

MalProtect: Stateful Defense Against Adversarial Query Attacks in ML-based Malware Detection

Aqib Rashid, Jose Such

Abstract—ML models are known to be vulnerable to adversarial query attacks. In these attacks, queries are iteratively perturbed towards a particular class without any knowledge of the target model besides its output. The prevalence of remotely-hosted ML classification models and Machine-Learning-as-a-Service platforms means that query attacks pose a real threat to the security of these systems. To deal with this, stateful defenses have been proposed to detect query attacks and prevent the generation of adversarial examples by monitoring and analyzing the sequence of queries received by the system. Several stateful defenses have been proposed in recent years. However, these defenses rely solely on similarity or out-of-distribution detection methods that may be effective in other domains. In the malware detection domain, the methods to generate adversarial examples are inherently different, and therefore we find that such detection mechanisms are significantly less effective. Hence, in this paper, we present MalProtect, which is a stateful defense against query attacks in the malware detection domain. MalProtect uses several threat indicators to detect attacks. Our results show that it reduces the evasion rate of adversarial query attacks by 80+% in Android and Windows malware, across a range of attacker scenarios. In the first evaluation of its kind, we show that MalProtect outperforms prior stateful defenses, especially under the peak adversarial threat.

Index Terms—Adversarial machine learning, Malware detection, Machine learning security, Deep learning

I. INTRODUCTION

ML has offered enormous capabilities, leading to the large-scale use of Machine-Learning-as-a-Service (MLaaS) [1], which allows users to leverage remotely-hosted ML models by requesting predictions from them [2]. The widespread use of these services means that their reliability and security is paramount, especially considering the threat posed by adversarial machine learning [3], [4], [5]. Recent work has shown that these systems are vulnerable to adversarial query attacks (e.g., [6], [7], [8]). With query attacks, an attacker iteratively perturbs a malware sample based on feedback from the target model that includes its predicted output. Using the feedback, the perturbations to the input sample are tuned across the queries to observe how the target model responds until, eventually, the truly malicious sample is classified as benign (i.e., goodware) [7], [8]. To deal with the adversarial ML threat, several defenses have been proposed based on a variety of approaches (e.g., [3], [9], [10], [11]). However, most of these approaches have been shown to be ineffective against *query* attacks in several domains [12], [6], [7], including ML-based malware detection [5], [8], [13].

Recent work has stressed that systems must monitor queries to identify hazards such as adversarial attacks [14]. In fact, ML-based anomaly detection has been regarded as essential to detect the misuse of ML-based systems [15]. To this end, *stateful defenses* have been proposed to protect ML prediction models against query attacks [16] by offering greater system awareness. This is achieved by monitoring queries received by the system. Researchers have hypothesized that sequences of adversarial queries are often abnormally similar to each other, unlike sequences of legitimate queries [16], [17], [18], and often do not fit their distribution [18], [19], [20]. Hence, stateful defenses analyze the sequence of queries made to the system, such as the similarity of feature vectors representing an image. Several stateful defenses have been proposed (e.g., [16], [17], [21], [18], [22]); however, prior to our work, they have not been tested in the malware detection domain, where query attacks can cripple prediction models [5], [8]. The malware detection domain is significantly different from other domains, with more constraints imposed on attackers regarding the discrete representation of feature vectors and the preservation of malicious functionality when generating adversarial examples [5], [23]. Therefore, attackers use techniques that are different from those in other domains to generate adversarial examples [8], [13]. Consequently, stateful defenses that have been applied to other domains may, in fact, be ineffective when tested in the malware detection domain.

Hence, in this paper, we present *MalProtect*, which is a model-agnostic stateful defense against query attacks in the ML-based malware detection domain. MalProtect employs several threat indicators that use different data-driven statistical and ML-based methods to analyze the sequence of queries received by the system. Upon receiving a new query and conducting an analysis, the threat indicators each produce scores reflecting the likelihood of an attack based on different criteria. These are then aggregated to predict if there is an attack in progress. This allows MalProtect to detect attacks in a manner that is reliable and interpretable. Across Android and Windows, we show that MalProtect reduces the evasion rate of query attacks by 80+%, and we compare this to the meager performance of prior stateful defenses across a range of attack scenarios in the first evaluation of its kind. We further show that, despite being designed to protect against query attacks, MalProtect affords a degree of protection against other types of attacks, such as transferability attacks. Furthermore, we demonstrate that an adaptive attacker — with complete knowledge of MalProtect — cannot achieve significant evasion. In summary, we make two key contributions:

Both authors are with the Department of Informatics, King's College London, Strand, London WC2R 2LS, United Kingdom. Jose Such is also with VRAIN, Universitat Politècnica de València, Spain.
E-mail: {aqib.rashid, jose.such}@kcl.ac.uk

- 1) We propose the first stateful defense against adversarial query attacks in ML-based malware detection. Our defense uses several threat indicators to prevent query attacks and the generation of adversarial examples.
- 2) We provide the first evaluation of several prior stateful defenses applied to the ML-based malware detection domain. We show that MalProtect reduces the evasion rate of query attacks across Android and Windows under a range of threat models and attacker scenarios.

The rest of this paper is organized as follows: Section III provides the background and related work. Section III defines the threat model. Section IV describes MalProtect. Section V details the experimental setting. In Sections VI-X we present our experimental results. In Section XI we conclude.

II. BACKGROUND & RELATED WORK

ML-based Malware Detection. ML-based malware detection has rapidly grown in popularity. It allows for unseen and unknown threats to be discovered and offers performance that surpasses other detection approaches [5]. ML-based malware detection classifiers decide if a query is benign (i.e., goodware) or malware. These models are trained on binary feature vectors representing the presence or absence of features in executables [5], [24]. The quality of the predictions is determined by these features, which can include the usage of API calls, the network addresses accessed, or the libraries used [25], [26], [5], [27]. **Adversarial ML Attacks.** However, ML models are vulnerable to adversarial ML attacks. Adversarial examples can be developed by an attacker, which are queries that are designed to evade a classifier. Even without direct access to the target model, an attacker can perturb a malware sample to have it classified as benign and evade the prediction model through *transferability attacks* and *query attacks*. With transferability attacks, an estimation of the target model is developed and attacked in anticipation that the generated adversarial examples will *transfer to* the target model [4]. This is based on the transferability property of adversarial examples [3].

Meanwhile, query attacks generate an adversarial example by iteratively perturbing an input sample towards the desired class based on feedback received from the target model until evasion is achieved [5], [7]. This attack can be conducted in a black-box manner without any knowledge about the target model besides the predicted outputs. In image recognition, recent work has introduced several query attacks using techniques such as gradient and decision boundary estimation that perturb the continuous feature-space of the domain [6], [7], [12]. However, in the malware detection domain, the functionality of software executables must be preserved, and discrete feature vectors must be used when generating adversarial examples [24], [5], [8]. For example, a feature representing an API call (e.g., `CreateFile()`) cannot be perturbed continuously by an attack (e.g., `CreateFile() + 0.05`); rather, a different feature must be used [8] that offers the same functionality. Therefore, software transplantation-based techniques have been proposed to achieve evasion that use features from benign samples to perturb a malware sample [8], [13].

Defenses Against Adversarial ML. Many defenses against adversarial ML have been proposed (e.g., [3], [9], [10], [11]). Most approaches are single-model defenses, where an individual model is made robust [28]. Ensemble defenses have also been proposed, which use several models with some particular method to return predictions [29], [30], [11], [31]. However, recent studies (e.g., [32], [33], [5]) have shown that most defenses are ineffective against query attacks. As these defenses are *stateless*, they cannot analyze the sequence of queries received and prevent query attacks.

Stateful Defenses. To defend against query attacks, *stateful defenses* have been proposed (e.g., [16], [17], [18], [19], [20], [34]). Stateful defenses conduct a form of anomaly detection, which has been recommended in recent studies as a method to combat adversarial query attacks [14], [15], [35]. It has been hypothesized that because query attacks iteratively perturb an input sample, they produce a sequence of abnormal queries that have high similarity or are outside the distribution of legitimate queries. Figure 1 shows that stateful defenses retain the sequence of queries received by the system and analyze them using different techniques to detect attacks. This is similar to intrusion detection systems and firewalls that monitor network activity to detect threats [36].




Fig. 1: Overview of a stateful defense.

However, stateful defenses that protect (remotely-hosted) ML models are different from traditional API protection mechanisms. To protect remotely-hosted services, traditional API protection mechanisms include simple rate-limiting, encryption, and input validation. Conversely, stateful defenses analyze the sequence of queries received by the system by examining their similarity, their distribution, and other factors before queries are passed to the underlying prediction model that is under protection. This means that attacks can be detected before they can cause harm to the ML model.

Limitations of Existing Stateful Defenses. A number of stateful defenses have been proposed in recent years [16], [18], [19], [21], [17]. Subsequent work has identified several limitations of these defenses, including insufficient similarity detection mechanisms, poor maintenance of the query history, and weaknesses of account-based detection, among others [17], [22], [21], [14]. In Table I we provide an overview of stateful defenses, covering the different approaches that have been employed in prior work. Subsequently, we provide detail about each of these approaches, discussing how they detect attacks and any key limitations.

More specifically, Chen et al.'s Stateful Detection approach (SD) is an account-based defense that maps a user's query to a feature vector and measures the average distance between it and its k-nearest neighbors. If this is below a threshold, an attack is detected, which resets the query history and cancels

	Chen et al. (SD) [16]	Li et al. (Blacklight) [17]	Azmoodeh et al. [34]	Kariyappa et al. [37]	Atli et al. [20]	Juuti et al. (PRADA) [18]
Maintains query history	✓	✓	✓			✓
Distance-based / Similarity detection	✓	✓	✓			✓
Account-based defense	✓		✓			✓
Out-of-distribution (OOD) detection				✓	✓	✓
OOD based on training data				✓	✓	

TABLE I: Overview of stateful defenses and attack detection for adversarial ML.

the user’s account. However, this can be defeated by a Sybil attack, where an attacker uses multiple accounts in the attack [38], [39]. Moreover, inspecting queries per-account limits the detection scope [17]. Alternatively, Li et al. present Blacklight [17], which measures the similarity of queries for all users using the L_2 distance (i.e., Euclidean distance), which is a type of L_p norm [40]. An attack is detected if the Euclidean distance between several queries (represented by hashes) is below the threshold. Other distance-based defenses have been proposed (e.g., [34]), though these can all be evaded when queries are intentionally designed to be dissimilar [14] or through query-blinding attacks [21].

Out-of-distribution (OOD) detectors have also been proposed to thwart query attacks and detect adversarial examples [35]. These check whether a query belongs to the distribution of the target model’s training data [35], [19], [20]. When OOD queries are detected, Kariyappa et al. propose returning inconsistent labels to those detected OOD queries [19]. However, an attacker may be able to construct adversarial queries that remain within the required distribution [21], [2]. PRADA [18] is a defense against model extraction attacks, which is a related problem. Prior work has found that it can also detect evasion attacks [16]. PRADA is based on the assumption that the L_2 distance among non-adversarial queries follows a normal distribution. This is monitored using the Shapiro-Wilk normality test. However, manipulation of the query distribution has been found to evade it [16], [2].

Our defense, MalProtect, is model-agnostic allowing it to be used with any underlying prediction model. To detect attacks, it utilizes techniques beyond basic similarity and OOD detection, such as analyzing the autoencoder loss of queries, the distribution of enabled and shared features across queries, and more. We next present the threat model we consider for our work, followed by a detailed description of MalProtect.

III. THREAT MODEL

In our work, attackers aim to evade a feature-based ML malware detection classifier so that their malicious queries are misclassified as benign. This is a well-established threat model in this domain [41], [42], [43], [25], [24], [23].

Target Model. The target model is a remotely-hosted malware classification system that predicts whether an input sample belongs to the benign or malware class. We refer to the classification system as the oracle O . The oracle has two principle components: the underlying classifier F (i.e., the prediction model) that predicts whether an input sample is malicious and the stateful defense that protects F by analyzing queries to

detect attacks (e.g., MalProtect — see Section IV later). In our work, we use a range of prediction models from prior work (see Section V later). Generally, to train such a prediction model F , input samples are represented as binary feature vectors using the features that are provided by a dataset. With the features $1..M$, a feature vector X can be constructed for each input sample such that $X \in \{0, 1\}^M$, similar to previous work [30], [44], [24]. The presence or absence of a feature i is represented by either 1 or 0 within X . Like prior work on ML-based malware detection [8], [24], [45], the features we employ in our work include, among others, libraries, API calls, permissions, and network addresses; these are provided by the datasets we use (see Section V later). With their class labels, several feature vectors can be used to train the binary classification model that we refer to as the classifier F . Then, when a user makes a query to O , a prediction is made and returned. For the predicted outputs, we use 1 for the malware class and 0 for the benign class. O returns scoreless feedback [12]. Recall that O consists of the prediction model and the independent stateful defense designed to protect it.

Attacker’s Goal. The goal of the attacker is to generate an adversarial example X' from a malware sample X to evade the oracle O and obtain a benign prediction. Suppose $O: X \in \{0, 1\}^M$ and we have a function $check()$ to check the malicious functionality of X . We can summarize this as:

$$check(X) = check(X'); O(X) = 1; O(X') = 0 \quad (1)$$

Other variations of this goal do exist (e.g., achieving evasion with minimal queries [46]). However, we focus on the more general goal above, as it is commonly explored in related work on stateful defenses and ML-based malware detection.

Attacker Capabilities & Knowledge. In order to evaluate MalProtect and other stateful defenses under different attack scenarios and conditions, we model three types of attackers that have been featured in prior work [12], [47], [48], [23] as is common in prior work. We consider the *gray-box attacker*, who has limited knowledge, has access to the same training data as the classifier F and is aware of the feature representation and the statistical representation of the features across the dataset. However, the attacker is unaware of F ’s parameters, configurations, and other pertinent information. This represents a scenario similar to when some model data has been leaked. To perform query attacks, this attacker applies perturbations to the feature vector through a software transplantation-based approach [8], [13], guided by their knowledge of the dataset. In contrast, the zero-knowledge *black-box attacker* can only observe the predictions received for their queries and has no information about the target model but does have some information pertaining to the kind of feature extraction performed (e.g., the static analysis that a malware detection classifier may consider). This attacker also uses a software transplantation-based approach, but with considerably less information. Neither attacker is aware that MalProtect is a part of O . We also consider an *adaptive attacker* (i.e., a *white-box attacker*) who has complete knowledge of the target model and knows that MalProtect is in place. Therefore, this attacker tries to evade all components of the oracle O with a specific, tailored attack.

We assume that all attackers possess ample computing power and resources to submit thousands of queries. Furthermore, unlike previous work [18], [16], we assume that attackers may control multiple user accounts and IP addresses, as in a Sybil attack [38], [39]. As attackers can submit many queries, they may be able to conduct a transferability attack by constructing substitute models [4], so that a black-box attack becomes a semi white-box one. However, we mainly focus on query attacks in this work, as stateful defenses are designed to primarily defend against them [16]. Moreover, were MalProtect to be deployed in a production environment, it would be unreasonable to assume that basic rate-limiting protections were not in place to prevent excessive queries by users, making such transferability attacks less practical.

IV. MALPROTECT

MalProtect is designed to protect ML models against adversarial query attacks in the malware detection domain’s feature-space. Figure 2 shows that MalProtect forms a layer of protection that monitors queries before they reach the prediction model F , which is the ML-based malware detection classifier to be protected. This allows MalProtect to detect adversarial queries that may evade the classifier (i.e., prediction model). The prediction model F can be a vanilla ML model or a defense (e.g., an adversarially-trained model, an MTD, etc.). By having a defense as the prediction model used with MalProtect, it can offer *defense in depth*, as multiple layers of security can contribute to the overall security of the system.

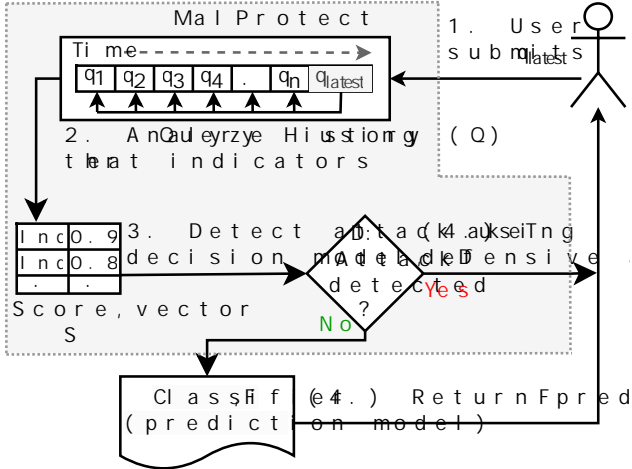


Fig. 2: Overview of MalProtect.

When a user submits a query (Step 1), MalProtect employs several *threat indicators* to analyze the sequence of queries for attack detection (Step 2). This is similar to a system that makes decisions using data from multiple sensors [14]. The threat indicators use different criteria, principally driven by an anomaly detection-based approach. This includes examining aspects such as the autoencoder loss of queries, the distance between queries, the features shared across queries, and whether other characteristics of queries fit within different distributions. Based on the state of the query history, each indicator produces a score that reflects the likelihood of an attack in progress according to that indicator, with adversarial queries expected to cause higher scores. A *decision model* then takes the indicator

scores as its input and predicts whether there is an attack in progress (Step 3). If an attack is detected, MalProtect takes some defensive action. If no attack is detected, the prediction model F returns a prediction for the user’s query (Step 4). Not only do the scores assist in interpreting how MalProtect produces a particular decision, but as we show later, each threat indicator can be analyzed to determine its influence on the final decision. Next, we provide details about each of MalProtect’s core components.

A. Query History

MalProtect analyzes the query history to detect adversarial behavior. When a user requests a prediction for their query q_{latest} , it is appended to the query history Q . Queries can be represented more sparsely in the malware detection domain to reduce storage overhead, which has been cited as an issue in other domains [17]. The depth of analysis performed by MalProtect and the size of Q are controlled by the resources possessed by the defender when MalProtect is deployed, but, as we show later in Section X, MalProtect needs very little spatial and computational resources to offer a very good performance. Importantly, the size of Q will obviously never be infinite [49], [50], though it can extend as far as the defender’s resources allow. Unlike other stateful defenses that reset the query history once they reach capacity (e.g., [16]), MalProtect retains queries as a sliding window to maximize its detection capability. Finally, MalProtect does not distinguish between queries from different users, as an attacker can conduct a Sybil attack to evade account-based defenses [38], [39].

B. Indicators for Analyzing Query History

After adding q_{latest} to Q , MalProtect analyzes the query history using multiple *threat indicators* to detect attacks through an anomaly detection-based approach. In conventional anomaly detection, thresholds are employed to assess whether a sample meets the criteria for being abnormal. However, with MalProtect, we use a data-driven approach where the weights of each indicator are learned by a *decision model* to predict an attack. That is, instead of each indicator deciding if an attack is taking place, each indicator produces a score between 0 and 1 that reflects the likelihood of an attack according to its criteria. Scores are then aggregated and passed to a decision model that decides if an attack is in progress. This allows for a degree of interpretability too, as we can understand which indicators play the biggest role in the decision model and which indicators produce the highest scores (giving clues about what is anomalous), as we show later. MalProtect is based on the principle that adversarial queries are expected to cause higher indicator scores as they are anomalous compared to legitimate queries and the training data distribution.

To conduct the analysis and produce scores, some indicators use standard anomaly detection techniques, such as the empirical rule [51], to assess whether a query fits the distribution of data. Meanwhile, other indicators use information about the training data to examine how much queries deviate from various measures. We next present a detailed description of each indicator:

Indicator 1. Distance between queries: This indicator assesses the similarity of q_{latest} to other queries in Q , which is the

primary method adopted by some stateful defenses (e.g., [16], [17]). A smaller distance between queries is an indication of iterative perturbations during a query attack, as an attacker makes modifications to their queries to assess how the target model responds [16], [7].

This indicator therefore examines how anomalous the minimum distance between q_{latest} and other queries in Q is (represented by $minDistQ$). As the malware detection domain uses binary feature vectors (unlike other domains, e.g., image recognition), the L_0 distance is the most appropriate distance metric [40], [44], as it measures the number of features that are different between two feature vectors. Other norms, such as L_∞ would always have a distance of 1 if at least one change was present between two feature vectors [44]. To assess the potential abnormality of $minDistQ$ without using a threshold and also to normalize it between 0 and 1, we put $minDistQ$ in the context of the average L_0 distance of all samples in the training data (represented by $avgDistD$). If $minDistQ$ is significantly below this, it implies that q_{latest} is similar to another query in Q beyond the norm.

The score for this indicator is based on the percentage change between $minDistQ$ and $avgDistD$. A smaller percentage change means that there exists a query whose similarity to q_{latest} is significantly lower than the average of the training data ($avgDistD$). So that a higher threat is represented by a higher score value, we negate the result. Thus, the score for this indicator is calculated as follows:

$$S_1 = -1 \cdot (minDistQ - avgDistD) / avgDistD \quad (2)$$

Indicator 2. Enabled features shared across queries: This indicator assesses whether q_{latest} shares a significant number of enabled features with another query in Q beyond the norm. The rationale behind this indicator is that an attacker could enable unused features in q_{latest} to intentionally increase the L_0 distance between queries in order to evade basic similarity detection (e.g., while trying to evade Indicator 1). This is an attack strategy specific to the malware detection domain. That is, the query distance can be increased (e.g., by adding features in bulk) without disrupting the original malicious functionality (e.g., through dead-code, opaque predicates). Conversely, in domains such as image recognition, this attack strategy would not be as useful since adding features to deliberately increase the query distance in this manner would lead to the original image being distorted.

Intriguingly, though, some core features of malicious queries must *always* remain enabled to preserve the original malicious functionality across a query attack. Thus, queries with the same enabled features (and therefore shared) would imply similarity specifically in this domain. Therefore, the highest number of enabled features shared between q_{latest} and another query in Q is calculated as $maxSharedQ$.

We then use the average number of enabled features shared between samples in the training data ($avgSharedD$) to assess to what extent $maxSharedQ$ may be abnormal. This is achieved by calculating the deviation of $maxSharedQ$ from this mean as a percentage change. A higher value indicates that q_{latest} shares an abnormally high number of enabled features with some other query in Q considering the training data’s

distribution. The score is calculated as follows:

$$S_2 = (maxSharedQ - avgSharedD) / avgSharedD \quad (3)$$

Indicators 3A & 3B. Number of enabled features: An attacker could also rapidly traverse the decision boundary by enabling a substantial number of benign features in q_{latest} at once. That is, by adding many features to a malware sample that are present in benign executables, a misclassification can be achieved. Therefore, an anomalous number of enabled features could indicate an attack attempt. Hence, the Indicator 3A is the percentage change between the number of enabled features in q_{latest} ($|q_{latest}|$) and the average number of enabled features in training data samples ($avgFeaturesD$). A large percentage change implies that far more features are enabled in q_{latest} compared to $avgFeaturesD$. The score is calculated as follows:

$$S_{3A} = (|q_{latest}| - avgFeaturesD) / avgFeaturesD \quad (4)$$

However, the Indicator 3A only assesses whether $|q_{latest}|$ is anomalous considering the training data and does not compare with other queries in Q . It is also useful to assess whether $|q_{latest}|$ is anomalous considering the distribution of Q regardless of the training data, similar to previous works (e.g., [19], [20], [18]), as most queries are expected to be legitimate.

Therefore, we also use the Indicator 3B, which assesses to what extent the number of enabled features in q_{latest} is anomalous considering the queries in Q . The Indicator 3B uses the empirical rule (i.e., 3- σ rule) [51], which is a standard technique in anomaly detection. The score is calculated as the percentage change from 3 standard deviations of the mean using the following equation, where $\mu_{FeaturesQ}$ is the mean number of enabled features of queries in Q and $\sigma_{FeaturesQ}$ is the standard deviation:

$$C = \mu_{FeaturesQ} + (\sigma_{FeaturesQ} \times 3) \quad (5)$$

$$S_{3B} = (|q_{latest}| - C) / C \quad (6)$$

Indicators 4A & 4B. Distribution similarity via autoencoder reconstruction loss:

An autoencoder is a neural network where the input and output are the same but the hidden layers have fewer neurons [52]. This limits the amount of information that can travel through the model, requiring it to learn a compressed version of the input. Interestingly, input samples that belong to the distribution of the autoencoder’s training data will produce a smaller distance between the input and output representations (known as the *reconstruction loss*) [53]. Thus, we hypothesize that an autoencoder trained on legitimate queries will produce a significantly higher reconstruction loss for an adversarial query. Other novelty detection techniques, such as one-class SVMs, have been found to be outperformed by autoencoders in this task in other domains [54].

Hence, Indicator 4A’s score is the percentage change between the reconstruction loss of q_{latest} ($RecLoss_{q_{latest}}$) and the maximum reconstruction loss observed in training data samples ($maxRecLossD$). A higher percentage change suggests a significant increase from the maximum reconstruction loss observed for training data samples, and that q_{latest} may

be adversarial. The score is calculated as follows:

$$S_{4A} = (RecLoss_{q_{latest}} - maxRecLossD) / maxRecLossD \quad (7)$$

As with Indicator 3B, we also consider the case where q_{latest} is only compared to other queries in Q regardless of the training data distribution. In this way, the Indicator 4B uses the empirical rule to assess whether the reconstruction loss of q_{latest} is anomalous considering the queries in Q . The score is calculated as a percentage change from 3 standard deviations using the following equation, where $\mu_{RecLossQ}$ is the mean reconstruction loss of queries in Q and $\sigma_{RecLossQ}$ is the standard deviation:

$$C = \mu_{RecLossQ} + (\sigma_{RecLossQ} \times 3) \quad (8)$$

$$S_{4B} = (RecLoss_{q_{latest}} - C) / C \quad (9)$$

C. Attack Detection

Once the scores are produced by each indicator, they are aggregated into the score vector S . The next step is to make a final decision using all the scores. For this purpose, static techniques can be used to combine several individual scores into an overall score or decision, such as a weighted sum model or a sum of squares model [55]. However, a key drawback of these static methods is that the weights for each indicator (i.e., how much each indicator contributes to the overall score) must be selected manually by the defender. Moreover, static aggregation models cannot detect any underlying trends or patterns in the data.

Therefore, MalProtect instead uses a pre-trained *decision model* (e.g., a neural network) that makes a prediction on the score vector S . Each threat indicator is a feature of this model, with the predicted output from the decision model representing whether there is an attack in progress constituting its decision. Later, in Section V, we provide details about how the decision models can be constructed, including the ones we use in this work (see ‘‘MalProtect Configurations’’). Importantly, MalProtect allows analysts to understand why it has made a particular decision with its scoring system. Analysts can *further* evaluate the influence of each threat indicator on the final prediction by examining the global feature importance of the decision model. Later, in Section V, we present the decision models that we evaluated, and we then discuss how each threat indicator influences the predictions made by these models in Section VIII.

D. Defensive Action & Prediction

If no attack is detected by the decision model, q_{latest} is passed on to the classifier F (also known as the *prediction model*). The classifier makes a prediction of whether q_{latest} is benign or malware, which is then returned to the user. That is, if the latest query received is not considered to be part of an attack by MalProtect, the query is then forwarded to the ML classifier trained to decide whether an input sample is benign or malware. Note that the classifier can itself be made more robust using techniques such as adversarial training (as shown later). MalProtect only acts as a filter, which only lets what it considers to be *legitimate* queries to be passed on to the ML model used for malware classification.

If MalProtect does detect an attack, a defensive action can be taken, such as returning a specific prediction, returning

inconsistent labels, notifying system administrators, banning user accounts, or rejecting further user queries. In the configuration of MalProtect that we use in this paper, if an attack is detected, we take the defensive action of returning a *malware* prediction. This is as if the query had actually been forwarded to the classifier and that had been its output. This means that an attacker would encounter failure throughout the course of the attack, believing that their query is being consistently classified as malware without necessarily knowing about MalProtect’s presence. An exciting line of future work would be to study what defensive actions are more or less effective depending on the setting, but this is out of the scope of this paper.

V. EXPERIMENTAL SETUP & PRELIMINARIES

Datasets. Sampling from the true distribution is a challenging and open problem in many ML-based security applications [56], [57]. Particularly, the lack of publicly accessible, up-to-date datasets in the malware detection domain is a well-known issue that limits the remit of academic research in this domain. To mitigate this, we use three datasets representative of different architectures as well as collection dates and methods that have been used in a variety of studies (e.g., [58], [24], [59], [13], [60], [61], [62]). The datasets we use are AndroZoo for Android malware [63], SLEIPNIR for Windows malware [26], and DREBIN for Android malware [64]. The consistency of our results across the datasets later show that MalProtect’s performance transcends the characteristics and nuances of the datasets.

The AndroZoo dataset [63] contains Android apps from 2017 to 2018, offering apps from different stores and markets with VirusTotal summary reports for each. Similar to prior work [13], [62], [49], we consider an app malicious if it has 4 or more VirusTotal detections, and benign if it has 0 VirusTotal detections (with apps that have 1-3 detections discarded). The dataset contains $\approx 150K$ recent applications, with 135,859 benign and 15,778 malicious samples. To balance the dataset, we use 15,778 samples from each class. Meanwhile, SLEIPNIR consists of 19,696 benign and 34,994 malware samples. This dataset is derived from Windows API calls in PE files parsed by LIEF. As our work is in the feature-space, we use SLEIPNIR as a representation of Windows out of simplicity due to its feature space, which is binary. This permits a more precise comparison between the Android and Windows datasets. We use 19,696 samples from each class. Meanwhile, DREBIN is based on extracted static features from Android APK files. The dataset contains 123,453 benign and 5,560 malware samples, from which we use 5,560 samples from each class. As in prior work [27], [24], and for completeness, we use a large number of features for each dataset, i.e., 10,000 for AndroZoo, 22,761 for SLEIPNIR, and 58,975 for DREBIN.

Initially, the datasets are partitioned with an 80:20 ratio for training and test data according to the Pareto principle. This training data is partitioned again into training and validation data using the same ratio, producing a 64:16:20 split that has been extensively used in prior work (e.g., [65], [66]). We use the training data to construct the prediction models used in our evaluation. Meanwhile, we use a partition of randomly-chosen malware samples from a subset of the test data as the input

samples for the query attacks, with 234 samples for AndroZoo, 230 for SLEIPNIR, and 229 for DREBIN.

We also consider the well-established guidelines for conducting malware experiments [67]. As our prediction models decide whether an input sample is benign or malicious, it is necessary to retain benign samples in the datasets. Furthermore, we do not strictly balance datasets over malware families but instead over the positive and negative classes. We randomly sample unique samples from each class to appear in the training and test data without repetition [44], [26], [5].

Prediction Models (Non-Stateful Defenses). Recall that stateful defenses themselves do not provide predictions, but rather they protect ML prediction models that decide whether an input sample is benign or malicious. Hence, for our experimental evaluation, a selection of prediction models is required. In our evaluation, we use six *non-stateful* defenses as the prediction models. Rather than using vanilla models, we use single-model defenses, ensemble defenses, and moving target defenses as the prediction models to demonstrate their vulnerability to attacks despite being defenses in their own right. By using MalProtect and other stateful defenses in conjunction with a defense as a prediction model (rather than a vanilla model), it can introduce defense in depth, which is a frequently-adopted strategy in information security.

We use prediction models from prior work and train them on the three datasets according to the procedures outlined earlier. Each prediction model is constructed according to the procedures outlined in their original papers using the training data for each dataset (see Appendix B for information about architectures). For single-model defenses, we test a neural network with defensive distillation (NN-DD) [9]. We use several white-box attacks to develop adversarial examples for a set of vanilla models (see later subsection), which are used to adversarially-train a neural network (NN-AT). We adversarially-train with a quantity of adversarial examples that is 25% of the size of the training data [10]. For ensemble defenses, we test majority voting and veto voting [30], [29], [31], [68], [30]. For moving target defenses, we evaluate Morphence [69] and StratDef [30]. Each prediction model is evaluated against attacks as is (i.e., a non-stateful setup). Then, each prediction model is combined with each stateful defense (e.g., MalProtect), allowing us to evaluate different setups and configurations. We next provide details about the MalProtect configurations we use and the other stateful defenses.

MalProtect Configurations. We develop two MalProtect configurations to showcase our defense’s capabilities. Each configuration uses a different decision model to predict an attack based on the indicator scores. For both configurations, we cap the query history size at 10,000. Later in Section X we evaluate and discuss the efficiency related to this.

For each configuration, we follow the steps described in Section IV for each dataset. For some indicators, we derive certain values from the dataset that are needed to assess queries (e.g., $avgDistD$, which is the average L_0 distance of samples in the training data). Then, to train each decision model (which are used to predict an attack given a score vector), a defender could use real-world attack and score data. However, in its absence, a synthetic dataset can be used. Therefore, in

this work, we develop a synthetic dataset of 1000+ labelled samples. For this, we initialize the query history to simulate past user activity with a randomly-chosen set of input samples from the training data. We then supply a range of queries (both legitimate and adversarial, as part of attacks) from our training data to have each indicator produce scores. This produces the synthetic dataset, where each threat indicator is a feature and the class labels represent the true condition of the system at the time (whether it was under attack (1) or not (0)). To predict if an attack is in progress when given a score vector, the decision models we train are a *logistic regression model* for one configuration (MalProtect-LR) and a *neural network* for another (MalProtect-NN) (see Appendix A for architectures). We use an 80:20 split for the training of decision models. In theory, both models should offer accurate predictions [70]. When deployed, the decision model may need regular maintenance to ensure that it performs well given the ever-changing threat landscape (see Section XI later).

Other Stateful Defenses. We compare MalProtect with three stateful defenses that have been applied in other domains to detect query attacks. The L_0 defense is based on similarity detection using L_p norms (e.g., [17]); PRADA [18]; and Stateful Detection (SD) [16]. We provide the technical details of each stateful defense in Appendix D. Although we keep the implementations of these stateful defenses as close to the original as possible (e.g., procedures and parameters), techniques applicable to other domains (e.g., encoding queries to generate a feature vector) are not applicable to the malware detection domain. Recall that stateful defenses do not provide the predictions but are designed to protect the underlying prediction models.

Performing Query Attacks. To evaluate MalProtect and the other defenses, we apply query attacks under the black-box, gray-box and adaptive attacker scenarios. When conducting attacks and generating an adversarial example in the malware detection domain, the original malicious functionality must remain intact, and the feature vectors must remain discrete (i.e., consist of 0s and 1s) as they represent the presence or absence of each feature. Query attacks designed for other domains (e.g., [7], [71]) do not consider these constraints and cannot produce effective adversarial examples in our domain, as we show in an experiment in Appendix H.

For our domain, query attacks can use software transplantation-based techniques to perturb the features of a malware sample using features from benign “donor” samples (which are the benign samples from the test set). In our evaluation, we modify query attack strategies that have been proven successful in this domain [8]. We refer the reader to Appendix F for full details of each attack strategy that we use in each scenario. Commonly under each strategy, we initialize the query history to simulate past user activity (as explained before) prior to starting an attack. Then, using malware samples from our test set, we apply each attack strategy accordingly under each threat model to perturb malware feature vectors. We use the parameter n_{max} to govern the maximum number of allowed queries permitted, where the transplantation of features continues until the target model is evaded, n_{max} is reached, or the donor features

are exhausted. We apply a procedure to provide a lower bound of functionality preservation of the original malware sample. For this, only valid perturbations for each dataset are allowed [45], [72], [73], [26], [13], [30]. If a feature has been modified invalidly — that is, the functionality would not be preserved in the problem-space — it is restored to its original value. This is to offer a lower bound of functionality preservation within the feature-space, similar to prior work [41], [24], [45], in which it was shown that 60+% of adversarial examples generated following this approach executed as expected in a sandbox environment. Recent work [45].

AndroZoo and DREBIN permit feature addition and removal (see Appendix E). For SLEIPNIR, only feature addition is possible because of the processing performed by LIEF to extract features when originally developing the dataset. While we remain in the feature-space (like recent work e.g., [45], [41]), the perturbations could be translated to the problem-space using techniques from previous work (e.g., [13], [8]). For example, feature addition can be achieved by adding dead-code or by using opaque predicates [74], [13]. Feature removal — which is more complex — can be performed by rewriting the dexcode, encrypting API calls and network addresses (removing features but retaining functionality).

Performing Other Attacks. Prior work [17], [16], [21] has used *query blinding attacks* in image recognition with domain-specific techniques (e.g., modifying contrast or brightness). Such techniques are not applicable to our domain. Instead, we later present an adaptive attack strategy, where the attacker has full knowledge of how MalProtect operates, and aims to evade both it and the underlying classifier.

We also use a *transferability attack strategy* to generate adversarial examples using substitute models. The adversarial examples are used to adversarially-train models as a defender; and to evaluate stateful defenses later under different system conditions under a transferability attack (in Section IX). For

this, we construct four vanilla models using the training data (see Appendix C for architectures). These models act as estimations of the prediction models, for they share the same training data, but not the same architectures. We apply a range of white-box attacks (BIM [75], Decision Tree attack [76], FGSM [77], [78], JSMA [79] and SVM attack [76]) against these models to generate adversarial examples, in anticipation that they transfer to the target models. However, these attacks perturb features without considering the constraints of this domain. Therefore, we preserve the functionality of the original malware sample in the feature-space and discretize the generated feature vectors. That is, once an adversarial example is produced, each value in the feature vector is discretized (i.e., if the value in the feature vector is < 0.5 , it is set to 0, else it is set to 1). Then, if a feature has been modified invalidly according to functionality constraints, it is restored to its original value (similar to the process for query attacks). We then ensure that the adversarial example can still achieve evasion; if not, it is discarded. This procedure is designed to offer a lower bound of functionality preservation within the feature-space.

VI. BLACK-BOX QUERY ATTACK RESULTS

In the black-box scenario, the attacker can only access the predicted output of the oracle. Hence, we conduct the black-box query attack following the black-box attack strategy in Appendix F and the procedure outlined in Section V for performing query attacks. That is, the black-box attack selects the features to perturb in malware samples in a randomized manner. We conduct the attack against each *non-stateful* defense as is, to establish the baseline performance of the prediction models. We then evaluate and compare each combination of stateful defense and prediction model, including both MalProtect configurations.

Figure 3 shows the evasion rate of the black-box query attack versus n_{max} (see Appendix G for extended results).

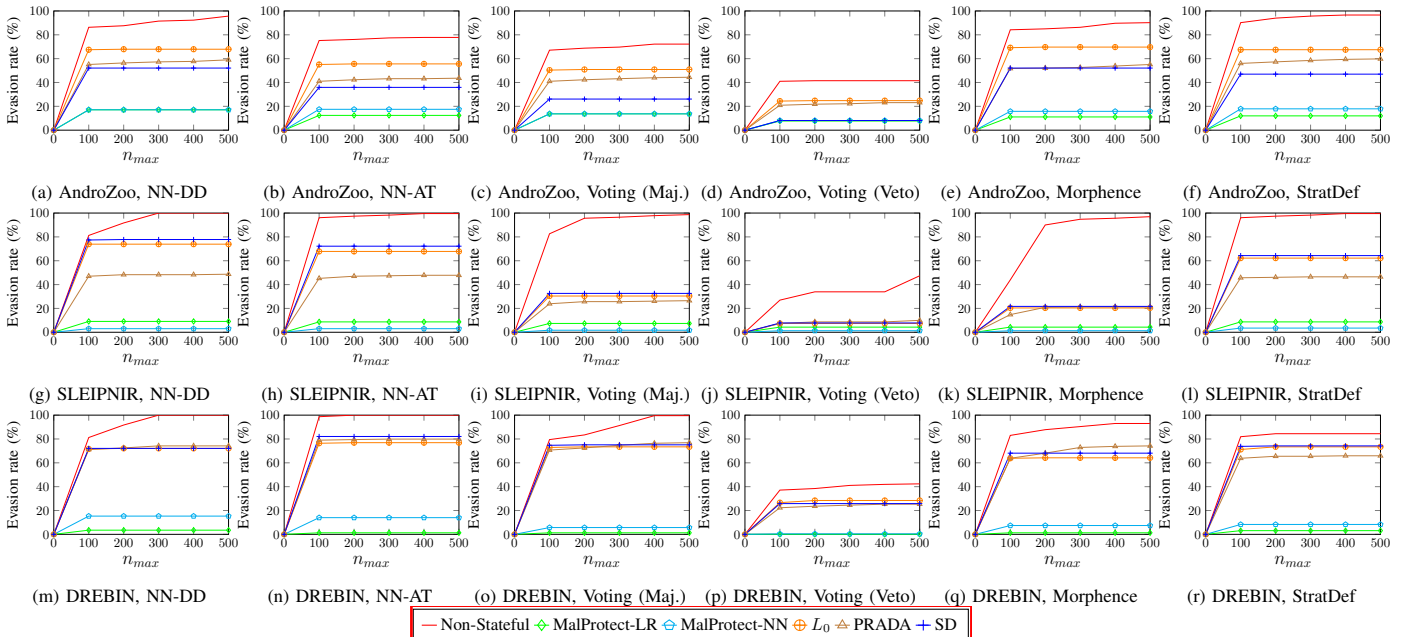


Fig. 3: Evasion rate vs. maximum number of queries (n_{max}) of black-box query attack against non-stateful defenses (prediction models) and each combination of prediction model and stateful defense.

Against non-stateful defenses (that is, the prediction models as they are), the black-box query attack can achieve a 100% evasion rate in some cases, with the average evasion rate sitting at 70+% across the datasets. These results clearly demonstrate that the non-stateful defenses do not provide adequate protection in this attack scenario, despite being defenses in their own right as they do not have any mechanism to detect attacks in progress. Thus, an attacker can continue to query them and iteratively develop successful adversarial examples. Out of the prediction models, only veto voting limits the effectiveness of the attack to some extent, likely because of a sensitive constituent model. Against veto voting, the attack achieves a maximum evasion rate of $\approx 50\%$ across the datasets. Despite this, around one in two to three queries can still achieve evasion against this model.

Meanwhile, our stateful defense, MalProtect, significantly decreases the attack success, with peak reductions in the evasion rates for AndroZoo, SLEIPNIR, and DREBIN of 84%, 96% and 98%, respectively. This highlights the benefits (and necessity) of a stateful defense for this domain, as attacks can be detected using different methods of analysis. Comparing the two MalProtect configurations, MalProtect-LR and MalProtect-NN exhibit similar performance, with both configurations able to detect attacks after 5, 7, and 2 queries for AndroZoo, SLEIPNIR, and DREBIN, respectively. This level of performance is achieved without greatly compromising other metrics, such as the accuracy or false positive rate on legitimate queries (as we show in Section IX later).

Meanwhile, the other stateful defenses that we evaluate are significantly less effective against this attack. For AndroZoo, L_0 offers the mildest performance in terms of adversarial robustness, while SD offers slightly improved performance, but only marginally. In the attack’s weakest performance against veto voting, we find that SD performs similarly to MalProtect for AndroZoo and SLEIPNIR, but its performance is much inferior in other cases (with other prediction models)

where MalProtect demonstrates robustness, likely because SD is unable to detect attacks using techniques beyond similarity detection. For SLEIPNIR, PRADA offers better performance than the prior stateful defenses, though only marginally in most cases, while L_0 and SD seem similar in their defensive robustness. Nonetheless, their performance is much inferior compared to MalProtect. For DREBIN, we observe that the three prior stateful defenses exhibit similar performance in all cases. Overall, the black-box query attack can achieve a 60+% evasion rate in most cases against prior stateful defenses, compared with a peak evasion rate of $\approx 18\%$ for MalProtect. Hence, MalProtect provides demonstrably better protection against the black-box query attack.

VII. GRAY-BOX QUERY ATTACK RESULTS

Prior work has stressed the importance of evaluating defenses against stronger adversaries [32]. Therefore, we next evaluate MalProtect’s performance against the gray-box query attack and compare it with other stateful defenses in that attack scenario. Under this scenario, the attacker has access to the training data of the prediction models and has more information about the features. The gray-box attack strategy that we use (see Appendix F) therefore utilizes information about the frequency of features in benign samples to determine the order of transplanted features, resulting in a more effective attack. As before, we conduct the attack against each non-stateful defense to establish a baseline performance, and against each combination of stateful defense and prediction model.

Figure 4 shows the evasion rate of the gray-box query attack against all combinations of stateful defenses and prediction models, versus n_{max} (see Appendix G for extended results). Against the non-stateful defenses, this stronger attack achieves significantly greater evasion across all the datasets, with 100% evasion rate in many cases. This is because the non-stateful defenses have no mechanism to detect that they are under attack; against a stronger attack, as in this scenario, they ex-

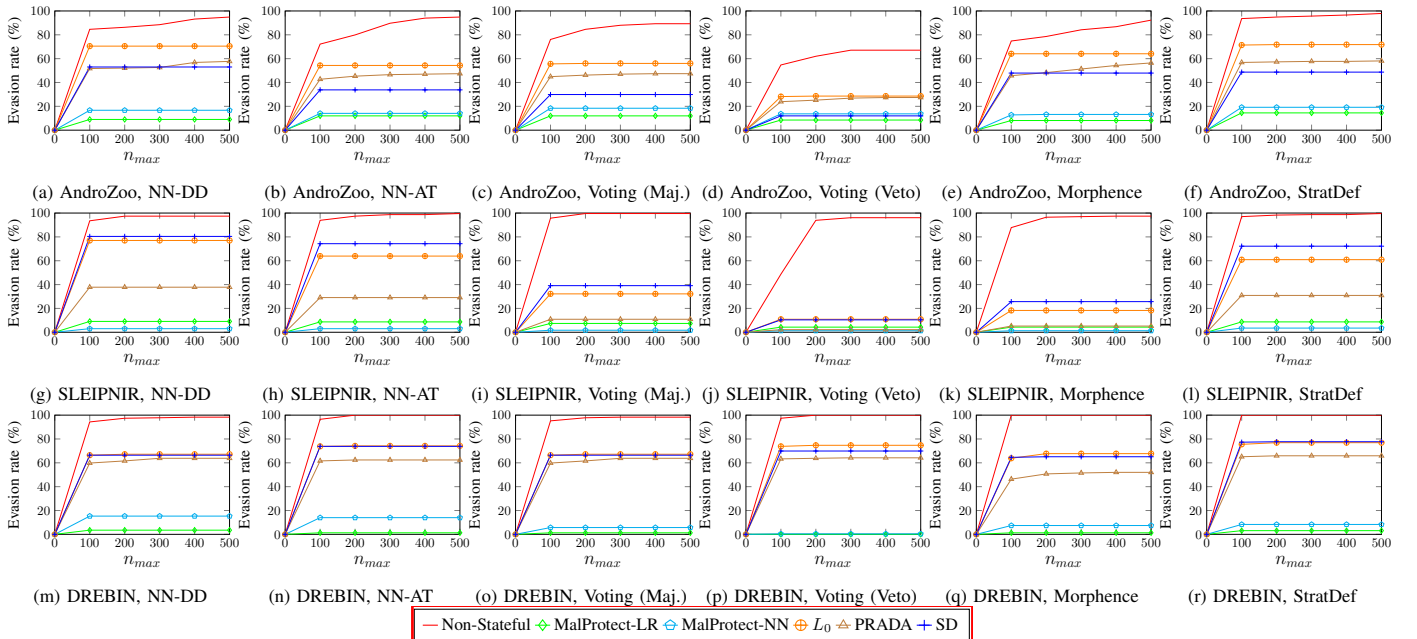


Fig. 4: Evasion rate vs. maximum number of queries (n_{max}) of gray-box query attack against non-stateful defenses (prediction models) and each combination of prediction model and stateful defense.

hibit greater vulnerability. Intriguingly, the greatest difference from the black-box query attack can be observed in the results against veto voting, which was the least evaded non-stateful defense in Section VI. Whereas the black-box query attack peaked at $\approx 40\%$ evasion rate against this model (without any stateful defense), the gray-box query attack achieves $60 + \%$ evasion rate.

As far as protection by stateful defenses is concerned, MalProtect performs the best. In the best case for the configurations, it reduces the average evasion rate across all models from a peak of 84% to 13% for AndroZoo, 94% to 5% for SLEIPNIR, and 91% to 6% for DREBIN, while only taking ≈ 6 , 6 , and 2 queries to detect an attack for AndroZoo, SLEIPNIR, and DREBIN, respectively. Like the black-box query attack, MalProtect-LR and MalProtect-NN offer similar performance, with MalProtect-NN slightly outperforming MalProtect-LR in some cases and vice-versa. In fact, the gray-box query attack only achieves a maximum evasion rate of $\approx 20\%$ against MalProtect across the board. This is higher than the black-box attack results, though expected, since the gray-box attack represents a more hostile threat where the attacker has more knowledge and information.

Meanwhile, for the other stateful defenses, we observe trends similar to those seen in the results for the black-box query attack. For AndroZoo, SD outperforms L_0 and PRADA, with average evasion rates of 57.5% , 46.9% , and 37.5% against L_0 , PRADA, and SD, respectively. For SLEIPNIR, PRADA outperforms L_0 and SD for the majority of prediction models, a trend that was also observed in the black-box query attack results. Here, the average evasion rates of the attack sit at 43.8% , 19.3% , and 50.3% for L_0 , PRADA, and SD, respectively. Although PRADA appears to offer some robustness, it underperforms when other metrics are also considered (especially in false positives, as we show later in Section IX). For DREBIN, the prior stateful defenses perform similarly with a more balanced performance. That is, they perform equally poorly, with the gray-box query attack able to achieve average evasion rates of 71.3% , 62% , and 69.9% for L_0 , PRADA, and SD, respectively, across all configurations. Interestingly, once again, PRADA (minimally) surpasses L_0 and SD in terms of robustness, though the poor general performance of these prior stateful defenses seen for both black-box and gray-box results reaffirms that relying on a single similarity or OOD detection mechanism is inadequate for this domain. In our domain, attackers use different techniques to generate adversarial examples that often include replacing features rather than making small detectable perturbations.

Further evaluating the evasion rate versus n_{max} , 100 or fewer queries are enough to achieve attack success in most cases. Query attacks in domains using continuous features (e.g., image recognition) may require substantially more queries [7] to achieve attack success compared with domains that use discrete features. This is because perturbations in a discrete feature-space (e.g., 0 to 1) have a greater effect on the final prediction, requiring fewer of them to accomplish evasion than perturbations made per query in a continuous feature-space (e.g., $+0.01$ to some feature). In fact, in Appendix H we demonstrate the ineffectiveness of query attacks for other

domains when applied to ML-based malware detection.

VIII. INTERPRETABILITY & ADAPTIVE ATTACK

A. Interpretability of MalProtect

As ML is being applied more to the cybersecurity domain, a key challenge is the interpretability of predictions made by models [80]. As an initial step towards addressing this in stateful defenses, MalProtect produces interpretable decisions. By observing the scores produced by each threat indicator, analysts can better understand how MalProtect made a particular decision, as higher scores for an indicator give clues about what is anomalous about queries.

The influence of each threat indicator on MalProtect’s decision can be examined by analyzing the *global feature importance* in MalProtect’s decision model. Recall that each threat indicator is a feature of the decision model. Therefore, we can assess the global feature importance using SHAP [81], which is a widely-used framework for interpreting ML predictions. Each feature of the decision model of the MalProtect configurations is given an importance value by SHAP at a global level, which produces the data observable in Figure 5.

Figure 5 shows that the Indicator 4A has the greatest influence on MalProtect’s predicted output across all decision models and datasets, which indicates that a significantly high autoencoder loss is more likely to affect the predicted output. Meanwhile, the importance of the other indicators is more balanced. Intriguingly, the indicators based on similarity detection seem less influential, which is interesting because other stateful defenses rely solely on similarity detection to detect attacks. In malware detection, not only do we show that attackers can evade stateful defenses relying solely on such detection methods, but we further show that MalProtect rightly considers those indicators as less important.

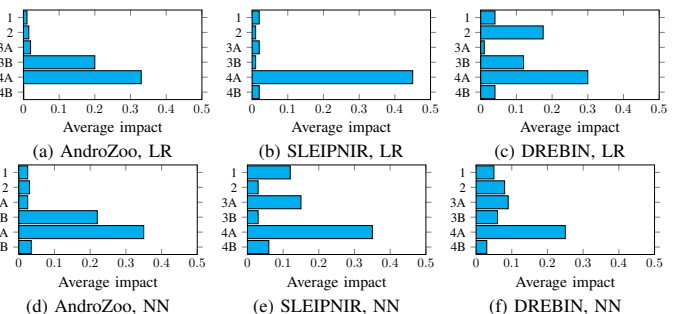


Fig. 5: Influence of MalProtect (LR & NN) threat indicators. The average impact on model output magnitude is shown for each threat indicator.

Intriguingly, a capable attacker could leverage this information to evade MalProtect. In an adaptive attack (i.e., white-box attack) scenario [32], [82], the attacker has complete knowledge of MalProtect and, thus, knowledge of how much each indicator contributes to the final decision. With this information, an attacker could craft an adversarial example that evades both MalProtect and the underlying prediction model. We next examine how such an attack could be performed.

Attack Rationale. In the adaptive attacker scenario, the attacker knows precisely how MalProtect operates, its internal workings, and other pertinent information related to the threat

indicators. At a high level, to evade MalProtect’s detection system, adversarial queries must appear as legitimate as possible and convince the indicators that the queries are distinct yet *normal*. Additionally, adversarial queries must remain within the distribution of legitimate queries. For example, we know that the Indicator 4A is the most influential in the final prediction, and therefore it would be essential to ensure that the autoencoder loss is not abnormally high. In essence, queries must remain far apart from each other while retaining their original malicious functionality and appearing legitimate. Importantly, any generated adversarial example must also evade the underlying prediction model; otherwise, such an attack would be of no use. Attackers must bypass all components of the target model to be successful [1].

Adaptive Attack Strategy. We modify the gray-box attack strategy to produce the adaptive attack strategy (see Appendix F). Firstly, we limit the number of features that can be perturbed in a single iteration of the attack so as not to exhaust possible perturbations early on. This frees up perturbations to be used in later iterations of the attack if earlier queries cannot achieve evasion. While this may increase the query distance and make queries appear less anomalous, other indicators (e.g., Indicator 4A) may still be able to trigger attack detection. To deal with this, the adaptive attack strategy also *removes* a proportion (p) of features at each iteration. This means that queries will be more distant, with fewer shared and enabled features, while conforming to the distribution of legitimate queries and the training data. For example, a *combination of features* that may cause an increase in the autoencoder loss (such that it appears anomalous) may be removed. Only the AndroZoo and DREBIN datasets support the removal of features while preserving functionality within the feature-space. Therefore, the adaptive attack strategy can only be applied to these datasets.

Results. Figure 6 shows the evasion rate of the adaptive attack against the stateful defenses (averaged across the maximum number of queries permitted, up to 500) versus the percentage of features removed (p). Recall that successful adversarial examples evade each stateful defense as well as the underlying prediction model, which is NN-AT.




Fig. 6: Average evasion rate vs. p for adaptive attack. The average is across the maximum number of queries permitted.

For both datasets, the adaptive attack fails to achieve significant increases in the evasion rate against either MalProtect configuration. In fact, in some cases, the average evasion rate decreases as p increases. This is likely because the (benign) features selected to cross the decision boundary are in fact removed. This attack causes other stateful defenses to exhibit similar performance to the gray-box query attack, with an average evasion rate of 72.9%, 61.9%, and 67.7% for L_0 ,

PRADA, and SD, respectively, and consistent evasion rate across different values of p . This is expected, as the adaptive attack is designed to target MalProtect and not the other stateful defenses (which are included as a comparison here).

IX. BEYOND ADVERSARIAL ROBUSTNESS

We next evaluate the stateful defenses under different system conditions using metrics beyond the evasion rate. To do this, we query stateful defenses with benign and malicious queries, some of which are adversarial. This is imperative, as the original task of classifying queries as benign or malicious must also be done properly.

Procedure. We examine the performance of stateful defenses in different system conditions with different levels of adversity. For this, we generate adversarial examples using the transferability attack strategy described in Section V. Each stateful defense (with NN-AT as the prediction model) is then queried 1000+ times with values of $0.1 \leq k \leq 0.9$, where k represents the adversarial intensity. For example, at $k = 0.1$ — which represents a less adversarial environment — 10% of all queries are adversarial examples, while the remaining 90% are equally split between benign and *non-adversarial* malware samples from the test set. As k increases, the system conditions become more hostile as the system faces more adversarial queries. We do not consider $k = 0$ (no adversarial examples) or $k = 1$ (only adversarial examples), as such environments are less likely. As before, the query history is initialized as necessary with random samples from the training data.

Results. Figure 7 shows the performance of each stateful defense versus k . In particular, MalProtect exhibits reasonably stable performance across the datasets considering all metrics, with the accuracy and F1 remaining greater than 90%, demonstrating MalProtect’s ability to properly classify user queries. Intriguingly, other stateful defenses also appear to provide decent performance in terms of these metrics in some cases. For AndroZoo and SLEIPNIR, L_0 , PRADA, and SD all exhibit relatively stable accuracy and F1 as k increases. This is possibly because the stateful defenses have it easier to detect the (less evasive) adversarial examples, which constitute most of the queries, as they are generated by attacks not for this domain. Recall that the perturbations applied by the attacks to generate the adversarial examples may be reversed if features are modified such that functionality is affected; this would limit attack success. Regardless, these stateful defenses do not perform as well as MalProtect. For DREBIN, we observe the decline of the non-stateful NN-AT model and stateful defenses except MalProtect when these metrics are considered.

Importantly, though, additional metrics must be examined to obtain a comprehensive assessment of performance. Analysis of the AUC reveals MalProtect’s superiority in distinguishing the benign and malware classes across the datasets without disrupting the classification of legitimate queries. Meanwhile, PRADA and SD exhibit significantly lower AUC in some cases (especially as k increases), implying that their decisions are no better than flipping a coin, which is contradictory to the higher accuracy and F1 that they exhibit. This is possibly because as k increases, the number of malware queries also increases in the test set because of more adversarial queries. If PRADA,

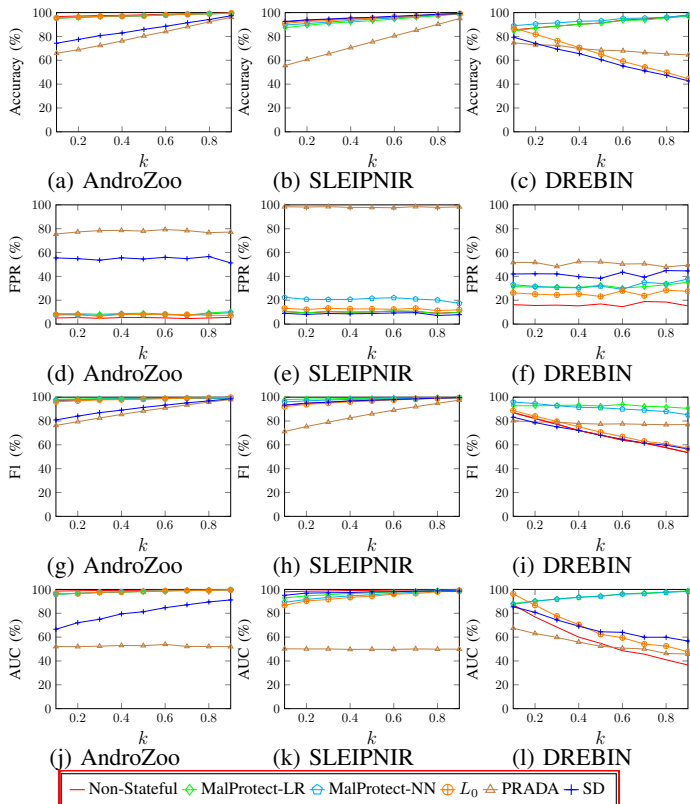


Fig. 7: Accuracy, FPR, F1 & AUC vs. k per stateful defense.

for example, decides that most queries are adversarial (and therefore malware), this would lead to high accuracy, even if it misclassifies benign queries. Therefore, PRADA seems less capable of separating between the classes in this domain. This is further evidenced through an examination of the false positive rate (FPR) next.

In the malware detection domain, as we have explained before, a low FPR is essential to ensure a reliable service for users [24], [83], [42], [84]. Although a stateful defense may improve protection against attacks, the FPR may also increase compared with the underlying prediction model. This can be attributed to the sensitivity of detection mechanisms; that is, benign queries may be more frequently misclassified as malicious. Importantly, though, MalProtect offers a low FPR that is as close as possible to the non-stateful prediction model for all datasets as it uses several indicators to predict an attack. Meanwhile, L_0 is the only other stateful defense that offers comparable performance to MalProtect in this regard. However, as we have seen, it does not work well at defending against query attacks in this domain. Conversely, PRADA exhibits a significantly high FPR for nearly all datasets, which supports the notion that it is merely classifying most queries as malware or adversarial rather than actually distinguishing queries between the classes in this experiment. This is somewhat resembled by SD, with its high FPR for the Android datasets. Overall, across the metrics and datasets, we showcase MalProtect’s ability to perform well across a range of different queries.

X. EFFICIENCY OF MALPROTECT VS. OTHER DEFENSES

Another important aspect is to evaluate MalProtect’s storage and time costs and compare them with those of other defenses.

Storage. The storage cost scales linearly with the size of the query history, Q , where queries are stored. The storage costs are the same for any stateful defense that uses a query history of the same size. In our evaluation, the size of Q is capped at 10,000, which only consumes ≈ 0.08 MB storage for all the datasets. If $|Q|$ is increased to 50K, only ≈ 0.4 MB would be required for all datasets, which is negligible.

Prediction Time. The *prediction time* refers to the total duration from when a query is received by the system until a prediction is returned to the user. For most non-stateful defenses and ML classifiers, this is nominal. However, there is latency associated with the analysis of the query history by stateful defenses, resulting in longer prediction times. In the worst-case, stateful defenses must analyze the entire query history, whose time cost is the *worst-case prediction time*. Figure 8 shows how the worst-case prediction time is affected by $|Q|$.




Fig. 8: Worst-case prediction time vs. size of Q ($|Q|$). Importantly, with optimization and distributed computing techniques, the time costs can be drastically reduced.

Figure 8 shows that the worst-case prediction-time scales linearly with $|Q|$ for all defenses. In our experiments where $|Q|$ is 10,000, the worst-case prediction time for MalProtect sits at ≈ 0.4 seconds for AndroZoo and SLEIPNIR, and ≈ 0.6 seconds for DREBIN. Both MalProtect configurations have the same overall cost as they have an identical analysis stage, while predictions from the decision model are near-instant. Meanwhile, other stateful defenses such as L_0 , PRADA, and SD have slightly lower worst-case prediction times (≈ 0.03 -0.6 seconds across the datasets) for this size of Q . In all cases, however, the worst-case prediction times increase linearly with $|Q|$. This means that with MalProtect, a significant improvement in robustness can be achieved (up to 80% reduction in evasion rate as shown in our experimental evaluation) at a slightly higher but still linear cost.

Optimization. MalProtect’s performance can be easily optimized. For example, by running MalProtect on an NVIDIA A100 GPU, the worst-case prediction time is reduced by over half (≈ 0.21 -0.41 seconds) across the datasets when $|Q| = 10,000$.

XI. CONCLUSION

In this paper, we presented MalProtect, which is the first stateful defense for adversarial query attacks in the ML-based malware detection domain. As we have shown, ML prediction models and defenses exhibit significant vulnerability to query attacks in this domain. Prior stateful defenses that have been applied to other domains provide little protection here either. Meanwhile, our defense, MalProtect, does not rely solely on a single form of similarity or out-of-distribution detection, as these prior stateful defenses do. Instead, we use several

threat indicators and a decision model to detect attacks more effectively. Our evaluation has shown that MalProtect performs well against attacks under various scenarios and also offers more reliable predictions for non-adversarial queries than prior stateful defenses. Furthermore, MalProtect displays resilience even against adaptive attackers.

In future work, we aim to explore a number of key aspects. Firstly, we aim to better understand the relationship between the size of the query history, the defender’s available resources, and the detection effectiveness. Since the query history cannot be infinite, it may be possible for attackers to submit multiple queries in order to move the sliding window to a particular point. An attacker may then dispatch adversarial queries at specific intervals so as to evade detection in a *spaced-out attack*. However, this attack may also be considered less practical for a number of reasons. The attacker would somehow need to know the size of the query history and then spread adversarial queries accordingly, notwithstanding the associated cost of doing so (e.g., time). Moreover, MalProtect is not a defense that is to be deployed in complete isolation. Hence, there is no guarantee that if a spaced-out attack were to succeed in evading the sliding window for detection, the underlying prediction model would be evaded.

Secondly, we seek to understand how additional indicators could be added to increase protection, such as cyber-threat intelligence [85], [86]. Moreover, there is an open research direction regarding *concept drift*. This relates to the constant evolution of malware, which makes it difficult to detect unseen behavior [60], [87], [88], leading to unsustainable models. Prior work has suggested retraining a model regularly [89], [90]. In the case of MalProtect, the defender may need to regularly evaluate the indicators and the decision model for predicting attacks. One potential way to do this would be to couple MalProtect with a detection framework for detecting when such modifications are necessary [60]. A further intriguing research direction is exploring the effectiveness of *online learning* in stateful defenses. With such a mechanism, the underlying prediction models could be updated on queries that are considered adversarial by the detection system in real-time.

REFERENCES

- [1] G. Apruzzese *et al.*, ““real attackers don’t compute gradients”: Bridging the gap between adversarial ml research and practice,” in *Proceedings of the 1st IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, 2023.
- [2] H. Yu *et al.*, “Cloudleak: Large-scale deep learning models stealing through adversarial examples,” in *NDSS*, 2020.
- [3] C. Szegedy *et al.*, “Intriguing properties of neural networks,” *arXiv preprint arXiv:1312.6199*, 2014.
- [4] N. Papernot *et al.*, “Practical black-box attacks against machine learning,” *Proceedings of the 2017 ACM on Asia framework for enhancing deep neural networks against adversaria conference on computer and communications security*, pp. 506–519, 2018.
- [5] I. Rosenberg *et al.*, “Generic black-box end-to-end attack against state of the art api call based malware classifiers,” in *International Symposium on Research in Attacks, Intrusions, and Defenses*. Springer, 2018, pp. 490–510.
- [6] W. Brendel *et al.*, “Decision-based adversarial attacks: Reliable attacks against black-box machine learning models,” *arXiv preprint arXiv:1712.04248*, 2017.
- [7] J. Chen *et al.*, “Hopskipjumpattack: A query-efficient decision-based attack,” in *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2020, pp. 1277–1294.
- [8] I. Rosenberg *et al.*, “Query-efficient black-box attack against sequence-based malware classifiers,” in *Annual Computer Security Applications Conference*, 2020, pp. 611–626.
- [9] N. Papernot *et al.*, “Distillation as a defense to adversarial perturbations against deep neural networks,” in *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2016, pp. 582–597.
- [10] F. Tramèr *et al.*, “Ensemble adversarial training: Attacks and defenses,” *arXiv preprint arXiv:1705.07204*, 2017.
- [11] S. Sengupta *et al.*, “Mtdeep: boosting the security of deep neural nets against adversarial attacks with moving target defense,” in *Workshops at the Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [12] A. Ilyas *et al.*, “Black-box adversarial attacks with limited queries and information,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 2137–2146.
- [13] Pierazzi *et al.*, “Intriguing properties of adversarial ml attacks in the problem space,” in *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 2020, pp. 1308–1325. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/SP40000.2020.00073>
- [14] D. Hendrycks *et al.*, “Unsolved problems in ml safety,” *arXiv preprint arXiv:2109.13916*, 2021.
- [15] M. Brundage *et al.*, “The malicious use of artificial intelligence: Forecasting, prevention, and mitigation,” *arXiv preprint arXiv:1802.07228*, 2018.
- [16] S. Chen *et al.*, “Stateful detection of black-box adversarial attacks,” in *Proceedings of the 1st ACM Workshop on Security and Privacy on Artificial Intelligence*, 2020, pp. 30–39.
- [17] H. Li *et al.*, “Blacklight: Scalable defense for neural networks against Query-Based Black-Box attacks,” in *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, Aug. 2022. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity22/presentation/li-huiying>
- [18] M. Juuti *et al.*, “Prada: protecting against dnn model stealing attacks,” in *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2019, pp. 512–527.
- [19] S. Kariyappa *et al.*, “Defending against model stealing attacks with adaptive misinformation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [20] B. G. Atli *et al.*, “Extraction of complex dnn models: Real threat or boogeyman?” in *International Workshop on Engineering Dependable and Secure Machine Learning Systems*. Springer, 2020, pp. 42–57.
- [21] Z. Zhang *et al.*, “Seat: Similarity encoder by adversarial training for detecting model extraction attack queries,” in *Proceedings of the 14th ACM Workshop on Artificial Intelligence and Security*, ser. AISec ’21. New York, NY, USA: Association for Computing Machinery, 2021, p. 37–48. [Online]. Available: <https://doi.org/10.1145/3474369.3486863>
- [22] S. Pal *et al.*, “Stateful detection of model extraction attacks,” *arXiv preprint arXiv:2107.05166*, 2021.
- [23] B. Biggio *et al.*, “Evasion attacks against machine learning at test time,” in *Joint European conference on machine learning and knowledge discovery in databases*. Springer, 2013, pp. 387–402.
- [24] K. Grosse *et al.*, “Adversarial examples for malware detection,” in *European symposium on research in computer security*. Springer, 2017, pp. 62–79.
- [25] —, “Adversarial perturbations against deep neural networks for malware classification,” *arXiv preprint arXiv:1606.04435*, 2016.
- [26] A. Al-Dujaili *et al.*, “Adversarial deep learning for robust detection of binary encoded malware,” in *2018 IEEE Security and Privacy Workshops (SPW)*. IEEE, 2018, pp. 76–82.
- [27] A. Demontis *et al.*, “Yes, machine learning can be more secure! a case study on android malware detection,” *IEEE Transactions on Dependable and Secure Computing*, 2017.
- [28] T. Pang *et al.*, “Improving adversarial robustness via promoting ensemble diversity,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 4970–4979.
- [29] M. Cheng *et al.*, “Voting based ensemble improves robustness of defensive models,” *arXiv preprint arXiv:2011.14031*, 2020.
- [30] A. Rashid *et al.*, “Stratdef: Strategic defense against adversarial attacks in ml-based malware detection,” *arXiv preprint arXiv:2202.07568*, 2022.
- [31] R. K. Shahzad *et al.*, “Comparative analysis of voting schemes for ensemble-based malware detection,” *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, vol. 4, no. 1, pp. 98–117, 2013.
- [32] N. Carlini *et al.*, “On evaluating adversarial robustness,” *arXiv preprint arXiv:1902.06705*, 2019.
- [33] A. Athalye *et al.*, “Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples,” in *International conference on machine learning*. PMLR, 2018, pp. 274–283.

- [34] M. Maghsoudimehrabani *et al.*, “Proactive detection of query-based adversarial scenarios in nlp systems,” in *Proceedings of the 15th ACM Workshop on Artificial Intelligence and Security*, ser. AISeC’22. New York, NY, USA: Association for Computing Machinery, 2022, p. 103–113. [Online]. Available: <https://doi.org/10.1145/3560830.3563727>
- [35] S. Bulusu *et al.*, “Anomalous example detection in deep learning: A survey,” *IEEE Access*, vol. 8, pp. 132 330–132 347, 2020.
- [36] R. Braden *et al.*, “Rfc1636: Report of iab workshop on security in the internet architecture - february 8-10, 1994,” USA, 1994.
- [37] S. Kariyappa *et al.*, “Improving adversarial robustness of ensembles with diversity training,” *arXiv preprint arXiv:1901.09981*, 2019.
- [38] J. R. Douceur, “The sybil attack,” in *Revised Papers from the First International Workshop on Peer-to-Peer Systems*, ser. IPTS ’01. Berlin, Heidelberg: Springer-Verlag, 2002, p. 251–260.
- [39] Z. Yan *et al.*, “Policy-driven attack: Learning to query for hard-label black-box adversarial examples,” in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=pzpytjk3Xb2>
- [40] N. Carlini *et al.*, “Towards evaluating the robustness of neural networks,” in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 39–57.
- [41] G. Severi *et al.*, “Explanation-guided backdoor poisoning attacks against malware classifiers,” in *30th {USENIX} Security Symposium ({USENIX} Security 21)*, 2021.
- [42] W. Yang *et al.*, “Malware detection in adversarial settings: Exploiting feature evolutions and confusions in android apps,” in *Proceedings of the 33rd Annual Computer Security Applications Conference*, 2017, pp. 288–302.
- [43] O. Suciú *et al.*, “Exploring adversarial examples in malware detection,” in *2019 IEEE Security and Privacy Workshops (SPW)*. IEEE, 2019, pp. 8–14.
- [44] R. Podschwadt *et al.*, “On Effectiveness of Adversarial Examples and Defenses for Malware Classification,” *International Conference on Security and Privacy in Communication Systems*, pp. 380–393, 2019.
- [45] D. Li *et al.*, “A framework for enhancing deep neural networks against adversarial malware,” *IEEE Transactions on Network Science and Engineering*, vol. 8, no. 1, pp. 736–750, 2021.
- [46] F. Ceschin *et al.*, “No need to teach new tricks to old malware: Winning an evasion challenge with xor-based adversarial samples,” in *Reversing and Offensive-Oriented Trends Symposium*, ser. ROOTS’20. New York, NY, USA: Association for Computing Machinery, 2021, p. 13–22. [Online]. Available: <https://doi.org/10.1145/3433667.3433669>
- [47] P. Laskov *et al.*, “Practical evasion of a learning-based classifier: A case study,” in *2014 IEEE Symposium on Security and Privacy*. IEEE, 2014, pp. 197–211.
- [48] N. Papernot *et al.*, “Sok: Security and privacy in machine learning,” in *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2018, pp. 399–414.
- [49] B. Miller *et al.*, “Reviewer integration and performance measurement for malware detection,” in *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*. Springer, 2016, pp. 122–141.
- [50] S. Afroz, “How to build realistic machine learning systems for security?” San Francisco, CA: USENIX Association, Jan. 2020.
- [51] F. Pukelsheim, “The three sigma rule,” *The American Statistician*, vol. 48, no. 2, pp. 88–91, 1994. [Online]. Available: <http://www.jstor.org/stable/2684253>
- [52] J. Schmidhuber, “Deep learning in neural networks: An overview,” *Neural networks*, vol. 61, pp. 85–117, 2015.
- [53] G. Pang *et al.*, “Deep learning for anomaly detection: A review,” *ACM Comput. Surv.*, vol. 54, no. 2, mar 2021. [Online]. Available: <https://doi.org/10.1145/3439950>
- [54] D. Gong *et al.*, “Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 1705–1714.
- [55] P. C. Fishburn, “Letter to the editor—additive utilities with incomplete product sets: Application to priorities and assignments,” *Operations Research*, vol. 15, no. 3, pp. 537–542, 1967. [Online]. Available: <https://doi.org/10.1287/opre.15.3.537>
- [56] D. Arp *et al.*, “Dos and don’ts of machine learning in computer security,” in *Proc. of the USENIX Security Symposium*, 2022.
- [57] O. A. Aslan *et al.*, “A comprehensive review on malware detection approaches,” *IEEE Access*, vol. 8, pp. 6249–6271, 2020.
- [58] Y. Wang *et al.*, “Robust learning against relational adversaries,” in *Advances in Neural Information Processing Systems*.
- [59] Dauodi *et al.*, “A deep dive inside drebin: An explorative analysis beyond android malware detection scores,” *ACM Trans. Priv. Secur.*, vol. 25, no. 2, may 2022. [Online]. Available: <https://doi.org/10.1145/3503463>
- [60] Jordaney *et al.*, “Transcend: Detecting concept drift in malware classification models,” in *26th USENIX Security Symposium (USENIX Security 17)*. Vancouver, BC: USENIX Association, Aug. 2017, pp. 625–642. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/jordaney>
- [61] F. Barbero *et al.*, “Transcending transcend: Revisiting malware classification in the presence of concept drift,” in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 805–823.
- [62] F. Pendlebury *et al.*, “Tesseract: Eliminating experimental bias in malware classification across space and time,” in *Proceedings of the 28th USENIX Conference on Security Symposium*, ser. SEC’19. USA: USENIX Association, 2019, p. 729–746.
- [63] K. Allix *et al.*, “Androzo: Collecting millions of android apps for the research community,” in *Proceedings of the 13th International Conference on Mining Software Repositories*, ser. MSR ’16. New York, NY, USA: ACM, 2016, pp. 468–471. [Online]. Available: <http://doi.acm.org/10.1145/2901739.2903508>
- [64] D. Arp *et al.*, “Drebin: Effective and explainable detection of android malware in your pocket,” in *Ndss*, vol. 14, 2014, pp. 23–26.
- [65] J. Ma *et al.*, “Partner-assisted learning for few-shot image classification,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10 573–10 582.
- [66] A. Li *et al.*, “Few-shot learning with global class representations,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 9715–9724.
- [67] C. Rossow *et al.*, “Prudent practices for designing malware experiments: Status quo and outlook,” in *2012 IEEE Symposium on Security and Privacy*, 2012, pp. 65–79.
- [68] S. Y. Yerima *et al.*, “Droidfusion: A novel multilevel classifier fusion approach for android malware detection,” *IEEE transactions on cybernetics*, vol. 49, no. 2, pp. 453–466, 2018.
- [69] A. Amich *et al.*, “Morphence: Moving target defense against adversarial examples,” in *Annual Computer Security Applications Conference*, ser. ACSAC. New York, NY, USA: Association for Computing Machinery, 2021, p. 61–75. [Online]. Available: <https://doi.org/10.1145/3485832.3485899>
- [70] S. Dreiseitl *et al.*, “Logistic regression and artificial neural network classification models: a methodology review,” *Journal of biomedical informatics*, vol. 35, no. 5-6, pp. 352–359, 2002.
- [71] P.-Y. Chen *et al.*, “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.
- [72] Z. Abaid *et al.*, “Quantifying the impact of adversarial evasion attacks on machine learning based android malware classifiers,” in *2017 IEEE 16th International Symposium on Network Computing and Applications (NCA)*, 2017, pp. 1–10.
- [73] D. Li *et al.*, “Enhancing deep neural networks against adversarial malware examples,” *arXiv preprint arXiv:2004.07919*, 2020.
- [74] A. Moser *et al.*, “Limits of static analysis for malware detection,” in *Twenty-third annual computer security applications conference (ACSAC 2007)*. IEEE, 2007, pp. 421–430.
- [75] A. Kurakin *et al.*, “Adversarial examples in the physical world,” 2016.
- [76] K. Grosse *et al.*, “On the (statistical) detection of adversarial examples,” *arXiv preprint arXiv:1702.06280*, 2017.
- [77] I. J. Goodfellow *et al.*, “Explaining and harnessing adversarial examples,” *arXiv preprint arXiv:1412.6572*, 2014.
- [78] J. Dhaliwal *et al.*, “Gradient similarity: An explainable approach to detect adversarial attacks against deep learning,” *arXiv preprint arXiv:1806.10707*, 2018.
- [79] N. Papernot *et al.*, “The limitations of deep learning in adversarial settings,” in *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2016, pp. 372–387.
- [80] A. Nadeem *et al.*, “Sok: Explainable machine learning for computer security applications,” *arXiv preprint arXiv:2208.10605*, 2022.
- [81] S. M. Lundberg *et al.*, “A unified approach to interpreting model predictions,” *Advances in neural information processing systems*, vol. 30, 2017.
- [82] F. Tramer *et al.*, “On adaptive attacks to adversarial example defenses,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 1633–1645, 2020.
- [83] J. W. Stokes *et al.*, “Attack and defense of dynamic analysis-based, adversarial neural malware classification models,” *arXiv preprint arXiv:1712.05919*, 2017.

- [84] D. Li *et al.*, “Arms race in adversarial malware detection: A survey,” *ACM Comput. Surv.*, vol. 55, no. 1, nov 2021. [Online]. Available: <https://doi.org/10.1145/3484491>
- [85] X. Shu *et al.*, “Threat intelligence computing,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, 2018, pp. 1883–1898.
- [86] Z. Zhu *et al.*, “Chainsmith: Automatically learning the semantics of malicious campaigns by mining threat intelligence reports,” in *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2018, pp. 458–472.
- [87] H. Cai *et al.*, “Towards sustainable android malware detection,” in *Proceedings of the 40th International Conference on Software Engineering: Companion Proceedings*, ser. ICSE ’18. New York, NY, USA: Association for Computing Machinery, 2018, p. 350–351. [Online]. Available: <https://doi.org/10.1145/3183440.3195004>
- [88] X. Fu *et al.*, “On the deterioration of learning-based malware detectors for android,” in *2019 IEEE/ACM 41st International Conference on Software Engineering: Companion Proceedings (ICSE-Companion)*, 2019, pp. 272–273.
- [89] F. Maggi *et al.*, “Protecting a moving target: Addressing web application concept drift,” in *Proceedings of the 12th International Symposium on Recent Advances in Intrusion Detection*, ser. RAID ’09. Berlin, Heidelberg: Springer-Verlag, 2009, p. 21–40. [Online]. Available: https://doi.org/10.1007/978-3-642-04342-0_2
- [90] Imam *et al.*, “A survey of attacks against twitter spam detectors in an adversarial environment,” *Robotics*, vol. 8, no. 3, 2019. [Online]. Available: <https://www.mdpi.com/2218-6581/8/3/50>

APPENDIX A MALPROTECT

Model	Parameters
Logistic Regression (MalProtect-LR)	No configurable parameters.
Neural Network (MalProtect-NN)	4 fully-connected layers (128 (Relu), 64 (Relu), 2 (Relu), 2 (Softmax))

TABLE II: Decision models for MalProtect configurations.

APPENDIX B PREDICTION MODELS

We use six non-stateful defenses as the prediction models. These are evaluated without any stateful protection and then evaluated in combination with each stateful defense.

Defense	Configuration/parameters/setup
Defensive distillation [9]	Neural Network (128 (Relu), 64 (Relu), 32 (Relu), 2 (Softmax) with defensive distillation applied.
Ensemble adversarial training [10], [30] (NN-AT)	Neural Network (128 (Relu), 64 (Relu), 32 (Relu), 2 (Softmax)). Adversarially-trained up to 25% size of training data.
Morphence [69]	$n = 4$, $p = 3$, $Q_{max} = 1000$.
StratDef [30]	Variety-GT using same models as original paper. Assumed strong attacker and $\alpha = 1$.
Voting (Majority & Veto)	Using same models as StratDef.

TABLE III: Architectures of prediction models.

APPENDIX C VANILLA MODELS

The following vanilla models are used in some instances (e.g., to generate adversarial examples, see Section V).

Model	Parameters
Decision Tree	max_depth=5, min_samples_leaf=1
Neural Network	4 fully-connected layers (128 (Relu), 64 (Relu), 32 (Relu), 2 (Softmax))
Random Forest	max_depth=100
Support Vector Machine	LinearSVC with probability enabled

TABLE IV: Architectures of vanilla models.

APPENDIX D OTHER STATEFUL DEFENSES

L_0 defense. The L_0 defense measures the similarity between feature vectors using the L_0 distance. This is akin to similarity detection schemes that use L_p norms in other domains (e.g.,

[17]). As binary feature vectors are used in the malware detection domain, L_0 is the most appropriate measure of distance between two such feature vectors, X and X' [40], [44]. It measures the number of instances such that $X_i \neq X'_i$. For the L_0 defense, an attack is detected if there are queries with an L_0 distance of less than 10 as lower thresholds may easily miss attacks [17], [18], [16].

PRADA [18]. For PRADA, we use $\delta = 0.9$ for the threshold, as per the original paper [18].

Stateful Detection (SD) [16]. For SD, the threshold for detecting an attack is derived by calculating the k-neighbor distance for the 0.1 percentile of the training set. At prediction-time, if the mean distance of the k-nearest queries falls below the calculated threshold, an attack is detected. For this, we use $k = 50$ (as per the original paper). The calculated thresholds for detection are as follows:

AndroZoo	SLEIPNIR	DREBIN
94.64749019607844	44.350758426966294	34.37145577151924

TABLE V: Thresholds for detection for SD as per calculations based on original paper.

APPENDIX E PERMITTED PERTURBATIONS FOR ANDROID DATASETS

The AndroZoo [63] and DREBIN [64] datasets are based on the Android platform. Both datasets can be divided into eight feature families comprised of extracted static features such as permissions, API calls, hardware requests, and URL requests. Based on their family, features may be addable or removable during attacks to traverse the decision boundary, according to prior work and industry documentation (e.g., [45], [72], [73], [26], [13], [30]). However, it is imperative to preserve malicious functionality in the feature-space as a core constraint in this domain. For example, attacks cannot remove features from the manifest file nor intent filter, and component names must be consistently named. Therefore, Table VI summarizes the permitted perturbations for each feature family for these datasets. This is used to determine whether the perturbations performed by an attack are valid. For example, if a feature belonging to the S1 family is removed by an attack, its original value is restored as it is not permitted to be removed (see Section V).

	Feature families	Addition	Removal
manifest	S1 Hardware	✓	✗
	S2 Requested permissions	✓	✗
	S3 Application components	✓	✓
	S4 Intents	✓	✗
dexcode	S5 Restricted API Calls	✓	✓
	S6 Used permission	✗	✗
	S7 Suspicious API calls	✓	✓
	S8 Network addresses	✓	✓

TABLE VI: Permitted perturbations for Android datasets.

APPENDIX F QUERY ATTACK STRATEGIES

We use variations of attack strategies from prior work in our domain [8]. These are based on a software transplantation approach where an attacker makes perturbations based on benign samples. We modify prior attack strategies by transplanting multiple features per iteration (rather than one perturbation

per iteration). The differences between the black-box and gray-box strategies lie in how perturbations are applied. The black-box attack selects the features to perturb in a randomized manner, while the gray-box attack perturbs features based on their frequency in benign samples in a heuristically-driven approach. Recall that the black-box query attack strategy is used in Section VI and the gray-box strategy in Section VII.

For the adaptive attack against MalProtect (in Section VIII), the number of features that are added in a single perturbation is capped, and from the *removable features* of the query, $p\%$ of features are removed. This is to make queries as distinct from each other as possible while remaining within the general distribution of other queries.

Algorithm 1 Black-box query attack: for oracle O , malware sample X , set of (randomly-ordered) benign features F , maximum permitted queries n_{max} .

Input: O, X, F, n_{max}

- 1: $X' \leftarrow X, n \leftarrow 0$
- 2: **while** $O(X') = 1$ & $n < n_{max}$ & $n < F.length$ **do**
- 3: $X' \leftarrow AddFeature(X', F[n])$
- 4: $r \leftarrow RandomInteger(0, Length(F))$
- 5: $F_r \leftarrow ChooseRandomFeatures(r, F)$
- 6: $X' \leftarrow AddFeatures(X', F_r)$
- 7: $X' \leftarrow ValidatePerturbations(X, X')$
- 8: $n \leftarrow n + 1$
- 9: **if** $O(X') = 0$ **then return** *Success*
- 10: **return** *Failure*

Algorithm 2 Gray-box query attack: for oracle O , malware sample X , vector of sorted benign features \vec{s} , maximum permitted queries n_{max} .

Input: O, X, \vec{s}, n_{max}

- 1: $X' \leftarrow X, n \leftarrow 0$
- 2: **while** $O(X') = 1$ & $n < n_{max}$ & $n < \vec{s}.length$ **do**
- 3: $X' \leftarrow AddFeature(X', \vec{s}[n])$
- 4: $r \leftarrow RandomInteger(0, Length(\vec{s}))$
- 5: $F_r \leftarrow ChooseRandomFeatures(r, \vec{s})$
- 6: $X' \leftarrow AddFeatures(X', F_r)$
- 7: $X' \leftarrow ValidatePerturbations(X, X')$
- 8: $n \leftarrow n + 1$
- 9: **if** $O(X') = 0$ **then return** *Success*
- 10: **return** *Failure*

Algorithm 3 Adaptive attack: for oracle O , malware sample X , vector of sorted benign features \vec{s} , maximum permitted queries n_{max} , p percentage of features to remove, m maximum features to add in each iteration.

Input: $O, X, \vec{s}, n_{max}, m$

- 1: $X' \leftarrow X, n \leftarrow 0$
- 2: **while** $O(X') = 1$ & $n < n_{max}$ & $n < \vec{s}.length$ **do**
- 3: $X' \leftarrow AddFeature(X', \vec{s}[n])$
- 4: $r \leftarrow RandomInteger(0, m)$
- 5: $F_r \leftarrow ChooseRandomFeatures(r, \vec{s})$
- 6: $X' \leftarrow AddFeatures(X', F_r)$
- 7: $X' \leftarrow RemoveFeatures(X', p)$
- 8: $X' \leftarrow ValidatePerturbations(X, X')$
- 9: $n \leftarrow n + 1$
- 10: **if** $O(X') = 0$ **then return** *Success*
- 11: **return** *Failure*

APPENDIX G EXTENDED RESULTS

The extended results are located at: https://osf.io/sfvyn/?view_only=2caef2fd7ae416a8891ce3f3bd50d2d

APPENDIX H EVALUATING OTHER QUERY ATTACKS

We also examine the Boundary [6], HopSkipJump [7], and ZOO [71] query attacks. We demonstrate their inability to generate adversarial examples in this domain, as these attacks are designed for other domains and do not consider the constraints of the malware detection domain (functionality preservation and discretization of features). We use each attack to try and generate adversarial examples against the NN-AT model. For this, once the feature vector of a malware sample has been perturbed, we discretize the feature vectors and evaluate whether the perturbations are permitted for each dataset. As usual, any invalid perturbations are reversed.

Figure 9 shows that these attacks are unable to achieve evasion *at all*. Evidently, the perturbations used to cross the decision boundary are reversed. That is, invalid perturbations are consistently made that must be discretized and reversed to ensure functionality preservation in the feature-space. Meanwhile, in Figure 9, we also show that our black-box query attack strategy achieves 80+% evasion rate across the datasets.




Fig. 9: Evasion rate of additional query attacks against NN-AT model vs. n_{max} .