL

Naman D. Singh
University of Tübingen

Nicolas Flammarion
EPFL

Matthias Hein
University of Tübingen

Abstract

A large body of research has focused on adversarial attacks which require to modify all input features with small l_2 - or l_∞ -norms. In this paper we instead focus on query-efficient *sparse* attacks in the black-box setting. Our versatile framework, Sparse-RS, based on random search achieves state-of-the-art success rate and query efficiency for different sparse attack models such as l_0 -bounded perturbations (outperforming established white-box methods), adversarial patches, and adversarial framing. We show the effectiveness of Sparse-RS on different datasets considering problems from image recognition and malware detection and multiple variations of sparse threat models, including targeted and universal perturbations. In particular Sparse-RS can be used for realistic attacks such as universal adversarial patch attacks without requiring a substitute model. The code of our framework is available at <https://github.com/fra31/sparse-rs>.

1 Introduction

The discovery of the vulnerability of neural networks to adversarial examples [9, 49] revealed that the decision of a classifier or a detector can be changed by small, carefully chosen perturbations of the input which do not modify its semantic content. Many efforts have been put into developing increasingly more sophisticated attacks to craft modifications imperceptible to humans but able to fool classifiers and bypass many defense mechanisms [13, 7]. Such imperceptibility is typically achieved by constraining or minimizing the l_p -norm of the perturbations, usually either l_∞ [49, 30, 13, 34, 18], l_2 [37, 13, 45, 18] or l_1 [14, 36, 18]. Metrics other than l_p -norms and more human perception-aligned have been also recently used, e.g. Wasserstein distance [53, 23]. However, it has been argued that, since all these attacks tend to modify all the elements of the input, they are not of practical concern.

A more realistic kind of adversarial perturbations are sparse attacks, where only a small percentage of components of the input are changed but possibly with large modifications. This can be done without imposing any relation among the modified features, the so-called l_0 -attacks [38, 13, 39, 46, 17, 41], or in a structured manner, e.g. with patches [11, 29, 31]. In this way, the perturbations are indeed visible but, since restricted to a small portion of the original image, do not hinder the ability of humans to correctly classify them, and can be applied in the physical world and not only at a digital level [31, 50, 32]. Moreover, such attacks generalize to tasks outside computer vision, such as malware detection or natural language processing, where the nature of the domain imposes to modify only a limited number of input features [21, 28].

A broadly usable attack should also work in the black-box score-based scenario, i.e. the attacker can only access the output of a classifier f , the score distribution over classes, for any input, but does not have access to the network weights parametrizing f , and in particular cannot use its gradient (as

in the white-box setup). We do not consider more restrictive (e.g. decision-based attacks [10, 12] where the adversary only knows the label assigned to each input) or more permissive (e.g. a surrogate model similar to the victim one is available [15, 25]) cases. However, while plenty of black-box attacks for standard threat models such as l_∞ or l_2 exist [27, 22, 47, 2, 35, 15, 4], few are aiming at l_0 -constrained perturbations [38, 46, 17], and none of them focuses on query efficiency and scales to datasets like ImageNet without suffering from prohibitive computational cost.

In this paper we propose `Sparse-RS`, a simple and flexible framework based on random search to generate sparse adversarial examples with a black-box score-based attack which is effective with minimal adaptations in a variety of tasks. First, we show for l_0 -attacks the superior performance of `Sparse-RS` in terms of success rate and query efficiency compared to existing black-box and even white-box attacks for quite different models, datasets and applications (image classification and malware detection), demonstrating its broad applicability. Then, we show results for the most challenging task: targeted universal adversarial patches and frames, a perturbation model that can potentially be realized by a physical attack. Up to our knowledge these are the first black-box attacks for these threat models. So far all other approaches for universal adversarial patches transferred the patch obtained by a white-box attack on a known surrogate model, e.g. [11]. Finally, we apply `Sparse-RS` in the scenario of image- and location-specific patches and show again its effectiveness in terms of query efficiency and success rate compared to a recently proposed black-box attack [54].

In the following we first recall the necessary background on adversarial attacks and random search, then we introduce our new framework, including a theoretical analysis for the l_0 -threat model, and finally present the results of `Sparse-RS` on several tasks, together with the description of the experimental setups and suitable competitors. Omitted proofs, experimental details, additional results and visualizations and ablation studies are available in the Appendix.

2 Black-box adversarial attacks

Let $f : \mathcal{S} \subseteq \mathbb{R}^d \rightarrow \mathbb{R}^K$ be a classifier which assigns an input $x \in \mathcal{S}$ to class $c = \arg \max_{r=1, \dots, K} f_r(x)$. The goal of an *untargeted* attack is to craft a perturbation $\delta \in \mathbb{R}^d$ such that

$$\arg \max_{r=1, \dots, K} f_r(x + \delta) \neq c, \quad x + \delta \in \mathcal{S} \quad \text{and} \quad \delta \in \mathcal{T}, \quad (1)$$

where \mathcal{S} is the input domain of the specific task and \mathcal{T} the constraints the adversarial perturbation has to fulfill (e.g. l_p -norm smaller than a threshold), while a *targeted* attack aims at δ such that

$$\arg \max_{r=1, \dots, K} f_r(x + \delta) = t, \quad x + \delta \in \mathcal{S} \quad \text{and} \quad \delta \in \mathcal{T}, \quad (2)$$

with t as target class. Generating such δ can be translated into an optimization problem as

$$\min_{\delta \in \mathbb{R}^d} L(f, x, \delta, c) \quad \text{s.t.} \quad x + \delta \in \mathcal{S} \quad \text{and} \quad \delta \in \mathcal{T} \quad (3)$$

by choosing a proper loss function L which minimized leads to the desired classification. We refer to the scenario in which an attack is carried out as a threat model: targeted vs untargeted, black- vs white-box, constraints set \mathcal{T} to satisfy, level of knowledge of the attacker.

Since a black-box attacker cannot rely on gradient-based methods to solve (3), many approaches have appeared: first, [26, 52] proposed to approximate the gradient through finite difference methods, later improved to reduce their high computational cost in terms of queries of the victim models [8, 51, 27]. Alternatively, [3, 33] use genetic algorithms in the context of image classification and malware detection respectively. A line of research has focused on rephrasing l_∞ -constrained attacks as discrete optimization problems [47, 2, 35], where specific techniques lead to significantly better query-efficiency. [19] sample candidate adversarial attacks via the Metropolis MCMC method, while [22] adopt a modification of canonical random search to produce perturbations with minimal l_2 -norm.

The method closest in spirit to our proposed one is the Square Attack of [4], which is state-of-the-art for l_∞ - and l_2 -bounded black-box attacks. It uses random search in order to iteratively generate samples on the surface of the l_p -ball of the chosen threat model. This, together with a carefully crafted sampling distribution and specific initialization, leads to a simple algorithm which outperforms more sophisticated and complex attacks in terms of success rate and query efficiency. In this paper we show that the idea of random search can be transferred to the case of sparse attacks, where the non-convex, combinatorial constraints are not easily handled even by gradient-based white-box attacks.

Algorithm 1: Sparse-RS

input : loss L , x_{orig} , max query N , sparsity k , input space \mathcal{S} **output** : minimizer of L

```
1  $M \leftarrow k$  indices of elements to be perturbed ; // initialize  $M$ 
2  $\Delta \leftarrow$  values of the perturbation to be applied ; // initialize  $\Delta$ 
3  $z \leftarrow x_{\text{orig}}$ ,  $z_M \leftarrow \Delta$  ; // set elements in  $M$  to values in  $\Delta$ 
4  $L^* \leftarrow L(z)$ ,  $i \leftarrow 0$  ; // initial value of the loss
5 while  $i < N$  and success not achieved do
6    $M' \leftarrow$  randomly sampled modification of  $M$  ; // new set of indices
7    $\Delta' \leftarrow$  randomly sampled modification of  $\Delta$  ; // new perturbation
8    $z \leftarrow x_{\text{orig}}$ ,  $z_{M'} \leftarrow \Delta'$  ; // create new candidate in  $\mathcal{S}$ 
9   if  $L(z) < L^*$  then
10    |  $L^* \leftarrow L(z)$ ,  $M \leftarrow M'$ ,  $\Delta \leftarrow \Delta'$  ; // if loss improves, update sets
11    |  $i \leftarrow i + 1$ 
12  $z \leftarrow x_{\text{orig}}$ ,  $z_M \leftarrow \Delta$  ; // return best solution found
13 return  $z$ 
```

3 Sparse-RS: a random search based framework for adversarial attacks

Random search (RS) is a well known scheme for derivative free optimization [43]. Given an objective function L to minimize, a starting point $x^{(0)}$ and a sampling distribution \mathcal{D} , RS iterations are given by

$$\delta \sim \mathcal{D}(x^{(i)}), \quad x^{(i+1)} = \arg \min_{y \in \{x^{(i)}, x^{(i)} + \delta\}} L(y). \quad (4)$$

This means that at every step an update of the current iterate $x^{(i)}$ is sampled according to \mathcal{D} and accepted only if it attains a lower value of the objective function than $x^{(i)}$, otherwise the procedure is repeated. Although not explicitly mentioned in Eq. (4), constraints on the space of the acceptable points can be integrated ensuring that δ is sampled so that $x^{(i)} + \delta$ is a feasible solution. Thus even complex constraints e.g. combinatorial ones can easily be integrated as RS just needs to be able to produce feasible points in contrast to gradient-based methods which depend on a continuous set to optimize over (or at least an embedding of the discrete set into a continuous one). While simple and flexible, RS is an effective tool in many tasks [55, 4], with the key ingredient for its success being a proper sampling distribution \mathcal{D} to guide the exploration of the space of possible solutions.

We summarize our general framework based on random search to generate sparse adversarial attacks, Sparse-RS, in Alg. 1, where the sparsity k indicates the maximum number of features that can be perturbed. Note that our algorithm does not try to minimize the amount of modified dimensions but rather exploits the whole budget k . In Alg. 1, since sparse attacks can be characterized by a pair (M, Δ) , where M is the set of components to be perturbed and Δ are the values to be inserted at M to form the adversarial input, also the sampling distribution \mathcal{D} consists of a two-stage process. First, we sample a random update of the locations M of the current perturbation (step 6) and, second, a random update of its values Δ (step 7). In some threat models (e.g. adversarial frames) the set M cannot be changed, so $M' \equiv M$ at every step. How Δ' is generated depends on the specific case, so we present the individual procedures in the next sections.

Common to all variations of Sparse-RS is that the whole budget for the perturbations is fully exploited both in terms of number of modified components and magnitude of the elementwise changes (constrained only by the limits of the input domain \mathcal{S}). This follows the intuition that larger perturbations should lead faster to an adversarial example. Moreover, the difference of the candidates M' and Δ' with M and Δ shrinks gradually with the iterations. The rationale behind this is to mimic the reduction of the step size in gradient-based optimization: initial large steps allow to quickly decrease the objective loss, but smaller steps are necessary to refine a close-to-optimal solution at the end of the algorithm. Finally, we impose a limit N on the maximum number of queries of the classifier, i.e. evaluations of the objective function.

As objective function L to be minimized, we use in the case of untargeted attacks the margin loss $L_{\text{margin}}(f(\cdot), y) = f_y(\cdot) - \max_{r \neq y} f_r(\cdot)$, where y is the correct class, so that $L < 0$ is equivalent to misclassification, whereas for targeted attacks we use the cross-entropy loss L_{CE} of the target class t , namely $L_{\text{CE}}(f(\cdot), t) = -f_t(\cdot) + \log \left(\sum_{r=1}^K e^{f_r(\cdot)} \right)$.

The code of the Sparse-RS framework is available at <https://github.com/fra31/sparse-rs>.

4 l_0 -bounded attacks

The first scenario we consider is that of l_0 -bounded adversarial examples, which means that only k pixels of a clean input $x_{\text{orig}} \in [0, 1]^{h \times w \times c}$ (width w , height h , color c) can be modified, but there are no constraints on the magnitude of the perturbations except for those of the input domain.

l_0 -RS algorithm: Let U be the set of the $h \cdot w$ pixels. In this case the set $M \subset U$ is initialized sampling uniformly k elements of U , while $\Delta \sim \mathcal{U}(\{0, 1\}^{k \times c})$, that is random values in $\{0, 1\}$ (every perturbed pixel gets one of the corners of the color cube $[0, 1]^c$). Then, at the i -th iteration, we randomly select $A \subset M$ and $B \subset U \setminus M$, with $|A| = |B| = \alpha^{(i)} \cdot k$, and create $M' = (M \setminus A) \cup B$. Δ' is formed by sampling random values from $\{0, 1\}^c$ for the elements in B , i.e. those which were not perturbed at the previous iteration. The quantity $\alpha^{(i)}$ controls how much M' differs from M and decays following a predetermined piecewise constant schedule rescaled according to the maximum number of queries N . The schedule is completely determined by the single value α_{init} , used to calculate $\alpha^{(i)}$ for every iteration i , which is also the only free hyperparameter of our scheme (details about the algorithm, schedule and values of α_{init} in App. A and B).

4.1 Theoretical analysis of l_0 -RS

We here analyze l_0 -RS on a binary classifier. While the analysis does not directly transfer to neural networks, most modern neural network architectures result in piecewise linear classifiers [5, 16], so that the result should approximately hold for a sufficiently small neighborhood of the target point x .

As in the malware detection task from Sec. 4.4, we assume that the input x has binary features, $x \in \{0, 1\}^d$, and we denote the label by $y \in \{-1, 1\}$ and the gradient of the linear model by $w_x \in \mathbb{R}^d$. Then the problem (3) of finding the optimal l_0 adversarial example is equivalent to:

$$\arg \min_{\substack{\|\delta\|_0 \leq k \\ x_i + \delta_i \in \{0, 1\}}} y \langle w_x, x + \delta \rangle = \arg \min_{\substack{\|\delta\|_0 \leq k \\ \delta_i \in \{0, 1 - 2x_i\}}} \langle y w_x, \delta \rangle = \arg \min_{\substack{\|\delta\|_0 \leq k \\ \delta_i \in \{0, 1\}}} \underbrace{\langle y w_x \odot (1 - 2x), \delta \rangle}_{\hat{w}_x}. \quad (5)$$

In the white-box case, i.e. when w_x is known, the solution is to simply set $\delta_i = 1$ for the k smallest weights of \hat{w}_x . The black-box case, where w_x is unknown and we are only allowed to query the model predictions $\langle \hat{w}_x, z \rangle$ for any $z \in \mathbb{R}^d$, is more complicated since the naive weight estimation algorithm requires $O(d)$ queries to first estimate \hat{w}_x and then to perform the attack by selecting the k minimal weights. This naive approach is prohibitively expensive for high-dimensional datasets (e.g., $d = 150,528$ on ImageNet assuming $224 \times 224 \times 3$ images). However, the problem of generating adversarial examples does not have to be always solved exactly, and often it is enough to find an approximate solution. Therefore we can be satisfied with only identifying k among the m smallest weights. Indeed, the focus is not on exactly identifying the solution but rather on having an algorithm that in expectation requires a *sublinear* number of queries. With this goal, we show that l_0 -RS satisfies this requirement for large enough m .

Proposition 4.1. *The expected number t_k of queries needed for l_0 -RS with $\alpha^{(i)} = 1/k$ to find a set of k weights out of the smallest m weights of a linear model is:*

$$\mathbb{E}[t_k] = (d - k)k \sum_{i=0}^{k-1} \frac{1}{(k - i)(m - i)} < (d - k)k \frac{\ln(k) + 2}{m - k}. \quad (6)$$

The proof is deferred to App. E and resembles that of the coupon collector problem [20]. In practice on non-linear models, l_0 -RS uses $\alpha^{(i)} > 1/k$ for better exploration initially, but then progressively reduces it. The main conclusion from Proposition 4.1 is that $\mathbb{E}[t_k]$ indeed becomes sublinear for large enough gap $m - k$, as illustrated in Fig. 1, where we plot the performance of l_0 -RS for $d = 150,528$ and $k = 150$ which is equivalent to our ImageNet experiments below with 50 pixels perturbed.




Figure 1: Comparison of the query efficiency of l_0 -RS to the performance of the naive weight estimation baseline.




Figure 2: Progression success rate vs number of queries for black-box untargeted l_0 -attacks (solid lines) on ImageNet. The dashed lines indicate the success rate of white-box attacks. At all sparsity levels l_0 -RS (red) outperforms PGD_0 with gradient estimation (blue), achieving higher success rate than SparseFool [36] (green) and better or similar results compared to white-box PGD_0 [17] (black).

4.2 ImageNet experiments

Experimental setup: ImageNet dataset contains RGB images resized to shape 224×224 , that is 50176 pixels. We test the robustness of VGG-16-BN and ResNet-50 from the PyTorch pretrained models on 500 images of the validation set. We modify $k \in \{50, 100, 150\}$ pixels to assess the effectiveness of the untargeted attacks at different thresholds (for black-box attacks, with a limit of 10,000 queries), while for targeted attacks we use $k \in \{150, 300\}$ (query limit 100,000).

Competitors: The existing black-box l_0 -attacks [38, 46, 17] do not aim at query efficiency but rather try to minimize the size of the perturbations. Among them, only CornerSearch [17] scales to ImageNet-size, but it requires $8 \times \#pixels$ queries only for the initial phase, exceeding the query limit we fix. Among white-box attacks, we compare to SparseFool [36] and PGD_0 [17], since [36, 17] showed them to be more effective than [39, 13] and we observe in our tests that they outperform also [41]. Finally, we introduce a black-box version of PGD_0 where the gradient is estimated by finite difference approximation as done in e.g. [26] (details about the attacks in App. A).

Results: For untargeted attacks, we show in Fig. 2 the progression of the success rate (on the initially correctly classified points) vs the number of queries used by the black-box attacks and (dashed lines) the final success rate of the white-box ones. In all cases l_0 -RS outperforms the competitors, including the white-box attacks, and needs only a limited amount of queries to get high fooling rate, e.g. on VGG with $k = 150$ it achieves 100% of success rate using on average 171 queries, with a median of 25 (the complete statistics about query efficiency is in App. A).

Table 1 shows the results of targeted attacks: l_0 -RS achieves higher success rate than the white-box PGD_0 of [17], especially for smaller k (we do not compare to the other methods as PGD_0 is the strongest competitor in the easier untargeted scenario).

Table 1: Success rate of targeted l_0 -attacks on VGG and ResNet modifying $k \in \{150, 300\}$ pixels. Note that PGD_0 is a white-box attack.

attack	$k = 150$		$k = 300$	
	VGG	ResNet	VGG	ResNet
PGD_0 [17]	55.6%	51.0%	95.4%	95.4%
l_0 -RS	98.2%	95.6%	99.6%	99.6%

4.3 MNIST robust models

We evaluate l_0 -RS on robust models on MNIST, which contains greyscale images $x \in [0, 1]^{28 \times 28 \times 1}$. [46] propose two models, ABS and Binary ABS, reported to have robust accuracy on 100 points of 69% and 77% respectively against l_0 -attacks with sparsity $k = 12$ (robust accuracy is defined as the classification error on the adversarial examples crafted by an attacker). These results are obtained by the Pointwise Attack with 10 restarts [46]. We run l_0 -RS on both ABS models in the same setup as in [46], using 10,000 queries and 5 random restarts, and achieve a robust accuracy of 45% for ABS and 49% on Binary ABS, which is a significant improvement over the Pointwise Attack in [46].




Figure 3: We plot success rate vs number of queries for JSMA with gradient estimation (blue) and our l_0 -RS (red), when used to fool a malware detector on Drebin [6]. For every sparsity level $k \in \{20, 30, 40\}$ l_0 -RS outperforms the competitor, with a large gap in the low query regime.

4.4 Malware detection

While in image classification the input space is continuous, other tasks are characterized by discrete features, e.g. in malware detection. We consider the Drebin dataset [6] consisting of 129,013 benign and malicious Android applications, with $d = 545,333$ features. A data point is encoded as a binary vector $x \in \{0, 1\}^d$ indicating whether each feature is present or not in x . Thus, we apply l_0 -RS as described at the beginning of the section for images as input by setting $h = d$ and $w = c = 1$ such that $x \in \{0, 1\}^{d \times 1 \times 1}$ (details in App. B). Then, we can modify malware applications with a limited number of insertions, without access to the internal weights of the classifier, to get them not detected.

Experimental setup: Following [21], we restrict the attacks to only adding features which preserve the functionality of the application (no features are removed). We trained a fully connected model which achieves a test accuracy of 98.85% and test its robustness on 500 malicious inputs from the test set, modifying $k \in \{20, 30, 40\}$ features, within a maximum of 10,000 queries.

Competitors: [21] successfully fooled similar models with a variant of the white-box JSMA [39], and [40] confirms that it is the most effective technique on Drebin, compared to the attacks of [48, 1, 24] including adaptation of FGSM [49] and PGD [30, 34]. We use JSMA in a black-box version with gradient estimation (details in App. B). [33] propose a black-box genetic algorithm with prior knowledge of the importance of the features for misclassification (via a pretrained random forest) which is not comparable to our threat model.

Results: Fig. 3 shows the progression of success rate (computed on the initially correctly classified points) of the attacks over number of queries used. At all three sparsity levels, our l_0 -RS attack (in red) outperforms JSMA with gradient estimation (in blue), especially in the low query regime, demonstrating the effectiveness of our framework beyond image classification.

5 Adversarial patches and frames

Adversarial patches [11] and frames [56] are sparse perturbations with a predefined structure. The most challenging and interesting type are *universal targeted attacks* where the classifier predicts a desired class regardless of the image (and location of the perturbation). Thus the goal is to produce a *single* perturbation which effectively deceives the classifier even when applied to *new* images not seen during its generation (see examples in Fig. 4). In Sec. 5.1 we show that Sparse-RS is able to craft such attacks only querying the target classifier and using a minimal amount of training images, while we defer the discussion of the untargeted case to App. C. We are not aware of any black-box method for this difficult task which is not based on transfer attacks. For the simpler task of image-specific adversarial patches (Sec. 5.2) we compare Patch-RS to a recent black-box attack [54].

5.1 Universal targeted attacks: patches and frames

Patches: The first type of universal perturbation we consider are squared patches, which should ideally fool a classifier independently from where they are applied. To achieve this, given a batch of $n = 100$ training images to create the patch, we copy it 2 times, so to have 300 images. Then we randomly select for each of them a location where the perturbation will be applied. These locations are then kept fixed throughout the iterations of the random search, i.e. the set M in the general Alg. 1

Table 2: We report the success rate of universal targeted attacks for 6 target classes when applied to unseen images. We compare the performance of our Sparse-RS, with a budget of either 20,000 or 100,000 queries for creating the perturbation, to that of the transfer-based PGD-attack.

		maypole	mailbox	peacock	traffic light	digital clock	wardrobe	mean
patches	Tr-PGD	1.1%	3.3%	1.7%	29.7%	0.6%	0.3%	6.1%
	RS _{20k}	0.6%	25.8%	5.0%	78.6%	79.0%	0.3%	31.6%
	RS _{100k}	25.5%	69.1%	51.5%	95.0%	93.9%	2.4%	56.2%
frames	Tr-PGD	30.4%	0.1%	9.3%	5.3%	0.0%	0.0%	7.5%
	RS _{20k}	85.9%	6.3%	74.0%	82.5%	8.0%	0.0%	42.8%
	RS _{100k}	93.9%	26.6%	88.2%	92.4%	26.5%	0.1%	54.6%

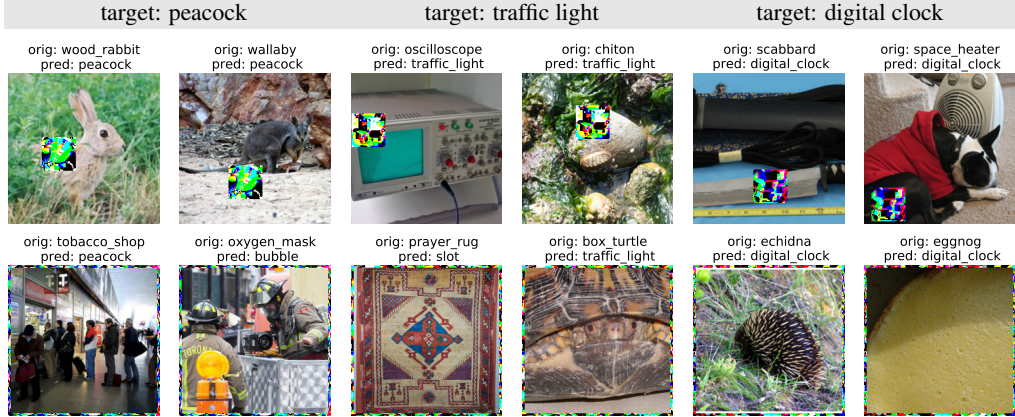


Figure 4: Examples of universal targeted patches (top) and frames (bottom) generated by Sparse-RS applied to random images (patches applied at random location) for 3 target classes (indicated on top): above each image we report the true class and the prediction after the application of the perturbation.

does not get updated (step 6 does not modify M). The patch is initialized and updated with the l_∞ Square Attack [4] where the only constraints are given by the box $[0, 1]^d$, so that all entries of the patch are in $\{0, 1\}$. Denoting $\{x_i, y_i\}_{i=1}^{3n}$ the images and correct labels in the expanded batch, t a target class and f the classifier, our algorithm minimizes the loss

$$L_{\text{targ}}(\{x_i, y_i\}_{i=1}^{3n}) = \sum_{i=1}^{3n} L_{\text{CE}}(f(x_i), t), \quad (7)$$

and a candidate update is only accepted if it reduces the loss without increasing the number of points where the attack is not successful.

Frames: The second type of universal perturbations consists in adversarial frames, introduced in [56], which modify only the pixels on the borders of the image: if the frame has width w , then all the pixels in the first/last w rows/columns can be perturbed (thus the set M of the modified pixels is fixed across iterations). For our attack **Frames-RS** we initialize the frame with random values in $\{0, 1\}$ and the proposed updates are built sampling $\alpha^{(i)} \cdot k$ (k is the total of pixels perturbed) random squares of size $w \times w$ (partially or totally) overlapping with the frame and changing all the corresponding pixels (intersecting the frame) to the same randomly chosen corner of the color cube $[0, 1]^3$. As usual $\alpha^{(i)}$ decreases with iterations and we use the same loss and acceptance criterion as for patches (see Eq. (7), where now the batch is not copied, so $i = 1, \dots, n$).

Competitors: As we are not aware of other black-box methods for this scenario not based on transfer attacks, we compare our universal perturbations via Sparse-RS to those obtained with white-box PGD on a surrogate model and then transferred to the victim model. At each iteration of PGD we sample a random location for each image in the training batch, apply the current patch and compute the gradient of the cross-entropy loss wrt each input. Then we sum the gradients at the patch locations, update the patch with a step in the direction of the sign of the overall gradient and clip it to $[0, 1]^d$. For frames we use the same procedure without sampling the location. We use the sign of the gradient

as the only constraints on the perturbations are the componentwise limits of the input domain $[0, 1]^d$. We note that [29] and [56] used similar techniques, but tested the resulting perturbations on the same model used at training time, while [11] used 4 surrogate models to craft an adversarial patch and then transferred it to the victim model. We provide only one surrogate classifier.

Results: We test the attacks on VGG on ImageNet using for both the same 100 training images. For the transfer-based PGD attack we craft the universal perturbations on ResNet, while Sparse-RS attacks directly VGG (as it is black-box). To evaluate *universality* of the patches (size 50×50), we compute the success rate of the patch applied to 500 unseen images at 1000 random locations each. Similarly, we compute the success rate of adversarial frames ($w = 4$) on 5000 images. We report the results in Table 2 for 6 target classes, 3 hand picked among classes commonly predicted by untargeted attacks and 3 randomly chosen. Sparse-RS outperforms PGD significantly, even when only 20,000 queries are used for training instead of the standard 100,000 (this is the number of queries usually given to non-universal targeted attacks in the l_∞ - and l_2 -threat models, see e.g. [47, 35, 4]). Note that PGD achieves success rate close to 100% for almost all the classes (the target class "wardrobe" is an exception) when evaluated on unseen images on ResNet, which shows that its perturbations are model-specific and do not transfer to VGG. The low success rate of the transfer attack is in line with [11] who reported for the class "toaster" a success rate of $\approx 6\%$ for 5% changed pixels. Thus it is important to have a black-box attack like Sparse-RS which does not rely on a surrogate model. The success rate of Sparse-RS significantly varies with the target class, from 0.1% of "wardrobe" to 95.0% of "traffic light", and some of the classes are better suited for patches or frames.

5.2 Image-specific untargeted adversarial patches

Attacks: A black-box (not transfer-based) method to produce image- and location-specific patches (both perturbation and position are optimized for each input independently) has been recently introduced in [54]. While [54] allows multiple patches on an image, we use their attack, TPA, in the standard setting of a single patch. [42] has proposed a PGD-based white-box attack for this scenario which we combine with the gradient estimation technique of [27]. For Patch-RS we alternate iterations where updates of the patch are sampled (step 7 in Alg. 1) and where a new position is tested (step 6 in Alg. 1). Please see App. D for details about our algorithm and other attacks.

Results: In Table 3 we report success rate, mean and median number of queries used (considering either only successful samples or all) by untargeted attacks with patches of size 30×30 on 500 images of ImageNet with query limit of 10,000 (VGG and ResNet as target models). Patch-RS achieves the best results. We note that the metrics only on the successful points are highly biased by the success rate (other attacks likely fail on the difficult points where Patch-RS succeeds after many queries). Finally, Fig. 5 shows some resulting adversarial examples of Patch-RS: random search is able to optimize the position of the patch so to cover semantic features of the original subject.

Table 3: Success rate and query statistics of image-specific 30×30 patches for VGG and ResNet. The query statistics are computed once only on the successful points and once on all points (the query limit is 10,000). Patch-RS outperforms TPA both in success rate and query efficiency.

attack	success rate		successful points				all points			
	VGG	ResNet	average queries		median queries		average queries		median queries	
			VGG	ResNet	VGG	ResNet	VGG	ResNet	VGG	ResNet
PGD w/ GE	46.7%	37.3%	522	595	30	30	5574	6494	10000	10000
TPA [54]	60.2%	59.9%	1760	1830	1200	1200	4536	4757	1600	2000
Patch-RS	95.8%	90.2%	886	996	184	162	1270	1875	201	230



Figure 5: Image and location-specific untargeted patches of size 30×30 generated by our Patch-RS.

Broader Impact

In this paper we discuss different variants of sparse black-box adversarial attacks where universal adversarial patches/frames have the potential to be used as a physical attack.

On the positive side powerful adversarial attacks are necessary to test security and safety of machine learning systems e.g. vulnerability to l_0 -attacks is a sign that an image recognition system might have unexpected behavior in case of pixel failures. Moreover, in particular (universal) adversarial patches/frames challenge our understanding what current neural networks really learn when they classify an image. In particular adversarial frames do not change at all the central part of the image content so it is quite disturbing that current architectures fail at all. Thus these kind of attacks should trigger the development of new architectures which are inherently robust to such changes.

On the negative side in particular black-box attacks have the potential to be used by a real attacker on some machine learning system which could cause potential harm to people. However, we think that it is impossible to construct secure and safe systems without discussing possibilities to challenge their robustness.

Acknowledgements

We thank Yang et al. [54] for quickly releasing their code and answering our questions. F.C., N.S. and M.H. acknowledge support from the German Federal Ministry of Education and Research (BMBWF) through the Tübingen AI Center (FKZ: 01IS18039A), the DFG Cluster of Excellence “Machine Learning – New Perspectives for Science”, EXC 2064/1, project number 390727645, and by DFG grant 389792660 as part of TRR 248.

References

- [1] A. Al-Dujaili, A. Huang, E. Hemberg, and U.-M. O’Reilly. Adversarial deep learning for robust detection of binary encoded malware. In *IEEE Security and Privacy Workshops (SPW)*, pages 76–82, 2018.
- [2] A. Al-Dujaili and U.-M. O’Reilly. There are no bit parts for sign bits in black-box attacks. In *ICLR*, 2020.
- [3] M. Alzantot, Y. Sharma, S. Chakraborty, and M. Srivastava. Genattack: practical black-box attacks with gradient-free optimization. *Genetic and Evolutionary Computation Conference (GECCO)*, 2019.
- [4] M. Andriushchenko, F. Croce, N. Flammarion, and M. Hein. Square attack: a query-efficient black-box adversarial attack via random search. arXiv preprint arXiv:1912.00049, 2019.
- [5] R. Arora, A. Basuy, P. Mianjyz, and A. Mukherjee. Understanding deep neural networks with rectified linear unit. In *ICLR*, 2018.
- [6] D. Arp, M. Spreitzenbarth, M. Hubner, H. Gascon, K. Rieck, and C. Siemens. Drebin: Effective and explainable detection of android malware in your pocket. In *NDSS*, volume 14, pages 23–26, 2014.
- [7] A. Athalye, N. Carlini, and D. A. Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *ICML*, 2018.
- [8] A. N. Bhagoji, W. He, B. Li, and D. Song. Practical black-box attacks on deep neural networks using efficient query mechanisms. In *ECCV*, 2018.
- [9] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Srdic, P. Laskov, G. Giacinto, and F. Roli. Evasion attacks against machine learning at test time. In *ECML/PKDD*, 2013.
- [10] W. Brendel, J. Rauber, and M. Bethge. Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. In *ICLR*, 2018.
- [11] T. B. Brown, D. Mané, A. Roy, M. Abadi, and J. Gilmer. Adversarial patch. In *NeurIPS 2017 Workshop on Machine Learning and Computer Security*, 2017.
- [12] T. Brunner, F. Diehl, M. T. Le, and A. Knoll. Guessing smart: biased sampling for efficient black-box adversarial attacks. *ICCV*, 2019.

- [13] N. Carlini and D. Wagner. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy*, 2017.
- [14] P. Chen, Y. Sharma, H. Zhang, J. Yi, and C. Hsieh. Ead: Elastic-net attacks to deep neural networks via adversarial examples. In *AAAI*, 2018.
- [15] S. Cheng, Y. Dong, T. Pang, H. Su, and J. Zhu. Improving black-box adversarial attacks with a transfer-based prior. In *NeurIPS*, 2019.
- [16] F. Croce, M. Andriushchenko, and M. Hein. Provable robustness of relu networks via maximization of linear regions. In *AISTATS*, 2019.
- [17] F. Croce and M. Hein. Sparse and imperceptible adversarial attacks. In *ICCV*, 2019.
- [18] F. Croce and M. Hein. Minimally distorted adversarial examples with a fast adaptive boundary attack. In *ICML*, 2020.
- [19] A. Fawzi and P. Frossard. Measuring the effect of nuisance variables on classifiers. In *British Machine Vision Conference (BMVC)*, 2016.
- [20] P. Flajolet, D. Gardy, and L. Thimonier. Birthday paradox, coupon collectors, caching algorithms and self-organizing search. *Discrete Applied Mathematics*, 39(3):207–229, 1992.
- [21] K. Grosse, N. Papernot, P. Manoharan, M. Backes, and P. McDaniel. Adversarial perturbations against deep neural networks for malware classification. *arXiv preprint arXiv:1606.04435*, 2016.
- [22] C. Guo, J. R. Gardner, Y. You, A. G. Wilson, and K. Q. Weinberger. Simple black-box adversarial attacks. In *ICML*, 2019.
- [23] J. E. Hu, A. Swaminathan, H. Salman, and G. Yang. Improved image wasserstein attacks and defenses. *arXiv preprint arXiv:2004.12478*, 2020.
- [24] W. Hu and Y. Tan. Generating adversarial malware examples for black-box attacks based on gan. *arXiv preprint arXiv:1702.05983*, 2017.
- [25] Z. Huang and T. Zhang. Black-box adversarial attack with transferable model-based embedding. In *ICLR*, 2020.
- [26] A. Ilyas, L. Engstrom, A. Athalye, and J. Lin. Black-box adversarial attacks with limited queries and information. *ICML*, 2018.
- [27] A. Ilyas, L. Engstrom, and A. Madry. Prior convictions: Black-box adversarial attacks with bandits and priors. *ICLR*, 2019.
- [28] D. Jin, Z. Jin, J. T. Zhou, and P. Szolovits. Is bert really robust? natural language attack on text classification and entailment. *arXiv preprint arXiv:1907.11932*, 2019.
- [29] D. Karmon, D. Zoran, and Y. Goldberg. Lavan: Localized and visible adversarial noise. In *ICML*, 2018.
- [30] A. Kurakin, I. J. Goodfellow, and S. Bengio. Adversarial examples in the physical world. In *ICLR Workshop*, 2017.
- [31] M. Lee and Z. Kolter. On physical adversarial patches for object detection. *ICML Workshop on Security and Privacy of Machine Learning*, 2019.
- [32] J. Li, F. Schmidt, and Z. Kolter. Adversarial camera stickers: A physical camera-based attack on deep learning systems. In *ICML*, pages 3896–3904, 2019.
- [33] X. Liu, X. Du, X. Zhang, Q. Zhu, H. Wang, and M. Guizani. Adversarial samples on android malware detection systems for IoT systems. *Sensors*, 19(4):974, 2019.
- [34] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Valdu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018.
- [35] L. Meunier, J. Atif, and O. Teytaud. Yet another but more efficient black-box adversarial attack: tiling and evolution strategies. *arXiv preprint, arXiv:1910.02244*, 2019.
- [36] A. Modas, S. Moosavi-Dezfooli, and P. Frossard. Sparsefool: a few pixels make a big difference. In *CVPR*, 2019.
- [37] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *CVPR*, pages 2574–2582, 2016.

- [38] N. Narodytska and S. Kasiviswanathan. Simple black-box adversarial attacks on deep neural networks. In *CVPR Workshops*, 2017.
- [39] N. Papernot, P. McDaniel, S. Jha, M. Fredrikson, Z. B. Celik, and A. Swami. The limitations of deep learning in adversarial settings. In *2016 IEEE European symposium on security and privacy (EuroS&P)*, pages 372–387. IEEE, 2016.
- [40] R. Podschwadt and H. Takabi. On effectiveness of adversarial examples and defenses for malware classification. In *International Conference on Security and Privacy in Communication Systems*, pages 380–393. Springer, 2019.
- [41] A.-A. Pooladian, C. Finlay, T. Hoheisel, and A. M. Oberman. A principled approach for generating adversarial images under non-smooth dissimilarity metrics. *arXiv preprint arXiv:1908.01667*, 2019.
- [42] S. Rao, D. Stutz, and B. Schiele. Adversarial training against location-optimized adversarial patches. *arXiv preprint arXiv:2005.02313*, 2020.
- [43] L. Rastrigin. The convergence of the random search method in the extremal control of a many parameter system. *Automaton & Remote Control*, 24:1337–1342, 1963.
- [44] J. Rauber, W. Brendel, and M. Bethge. Foolbox: A python toolbox to benchmark the robustness of machine learning models. In *ICML Reliable Machine Learning in the Wild Workshop*, 2017.
- [45] J. Rony, L. G. Hafemann, L. S. Oliveira, I. B. Ayed, R. Sabourin, and E. Granger. Decoupling direction and norm for efficient gradient-based l2 adversarial attacks and defenses. In *CVPR*, pages 4322–4330, 2019.
- [46] L. Schott, J. Rauber, M. Bethge, and W. Brendel. Towards the first adversarially robust neural network model on MNIST. In *ICLR*, 2019.
- [47] M. Seungyong, A. Gaon, and O. S. Hyun. Parsimonious black-box adversarial attacks via efficient combinatorial optimization. In *ICML*, 2019.
- [48] J. W. Stokes, D. Wang, M. Marinescu, M. Marino, and B. Bussone. Attack and defense of dynamic analysis-based, adversarial neural malware classification models. *arXiv preprint arXiv:1712.05919*, 2017.
- [49] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. In *ICLR*, pages 2503–2511, 2014.
- [50] S. Thys, W. Van Ranst, and T. Goedemé. Fooling automated surveillance cameras: adversarial patches to attack person detection. In *CVPR Workshops*, 2019.
- [51] C.-C. Tu, P. Ting, P.-Y. Chen, S. Liu, H. Zhang, J. Yi, C.-J. Hsieh, and S.-M. Cheng. Autozoom: Autoencoder-based zeroth order optimization method for attacking black-box neural networks. In *AAAI*, 2019.
- [52] J. Uesato, B. O’Donoghue, A. Van den Oord, and P. Kohli. Adversarial risk and the dangers of evaluating against weak attacks. In *ICML*, 2018.
- [53] E. Wong, F. R. Schmidt, and J. Z. Kolter. Wasserstein adversarial examples via projected sinkhorn iterations. In *ICML*, 2019.
- [54] C. Yang, A. Kortylewski, C. Xie, Y. Cao, and A. Yuille. Patchattack: A black-box texture-based attack with reinforcement learning. *arXiv preprint arXiv:2004.05682*, 2020.
- [55] Z. B. Zabinsky. Random search algorithms. *Wiley encyclopedia of operations research and management science*, 2010.
- [56] M. Zajac, K. Zołna, N. Rostamzadeh, and P. O. Pinheiro. Adversarial framing for image and video classification. In *AAAI*, pages 10077–10078, 2019.

A l_0 -bounded attacks: image classification

Details of l_0 -RS algorithm: As mentioned in Sec. 4, at iteration i the new set M' is formed modifying $\alpha^{(i)} \cdot k$ (rounded to the closest positive integer) elements of M containing the currently modified dimensions (see step 6 in Alg. 1). Inspired by the step-size reduction in gradient-based optimization methods, we progressively reduce $\alpha^{(i)}$. Assuming $N = 10,000$, the schedule of $\alpha^{(i)}$ is piecewise constant where the constant segments start at iterations $j \in \{0, 50, 200, 500, 1000, 2000, 4000, 6000, 8000\}$ with values $\alpha_{\text{init}}/\beta_j, \beta_j \in \{2, 4, 5, 6, 8, 10, 12, 15, 20\}$. For a different maximum number of queries N , the schedule is linearly rescaled accordingly. In practice, we use $\alpha_{\text{init}} = 0.3$ and $\alpha_{\text{init}} = 0.1$ for the untargeted and targeted scenario respectively.

Additional results: We here report more detailed statistics about the comparison of different attacks on the task of l_0 -bounded perturbations (see Sec. 4.2). In particular, Table 4 shows the success rate (on the initially correctly classified points out of 500 of the validation set of ImageNet) of both white- and black-box attacks, for the three sparsity level considered, with l_0 -RS outperforming the competitors. In Table 5 we report the statistics about the query efficiency of black-box untargeted and targeted attacks. We compute first the average and median number of queries necessary to generate a successful adversarial example (this means that for each attack only the successful points are considered in such statistics) and second the same metrics including all points (i.e. when an adversarial example is not found we use the query limit to calculate the statistics). We recall that we set query limits of 10,000 and 100,000 for untargeted and targeted scenario respectively, consistently with previous works e.g. [47, 35, 4]. Our attack is able to achieve high fooling rate in all threat models requiring a small computational effort, especially in the easier untargeted scenario.

Table 4: Success rate of l_0 -attacks on VGG-16-BN and ResNet-50 perturbing at most $k \in \{50, 100, 150\}$ pixels. For black-box attacks we set a maximum of 10,000 queries. Our Sparse-RS outperforms both white- and black-box attacks at every threshold.

attack	scenario	$k = 50$		$k = 100$		$k = 150$	
		VGG	ResNet	VGG	ResNet	VGG	ResNet
SparseFool [36]	white-box	27.2%	26.5%	41.7%	46.2%	52.0%	56.6%
PGD ₀ [17]	white-box	77.8%	72.8%	96.0%	95.1%	99.5%	99.2%
PGD ₀ w/ GE	black-box	49.1%	38.0%	58.6%	50.9%	68.1%	59.4%
l_0 -RS	black-box	97.6%	94.6%	99.7%	100%	100%	100%

Table 5: Success rate and query metrics of targeted l_0 -attacks on VGG-16-BN (V) and ResNet-50 (R) at sparsity levels k . Mean and median queries are calculated once only on successful samples, once on all points using the query limit for the unsuccessful ones.

k	attack	succ. rate %		successful points				all points				
		V	R	avg. queries	med. queries	avg. queries	med. queries	avg. queries	med. queries	V	R	
untarg.	50	PGD ₀ w/GE	49.1%	38.0%	958	691	29	24	5563	6458	10000	10000
		l_0 -RS	97.6%	94.6%	511	672	82	133	737	1176	88	150
	100	PGD ₀ w/GE	58.6%	50.9%	276	511	12	14	4304	5170	256	7794
		l_0 -RS	99.7%	100%	282	578	42	81	308	578	44	81
150	PGD ₀ w/GE	68.1%	59.4%	373	338	6	8	3447	4263	42	240	
	l_0 -RS	100%	100%	171	359	25	49	171	359	25	49	
targ.	150	l_0 -RS	98.2%	95.6%	8243	10534	4542	6529	9895	14470	4914	6960
	300	l_0 -RS	99.6%	99.6%	7472	10768	3574	5771	7842	11125	3676	5869

Details about the attacks: We use SparseFool [36] as implemented in Foolbox [44] and optimize the hyperparameter which controls the sparsity of the solutions (called λ , setting finally $\lambda = 1$) to achieve the best results. The implementation in Foolbox has other two hyperparameters: we use 60 steps and 10 subsamples (default values are 30 and 10), after checking that higher values do not lead to an improved performance but significantly increase the computational cost. For PGD₀ we use 500 iterations (5,000 for targeted attacks) and step size $0.05 \cdot d$ (with $d = 224 \times 224 \times 3$ the input dimension). PGD₀ with GE uses the same gradient step of PGD₀ but with step size $5 \cdot d$. To estimate

the gradient we use finite difference, similarly to [26], as shown in steps 4-7 of Alg. 2, with $x^{(i)}$, the current iterate, instead of x_{orig} (we use $\sigma = \sqrt{d}$ and $\eta = 0.01$, after optimizing them, which is in line with what suggested in [27] for this algorithm). We also tested that sampling more points to better approximate the gradient at each iteration leads to similar success rate with worse query consumption.

B l_0 -bounded attacks: malware detection

Dataset: The Drebin dataset [6] consists of 129,013 Android applications, among which 123,453 are benign and 5,560 malicious, with $d = 545,333$ features divided into 8 families. Data points are represented by a binary vector $x \in \{0, 1\}^d$ indicating whether each feature is present or not in x .

Experimental setup: As [21], we restrict the attacks to only adding features from the first 4 families, that is modifications to the manifest of the applications, to preserve the functionality of the samples (no feature present in the clean data is removed), which leaves a maximum of 233,727 alterable dimensions. Only adding features means that all values in Δ equal 1, thus $\Delta' \equiv \Delta$ at every iteration (step 7 in Alg. 1) and only the set M is updated. For our attack we use $\alpha_{\text{init}} = 1.6$ and the same schedule of $\alpha^{(i)}$ of l_0 -RS on image classification tasks (see Sec. A).

Model: We trained the classifier, which has 1 fully-connected hidden layer with 200 units and uses ReLU as activation function, with 20 epochs of SGD minimizing the cross-entropy loss, with learning rate of 0.1 reduced by a factor of 10 after 10 and 15 epochs. We use batches of size 2000, consisting of 50% of benign and 50% of malicious examples. For training we merged the training and validation sets of one of the splits provided by [6]. It achieves test accuracy of 98.85%, with false positive rate of 1.01%, false negative rate 4.29%.

JSMA with gradient estimation: The idea of the white-box attack of [21] is, given a sparsity level of k , to perturb iteratively the feature of the iterate $x^{(i)}$ which corresponds to the largest value (if positive) of $\nabla_x L_{\text{CE}}(f(x^{(i)}), y)$, with f the classifier, y the correct label and L_{CE} the cross-entropy loss, until either the maximum number of modifications are made or misclassification is achieved. With only approximate gradients this approach is not particularly effective. However, since $k \ll d$ and only additions can be made, we aim at estimating the gradient of the cross-entropy loss at x_{orig} and then set to 1 the k elements (originally 0) of x_{orig} with largest component in the approximated gradient. Chasing query-efficiency, every m iterations of gradient estimation through finite difference we check if an adversarial example has been found (we do not count these queries in the total for the query limit). Alg. 2 shows the procedure, and for the main experiments in Sec. 4.4 we set $m = 5$ and $\sigma = \eta = 1$ (we tested other values which achieved worse or similar performance).

Algorithm 2: JSMA w/ GE

input : $L, x_{\text{orig}}, N, k, m, \sigma, \eta$

output : z

```

1  $g \leftarrow \mathbf{0}, \quad q \leftarrow 2 \cdot m$ 
2 while  $q \leq N$  do
3   for  $i = 1, \dots, m$  do
4      $s \leftarrow \mathcal{N}(0, I) / \sigma$ 
5      $l_1 \leftarrow L(x_{\text{orig}} + \eta \cdot s)$ 
6      $l_2 \leftarrow L(x_{\text{orig}} - \eta \cdot s)$ 
7      $g \leftarrow g + (l_1 - l_2) / (2\eta) \cdot s$ 
8    $z \leftarrow x_{\text{orig}}$ 
9    $A \leftarrow$  indices of  $k$  largest positive
   components of  $g$  (if there are)
10   $z_A \leftarrow 1$ 
11   $q \leftarrow q + 2 \cdot m$ 
12  if  $z$  is adversarial then return;
```

C Universal attacks: patches and frames

Details of Patch-RS and Frame-RS algorithms: In these threat models, the value of $\alpha^{(i)}$ starts with $\alpha^{(0)} = \alpha_{\text{init}}$ and then is halved at iteration $j \in \{10, 50, 200, 500, 1000, 4000, 8000\}$ if the query limit is $N = 10,000$, otherwise the values of j are linearly rescaled according to the new N . We initialize the patch with 1 pixel wide vertical stripes, each of a random color in $\{0, 1\}^c$, as in [4]. Then, to update the content of the patch we use the same sampling distribution of the l_∞ Square Attack [4]: given a patch of size $s \times s$, a square with side of size $\sqrt{\alpha^{(i)}} \cdot s$ (subset of the patch) is sampled and all its pixels set to the same color of the 2^c in $\{0, 1\}^c$ (we use $\alpha_{\text{init}} = 0.3$). For frames of width w , gathering in total k pixels, $\alpha^{(i)} \cdot k$ indicates the number of squares of size $w \times w$ sampled at

iteration i and used to update the adversarial frame ($\alpha_{\text{init}} = 0.005$). For crafting the perturbations we use a batch of 100 images from the validation set, so that in the case of patches the expanded batch of 300 data points (see Sec. 5.1) fits into GPU memory (16 GB) and forward passes for all the images can be executed with a single run.

Untargeted attacks: In case of untargeted attacks the loss minimized to create the perturbations is

$$L_{\text{untarg}}(\{x_i, y_i\}_{i=1}^m) = - \sum_{i=1}^m L_{\text{CE}}(f(x_i), y_i), \quad (8)$$

with the same notation of Sec. 5.1 and m the total number of images used at training time (we use the same 100 images for all the attacks but for patches the batch is copied 2 times, so in the end $m = 300$). In fact, we found maximizing the cross-entropy loss of the correct class more effective than using the margin loss as for the untargeted attacks in the l_0 -bounded threat model. Moreover, we keep the acceptance rule used for the universal targeted setting. We report the comparison of the success rate achieved by white-box PGD through transfer attack and Sparse-RS in Table 6, computed as for Table 2 on VGG, showing that both Patch-RS and Frame-RS outperform PGD. Finally, we observe that our untargeted attacks, when cause a misclassification, lead the model to predict in the large majority of the cases the classes which are easiest for targeted perturbations, i.e. "traffic light" for patches and "maypole" for frames.

Additional visualizations: We show in Fig. 6 the universal patches for targeted and untargeted attacks by white-box PGD (top) and Patch-RS (bottom). Moreover, Figures 9 and 10 show examples of how universal targeted patches and frames generated with Sparse-RS are used and affect the classification of random images for the target classes used in Sec. 5.1.




Figure 6: Universal patches of size 50×50 crafted by PGD (top row) and RS (bottom row): above each patch we indicate the target class (or untargeted attacks).

Additional results and details: We evaluate the performance of patches and frames generated by PGD when used to modify unseen images to attack ResNet, that is the model on which the perturbations were generated. As shown in Table 7, when the adversarial attacks are transferred to unseen images but evaluated on the same classifier used at training time, PGD has very high success rate, with the exception of the class "wardrobe", which results being particularly hard to use as target (the statistics are computed as in Table 2 in Sec. 5.1). Finally, for generating the adversarial patches and frames we use 1000 iterations of PGD and step size of 0.05 in the direction of the sign of the gradient (as explained in Sec. 5.1).

Table 7: We report the success rate of the universal targeted and untargeted attacks given by PGD on ResNet, that is the same model used to generate patches and frames.

	maypole	mailbox	peacock	traffic light	digital clock	wardrobe	untargeted
patches	90.2%	93.3%	94.4%	97.5%	99.8%	6.3%	99.9%
frames	99.9%	92.9%	99.8%	99.8%	99.9%	71.9%	99.9%

Table 6: Success rate of untargeted universal patches and frames given by transfer-based PGD and Sparse-RS on unseen images from ImageNet on VGG.

	patches 50×50	frames $w = 4$
Tr-PGD	20.0%	46.8%
RS	95.4%	98.3%

D Image- and location-specific adversarial patches

Patch-RS algorithm: As mentioned in Sec. 5.2, in this scenario we optimize via random search the attacks for each image independently. In Patch-RS we alternate iterations where a candidate update of the patch is sampled with others where a new location is sampled. We have a location update every m iterations, with $m = 5 \cdot (1 + \lfloor i/1000 \rfloor)$ and i the current iteration. The values of pixels of the patch are updated according to the same sampling distribution used for universal patches. The position of the patch is updated with an uniformly sampled shift in $[-h^{(i)}, h^{(i)}]$ for each direction (plus clipping to the image size if necessary), where $h^{(i)}$ is linearly decreased from $h^{(0)} = 0.75 \cdot s_{\text{image}}$ to $h^{(N)} = 0$ (s_{image} indicates the side of the squared images). In this way, initially the patch can be easily moved on the image, while towards the final iterations it is kept (almost) fixed and only its values are optimized.

Competitors: The first competitor we consider is a black-box version of the white-box PGD-based attack of [42]. It updates the patch with a step in the direction of the sign of the gradient in order to maximize the cross-entropy function at the current iterate. Then it tests if shifting of s_{shift} pixels the patch in one of the four possible directions improves the loss. If so, the location is updated otherwise kept ([42] have a version where only one random shift is tested, but this leads to lower success rate in our experiments). We use the technique introduced by [27] in the context of l_∞ - and l_2 -bounded attacks to approximate the gradient via finite differences. This implies that every iteration costs 6 queries (2 for gradient estimation, 4 for the location update). We optimized the large number of hyperparameters to achieve the best success rate: for PGD of [42] we use step size of 0.3 and $s_{\text{shift}} = 10$, for the method of [27] we set "tile size" to 56, "exploration" to 100, "online learning rate" to 100, η for finite difference to 10. The second method we compare to is TPA [54], which is based on reinforcement learning and exploits a dictionary of patches. Note that in [54] TPA was used primarily putting multiple patches on the same image to achieve misclassification, while our threat model does not include this option. We obtained the optimal values for the hyperparameters ("rl batch" 400 and number of steps 25) via personal communication with the authors. TPA has a mechanism of early stopping, which means that it might happen that not the whole budget of queries is exploited even for unsuccessful points. Finally, [54] show that TPA significantly outperforms the methods of [19], which also generates image-specific patches, although optimizing the shape and the location but not the values of the perturbation. Thus we do not compare directly to [19] in our experiments.

E Theoretical analysis

Proposition 4.1. *The expected number t_k of queries needed for l_0 -RS with $\alpha^{(i)} = 1/k$ to find a set of k weights out of the smallest m weights of a linear model is:*

$$\mathbb{E}[t_k] = (d - k)k \sum_{i=0}^{k-1} \frac{1}{(k - i)(m - i)} < (d - k)k \frac{\ln(k) + 2}{m - k}. \quad (9)$$

Proof. According to the l_0 -RS algorithm, we have d features grouped in a set U and the goal is to find a set $M \subset U$ containing k elements among the m smallest elements of U . Since $\alpha^{(i)} = 1/k$, at every iteration we pick one element $p \in M$ to remove from M and one element $q \in U \setminus M$ to add to M . This results in a binary vector $z_{\text{new}} \in \{0, 1\}^d$ indicating which are the features in M . Then we query the black-box linear model to determine whether the loss at a new point z_{new} improves compared to the point on the previous iteration z_{current} , i.e. $L(z_{\text{new}}) < L(z_{\text{current}})$, which gives us information whether $q < p$.

If the current set M contains i elements which belong to the smallest m , the probability of increasing it to $i + 1$ elements with the next pick is equal to

$$\mathcal{P}[i \rightarrow i + 1] = \mathcal{P}[p \notin \text{smallest } m, q \in \text{smallest } m] = \frac{k - i}{k} \cdot \frac{m - i}{d - k}.$$

Then the expected number of iterations for the step $i \rightarrow i + 1$ is equal to

$$\mathbb{E}[t_{i+1} - t_i] = \frac{(d - k)k}{(k - i)(m - i)}$$

since all the draws are independent. Assuming that we begin with none of the smallest m elements in M , the expected number of iterations t_k needed to find a set of k weights out of the smallest m weights is given by

$$\mathbb{E}[t_k] = \sum_{i=0}^{k-1} \mathbb{E}[t_{i+1} - t_i] = \sum_{i=0}^{k-1} \frac{(d-k)k}{(k-i)(m-i)} = (d-k)k \sum_{i=0}^{k-1} \frac{1}{(k-i)(m-i)}. \quad (10)$$

Now assuming that $m > k$, we can write the summation as:

$$\begin{aligned} \sum_{i=0}^{k-1} \frac{1}{(k-i)(m-i)} &= \sum_{j=1}^k \frac{1}{j(j+m-k)} = \frac{1}{m-k} \sum_{j=1}^k \left(\frac{1}{j} - \frac{1}{j+m-k} \right) = \\ &= \frac{1}{m-k} (H_k - H_m + H_{m-k}), \end{aligned} \quad (11)$$

where H_k is the k -th harmonic number. Using the fact that $\ln(k) < H_k \leq \ln(k) + 1$, we have:

$$H_k - H_m + H_{m-k} < \ln(k) - \ln(m) + \ln(m-k) + 2 < \ln(k) + 2.$$

If we combine this result with Eq. (10) and Eq. (11), we obtain the desired upper bound on $\mathbb{E}[t_k]$:

$$\mathbb{E}[t_k] < (d-k)k \frac{\ln(k) + 2}{m-k}.$$

□

We further remark that for $m = k$ we get that $\mathbb{E}[t_k] = (d-k)k \sum_{j=1}^k \frac{1}{j^2} < \frac{\pi^2}{6} (d-k)k$, however we are interested in the setting when the gap $m - k$ is large enough so that $\mathbb{E}[t_k]$ becomes sublinear.

F Ablation studies

We here present a series of studies to better understand the robustness of our algorithm to different random seeds, how its performance changes depending on α_{init} and to justify the algorithmic choice of the piecewise decaying schedule for $\alpha^{(i)}$. We focus on l_0 -RS on ImageNet as in this context we can compare both success rate and query efficiency with different sparsity levels.

Different random seeds: First, we study how the performance of l_0 -RS varies when using different random seeds, which influence the stochastic component inherent in random search. In Table 8 we report mean and standard deviation over 10 runs (with different seeds) of success rate, average and median queries of l_0 -RS on VGG and ResNet with sparsity levels $k \in \{50, 100, 150\}$ (the same setup of Sec. 4.2). One can observe that the success rate is very stable in all the cases, and the statistics of query consumption consistent across different runs.

Table 8: Mean and standard deviation of the performance (success rate and query efficiency) of l_0 -RS repeated with 10 different random seeds.

model	k	success rate (%)	successful points		all points	
			avg. queries	med. queries	avg. queries	med. queries
VGG	50	97.4 ± 0.42	497 ± 28	80 ± 4	745 ± 43	86 ± 5
	100	99.8 ± 0.14	320 ± 20	42 ± 3	335 ± 17	43 ± 3
	150	100.0 ± 0.00	193 ± 16	26 ± 2	193 ± 16	26 ± 2
ResNet	50	94.6 ± 0.70	686 ± 47	135 ± 6	1187 ± 48	166 ± 11
	100	99.7 ± 0.15	544 ± 36	73 ± 5	571 ± 26	74 ± 5
	150	100.0 ± 0.00	365 ± 19	49 ± 3	365 ± 19	49 ± 3

Different values of α_{init} : In second place, we analyze the behaviour of our algorithm with different values of α_{init} , since it is the only free hyperparameter of Sparse-RS. Let us recall that it is used to regulate how much M' and Δ' differ from M and Δ respectively in steps 6-7 of Alg. 1 at each iteration. Fig. 7 shows the success rate and query usage (computed on the successful samples) of our untargeted l_0 -bounded attack on VGG at the usual three sparsity levels k for runs with

$\alpha_{\text{init}} \in \{0.01, 0.05, 0.1, 0.2, 0.3, 0.4, 0.6, 0.8, 1\}$ (for our experiments in Sec. 4.2 we use $\alpha_{\text{init}} = 0.3$). We observe that the success rate is similar (close to 100%) for all the values, with a slight degradation for the largest ones. In order to minimize the queries of the classifier, the runs with α_{init} between 0.1 and 0.4 are comparably good, with small differences in the tradeoff between average and median number of queries.

Constant vs decaying $\alpha^{(i)}$: In order to demonstrate the role of decaying the difference between the candidate updates M' and Δ' and the current iterates M and Δ over iterations (see steps 6-7 of Alg. 1) to achieve good performance, we run our attack with constant $\alpha^{(i)}$ schedule instead of the piecewise constant schedule with decreasing values. We fix $\alpha^{(i)} = c \in [0, 1]$ for every i so that for the whole algorithm M' and M differ in $\max\{c \cdot k, 1\}$ elements. In Fig. 8 we report the results achieved by $c \in \{0, 0.05, 0.15, 0.3, 0.5, 1\}$ on VGG at $k \in \{50, 100, 150\}$, together with the baseline (black dashed line) of the standard version of l_0 -RS. One can observe that small constant values c for $\alpha^{(i)} = c$ achieve good success rate but suffer in query efficiency, in particular computed regarding the median and for larger k , while the highest values of c lead to significantly worse success rate (note that average and median queries are computed only on the successful points) than the baseline. These results highlight how important it is to have an initial exploration phase, with a larger step size, and at a later stage a more local optimization.

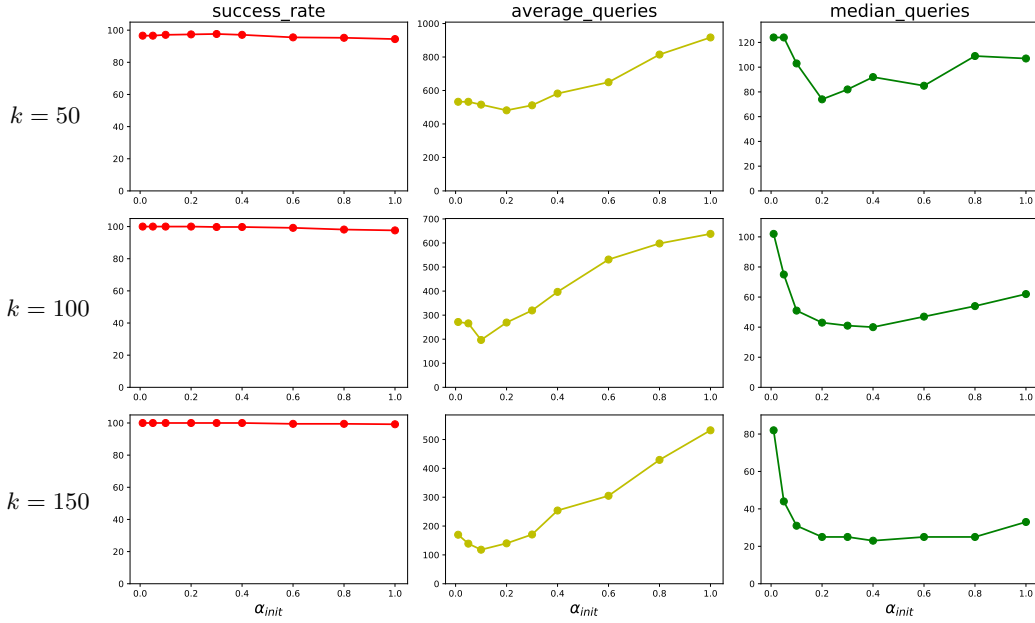


Figure 7: Ablation study on the influence of α_{init} , the hyperparameter which regulates the size of the updates at each iteration. We show success rate (first column), average (second) and median queries (third) achieved by l_0 -RS on VGG at sparsity levels $k = \{50, 100, 150\}$. Considering jointly the three statistics values in $[0.1, 0.4]$ are preferable for this threat model and schedule.

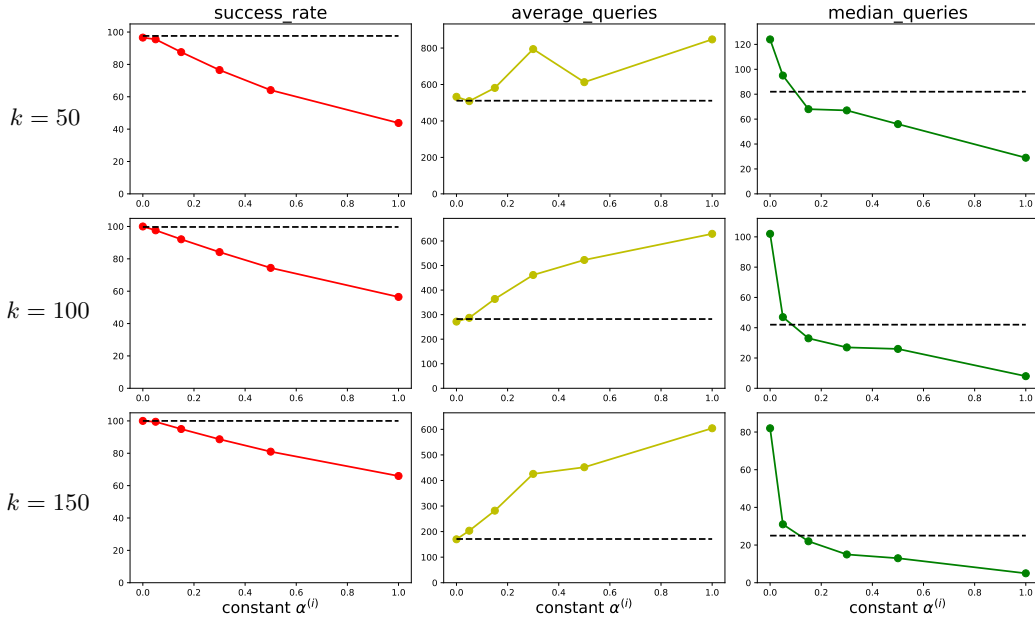


Figure 8: Performance of l_0 -RS on VGG when using a constant schedule for $\alpha^{(i)}$, that is the size of $|A| = |B| = c \cdot k$ (see Sec. 4.2) at every iteration, or equivalently $\alpha^{(i)} = c$ for every $i = 1, \dots, N$. The dashed line is the reference of the results achieved with the piecewise constant schedule to decay $\alpha^{(i)}$. While constantly small updates lead to good success rate, the median number of queries used increases notably, especially with larger sparsity levels k .

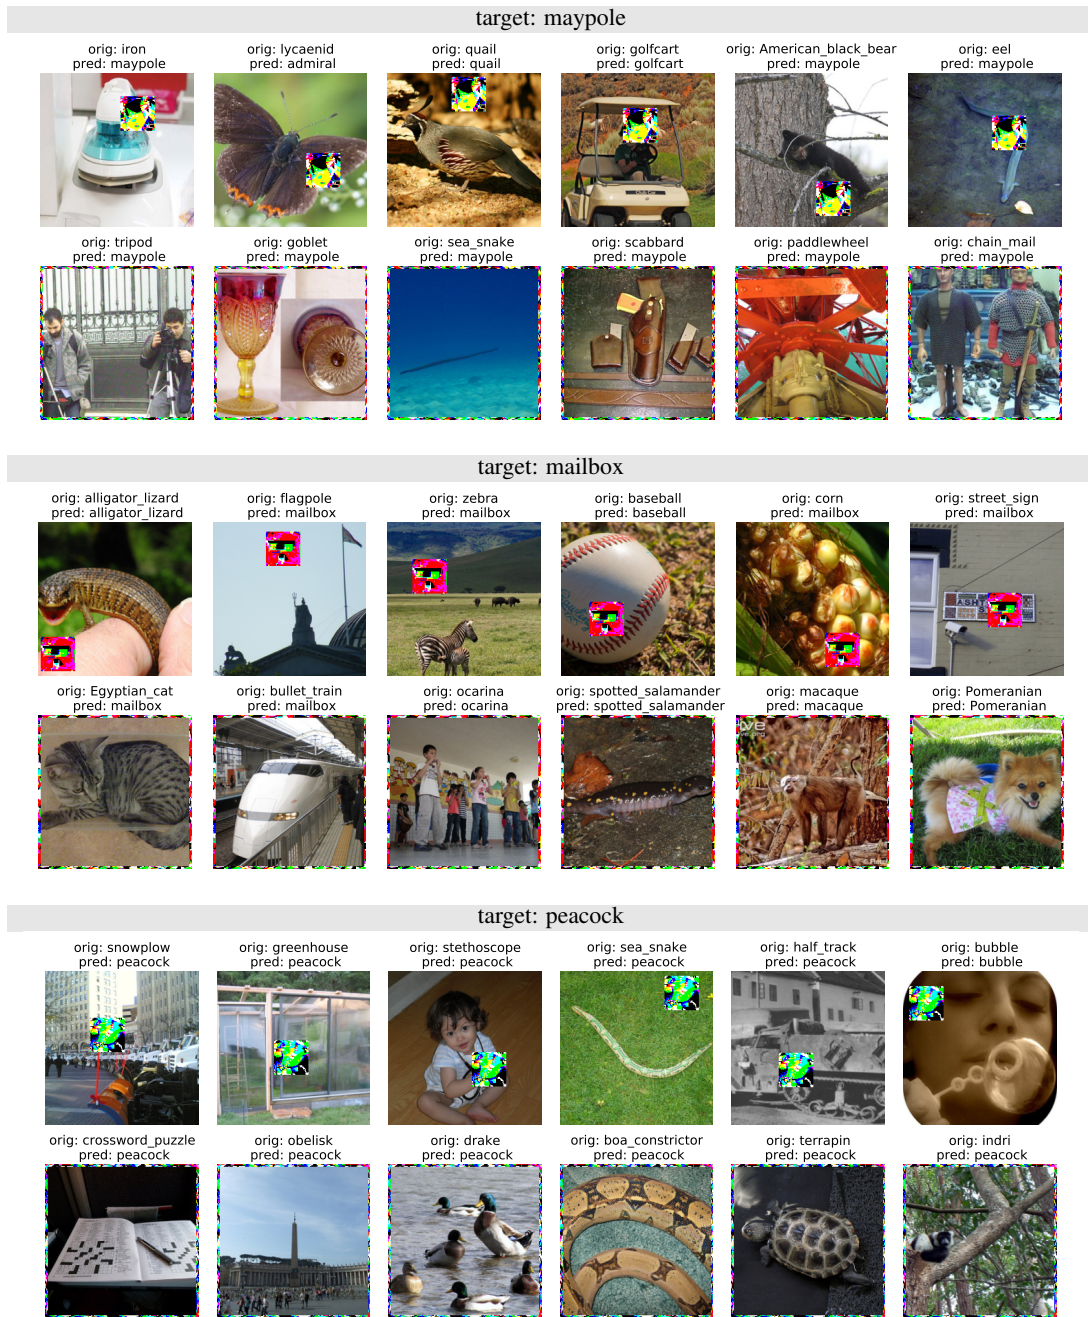


Figure 9: Example applications of universal targeted patches and frames for target classes "maypole", "mailbox" and "peacock", produced with our attack Sparse-RS.



Figure 10: Example applications of universal targeted patches and frames for target classes "traffic light", "digital clock" and "wardrobe", produced with our attack Sparse-RS.