

AdvMind: Inferring Adversary Intent of Black-Box Attacks

Ren Pang

rbp5354@psu.edu

Pennsylvania State University

Xinyang Zhang

xqz5366@psu.edu

Pennsylvania State University

Shouling Ji

svi@zju.edu.cn

Zhejiang University, Ant Financial

Xiapu Luo

csxluo@comp.polyu.edu.hk

Hong Kong Polytechnic University

Ting Wang

inbox.ting@gmail.com

Pennsylvania State University

ABSTRACT

Deep neural networks (DNNs) are inherently susceptible to adversarial attacks even under black-box settings, in which the adversary only has query access to the target model. In practice, while it may be possible to effectively detect such attacks (e.g., observing massive similar but non-identical queries), it is often challenging to exactly infer the adversary’s intent (e.g., the target class of the adversarial example the adversary attempts to craft) especially during early stages of the attacks, which is crucial for performing effective deterrence and remediation of the threats in many scenarios.

In this paper, we present AdvMind, a new class of estimation models that infer the adversary intent of black-box adversarial attacks in a *robust* and *prompt* manner. Specifically, to achieve robust detection, AdvMind accounts for the adversary adaptiveness such that her attempt to conceal the target will significantly increase the attack cost (e.g., the number of queries); to achieve prompt detection, AdvMind proactively synthesizes plausible query results to solicit subsequent queries from the adversary that maximally expose her intent. Through extensive empirical evaluation on benchmark datasets and state-of-the-art black-box attacks, we demonstrate that on average AdvMind detects the adversary intent with over 75% accuracy after observing less than 3 query batches and meanwhile increases the cost of adaptive attacks by over 60%. We further discuss the possible synergy between AdvMind and other defenses against black-box adversarial attacks, pointing to several promising research directions.

CCS CONCEPTS

• Security and privacy; • Computing methodologies → Machine learning;

KEYWORDS

adversarial examples; neural networks; black-box; defense; intent inference

ACM Reference Format:

Ren Pang, Xinyang Zhang, Shouling Ji, Xiapu Luo, and Ting Wang. 2020. AdvMind: Inferring Adversary Intent of Black-Box Attacks. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '20)*, August 23–27, 2020, Virtual Event, CA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3394486.3403241>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '20, August 23–27, 2020, Virtual Event, CA, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7998-4/20/08...\$15.00

<https://doi.org/10.1145/3394486.3403241>

the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '20), August 23–27, 2020, Virtual Event, CA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3394486.3403241>

1 INTRODUCTION

The recent advances in deep learning [16] have led to breakthroughs in a range of long-standing machine learning tasks (e.g., image classification [9], natural language processing [26], and even playing Go [30]), enabling many scenarios previously considered strictly experimental. However, it is well known that deep neural network (DNN) models are inherently vulnerable to adversarial inputs – those maliciously crafted samples to force the target DNNs to misbehave [32] – which significantly hinder their use in security-sensitive domains. Typically, adversarial inputs are crafted by carefully perturbing legitimate samples under the guidance of the gradient information of target DNNs (i.e., “white-box” attacks) [5, 11].

Meanwhile, many cloud-based service providers, including Amazon, Google, Microsoft, BigML, and others all have arisen to provide Machine Learning-as-a-service (MLaaS) platforms. On such platforms, increasingly many commercial and proprietary DNN models are being deployed with publicly accessible interfaces (“predictive APIs”), which allow users to query the backend models with inputs of interests and charge users on a pay-per-query basis. For instance, the CLARIFAI NSFW¹ (“not safe for work”) detection API returns probability scores that a given image contains nudity. The first 2,500 queries are free, it then charges \$2.4 per 1,000 queries.

Under such settings, the adversary must (i) construct adversarial inputs with only query access to the target DNNs (i.e., without gradient information) and (ii) also minimize the attack cost in terms of the number of queries, which spurs the research on “black-box” adversarial attacks. The existing black-box attacks in the literature can be categorized in two classes. The first one trains a surrogate DNN to emulate the target model and attacks the surrogate model using first-order white-box methods [25]. However, it often requires on the order of millions of queries to construct surrogate models for DNNs used in practice [14]; the second one constructs adversarial inputs by estimating the gradient of target DNNs via coordinate-wise finite difference methods [6, 14, 24], as illustrated in Figure 1. Due to its practical feasibility, in the following, we focus on the second class of black-box adversarial attacks.

Given their reliance on zeroth order methods, such black-box attacks often require issuing a number of queries in the vicinity of a point of interest to accurately estimate its gradient information.

¹<https://www.clarifai.com/models>




Figure 1: Illustration of black-box adversarial attacks.

Thus, it may be straightforward to detect the undergoing black-box attacks (e.g., by observing massive similar but non-identical queries [7]). Yet, it is often challenging to accurately infer the adversary’s intent – the target class of the adversarial input the adversary attempts to craft – especially during early stages of the attack. Identifying the adversary’s target in many scenarios, such as the person the adversary impersonates in face recognition [3], the insurance category the adversary claims in healthcare fraud [28], and the legitimate source the adversary disguises as in fake news detection [17], is critical for deploying proper mitigation strategies or performing prompt remediation against such threats.

Yet, despite the plethora of work on black-box adversarial attacks, the research on understanding the adversary’s intent is still fairly limited. To our best knowledge, this work represents a solid step towards bridging this gap. We present AdvMind, a novel framework for early-stage detection of the adversary’s intent in black-box adversarial attacks. At a high level, AdvMind achieves *robust* and *prompt* inference of adversary intent with two key strategies:

Robust intent estimation – Taking into account that the adversary may purposely issue fake queries to conceal her true intent, AdvMind employs an intent estimation model that is (i) reliable even in the presence of adversarial noise and (ii) agnostic to the underlying distribution from which the adversary samples the queries. Furthermore, by leveraging the fact that camouflage requires significantly increasing the attack cost (e.g., the number of queries issued), AdvMind creates for the adversary the dilemma between intent disclosure and attack cost.

Proactive intent solicitation – To achieve early-stage detection, AdvMind proactively synthesizes plausible query results to solicit subsequent queries from the adversary that maximally exposes her intent. Moreover, generated by slight perturbation on first-order information, the synthetic query results retain sufficient accuracy for legitimate use but can cause significant deviation if the results are used for gradient estimation as in black-box adversarial attacks, which further deters the adversary from successful attacks.

We extensively evaluate the efficacy of AdvMind with respect to benchmark datasets, popular DNNs, and state-of-the-art black-box attacks. We show that across all cases, AdvMind manages to detect the adversary intent with over 75% accuracy after observing less than 3 query batches and meanwhile increases the query cost of adaptive attacks by over 60%. We further discuss the synergy between AdvMind and existing defenses against black-box adversarial attacks, pointing to several promising research directions.²

²The code and data used in the paper is released at <https://github.com/alps-lab/advmind>.

2 FUNDAMENTALS

We begin with introducing a set of concepts and assumptions used throughout the paper.

2.1 Threat Models

In the following, we primarily focus on predictive tasks (e.g., image classification [9]). Under this setting, a DNN f represents a function $f : \mathcal{X} \rightarrow \mathcal{C}$, which assigns a given input $x \in \mathcal{X}$ to one of a set of predefined classes \mathcal{C} , $f(x) = c \in \mathcal{C}$. DNNs are inherently vulnerable to adversarial inputs, which are maliciously crafted to force target DNNs to misbehave [8, 22, 23, 32]. Here we focus on the setting of targeted attacks. Specifically, an adversarial input x_* is often generated by slightly modifying a benign input x_0 , with the objective of forcing f to misclassify x_* into a target class c_* , $f(x_*) = c_* \neq f(x_0)$. To ensure that x_* is perceptually similar to x_0 , the perturbation is constrained to a set of allowed perturbations (e.g., a norm ball $\mathcal{B}_\epsilon(x_0) = \{x \mid \|x - x_0\|_\infty \leq \epsilon\}$).

Consider project gradient descent (PGD) [22], a universal first-order adversarial attack, as an example. It is implemented as iterative projected gradient descent on the negative loss function:

$$x^{(i+1)} = \Pi_{\mathcal{B}_\epsilon(x_0)}(x^{(i)} - \alpha \operatorname{sgn}(\nabla \ell(f(x^{(i)}), c_*))) \quad (1)$$

where Π denotes the projection operator, α ($\alpha \geq 0$) represents the learning rate, the loss function ℓ measures the difference of the model prediction $f(x)$ and the class c_* desired by the adversary (e.g., cross entropy), and $x^{(0)}$ is initialized as x_0 .

Algorithm 1: Black-Box Attack

Input: original input $x^{(0)} = x_0$, target class c_* , target DNN f , learning rate α , norm constraint ϵ , and number of iterations n_{iter}
Output: adversarial input x_*

```

1 for  $i = 1, \dots, n_{\text{iter}}$  do
2    $\hat{g}(x^{(i)}) \leftarrow \text{GradEstimate}(x^{(i)}, c_*)$ ;
3    $x^{(i+1)} \leftarrow \Pi_{\mathcal{B}_\epsilon(x_0)}(x^{(i)} - \alpha \operatorname{sgn}(\hat{g}(x^{(i)})))$ ;
4   if  $f(x^{(i+1)}) = c_*$  then return  $x^{(i+1)}$  as  $x_*$ ;
5 return Fail;

```

Under the black-box setting, the adversary has only query access to the target DNN f (i.e., without the gradient information), as sketched in Algorithm 1. Note that here we focus on an account-oriented setting, where users must create an account before they can query the model [7]. The adversary typically uses coordinate-wise finite difference methods to estimate the gradient and then applies gradient descent as in Eqn (1) to generate adversarial inputs x_* . Different black-box adversarial attacks mainly differ in their gradient estimation method $\text{GradEstimate}(x, c_*)$ in Algorithm 1. In this work, we consider three representative black-box attacks.

To clarify, $f(x)$ returns the classification label c of image x in the above content. However, we’ll introduce different attack methods in the following section, within which $f(x)$ is the probability confidence vector after softmax layer. For simplicity, in the remainder of the paper, $f(x)$ stands for the classification label in $\ell(f(x), c_*)$, and stands for the probability vector when it’s alone.

Zerth Order Natural Evolutionary Strategy. In [14], Ilyas *et al.* apply natural evolutionary strategies (NES), a derivative-free

optimization method [35] to estimate the gradient of $\hat{g}(x^{(i)})$. Specifically, NES generates n_{query} data points in the neighborhood of $x^{(i)}$ by sampling from a normal distribution, retrieves their predictions from the target DNN f , and estimates the gradient $\hat{g}(x^{(i)})$ as follow:

$$\hat{g}(x^{(i)}) = \frac{1}{\sigma n_{\text{query}}} \sum_{j=1}^{n_{\text{query}}} f(x^{(i)} + \sigma u^{(j)}) u^{(j)} \quad (2)$$

where each sample $u^{(j)}$ is sampled from the standard normal distribution $\mathcal{N}(0, I)$, and σ is the sampling variance.

To accelerate the sampling, it is possible to generate only half of the samples from the distribution and set the other half symmetrically. The gradient estimate is then formulated as follows:

$$\hat{g}(x^{(i)}) = \frac{1}{\sigma n_{\text{query}}} \sum_{j=1}^{[n_{\text{query}}/2]} (f(x^{(i)} + \sigma u^{(j)}) - f(x^{(i)} - \sigma u^{(j)})) u^{(j)} \quad (3)$$

Zeroth Order signSGD. Liu *et al.* [19] integrate signSGD [2], a compressed optimization method only using gradient sign, with zeroth order optimization [6], and propose one- and two-sided estimators of $\hat{g}(x^{(i)})$. The one-sided estimator is given by:

$$\begin{aligned} \hat{g}(x^{(i)}) &= \frac{1}{\sigma n_{\text{query}}} \sum_{j=1}^{n_{\text{query}}} (f(x^{(i)} + \sigma u^{(j)}) - f(x^{(i)})) u^{(j)} \\ &= \frac{1}{\sigma n_{\text{query}}} \sum_{j=1}^{n_{\text{query}}} f(x^{(i)} + \sigma u^{(j)}) u^{(j)} - \frac{f(x^{(i)})}{\sigma n_{\text{query}}} \sum_{j=1}^{n_{\text{query}}} u^{(j)} \end{aligned} \quad (4)$$

where the NES estimator in Eqn (2) can be used to compute the first term above. Since the expectation of each sample $u^{(j)}$ is 0, Eqn (4) and Eqn (2) share the same expectation. Yet, compared with the NES estimator, the second term in Eqn (4) reduces the error caused by the variance of $u^{(j)}$, leading to faster convergence.

The two-sided estimator in [19] is essentially identical to the improved NES estimator in Eqn (3).

Zeroth Order Hessian-Aware. Note that the NES and signSGD estimators only use first-order information in estimating $\hat{g}(x^{(i)})$. In [36], Haishan *et al.* argue that it is more efficient to estimate $\hat{g}(x^{(i)})$ if the second-order information (Hessian) is considered. They propose the following gradient estimator:

$$\hat{g}(x^{(i)}) = \frac{1}{\sigma n_{\text{query}}} \sum_{j=1}^{n_{\text{query}}} \left(f(x^{(i)} + \sigma \tilde{H}^{-\frac{1}{2}} u^{(j)}) - f(x^{(i)}) \right) \tilde{H}^{-\frac{1}{2}} u^{(j)} \quad (5)$$

where $\tilde{H}^{-\frac{1}{2}}$ is the Cholesky inverse of the Hessian estimate \tilde{H} . Further, \tilde{H} is estimated as follows:

$$\tilde{H} = \frac{\sum_{j=1}^{n_{\text{query}}} |f(x^{(i)} + \sigma u^{(j)}) + f(x^{(i)} - \sigma u^{(j)}) - 2f(x^{(i)})| u^{(j)} \otimes u^{(j)}}{2\sigma^2 n_{\text{query}}} \quad (6)$$

where \otimes represents the outer product of two vectors. Note that to ensure the feasibility of Cholesky decomposition and invertibility, a small identity matrix (e.g., τI) is often added to \tilde{H} .

2.2 Overview of AdvMind

Despite their apparent variations, the query-based black-box adversarial attacks share common patterns. Each attack consists of a sequence of iterations. At the i -th iteration, the adversary issues queries regarding the current input $x^{(i)}$, which we refer to as the query of interest (QOI), and the auxiliary data points in the vicinity of $x^{(i)}$, which are used to estimate $\hat{g}(x^{(i)})$.

Therefore, to infer the adversary's intent (i.e., the target class c_*), AdvMind adopts the following strategies: (i) grouping the incoming queries into a sequence of query batches, each corresponding to an iteration (the feasibility of grouping queries into different iterations is detailed in technical report [27]; (ii) estimating the QOI at each iteration – As the auxiliary queries appear in the close vicinity of $x^{(i)}$, it is feasible to accurately estimate $\hat{x}^{(i)}$ by clustering the queries; and (iii) estimating the gradient direction followed by the adversary – By linking consecutive QOIs (e.g., $\hat{x}^{(i)}$ and $\hat{x}^{(i+1)}$), it is feasible to infer the gradient direction the adversary is following to craft the adversarial input. Intuitively, as the adversary attempts to minimize the loss function with respect to the target class c_* , the gradient tends to point to that direction. Further, by combining the inference at multiple iterations, it is possible to reliably infer the adversary's intent c_* . This scheme is illustrated in Figure 2.




Figure 2: Illustration of AdvMind framework.

Yet, this naïve inference model suffers the following drawbacks. First, at each iteration, the adversary may attempt to inject fake queries (i.e., irrelevant to the QOI) to conceal her intent. As outliers, the fake queries may significantly deviate AdvMind's estimate $\hat{x}^{(i)}$ from the ground-truth QOI $x^{(i)}$. To account for the possible adversarial noise, AdvMind adopts a robust QOI estimator. With reasonable fake query proportion p_{fake} (e.g., $p_{\text{fake}} \leq 40\%$), the robust estimator is able to estimate the QOI with bounded bias. Further, it creates for the adversary the dilemma between intent disclosure and attack cost (e.g., in terms of the number of queries).

Second, as the adversary performs projected descent based on the estimated gradient (Algorithm 1), the descent direction she follows may deviate from the true gradient direction, which in turn causes errors in AdvMind's inference of the adversary's target class c_* . To enable the early-stage inference and minimize the inference uncertainty, AdvMind adopts proactive intent solicitation. Specifically, by slightly perturbing the query answers, AdvMind solicits subsequent queries from the adversary that maximally expose the target class. Note that such perturbation retains sufficient accuracy in the query answers for legitimate use but causes significant deviation for gradient estimation.

We elaborate on the strategies of robust intent estimation and proactive intent solicitation in § 3 and § 4 respectively.

3 QOI ESTIMATION

Here we focus on the QOI estimation at each iteration. We first introduce a naïve estimator, which is however vulnerable to adversarial noise (i.e., fake queries), and then present a robust estimator, which bounds the estimation bias under reasonable noise ratio.

3.1 Naïve Estimator

Recall that at i -th iteration, the adversary samples a set of auxiliary queries $\mathcal{X}^{(i)}$ in the vicinity of the QOI $x^{(i)}$ following a certain distribution (e.g., Gaussian). We may estimate $\tilde{x}^{(i)}$ using the mean or median of the received queries $\mathcal{X}^{(i)}$.³ Note that compared with the mean estimate, the median estimate is more robust to outliers but has a larger bias as it does not utilize the data points on edges.




Figure 3: Naïve versus robust estimation of QOI in the presence of adversarial noise.

However, to conceal her intent, the adversary may adaptively inject fake queries into \mathcal{X} to deviate \hat{x} from the true QOI x . As illustrated in Figure 3, the presence of such adversarial noise significantly affects the naïve estimator, even for the median estimate. The adversary may select from a variety of models to generate fake queries, including random sampling, duplicating existing queries, and query blinding [7], which lead to AdvMind’s varying estimation error (e.g., measured by $\|\hat{x} - x\|$).

To be concise, in the following we consider Huber’s contamination model [13], which subsumes a number of models above. Specifically, it assumes that the set of queries $x_1, \dots, x_{n_{\text{query}}}$ are randomly sampled from a mixture distribution, with the cumulative distribution defined by:

$$F = (1 - p_{\text{fake}})\Phi + p_{\text{fake}}H \quad (7)$$

where p_{fake} is the fraction of fake queries ($0 \leq p_{\text{fake}} < 0.5$) among all the queries, H is the (unknown) cumulative distribution of fake queries, and Φ is the cumulative distribution of true auxiliary queries (e.g., Gaussian or uniform).

Next we develop a robust QOI estimator against the influence of fake queries and bounds the estimation bias even under the worst-case distribution H .

3.2 Robust Estimator

We extend the robust mean estimator [13, 29] to the context of QOI estimation. Specifically, given the set of queries \mathcal{X} ($|\mathcal{X}| = n_{\text{query}}$), we estimate the QOI in an iterative manner, in which the k -th round estimator $\tilde{x}^{(k)}$ is updated as follows:

$$\tilde{x}^{(k)} = \tilde{x}^{(k-1)} + \frac{\text{MAD}(\mathcal{X})}{\mathbb{E}_{\Phi}[\psi']} \cdot \frac{\sum_{i=1}^{n_{\text{query}}} \psi\left(\frac{x_i - \tilde{x}^{(k-1)}}{\text{MAD}(\mathcal{X})}\right)}{n_{\text{query}}} \quad (8)$$

³When the context is clear, we omit the superscript (i) in the notations.

where the function ψ is defined as $\psi(x) = (e^x - 1)/(e^x + 1)$, $\text{MAD}(\mathcal{X})$ is the median average deviation of the query set \mathcal{X} . Concretely, with Φ instantiated as the standard Gaussian distribution, we may estimate the expectation $\mathbb{E}_{\Phi}[\psi']$ as follows:

$$\mathbb{E}_{\Phi}[\psi'] = \int_{-\infty}^{+\infty} \psi'(s) d\Phi(s) \approx 0.4132 \quad (9)$$

The value of $\tilde{x}^{(0)}$ is initialized as the median of the query set \mathcal{X} . As observed in our empirical evaluation (which is also consistent with the results of [29]), one-iteration estimation $\tilde{x}^{(1)}$ is typically accurate enough to approximate the QOI x . Thus we set the number of iterations $k = 1$ by default.

Further, we may also estimate the bias of \tilde{x} under the worst-case distribution of adversarial noise, which is needed for inferring the adversary’s target in § 4. Specifically, the bias of the k -th iteration estimator, $B(\tilde{x}^{(k)})$, is updated as:

$$B(\tilde{x}^{(k)}) = B(\tilde{x}^{(k-1)}) + \frac{(1 - p_{\text{fake}})\mathbb{E}_{\Phi}[\psi(X - B(\tilde{x}^{(k-1)}))] + p_{\text{fake}}\psi(\infty)}{\mathbb{E}_{\Phi}[\psi']} \quad (10)$$

with the initial value given by $B(\tilde{x}^{(0)}) = \Phi^{-1}\left(\frac{1}{2(1-p_{\text{fake}})}\right)$, where Φ^{-1} is the inverse function of the Gaussian distribution.

4 INTENT INFERENCE

In this section, we introduce AdvMind’s intent inference model, which based on the QOIs estimated at consecutive iterations, identifies the adversary’s target class. Moreover, to infer the adversary’s intent at an early attack stage and to minimize the inference uncertainty, AdvMind employs a proactive solicitation strategy, which solicits subsequent queries from the adversary that maximally expose her target class.

4.1 Passive Inference

We begin with the passive inference model, which identifies the target class c_* directly from the QOI estimation.

Specifically, assuming $\tilde{x}^{(i)}$ and $\tilde{x}^{(i+1)}$ as the estimated QOIs at two consecutive iterations, we compute the descent direction vector $d^{(i)} = \tilde{x}^{(i+1)} - \tilde{x}^{(i)}$ and the gradient vector $g_c(\tilde{x}^{(i)})$ at $\tilde{x}^{(i)}$ with respect to each class $c \in C$, and regards the cosine similarity of $d^{(i)}$ and $g_c(\tilde{x}^{(i)})$ as the score of c :

$$\text{score}(c) = \frac{d^{(i)} \cdot g_c(\tilde{x}^{(i)})}{\|d^{(i)}\| \|g_c(\tilde{x}^{(i)})\|} \quad (11)$$

We further normalize the scores of C and transforms them into a probability vector using the softmax function. Let $p^{(i)}$ denote the probability vector estimated at the i -th iteration. We then infer the most likely target class \hat{c}_* as the class c leading to the overall largest probability:

$$\hat{c}_* = \arg \max_{c \in C} \Pi_i p_c^{(i)} \quad (12)$$

Algorithm 2 sketches the procedure of adversary intent inference for n_{iter} iterations.

Algorithm 2: Passive Intent Inference

Input: number of iterations n_{iter} , QOI estimates $(\tilde{x}^{(1)}, \dots, \tilde{x}^{(n_{\text{iter}})})$, confidence threshold κ
Output: target class estimate \hat{c}_*
// initialization

- 1 $p \leftarrow 1;$
- 2 **for** $i = 2, \dots, n_{\text{iter}}$ **do**
- 3 $d \leftarrow \tilde{x}^{(i+1)} - \tilde{x}^{(i)}$;
- 4 **for** $c \in C$ **do**
- 5 // descent direction estimation
- 6 $g_c = \nabla \ell(f(\tilde{x}^{(i)}), c)$;
- 7 $s_c \leftarrow \frac{d \cdot g_c}{\|d\| \|g_c\|}$;
- 8 // normalization
- 9 $p \leftarrow (p \odot \text{softmax}(s))$;
- 10 **if** $\max_c p_c \geq \kappa$ **then return** $\arg \max_c p_c$ as \hat{c}_* ;
- 11 **return** $\arg \max_c p_c$ as \hat{c}_* ;

4.2 Proactive Solicitation

As illustrated in Figure 4, because the adversary relies on the estimated gradient (Algorithm 1), the descent direction $\hat{g}(x)$ may deviate from the true gradient direction $g(x)$, causing possible errors in AdvMind’s estimation of c_* . The estimate error of $\hat{g}(x)$ mainly results from the bias between $\hat{g}(x)$ and $g(x)$. Takeing signSGD in Eqn (4) as a concrete example, the expectation of $\hat{g}(x)$ is given by:

$$\begin{aligned} \mathbb{E}[\hat{g}(x)] &= \frac{1}{\sigma} \mathbb{E}_u [(f(x + \sigma u) - f(x)) u] \\ &= \mathbb{E}_u [(g(x) \cdot u) u] + \mathbb{E}_u [R(x, u) u] \end{aligned} \quad (13)$$

Where $u \sim \mathcal{N}(0, I)$ and $R(x, u)$ is the second order remainder of $f(x)$ in its Taylor expansion. We have the following proposition (proof in technical report [27]).

PROPOSITION 4.1. *For any vector v , we have $\mathbb{E}_u[(v \cdot u) u] = v$, where $u \sim \mathcal{N}(0, I)$.*

Applying this proposition to Eqn (13), the first term becomes $g(x)$, which is the true gradient. Yet, due to the high-order term $R(x, u)$, Eqn (13) is a biased estimate of $g(x)$. Meanwhile, if we remove the high-order information in the query answers, the adversary would be able to estimate the gradient more accurately, and AdvMind would also obtain a more accurate estimation of the adversary’s intent. To this end, we define the prediction for a query x_i as:

$$\hat{f}(x_i) = f(\tilde{x}) + G \cdot (x_i - \tilde{x}) \quad (14)$$

where \tilde{x} is the estimated QOI, $\hat{f}(x_i)$ is the modified prediction on x_i , G is a gradient matrix ($G \in \mathbb{R}^{|C| \times n}$), in which the c -th row is the gradient vector of the class c .

We have the following proposition that ensures the similarity of $\hat{f}(x_i)$ and its true value $f(x_i)$ (proof in technical report [27]):

PROPOSITION 4.2. *If the model f is Lipschitz continuous with $\|f(x) - f(y)\| \leq K \|x - y\|$, we have*

$$\|\hat{f}(x_i) - f(x_i)\| \leq (K + \|G\|)(\sigma \|u\| + B(\tilde{x})) \quad (15)$$

where $u \sim \mathcal{N}(0, I)$ and $B(\tilde{x})$ is the bias of the QOI estimate in § 3.

Applying Eqn (14), the adversary’s gradient descent direction now shares the same expectation with the true gradient direction




Figure 4: Passive inference versus proactive solicitation

with respect to the target class c_* . Yet, due to the estimation variance, the gradient of other classes may be even closer to the adversary’s estimate, as shown in Figure 4. To address this issue, we force the gradient directions of different classes to be orthogonal to each other so that if the adversary follows a descent direction, the target class can be easily distinguished from other classes, as illustrated in Figure 4.

To this end, the new gradient matrix is defined as follows:

$$G = (1 - \mu)g(\tilde{x}) + \mu M \quad (16)$$

where the parameter $\mu \in [0, 1]$ (dispersion coefficient) controls the magnitude of direction dispersion, and M is an orthogonal matrix set by AdvMind (an example given below). Specifically, M separates gradients of different classes since each row of G is the gradient of the corresponding class.

$$M = \begin{bmatrix} 1 & 0 & \dots & 0 & 1 & 0 & \dots & 0 & \dots \\ 0 & 1 & \dots & 0 & 0 & 1 & \dots & 0 & \dots \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & 0 & 0 & \dots & 1 & \dots \end{bmatrix} \quad (17)$$

Note that Eqn (17) is only one possible instantiation of M . Let c be the number of classes and d be the data dimensionality. There are essentially $(c!)^{(d/c)}$ different permutation matrices from which the defender is able to arbitrarily select, creating a prohibitively large number of possibilities for the attacker to explore in order to recover the genuine gradient information.

In conclusion, the proactive solicitation modifies the first-order information (i.e., gradient) while preserving the zeroth-order information. By properly controlling the proportion of gradient information released to the adversary, AdvMind is able to expose the adversary’s intent at an early stage while also deterring her from achieving successful attacks.

5 EMPIRICAL EVALUATION

Next we empirically evaluate the performance of AdvMind with respect to benchmark datasets, popular DNNs, and state-of-the-art black-box adversarial attacks. The experiments are designed to answer the following questions:

- RQ1: Is AdvMind effective to detect the attacker’s intent during an early stage of the attack?
- RQ2: Is AdvMind effective against the attacker who purposely attempts to conceal her intent?
- RQ3: Is AdvMind effective against various black-box attacks?

5.1 Experimental Setting

Datasets. We primarily use four image classification datasets.

- CIFAR10 [15] – it includes 60K 32×32 colored images from 10 classes (e.g., ‘ship’);
- CIFAR100 [15] – it is essentially CIFAR10 but categorized into 100 fine-grained classes;
- ISIC [10] – it represents the skin cancer screening task from the ISIC 2018 challenge, in which given 600×450 skin lesion images are categorized into a 7-disease taxonomy (e.g., melanoma);
- Mini-VGGface2 – it is a subset of the VGGface2 dataset [4], which consists of 224×224 (center-cropped) color images drawn from 50 individuals (e.g., ‘Aaron Stanford’);

DNNs. We consider 4 DNNs as the backend models, VGG13 [31] for CIFAR10, VGG16 for CIFAR100, ResNet101 [12] for ISIC, and ResNet18 for VGGface2, which respectively attain 92.440%, 70.470%, 88.176%, and 96.175% accuracy on corresponding testing sets.

Attack and Inference Models. We consider 3 state-of-the-art black-box adversarial attacks, NES [14], signSGD [2], and HessAware [36], and their adaptive variants (indicated by the superscript of “A”). More details of the attacks can be found in § 2. We build 3 variants of AdvMind: basic (AdvMind) – which passively observes the attacker’s queries; robust (AdvMind^R) – which employs robust intent estimation; and robust + proactive (AdvMind^{RP}) – which adopts both robust intent estimation and proactive intent solicitation.

All the algorithms are implemented in PyTorch. The default parameter setting is summarized in Table 1. Note that the HessAware attack is only conducted on the CIFAR datasets due to the infeasible cost of computing Hessian matrices for 224×224 images.

| Parameter | Definition | Dataset | | | |
|--------------------|---------------------------------|---------|----------|------|----------|
| | | CIFAR10 | CIFAR100 | ISIC | VGGface2 |
| α | adversary’s learning rate | 0.01 | | | |
| ϵ | norm constraint of perturbation | 0.03 | | | |
| n_{iter} | number of attack iterations | 10 | | | |
| n_{query} | queries per iteration | 100 | | | |
| p_{fake} | proportion of fake queries | 0.5 | | | |
| σ | sampling variance | 0.001 | | | |
| τ | HessAware parameter | 1 | | | |
| k | QOI estimation iterations | 1 | | | |
| μ | dispersion coefficient | 0.1 | 0.1 | 0.3 | 0.3 |
| κ | inference confidence | | | 0.6 | |

Table 1. Default setting of key parameters.

5.2 A CTF Game

To evaluate the performance of AdvMind (the defender) against different black-box attacks and their variants (the attacker). We define the following capture-the-flag (CTF) game.

Setting. Given the backend DNN f , the attacker randomly selects one input x_o from the testing set and a target class $c_* \neq f(x_o)$, and aims to generate an adversarial input x_* that satisfies (i) $x_* \in \mathcal{B}_\epsilon(x_o)$ and (ii) $f(x_*) = c_*$. At the k -th iteration of the attack, the attacker estimates the gradient with respect to the current input $x^{(k-1)}$ and then perturbs $x^{(k-1)}$ to generate $x^{(k)}$; meanwhile, the AdvMind attempts to infer the attacker’s target class c_* based on the previous queries. We define that the attacker wins the game if she successfully generates the adversarial input x_* before AdvMind is able to correctly infer c_* with high confidence, and AdvMind wins the game

if she correctly infers c_* before the attacker generates x_* . Note that if none of them succeeds, it is a “draw”.

Observations. We measure up to each iteration the percentage of games won by the attacker or the defender, under different configurations for the attacker and the defender. The results on different datasets and DNNs are illustrated in Figure 5 and Figure 6, from which we have the following key observations. Note that we aim at early detection so we only focus on the first 10 iterations, during which most attacks haven’t succeeded, especially for VGGface2.

O1: Basic AdvMind is able to infer the attacker’s intent in non-adaptive attacks to a large extent. AdvMind defeats the attacker in around 50% of the games up to the 10-th iteration on CIFAR10 and CIFAR100, and in around 90% of the games on ISIC and VGGface2. Its better performance on higher-dimensional datasets may be attributed to the larger divergence of gradient directions among different classes in higher-dimensional spaces. Yet, the success rate of AdvMind increases gradually with the number of iterations, indicating that the passive inference requires a large number of observations to make reliable estimation. Furthermore, basic AdvMind is highly sensitive to the fake queries injected by the attacker. Under the adaptive attacks, AdvMind fails to correctly infer the attacker’s target before the attacker successfully generates the adversarial examples. The detection rate drops around 10% on CIFAR10 and CIFAR100 and about 30% on ISIC and VGGface2. This may be explained by that basic AdvMind relies on naïve mean estimation, which tends to significantly deviate from the real value, under the influence of the injected noise. This deviation is especially evident for higher-dimensional data (e.g., ISIC and VGGface2).

O2: The robust intent estimation of AdvMind^R effectively mitigates the influence of fake queries by focusing on the subset of coherent queries to estimate the query of interest. Meanwhile, at each iteration, the attacks attain lower success rates than the case of non-adaptive attacks against basic AdvMind, which is explained by the smaller number of queries useful for gradient estimation.

O3: The proactive intent solicitation of AdvMind^{RP} further significantly improves the inference accuracy, especially during the early stages of the attacks. Across most cases, by the 3rd iteration, AdvMind^{RP} is capable of identifying the attacker’s true target with approximately 75% accuracy even against the adaptive attacks. Meanwhile, because of the dominant performance of AdvMind^{RP}, the attacker’s winning percentage is kept close to 0 during the first 10 iterations under all the settings.

O4: Overall, AdvMind’s performance seems agnostic to different datasets, DNNs, or attacks. In comparison, the attacker’s effectiveness varies significantly with concrete settings. For instance, at the 10-th iteration, the winning percentages of most adaptive attacks on ISIC2018 decrease by around 12.5% compared with their counterparts on CIFAR10. This variations may be explained by that the adaptive attacks require more iterations to craft adversarial examples under black-box setting in high-dimensional space.

5.3 Impact of Key Factors

Next we evaluate the impact of different factors on AdvMind’s performance, which shed light on the optimal operation of AdvMind. By default, we use the robust+proactive variant of AdvMind (AdvMind^{RP})




Figure 5: Winning percentages of the attacker and AdvMind in the CTF competition on first order methods with respect to different datasets and DNN models.




Figure 6: Winning percentages of the attacker and AdvMind in the CTF competition on the second order method with respect to different datasets and DNN models.

in our evaluation. Unless noted otherwise, all the parameters are set with their default values in Table 1. To evaluate the inference of AdvMind and the attack effectiveness, we evaluate the inference success rate and the attack success rate up to the 10-th iteration. Note that to factor out the mutual influence, here we evaluate the attacker’s and AdvMind’s success rates independently.

Number of Queries per Iteration n_{query} . We first measure the impact of the number of queries issued by the attacker per iteration (n_{query}). Figure 7 illustrates the attacker’s and AdvMind’s success rates on different datasets and DNNs as n_{query} varies from 20 to 200. Note that we keep the proportion of fake queries p_{fake} constant ($p_{\text{fake}} = 0.5$) in all the experiments here.

As expected, across all the settings, a large number of queries allow the attacker to generate adversarial inputs more effectively. For instance, on CIFAR10, as n_{query} grows from 20 to 200, the attack success rate increases from about 0.1 to 0.4. At the same time, with more queries, AdvMind is also enabled to estimate the attacker’s query of interest more accurately, leading to higher inference accuracy. For instance, on CIFAR10, the inference accuracy increases from around 0.7 to 1.0. Therefore, the attacker faces the dilemma between attack effectiveness and intent exposure.

Proportion of Fake Queries p_{fake} . With $n_{\text{query}} = 100$ fixed, we evaluate the impact of the proportion of fake queries p_{fake} on the attacker’s and AdvMind’s performance (Figure 8).

It is observed that for both attacker and AdvMind, there exists a critical threshold of p_{fake} : once p_{fake} exceeds this threshold, the attack/inference success rate sharply decreases. For instance, on CIFAR10, if $p_{\text{fake}} \geq 0.6$, the success rate of AdvMind drops to around 0.1; once $p_{\text{fake}} \geq 0.6$, the attack success rate drops to around 0. Interestingly, the accuracy drop is larger for AdvMind than that for the attacker. This may be intuitively explained as follow: p_{fake} mainly affects the attacker’s estimation about the gradient information and AdvMind’s estimation about the query of interest; yet, as AdvMind needs first to differentiate fake and true queries (which is set by the attacker), p_{fake} tends to have a larger impact on AdvMind. Thus, to hide her intent, the adversary may set p_{fake} sufficiently large (e.g., $p_{\text{fake}} \geq 0.6$), which however significantly increases the attack cost (number of queries) and reduces the attack success rate.

Dispersion Coefficient μ . For AdvMind, the parameter μ controls the weight of the dispersion matrix M in Eqn (16). A large μ implies that the gradient directions with respect to different classes tend to be approximately orthogonal to each other, therefore highly indicative of the adversary’s target.

Figure 9 shows the attack and inference success rates as functions of the dispersion coefficient. It is noted that across all the settings, AdvMind’s effectiveness increases with μ , while the attack’s success rate drops accordingly. This is due to that (i) the orthogonality of the gradient directions increases with μ , making it easy for AdvMind to identify the adversary’s gradient descent direction, and (ii) meanwhile the gradient estimation error of the adversary grows with μ , making it difficult to find adversarial inputs.

However, it is worth pointing out that as μ increases, the estimated gradient using the query results becomes less and less informative for the adversary to craft adversarial inputs. Therefore, the adversary may abandon the estimate as μ reaches a critical




Figure 7: Attack/inference success rate as a function of the number of queries per iteration n_{query} .




Figure 8: Attack/inference success rate as a function of the proportion of fake queries p_{fake} .




Figure 9: Attack/inference success rate as a function of the dispersion coefficient μ .

threshold. To incentivize the adversary to estimate the gradient, the defender should keep μ small (e.g., $\mu = 0.1$) in AdvMind.

6 RELATED WORK

Due to their use in security-critical domains, deep neural network (DNN) models are increasingly becoming the targets of malicious attacks. Most existing work focuses on the white-box setting, in which the adversary has full access to the model information (e.g., its architecture and parameters). One line of work develops new evasion attacks against DNN models [11, 22]. Another line of work attempts to improve DNN resilience against such attacks by inventing new training and inference strategies [21]. Yet, such defenses are often circumvented by even powerful attacks [5] or adaptively engineered adversarial inputs [1], resulting in a constant arms race between adversaries and defenders [18].

The popularity of publicly accessible predictive APIs [34] has spurred research on black-box adversarial attacks in which the adversary has limited or no knowledge about the target DNN. Early attack methods leverage the property of transferability [32]: certain adversarial inputs crafted against one DNN model are effective against another model; thus the adversary can attack a surrogate model and apply them on the target model [20]. To defend against

such attacks, the method of ensemble adversarial training [33] has been proposed recently, which trains a DNN model using data augmented with adversarial inputs crafted on other models.

Recent work notes that adversarial inputs for surrogate do not always transfer to the target model, especially when conducting targeted attacks [6, 24]. This line of attacks instead constructs adversarial inputs by estimating the gradient through the target DNNs with coordinate-wise finite difference methods [14, 19, 36]. The research on defending against query-based black-box attacks is still limited. The recent work [7] proposes a stateful method to detect query-based black-box attacks by measuring the relationships (e.g., similar but non-identical) of a sequence of queries.

This work complements the existing work by focusing on identifying the adversary’s intent during an early stage of the attack, with new findings to mitigate black-box adversarial attacks.

7 DISCUSSION AND CONCLUSION

In this paper, we present AdvMind, a new class of models for inferring adversary intent in black-box adversarial attacks. Combining robust intent estimation and proactive intent solicitation, AdvMind is able to reliably identify the adversary’s query of interest and accurately detect the adversary’s target class in an early stage of the

attack, which facilitates to deploy proper mitigation strategies and to perform prompt remediation against such threats. The empirical evaluation with respect to benchmark datasets, popular DNNs, and state-of-the-art attacks validates the efficacy of AdvMind.

This work also opens several avenues for further research. First, we assume that a single adversary launches the attack and thus can be easily detected [7]. In practice, the attack can be performed by multiple colluding adversaries. It is critical to extend AdvMind to such complicated settings. Second, we focus on query-based attacks. There are other types of black-box attacks (e.g., surrogate model [20]). As they do not require query access during drafting adversarial inputs, AdvMind is unable to infer the adversary’s target in such attacks. Finally, after identifying the adversary’s intent, the next step is to perform effective mitigation. It is beneficial to integrate AdvMind with model enhancement methods (e.g., adversarial training [22]) to provide end-to-end protection.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1910546, 1953813, and 1846151. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. Shouling Ji was partly supported by NSFC under No. U1936215, 61772466, and U1836202, the National Key Research and Development Program of China under No. 2018YFB0804102, the Zhejiang Provincial Natural Science Foundation for Distinguished Young Scholars under No. LR19F020003, the Zhejiang Provincial Key R&D Program under No. 2019C01055, and the Ant Financial Research Funding. Xiapu Luo was partly supported by HK RGC Project (PolyU 152239/18E) and HKPolyU Research Grant (ZVQ8).

REFERENCES

[1] Anish Athalye, Nicholas Carlini, and David Wagner. 2018. Obfuscated Gradients Give a False Sense of Security: Circumventing Defenses to Adversarial Examples. In *Proceedings of IEEE Conference on Machine Learning (ICML)*.

[2] Jeremy Bernstein, Yu-Xiang Wang, Kamyar Azizzadenesheli, and Animashree Anandkumar. 2018. signSGD: Compressed Optimisation for Non-Convex Problems. In *Proceedings of IEEE Conference on Machine Learning (ICML)*.

[3] Battista Biggio, Giorgio Fumera, Fabio Roli, and Luca Didaci. 2012. Poisoning Adaptive Biometric Systems. In *Proceedings of Joint IAPR International Workshop on Structural, Syntactic, and Statistical Pattern Recognition (SSPR&SPR)*.

[4] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. 2018. Vggface2: A dataset for recognising faces across pose and age. In *13th IEEE International Conference on Automatic Face & Gesture Recognition*.

[5] Nicholas Carlini and David A. Wagner. 2017. Towards Evaluating the Robustness of Neural Networks. In *Proceedings of IEEE Symposium on Security and Privacy (S&P)*.

[6] Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. 2017. Zoo: Zeroth Order Optimization based Black-Box Attacks to Deep Neural Networks without Training Substitute Models. In *Proceedings of ACM Workshop on Artificial Intelligence and Security (AISeC)*.

[7] Steven Chen, Nicholas Carlini, and David Wagner. 2019. Stateful Detection of Black-Box Adversarial Attacks. *ArXiv e-prints* (2019).

[8] Niles Dalvi, Pedro Domingos, Mausam, Sumit Sanghai, and Deepak Verma. 2004. Adversarial Classification. In *Proceedings of ACM International Conference on Knowledge Discovery and Data Mining (KDD)*.

[9] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A Large-scale Hierarchical Image Database. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

[10] Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. 2017. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 542, 7639 (2017), 115–118.

[11] Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and Harnessing Adversarial Examples. In *Proceedings of International Conference on Learning Representations (ICLR)*.

[12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

[13] Peter J. Huber. 1964. Robust Estimation of a Location Parameter. *Ann. Math. Statist.* 35, 1 (1964), 73–101.

[14] Andrew Ilyas, Logan Engstrom, Anish Athalye, and Jessy Lin. 2018. Black-box Adversarial Attacks with Limited Queries and Information. In *Proceedings of IEEE Conference on Machine Learning (ICML)*.

[15] Alex Krizhevsky and Geoffrey Hinton. 2009. Learning Multiple Layers of Features from Tiny Images. *Technical report, University of Toronto* (2009).

[16] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521, 7553 (2015), 436–444.

[17] Bin Liang, Miaoqiang Su, Wei You, Wenchang Shi, and Gang Yang. 2016. Cracking Classifiers for Evasion: A Case Study on the Google’s Phishing Pages Filter. In *Proceedings of International Conference on World Wide Web (WWW)*.

[18] X. Ling, S. Ji, J. Zou, J. Wang, C. Wu, B. Li, and T. Wang. 2019. DEEPSEC: A Uniform Platform for Security Analysis of Deep Learning Model. In *Proceedings of IEEE Symposium on Security and Privacy (S&P)*.

[19] Sijia Liu, Pin-Yu Chen, Xiangyi Chen, and Mingyi Hong. 2019. signSGD via Zeroth-Order Oracle. In *Proceedings of International Conference on Learning Representations (ICLR)*.

[20] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. 2016. Delving into Transferable Adversarial Examples and Black-Box Attacks. *ArXiv e-prints* (2016).

[21] Xingjun Ma, Bo Li, Yisen Wang, Sarah M. Erfani, Sudanthi Wijewickrema, Grant Schoenebeck, Dawn Song, Michael E. Houle, and James Bailey. 2018. Characterizing Adversarial Subspaces Using Local Intrinsic Dimensionality. In *Proceedings of International Conference on Learning Representations (ICLR)*.

[22] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *Proceedings of International Conference on Learning Representations (ICLR)*.

[23] S. Moosavi-Dezfooli, A. Fawzi, O. Fawzi, and P. Frossard. 2017. Universal Adversarial Perturbations. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

[24] Nina Narodytska and Shiva Prasad Kasiviswanathan. 2017. Simple Black-Box Adversarial Perturbations for Deep Networks. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

[25] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. 2017. Practical Black-Box Attacks Against Machine Learning. In *Proceedings of ACM Symposium on Information, Computer and Communications Security (AsiaCCS)*.

[26] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. In *Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

[27] Pang Ren, Zhang Xinyang, Ji Shouling, Luo Xiapu, and Wang Ting. 2020. AdvMind: Inferring Adversary Intent of Black-Box Attacks. *ArXiv e-prints* (2020).

[28] Orbis Research. 2019. Healthcare Fraud Detection Market to grow at 24.59% CAGR by 2024. <https://www.globenewswire.com/>.

[29] Peter J. Rousseeuw and Sabine Verboven. 2002. Robust estimation in very small samples. *Computational Statistics & Data Analysis* 40, 4 (2002), 741–758.

[30] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. 2016. Mastering the Game of Go with Deep Neural Networks and Tree Search. *Nature* 5787 (2016), 484–489.

[31] Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *Proceedings of International Conference on Learning Representations (ICLR)*.

[32] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2014. Intriguing Properties of Neural Networks. In *Proceedings of International Conference on Learning Representations (ICLR)*.

[33] F. Tramèr, A. Kurakin, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel. 2018. Ensemble Adversarial Training: Attacks and Defenses. In *Proceedings of International Conference on Learning Representations (ICLR)*.

[34] Florian Tramèr, Fan Zhang, Ari Juels, Michael K. Reiter, and Thomas Ristenpart. 2016. Stealing Machine Learning Models via Prediction APIs. In *Proceedings of USENIX Security Symposium (SEC)*.

[35] Daan Wierstra, Tom Schaul, Tobias Glasmachers, Yi Sun, Jan Peters, and Jürgen Schmidhuber. 2014. Natural Evolution Strategies. *J. Mach. Learn. Res.* 15, 1 (2014), 949–980.

[36] Haishan Ye, Zhichao Huang, Cong Fang, Chris Junchi Li, and Tong Zhang. 2018. Hessian-Aware Zeroth-Order Optimization for Black-Box Adversarial Attack. *ArXiv e-prints* (2018).