

FOAP: Fine-Grained Open-World Android App Fingerprinting

Jianfeng Li¹, Hao Zhou¹, Shuohan Wu¹, Xiapu Luo^{1*}, Ting Wang², Xian Zhan¹, Xiaobo Ma³

¹*Department of Computing, The Hong Kong Polytechnic University*

²*College of Information Sciences and Technology, Pennsylvania State University*

³*School of Computer Science and Technology, Xi'an Jiaotong University*

Abstract

Despite the widespread adoption of encrypted communication for mobile apps, adversaries can still identify apps or infer selected user activities of interest from encrypted mobile traffic via app fingerprinting (AF) attacks. However, most existing AF techniques only work under the closed-world assumption, thereby suffering potential precision decline when faced with apps unseen during model training. Moreover, serious privacy leakage often occurs when users conduct some sensitive operations, which are closely associated with specific UI components. Unfortunately, existing AF techniques are too coarse-grained to acquire such fine-grained sensitive information. In this paper, we take the first step to identify method-level fine-grained user action of Android apps in the open-world setting and present a systematic solution, dubbed FOAP, to address the above limitations. First, to effectively reduce false positive risks in the open-world setting, we propose a novel metric, named structural similarity, to adaptively filter out traffic segments irrelevant to the app of interest. Second, FOAP achieves fine-grained user action identification via synthesizing traffic and binary analysis. Specifically, FOAP identifies user actions on specific UI components through inferring entry point methods correlated with them. Extensive evaluations and case studies demonstrate that FOAP is not only reasonably accurate but also practical in fine-grained user activity inference and user privacy analysis.

1 Introduction

Mobile devices in markets have bloomed unprecedentedly in the last decade [37]. The ubiquitous mobile devices, such as smartphones and tablets, have become necessities in modern life. Their prosperity and success are largely attributed to diverse apps running on them [35, 36]. These apps substantially extend the capability of off-the-shelf mobile devices and increasingly improve modern life in different aspects.

While offering convenience, mobile apps also raise privacy concerns. Massive amounts of private user data are transmitted to and stored in cloud servers. Once these servers are compromised, catastrophic privacy leakage will happen subsequently [6]. The adversary can also indirectly infer sensitive private information, such as diseases, religious preferences, and individual location, through side-channel attacks [26, 34].

Despite the widespread adoption of encrypted communication, mobile apps are still susceptible to app fingerprinting (AF) attacks [2, 3, 14, 39–41, 43], which are essentially side-channel attacks. In such attacks, adversaries recognize apps or selected user activities of interest using pre-trained machine learning models without inspecting the packet payload plaintext, thereby immune to traffic encryption. However, existing AF techniques face two major limitations.

Closed-World Assumption. The vast majority of existing AF techniques (e.g., [15, 22, 29]) formulate app identification as a multi-class classification problem. It implies that they cannot correctly identify apps that are unseen during model training, because these apps will be erroneously classified into known classes. A straightforward solution is training a one-vs.-rest binary classifier for each app. Nevertheless, such a plausible solution only works under the closed-world assumption, i.e., apps as the negative class in testing stage should be presented during model training. If not, false positive cases may tremendously increase. Unfortunately, involving all rest apps in the training dataset is impossible as Android and iOS both host more than 3.4 million third-party apps [35, 36]. Even worse, recent studies [43, 51] pointed out that apps extensively use third-party libraries, leading to similar network behaviors, which may increase the risks of false positives.

Identification Granularity. Serious privacy leakage often occurs when users conduct some sensitive operations on certain UI components. For instance, in some COVID-19 contact tracing apps, the click action on the button to report positive testing can be a strong indicator of COVID-19 infection. However, information obtained through existing AF techniques is too coarse-grained to infer the above user privacy because these techniques focus on either app identification [39, 43, 49]

*The corresponding author.

or inferring selected user activities of interest [15, 22, 29], and none of them can identify user actions on specific UI components. Additionally, previous works need labor-intensive efforts to manually label selected user activities, limiting their scalability in the wild.

In this paper, we present FOAP, a novel Fine-grained Open-world App fingerPrinting technique, to address the above limitations. First, FOAP carries out *open-world* app recognition, obviating the need for the closed-world assumption. The core challenge is the risks of false positives in the open-world setting. We solve it by exploiting the fact that two different apps seldom exhibit completely identical network behaviors. Specifically, we propose a novel metric named *structural similarity* to characterize how network flows within a traffic segment behave similarly to those generated by the app of interest. By leveraging this metric, FOAP effectively reduces false positives by adaptively filtering out traffic segments that are irrelevant to the app of interest.

Second, FOAP aims at *method-level* fine-grained user action identification via synthesizing traffic and binary analysis. We focus on identifying user actions on specific UI components, e.g., clicking a certain button. To this end, we characterize user actions using the corresponding entry point (EP) methods, including callbacks of UI components and lifecycle methods of Android components, and transform user action identification to EP method identification. In the training stage, the major challenge is how to automatically label network flows with EP methods in the face of various ways to trigger network flows. We solve it through extensive app analysis and implement the automatic labeling based on Android framework instrumentation. Automatic labeling brings two salient advantages: i) it endows FOAP with the capability of exploring all possible user actions that may generate network flows, and thus FOAP can identify not only selected user activities but also other user actions that trigger network traffic; ii) it obviates the need of manual labeling, facilitating the scalability of FOAP. In the identification stage, the major challenge is the multi-label flow. A network flow is often correlated with multiple EP methods, leading to a multi-label classification problem. We relax this problem to a series of multi-class classification sub-problems by extracting in-flow bursts. To identify EP methods associated with a network flow, we model the spatial-temporal contextual dependency of in-flow bursts within it and infer the EP method associated with each of them based on the conditional random field.

In summary, this paper makes the following contributions:

- To the best of our knowledge, FOAP is the first approach for conducting method-level fine-grained app user action recognition in the open-world setting. We release its source code and the datasets at <https://github.com/jflixtu/FOAP>.
- We address several challenging issues in the design of FOAP. First, we design structural similarity, a novel metric




Figure 1: The threat model.

for effectively reducing the false positive risks in open-world app traffic recognition. Second, we propose a new approach to automatically label network flows with EP methods and build a spatial-temporal context model to accurately identify EP methods from network flows.

- We implement a prototype of FOAP and conduct extensive experiments to evaluate it. The results show that FOAP outperforms baseline approaches in open-world app recognition, improving the F1-score from 0.679 to 0.911. It is also reasonably accurate in EP method identification with an average F1-score of 0.885. We present three case studies to demonstrate FOAP’s practicality in fine-grained user activity inference and user privacy analysis.

2 Overview

2.1 Threat Model

As shown in Figure 1, the adversary considered in this paper sniffs network traffic on a wireless access point. His goal is to recognize network flows generated by the app of interest and identify what EP methods of this app trigger them to infer fine-grained user actions. We define a network flow as a sequence of packets corresponding to a socket-to-socket communication identified by a unique combination of source and destination addresses and port numbers, together with transport protocols. In this paper, we focus on TCP flows, because i) TCP packets are dominating in Android app traffic (98.6% in our dataset), and ii) we mainly consider encrypted network traffic based on TCP protocol (e.g., HTTPS, etc.).

We assume that the adversary can trace back all network flows to different devices according to source addresses. We also assume the adversary *cannot* exploit i) packet plaintext payload and ii) destination feature of network flows. Our assumptions match the real-world scenario of AF attacks. First, since app network flows are commonly encrypted [43], the adversary cannot access the plaintext payload in most cases. Second, due to the wide use of encrypted proxy agents (e.g., ShadowSocks), the real destinations of network flows are often invisible. Even if the adversaries can obtain the real destinations, such features are not reliable since they may change across networks due to the widespread adoption of CDN. Moreover, the share of third-party libraries among apps leads to the share of destinations across apps.

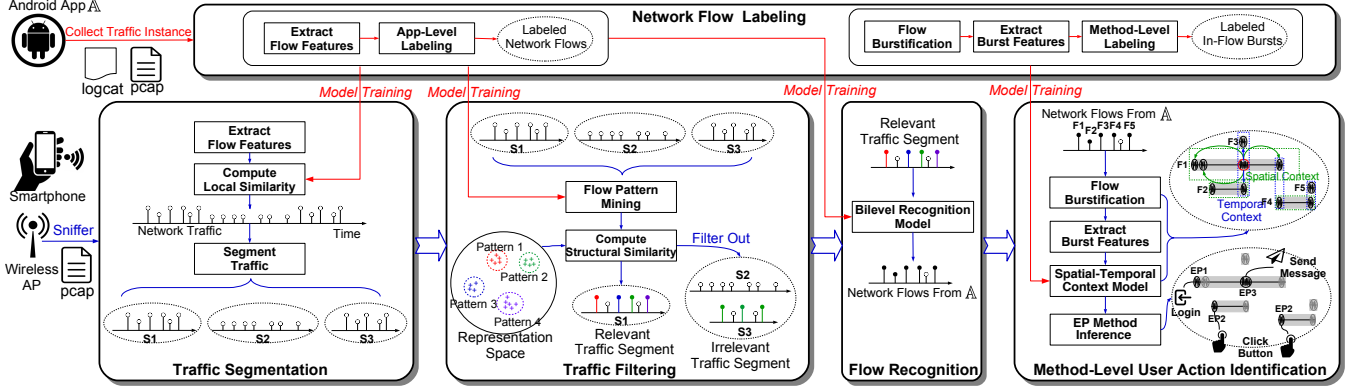


Figure 2: The workflow of FOAP. Training stage (resp. identification stage) is colored in red (resp. blue).

2.2 Workflow of FOAP

Without loss of generality, we assume A is one of the apps of interest. FOAP identifies method-level user actions of A from encrypted network traffic in the open-world setting. Figure 2 illustrates the workflow of FOAP.

Training Stage:

- **Network Flow Labeling** (§ 3) plays a core role in constructing training datasets. By either automatically or manually running A , we collect A 's traffic instances, each of which contains a log file (in Android, we call it logcat) and a traffic file in pcap format. We first label whether a network flow is from A , i.e., app-level labeling. If yes, we next label when and what EP methods of A trigger this network flow, i.e., method-level labeling. We extract a feature vector for every network flow after app-level labeling to construct the training dataset for i) *Traffic Segmentation*, ii) *Traffic Filtering*, and iii) *Flow Recognition*. Since a network flow usually contains multiple in-flow bursts triggered by different EP methods, we divide a network flow into a series of in-flow bursts and extract a feature vector for each one to construct the training dataset for *Method-Level User Action Identification*.

Identification Stage:

- **Traffic Segmentation** (§ 5.1) divides network traffic into different time periods, such that there are some time periods during which A is active (if they exist). To this end, we extract a feature vector for each network flow and compute its *local similarity* with A to characterize how likely it is generated by A , from the perspective of a single flow. Based on the local similarity, we segment network traffic to locate all possible time periods when A is probably active. In Figure 2, we represent each network flow with a bar. The time-axis location (resp. height) of a bar represents the start time (resp. local similarity) of the corresponding flow. Network traffic is divided into a series of traffic segments, i.e., s_1 , s_2 , and s_3 .
- **Traffic Filtering** (§ 5.2) is designed for mitigating false positives in the open-world setting. By leveraging flow pattern mining, we quantify how likely a network flow matches a certain pattern of A 's network flows in the feature representation

space and then use this pattern-related information to profile structural characteristics of a traffic segment. For each traffic segment, we compute its *structural similarity* with A by synthesizing pattern-related information from all network flows within it. Structural similarity is informative to reflect how likely A is active during the time period corresponding to a traffic segment. We take advantage of such a metric to identify traffic segments relevant to A (labeled by A) and filter out other irrelevant traffic segments (labeled by non- A) without further analysis. In Figure 2, s_1 has a high structural similarity with A because network flows within it match various patterns of A , whereas s_2 has a very low structural similarity with A because network flows within it match no pattern of A . Even though some network flows within s_3 have high local similarities with A , s_3 has a relatively low structural similarity with A because these network flows only match one pattern of A (colored in green), implying s_3 is structurally different from network traffic generated by A . Finally, s_1 is identified as a relevant traffic segment of A , whereas s_2 and s_3 are identified as irrelevant traffic segments of A . An irrelevant traffic segment of A may contain network traffic generated by more than one apps except A , while relevant traffic segments of different apps may overlap.

- **Flow Recognition** (§ 5.3) constructs a bilevel recognition model to identify whether a network flow in the relevant traffic segment is generated by A . This model considers not only the feature vector of a network flow but also contextual information from surrounding network flows in favor of better recognition accuracy.

- **Method-Level User Action Identification** (§ 6) infers which entry point (EP) methods trigger a network flow of A to characterize fine-grained user actions. Network flows corresponding to persistent connections often contain multiple packet bursts, i.e., in-flow bursts, each of which may be triggered by different EP methods. Network flows corresponding to short connections often contain only one in-flow burst that is triggered by an EP method. To identify EP methods, we first extract in-flow bursts and their feature vectors. Next, we construct a spatial-temporal context model to characterize the

correlation between in-flow bursts. Based on this model, we identify EP methods to infer fine-grained user actions, such as clicking on a button.

3 Network Flow Labeling

To construct the training dataset for open-world app traffic recognition and method-level user action identification, we capture network flows generated by the tested apps and label each network flow with the app and EP methods (i.e., the component lifecycle or the UI callback) that generate it. To this end, we instrument Android framework to i) extract socket information to correlate network flows, ii) retrieve process identifier (PID) to obtain the package name and collect stack trace to determine the EP method. The collected data is written to `logcat`.

3.1 Resolving Network Socket

Since Android apps commonly use network sockets to generate network flows, we resolve each network socket operation performed by the tested apps to correlate the captured network flows. According to the observation that framework APIs for performing socket operations (e.g., `connect` defined in `Socket`) rely on socket related native system functions (e.g., `connect` exported by `libc.so`) to complete tasks, we instrument socket related system functions to get a unique 4-tuple for each TCP socket. Moreover, since identifiers (i.e., PID and TID) of the process and thread, executing the socket related system functions, are essential for obtaining the package name and correlating the EP method, we make the instrumented socket related functions invoke `getpid` and `gettid` to retrieve the PID and TID values.

It is worth noting that, the state-of-the-art work [20] cannot correlate the socket operations implemented in apps' native code to their corresponding EP methods because it just analyzes socket related framework APIs. Since our approach analyzes the socket related system functions, which are internally called by socket related framework APIs, we can associate socket operations implemented in both Java code and native code with their EP methods.

3.2 Analyzing Network Socket Operation

In order to correlate each network flow to an app as well as EP methods of this app, we further analyze each network operation performed by the tested apps.

Correlating Network Flow to App: Relying on PIDs retrieved when the socket related system functions are invoked, we get the unique identifier of the app, performing the socket operations, by accessing the `/proc/PID/cmdline` file, which keeps app's package name.

Correlating Network Flow to EP Method: Since socket operations must be carried out in non-UI threads of apps [1],

apps can create a new thread by internally calling the `clone` function in `libc.so` to perform network operations (case-1). Therefore, we instrument `clone` to get the stack trace and TID of each created thread, so that we can correlate the thread with EP method, which is presented in the stack trace and leads to the creation of the thread. If the TID equals the one retrieved when socket related functions are invoked, we correlate the EP method with generated network flows.

Instead of creating a new thread, apps can reuse a thread to perform different network operations. For example, apps can send messages to notify the handler running in a non-UI thread (case-2) or submit tasks to an idle thread managed by the thread pool (case-3) to conduct socket operations. It is non-trivial to handle these cases because we need to determine the EP method for each network operation conducted in the reused thread. In the following, we detail how we label network flows in these cases.

For case-2, we use the `Message` objects, which notify the message handler to perform socket operations, to distinguish every network operation conducted by the handler. Specifically, we instrument `enqueueMessage` defined in the `Handler` class to get the stack trace and the message (i.e., the `Message` object) sent to the handler so that we can correlate the message with the EP method that results in the sending of the message. Meanwhile, to correlate each message with the thread that runs the handler, we instrument `dispatchMessage` of the `Handler` class to get the received message and thread's TID. If the TID equals the one retrieved when socket related functions are invoked, we correlate the received message with generated network flows. Since the received and sent messages refer to the same `Message` object, we then correlate the EP method to the network flows.

For case-3, we use the `Runnable` objects, which are executed in the thread managed by the thread pool, to differentiate every network operation performed by the same thread in the thread pool. Specifically, we instrument `execute` defined in the `ThreadPoolExecutor` class to get the stack trace and the task (i.e., the `Runnable` object) submitted to the thread pool, so that we can correlate the task with the EP method that causes the submission of the task. Meanwhile, to correlate each task with the thread that runs the task, we instrument `runworker` of the `ThreadPoolExecutor` class to get the executed task and the thread's TID. If the TID equals the one retrieved when socket related functions are invoked, we correlate the executed task with the network flows. Since the submitted and the executed tasks refer to the same `Runnable` object, we then correlate the EP method to the network flows.

Note that, for those UI callbacks which are identified as EP methods, we associate them with their corresponding resource IDs. For example, we correlate each `onClick` callback with the resource ID of the relevant UI component that registered the click event listener.

4 Traffic Feature Extraction

Although some studies [5, 32, 33] take advantage of deep learning models to automatically extract traffic features, deep learning models require large amounts of training data to avoid model overfitting. In our problem, the number of network flows/in-flow bursts triggered by different EP methods is unbalanced. Some apps/EP methods only have few samples. Therefore, deep-learning-based feature extraction is not suitable for our problem. Either a network flow or an in-flow burst is essentially a packet sequence. We extract 123-dimensional features for both of them from five perspectives.

- **General characteristic** (8 features). We extract inbound/outbound packet number, bidirectional packet number, inbound/outbound packet percentage, inbound/outbound bytes, and the duration of packet sequence.
- **Interactive pattern** (20 features). We characterize the interactive pattern between endpoints using a function $f(x) = (i-x) \hat{\Delta}_{j=1}^{i-1} d_j + (x-i+1) \hat{\Delta}_{j=1}^i d_j$ for $i-1 < x \leq i$, where $d_j = 1$ (resp. $d_j = -1$) if the j th packet is an inbound (resp. outbound) packet. Assume the packet sequence contains n packets. The interval $(0;n)$ is divided into 20 bins with the size of $n=20$. We pick $f(x)$ at the center of each bin as a feature and totally get 20 features to approximate $f(x)$ (or, equivalently, the interactive pattern).
- **Packet rate characteristic** (5 features). The duration of a packet sequence is divided into a series of time windows with the size of 1 second. Mean, maximum, minimum, median, and standard deviation of packet number in a time window (i.e., packet rate), are computed as features.
- **Temporal characteristic** (39 features). We consider bidirectional, inbound, and outbound packet arrivals respectively. For each one, mean, maximum, minimum, and standard deviation of packet interval time, and 9 percentiles (from 10% to 90%) of relative arrival time, i.e., the time lag relative to the first packet, are extracted as features.
- **Packet size characteristic** (51 features). We extract various packet size statistics for inbound, outbound, and bidirectional packets respectively, including mean, maximum, minimum, median absolute deviation, standard deviation, variance, skew, kurtosis, and 9 percentiles (from 10% to 90%).

5 Open-World App Traffic Recognition

Open-world app traffic recognition aims to recognize app-specific traffic. More specifically, given a network flow, FOAP should identify whether it is from the app of interest, say A .

Open world vs. closed world. We define open-world setting along with closed-world setting in this paper. Let $S_T = \{\text{app}_T^i\}_{i=1}^{m_T}$ (resp. $S_I = \{\text{app}_I^i\}_{i=1}^{m_I}$) be the set comprised of all apps in the training stage (resp. the identification stage). In the closed-world setting, we have $S_I \subseteq S_T$. In the open-world setting, S_I might not be a subset of S_T but it could be.

In the open-world setting, a key challenge is how to effectively reduce false positives in face of apps that are unseen during model training. FOAP tackles it in two steps, i.e., traffic segmentation and traffic filtering. After that, FOAP constructs a bilevel recognition model to identify whether a network flow is generated by A by taking advantage of not only its feature vector but also its contextual information.

5.1 Traffic Segmentation

The goal of traffic segmentation is dividing network traffic into different time periods, such that there are some time periods during which A is active (if they exist). To this end, we first propose a metric, named by *local similarity*, to quantify how likely a network flow is generated by A from the perspective of a single flow.

Definition 1. *Local similarity is a metric to quantify how a network flow f is similar to network flows from A in terms of the intrinsic features from f itself. We compute f 's local similarity with A by $S_l(f; A) = \Pr\{f \in A | \mathbf{F}(f)\}$, where $\mathbf{F}(\cdot)$ is a function that returns the feature vector of f .*

We denote by $\mathbf{f} = (f_1; f_2; \dots; f_M)$ the sniffed network traffic consisting of M network flows, where the i th flow is denoted by f_i and its start time is t_i . Let q_i be f_i 's local similarity with A , i.e., $S_l(f_i; A)$. To compute q_i , we first extract f_i 's feature vector (see § 4) and feed it to a pre-trained random forest (RF) classifier. q_i is estimated as the percentage of decision trees that predict f_i is from A . Next, we segment network traffic based on $q_1; q_2; \dots; q_M$. Based on the measurement over our dataset, network flows generated by an app, say A , have significantly higher local similarity with this app, compared with network flows generated by other apps, thereby resulting in obvious difference of local similarity between time periods dominated by A and time periods dominated by other apps. Such a difference is about 4.51 times the local similarity variance for different network flows from an app. Inspired by this observation, we expect network flows within the same traffic segment have similar local similarity with A (i.e., the small variance within a segment), whereas network flows belonging to adjacent traffic segments have significantly different local similarities with A (i.e., large variance between adjacent segments). By doing so, we are able to extract all possible traffic segments when A is active. Formally, we formulate traffic segmentation as a combinational optimization problem.

$$\begin{aligned}
 & \min_{\mathbf{z}} \hat{\mathbf{a}} \max_{c=0}^{\max(\mathbf{z})} |Q_c| \cdot \text{Var}(Q_c) + \mathbf{l} \max(\mathbf{z}); \\
 \text{s.t. } & z_i \in \mathbb{N}; z_1 = 0; 0 \leq z_{i+1} - z_i \leq 1 \text{ for } 1 \leq i < M; \\
 & Q_c = \{q_i | z_i = c\}; |Q_c| \geq n_{\min}; \\
 & \max\{t_i | z_i = c\} - \min\{t_i | z_i = c\} \geq \tau_{\min};
 \end{aligned} \tag{1}$$

where z_i is a variable indicating which segment the i th network flow belongs to, $\mathbf{l} \max(\mathbf{z})$ is a regularization term that

penalizes the complexity of results. s_{\min} is the minimum number of network flows that a segment should contain, and t_{\min} is the minimum time span of a segment. The above optimization problem can be solved based on integer programming [19]. To speed up the traffic segmentation, we propose a divisive-agglomerative tree method to solve it in a greedy fashion. Divisive tree recursively generates small segment candidates while agglomerative tree merges some of them to minimize the loss in (1). We elaborate this method in [21].

5.2 Traffic Filtering

If A is active during the time period of a traffic segment, we refer to it as a relevant traffic segment and otherwise an irrelevant traffic segment. Traffic filtering aims to filter out all irrelevant traffic segments.

5.2.1 Flow Pattern Mining

In this step, we quantify how likely a network flow matches a certain pattern of network flows generated by A in the feature representation space. This pattern-related information will be used for profiling structural characteristics of a traffic segment. To this end, we need to identify different patterns by clustering A's network flows over the training dataset. Unfortunately, the original feature space is irregular because i) both continuous and discrete features are considered and ii) the range and variance of different features are tremendously diverse. It is difficult to find a proper distance metric for reasonable clustering analysis in the original feature space. To tackle this problem, we learn a mapping function based on metric learning [45] to map original feature vectors to a low-dimensional representation space, where we expect the euclidean distance between network flows in the same pattern, i.e., the same cluster, will be minimized whereas that between network flows in different patterns, i.e., different clusters, will be maximized. We refer the interested readers to [19] for the learning of mapping function.

Clustering Analysis. We map A's network flows in the training dataset to the representation space and conduct clustering analysis to identify flow patterns as different clusters. Hierarchical clustering is employed for this task as it obviates the need to specify the cluster number, which is unknown a priori. We specify the clustering threshold δ and obtain clusters $P_1; P_2; \dots; P_m$. We refer to these clusters as patterns of A's network flows. Next, we train a multi-class k-nearest neighbors classifier in the representation space to identify which pattern a new network flow belongs to. Let P_i be the identified pattern that f_i belongs to. We reuse f_i 's local similarity q_i to quantify the likelihood of f_i matching pattern P_i .

5.2.2 Filtering Out Irrelevant Traffic Segment

To filter out irrelevant traffic segments, we compute traffic segments' structural similarity with A.

Definition 2. Structural similarity characterizes how network flows within a traffic segment s_k are similar to A's network flows in various patterns. It is computed by

$$S_s(s_k; A) = \frac{\int_{q_{\min}}^{R_1} |H(s_k; q)| dq}{|H(s_k; 0)| (1 - q_{\min})}; \quad (2)$$

where $H(s_k; q) = \{f_i | q_i > q; f_i \in s_k\}$ is a function that returns a set containing all patterns that network flows within s_k belong to with the likelihood greater than q .

We consider network flows match certain patterns only if their likelihood are greater than q_{\min} . In this paper, we set $q_{\min} = 0.5$ by default. The value of $S_s(s_k; A)$ ranges from 0 to 1. A smaller value of $S_s(s_k; A)$ indicates s_k is more likely an irrelevant traffic segment. Formally, s_k will be deemed to be an irrelevant traffic segment and filtered out if $S_s(s_k; A) < S_{\min}^A$. In here, $S_{\min}^A \in [S_{\min}, S_{\max}]$ is an app-specific similarity threshold bounded by S_{\min} and S_{\max} . We specify it to be the value such that 95% network flows from A over the training dataset will not be erroneously filtered out. Despite a potential slight decline of recall for traffic recognition, such a similarity threshold can effectively reduce false positives by filtering out irrelevant traffic segments as many as possible.

5.3 Flow Recognition

Through traffic filtering, we obtain relevant traffic segments. However, not all network flows within them must be from A because they may contain network flows from background system services and other apps. To recognize whether a network flow is from A, we construct a bilevel recognition model. The low-level model is comprised of two binary random forest classifiers. The first classifier is used for computing the local similarity and has been trained before. The second classifier predicts the probability that a network flow is a background flow, (i.e., network flows generated by background system services). For a network flow f_i , the first classifier outputs its local similarity with A, denoted by q_i , while the second classifier outputs the probability that it is a background flow, denoted by r_i . Both q_i and r_i play important roles in the high-level model. The high-level model is a logistic regression classifier. We choose this model because of its excellent generalization capability due to a small number of parameters. For a network flow f_i , the high-level model takes advantage of features not only from itself but also from its contextual network flows. Specifically, we consider network flows within N_i , a time window centered at the start time of f_i with the size of T , as f_i 's contextual network flows. We construct f_i 's feature vector for high-level model as $\mathbf{x}_i = (q_i; \bar{q}_{N_i}; q_i - r_i)$, where $\bar{q}_{N_i} = \frac{1}{|N_i|} \sum_{f_j \in N_i} q_j$. For any network flow in relevant traffic segments, we recognize whether it is from A according to the output of the high-level model. To avoid mutual interference, we decouple the training of low-level and high-level models by leveraging bootstrap resampling method [12, 44].

6 Method-Level User Action Identification

After recognizing network flows from \mathcal{A} , we further infer what EP methods trigger them. Such information plays a critical role in characterizing fine-grained user actions.

Different from the binary classification in recognizing specific network flows, EP method identification can be formulated as a multi-label classification because a network flow may correspond to multiple EP methods. Network flows in app traffic can be roughly grouped into two categories. The first category corresponds to persistent connection network flows in this category often contain multiple packet bursts associated with sequentially invoked EP methods. The second category corresponds to short connection network flows in this category often contain only one packet burst associated with a single EP method. Instead of regarding a network flow as a whole to extract its feature vector and identify what EP methods trigger it, we identify EP methods associated with all packet bursts within it. By doing so, our method has three-fold advantages. First, we relax the multi-label classification to a series of multi-class classification subproblems. Second, we can make use of contextual dependency between bursts to improve identification accuracy. Third, we obtain a higher temporal resolution in characterizing user actions.

6.1 Flow Burstification

The first step is extracting in-flow bursts from network flows.

Definition 3. Given a burst threshold δ , an in-flow burst is a subsequence of packets within a network flow. Time interval between two packets in an in-flow burst is less than δ while that between two in-flow bursts is not less than δ .

Figure 3(a) illustrates an example, where we extract four in-flow bursts from a network flow. For each in-flow burst, we extract a 123-dimensional feature vector (see § 4).

Burst Labeling. To construct the training dataset for EP method identification, we need to label in-flow bursts. Assume that the network flow labeling module has obtained three EP methods, i.e. m_a , m_b , and m_c , and their invocation times t_1 , t_2 , and t_3 as shown in Figure 3(a). Before labeling in-flow burst, we first label each packet to indicate which method triggers it. Specifically, a packet is labeled by the first EP method invoked before it. We then label an in-flow burst by aggregating the label of packets belonging to this burst (majority voting). Finally, we obtain four in-flow bursts $(m_a; b_1)$, $(m_a; b_2)$, $(m_b; b_3)$, and $(m_c; b_4)$.

6.2 Spatial-Temporal Context Model

We model contextual dependency of in-flow bursts based on conditional random field [38] due to its flexibility to capture both spatial and temporal context.

Spatial Context: It reflects the contextual relationship between bursts within the same network flow. In Figure 3(b),

(a) Label within-flow bursts. (b) Spatial-temporal context.

Figure 3: Flow burstification and spatial-temporal context.

bursts b_1, b_2, b_3 , and b_4 are within the same network flow and they are mutually spatially contextual bursts.

Temporal Context: An EP method may trigger multiple network flows simultaneously, yielding concurrent in-flow bursts. To capture the temporal context, we specify an interval threshold δ_{burst} . Two in-flow bursts are temporally contextual if they belong to different network flows but the time interval between them is less than δ_{burst} . In Figure 3(b), bursts b_3, b_5 , and b_6 are mutually temporally contextual bursts.

Formally, let $M_A = \{m_k, g_{k=1}^K\}$ be a set comprised of EP methods of A that may trigger network flows. We denote by $b = (b_1; b_2; \dots; b_N)$ a sequence of in-flow bursts extracted from A 's network flows and $y = (y_1; y_2; \dots; y_N) \in M_A^N$ the underlying EP method invocations that trigger these in-flow bursts. The posterior distribution p can be expressed by

$$p(y|b) = \frac{1}{Z(b)} \prod_{i=1}^N W_i(y_i; b_i) \prod_{i_1, i_2 \in C^S} Y_{ij}(y_i; y_j; b) \prod_{i_1, i_2 \in C^T} F_{ij}(y_i; y_j; b); \quad (3)$$

where $W_i(y_i; b_i)$ is a unary potential characterizing the relationship between y_i and b_i , $Y_{ij}(y_i; y_j; b)$ (resp. $F_{ij}(y_i; y_j; b)$) is a spatially (resp. temporally) pairwise potential characterizing the relationship between y_i and its spatially (resp. temporally) contextual burst b_j , C^S (resp. C^T) is a set comprised of all spatially (resp. temporally) contextual bursts of b_i , and $Z(b)$ is a normalizing constant. We refer the interested readers to [21] for the construction and parameter learning of $W_i(y_i; b_i)$, $Y_{ij}(y_i; y_j; b)$, and $F_{ij}(y_i; y_j; b)$.

6.3 Iterative Inference of EP Method

Given the CRF-based spatial-temporal context model, the problem of EP method identification is transformed into inferring the hidden variables (i.e., EP method invocations) from the observable variables (i.e., in-flow bursts). For general graphs, like our spatial-temporal context model, the exact inference of CRFs is intractable because it requires exponential time in the worst case [38]. Therefore, we propose a greedy EP method inference algorithm to conduct approximate inference. Instead of estimating the posterior marginal distribution, this algorithm derives the most likely EP method sequence, i.e. $\hat{y} = \arg \max_y p(y|b)$ in an iterative fashion. Let $y^{(k)} = (y_1^{(k)}; y_2^{(k)}; \dots; y_N^{(k)})$ be the estimation of y in the k th round. We derive $y_1^{(k)}$ as the EP method that maximizes the

probability $p(y_i|b; y^{(k-1)} \Rightarrow y_i^{(k-1)})$. We record $y^{(k)}$ in the set Y and $p^{(k)} = Z(b)p(y^{(k)}|b)$ will also be stored. The iteration process repeats until it exceeds the maximum iteration number or it meets the early stopping condition, $y_i^{(k)} \in Y$. The resulting inference is derived as $\hat{y} = \arg \max_{y^{(k)} \in Y} p^{(k)}$. We refer the interested readers [21] for details about this algorithm. Given EP methods associated with all in-flow bursts, we infer what EP methods trigger a network flow by aggregating EP methods associated with in-flow bursts within it.

7 Evaluation

We first evaluate FOAP in open-world app traffic recognition and then evaluate it in EP method identification. We next analyze various factors that may influence the performance of FOAP, such as experimental data size and third-party libraries. To investigate how FOAP trained over the dataset generated by automatic test tools performs for network traffic generated by human users, we conduct cross-dataset evaluation. Finally, we analyze the running time of FOAP.

7.1 Dataset Construction

Because our work constitutes the first effort towards method-level open-world app fingerprinting, there is no available public dataset where method-level labels of network flows are available. Therefore, we download 1000 apps from Google Play Store and collect app traffic in our testbed. When selecting apps, we give consideration to both popularity and comprehensiveness. First, we select the most popular apps according to the rank in Google Play Store in favor of better popularity. Second, the selected apps fall into 46 different categories to guarantee a reasonable comprehensiveness. Note that when selecting apps, we skip various mobile browsers and apps that need to sign up with a credit card to avoid potential legal risk. We generate app traffic instances with the aid of open-source tools Monkey and RERAN [16]. Monkey is an automated test tool, while RERAN is a record and replay tool for Android. For every app, we collect 50 traffic instances. To generate each traffic instance, we run apps on Google Pixel 2 smartphones with Monkey to generate pseudo-random streams of user events and simulate user operations. The maximum number of user events is set to 5000. Meanwhile, network traffic of smartphones is captured by Tcpdump and stored in pcap format for network flow extraction; logcat files are saved for network flow labeling. In a nutshell, we collect 50,000 traffic instances, consisting of 701,931 network flows. The total file size is more than 29 TB.

7.2 Open-World App Traffic Recognition

We first evaluate FOAP in recognizing network flows.

Experimental Setup. Since FOAP establishes a model for every individual app of interest, we construct a training dataset and testing dataset for each app. Specifically, we randomly split traffic instance set into two subsets: 40 instances for training and 10 instances for testing. To construct the training dataset of an app, say, besides A's traffic instances, we randomly choose n_T other apps to involve their traffic instances as negative samples in the training dataset. As for testing dataset, similar to the experimental setting [16], we consider another 20 apps, which are unseen A's training dataset. In other words, apps considered as negative class in testing stage are completely different from those for model training. To mimic human behavior, we simulate the use of different apps as Poisson process, which is commonly used to model human behaviors [31, 50]. Specifically, we construct network traces by merging traffic instances from different apps in random order. The time interval between traffic instances is exponentially distributed with an average time interval 1 second. Additionally, we also change the value of n_T to evaluate how it influences the accuracy of app recognition. We elaborate on the process of hyperparameter tuning in § A.

Baseline. We compare FOAP with the state-of-the-art tool AppScanner [39]. Other works, such as [10, 43], have different threat models. For example, [10] only handle unencrypted data and [43] consider destination features, which are unavailable in our threat model. Another line of works, such as [15, 22, 29] focus on different objects for analysis. For example, [29] analyzes packets in a series of time windows to identify user activities and [22] divides packet arrivals into segments and analyzes user activities within each segment. Therefore, we skip a direct comparison with them in our experiment. The feature set of the vanilla AppScanner is just a subset of FOAP's feature set, i.e., 51 features related to packet size characteristic and 6 features related to general characteristic. To conduct a fair comparison, we also extend AppScanner by using the full feature set of FOAP. AppScanner is a modular tool, comprised of different classifiers. Among these classifiers we choose Single Random Forest Classifier Per App to evaluate because i) it achieves the best accuracy and ii) it can be evaluated in the open-world setting. All parameters are consistent with those in the original paper [39].

Result. Table 1 reports the experimental result. Generally speaking, FOAP consistently outperforms AppScanner in terms of precision, recall, and F1-score for different n_T . For example, FOAP improves average precision with 0.3, average recall with 0.14, and average F1-score with 0.23, compared to AppScanner for $n_T = 20$. FOAP achieves the highest F1-score 0.911 when $n_T = 20$. Compared with AppScanner, a significant advantage of FOAP is its much higher precision. It is because FOAP takes advantage of inter-flow information while AppScanner doesn't. Such information is important for FOAP to reduce false positives in the open-world setting. Specifically, FOAP makes use of both intrinsic features from a network flow itself and inter-flow information

¹<https://developer.android.com/studio/test/monkey>

Table 1: Evaluating FOAP in open-world app traf c recognition (mean and standard deviation).

Tr	AppScanner						AppScanner (extended)						FOAP (our approach)					
	Precision		Recall		F1-Score		Precision		Recall		F1-Score		Precision		Recall		F1-Score	
5	0.315	0.171	0.866	0.105	0.439	0.174	0.307	0.158	0.840	0.111	0.427	0.166	0.759	0.253	0.925	0.094	0.803	0.200
10	0.476	0.183	0.814	0.127	0.582	0.161	0.470	0.178	0.774	0.128	0.568	0.159	0.886	0.170	0.906	0.111	0.881	0.140
15	0.568	0.189	0.782	0.137	0.641	0.158	0.575	0.182	0.729	0.144	0.627	0.155	0.926	0.123	0.897	0.114	0.902	0.111
20	0.641	0.180	0.755	0.146	0.679	0.151	0.646	0.164	0.700	0.148	0.659	0.143	0.945	0.102	0.897	0.126	0.911	0.115

to improve recognition accuracy. FOAP computes structural similarity to characterize how network flows within a time period are structurally similar to those from the app of interest. Based on this metric, FOAP can effectively filter out irrelevant network flows and significantly reduce false positives. On the contrary, AppScanner recognizes a network flow by only using its intrinsic features and thus it is prone to generate false positives caused by similar network flows. Additionally, AppScanner only recognizes which app a network flow comes from, while FOAP not only recognizes the app that generates a network flow but also identifies the specific EP methods that generate the traffic to infer fine-grained UI operations.

Case Studies of False Positives We observed two salient factors that increase false positive rate.

i) Share of third-party libraries Different apps often use the same third-party libraries (TPLs) to shorten development cycle, potentially resulting in similar network flows generated by them. For example, the comic app `com.tencent.wecom` uses the TPL `com.facebook`, which is widely used by other apps. When FOAP recognizes network traffic generated by this app, the average false positive rate (FPR) is 4.56×10^{-3} , while its FPR associated with the TPL `com.facebook` is 1.10×10^{-2} , significantly higher than the average FPR. A similar situation is found for the sport app `com.bundesliga`, which uses the TPL `com.bumptech.glide`. We observed its average FPR is 4.63×10^{-3} , while its FPR associated with the TPL `com.bumptech.glide` is 1.75×10^{-1} .

ii) Apps from the same developer Developers may reuse their codes when developing various apps. It aggravates the false positive risk when distinguishing traffic from these apps. For example, both `com.google.android.apps.meetings` and `com.google.android.apps.classroom` are developed by Google. To recognize network flows generated by the former app, FOAP reports more false positives caused by the latter app (FPR = 1.13×10^{-2}) than other apps not developed by Google (FPR = 2.36×10^{-4}). Another example is found for two apps developed by Microsoft. When recognizing network flows of `com.microsoft.office.outlook`, the app `com.microsoft.translator` produces more false positives (FPR = 1.62×10^{-2}) than other apps (FPR = 6.57×10^{-4}).

Case Studies of False Negatives We identified two key factors leading to false negatives.

i) Dynamic nature Diverse user behaviors bring highly dynamic traffic, especially for information-intensive apps, like news apps. In these apps, clicking on various articles/videos may trigger diverse network flows, thereby increasing false negative risk. Take an example from the news

app `bbc.mobile.news.ww`. FOAP recognizes its network flows with a false negative rate (FNR) of 5.8×10^{-1} . A closer look reveals the network flows missed by FOAP (i.e., FN) indeed exhibit substantial difference from network flows in the training dataset, since their average local similarity with this app is 0.291, much lower than that for network flows successfully recognized by FOAP, i.e., 0.753. Similar situations can be observed for map apps. For example, FOAP recognizes network flows generated by `com.baidu.BaiduMap` with an FNR of 0.107.

The network flows missed (resp. recognized) by FOAP have an average local similarity of 0.381 (resp. 0.832) with this app.

ii) Low-traffic apps We observed that some false negative cases for low-traffic apps are caused by traffic filtering. Specifically, some traffic instances of these apps may contain only a few network flows, and FOAP may erroneously filter out these network flows if they are mixed with other apps' network flows. For example, when recognizing network traffic generated by the app `official.pdfmaker`, FOAP suffers an FNR of 0.136 as it erroneously filters out network flows in 3 traffic instances, each of which contains only one network flow. In another example, we recognize network traffic generated by the app `tw.com.trtc.is.androideng`. The FNR is 0.160, and FOAP erroneously filters out network flows in 2 traffic instances of this app, each of which contains 2 network flows.

Generalized Open-World Setting In a more generalized open-world setting, the adversary will handle both app traffic and non-app traffic. To further evaluate how FOAP performs in such an setting, we collect various non-app traffic including

i) Mobile website traffic: It contains 580,335 network flows that are generated by visiting Alexa top 10000 websites using the android browser.

ii) IoT traffic from PINGPONG dataset [42]: It contains 378,797 network flows that are captured from 19 IoT devices.

iii) PC traffic: It is captured from 6 PCs, including 2 windows PCs, 3 macOS PCs, and 1 Linux Ubuntu PC for 1 day. There are 16801 network flows in total.

Table 2: False positive rate for open-world network traffic recognition (lower is better).

Open-World Traffic	AppScanner	AppScanner (extended)	FOAP (our method)
Mobile Web	241 10^{-2}	2.18 10^{-2}	9.93 10^{-4}
IoT	3.39 10^{-2}	2.85 10^{-2}	4.36 10^{-5}
PC	1.21 10^{-2}	2.02 10^{-2}	3.54 10^{-4}
SC-App	324 10^{-2}	3.09 10^{-2}	1.45 10^{-2}
DC-App	212 10^{-2}	1.99 10^{-2}	2.64 10^{-3}

All these non-app traffic will naturally be viewed as negative samples, because the goal of FOAP is recognizing the app of interest from network traffic. Therefore, we employ the metric false positive rate (FPR) to evaluate whether FOAP will generate false positives over these non-app traffic. Table 2 reports the experimental results. We found the false positive rates of FOAP for all non-app traffic are extremely low and much lower than those of baseline methods. It again verifies FOAP's advantage in the open-world setting. Other than non-app traffic, Table 2 also reports the FPR for app traffic. Intuitively, apps from the similar category may share common network behaviors, potentially resulting in more false positives. To test this hypothesis, we group apps used for negative class in testing into two categories, SC-App and DC-App. SC-App refers to apps belonging to the same category with A, while DC-App refers to apps whose categories differ from that of A. As shown in Table 2, FPR for SC-App is higher than that for DC-App, indicating more false positives. This result is consistent with the above hypothesis.

7.3 EP Method Identification

We next evaluate FOAP in identifying EP methods.

Experimental Setup. EP method identification is an app-specific multi-label classification task. Therefore, FOAP establishes a model for each app. Similar to the previous setting, we randomly split traffic instance set of each app into two subsets: 40 instances for training and 60 instances for testing.

Result. Multi-label classification can be viewed as the combination of multiple binary classification tasks. To evaluate the performance of FOAP in EP method identification, we consider a group of binary classification tasks corresponding to different EP methods. Specifically, we record the number of true positives, false positives, true negatives, and false negatives for each EP method and aggregate results of all EP methods in an app to compute the precision, recall, and F1-score. We employ micro-average method to compute these metrics due to the class imbalance in our problem.

Table 3: Evaluating FOAP in EP method identification (mean standard deviation).

e	Precision		Recall		F1-Score	
0.5	0.888	0.102	0.883	0.105	0.885	0.101
1	0.886	0.104	0.868	0.114	0.876	0.106
2	0.885	0.111	0.853	0.126	0.866	0.114
5	0.892	0.105	0.828	0.137	0.856	0.116

Table 3 reports the experimental results. FOAP achieves the highest average F1-score of 0.885 when $e = 0.5$ second. As the value of burst threshold τ increases, the precision of FOAP improves. The underlying reason is a large number of more packets (therefore more information) involved in each in-flow burst, which helps reduce false positives. The highest average precision of 0.892 is achieved when $e = 5$ seconds. A

(a) False positive rate (FPR). (b) False negative rate (FNR).

Figure 4: Impact of experimental data size.

side-effect of increasing n_1 is recall decline, which is caused by the majority voting mechanism used in burst labeling. We relax the multi-label classification problem to a series of multi-label classification subproblems at the cost of labeling a in-flow burst with only a single EP method, which implies a potential information loss of label. A larger n_1 will lead to more significant information loss and thus a lower recall. Consequently, the highest average recall of 0.883 is achieved when $e = 0.5$ second.

7.4 Impact of Experimental Data Size

To analyze how experimental data size influences the performance, we compare FOAP with baseline methods in both closed-world setting and open-world setting with varied app number in testing (i.e., $n_1 = 1; 5; 25; 125; 625$).

Experimental Setup. As for the closed-world setting, apps for model training are the same as those in testing. As for the open-world setting, apps that are considered as the negative class for model training and those in testing belong to two disjoint sets. We fix the number of apps that are considered as the negative class during model training to be 20 and change that in testing to $n_1 = 1; 5; 25; 125; 625$.

Result. We report the experimental results in Figure 4. Figure 4(a) presents a log-log plot about how FPR changes when n_1 increases. In general, FOAP consistently outperforms baseline methods as it has lower FPR. In the open-world setting, unlike FPR for baseline methods, which tends to be constant, FPR of FOAP gradually decreases as n_1 increases. The underlying reason is related to traffic filtering. When n_1 is small, false positives caused by "noise", e.g., background network flows, within relevant traffic segments dominate FPR. When n_1 is large, false positives are expected to be mainly caused by mis-classification of network flows from other apps. Thanks to the traffic filtering module of FOAP, most network flows from other apps are filtered out. Consequently, the FPR of FOAP decreases when n_1 increases. In the closed-world setting, FPR of all methods first drops because the diversity of negative samples for model training increases with n_1 (Recall that $n_T = n_1$ in the closed-world setting.) A large value of

results in extreme class imbalance, which is commonly recognized as a negative factor in machine learning. To mitigate the class imbalance, we employ the resampling strategy. Its side-effect is FPR for all methods rises again when Figure 4(b) shows how FNR changes with In the open-world setting, FNR of all methods tends to be constant, while in the closed-world setting FNR increases with indicating more false negatives. It is caused by class imbalance during model training, despite a resampling strategy has been employed. Compared with baseline methods, FOAP shows advantage again in both settings.

7.5 Impact of Third-Party Libraries

The share of third-party libraries, such as advertisement library and third-party app authorization, among different apps potentially increases false positives, because the same third party library may trigger similar network flows. We investigate this hypothesis in this experiment.

Experimental Setup. We analyze how third-party libraries influence the performance of FOAP and baseline methods in the open-world setting. τ and n_l are fixed to be 20. For each app, we first identify all TPLs used by it with the aid of Libradar [25]. Next, we correlate its EP methods with these TPLs. If the package name of an EP method matches a TPL we regard this EP method as an EP method belonging to the TPL. Recall that we have labeled network flows with EP methods by leveraging method-level labeling. Combining the above two kinds of information, we can label a network flow with a TPL if any EP method belonging to this TPL triggers this network flow. Otherwise, we label it as a network flow triggered by the app-specific library (i.e., app class).

Result. Table 4 reports FPR and FNR related with the top 5 third-party libraries in our dataset as well as the average performance for third-party libraries. As a comparison, we also present FPR and FNR related with app-specific libraries, which we refer to as app class. For all methods, the FPR related with third-party libraries is higher than that related with app class. Particularly, the FPR of FOAP in recognizing network flows associated with app class is 3.65×10^{-3} , while FPR rises to 6.80×10^{-3} for third-party libraries, which implies more false positives are generated. Another observation is third-party libraries do not increase false negatives. In fact, we even observe a slight FNR decline for FOAP. We conjecture a possible reason is when compared with app-specific libraries, third-party libraries may exhibit less diverse network behaviors because each of them is probably invoked for automatic test tools and humans, because i) it is less labor-intensive for automatic test tools to generate a large number of despite third-party libraries increase false positives, the FPR of FOAP is still reasonably low. In addition to third-party libraries, other factors may also influence the FPR of FOAP. We analyze the impact of automatic app generators in § C.

7.6 Cross-Dataset Evaluation

In this experiment, we evaluate to what extent FOAP trained based on a Monkey-generated dataset is able to recognize app traffic generated by humans.

Experimental Setup. We randomly choose 100 out of 1000 apps to construct a dataset that reflects the true human-generated behaviors. Specifically, we recruit 5 volunteers to manually operate these apps and generate 100 traffic instances, each of which lasts for about 5 minutes. We compare three transfer settings. In the “M M” setting, FOAP is trained over Monkey-generated dataset and tested over the same dataset. In the “M H” setting, FOAP is trained over Monkey-generated dataset and tested over human-generated dataset. In the “EM H” setting, the Monkey-generated dataset is enhanced by traffic instances generated by volunteers and FOAP recognizes the traffic instance generated by the remaining volunteer. For statistical soundness, we rotate the volunteer for testing and report the average performance.

Result. Table 5 reports experimental results for open-world app traffic recognition. Compared with “M M” setting, the recall of FOAP in “M H” setting reduces from 0.886 to 0.760, indicating some network flows generated by humans cannot be recognized and reported as false negatives. Nonetheless, it still outperforms baseline methods and even exhibits more significant advantage. Table 6 reports experimental results for EP method inference. Both precision and recall of FOAP exhibit a decline trend. We conjecture that the underlying reasons are two-fold. On one hand, Monkey test can hardly cover all UI operations that human may trigger in the wild. On the other hand, Monkey’s behavior is highly random and thus UI operations triggered by monkey tend to be contextually independent. On contrary, UI operations triggered by humans are contextually correlated. Such a difference potentially results in different code paths and different network flows when the same UI operation (e.g., clicking a button) is triggered by Monkey and humans respectively.

To further explore how to improve the performance of FOAP, we enhance Monkey-generated dataset by involving human-generated traffic instances in the “EMH” setting. Even though only 4 traffic instances are added to training dataset, the recall for open-world app traffic recognition is improved from 0.760 to 0.873. Likewise, both precision and recall for EP method inference rise. This experiment reveals a practical manner to construct training dataset when FOAP is applied in the wild. That is constructing FOAP’s training dataset by combining traffic instances generated by both humans and automatic test tools, because i) it is less labor-intensive for automatic test tools to generate a large number of traffic instances, which empower FOAP with a high precision and fewer false positives, and ii) traffic instances generated by humans are beneficial to reducing false negatives and improving the recall of FOAP.

Table 4: Impact of third-party libraries.

Library	Coverage	AppScanner		AppScanner (extended)		FOAP (our approach)	
		FPR	FNR	FPR	FNR	FPR	FNR
App Class	100%	26	3:10	2:18	3:04	3:65	1:26
TPL (average performance)	88%	357	3:30	3:27	3:29	5:80	1:21
com.google.android.gms	65%	515	2:75	4:26	3:48	9:32	2:20
com.facebook	54%	317	3:26	3:13	2:99	3:99	6:59
com.google.rebase	29%	257	4:43	2:51	4:15	2:35	9:99
com.bump.tech.glide	10%	385	2:29	2:18	2:09	4:53	8:15
com.urbanairship	9%	922	3:93	7:02	3:49	1:58	2:81

Table 5: Evaluating open-world app traf c recognition in cross-dataset experimental settings (standard deviation).

Transfer Setting	AppScanner			AppScanner (extended)			FOAP (our approach)			
	Precision	Recall	F1-Score	Precision	Recall	F1-Score	Precision	Recall	F1-Score	
M! M	0:625	0:167	0:775	0:681	0:135	0:634	0:148	0:700	0:136	0:655
M! H	0:506	0:222	0:521	0:481	0:185	0:436	0:212	0:508	0:201	0:439
EM! H	0:514	0:213	0:649	0:547	0:180	0:470	0:199	0:656	0:173	0:525

Table 6: Evaluating EP method identi cation in cross-dataset experimental settings (mean standard deviation).

Transfer Setting	Precision	Recall	F1-Score
M! M	0:879	0:098	0:874
M! H	0:771	0:143	0:783
EM! H	0:830	0:145	0:831

7.7 Running Time of FOAP

In addition to accuracy, system ef ciency also plays an important role. To this end, we evaluate the running time of FOAP by running it on a single core of a desktop equipped with an Intel Core i7-7700 CPU processor. FOAP takes an average of $1:77 \cdot 10^{-2}$ second to recognize whether a network ow is from the app of interest. To further identify method-level user actions, FOAP spends another $5:01 \cdot 10^{-2}$ second processing each network ow from the app of interest. The above experimental results indicate FOAP is highly ef cient when applied in practice. Table 7 reports ne-grained running time for different modules of FOAP.

Table 7: Running time of FOAP.

Task	Major Steps	Running Time (second/ ow)
Open-World App Traf c Recognition	Extract Flow Features	1:71 10^{-2}
	Traf c Segmentation	5:10 10^{-4}
	Traf c Filtering	3:61 10^{-6}
	Flow Recognition	7:73 10^{-5}
	Total	1:77 10^{-2}
Method-Level User Action Identi cation	Extract Burst Features	1:16 10^{-2}
	EP Method Inference	4:74 10^{-2}
	Total	5:91 10^{-2}

8 Application

We further demonstrate how FOAP can be applied in ne-grained user activity inference and user privacy analysis.

8.1 Fine-Grained User Activity Inference

A direct application of FOAP is inferring ne-grained semantic user activities from encrypted mobile traf c. We demon-

strate this application through the case study of Twitter. To this end, we rst map EP methods of Twitter to various user actions with the aid of UI Automator². Speci cally, we extract resource IDs of UI components that are associated with user actions listed in Table 8, (e.g., the button to share a tweet), using UI Automator and then compare them with the resource ID of EP methods obtained by FOAP’s network ow labeling module to achieve a one-to-one mapping between user action and EP method. By doing so, we can reasonably infer user activities characterized by these user actions through identifying EP method invocation from encrypted mobile traf c.

We recruit two volunteers for this experiment. To construct the training dataset, the rst volunteer operates Twitter to generate 20 traf c instances. The reason why we construct training dataset manually instead of using an automatic test tool (e.g., Monkey) is we consider many complicated user actions (e.g., retweet) in this case study and they are extremely dif cult, if not entirely impossible, to trigger these operations with an automatic test tool. However, designing a smarter automatic test tool is out of the scope of this paper. To test how FOAP performs, the second volunteer is required to operate Twitter and meanwhile record his operations, i.e., user actions and timestamps, to establish the ground truth. Additionally, he also randomly chooses another 20 apps and operates them to test whether FOAP generates false positives.

FOAP spends about 40 seconds processing all 1706 network ows (31 TP, 1673TN, 0 FP, 2 FN). Figure 5 illustrates the inference result. For ease of representation, we depict all network ows associated with user actions that have been successfully identi ed (denoted by F1, F2, F5) but omit other network ows even if they are from Twitter. FOAP successfully identifies 12 user actions. It also makes some mistakes. Speci cally, FOAP fails to report “search” at 12:05 second but misreports it at 13:52 second. At the 17 second, it erroneously identi es “retweet” as “favorite”. To summarize, FOAP achieves a high precision in recognizing network ows generated by Twitter. Besides, the accuracy for identifying ne-grained user actions is also reasonable.

²<https://developer.android.com/training/testing/ui-automator>

Table 8: EP methods for inferring Twitter user activities.

User Action	EP Method
Login	com.twitter.android.p1.onClick
Like	com.twitter.ui.tweet.inlineactions.InlineActionView\$b.onAnimationEnd
Favorite	com.twitter.android.va\$k.onClick
Retweet	xb.b.onClick
Follow	com.twitter.app.pro les.y0.onClick
Notification	tcb.onClick
Send Message	com.twitter.app.dm.widget.DMConversationMessageComposer.onClick
Navigation	com.google.android.material.tabs.TabLayout\$i.performClick
Search	km4.onClick

Figure 5: Inferring Twitter user activities.

Table 9: Packet sequences associated with searching doctor of different hospital departments in ECEasyBook.

Department	Packet Sequence (Byte)
General Clinics	673" 66#97#78" 1180#66" 97" 66" 66#
Gynaecology	673" 66#97#78" 1272#66" 97" 66" 66#
Dermatology	673" 66#97#78" 1266#66" 97" 66" 66#
Chinese Medicine	674" 66#97#78" 1306#66" 97" 66" 66#
Paediatrics	673" 66#97#78" 1238#66" 97" 66" 66#
Geriatrics	673" 66#97#78" 1292#66" 97" 66" 66#
Psychiatry	674" 66#97#78" 1247#66" 97" 66" 66#

"#" stands for inbound packets and " stands for outbound packets.

for Twitter. FOAP spends about 0.09 seconds in processing all 2250 network flows and recognizes 5 network flows generated by TraceTogether (no FP and FN). FOAP successfully identifies that the TraceTogether user (i.e., the second volunteer) clicks the button for positive test reporting, which invokes the EP method `gov.tech.bluetrace.fragment.EnterPinFragment$onViewCreated$2.onClick`. FOAP captures this sensitive information by identifying in-flow burst associated with this method.

In-Method Side-Channel Leaks. FOAP also helps the analysis of in-method side-channel leaks, where diverse packet sequences triggered by the same EP method may leak sensitive private information. To demonstrate how the adversary may infer user privacy, e.g., the illness of the user, from in-method side-channel, we present an example about ECEasyBook⁴, a clinic appointment app. In this app, the user searches doctors in different hospital departments will invoke the same EP method `com.ecbook.eceasy.ui.main.c.c$.onClick` but trigger network flows with some differences in packet sequence. For example, the difference between searching doctors of dermatology and searching doctors of paediatrics occurs in the twelfth packet. The packet size for dermatology is 1266 bytes, while that for paediatrics is 1238 bytes. By leveraging this difference, the adversary may infer the illness of the user. We list packet sequences associated with searching doctor of different hospital departments and highlight the major differences in Table 9. Although FOAP did not directly infer user privacy in this example, it plays an important role in locating privacy-sensitive methods, enabling the adversary to zero in on network flows that leak out user privacy.

Figure 6: Identifying positive reporting in TraceTogether.

8.2 User Privacy Analysis

FOAP facilitates user privacy analysis by locating network flows associated with privacy-sensitive EP methods.

Sensitive Method Identification. The invocation of sensitive EP methods in an app may leak out private user information, such as health status, sexual preference, and search queries. FOAP is able to identify them to infer user privacy.

The COVID-19 pandemic has rapidly spread across the world. To fight against the pandemic, contact tracing has been proven to be an effective strategy in monitoring virus spreading. To this end, lots of contact tracing apps are developed and published. Unfortunately, some of these apps may suffer the risk of privacy leakage. By analyzing the encrypted mobile traffic, one may infer if a user tests positive. Figure 6 presents an example about COVID-19 contact tracing appssizes [24, 47]. The experimental setting is similar to that

³<https://www.tracetgether.gov.sg/>

⁴<https://eceasybook.com/>

9 Discussion

9.1 Countermeasure

To defend against traffic analysis attacks, existing countermeasures aim to obfuscate traffic features using different strategies. There are three common obfuscation strategies: i) padding packets to obfuscate features related to packet sizes [24, 47], ii) injecting dummy packets to obfuscate features related to packet directions [17], iii) delaying packets

Figure 7: Design of our AProxy countermeasure.

to obfuscate features related to packet timing [24]. Inspired by IMProxy [7], a proxy-based obfuscation system against traffic analysis attacks on messaging apps, we design an obfuscation system, named by AProxy, to implement traffic obfuscation in the form of proxy-based relayers and incorporate two obfuscation strategies that have been proven to be effective in defending against WF attacks, i.e., packet padding and injection of dummy packets. We do not consider the packet delaying strategy because it greatly degrades user experience if applied to apps. As illustrated in Figure 7, AProxy is comprised of an AProxy client and an AProxy server. The AProxy client is deployed on the Android smartphone while the AProxy server is a remote proxy server. Before relaying the outbound traffic to the AProxy server, the AProxy client first adds padding to the outbound packets so that all packet sizes equal 1500 bytes (i.e., MTU). It also injects outbound dummy packets with MTU size into each flow. The arrival time of dummy packets follows a Poisson process. After receiving the manipulated traffic, the AProxy server removes padding and dummy packets, and further relays it to the app server. As for the inbound traffic, the AProxy client and server switch roles. The AProxy server adds padding and injects inbound dummy packets while the AProxy client removes them and relays purified inbound packets to the app.

We experimentally evaluate the effectiveness of AProxy. Figure 8 reports the defense effect. “Padding” represents only packet padding is conducted and “P+D(x)” represents the dummy packet number is x times the original packet number. Generally, more dummy packets result in a worse performance of FOAP. For example, when the dummy packet number is twice the original packet number, the average F1-score for app traffic recognition drops from 0.911 to 0.714. A counterintuitive example is the precision for P+D(2) is higher than that for P+D(1) in EP method identification. A closer look reveals the precision decline is mainly caused by dummy bursts. Specifically, dummy packets may form dummy bursts which are injected into original network flows and potentially undermine EP method identification as noises. FOAP will not distinguish these noises from original in-flow bursts and still predicts their labels, i.e., EP methods. Once the predicted EP methods are inconsistent with EP methods that exactly trigger the network flows where these noises are injected, false positives are generated. Therefore, more dummy bursts are expected to result in larger precision decline. At first glance,

(a) Open-world app traffic recognition. (b) EP method identification.

Figure 8: Evaluating AProxy against FOAP.

P+D(2) injects more dummy packets than P+D(1). However, P+D(2) injects about 0.02 dummy bursts per flow, while P+D(1) injects about 2.1 dummy bursts per flow. It is believed that denser dummy packets reduce inter-packet interval and extend to form bigger but fewer dummy bursts. Consequently, P+D(1) leads to a lower precision of FOAP than P+D(2).

9.2 Limitation

Separability of Flow. Since FOAP recognizes apps by analyzing network flows, the prerequisite for doing so is different network flows are separable. In other words, packets from the same network flow can be grouped together to extract flow features. It implies FOAP’s incompetence in the scenarios where network flows are inseparable, such as Tor.

Offline Recognition. FOAP works in an offline fashion, because some procedures, e.g., traffic segmentation and EP method inference, are not stream-oriented. To enable quasi-real-time analysis, an immediate improvement is specifying sliding time windows and analyzing network traffic within

each window independently. A larger window size potentially helps improve recognition accuracy, but introduces a higher computational overhead because more information will be involved in each time window. A smaller sliding step facilitates timeliness at the cost of more frequent analysis. Therefore, the proper window size and sliding step are chosen based on a trade-off between accuracy, timeliness, and overhead.

Model Integration. To achieve fine-grained open-world app fingerprinting, FOAP needs to integrate multiple machine learning (ML) models. On one hand, these models are necessary to tackle various challenging issues in our problem. On the other hand, each model may need separate data science effort for model training and parameter tuning, which increases difficulty, albeit controllable, to apply FOAP in the wild. Specifically, for each app of interest, FOAP collects its traffic instances and constructs an app-level training dataset to train ML models for traffic segmentation, traffic filtering, and flow recognition. FOAP also needs to construct a method-level training dataset to train the spatial-temporal context model for EP method inference. Fortunately, FOAP features

the capability of automatic labeling of training data, thereby drastically minimizing the overhead of training data preparation. Additionally, different models may need separate effort for parameter fine-tuning. Nevertheless, our experiments demonstrate default parameters of these models are extensively suitable across different apps, hence alleviating the difficulty in parameter fine-tuning.

Limited Open-Worldness. FOAP aims at open-world app traffic recognition. The number of available apps in the Google Play Store is about 3.5 million [36]. We evaluate FOAP using 1000 apps, around 29% of apps hosted in Google Play store. Such a percentage is comparable to those for open-world WF attacks in [33, 46]. There are around 1.9 billion websites in total [23], among which 46% considered 80,100 websites (around 0.04% of total websites) in the open-world evaluation, while [33] considered about 400 thousand websites (around 0.02% of total websites). Despite the number of apps used in this paper is substantially larger than that in most existing works [5, 14, 22, 29, 39], there is still a gap between our evaluation and realistic open world. We will collect more apps for evaluation in future work.

Synthesized Traffic. To reduce the cost for data collection, we use traffic instances as negative samples when constructing the testing data for different apps. Consequently, we generate synthesized testing traffic by merging traffic instances from different apps and employing Poisson process to mimic a random use of apps. While Poisson process has been widely used to model human behaviors [18, 31, 50], further investigation is needed to evaluate whether Poisson process accurately characterize the use of apps. First, the use of apps depends on human dynamic, e.g., work and rest, and thus the use of an app may exhibit temporally heterogeneous property which may not be well characterized by Poisson process. Second, the use of some apps may be temporally correlated for real human. However, we randomly arrange the order of different apps without considering their correlation. Therefore, the synthesized traffic resulting from merging traffic instances based on Poisson process might not exactly reproduce a real environment, and thus the performance of FOAP in the real world might be affected. We will further evaluate FOAP using non-synthetic traffic in future work.

10 Related Work

As an important branch of traffic analysis, traffic fingerprinting techniques aim to infer user behavior by leveraging statistical features of network traffic without accessing packet payload plaintext, thereby immune to traffic encryption.

App Fingerprinting. Our work falls into the category of app fingerprinting (AF), which focuses on the analysis of mobile network traffic. There are lots of efforts towards the identification of apps [1-4, 14, 39-41, 43]. For example, Taylor et al. proposed a modular framework AppScanner for automatic fingerprinting and real-time identification of Android apps

from encrypted network traffic [9, 40]; Aceto et al. proposed a multimodal deep learning framework, dubbed MIMETIC, for mobile encrypted traffic classification to identify different apps [5]. To carry out more fine-grained characterization of user behaviors, some works [5, 22, 29] aim at the identification of selected user activities of interest. For example, Saltaformaggio et al. presented NetScope, a technique that utilizes traffic behavioral clues to automatically build a detector for smartphone app activities [29]; Fu et al. proposed CUMMA to classify service usages of mobile messaging apps by jointly modeling user behavioral patterns, network traffic characteristics, and temporal dependencies [16]; Liu et al. developed a recursive time continuity constrained KMeans clustering algorithm for traffic flow segmentation and classified the segmented traffic to identify in-app user activities [24].

Compared with existing AF techniques, FOAP is more fine-grained because we conduct method-level fingerprinting and infer EP methods associated with various UI components. By doing so, FOAP is able to not only identify fine-grained user activities but also infer sensitive private information (e.g., illness). Besides, the vast majority of existing AF techniques (e.g., [15, 22, 29, 39, 40]) work under the closed-world assumption. In this paper, we consider a more practical but challenging scenario, i.e., open-world setting. In literature, [9, 10, 43] also consider the open-world scenario. Unfortunately, their methods cannot solve our problem because they have different threat models [43] took advantage of destination features while [9, 10] only handle unencrypted data.

Website Fingerprinting. Website fingerprinting (WF) focuses on inferring sensitive websites visited by users via encrypted proxies or anonymity networks, e.g., Tor. In the light of how to extract traffic features, existing WF techniques (e.g., [19, 27, 28, 30, 32, 33, 46, 48]) can be roughly categorized into the feature-engineering-based and the deep-learning-based. The former category [19, 27, 30, 46, 48] requires crafted features extracted using specific domain knowledge and employs a variety of machine learning algorithms, such as random forest, SVM, and KNN, to train classifiers for website identification, while the latter category needs either an enormous training set for automatic feature extraction [28, 32] or samples from a large number of similar tasks to conduct meta learning [33]. Due to different threat models, WF cannot be directly applied to our problem. Nevertheless, some of these works inspire us in traffic feature extraction. For example, we extract traffic features about interactive pattern inspired by [27]. In a recent work [46], Wang proposed three novel classes of precision optimizers to improve the precision of open-world website fingerprinting (WF) at the cost of recall decline. He also argued that it is more important to optimize the precision of WF attacks than their recall. Unfortunately, such a method is not applicable for our problem where precision and recall are comparably important.

11 Conclusion

In this paper, we took the first step to identify method-level fine-grained user action of Android apps in the open-world setting. To this end, a systematic solution dubbed FOAP was proposed. It features i) the capability of effectively reducing the false positive risk in open-world app recognition through adaptive traffic filtering and ii) method-level user action identification via synthesizing traffic and binary analysis. We have evaluated the effectiveness of FOAP through extensive experiments. FOAP significantly outperforms the baseline approaches by improving the F1-score from 0.679 to 0.911 for app recognition. It also achieves a reasonable accuracy in EP method identification. The average F1-score is up to 0.885. We also demonstrated the practicality of FOAP in fine-grained user activity inference and user privacy analysis.

Acknowledgements

We sincerely thank Prof. Andrea Continella for shepherding our paper and the anonymous reviewers for their constructive comments. This work is partly supported by Hong Kong RGC Projects (No. PolyU15223918, PolyU15227916) and National Science Foundation under Grant No. 1951729, 1953813, and 1953893.

References

- [1] “Introduce network operations on a separate thread,” <https://developer.android.com/training/basics/network-ops/connecting>, 2020.
- [2] G. Aceto, D. Ciuonzo, A. Montieri, and A. Pescapé, “Traffic classification of mobile apps through multi-classification,” in *Proc. GLOBECOM*, 2017.
- [3] G. Aceto, D. Ciuonzo, A. Montieri, and A. Pescapé, “Mobile encrypted traffic classification using deep learning,” in *Proc. TMA*, 2018.
- [4] G. Aceto, D. Ciuonzo, A. Montieri, and A. Pescapé, “Multi-classification approaches for classifying mobile app traffic,” *Journal of Network and Computer Applications*, vol. 103, pp. 131–145, 2018.
- [5] G. Aceto, D. Ciuonzo, A. Montieri, and A. Pescapé, “Mimetic: Mobile encrypted traffic classification using multimodal deep learning,” *Computer Networks*, vol. 165, p. 106944, 2019.
- [6] Arthur and Charles, “Naked celebrity hack: security experts focus on icloud backup theory,” <http://www.theguardian.com/technology/2014/sep/01/naked-celebrity-hack-icloud-backup-jennifer-lawrence>.
- [7] A. Bahramali, R. Soltani, A. Houmansadr, D. Goeckel, and D. Towsley, “Practical traffic analysis attacks on secure messaging applications,” in *Proc. NDSS*, 2020.
- [8] C. Barnhart, E. L. Johnson, G. L. Nemhauser, M. W. Savelsbergh, and P. H. Vance, “Branch-and-price: Column generation for solving huge integer programs,” *Operations research*, vol. 46, no. 3, pp. 316–329, 1998.
- [9] R. Bortolameotti, T. van Ede, M. Caselli, M. H. Everts, P. Hartel, R. Hofstede, W. Jonker, and A. Peter, “Decanter: Detection of anomalous outbound http traffic by passive application fingerprinting,” in *Proc. ACSAC*, 2017.
- [10] R. Bortolameotti, T. Van Ede, A. Continella, T. Hupperich, M. H. Everts, R. Rafati, W. Jonker, P. Hartel, and A. Peter, “Headprint: detecting anomalous communications through header-based application fingerprinting,” in *Proc. SAC*, 2020.
- [11] X. Cai, R. Nithyanand, T. Wang, R. Johnson, and I. Goldberg, “A systematic approach to developing and evaluating website fingerprinting defenses,” in *Proc. CCS*, 2014.
- [12] M. R. Chernick, W. González-Manteiga, R. M. Crujeiras, and E. B. Barrios, “Bootstrap methods,” 2011.
- [13] J. Cho and H. Garcia-Molina, “Estimating frequency of change,” *ACM Transactions on Internet Technology (TOIT)*, vol. 3, no. 3, pp. 256–290, 2003.
- [14] M. Conti, L. V. Mancini, R. Spolaor, and N. V. Verde, “Can’t you hear me knocking: Identification of user actions on android apps via traffic analysis,” in *Proc. CODASPY*, 2015.
- [15] Y. Fu, H. Xiong, X. Lu, J. Yang, and C. Chen, “Service usage classification with encrypted internet traffic in mobile messaging apps,” *IEEE Transactions on Mobile Computing*, vol. 15, no. 11, pp. 2851–2864, 2016.
- [16] L. Gomez, I. Neamtiu, T. Azim, and T. Millstein, “Reran: Timing-and touch-sensitive record and replay for android,” in *Proc. ICSE*, 2013.
- [17] J. Gong and T. Wang, “Zero-delay lightweight defenses against website fingerprinting,” in *Proc. USENIX Security*, 2020.
- [18] A. Gupta and R. Sekar, “An approach for detecting self-propagating email using anomaly detection,” in *International Workshop on Recent Advances in Intrusion Detection*. Springer, 2003, pp. 55–72.
- [19] J. Hayes and G. Danezis, “k-fingerprinting: A robust scalable website fingerprinting technique,” in *Proc. USENIX Security*, 2016.
- [20] J. Kim, J.-h. Park, and S. Son, “The Abuser Inside Apps: Finding the Culprit Committing Mobile Ad Fraud,” in *Proc. NDSS*, 2021.
- [21] J. Li, H. Zhou, S. Wu, X. Luo, T. Wang, X. Zhan, and X. Ma, “Towards fine-grained open-world android app fingerprinting,” https://github.com/jflixjtu/FOAP/blob/main/FOAP_full.pdf, 2021.

- [22] J. Liu, Y. Fu, J. Ming, Y. Ren, L. Sun, and H. Xiong, "Effective and real-time in-app activity analysis in encrypted internet traffic streams," in *Proc. SIGKDD*, 2017.
- [23] I. live stats, "Total number of Websites," <https://www.internetlivestats.com/total-number-of-websites/>, 2021.
- [24] X. Luo, P. Zhou, E. W. Chan, W. Lee, R. K. Chang, R. Perdisci *et al.*, "Httpos: Sealing information leaks with browser-side obfuscation of encrypted flows." in *Proc. NDSS*, vol. 11, 2011.
- [25] Z. Ma, H. Wang, Y. Guo, and X. Chen, "Libradar: fast and accurate detection of third-party libraries in android apps," in *Proc. ICSE*, 2016.
- [26] W. Meng, R. Ding, S. P. Chung, S. Han, and W. Lee, "The price of free: Privacy leakage in personalized mobile in-apps ads." in *Proc. NDSS*, 2016.
- [27] A. Panchenko, F. Lanze, J. Pennekamp, T. Engel, A. Zinnen, M. Henze, and K. Wehrle, "Website fingerprinting at internet scale." in *Proc. NDSS*, 2016.
- [28] V. Rimmer, D. Preuveneers, M. Juarez, T. Van Goethem, and W. Joosen, "Automated website fingerprinting through deep learning," *arXiv preprint arXiv:1708.06376*, 2017.
- [29] B. Saltaformaggio, H. Choi, K. Johnson, Y. Kwon, Q. Zhang, X. Zhang, D. Xu, and J. Qian, "Eavesdropping on fine-grained user activities within smartphone apps over encrypted network traffic," in *Proc. WOOT*, 2016.
- [30] A. Shusterman, L. Kang, Y. Haskal, Y. Meltser, P. Mittal, Y. Oren, and Y. Yarom, "Robust website fingerprinting through the cache occupancy channel," in *Proc. USENIX Security*, 2019.
- [31] K. C. Sia, J. Cho, and H.-K. Cho, "Efficient monitoring algorithm for fast news alerts," *IEEE Transactions on Knowledge and Data Engineering*, vol. 19, no. 7, pp. 950–961, 2007.
- [32] P. Sirinam, M. Imani, M. Juarez, and M. Wright, "Deep fingerprinting: Undermining website fingerprinting defenses with deep learning," in *Proc. CCS*, 2018.
- [33] P. Sirinam, N. Mathews, M. S. Rahman, and M. Wright, "Triplet fingerprinting: More practical and portable website fingerprinting with n-shot learning," in *Proc. CCS*, 2019.
- [34] R. Spreitzer, V. Moonsamy, T. Korak, and S. Mangard, "Systematic classification of side-channel attacks: A case study for mobile devices," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 1, pp. 465–488, 2017.
- [35] Statista, "Number of available apps in apple app store," <https://www.statista.com/statistics/779768/number-of-available-apps-in-the-apple-app-store-quarter/>.
- [36] Statista, "Number of available apps in google-play store," <https://www.statista.com/statistics/289418/number-of-available-apps-in-the-google-play-store-quarter/>.
- [37] Statista, "Number of smartphones sold to end users worldwide from 2007 to 2021," <https://www.statista.com/statistics/263437/global-smartphone-sales-to-end-users-since-2007/>.
- [38] C. Sutton and A. McCallum, "An introduction to conditional random fields for relational learning," *Introduction to statistical relational learning*, vol. 2, pp. 93–128, 2006.
- [39] V. F. Taylor, R. Spolaor, M. Conti, and I. Martinovic, "Appscanner: Automatic fingerprinting of smartphone apps from encrypted network traffic," in *Proc. EuroS&P*, 2016.
- [40] V. F. Taylor, R. Spolaor, M. Conti, and I. Martinovic, "Robust smartphone app identification via encrypted network traffic analysis," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 1, pp. 63–78, 2017.
- [41] M. Tian, P. Chang, Y. Sang, Y. Zhang, and S. Li, "Mobile application identification over https traffic based on multi-view features," in *Proc. ICT*, 2019.
- [42] R. Trimananda, J. Varmarken, A. Markopoulou, and B. Demsky, "Packet-level signatures for smart home devices," in *Proc. NDSS*, 2020.
- [43] T. van Ede, R. Bortolameotti, A. Continella, J. Ren, D. J. Dubois, M. Lindorfer, D. Choffnes, M. van Steen, and A. Peter, "Flowprint: Semi-supervised mobile-app fingerprinting on encrypted network traffic," in *Proc. NDSS*, 2020.
- [44] H. Varian, "Bootstrap tutorial," *Mathematica Journal*, vol. 9, no. 4, pp. 768–775, 2005.
- [45] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra *et al.*, "Matching networks for one shot learning," in *Proc. NIPS*, 2016.
- [46] T. Wang, "High precision open-world website fingerprinting," in *Proc. S&P*, 2020.
- [47] T. Wang, X. Cai, R. Nithyanand, R. Johnson, and I. Goldberg, "Effective attacks and provable defenses for website fingerprinting," in *Proc. USENIX Security*, 2014.
- [48] T. Wang and I. Goldberg, "On realistically attacking tor with website fingerprinting," *Proc. Priv. Enhancing Technol.*, vol. 2016, no. 4, pp. 21–36, 2016.
- [49] Q. Xu, J. Eрман, A. Gerber, Z. Mao, J. Pang, and S. Venkataraman, "Identifying diverse usage behaviors of smartphone apps," in *Proc. IMC*, 2011.
- [50] C. C. Zou, D. Towsley, and W. Gong, "Email worm modeling and defense," in *Proc. ICCCN*. IEEE, 2004.
- [51] O. Zungur, G. Stringhini, and M. Egele, "Libspector: Context-aware large-scale network traffic analysis of android applications," in *Proc. DSN*, 2020.

