

An Evolutionary, Gradient-Free, Query-Efficient, Black-Box Algorithm for Generating Adversarial Instances in Deep Networks

Raz Lapid, Zvika Haramaty, Moshe Sipper

Abstract—Deep neural networks (DNNs) are sensitive to adversarial data in a variety of scenarios, including the black-box scenario, where the attacker is only allowed to query the trained model and receive an output. Existing black-box methods for creating adversarial instances are costly, often using gradient estimation or training a replacement network. This paper introduces Query-Efficient Evolutionary Attack, *QuEry Attack*, an untargeted, score-based, black-box attack. *QuEry Attack* is based on a novel objective function that can be used in gradient-free optimization problems. The attack only requires access to the output logits of the classifier and is thus not affected by gradient masking. No additional information is needed, rendering our method more suitable to real-life situations. We test its performance with three different state-of-the-art models—Inception-v3, ResNet-50, and VGG-16-BN—against three benchmark datasets: MNIST, CIFAR10 and ImageNet. Furthermore, we evaluate *QuEry Attack*’s performance on non-differential transformation defenses and state-of-the-art robust models. Our results demonstrate the superior performance of *QuEry Attack*, both in terms of accuracy score and query efficiency.

Index Terms—Deep learning, computer vision, adversarial attack, evolutionary algorithm.

I. INTRODUCTION

DEEP neural networks (DNNs) have become the central approach in modern-day artificial intelligence (AI) research. They have attained superb performance in multifarious complex tasks and are behind fundamental breakthroughs in a variety of machine-learning tasks that were previously thought to be too difficult. Image classification, object detection, machine translation, and sentiment analysis are just a few examples of domains revolutionized by DNNs.

Despite their success, recent studies have shown that DNNs are vulnerable to adversarial attacks. A barely detectable change in an image, for example, can cause a misclassification in a well-trained DNN.

The authors are with the Department of Computer Science, Ben-Gurion University, Beer Sheva 84105, Israel. Corresponding author: R. Lapid, razla@post.bgu.ac.il.

Targeted adversarial examples can even evoke a misclassification of a specific class (e.g., misclassify a car as a cat). Researchers have demonstrated that adversarial attacks are successful in the real world and may be produced for data modalities beyond imaging, e.g., natural language and voice recognition [1], [2], [3], [4]. DNNs’ vulnerability to adversarial attacks has raised concerns about applying these techniques to safety-critical applications.

To discover effective adversarial instances, most past work on adversarial attacks has employed gradient-based optimization [5], [6], [7], [8], [9]. Gradient computation can only be executed if the attacker is fully aware of the model architecture and weights. Thus, these approaches are only useful in a white-box scenario, where an attacker has complete access and control over a targeted DNN. Attacking real-world AI systems, however, might be far more arduous. The attacker must consider the difficulty of implementing adversarial instances in a black-box setting, in which no information about the network design, parameters, or training data is provided. In this situation, the attacker is exposed only to the classifier’s input-output pairs. In this context, a typical strategy has been to attack trained replacement networks and hope that the generated examples transfer to the target model [10]. The substantial mismatch of the model between the alternative model and the target model, as well as the significant computational cost of alternative network training, often renders this technique ineffective.

In our work we assume a real-world, black-box attack scenario, wherein a DNN’s input and output may be accessed but not its internal configuration. We focus on a scenario in which a specific DNN is an image classifier, specifically, a convolutional neural network (CNN), which accepts an image as input and outputs a probability score for each class.

Herein, we present an evolutionary, gradient-free optimization approach for generating adversarial instances. Our proposed attack can deal with either constrained (ϵ value that constrains the norm of the



Fig. 1. Examples of adversarial attacks generated by QuEry Attack. Top row: Imagenet ($l_\infty = 6/255$). Middle row: CIFAR10 ($l_\infty = 6/255$). Bottom row: MNIST ($l_\infty = 60/255$). Left: the original image. Middle: a successful attack. Right: A failed attack.

allowed perturbation) or unconstrained (no constraint on the norm of the perturbation) problems, and focuses on constrained, untargeted attacks. We believe that our framework can be easily adapted to the targeted setting.

In the next section we review the literature on adversarial attacks. Section III summarizes the threat model we assume for our proposed evolutionary attack algorithm. The algorithm itself—*QuEry Attack* (for **Query-Efficient Evolutionary Attack**)—is delineated in Section IV. The experiments conducted to test the method, along with results, are described in Section V. We discuss our findings and present concluding remarks in Section VI.

Figure 1 shows examples of successful and unsuccessful instances of images generated by QuEry Attack, evaluated against ImageNet, CIFAR10, and MNIST.

II. RELATED WORK

Adversarial attacks against DNNs have become an important research field in the last few years. For a comprehensive survey, we refer the reader to [11].

An important distinction is between ‘white box’ and ‘black box’ attacks. In white-box attacks, the attacker has knowledge of the attacked model’s internal structure and parameters, and exploits that knowledge. It is commonly done by using the model’s gradients [5], [7], [12].

Although white-box methods achieved good results they usually do not represent a real-world scenario. More realistic is a black-box attack, where the attacker has no knowledge of the model’s structure and parameters. The attacker can only query the model—and act upon its outputs. We can further distinguish between the more common ‘light black box’ attacks, where the model gives the prediction’s confidence (which can be exploited), and ‘dark black box’ attacks where the model only gives the final class prediction (e.g., [10]).

The first effective black-box attack traded a-priori information of the model with extensive runtime querying [13]. Using a large number of queries it builds a substitute model, and attacks it with traditional white-box methods. It uses the transferability property, namely, an attack that succeeds on one model will likely succeed on a similar—though not identical—model.

Other works used the more permissive ‘light black box’ scenario, which can use the prediction’s confidence value. Some works estimate the gradient with this information and then use traditional gradient-based attacks [14].

All these black-box methods rely on gradients, and thus are sensitive to many defense methods that obscure gradients [15], [16], [17]. This has given rise to methods that do not rely on gradients at all, e.g., [18], which uses random search and is also query-efficient.

Instead of randomness, one can use evolutionary methods. In *Evolutionary Algorithms* (EAs), core concepts from evolutionary biology—inheritance, random variation, and selection—are harnessed in algorithms that are applied to complex computational problems. EA techniques have been shown to solve numerous difficult problems from widely diverse domains, and also to produce human-competitive machine intelligence [19]. EAs also present many important benefits over popular Machine Learning methods, including [20]: less reliance on the existence of a known or discoverable gradient within the search space; ability to handle design problems, where the objective is to design new entities from scratch; fewer required a priori assumptions about the problem at hand; seamless integration of human expert knowledge; ability to solve problems where human expertise is very limited; support of interpretable solution representations; support of multiple objectives.

The evolutionary method GenAttack is a *targeted* attack (thus not directly comparable to ours, which is *untargeted*) that used a fitness function that tries

to increase the probability of the target class and decrease the probability of the other classes [21]. Its fitness function ignored the distance between the images. Interestingly, GenAttack uses fitness-proportionate selection, which is employed less often nowadays due to known problems. It uses an adaptive mutation rate to balance between exploration in early phases and exploitation in later phases.

[22] treated the adversarial problem as one of multi-objective optimization: minimize the class prediction's score on one hand, and the distance between the original image and a perturbed one on the other hand.

Another attack method changes a single pixel [23]. This method uses differential evolution (DE) [24] without crossover. However, it sometimes required thousands of queries.

[25] also uses DE, but unlike the other evolutionary computation (EC) methods reviewed, it uses EC to approximate the gradients.

Unlike the above methods, which tried to minimize the perturbation as much as possible and make it as unnoticeable to the human eye as possible, [26] makes a small but noticeable change, which looks like a regular scratch (a similar approach in the domain of Natural Language Processing creates sentences that do not make sense to a human reader [27]). The approach uses DE as well, and also Covariance-Matrix Adaptation Evolution Strategies (CMA-ES) [28]. Unlike most attacks, which use the l_∞ or l_2 norms, this one is based on the l_0 norm.

In the EC methods seen so far, evolution is run against a *single* image, and each individual is a perturbation added to that image. [29], on the other hand, used the transferability property mentioned earlier to evolve a universal perturbation. An individual is an image mask that can be applied as a perturbation to any image.

Our extensive scrutiny of the literature and software repositories revealed that many authors compare their work to prior works that do not use the same threat model: There might be a mismatch in norms (e.g., l_2 vs. l_∞), white box vs. black box, or other subtle differences. Moreover, having investigated numerous software repositories, we found that running the code of many papers is far from straightforward.

III. THREAT MODEL

In the black-box attack setting, queries to the network are permitted but access to internal states is prohibited (e.g., executing backpropagation). Hence, our threat model, or scenario, is as follows:

- The attacker is unaware of the network's design, settings, or training data.
- The attacker has no access to intermediate values in the target model.
- The attacker can only use a black-box function to query the target model.

Note that the above threat model determines the comparisons we perform, which focus on attacks that are:

- 1) black-box,
- 2) untargeted,
- 3) l_∞ -norm bounded.

We can consider a network model to be a function:

$$f : [0, 1]^d \rightarrow \mathbb{R}^C, \quad (1)$$

where d is the number of input features and C is the number of classes. The c -th value $f_c(x) \in \mathbb{R}$ specifies the predicted score of classifying input image x as class c . The classifier assigns class $y = \operatorname{argmax}_{c=1, \dots, C} f_c(x)$ to the input x .

A *targeted* attack aims to create an image that will be incorrectly classified into a given (incorrect) class. An *untargeted* attack aims to create an image that will be incorrectly classified into *any* class except the correct one. An image \hat{x} is termed an adversarial example, with an l_p -norm bound of ϵ for x , if:

$$\begin{aligned} & \operatorname{argmax}_{c=1, \dots, C} f_c(\hat{x}) \neq y, \\ & \text{s.t. } \|\hat{x} - x\|_p \leq \epsilon \text{ and } \hat{x} \in [0, 1]^d. \end{aligned} \quad (2)$$

To wit, the model should classify \hat{x} incorrectly, while preserving similarity between x and \hat{x} under an l_p distance metric.

We focus on a black-box, score-based attack, wherein the only information of the threat model is the raw output (logits).

Our suggested black-box approach may theoretically be used in conjunction with classic machine-learning models, with the same input-output relationship. Because DNNs have reached state-of-the-art performance in a variety of image tasks, we focus on them in this paper.

IV. QUERY ATTACK

QuEry Attack is an evolutionary algorithm (EA) that explores a space of images, defined by a given input image and a given input model, in search of adversarial instances. It ultimately generates an attacking image for the given input image. Unlike white-box approaches, we make no assumptions about the targeted model, its architecture, dataset, or

training procedure. We assume that we have an image x , which a black-box neural network, f , classifies by outputting a probability distribution over the set of classes, as stated in Section III. The actual label y is computed as $y = \operatorname{argmax} f(x)$.

Our objective is to find a perturbed image, \hat{x} , of image x , such that, $\|\hat{x} - x\|_\infty \leq \epsilon$, which causes the network to predict $y' = \operatorname{argmax} f(\hat{x})$, such that $y' \neq y$. Finding \hat{x} may be cast as a constrained optimization problem:

$$\min_{\hat{x} \in [0,1]^d} \mathcal{L}(f(\hat{x}), y), \quad s.t. \quad \|\hat{x} - x\|_\infty \leq \epsilon, \quad (3)$$

for a given loss function \mathcal{L} .

We use loss \mathcal{L} as the fitness function, defined in our case as:

$$\text{fitness}(\hat{x}) = f_y(\hat{x}) - \max_{c \neq y} f_c(\hat{x}) + \lambda \|\hat{x} - x\|_2, \quad (4)$$

where \hat{x} is a perturbed image, f_y is the predicted score that \hat{x} belongs to class y , f_c is the predicted score that \hat{x} belongs to class $c \neq y$. In order to guarantee that the adversarial perturbation is as imperceptible as possible we penalize the l_2 distortion of the perturbation by including a regularization component in the fitness function. We use l_2 regularization because we noticed that most of the evolved adversarial examples were on the edges of the ϵ -ball, and we wanted to give precedence to examples which were closer to the original input. This penalization is determined by the λ value, which is the regularization strength. In our experiments we used $\lambda = 1$.

The ultimate goal is to minimize the fitness value: Essentially, the lower the logit of the correct class and the higher the maximum logit of the other classes—the better the value.

Algorithm 1 provides the pseudo-code of QuEry Attack. The original image x , along with a number of hyperparameters, are given as input to the algorithm. QuEry Attack generates an adversarial image \hat{x} , with the model classifying \hat{x} as y' such that $y' \neq y$ and $\|\hat{x} - x\|_\infty \leq \epsilon$.

The main goal of QuEry Attack is to produce a successful attack, using as few queries to the model as possible. The maximal number of queries equals generation count (G) \times population size (N).

a) Initialization: Initialization is crucial for optimization problems, e.g., in deep-learning training, gradient descent reaches a local minimum that is significantly determined by the initialization technique [30], [31]. QuEry Attack generates an initial population of perturbed images by randomly selecting images from the edges of the sphere centered on the original image x with radius $= \epsilon$. This is

accomplished by adding vertical stripes of width 1 along the image, with the color of each stripe sampled uniformly at random from $\{-\epsilon, \epsilon\}$ per channel (i.e., the pixels of each stripe can be either $-\epsilon$ or ϵ); in [18], they discovered that convolutional neural networks are especially vulnerable to such perturbations.

b) Mutation: Considering the use of (square-shaped) convolutional filters by convolutional neural networks, we used square-shaped perturbations. Specifically, we employed [18]’s technique for determining square size. Let $p \in [0, 1]$ be the proportion of elements to be perturbed for an image of shape $h \times w$. The nearest positive integer to $\sqrt{p \cdot h \cdot w}$ determines the length of the square’s edge, with p being a hyperparameter. We set it initially to $p = 0.1$, then halved it after $\{40, 200, 800, 4000, 8000, 16000, 24000, 32000\}$ queries, respectively (similar to [18]).

c) Crossover: We experimented both with single-point and two-point crossover, eventually settling on the latter as it performed better. The operator works by flattening both (two-dimensional image) parents, randomly picking two indices, then swapping the pixels between the chosen pixels.

The EA then proceeds by evaluating the fitness of each individual, selecting parents, and performing crossover and mutation to generate the next generation. This latter is obtained by adding one elite individual from the current generation, with all other next-generation individuals derived through crossover and mutation. The process is repeated until a successful perturbation is found or until the termination condition is met.

A major advantage of QuEry Attack is its amenability to parallelization—due to being evolutionary—in contrast to most other adversarial, iterative (non-evolutionary) attacks in this field.

V. EXPERIMENTS AND RESULTS

To evaluate QuEry Attack we set out to collect state-of-the-art algorithms for comparative purposes. A somewhat disconcerting reality we then encountered involved our struggle to find good benchmarks and software for comparison purposes. Sadly, we found ourselves wasting many a day (which, alas, turned into weeks) trying to run buggy software, chasing down broken links, issuing GitHub issues, and so forth. Perhaps this is due in part to the field of adversarial attacks being young.

Our experimental results are summarized in Table I. The code will be available at github.com/razla.

Algorithm 1 QuEry Attack**Input:**

$x \leftarrow$ original image
 $y \leftarrow$ original label
 $\epsilon \leftarrow$ maximum l_∞ distance
 $p \leftarrow$ proportion of elements in x to be perturbed
 $N \leftarrow$ population size
 $G \leftarrow$ maximum number of generations
 $T \leftarrow$ tournament size

Output:

$\hat{x} \leftarrow$ adversarial image

Main loop

```

1:  $gen \leftarrow 0$ 
2:  $pop \leftarrow \text{INIT}()$ 
3: while not  $\text{TERMINATION\_CONDITION}(pop, gen)$  do
4:   for  $\hat{x} \in pop$  do
5:     compute fitness of  $\hat{x}$  using Equation 4
6:    $new\_pop \leftarrow \emptyset$ 
7:    $elite \leftarrow \text{ELITISM}(pop)$ 
8:   add  $elite$  to  $new\_pop$ 
9:   for  $i \leftarrow 1$  to  $\frac{P-1}{3}$  do
10:     $parent_1 \leftarrow \text{SELECTION}(pop)$ 
11:     $parent_2 \leftarrow \text{SELECTION}(pop)$ 
12:     $offspring_1, offspring_2 \leftarrow \text{CROSSOVER}(parent_1, parent_2)$ 
13:     $mut_1 \leftarrow \text{SQUARE\_MUTATION}(offspring_1)$ 
14:     $mut_2 \leftarrow \text{SQUARE\_MUTATION}(offspring_2)$ 
15:    add  $offspring_1, mut_1, mut_2$  to  $new\_pop$ 
16:    $pop \leftarrow new\_pop$ 
17:    $gen \leftarrow gen + 1$ 
18:
19: return best  $\hat{x}$  from  $pop$  # QuEry Attack's final output

20: function  $\text{INIT}()$ 
21:    $pop \leftarrow \emptyset$ 
22:   for  $i \leftarrow 1$  to  $N$  do
23:      $\hat{x} \leftarrow \text{STRIPES\_INIT}(x)$ 
24:     add  $\hat{x}$  to  $pop$ 
25:   return  $pop$ 

26: function  $\text{ELITISM}(pop)$ 
27:   return best  $\hat{x}$  from  $pop$ 

```

28: **function** $\text{TERMINATION_CONDITION}(pop, gen)$

```

29:   if  $gen = G$  then
30:     return true
31:   for  $\hat{x} \in pop$  do
32:      $\hat{y} \leftarrow$  predicted label of  $\hat{x}$ 
33:     if  $\hat{y} \neq y$  then
34:       return true
35:   return false

```

36: **function** $\text{SELECTION}(pop)$

```

37:    $tournament \leftarrow$  randomly and uniformly pick  $T$  individuals from  $pop$ 
38:   return best  $\hat{x}$  from  $tournament$ 

```

39: **function** $\text{STRIPES_INIT}(\hat{x})$

```

40:   for  $i \leftarrow 1$  to  $c$  do #  $c$  is the image's number of channels
41:      $stripe \leftarrow$  create a vertical stripe of width 1, randomly positioned, with random values  $\in \{-\epsilon, \epsilon\}$ 
42:      $\hat{x} \leftarrow \hat{x} + stripe$ 
43:    $\hat{x} \leftarrow \Pi_\epsilon(\hat{x})$  #  $\Pi_\epsilon$ : clipping operator to ensure pixel values are within  $\epsilon$ -ball
44:   return  $\hat{x}$ 

```

45: **function** $\text{SQUARE_MUTATION}(\hat{x})$

```

46:    $c \leftarrow$  number of channels
47:    $f \leftarrow$  number of features ( $h \times w$ ) #  $h$ : height,  $w$ : width
48:    $k \leftarrow \lceil \sqrt{p \cdot f} \rceil$ 
49:    $\delta \leftarrow$  array of ones of size  $k \times k \times c$ .
50:    $row, col \leftarrow \mathcal{U}(\{0, \dots, w - h\})$  #  $\mathcal{U}$  randomly and uniformly selects from given set
51:   for  $i \leftarrow 1$  to  $c$  do
52:      $\tau \leftarrow \mathcal{U}(\{-2\epsilon, 2\epsilon\})$ 
53:      $\delta_{row+1:row+h, col+1:col+h, i} \leftarrow \tau \cdot \delta$ 
54:      $\hat{x} \leftarrow \hat{x} + \delta$ 
55:    $\hat{x} \leftarrow \Pi_\epsilon(\hat{x})$ 
56:   return  $\hat{x}$ 

```

57: **function** $\text{CROSSOVER}(parent_1, parent_2)$

```

58:   Flatten  $parent_1$  and  $parent_2$ 
59:   Perform standard two-point crossover (as explained in text), creating  $offspring_1, offspring_2$ 
60:    $offspring_1, offspring_2 \leftarrow \Pi_\epsilon(offspring_1), \Pi_\epsilon(offspring_2)$ 
61:   return  $offspring_1, offspring_2$ 

```

We evaluated QuEry Attack by executing experiments against three state-of-the-art image classification models—Inception-v3 [32], ResNet-50 [33], and VGG-16-BN [34]—over three image datasets: ImageNet, CIFAR-10, and MNIST. We employed 200 randomly picked and correctly classified images from the test sets. For ImageNet, Inception-v3 has an accuracy of 78.8%, ResNet-50 has an accuracy of 76.1%, and VGG-16-BN has an accuracy of 73.3% (these are top-1 accuracy values; for ImageNet, top-5 accuracy values are also sometimes given, which in our case are: Inception-v3 – 94.4%, ResNet-50 – 92.8%, VGG-16-BN – 91.5%). CIFAR10 accuracy values are: Inception-v3 – 93.7%, ResNet-50 – 93.6%, and VGG-16-BN – 94.0%. For the MNIST

dataset we trained a convolutional neural network (CNN), whose architecture is delineated in Table II, which attained 98.9% accuracy.

All accuracy values are over test images. We used the Adversarial Robustness Toolbox (ART) [35] to evaluate QuEry Attack against other attacks. We restricted all attacking algorithms to a maximum of 42K queries to the model ($N = 70$, $G = 600$) for MNIST and CIFAR10, and 84K queries ($N = 70$, $G = 1200$) for ImageNet. A query refers to a prediction supplied by the model for a given image. To make the most of the computational resources we had available we prioritized actual, experimental runs over hyperparameter runs, so hyperparameters were chosen through limited trial and error. In the future

TABLE I

EXPERIMENTAL RESULTS. ASR: ATTACK SUCCESS RATE. QUERIES: NUMBER OF MODEL QUERIES. EACH VALUE REPRESENTS THE MEDIAN OF 200 RUNS (IMAGES). EPSILON: $\mathcal{E} \in \{0...255\}$ (8-BITS PIXEL VALUES). TOP RESULTS ARE MARKED IN BOLDFACE.

ImageNet

Model	\mathcal{E}	QuEry Attack		Square		AdversarialPSO	
		ASR	Queries	ASR	Queries	ASR	Queries
Inception-v3	24	100%	1	100%	5	98.5%	51
	18	100%	2	100%	8	98.5%	69
	12	100%	22	95.5%	32	97%	102
	6	99.5%	276	99.0%	263	95%	285
ResNet-50	24	100%	1	99.5%	5	98.5%	51
	18	100%	1	100%	5	98.5%	69
	12	100%	13	100%	26	97%	102
	6	100%	211	99.0%	248	95%	285
VGG-16-BN	24	100%	1	100%	5	98.5%	51
	18	100%	1	100%	5	98.5%	69
	12	100%	1	100%	5	97%	102
	6	100%	77	100%	86	95%	285

CIFAR10

Model	\mathcal{E}	QuEry Attack		Square		AdversarialPSO	
		ASR	Queries	ASR	Queries	ASR	Queries
Inception-v3	24	100%	1	89.5%	5	97.5%	31
	18	100%	2	92.5%	17	96.0%	41
	12	97.5%	20	94.5%	77	95.0%	54
	6	91.0%	428	95.5%	776	94.5%	223
ResNet-50	24	100%	2	92.0%	8	97.5%	31
	18	100%	8	91.0%	23	96.0%	41
	12	100%	96	87.0%	110	95.0%	54
	6	99.0%	565	87.0%	449	94.5%	223
VGG-16-BN	24	100%	1	89.5%	5	97.5%	31
	18	99.0%	2	87.0%	14	96.0%	41
	12	98.0%	60	87.0%	86	95.0%	54
	6	95.5%	741	88.5%	890	94.5%	223

MNIST

Model	\mathcal{E}	QuEry Attack		Square		AdversarialPSO	
		ASR	Queries	ASR	Queries	ASR	Queries
Conv Net	80	100%	5	86.0%	14	76%	2675
	60	93.5%	72	93.0%	77	99%	292

we plan to perform a more thorough hyperparameter sweep using Optuna [36]. The only hyperparameters we set were the population size $N = 70$, tournament size $T = 25$, and $p = 0.1$; these are used for all experiments reported herein. The number of generations (G) was derived from the query budget.

AdversarialPSO [37] results were obtained by running the code in the GitHub repository referred to in their paper. Due to technical difficulties it was run against the original models that this attack was planned to run against, namely, Inception-v3 for ImageNet, and their own trained networks for CIFAR-10 and MNIST. We duplicated these results in the table for all models.

A. Attacking defenses

We show how QuEry Attack breaks a collection of defense strategies designed to boost the robustness of models against adversarial attacks.

a) *Attacking non-differentiable transformations:* Gradient masking is achieved via non-differentiable input transformations, which rely on manipulating gradients to defeat gradient-based attackers [38], [39]. Further, randomized transformations make it more difficult for the attacker to be certain of success. It is possible to foil such a defense by altering the defense module that performs gradient masking, but this is not an option within the black-box scenario. Herein, we

TABLE II
CNN USED FOR MNIST.

Conv Block		Hyperparameters	
Layer	Layer type	Hyperparameter	Value
1	Convolution	Epochs	300
2	BatchNorm2d	Batch size	64
3	ReLU	Optimizer	<i>Adam</i>
		Learning rate	0.01
		Weight decay	$1e-6$

CNN Architecture

Layer	Layer type	No. channels	Filter size	Stride
1	Conv Block	32	3×3	1
2	Max Pooling	N/A	2×2	2
3	Conv Block	64	3×3	1
4	Max Pooling	N/A	2×2	2
5	Conv Block	128	3×3	1
6	Max Pooling	N/A	2×2	2
7	Dropout ($p = 0.5$)	N/A	N/A	N/A
8	Fully Connected	128	N/A	N/A
9	Fully Connected	10	N/A	N/A



Fig. 2. JPEG compression examples, sorted from left to right by quality value, ranging from 10 to 100 (original image). Top images: from ImageNet, middle: CIFAR10, bottom: MNIST.

investigated three non-differentiable transformations against QuEry Attack: JPEG compression, bit-depth reduction (also known as feature squeezing), and spatial smoothing. We show that QuEry Attack can defeat these input modifications, due to its gradient-free nature.

JPEG compression [40] tries to generate patterns in color values to minimize the amount of data that has to be captured, resulting in a smaller file size. Some color values are estimated to match those of surrounding pixels in order to produce these patterns. This compression means that slight imperfections in the quality of the image will not be as noticeable. The degree of compression may be tweaked, providing a customizable trade-off between image quality and storage space. An example of the different compression degrees is shown in Figure 2. The results in Table III were evaluated with image quality $q = 70$.

Bit-depth reduction [41] can be done both by reducing the color depth of each pixel in an image and using spatial smoothing to smooth out individual pixel discrepancies. By merging samples that correspond to many different feature vectors in the original space into a single sample, bit-depth reduction decreases the search space accessible to an opponent.



Fig. 3. Bit-depth compression examples, sorted from left to right by bit-depth values, ranging from 1 to 8 (original image). Top images: from ImageNet, middle: CIFAR10, bottom: MNIST.




Fig. 4. Examples of spatial-smoothing compression, sorted from left to right by window-size values, ranging from 10 to 1 (original image). Top images: from ImageNet, middle: CIFAR10, bottom: MNIST.

An example of different bit-depth values is shown in Figure 3. The results in Table III were evaluated with bit depth $d = 3$.

The term “spatial smoothing” refers to the averaging of data points with their neighbors [42]. This has the effect of a low-pass filter, with high frequencies of the signal being eliminated from the data while low frequencies are enhanced. As a result, an image’s crisp “edges” are softened, and spatial correlation within the data becomes more prominent, as shown in Figure 4. Data averaging is determined according to a given window size. The results in Table III were evaluated with window $w = 5$.

Our results are delineated in Table III. For this experiment we used a total budget of 82K queries to the model ($N = 70$, $G = 1200$)—which was Inception-v3. For each given image, we first checked that it was correctly classified after applying the defense on the image, then we applied QuEry Attack on it. The different input values for the transformations were chosen such that applying them would not be destructive.

For these experiments we used the same budget of queries as in the previous experiments. For both CIFAR-10 and ImageNet, QuEry Attack has a high success rate against all non-differentiable transformations.

b) Attacking robust models: A model is considered to be robust when some of the input variables are largely perturbed, but the model still makes correct predictions. Recently, several techniques have

TABLE III
 QUERY ATTACK’S RESISTANCE TO NON-DIFFERENTIABLE TRANSFORMATION DEFENSES. JPEG COMPRESSION, BIT-DEPTH REDUCTION, AND SPATIAL SMOOTHING WERE TESTED ON CIFAR10 AND ImageNet.

Defense	\mathcal{E}	CIFAR10		ImageNet	
		ASR	Queries	ASR	Queries
JPEG Compression ($q = 70$)	18	100%	2	100%	2
	12	98.0%	13	97.5%	14
	6	93.5%	287	99.5%	311
Bit-Depth Reduction ($d = 3$)	18	100%	2	100%	2
	12	99.5%	5	100%	27
	6	96.0%	72	99.5%	145
Spatial Smoothing ($w = 5$)	18	100%	1	100%	1
	12	100%	1	100%	2
	6	99.0%	6	99.5%	144

TABLE IV
 QUERY ATTACK’S PERFORMANCE ON STATE-OF-THE-ART ROBUST MODELS OVER CIFAR10 AND ImageNet.

Model	\mathcal{E}	ImageNet	
		ASR	Queries
Wide ResNet-50-2	12	98.0%	98
	6	93.5%	1187
		CIFAR10	
Wide ResNet-70-16	18	94.5%	156
	12	85.0%	487

been proposed to render the models more robust to adversarial attacks. One commonly used technique to improve model robustness is adversarial training. Adversarial training integrates adversarial inputs—generated with other trained models—into the models’ training data. Adversarial training has been proved to be one of the most successful defense mechanisms for adversarial attacks [43], [44], [45], [46].

We conducted an experiment to see how well QuEry Attack performs on robust models. For CIFAR10 we used the robust model, WideResNet-70-16 [47], wherein they used generative models trained only on the original training set in order to enhance adversarial robustness to l_p norm-bounded perturbations. For ImageNet we used WideResNet-50-2 [48], wherein they used adversarially robust models for transfer learning. Both of these top models were taken from the RobustBench repository [49]. We used the same 200 randomly selected images from our previous experiments, a budget of 84K queries ($N = 70$, $G = 1200$) for CIFAR10, and a budget of 126K queries ($N = 70$, $G = 1800$) for ImageNet. As seen in Table IV, QuEry Attack succeeds at breaking those strongly defended models.

B. Transferability

An adversarial example for one model can often serve as an adversarial example for another model,

even if the two models were trained on different datasets, using different techniques; this is known as transferability [50]. White-box attacks may overfit on the source model, as evidenced by the fact that black-box success rates for an attack are almost always lower than those of white-box attacks [7], [51], [52], [53]. Herein, we checked transferability of our proposed black-box attack on 200 correctly classified ImageNet images by both the source model and the target model, using different ϵ values. The results, summarized in Table V, show a positive correlation between the ϵ values and the transferability success rate. We noted that attacks are better transferred between ResNet-50 to VGG-16-BN models and surmise this is due to the fact that ResNets models were mostly inspired by the philosophy of VGG models, wherein they use relatively small 3×3 convolutional layers.

VI. DISCUSSION AND CONCLUDING REMARKS

We presented an evolutionary, score-based, black-box attack, showing its superiority in terms of ASR (attack success rate) and number-of-queries over previously published work. QuEry Attack is a strong and fast attack that employs a gradient-free optimization strategy. We tested QuEry Attack against MNIST, CIFAR10, and ImageNet models, comparing it to other state-of-the-art algorithms. We evaluated

TABLE V
QUERY ATTACK’S TRANSFERABILITY ON IMAGENET MODELS. TSR: TRANSFERABILITY SUCCESS RATE

Source Model \rightarrow Target Model	\mathcal{E}	TSR
Inception-v3 \rightarrow ResNet-50	24	67.0%
	18	47.5%
	12	33.5%
	6	13.5%
ResNet-50 \rightarrow VGG-16-BN	24	90.0%
	18	81.5%
	12	61.5%
	6	28.5%
VGG-16-BN \rightarrow Inception-v3	24	59.0%
	18	46.5%
	12	30.0%
	6	12.5%

QuEry Attack’s performance against non-differential transformations and robust models, and it proved to succeed in both scenarios.

As noted, we discovered that the software scene in adversarial attacks is a tad bit muddy. We encourage researchers to place executable code on public repositories—code that can be used with ease. Furthermore, we feel that the field lacks standard means of measuring and comparing results. We encourage the community to establish common baselines for these purposes.

We came to realize the importance of a strong initialization procedure. While this is true of many optimization algorithms, it seems doubly so where adversarial optimization is concerned. Table I shows that successful attacks are sometimes found during initialization—the vertical-stripes initialization in particular proved highly potent—and even if not, the number of queries (and generations) is significantly curtailed.

Figures 5 and 6 show that adversarial examples are barely distinguishable to the human eye. Clearly, neural networks function quite differently than humans, capturing entirely different features. More work is needed to create networks that are robust in a human sense.

We think that evolutionary algorithms are well-suited for this kind of optimization problems and our findings imply that evolution is a potential research avenue for developing gradient-free black-box attacks. Furthermore, evolution needs to be evaluated against a fully black-box model.

ACKNOWLEDGEMENTS

We thank Ofer Hadar and Itai Dror for helpful discussions. This research was partially supported by the Israeli Innovation Authority and the Trust.AI consortium.




Fig. 5. Adversarial examples generated by QuEry Attack on MNIST, with $\epsilon = 0.3$. An image at $row, col = i, i$ shows the original image for class i . An image at $row, col = i, j, i \neq j$ shows a targetted attack on class i , with the target being class j .

REFERENCES

- [1] X. Wang, H. Jin, and K. He, “Natural language adversarial attack and defense in word level,” 2019.
- [2] J. X. Morris, E. Liffand, J. Y. Yoo, and Y. Qi, “Textattack: A framework for adversarial attacks in natural language processing,” *Proceedings of the 2020 EMNLP, Arxiv*, 2020.
- [3] N. Carlini and D. Wagner, “Audio adversarial examples: Targeted attacks on speech-to-text,” in *2018 IEEE Security and Privacy Workshops (SPW)*. IEEE, 2018, pp. 1–7.
- [4] L. Schönherr, K. Kohls, S. Zeiler, T. Holz, and D. Kolossa, “Adversarial attacks against automatic speech recognition systems via psychoacoustic hiding,” *arXiv preprint arXiv:1808.05665*, 2018.
- [5] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” *arXiv preprint arXiv:1412.6572*, 2014.
- [6] N. Papernot, P. McDaniel, S. Jha, M. Fredrikson, Z. B. Celik, and A. Swami, “The limitations of deep learning in

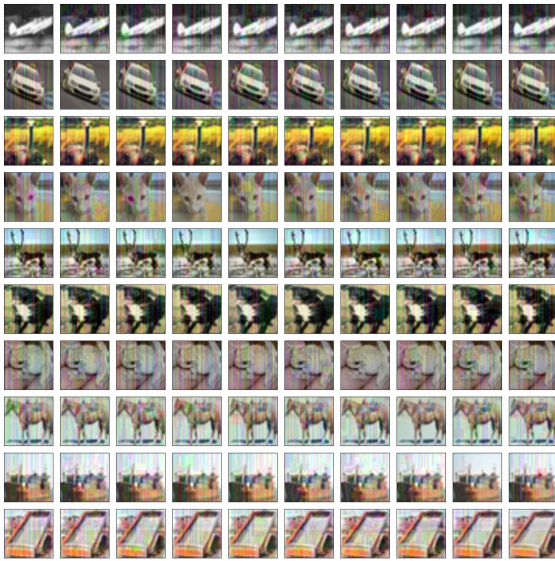


Fig. 6. Adversarial examples generated by QuEry Attack on CIFAR10, with $\epsilon = 12$. An image at $row, col = i, i$ shows the original image for class i . An image at $row, col = i, j, i \neq j$ shows a targetted attack on class i , with the target being class j .

- adversarial settings,” in *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2016, pp. 372–387.
- [7] N. Carlini and D. Wagner, “Towards evaluating the robustness of neural networks,” in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 39–57.
 - [8] S. Gu and L. Rigazio, “Towards deep neural network architectures robust to adversarial examples,” *arXiv preprint arXiv:1412.5068*, 2014.
 - [9] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, “Deepfool: a simple and accurate method to fool deep neural networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2574–2582.
 - [10] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, “Practical black-box attacks against machine learning,” in *Proceedings of the 2017 ACM Asia Conference on Computer and Communications Security*, 2017, pp. 506–519.
 - [11] H. Xu, Y. Ma, H. C. Liu, D. Deb, H. Liu, J. L. Tang, and A. K. Jain, “Adversarial Attacks and Defenses in Images, Graphs and Text: A Review,” *Int. J. Autom. Comput.* 2020 172, vol. 17, no. 2, pp. 151–178, mar 2020. [Online]. Available: <https://link.springer.com/article/10.1007/s11633-019-1211-x>
 - [12] A. U. H. Qureshi, H. Larijani, M. Yousefi, A. Adeel, and N. Mtetwa, “An adversarial approach for intrusion detection systems using jacobian saliency map attacks (jsma) algorithm,” *Computers*, vol. 9, no. 3, 2020. [Online]. Available: <https://www.mdpi.com/2073-431X/9/3/58>
 - [13] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, “Practical black-box attacks against machine learning,” *ASIA CCS 2017 - Proc. 2017 ACM Asia Conf. Comput. Commun. Secur.*, pp. 506–519, apr 2017.
 - [14] P. Y. Chen, H. Zhang, Y. Sharma, J. Yi, and C. J. Hsieh, “ZOO: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models,” *AISeC 2017 - Proc. 10th ACM Work. Artif. Intell. Secur. co-located with CCS 2017*, pp. 15–26, nov 2017.
 - [15] J. Buckman, A. Roy, C. Raffel, and I. Goodfellow, “Thermometer encoding: One hot way to resist adversarial examples,” in *International Conference on Learning Representations*, 2018.
 - [16] C. Guo, M. Rana, M. Cisse, and L. Van Der Maaten, “Countering adversarial images using input transformations,” *arXiv preprint arXiv:1711.00117*, 2017.
 - [17] G. S. Dhillon, K. Azizzadenesheli, Z. C. Lipton, J. Bernstein, J. Kossaifi, A. Khanna, and A. Anandkumar, “Stochastic activation pruning for robust adversarial defense,” *arXiv preprint arXiv:1803.01442*, 2018.
 - [18] M. Andriushchenko, F. Croce, N. Flammarion, and M. Hein, “Square attack: a query-efficient black-box adversarial attack via random search,” in *European Conference on Computer Vision*. Springer, 2020, pp. 484–501.
 - [19] K. Kannappan, L. Spector, M. Sipper, T. Helmuth, W. La Cava, J. Wisdom, and O. Bernstein, “Analyzing a decade of human-competitive (“HUMIE”) winners: What can we learn?” in *Genetic Programming Theory and Practice XII*, R. Riolo, W. P. Worzel, and M. Kotanchek, Eds. Cham: Springer International Publishing, 2015, pp. 149–166.
 - [20] M. Sipper, R. S. Olson, and J. H. Moore, “Evolutionary computation: the next major transition of artificial intelligence?” *BioData Mining*, vol. 10, no. 1, p. 26, Jul 2017.
 - [21] M. Alzantot, Y. Sharma, S. Chakraborty, H. Zhang, C.-J. Hsieh, and M. B. Srivastava, “GenAttack: Practical Black-Box Attacks with Gradient-Free Optimization,” in *Proc. Genet. Evol. Comput. Conf.*, ser. GECCO ’19. New York, NY, USA: Association for Computing Machinery, 2019, pp. 1111–1119. [Online]. Available: <https://doi.org/10.1145/3321707.3321749>
 - [22] S. Prochazka and R. Neruda, “Black-box Evolutionary Search for Adversarial Examples against Deep Image Classifiers in Non-Targeted Attacks,” *Proc. Int. Jt. Conf. Neural Networks*, jul 2020.
 - [23] J. Su, D. V. Vargas, and K. Sakurai, “One Pixel Attack for Fooling Deep Neural Networks,” *IEEE Trans. Evol. Comput.*, vol. 23, no. 5, pp. 828–841, oct 2019.
 - [24] S. Das and P. N. Suganthan, “Differential evolution: A survey of the state-of-the-art,” *IEEE Trans. Evol. Comput.*, vol. 15, no. 1, pp. 4–31, feb 2011.
 - [25] J. Lin, L. Xu, Y. Liu, and X. Zhang, “Black-box Adversarial Sample Generation Based on Differential Evolution,” *J. Syst. Softw.*, vol. 170, jul 2020. [Online]. Available: <https://arxiv.org/abs/2007.15310v1>
 - [26] M. Jere, L. Rossi, B. Hitaj, G. Ciocarlie, G. Boracchi, and F. Koushanfar, “Scratch that! An Evolution-based Adversarial Attack against Neural Networks,” *arXiv1912.02316*, dec 2019. [Online]. Available: <https://arxiv.org/abs/1912.02316v3>
 - [27] M. Di Giovanni and M. Brambilla, “EFSG: Evolutionary Fooling Sentences Generator,” *Proc. - 2021 IEEE 15th Int. Conf. Semant. Comput. ICSC 2021*, pp. 171–178, jan 2021.
 - [28] N. Hansen and A. Ostermeier, “Completely derandomized self-adaptation in evolution strategies,” *Evolutionary Computation*, vol. 9, no. 2, pp. 159–195, 2001.
 - [29] S. Wang, Y. Shi, and Y. Han, “Universal Perturbation Generation for Black-box Attack Using Evolutionary Algorithms,” *Proc. - Int. Conf. Pattern Recognit.*, vol. 2018-Augus, pp. 1277–1282, nov 2018.
 - [30] S. K. Kumar, “On weight initialization in deep neural networks,” *arXiv preprint arXiv:1704.08863*, 2017.
 - [31] S. Koturwar and S. Merchant, “Weight initialization of deep neural networks (DNNs) using data statistics,” *arXiv preprint arXiv:1710.10570*, 2017.
 - [32] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2818–2826.
 - [33] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE*

- Conference On Computer Vision And Pattern Recognition*, 2016, pp. 770–778.
- [34] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
 - [35] M.-I. Nicolae, M. Sinn, M. N. Tran, B. Buesser, A. Rawat, M. Wistuba, V. Zantedeschi, N. Baracaldo, B. Chen, H. Ludwig *et al.*, “Adversarial robustness toolbox v1. 0.0,” *arXiv preprint arXiv:1807.01069*, 2018.
 - [36] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, “Optuna: A next-generation hyperparameter optimization framework,” in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 2623–2631.
 - [37] R. Mosli, M. Wright, B. Yuan, and Y. Pan, “They Might NOT Be Giants Crafting Black-Box Adversarial Examples Using Particle Swarm Optimization,” 2020. [Online]. Available: https://doi.org/10.1007/978-3-030-59013-0_{_}22
 - [38] A. Athalye, N. Carlini, and D. Wagner, “Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 274–283.
 - [39] H. Qiu, Y. Zeng, Q. Zheng, T. Zhang, M. Qiu, and G. Memmi, “Mitigating advanced adversarial attacks with more advanced gradient obfuscation techniques,” *arXiv preprint arXiv:2005.13712*, 2020.
 - [40] G. K. Dziugaite, Z. Ghahramani, and D. M. Roy, “A study of the effect of jpg compression on adversarial images,” *arXiv preprint arXiv:1608.00853*, 2016.
 - [41] W. Xu, D. Evans, and Y. Qi, “Feature squeezing: Detecting adversarial examples in deep neural networks,” *arXiv preprint arXiv:1704.01155*, 2017.
 - [42] M. Mikl, R. Mareček, P. Hlušík, M. Pavlicová, A. Drastich, P. Chlebus, M. Brázdil, and P. Krupa, “Effects of spatial smoothing on fmri group inferences,” *Magnetic Resonance Imaging*, vol. 26, no. 4, pp. 490–503, 2008.
 - [43] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, “Domain-adversarial training of neural networks,” *Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2096–2030, 2016.
 - [44] F. Tramèr, A. Kurakin, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel, “Ensemble adversarial training: Attacks and defenses,” *arXiv preprint arXiv:1705.07204*, 2017.
 - [45] A. Shafahi, M. Najibi, M. A. Ghiasi, Z. Xu, J. Dickerson, C. Studer, L. S. Davis, G. Taylor, and T. Goldstein, “Adversarial training for free!” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
 - [46] E. Wong, L. Rice, and J. Z. Kolter, “Fast is better than free: Revisiting adversarial training,” *arXiv preprint arXiv:2001.03994*, 2020.
 - [47] S. Goyal, S.-A. Rebuffi, O. Wiles, F. Stimberg, D. A. Calian, and T. A. Mann, “Improving robustness using generated data,” *Advances in Neural Information Processing Systems*, vol. 34, 2021.
 - [48] H. Salman, A. Ilyas, L. Engstrom, A. Kapoor, and A. Madry, “Do adversarially robust imagenet models transfer better?” *Advances in Neural Information Processing Systems*, vol. 33, pp. 3533–3545, 2020.
 - [49] F. Croce, M. Andriushchenko, V. Schwag, E. Debenedetti, N. Flammarion, M. Chiang, P. Mittal, and M. Hein, “Robust-Bench: a standardized adversarial robustness benchmark,” in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2021. [Online]. Available: <https://openreview.net/forum?id=SSKZPJc7B>
 - [50] N. Papernot, P. McDaniel, and I. Goodfellow, “Transferability in machine learning: from phenomena to black-box attacks using adversarial samples,” *arXiv preprint arXiv:1605.07277*, 2016.
 - [51] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” *arXiv preprint arXiv:1706.06083*, 2017.
 - [52] W. Brendel, J. Rauber, and M. Bethge, “Decision-based adversarial attacks: Reliable attacks against black-box machine learning models,” *arXiv preprint arXiv:1712.04248*, 2017.
 - [53] P.-Y. Chen, H. Zhang, Y. Sharma, J. Yi, and C.-J. Hsieh, “Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models,” in *Proceedings of the 10th ACM workshop on artificial intelligence and security*, 2017, pp. 15–26.