

Detection under Privileged Information (Full Paper)*

Z. Berkay Celik
Pennsylvania State University
zbc102@cse.psu.edu

Patrick McDaniel
Pennsylvania State University
mcdaniel@cse.psu.edu

Rauf Izmailov
Vencore Labs
rizmailov@appcomsci.com

Nicolas Papernot,
Ryan Sheatsley, Raquel Alvarez
Pennsylvania State University
{ngp5056,rms5643,rva5120}@cse.psu.edu

Ananthram Swami
Army Research Laboratory
ananthram.swami.civ@mail.mil

ABSTRACT

For well over a quarter century, detection systems have been driven by models learned from input features collected from real or simulated environments. An artifact (e.g., network event, potential malware sample, suspicious email) is deemed malicious or non-malicious based on its similarity to the learned model at runtime. However, the training of the models has been historically limited to only those features available at runtime. In this paper, we consider an alternate learning approach that trains models using “privileged” information—features available at training time but not at runtime—to improve the accuracy and resilience of detection systems. In particular, we adapt and extend recent advances in knowledge transfer, model influence, and distillation to enable the use of forensic or other data unavailable at runtime in a range of security domains. An empirical evaluation shows that privileged information increases precision and recall over a system with no privileged information: we observe up to 7.7% relative decrease in detection error for fast-flux bot detection, 8.6% for malware traffic detection, 7.3% for malware classification, and 16.9% for face recognition. We explore the limitations and applications of different privileged information techniques in detection systems. Such techniques provide a new means for detection systems to learn from data that would otherwise not be available at runtime.

KEYWORDS

Detection systems; privileged information; machine learning

1 INTRODUCTION

Detection systems based on machine learning are an essential tool for system and enterprise defense [49]. Such systems provide predictions about the existence of an attack in a target domain using information collected in real-time. The detection system uses this information to compare the runtime environment against known normal or anomalous states. In this way, the detection system “recognizes” when the environmental state becomes—at least probabilistically—dangerous. What constitutes dangerous is learned; detection algorithms construct models of attacks (or non-attacks) from past observations using a training algorithm. Thereafter, the detection systems use that model for detection at runtime. A limitation of this traditional approach is that it relies solely on the features (also referred to as inputs) that are available at runtime. In practice, many

features are too expensive to collect in real-time or only available after the fact—and are thus ignored for the purposes of detection.

Consider a recent event that occurred in the United States. In the Summer of 2017, the credit reporting agency Equifax fell victim to sophisticated cyber attacks that resulted in substantial exfiltration of personal information and intellectual property [8]. Working with the government staff and security analysts conducted a clandestine investigation. During that time, a vast amount of information was collected from networks and systems across the agency, e.g., network flows, system logs files and user activity. An analysis of the collected data revealed the presence of previously undetected advanced persistent threat (APT) actors on the agency’s network. Yet, the collected analysis is largely non-actionable by detection systems post investigation; because the vast array of derived features would not be available at runtime, they cannot be used to train Equifax’s detection systems.

In other contexts, features may be available at runtime but infeasible or undesirable to collect because of environmental or system constraints. For example, the collection of a large numbers of features in environments of mobile phones [6], Internet of Things (IoT) [45], sensor networks [23], embedded control systems [46], or ad-hoc networks [9] is often too slow or requires too many resources to be feasible in practice.

These examples highlight a challenge for future intrusion detection: *how can detection systems integrate intelligence relevant to an attack that is not available at runtime?* Here, we turn to recent advances in machine learning that support models that learn on a superset of features used at runtime [54, 55]. The goal of the work described in this paper is to leverage these additional features, called *privileged information* (features available at training time, but not at runtime), to improve the accuracy of detection. Using this approach designers and operators of detection systems can leverage additional effort during system calibration to improve detection models without inducing additional runtime costs.

Pioneered recently by Vapnik, Izmailov, and others, learning under privileged information eliminates the need for symmetric features in training and runtime—thereby expanding the space of learning models to include “ancillary” and “non-runtime” information. However, to date, the application of these techniques in practical domains has been limited, and within the context of security non-existent. In this work, we explore how this new mode of learning can be leveraged in detection systems. This requires an exploration of not only use of these new learning models but also their applicability to security domains and the requirements of

*A short version of this paper appears in ACM Asia Conference on Computer and Communications Security (ASIACCS) 2018.

those domains on feature engineering. Our experience in this effort over the last two years has demonstrated that blind application of privileged application can lead to poor detection—yet judicious and careful use can substantially improve detection quality.

More concretely, in this paper, we explore an alternate approach to training intrusion detection systems that exploit privileged information. We design algorithms for three classes of privileged-augmented detection systems that use: (1) *Knowledge transfer*, a general technique of extracting knowledge from privileged information by estimation from available information, (2) *Model influence*, a model-dependent approach of influencing the model optimization with additional knowledge obtained from privileged information; and (3) *Distillation*, an approach of summarizing the additional knowledge about privileged samples as class probability vectors. We further explore feature engineering in a privileged setting. To this end, we measure the potential impacts of privileged features on runtime models. Here, we use the degree to which a feature improves a model (*accuracy gain*—a feature’s additive contribution to accuracy) as a quality metric for selecting privileged features for a target model. We develop an algorithm and system that selects features that maximize model accuracy in the presence of privileged information. Finally, we compare the performance of privileged-augmented systems with systems without privileged information. We evaluate four recently proposed detection systems: (1) face authentication, (2) fast-flux bot detection, (3) malware traffic detection, and (4) malware classification. Our contributions are:

- We augment several diverse detection systems using three classes of privileged information techniques and explore the strength and weaknesses of these techniques.
- We present the first methods for feature engineering in privileged-augmented detection for security domains and identify inherent tensions between information utilization, detection accuracy, and model robustness.
- We provide an evaluation of techniques on a variety of existing detection systems on real-world data. We show that privileged information decreases the relative detection error of traditional detection systems up to 16.9% for face authentication, 7.7% for fast-flux bot detection, 8.6% for malware traffic detection, and 7.3% for malware classification.
- We analyze dataset properties and algorithm parameters that maximize detection gain, and present guidelines and cautions for the use of privileged information in realistic deployment environments.

After introducing the technical approach for detection in the next section, we consider several key questions:

- (1) How can the best features for a specific detection task be identified? (Section 4)
- (2) How does privileged-base detection perform against traditional systems? (Section 5)
- (3) How can we select the best privileged algorithm for a given domain and detection task? (Section 6)
- (4) What are the practical concerns in using privileged information for detection? (Section 7)

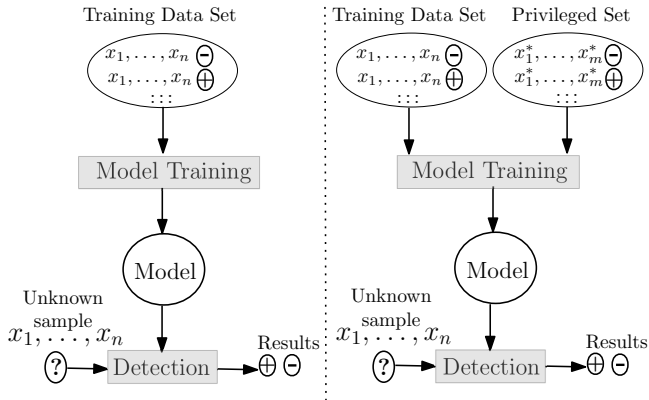


Figure 1: Overview of a typical detection system (left) and proposed solution (right): Given training data including both malicious (+) and benign samples (−), modern detection systems use training data to construct a model. To predict the class of an unknown sample, the system examines its features using the model to predict the sample as benign or malicious. In contrast, we construct a model both using training and privileged data, yet the model only requires the training data features to detect unknown sample as benign or malicious.

2 PROBLEM STATEMENT

Detection systems use traditional learning algorithms such as support vector machines or multilayer perceptrons (neural networks) to construct *detection models*. These models aim at learning patterns from historical data (also referred to as training data) to estimate an underlying dependency, structure or behavior of a system, process or environment. This *training data* is a collection of samples that includes a vector of *features* (e.g., packets per second, port number) and class *labels* (e.g., anomalous or normal). Once trained, runtime events (e.g., network event) are compared to the learned model. Without loss of generality, the model outputs a label (or label confidence) that most closely fits with those of the training data. The percentage of output labels that are correctly predicted for a sample set is known as its accuracy.

The quality of the detection system largely depends on features used to train models. In turn, the success of detection depends on explanatory variation behind the features that are used to separate an attack and benign sample. However, modern detection systems by construction assume that the features used to make predictions at runtime would be identical to those used for training (See Figure 1 (left)). This assumption restricts the model training to the features that are available at runtime to make predictions. As highlighted above, intelligence obtained from forensic investigations [56], data obtained through a human expert analysis [10], or features unable to be feasibly collected at runtime is simply not actionable. Juxtapose this with our goal of leveraging features in model training that are not available at runtime to improve detection accuracy (See Figure 1 (right)). Note that we do not focus on a specific detection task or domain. We begin in the following section by introducing the three approaches we use to integrate privileged information into detection models.

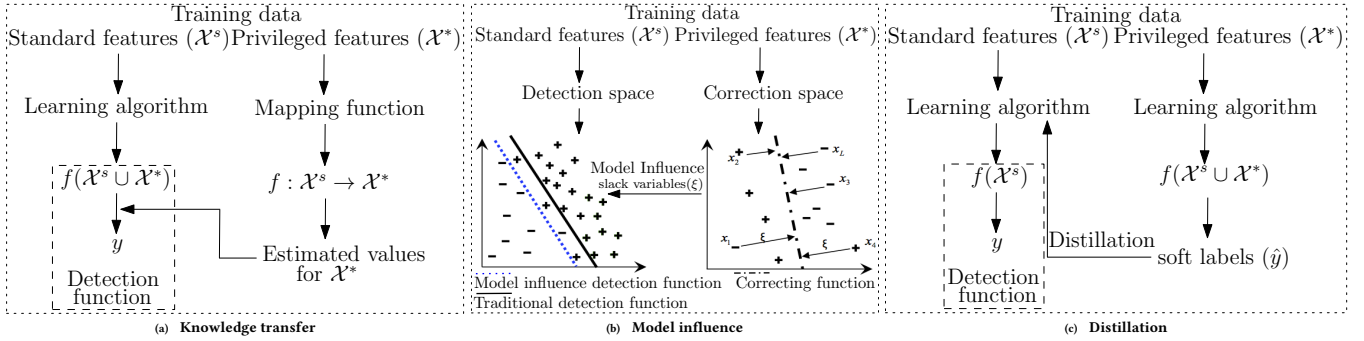


Figure 2: Schematic overview of approaches. Instead of making a detection model solely rely on standard features of a detection system, we also integrate privileged features into detection algorithms.

3 PRIVILEGED INFORMATION DETECTION

This section introduces three approaches to integrate privileged features into detection algorithms: knowledge transfer, model influence, and distillation. Figure 2 presents the schematic overview of approaches. Stated formally, we consider a conventional detection algorithm as a function $f: \mathcal{X} \rightarrow \mathcal{Y}$ that aims at predicting targets $y \in \mathcal{Y}$ given some explanatory features $x \in \mathcal{X}$. The models are built using a dataset containing pairs of features and targets, denoted by $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$. Following the definition of privileged information, we consider a detection setting where the features \mathcal{X} used for detection are split into two sets to characterize their availability for detection at training time. The *standard features* $\mathcal{X}^s = \{x_i, i = 1, \dots, n, x_i \in \mathbb{R}^d\}$ includes the features that are reliably available both at training and runtime (as in conventional systems), while *privileged features* $\mathcal{X}^* = \{x_i^*, i = 1, \dots, n, x_i^* \in \mathbb{R}^m\}$ have constraints that prevent using them for detection at runtime. More formally, we assume that detection models will be trained on some data $\{(x_i, x_i^*, y_i)\}_{i=1}^n$, and they will make detection on some data $\{x_j\}_{j=n+1}^{n+m}$. Therefore, our goal is creation of algorithms that efficiently integrate privileged features $\{x_i^*\}_{i=1}^n$ into detection models, without requiring them at runtime.

3.1 Knowledge Transfer

We consider a general algorithm to transfer knowledge from the space of privileged information to the space where the detection model is constructed [54]. The algorithm works by deriving a *mapping function* to estimate each privileged feature from a subset of standard features. The algorithm used to identify a mapping function f_i is described in Algorithm 1. The estimation is straightforward: the relationship between standard and privileged features is learned by defining each privileged feature x_i^* as a target and standard set as an input to a mapping function $f_i \in \mathcal{F}$ (lines 1-3). The mapping functions can be defined in the form of any function such as regression or similarity-based (we give examples in Section 5.1). The use of the mapping function allows a system to apply the same model learned from the complete set at training time with the union of standard and estimated privileged features on unknown samples (See Figure 2a). By using f_i , detection systems are able to construct the complete features with the correct values of privileged features

Algorithm 1: Knowledge Transfer Algorithm

Input : Standard training set $\mathcal{X}^s = \vec{x}_1, \dots, \vec{x}_L, \vec{x}_i \in \mathbb{R}^d$
Privileged training set $\mathcal{X}^* = \vec{x}_1^*, \dots, \vec{x}_L^*, \vec{x}_i^* \in \mathbb{R}^m$
 $f_i \in \mathcal{F} \leftarrow$ mapping function
 $\theta \leftarrow$ mapping function parameters

- 1 Find an optimal standard set $\hat{\mathcal{X}} \subseteq \mathcal{X}^s$, for all $x_i^* \in \mathcal{X}^*$
- 2 Select set $\hat{\mathcal{X}}_i$ for x_i^*
- 3 Mapping function f evaluates $(\hat{\mathcal{X}}_i, x_i^*)$ as given
 $x_i^* = f_i(\hat{\mathcal{X}}_i, \beta | \theta)$, $f_i \in \mathcal{F}_m$
- 4 Output β for x_i^* that minimizes f_i
- 5 At runtime use f_i to estimate x_i^*

at runtime—intuitively, each f_i is used to generate a synthetic feature that represents an estimate of a privileged feature (line 4). As a result, the accurate estimation of privileged features contributes to using complete and relevant features in a model training and, therefore, enhances the generalization of models compared to those trained solely on standard features. Note that the estimating power of f_i is bounded by the size and completeness of the training data (with respect to the privileged features), and thus the use of f_i in the model should be calibrated based on measurements of estimation quality (See Section 5.4 for details).

3.2 Model Influence

Model influence incorporates the useful information obtained from privileged features to the correction space of the detection model by defining additional constraints on the training errors (See Figure 2b) [54, 55]. Intuitively, the algorithm learns how privileged information influences outputs on training input feature vectors towards building a set of corrections for the space of inputs—in essence creating a correction function that takes as input runtime features and adjusts model outputs. Note that while we adapt model influence to support vector machines (SVM) herein, it is applicable to other ML techniques. More formally, consider a training datum that is generated from an unknown probability distribution $p(\mathcal{X}^s, \mathcal{X}^*, y)$. Our goal is to find a model $f: \mathcal{X}^s \rightarrow \mathcal{Y}$ such that the

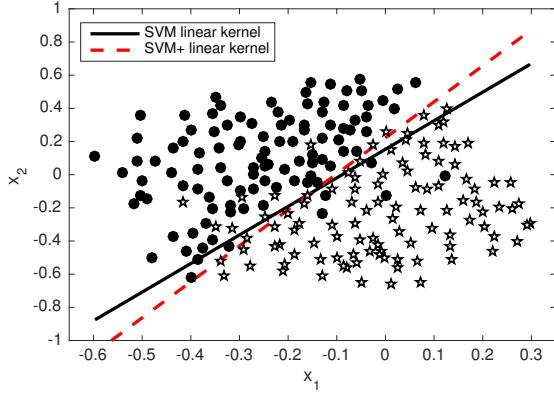


Figure 3: Model influence: The SVM+ uses privileged information to correct the decision boundary of the model.

expected loss is defined as:

$$R(\theta) = \int_{\mathcal{X}^s} \int_{\mathcal{Y}} \mathcal{L}((x^s, \theta), y) p(x^s, x^*, y) dx^s dy \quad (1)$$

Here the system is trained on standard and privileged features but only uses standard features at runtime. We consider the optimization problem of SVM in its dual form as shown in Equation 2 (labeled as primary detection objective) where α is Lagrange multiplier.

$$\begin{aligned} \mathcal{L}(w, w^*, b, b^*, \alpha, \delta) = & \underbrace{\frac{1}{2} \|w\|^2 + \sum_{i=1}^m \alpha_i - \sum_{i=1}^m \alpha_i y_i f_i}_{\text{main detection objective}} \\ & + \underbrace{\frac{\gamma}{2} \|w^*\|^2 + \sum_{i=1}^m (\alpha_i + \delta_i - C) f_i^*}_{\text{influence from privileged features}} \quad (2) \end{aligned}$$

We influence the detection boundary of a model trained on standard features $f_i = w^\top x_i^s + b$ at x_i^s with the correction function defined by the privileged features $f_i^* = w^{*\top} x_i^* + b^*$ at the same location (labeled as influence from privileged features). In this manner, we use privileged features as a correction function of the slack variables ξ_i defined in the objective of SVM. In turn, the useful information obtained from privileged features is incorporated as a measure of confidence for each labeled standard sample. The formulation is named as SVM+ and requires $O(\sqrt{n})$ samples to converge compared to $O(n)$ samples for SVM which is useful for systems with a sparse data collection [12, 54]. We refer readers to Appendix A for a complete formulation and implementation.

To illustrate, consider the 2-dimensional synthetic dataset presented in Figure 3, as well as the decision boundaries of two detection algorithms SVM (an unmodified Support Vector Machines) and SVM+ (the same SVM augmented with model influence correction). The use of privileged information in model training separates the classes more accurately because privileged features accurately transfer information to standard space, and the resulting model becomes more robust to the outliers. This approach may provide even better class separation in datasets with higher dimensionality.

To summarize, as opposed to knowledge transfer, we eliminate the task of finding the mapping functions between standard and

Algorithm 2: Distillation Algorithm

- Input** : Standard training set $\mathcal{X}^s = \vec{x}_1^s, \dots, \vec{x}_L^s, \vec{x}_i^s \in \mathbb{R}^d$
Privileged training set $\mathcal{X}^* = x_1^*, \dots, x_L^*, x_i^* \in \mathbb{R}^m$
Training labels \mathcal{Y}
 $T \leftarrow$ Temperature parameter, $T > 0$
 $\lambda \leftarrow$ Imitation parameter, $\lambda \in [0, 1]$
- 1 Learn a model f_s using $(\mathcal{X}^*, \mathcal{Y})$
 - 2 Compute soft labels as $s_i = \sigma(f_s(x_i^*)/T)$ for all $1 \leq i \leq L$
 - 3 Learn detection model using Equation 3 with the given data as $\{(\mathcal{X}^s, \mathcal{Y}), (\mathcal{X}^s, \mathcal{S})\}$
 - 4 Make detection using θ minimizing f_t for standard features x^s
-

privileged features. Thus, we reduce the problem of model training to a single unified task.

3.3 Distillation

Model compression or distillation are techniques to transfer knowledge from a complex Deep Neural Network (DNN) to a smaller one without loss of accuracy [24]. The motivation behind the idea suggested in [3] is closely related to knowledge transfer. The goal of the distillation is to use the class knowledge from both class labels (i.e., hard labels) and probability vectors of each (i.e., soft labels). The benefit of using class probabilities in addition to the hard labels is intuitive because probabilities of each class define a similarity metric over the classes apart from the samples' correct classes. Lopez-Paz et al. recently introduced an extension of model distillation used to compress models built on a set of features into models built on a different set of features [33]. We adapt this technique to detection algorithms.

We address the problem of privileged information using distillation as follows. First, we train a “privileged” model on the privileged set and labels whose output of the model is the vector of soft labels S . Second, we train a *distilled* model (used at runtime) by minimizing Equation 3, which learns a detection model by simultaneously imitating the privileged predictions of the privileged model and learning the targets of the standard set. The algorithm for learning such a model is presented in Algorithm 2 and outlined as follows:

$$f_t(\theta) = \operatorname{argmin}_{f \in F_t} \frac{1}{n} \sum_{i=1}^n \left(\overbrace{(1 - \lambda) \mathcal{L}(y_i, \sigma(f(x_i^s)))}^{\text{detection}} + \lambda \overbrace{\mathcal{L}(s_i, \sigma(f(x_i^s)))}^{\text{imitate privileged set}} \right) \quad (3)$$

We learn a privileged model $f_s \in F_s$ by using the privileged samples available at training time (line 1). We then compute the soft labels by applying the softmax function (i.e., normalized exponential) $s_i = \sigma(f_s(x_i^*)/T)$ (line 2). The output is a vector which assigns a probability to each class of the privileged samples. We note that class probabilities obtained from privileged model provide additional information for each class. Here, temperature parameter T controls the degree of class prediction smoothness. Higher T enables softer probabilities over classes and vice versa. As a final step, Equation 3 is sequentially minimized to distill the knowledge transferred from privileged features as a form of probability vectors (soft labels) into the standard sample classes (hard labels) (line 3). In Equation 3, the λ parameter controls the trade-off between privileged and standard features. For $\lambda \approx 0$, the objective approaches

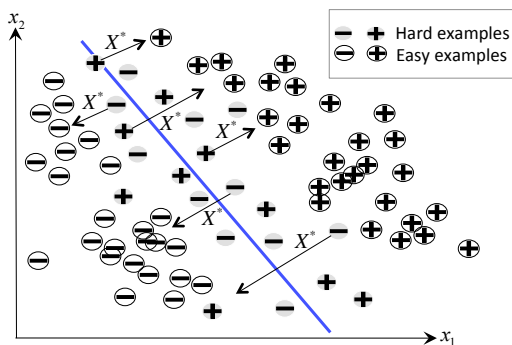


Figure 4: Visualizing hard and easy benign (-) and malicious (+) examples. We select privileged features that aims at increasing the detection of hard examples.

the standard set objective, which amounts to detection solely on standard features. However, as $\lambda \rightarrow 1$, the objective transfers the knowledge acquired by the privileged model into the resulting detection model. Therefore, learning from the privileged model does, in many cases, significantly improve the learning process of a detection model.

Distillation differs from model influence and knowledge transfer in at least two ways. First, while knowledge transfer attempts to estimate the privileged features with a representation of a mapping function, distillation is a trade-off between the privileged sample probabilities and standard sample class labels. Second, in contrast to model influence, distillation is independent of the machine learning algorithm (model-free), and its objective function can be minimized using a model of choice.

4 PRIVILEGED INFORMATION SYSTEMS

In this section, we explore algorithms for feature engineering (selecting privileged features for a detection task) and demonstrate their use in diverse experimental systems.

4.1 Selecting Privileged and Standard Features

The first challenge facing our model of detection is deciding which features should be used as privileged information. Asked another way, *given some potentially large universe of offline features, which are the most likely to improve detection?* To address this, we develop an iterative algorithm that selects features that maximize model accuracy. Selection is made on the calculated *accuracy gain* of each feature—a measure of the additive value of the features concerning an existing feature set for detection accuracy.

Our feature selection algorithm measures the potential impacts of privileged features that help detect the hard-to-classify examples. Generally speaking, easy examples fall in a distribution that can be explained by some set of model parameters, and hard examples do not precisely fit the model—and are either misclassified or near a decision boundary (See Figure 4) [50]. As a consequence, accurate classification of hard examples is one of the main challenges of practical systems, as they are the main source of detection errors due to the incorrect or insufficient information about normal or anomalous states of a system.

Algorithm 3: Selecting privileged features

Input : Standard feature set x^s
 Related privileged feature set x^*
 SVM Detection model J evaluating accuracy of hard-to-classify examples

- 1 Start with the standard set $Y_0 = x^s$
 - 2 Select the next privileged feature
 $x^+ = \operatorname{argmax}_{x^* \notin Y_k} [\mathcal{J}(Y_k + x^*)]$
 - 3 Update $Y_{k+1} = Y_k + x^+$; $k = k + 1$
 - 4 Go to 2
 - 5 Output Y including standard and selected privileged features
-

The first step of feature engineering—as is true of any detection task—is identifying all of the available features that potentially may be used for detection. Specifically, we collect the set of domain-specific features based on using domain knowledge and surveying the recent efforts in that domain. It is from that set that we will identify the privileged features to be used for training. Note that defining privileged features sometimes requires a level of domain expertise. However, trained security experts will find most privileged features straightforward after defining the runtime constraints on features. Additionally, the initial privileged set may include irrelevant features that carry little or no useful information for the target detection task; thus we identify the privileged set using Algorithm 3. The algorithm starts with standard features of a detection system and sequentially adds one privileged feature from the set which maximizes correct classification of hard examples, i.e., the feature whose addition to the existing set has the greatest positive impact on accuracy (measured accuracy gain). The accuracy gain of hard examples is found using SVM classifier (model J in algorithm 3). This process is repeated until the potential feature set is empty, a maximum number of features is reached, or the accuracy gain is below a threshold for usefulness.

Note that the quality of the selection process is a consequence of the training data used to calculate accuracy gain. If the training data is not representative of the runtime input distribution, the algorithm could inadvertently over or under-estimate the accuracy gain of a feature and thereby weaken the detection system [15]. Additionally, the accuracy gain can be computed via different feature selection algorithms. For instance, a more brute force approach that evaluates the impact of all standard and privileged feature combination can be used. This drastically increases computation overhead with the high dimensional datasets used in our evaluation. Note that this limitation is not unique to feature selection in this context, but applies to all feature engineering in extant detection systems.

4.2 Experimental Systems

In this section, we introduce four security-relevant systems for face authentication, fast-flux bot detection, malware traffic detection, and malware classification. We selected these experimental systems based on their appropriateness and diversity of their detection task. This diverse set of detection systems serves as a representative benchmark suite for our approaches. The following are the steps involved in constructing each system (discussed below):

- (1) Extract standard features of existing detection systems

System	Datasets and	Standard features	Incorporated privileged features	Detection time constraints on privileged features
❶	[27, 30, 58]	-Raw face images	-Bounding boxes and cropped versions of facial images	-Need of additional software for processing -Infeasible in energy and processing constrained sensors
❷	[16, 25, 29, 42, 59]	-Number of unique A and NS records in DNS packets -Network, processing and document fetch delay	-Edit distance, KL divergence and Jaccard index (domain names) -Time zone entropy of A and NS records in DNS packets -Euclidean distance between server IP and NS address -Number of distinct Autonomous systems and networks	-Processing overhead of whitelist domains -IP coordinate database processing overhead -Time consuming WHOIS processing
❸	[17, 18, 36, 60]	-Data bytes divided by the total number of packets -Total number of RTT samples found -The count of all packets with at least a byte payload -The median of total IP packets	-Source and destination port numbers -Byte frequency distribution in packet payload -Total connection time -Total number of packets with URG and PUSH flag set	-Adversary easily change them in subsequent malware versions -Payload encryption in subsequent malware versions
❹	[2, 19]	-Frequency count of hexadecimal duos in binary files	-Frequency count of distinct tokens in metadata log	-Software-dependency of obtaining assembly source code -Computational overhead and error-prone feature acquisition

Table 1: Description of detection systems: ❶ Face Authentication, ❷ Fast-flux Bot Detection, ❸ Malware Traffic Detection, and ❹ Malware Classification. Appendix B provides details on the standard and privileged features used throughout.



Figure 5: Example of face authentication features. Original image for standard features (Left), cropped (Middle), and funneled (Right) images used for privileged features.

- (2) Add new privileged features by identifying detection-time constraints on the features
- (3) Use the algorithm in preceding section to calibrate the detection system with standard and privileged features

Through this process, we construct their privileged-augmented systems with application of approaches that is used for the validation in the following section. Table 1 summarizes the experimental systems and the standard and privileged features selected. Additional details about these systems and their features are presented in Appendix B.

Experimental System: Face Authentication - To explore the efficiency of approaches in image domains, we modeled a user authentication system based on recognition of facial images. Our goal is to recognize an image containing a face with an identifier corresponding to the individual depicted in the image. We use images from a public dataset that includes face images labeled with each person’s name [27]. We build the features from 1348 facial images with at least 50 images per user.

Privileged information. It is recently found that face recognition systems used for access control in particular energy and computation constrained camera sensors can be easily bypassed by an attacker [20]. In this, the lack of useful features or number of images used to train the systems is the main reason of duping the systems into falsely authenticate/recognize users. We use two types of privileged features for each image in addition to the original images in model training: cropped and funneled versions of the images (See Figure 5) [30, 58]. These images provide additional information for a users’ face by image aligning and localizing [26]. While it is technically possible these features can be obtained by an aid of software or human expert at runtime, they are much more




Figure 6: Excerpt from hexadecimal representations (right), and assembly view (left) of an example malware. Selected byte bigrams and tokens for this malware is shown in boxes.

likely to not be available in low energy and slow processing sensors (and thus we define them as privileged).¹

Experimental System: Fast-flux Bot Detection - The Fast-flux bot detector is used to identify hosts that use fast-changing DNS entries to hide the existence of server hosts used for malicious activities. The raw data consists of 4GB DNS requests of benign and active fast-flux servers collected in early 2013 [16]. We build a detection system by using the 19 features used in recently proposed botnet detectors [29, 42, 59]. This system relies on features obtained from domain names, DNS packets, packet timing intervals, WHOIS domain lookup and IP coordinate database. The resulting dataset includes many features to increase separation of Content Delivery Networks (CDNs) from fast-flux servers, as similarities between them are the main source of detection errors.

Privileged information. In this system, even though the complete features are relevant for fast-flux detection, obtaining some features at runtime entails computational delays. For example, processing of WHOIS records, maintaining up-to-date IP coordinate database and whitelist of domain names takes several minutes/hours. Thus, we define eleven features obtained from these sources as privileged to assure real-time detection.

Experimental System: Malware Traffic Detection - Next, we modeled a malware traffic anomaly detection system based on network flow statistics used in recent detection systems [17, 22, 36]. The system aggregates 20 flow features for detecting botnet command and control (C&C) activity among benign applications. For instance, the ratio between maximum and minimum packet size from server to client and client to server find out to be a distinctive observation between benign and malicious samples. We add 173 botnet traffic of Zeus variants that is used for spam distribution, DDoS

¹We interpret the accuracy gain as a defense for hardening misclassification of users.

System	Approach	Relative Gain over Traditional Detection			Parameter Optimization?	Detection time Overhead?
		Accuracy	Precision	Recall		
Fast-flux Bot Detection	Knowledge Transfer	7.7%	5.3%	3.4%	✓	✓
Malware Classification	Model Influence	7.3%	2.2%	3.7%	✓	✗
Malware Traffic Analysis	Distillation	8.6%	2.2%	5%	✓	✗
Face Authentication	Distillation	16.9%	9.3%	9.2%	✓	✗

Table 2: Summary of accuracy, precision and recall gain via privileged information-augmented systems. The best resulting approach for each detection system is illustrated. (See the next two sections for more details on approach comparison and runtime overhead of approaches.)

attacks, and click fraud [21, 48] to the 1553 benign applications (web browsing, chat, email, etc.) of Lawrence Berkeley National Laboratory (LBNL) [39] and University of Twente [5].

Privileged information. In this system, the authors eliminate the features that can be readily altered by an attacker, as the model trained with tampered features allows an attacker to manipulate the detection results easily [35, 60]. The impact of altering the features at run time on detection accuracy is recently studied [17]. For instance, consider that destination port numbers or packet interval time are used as a feature. An adversary may “easily” change them in subsequent malware versions to evade detection systems. Also, the authors do not use payload content to obtain features because the attacker can use encrypted traffic to prevent deep packet inspection. Thus, we deem such eight features as a privileged, as inference does not consider their tampered values at runtime.

Experimental System: Malware Classification - The Microsoft malware dataset [19] is an up-to-date publicly available corpus. The dataset includes nine malware classes including hexadecimal representation of the malware’s binary content, and a class representing one of nine family names. The dataset used in our experiments includes 1746 malware samples extracted from 200GB malware files. We build a real-time malware classification system by using the binary content file. Following a recent malware classification system [2], we construct features by counting frequencies of each hexadecimal duos (i.e., byte bigrams). These features found out to provide distinctive between different families because of exploiting the code dissimilarities among families.

Privileged information. This dataset also includes a metadata manifest log file. The log file contains information such as memory allocation, function calls, strings, etc.. The logs along with the malware files can be used for classifying malware into their respective families. Thus, similar to the byte files, we obtain the frequency count of distinct tokens from asm files such as `mov()`, `cmp()` in the text section (See Figure 6). These tokens allow us to capture execution differences between different families [2]. However, in practice, obtaining features from log files introduces significant overheads in the disassembly process. Further, various types or versions of a disassembler may output byte sequences differently. Thus, this process may result in inaccurate and slow feature processing in real-time automated systems [37]. To address these limitations, we include features from disassembler output as privileged for accurate and fast classification.

5 EVALUATION

In this section, we explore the following questions:

	Knowledge Transfer	Model Influence	Distillation
Fast-flux Bot Detection	Section 5.1	Section 5.4	Section 5.4
Malware Traffic Detection	Section 5.4	Section 5.2	Section 5.4
Face Authentication	—	—	Section 5.3
Malware Classification	Section 5.4	Section 5.4	Section 5.4

Table 3: Summary of validation experiments.

- (1) *How much does privileged-augmented detection improve performance over systems with no privileged information?* We evaluate the accuracy, precision, and recall of approaches, and demonstrate the detection gain of including privileged features.
- (2) *How well do the approaches perform for a given domain and detection task?* We answer this question by comparing the results of approaches and present guidelines and cautions for appropriate approach calibration to maximize the detection gain.
- (3) *Do approaches introduce training and detection overhead?* We report model learning and runtime overhead of approaches for realistic environments.

Table 2 shows the summary of the privileged information augmented systems, and Table 3 identifies the validation experiments described throughout. As detailed throughout, we find that the use of privileged information can improve—often substantially—detection performance in the experimental systems.

Overview of Experimental Setup. We compare performance of privileged-augmented systems against two baseline (non-privileged) models: the standard set model and the complete set model. The *standard set model* is a conventional detection system that does not include the privileged features for training or runtime, but uses all of the standard features. The *complete set model* is a conventional system that includes all the privileged and standards features for training or runtime. Note that the ideal privileged information approach would have similar performance as the complete set.

To learn standard and complete set models, we use classifiers of Random Forest (RF) and Support Vector Machines (SVM) with a radial basis function kernel. These classifiers give better performance in the previously introduced systems and are also preferred by the system authors. The parameters of the models are optimized with exhaustive or randomized parameter search based on the dataset size. All of our experiments are implemented in Python with the scikit-learn machine learning library or MATLAB with the optimization toolbox and run on Intel i5 computer with 8 GB RAM. We give the details of the implementation of privileged-augmented systems while presenting the calibration of approaches in Section 6.

		Fast-flux Bot Detection (FF)		
	Model	Accuracy	Precision	Recall
Complete Set	RF	99±0.9	99.4	99.4
	SVM	99.4±0.3	99.3	100
Standard Set	RF	96.5±2.6	98.7	96.8
	SVM	95±2.3	94.4	95.5
Similarity (KT)	RF	98.6±1.3	99.3	98.7
	SVM	96.7±1.2	98.7	97.1
Regression (KT)	RF	98.3±1.4	99	98.7
	SVM	96±1.1	98.7	96.1

Table 4: Fast-flux Bot Detection knowledge transfer (KT) results. All numbers shown are percentages. The best result of the knowledge transfer options is highlighted in boldface.

We show detection performance of complete and standard set models and compare their results with our approaches based on three metrics: accuracy, recall, and precision. We also present the false positives and negatives when relevant. *Accuracy* is the sum of the true positive and true negatives over a total number of samples. *Recall* is the number of true positives over the sum of false negatives and true positives, and *precision* is the number of true positives over the sum of false positives and true positives. Higher values of accuracy, precision, and recall indicates a higher quality of the detection output.

5.1 Knowledge Transfer

Our first set of experiments compares the performance of privileged-augmented detection system using knowledge transfer (KT) over standard and complete set models. In this experiment, we classify domain names into benign or malicious in fast-flux bot detection system (FF).

To realize KT, we implement two mapping functions to estimate the privileged features from a subset of standard features: *regression-based* and *similarity-based*. We find that both mapping functions learn the patterns of a training data and mostly suffice for the derivation of the nearly precise estimated privileged features. First, a polynomial regression function is built to find a coefficient vector $\beta \in \mathbb{R}^d$ such that there exists a $x_i^* = f_i(x^s, \beta) + b + \epsilon$ for some bias term b and random residual error ϵ . The resulting function is then used to estimate each privileged feature at detection time given the standard features as an input. We use polynomial regression that fits a nonlinear relationship to each privileged feature and picks the one that minimizes the sum of squares error. To evaluate the effectiveness of the regression, we implement a second mapping function named weighted-similarity. This function is used to estimate the privileged features from the most similar samples in the training set. We first find the k most similar subset of standard samples that are selected by using the Euclidean distance between an unknown sample and training instances. Then, the privileged features are replaced by assigning weights that are inversely proportional to the similarities of their neighbors. We note that other distance metrics gives worse accuracy than Euclidean distance for the studied datasets.

Table 4 shows the accuracy of Random Forest Classifier (RF) and Support Vector Machines (SVM) on standard, complete models, and KT in the form of multiple regression and weighted-similarity.

		Malware Traffic Detection		
	Model	Accuracy	Precision	Recall
Complete Set	RF	98.7±0.3	99.7	98.9
	SVM	95.6±1.2	98.8	94.6
Standard Set	RF	92±3	97.4	95.3
	SVM	89.2±0.6	94	94.6
Model Influence	SVM+	94±1.4	94.8	98.8

Table 5: Malware Traffic Detection model influence results. All numbers shown are percentages.

The average accuracy of ten independent runs of stratified cross-validation is given by measuring the difference between training and validation performance with parameter optimization (e.g., k parameter in similarity). The complete model accuracy of both classifiers is close to 99%. Note that baseline performance is obtained by always guessing the most probable class yields 68% accuracy.

We found that mapping functions are effective in finding a nearly precise relation between standard and privileged features. This decreases the expected misclassification rate on average both in false positives and negatives over benchmark detection with no privileged features. Both KT mapping options come close to the complete model accuracy on FF dataset (1% less accurate), and significantly exceeds the standard set accuracy (2% more accurate). The results confirm that regression and similarity are more effective at estimating privileged features than solely using the standard features available at runtime.

5.2 Model Influence

Next, we evaluate the performance of model influence-based privileged information in detecting Zeus botnet in real-world web (HTTP(S)) traffic. Here, the system attempts to detect the malicious activity of a Zeus botnet that connects to C&C centers and filters private data. Note that the Zeus botnet uses HTTP mimicry to avoid detection. As a consequence, the sole use of standard features makes detection of Zeus difficult, resulting in high detection error (where Zeus traffic is mostly classified as legitimate web traffic). To this end, we include privileged features of packet flags, port numbers, and packet timing information from packet headers (See Appendix B for the complete list of features). We observe that while these features can be spoofed by adversaries under normal conditions, using them as privileged information may counteract spoofing (because inference does not consider their runtime value).

We evaluate accuracy gain of model influence over the standard model. We use a polynomial kernel in the objective function to perform a non-linear classification by implicitly mapping the features into a higher dimensional feature space. We note that we avoid overfitting by tuning of the regularization parameter. Table 5 presents the impact of model influence on accuracy, precision, and recall and compares with standard and complete models. We found that using the privileged features inherent to malicious and benign samples in model training systematically better separates the classes. This positive effect substantially improves both false negative and false positive rates. The accuracy of model influence is close to the optimal accuracy and reduces the detection error on average 2% over RF trained on the standard set. This positive effect is more observable in SVM, and the accuracy gain yields 4.8%.

5.3 Distillation

We evaluate distillation on the experimental face authentication system. The standard features of a face authentication system consist of original images of users (i.e., background included) with 3 RGB channels. We obtain the privileged features of each image from funneled and downscaled bounding boxes of images. In this way, we better characterize each user’s face by specifying the face localization and eliminating the background noise. It is important to note that background of images may unrealistically increase the accuracy because background regions may contribute to the distinction between images. However, we verify by manual checking that the images in our training set do not suffer from this effect. In the following experiments, we construct both standard and privileged models using a deep neural network (DNN) with two hidden layers of 10 rectifiers linear unit with a softmax layer for each class. This type of architecture is commonly applied in computer vision and provides superior results for image-specific applications. We train the network with ten runs of 400 random training samples.

Figure 7 plots the average distillation accuracy with various temperature and imitation parameters. We show the accuracy of standard (dotted) and privileged set (dotted-dashed) models as a baseline. The resulting model achieves an average of 89.2% correct classification rate on the privileged set, which is better than the standard set with 66.5%. We observe that distilling the privileged set features into our detection algorithm gives better accuracy than standard set accuracy with optimal T and λ parameters. The accuracy is maximized when $T = 1$, the gain is on average 6.56%. The best improvement is obtained when $T = 1$ and $\lambda = 0.2$ with 11.2% increase over the standard set model accuracy. However, increases in T negatively affect detection accuracy. This is because as T increases, the objective function puts more weight on learning from the standard features which upsets the trade-off between standard and privileged features.

5.4 Comparison of Approaches

Next, we compare the relative performance of the approaches on three data sets.² Distillation is implemented using Deep Neural Networks (DNN), and regression and weighted similarity mapping functions are used for knowledge transfer. Table 6 presents the results of knowledge transfer (We report similarity results as it yields better results than regression), model influence and distillation and compares against complete and standard models. The accuracy, precision and recall gain of the best resulting approach for each detection system is summarized in Table 2.

The accuracy of model influence, distillation and knowledge transfer on the fast-flux detector and malware traffic detection is stable. All approaches yield accuracy similar to the ideal accuracy of the complete set model, and often the increased accuracy is the result of the correct classification of true positives (intrusions). This results in on average up to 5% relative gain in recall with the model influence and distillation over conventional models. In contrast, knowledge transfer often increases the number of samples detected by systems as being actually malicious (e.g., 99.3% precision in the fast-flux detector), meaning that the number of false alarms is

²We do not compare performance on face recognition because processing the number of the input features (e.g., pixels) was intractable for several solutions.




Figure 7: Distillation impact on accuracy with privileged features in face authentication system. We plot standard and privileged set accuracy as a baseline. Temperature (T) and imitation (λ) parameters are quantified on various values to show their impact on accuracy.

reduced over conventional detection. The results confirm that the approaches are able to balance the conventional detection and its accuracy by using privileged features inherent to both benign and malicious samples: either reducing the false positives or negatives. We note that these results are obtained after carefully tuning the model parameters. We further discuss their parameter tuning in the following section and impact of results on systems in section 7.

Distillation is easy to apply as its objective is independent of the machine learning algorithm (model-free) and often yields better results than other approaches. Its quality as a detection mechanism becomes more apparent when its objective function is implemented with deep neural networks (DNN) with a nonlinear objective [33]. This makes distillation give better results on average than other approaches. On the other hand, the design and calibration of model influence detection require additional effort and care in tuning parameters—in the experiments, this additional effort yields strong detection (as 94.6% in malware classification). Note that when the dataset includes a large number of privileged features or samples, training of model influence takes significantly more time compared to other approaches (See next section).

Finally, it is important to note that while knowledge transfer accuracy gain for fast-flux detection and malware traffic analysis is similar to other approaches, its malware classification results are inconsistent (i.e., 83.5% average accuracy with an 11.2% standard deviation). Neither regression nor similarity mapping functions were able to predict the privileged features near precisely, in turn, they slightly degrade the accuracy (7-8%) on both RF and SVM standard set models. This observation confirms the need to find and evaluate an appropriate mapping function for the transfer of knowledge discussed in Section 3.1. In this particular dataset, the mapping functions fail to find a good relation between standard and privileged features. Regression suffers from overfitting to uncommon data points and similarity lacks fitting data points that distinctly lie an abnormal distance from the range of standard features (confirmed by an increase in the sum of square errors of estimated and true values for the privileged features). We remark that derivation of more advanced mapping functions may solve this problem. Further, model influence and distillation solve this by eliminating the use of

	Fast-flux Bot Detection			Malware Traffic Detection			Malware Classification			
	Model	Acc.	Pre.	Rec.	Acc.	Pre.	Rec.	Acc.	Pre.	Rec.
Complete Set	RF	99±0.9	99.4	99.4	98.7±0.3	99.7	98.9	96.6±1.2	99.3	95.2
	SVM	99.4±0.3	99.3	100	95.6±1.2	98.8	94.6	95.7±1	98.6	94.6
Standard Set	RF	96.5±2.6	98.7	96.8	92.9±3	97.4	95.3	91.2±1	91.3	94.4
	SVM	95±2.3	94.4	95.5	89.2±0.6	94	94.6	91.8±1.1	93.2	93.6
KT (Similarity)	RF	98.6±1.3	99.3	98.7	93.3±2.1	94.4	98.4	90.1±2.2	97.3	88
	SVM	96.7±1.2	98.7	97.1	92.6±0.9	95.1	96	83.5±11.2	88.3	85.6
Model Influence	SVM+	97.3±1.3	97	99.3	94±1.4	94.8	98.8	94.6±2.3	93.3	97.8
Distillation	DNN	97.5±0.3	97.4	99.3	95.7±0.6	96.1	99.3	92.6±0.7	92.6	95.3

Table 6: Summary of results: Accuracy (Acc.), Precision (Pre.) and Recall (Rec.) The best result for each detection system is highlighted in bold.

mapping functions and including standard and privileged feature dependency into their objectives.

Therefore, based on the above observations, the approaches are in need for a calibration based on the domain and task-specific properties to maximize the detection gain, as explored next.

6 LIMITATIONS

In this section, we discuss the required dataset properties, algorithm parameters, training, and runtime overhead of using privileged information for detection. We also present guidelines and cautionary warnings for use of privileged information in realistic deployment environments. A summary of an approach selection criteria is presented in Table 7.

Model Dependency. Model selection is a task of picking an appropriate model (e.g., classifier) to construct a detection function from a set of potential models. Knowledge transfer can be applied to a model of choice, as privileged features are inferred with any accurately selected mapping function. Distillation requires a model with a softmax output layer for obtaining probability vectors. However, we adapt model influence to SVM’s objective function.

Detection Overhead. The mapping functions used in knowledge transfer may introduce detection delays while estimating the privileged features. For instance, weighted similarity introduced in Section 3.1 defers estimation until detection without learning a function at training time (i.e., lazy learner). This may introduce a detection bottleneck if dataset includes a large number of samples. To solve this problem, we apply stratified sampling to reduce the size of the dataset. Furthermore, constructing mapping functions at training time such as regression-based minimize the delay of estimating privileged features. For instance, in our experiments, weighted-similarity is used to estimate ten privileged features of 5K training samples in less than a second delay on 2.6GHz 2-core Intel i5 processor with 8GB RAM. Regression reduces this value to milliseconds. Therefore, if delay at runtime is the primary concern, we suggest using model influence and distillation for learning the detection model, as they introduce no overhead at runtime.

Model Optimization. To obtain the best performance results, the parameters and hyperparameters of approaches need to be carefully tuned. For instance, fine-tuning of temperature and imitation parameters in distillation and kernel function hyperparameters in model influence approaches may increase the detection performance. Similar to conventional detection, the number of parameters

Approach	Model dependency	Detection time overhead	Model optimization	Training time overhead
Knowledge Transfer	✗	✓	⊖	⊖
Model Influence	✓	✗	●	●
Distillation	✗	✗	⊖	⊖

Legend: ✓ yes ✗ no ⊖ model dependent ● relatively higher

Table 7: Guideline for approach selection.

required to be optimized both for knowledge transfer and generalized distillation can be determined a priori based on the selected model. However, model influence has twice as many parameters as SVM—two kernel functions are used simultaneously to learn detection boundary in standard and privileged feature spaces. We apply grid search for small training sets and evolutionary search for large-scale datasets for parameter optimization.

Training Overhead. Training set size affects the time required by model learning. The amount of additional time needed to run both knowledge transfer and generalized distillation is negligible, as they require similar models as existing systems apply. However, the objective function of model influence may become infeasible or take a long time when the dimension of the feature space is very small, or dataset size is quite large. For instance, in our experiments, distillation and knowledge transfer train 1K samples with 50 standard and privileged features on the same machine used for one minute including optimal parameter search. Model influence takes on average 30 minutes on the same machine used for measuring detection overhead. Packages that are designed specifically for solving the quadratic programming (QP) problems (e.g., MATLAB quadprog() function) can be used instead of general solvers such as convex optimization package CVX to reduce the training time. Further, specialized spline kernels can be used to accelerate the computation [54]. We give a specific implementation of model influence in such packages in Appendix A.

7 DISCUSSION

Our empirical results show that approaches reduce both false positives and negatives over the systems solely built on their standard features. In a security setting, a false positive makes it extremely difficult for the analyst examining the reported incidents only to identify the mistakenly triggered benign events correctly. It is not surprising, therefore, that recent research focuses on post-processing of the alerts to produce a more qualitative alert set useful to the human analyst [51]. False negatives, on the other hand, have the

potential to cause catastrophic damage to both users and organizations: even a single compromised system can cause serious security breaches. For instance, in malware traffic and fast-flux bot detection, a false negative may cause a bot filter private data to a malicious server. In the case of malware classification and face authentication, it undermines the integrity of a system by misclassifying malware into another family or recognizing the wrong user. Thus, in no small way, improvement in false positives and negatives of these systems does matter in operational settings, improving reliable detection.

7.1 Uses of Privileged Information

Privileged information is not restricted to the domains discussed above, but are readily adaptable problems and machine learning settings. For example, privileged information can be adapted to settings with unsupervised, regression, and metric learning [31]. With respect to detection, we consider several illustrative uses of privileged information below:

Mobile Device Security - The growth of mobile malware requires robust malware detectors on mobile devices. Current systems collect data for numerous type of attacks; however, exhaustive data collection at runtime can have high energy costs and induce noticeable interface lag. As a consequence, users may disable the detection mechanism [6]. We note that high-cost features can be defined as privileged information to combat this problem.

Enterprise Security - Enterprise systems use audit data generated from a diverse set of devices and information sources for analysis [10]. For instance, SIEM products collect data from hosts, applications, and network devices in incredible volumes (e.g., 100K events per second yielding to 42 TB of compressed data). These massive datasets are mined for patterns identifying sophisticated threats and vulnerabilities. However, systems may be overwhelmed by feature collection and processing at runtime which makes their collection impractical for many settings. In such cases, features involving complex and expensive data collection can be defined as privileged to balance the real-time costs and accuracy.

Privacy Enhanced Detection - Many detection processes require the collection of privacy-relevant features, e.g., pattern and substance of user network traffic, use of the software [11, 13]. Hence, it is important to reduce the collection and exposure of such data-legal and ethical issues may prevent continuously monitoring them in their original form. In these cases, a set of features can be defined as privileged to eliminate the requirement of obtaining and potentially retaining privacy-sensitive features at runtime from users and environments.

7.2 Privileged Information as a Defense

We also posit that privileged information can be used as a defense mechanism in adversarial settings. More specifically, the key attacks targeting machine learning are organized into two categories based on adversarial capabilities [28]: (1) *causative (poisoning) attacks* in which an attacker controls the training data by injecting well-crafted attack samples to control the prediction results, and (2) *explanatory (evasion) attacks* in which attacker manipulates the malicious samples to evade detection. For the former, privileged features adds an extra step for the attacker to pollute the training data because the attacker needs to dupe the data collection into

including polluted privileged samples in addition to the standard samples—which for many systems including online learning would potentially be much more difficult. For the latter, privileged features may make detection systems more robust to the adversarial samples because privileged features cannot be controlled by the adversary when producing malicious samples [32]. Moreover, because the model is hidden from the adversary they cannot know the influence of these features on the model [7]. As a proof of concept, recent works have used distillation of standard features as a defense mechanism against adversarial perturbations in DNNs [40]. In future work, we plan to further evaluate privileged information as a mechanism to harden machine learning systems.

8 RELATED WORK

Domain-specific feature engineering has been a key effort within the security communities. For example, researchers have previously used specific patterns to group malware samples into families [2, 38, 43], have explored using DNS information to understand and predict botnet domains [1, 4, 14, 59], and have analyzed network and system level features to identify previously unknown malware traffic [17, 36, 44]. Other works have focused on user authentication from facial images [41, 52]. We view our efforts in this paper to be complementary to these and related works. Features in these works can be easily enhanced with privileged information in detection algorithms to strike a balance between accuracy and the cost or availability constraints at runtime.

The use of privileged information has recently attracted attention in a few others areas such as computer vision, image processing, and even finance. Wang et al. [57] and Sharmanska et al. [50] derived privileged features from images in the form of annotator rationales, object bounding boxes, and textual descriptions. Ribeiro et al. used annual turnover and global balance values as privileged features for enhancing the financial decision-making [47]. However, their approaches are not designed to model security-relevant data and do not consider feature engineering but rather to determine if there is a possibility of application to a domain specific information.

9 CONCLUSIONS

We have presented a range of techniques to train detection systems with privileged information. All approaches use features available only at training time to enhance the accuracy of detection models. We consider three approaches: (a) knowledge transfer to construct mapping functions to estimate the privileged features, (b) model influence to smooth the detection model with the useful information obtained from the privileged features, and (c) distillation using probability vector outputs obtained from the privileged features in the detection objective function. Our evaluation of several detection systems shows that we can we improve the accuracy, recall, and precision regardless of their high detection performance using privileged features. We also presented guidelines for approach selection in realistic deployment environments.

This work is the first effort at developing detection under privileged information by exploring feature engineering, algorithms, and environmental calibration. The capability afforded by this approach will allow us to integrate forensic and other auxiliary information that, to date, has not been actionable for detection. In the future, we

will explore various environments and evaluate its ability to promote resilience to adversarial manipulation in detection systems. In this way, we will explore new models and systems using privileged features to promote lightweight, accurate and robust detection.

ACKNOWLEDGMENTS

Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

REFERENCES

- [1] Manos A. et al. 2012. From throw-away traffic to bots: detecting the rise of DGA-based malware. In *USENIX Security*.
- [2] Mansour A. et al. 2015. Novel feature extraction, selection and fusion for effective malware family classification. *arXiv preprint arXiv:1511.04317*.
- [3] Jimmy B. and Rich C. 2014. Do deep nets really need to be deep?. In *Advances in neural information processing systems*.
- [4] Leyla B., E. Kirda, C. Kruegel, and Marco B. 2011. EXPOSURE: Finding malicious domains using passive DNS analysis. In *NDSS*.
- [5] R. Barbosa, R. Sadre, A. Pras, and R. Meent. 2010. University of Twente traffic traces data repository. In *University of Twente Tech Report*.
- [6] J. Bickford et al. 2011. Security versus energy tradeoffs in host-based mobile malware detection. In *Mobile systems, applications, and services*.
- [7] B. Biggio, G. Fumera, and F. Roli. 2014. Pattern recognition systems under attack: Design issues and research challenges. *International Journal of Pattern Recognition and Artificial Intelligence*.
- [8] Equifax Data Breach. 2018. <https://en.wikipedia.org/wiki/Equifax>. (2018). [Online; accessed 15-January-2018].
- [9] I. Butun et al. 2014. A survey of intrusion detection systems in wireless sensor networks. *IEEE Communications Surveys & Tutorials*.
- [10] A. A. Cardenas, P. K. Manadhata, and S. P. Rajan. 2013. Big data analytics for security. *IEEE System Security*.
- [11] Z. B. Celik, H. Aksu, A. Acar, R. Sheatsley, A. S. Uluagac, and P. McDaniel. 2017. Curie: Policy-based Secure Data Exchange. *arXiv:arXiv:1702.08342*
- [12] Z. B. Celik, R. Izmailov, and P. McDaniel. 2015. *Proof and implementation of algorithmic realization of learning using privileged information (LUPI) paradigm: SVM+*. Technical Report NAS-TR-0187-2015. NSCR, CSE, PSU.
- [13] Z. Berkay Celik, David Lopez-Paz, and Patrick McDaniel. 2016. Patient-driven privacy control through generalized distillation. *IEEE Symposium on Privacy-Aware Computing (PAC)*.
- [14] Z. B. Celik, P. McDaniel, and T. Bowen. 2017. Malware modeling and experimentation through parameterized behavior. *The Journal of Defense Modeling and Simulation*.
- [15] Z. B. Celik, P. McDaniel, and R. Izmailov. 2017. Feature cultivation in privileged information-augmented detection. In *ACM CODASPY International Workshop on Security And Privacy Analytics*.
- [16] Z. B. Celik and S. Oktug. 2013. Detection of fast-flux networks using various DNS feature sets. In *ISCC*.
- [17] Z. B. Celik, J. Raghuram, G. Kesidis, and D. J. Miller. 2011. Salting public traces with attack traffic to test flow classifiers. In *Usenix CSET*.
- [18] Z. B. Celik, R. J. Walls, P. McDaniel, and A. Swami. 2015. Malware traffic detection using tamper resistant features. In *IEEE MILCOM*.
- [19] Microsoft Malware Classification Challenge. 2017. <https://www.kaggle.com/c/malware-classification/>. (2017). [Online; accessed 10-May-2017].
- [20] N. M. Duc and B. Q. Minh. 2009. Your face is not your password face authentication bypassing lenovo-asus-toshiba. *Black Hat Briefings*.
- [21] S. Garcia, M. Grill, J. Stiborek, and A. Zunino. 2014. An empirical comparison of botnet detection methods. In *Computers & Security*.
- [22] G. Gu et al. 2008. BotMiner: Clustering Analysis of Network Traffic for Protocol- and Structure-Independent Botnet Detection. In *USENIX Security*.
- [23] A. Hassanzadeh et al. 2013. PRIDE: Practical intrusion detection in resource constrained wireless mesh networks. In *Information and Communications Security*.
- [24] G. Hinton, O. Vinyals, and J. Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- [25] C. Hsu, C. Huang, et al. 2010. Fast-flux bot detection in real time. In *RAID*.
- [26] Gary Huang et al. 2012. Learning to align from scratch. In *Advances in Neural Information Processing Systems*.
- [27] G. B. Huang et al. 2014. *Labeled faces in the wild: Updates and new reporting procedures*. Technical Report UM-CS-2014-003. UMASS.
- [28] L. Huang et al. 2011. Adversarial machine learning. In *ACM security and artificial intelligence workshop*.
- [29] S. Huang et al. 2010. Fast-flux service network detection based on spatial snapshot mechanism for delay-free detection. In *ASLACCS*.
- [30] Vidit J. and E. Learned-Miller. 2010. *Fddb: A Benchmark for Face Detection in Unconstrained Settings*. Technical Report UM-CS-2010-009. UMass.
- [31] Rico Jonschkowski, Sebastian Höfer, and Oliver Brock. 2015. Patterns for learning with side information. *arXiv:arXiv:1511.06429*
- [32] P. Laskov et al. 2014. Practical evasion of a learning-based classifier: A case study. In *IEEE Security and Privacy*.
- [33] D. Lopez-Paz, L. Bottou, B. Schölkopf, and V. Vapnik. 2015. Unifying distillation and privileged information. *arXiv preprint arXiv:1511.03643*.
- [34] MATLAB documentation of quadprog function. 2018. www.mathworks.com/help/optim/ug/quadprog.html. [Online; accessed 15-March-2018].
- [35] P. McDaniel, N. Papernot, and Z. B. Celik. 2016. Machine learning in adversarial settings. *Security & Privacy Magazine*.
- [36] D. J. Miller et al. 2012. Sequential anomaly detection in a batch with growing number of tests: Application to network intrusion detection. In *MLSP*.
- [37] L. Nataraj, V. Yegneswaran, P. Porras, and J. Zhang. 2011. A comparative assessment of malware classification using binary texture analysis and dynamic analysis. In *Security and Artificial Intelligence Workshop*. ACM.
- [38] N. Nissim et al. 2014. Novel active learning methods for enhanced PC malware detection in windows OS. *Expert Systems with Applications*.
- [39] R. Pang, M. Allman, M. Bennett, J. Lee, V. Paxson, and B. Tierney. 2005. A first look at modern enterprise traffic. In *Internet Measurement Conference*.
- [40] N. Papernot et al. 2016. Distillation as a defense to adversarial perturbations against deep neural networks. *IEEE S&P*.
- [41] O. M. Parkhi et al. 2015. Deep face recognition. *British Machine Vision*.
- [42] E. Passerini, R. Paleari, L. Martignoni, and D. Bruschi. 2008. Fluxor: Detecting and monitoring fast-flux service networks. In *RAID*.
- [43] M. Z. Rafique and J. Caballero. 2013. Firma: Malware clustering and network signature generation with mixed network behaviors. In *RAID*.
- [44] B. Rahbarinia et al. 2015. Segugio: Efficient behavior-based tracking of malware-control domains in large ISP networks. In *IEEE DSN*.
- [45] S. Raza, L. Wallgren, and T. Voigt. 2013. SVELTE: Real-time intrusion detection in the Internet of Things. *Ad hoc networks*.
- [46] J. Reeves, A. Ramaswamy, M. Locasto, S. Bratus, and S. Smith. 2011. Lightweight intrusion detection for resource-constrained embedded control systems. In *Critical Infrastructure Protection*.
- [47] Bernardete Ribeiro et al. 2012. Enhanced default risk models with SVM+. *Expert Systems with Applications*.
- [48] C. Rossow et al. 2013. Sok: P2pwned-modeling and evaluating the resilience of peer-to-peer botnets. In *IEEE Security and Privacy*.
- [49] K. Scarfone and P. Mell. 2007. Guide to intrusion detection and prevention systems (IDPS). *NIST special publication*.
- [50] V. Sharmanska, N. Quadrianto, and C. H. Lampert. 2013. Learning to rank using privileged information. In *International Conference on Computer Vision*.
- [51] R. Shittu et al. 2015. Intrusion alert prioritisation and attack detection using post-correlation analysis. *Computers & Security*.
- [52] Y. Sun, D. Liang, X. Wang, and X. Tang. 2015. Deepid3: Face recognition with very DNNs. *arXiv preprint arXiv:1502.00873*.
- [53] B. A. Turlach and A. Weingessel. 2007. quadprog: Functions to solve quadratic programming problems. *R package version*.
- [54] V. Vapnik and R. Izmailov. 2015. Learning using privileged information: Similarity control and knowledge transfer. *Journal of Machine Learning Research*.
- [55] V. Vapnik and A. Vashist. 2009. A new learning paradigm: Learning using privileged information. *Neural Networks*.
- [56] R. J. Walls, E. G. Learned-Miller, and B. N. Levine. 2011. Forensic triage for mobile phones with DECODE. In *USENIX Security*.
- [57] Z. Wang and Q. Ji. 2015. Classifier learning with hidden information. In *IEEE Computer Vision and Pattern Recognition*.
- [58] L. Wolf, T. Hassner, and Y. Taigman. 2011. Effective unconstrained face recognition by combining multiple descriptors and learned background statistics. *Pattern Analysis and Machine Intelligence*.
- [59] S. Yadav et al. 2010. Detecting algorithmically generated malicious domain names. In *Internet Measurement Conference*.
- [60] G. Zou, G. Kesidis, and D. J. Miller. 2011. A flow classifier with tamper-resistant features and an evaluation of its portability to new domains. *JSAC*.

A MODEL INFLUENCE OPTIMIZATION

In this Appendix, we present the formulation of model influence approach introduced in Section 3.2 (Section A.1) and its implementation in MATLAB (Section A.2) in order to realize the paradigm in detection systems.

A.1 Model Influence Formulation

We can formally divide the feature space into two spaces at training time. Given L standard vectors $\vec{x}_1, \dots, \vec{x}_L$ and L privileged vectors $\vec{x}_1^*, \dots, \vec{x}_L^*$ with a target class $y = \{+1, -1\}$, where $\vec{x}_i \in \mathbb{R}^N$ and $\vec{x}_i^* \in \mathbb{R}^M$ for all $i = 1, \dots, L$. The kernels $K(\vec{x}_i, \vec{x}_j)$ and $K^*(\vec{x}_i^*, \vec{x}_j^*)$ are selected along with positive parameters κ , and γ . Our goal is finding a detection model $f \rightarrow \mathbf{x}^s : y$. The optimization problem is formulated as [54]:

$$\left\{ \begin{array}{l} \sum_{i=1}^L \alpha_i - \frac{1}{2} \sum_{i,j=1}^L y_i y_j \alpha_i \alpha_j K(\vec{x}_i, \vec{x}_j) \\ - \frac{\gamma}{2} \sum_{i,j=1}^L y_i y_j (\alpha_i - \delta_i)(\alpha_j - \delta_j) K^*(\vec{x}_i^*, \vec{x}_j^*) \rightarrow \max \\ \sum_{i=1}^L \alpha_i y_i = 0, \quad \sum_{i=1}^L \delta_i y_i = 0 \\ 0 \leq \alpha_i \leq \kappa C_i, \quad 0 \leq \delta_i \leq C_i, \quad i = 1, \dots, L \end{array} \right. \quad (1)$$

The detection rule f for vector \vec{z} is defined as:

$$f(\vec{z}) = \text{sign} \left(\sum_{i=1}^L y_i \alpha_i K(\vec{x}_i, \vec{z}) + B \right) \quad (2)$$

where to compute B , we first derive the Lagrangian of (1):

$$\begin{aligned} \mathcal{L}(\vec{\alpha}, \vec{\beta}, \vec{\phi}, \vec{\lambda}, \vec{\mu}, \vec{v}, \vec{\rho}) &= \sum_{i=1}^L \alpha_i - \frac{1}{2} \sum_{i,j=1}^L y_i y_j \alpha_i \alpha_j K(\vec{x}_i, \vec{x}_j) \\ &- \frac{\gamma}{2} \sum_{i,j=1}^L y_i y_j (\alpha_i - \delta_i)(\alpha_j - \delta_j) K^*(\vec{x}_i^*, \vec{x}_j^*) \\ &+ \phi_1 \sum_{i=1}^L \alpha_i y_i + \phi_2 \sum_{i=1}^L \delta_i y_i + \sum_{i=1}^L \lambda_i \alpha_i \\ &+ \sum_{i=1}^L \mu_i (\kappa C_i - \alpha_i) + \sum_{i=1}^L v_i \delta_i \\ &+ \sum_{i=1}^L \rho_i (C_i - \delta_i) \end{aligned} \quad (3)$$

with Karush-Kuhn-Tucker (KKT) conditions (for each $i = 1, \dots, L$), we rewrite

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \alpha_i} &= -K(\vec{x}_i, \vec{x}_i) \alpha_i - \gamma K^*(\vec{x}_i^*, \vec{x}_i^*) \alpha_i + K^*(\vec{x}_i^*, \vec{x}_i^*) \gamma \delta_i \\ &- \sum_{k \neq i} K(\vec{x}_i, \vec{x}_k) y_i y_k \alpha_k - \gamma \sum_{k \neq i} K^*(\vec{x}_i^*, \vec{x}_k^*) y_i y_k (\alpha_k - \delta_k) \\ &+ 1 + \phi_1 y_i + \lambda_i - \mu_i = 0 \\ \frac{\partial \mathcal{L}}{\partial \delta_i} &= -K^*(\vec{x}_i^*, \vec{x}_i^*) \gamma \delta_i + K^*(\vec{x}_i^*, \vec{x}_i^*) \gamma \alpha_i \\ &+ \sum_{k \neq i} K(\vec{x}_i, \vec{x}_k) y_i y_k \gamma (\alpha_k - \delta_k) \\ &+ \phi_2 y_i + v_i - \rho_i = 0 \end{aligned} \quad (4)$$

where

$$\begin{aligned} \lambda_i &\geq 0, \quad \mu_i \geq 0, \quad v_i \geq 0, \quad \rho_i \geq 0, \\ \lambda_i \alpha_i &= 0, \quad \mu_i (C_i - \alpha_i) = 0, \\ v_i \delta_i &= 0, \quad \rho_i (C_i - \delta_i) = 0, \\ \sum_{i=1}^L \alpha_i y_i &= 0, \quad \sum_{i=1}^L \delta_i y_i = 0 \end{aligned} \quad (5)$$

We denote for $i = 1, \dots, L$

$$\begin{aligned} F_i &= \sum_{k=1}^L K(\vec{x}_i, \vec{x}_k) y_k \alpha_k, \\ f_i &= \sum_{k=1}^L K^*(\vec{x}_i^*, \vec{x}_k^*) y_k (\alpha_k - \delta_k) \end{aligned} \quad (6)$$

and rewrite (4) in the form

$$\left\{ \begin{array}{l} \frac{\partial \mathcal{L}}{\partial \alpha_i} = -y_i F_i - \gamma y_i f_i + 1 + \phi_1 y_i + \lambda_i - \mu_i = 0 \\ \frac{\partial \mathcal{L}}{\partial \delta_i} = \gamma y_i f_i + \phi_2 y_i + v_i - \rho_i = 0 \\ \lambda_i \geq 0, \quad \mu_i \geq 0, \quad v_i \geq 0, \quad \rho_i \geq 0, \\ \lambda_i \alpha_i = 0, \quad \mu_i (C_i - \alpha_i) = 0, \quad v_i \delta_i = 0, \quad \rho_i (C_i - \delta_i) = 0 \\ \sum_{i=1}^L \alpha_i y_i = 0, \quad \sum_{i=1}^L \delta_i y_i = 0 \end{array} \right. \quad (7)$$

The first equation in (7) implies

$$\phi_1 = -y_j (1 - y_j F_j - \gamma y_j f_j + \lambda_j - \mu_j) \quad (8)$$

for all j . If j is selected such that $0 < \alpha_j < \kappa C_j$ and $0 < \delta_j < C_j$, then (7) implies $\lambda_j = \mu_j = v_j = \rho_j = 0$ and (8) has the following form

$$\begin{aligned} \phi_1 &= -y_j (1 - y_j F_j - \gamma y_j f_j) \\ \phi_1 &= -y_j \left(\left(1 - \sum_{i=1}^L y_i y_j K(\vec{x}_i, \vec{x}_j) \right) \alpha_i \right. \\ &\quad \left. - \gamma \sum_{i=1}^L y_i y_j K^*(\vec{x}_i^*, \vec{x}_j^*) (\alpha_i - \delta_i) \right) \end{aligned}$$

Therefore, B is computed as $B = -\phi_1$:

$$B = y_j \left(1 - \sum_{i=1}^L y_i y_j K(\vec{x}_i, \vec{x}_j) (\alpha_i) - \gamma \sum_{i=1}^L y_i y_j K^*(\vec{x}_i^*, \vec{x}_j^*) (\alpha_i - \delta_i) \right) \quad (9)$$

where j is such that $0 < \alpha_j < \kappa C_j$ and $0 < \delta_j < C_j$.

A.2 Model Influence Implementation

We present implementation of model influence by solving its quadratic programming problem using MATLAB quadprog function provided by the optimization toolbox. Other equivalent functions in R [53] or similar software can be easily adapted.

MATLAB function quadprog($H, \vec{f}, A, \vec{b}, Aeq, \vec{beq}, \vec{lb}, \vec{ub}$) solves the quadratic programming problem in the form as follows:

$$\begin{cases} \frac{1}{2} \vec{z}^T H \vec{z} + f^T \vec{z} \rightarrow \min \\ A \cdot \vec{z} \leq \vec{b} \\ Aeq \cdot \vec{z} = \vec{beq} \\ \vec{lb} \leq \vec{z} \leq \vec{ub} \end{cases} \quad (10)$$

Here, H, A, Aeq are matrices, and $\vec{f}, \vec{b}, \vec{beq}, \vec{lb}, \vec{ub}$ are vectors. \vec{z} is defined as $\vec{z} = (\alpha_1, \dots, \alpha_L, \delta_1, \dots, \delta_L) \in R^{2L}$. We now rewrite (1) in the form of (10).

$$\frac{1}{2} \vec{z}^T H \vec{z} + f^T \vec{z} \rightarrow \min$$

where

$$f = (-1_1, -1_2, \dots, -1_L, 0_{L+1}, \dots, 0_{2L})$$

and

$$H_{ij} = \begin{pmatrix} H^{11} & H^{12} \\ H^{12} & H^{22} \end{pmatrix}$$

where, for each pair $i, j = 1, \dots, L$,

$$H_{ij}^{11} = K(\vec{x}_i, \vec{x}_j) y_i y_j + \gamma K^*(\vec{x}_i^*, \vec{x}_j^*) y_i y_j,$$

$$H_{ij}^{12} = -\gamma K^*(\vec{x}_i^*, \vec{x}_j^*) y_i y_j,$$

$$H_{ij}^{22} = +\gamma K^*(\vec{x}_i^*, \vec{x}_j^*) y_i y_j$$

The second line of (10) is absent. The third line of (10) corresponds to the second line of (1) when written as

$$Aeq \cdot \vec{z} = \vec{beq}$$

where

$$Aeq = \begin{pmatrix} y_1 & y_2 & \dots & y_L & 0 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 & y_1 & y_2 & \dots & y_L \end{pmatrix},$$

$$\vec{beq} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

The fourth line of (10) corresponds to the third line of (1) when written as

$$\vec{lb} \leq \vec{z} \leq \vec{ub}$$

where

$$\vec{lb} = (0_1, 0_2, \dots, 0_L, 0_{L+1}, \dots, 0_{2L}),$$

$$\vec{ub} = (\kappa C_1, \kappa C_2, \dots, \kappa C_L, C_1, C_2, \dots, C_L)$$

After all variables ($H, \vec{f}, A, \vec{b}, Aeq, \vec{beq}, \vec{lb}, \vec{ub}$) are defined, optimization toolbox guide [34] can be used to select quadprog() function options such as an optimization algorithm and maximum number of iterations. Then, output of the function can be used in detection function f for a new sample \vec{z} to make predictions as follows:

$$f(\vec{z}) = \text{sign} \left(\sum_{i=1}^L y_i \alpha_i K(\vec{x}_i, \vec{z}) + B \right) \quad (11)$$

B DETAILS OF DETECTION SYSTEMS

In this Appendix, we detail the standard and privileged features of fast-flux bot and malware traffic detection systems introduced in Section 4.2. Table 1 presents feature categories and definitions of fast-flux bot detector obtained from recent works [25, 29, 42, 59], and Table 2 presents the features of malware traffic detector obtained from recent works [17, 18, 36, 60]. The interested reader can refer to the references for the motivation of the feature selection.

Category	Definition	Feature dependency	Feature type
DNS Answer	Number of unique A records Number of NS records	DNS packet analysis	standard set
Timing	Network delay (μ and σ) Processing delay (μ and σ) Document fetch delay (μ and σ)	HTTP requests	standard set
Domain name	Edit distance Kullback-Leibler divergence (unigrams and bigrams) Jaccard similarity (unigrams and bigrams)	Whitelist of benign domain names	privileged set
Spatial	Time zone entropy of A records Time zone entropy of NS records Minimal service distances (μ and σ)	IP coordinate database lookup (external source)	privileged set
Network	Number of distinct autonomous systems Number of distinct networks	WHOIS processing (external source)	privileged set

Table 1: Fast-flux bot detection system standard and privileged feature descriptions (μ is mean and σ is std. dev.).

Abbreviation	Definition	Properties	Feature type
cnt-data-pkt	The count of all the packets with at least a byte of TCP data payload	-TCP length is observed -Client to server	standard set
min-data-size	The minimum payload size observed	-TCP length observed -Client to server -0 if there are no packets	standard set
avg-data-size	Data bytes divided by the total number of packets	-TCP length observed -Packets with payload observed -Server to client -0 if there are no packets	standard set
init-win-bytes	The total number of bytes sent in initial window	-Retransmitted packets not counted -Client to server & server to client -0 if no ACK observed -Frame length calculated	standard set
RTT-samples	The total number of RTT samples found	-Client to server	standard set
IP-bytes-median	Median of total IP packets	-IP length calculated -Client to server	standard set
frame-bytes-var	Variance of bytes in Ethernet packets	-Frame length calculated -Client to server	standard set
IP-ratio	Ratio between the maximum packet size and minimum packet size	-IP length calculated -Client to server & server to client -1 If a packet observed, and if no packets are observed 0 is reported	standard set
pushed-data-pkts	The count of all the packets seen with the PUSH set in TCP header	-Client to server & server to client	standard set
goodput	Total number of frame bytes divided by the differences between last packet time and first packet time	-Frame length calculated -Client to server -Retransmitted bytes not counted	standard set
duration	Total connection time	Time difference between the last packet and first packet (SYN flag is seen from destination)	privileged set
min-IAT	Minimum packet inter-arrival time for all packets of the flow	-Client to server & server to client	privileged set
urgent-data-pkts	The total number of packets with the URG bit turned on in the TCP header	-Client to server & server to client	privileged set
src-port	Source port number	-Undecoded	privileged set
dst-port	Destination port number	-Undecoded	privileged set
payload-info	Byte frequency distributions	-If not HTTPS at training time -If payloads are available -Client to server & server to client	privileged set

Table 2: Malware traffic detection standard and privileged feature descriptions.