

PAPER

Understanding Characteristics of Phishing Reports from Experts and Non-Experts on Twitter

Hiroki NAKANO^{†,††a)}, *Nonmember*, Daiki CHIBA[†], *Member*, Takashi KOIDE[†], Naoki FUKUSHI[†], *Nonmembers*, Takeshi YAGI^{†††}, Takeo HARIU[†], Katsunari YOSHIOKA^{††}, and Tsutomu MATSUMOTO^{††}, *Members*

SUMMARY The increase in phishing attacks through email and short message service (SMS) has shown no signs of deceleration. The first thing we need to do to combat the ever-increasing number of phishing attacks is to collect and characterize more phishing cases that reach end users. Without understanding these characteristics, anti-phishing countermeasures cannot evolve. In this study, we propose an approach using Twitter as a new observation point to immediately collect and characterize phishing cases via e-mail and SMS that evade countermeasures and reach users. Specifically, we propose CrowdCanary, a system capable of structurally and accurately extracting phishing information (e.g., URLs and domains) from tweets about phishing by users who have actually discovered or encountered it. In our three months of live operation, CrowdCanary identified 35,432 phishing URLs out of 38,935 phishing reports. We confirmed that 31,960 (90.2%) of these phishing URLs were later detected by the anti-virus engine, demonstrating that CrowdCanary is superior to existing systems in both accuracy and volume of threat extraction. We also analyzed users who shared phishing threats by utilizing the extracted phishing URLs and categorized them into two distinct groups - namely, experts and non-experts. As a result, we found that CrowdCanary could collect information that is specifically included in non-expert reports, such as information shared only by the company brand name in the tweet, information about phishing attacks that we find only in the image of the tweet, and information about the landing page before the redirect. Furthermore, we conducted a detailed analysis of the collected information on phishing sites and discovered that certain biases exist in the domain names and hosting servers of phishing sites, revealing new characteristics useful for unknown phishing site detection.

key words: phishing, Twitter, social network, cyber threat intelligence

1. Introduction

A phishing attack involves an attempt by an attacker to deceive a user into believing that a harmful website is authentic, with the aim of acquiring valuable information like account credentials or credit card details. Recently, phishing attacks have increased globally [1]–[4], especially attacks targeting mobile devices, with a 3.28-fold increase from 2020 Q2 to 2020 Q3 [3]. In addition to the traditional phishing attacks via e-mail and short message service (SMS) have been especially on the rise [5]. Smishing, a portmanteau of “SMS” and “phishing,” refers to phishing attacks that specifically exploit smartphone SMSs to deceive users into providing

sensitive information or clicking on malicious links. Attackers are exploiting SMS features for phishing: it can be sent with a phone number, with a much smaller namespace than an email address; it can be reliably pushed to cell phone subscribers when they are in range; and SMS is used for legitimate notifications and two-factor authentication, making it impossible to ignore completely.

The first step in timely combatting this ever-increasing number of phishing attacks is to collect a wider range of phishing cases that reach end users and continue understanding their characteristics. In fact, to that end, numerous studies have been conducted to measure and analyze phishing attacks [6]–[9]. The facts about phishing and the weaknesses of the countermeasures revealed by these studies at that time have helped improve the coverage of spam filters in email services (e.g., Gmail and Outlook), web browser blocklists (e.g., Google Safe Browsing [10] and Microsoft Defender SmartScreen [11], threat feeds (e.g., PhishTank [12] and OpenPhish [13]), and security analysis engines (e.g., VirusTotal [14] and urlscan.io [15]).

However, existing countermeasures are still insufficient when phishing messages reach end users and users encounter phishing sites. This raises the following question for us. *How can we collect phishing that reaches users bypassing existing countermeasures?*

In this study, we propose an approach that uses Twitter as a new observation point to immediately collect *actual phishing situations* encountered by users that have bypassed existing countermeasures and to understand the characteristics of such phishing. Some previous studies have also used Twitter as a source to extract *non-phishing* cyberattack information (e.g., vulnerability information and malware behavior information) [16]–[19] and limited phishing cyberattack information (e.g., search by fixed keywords or monitor only specific users) [18], [20], [21]. Specifically, these previous studies used Twitter posts of the cyberattack information by *security experts*, which allowed them to identify vulnerability information and indicator of compromises (IOCs) before they were published on databases that share vulnerability information (e.g., Common Vulnerabilities and Exposures numbers) and threat information (e.g., malicious URLs) such as the National Vulnerability Database [22] and VirusTotal [14]. While at first glance these studies appear to be close to what our study aims to do, they differ significantly in that our goal is to extract and analyze *phishing*-related information even from the actual situations that reach *non-experts*.

Manuscript received October 26, 2023.

Manuscript revised January 10, 2024.

Manuscript publicized March 1, 2024.

[†]The authors are with the NTT Security (Japan) KK, Tokyo, 101-0021 Japan.

^{††}The authors are with the Yokohama National University, Yokohama-shi, 240-8501 Japan.

^{†††}The author is with the NTT, Musashino-shi, 180-8585 Japan.

a) E-mail: hi.nakano.sec@gmail.com

DOI: 10.1587/transinf.2023EDP7221

Indeed a large number of non-experts have posted suspicious phishing attack-related cases on Twitter as alerts [23]. We are eager to immediately analyze the content of alerts they report as cases where phishing has reached users because existing countermeasures have been bypassed. These reports have the benefit of being more victim-centered and comprehensive than posts by security experts and potentially being used as new information for anti-phishing technology. Our challenge is to extract only phishing attack reports from a large number of irrelevant tweets in their everyday lives.

To this end, we propose CrowdCanary, a system capable of structurally and accurately extracting phishing information (e.g., URLs and domains) from tweets of experts and non-experts who have actually discovered or encountered phishing. CrowdCanary is a system that employs pre-selected keywords (e.g., phishing and scam) as input to identify and output phishing attack-related user reports. Additionally, CrowdCanary can collect a diverse set of tweets by automatically identifying and extracting new keywords that are often seen in such reports and adding them to the system. We evaluate the effectiveness of our malicious URL collection in CrowdCanary against security engines [14], as well as existing systems that collect attack information from Twitter [20], [24]. We also analyzed the differences between experts and non-experts and considered what approach should be taken to collect the information shared by non-experts. Finally, we discussed how the phishing information extracted by CrowdCanary could be analyzed to help protect actual end users.

Our primary contributions are as follows.

- We proposed CrowdCanary, a system that identifies reports of phishing attacks by both English and Japanese Twitter users with a high accuracy rate of 95% for evaluation data.
- We operated CrowdCanary for three months and were able to identify 38,935 phishing reports out of 19 million tweets and extract 35,432 phishing URLs. We confirmed that 31,960 (90.2%) of these phishing URLs were later detected by anti-virus engines, demonstrating the high accuracy of CrowdCanary's threat intelligence extraction
- We analyzed users who shared phishing reports and discovered that the majority of phishing reports detected by CrowdCanary were shared by non-experts. We showed that the threat intelligence reported by non-experts includes many URLs not included in the intelligence shared by experts, making it useful as a new observation point for phishing attacks from a more victim-friendly perspective.

This paper is an extended version of our paper presented at ARES 2023 [25]. Our previous paper proposed a system to detect reports of phishing attacks on Twitter by both experts and non-experts, evaluated its ability to detect them with high accuracy, and analyzed the differences between expert and non-expert reports. However, we did not perform a comprehensive analysis on the intelligence within the detected

reports, such as comparing the extracted useful information to actual phishing-specific data feeds, examining the actual attack infrastructure based on the detected phishing attacks. In this study, we analyzed information on phishing attack reports from a new perspective and uncovered previously unidentified insights (Sect. 7). The new contributions of this paper are as follows.

- We compared the phishing URLs detected by CrowdCanary to those from two data feeds specialized in phishing attacks and found that more than half of the phishing URLs detected by CrowdCanary represented unique threat intelligence. Furthermore, we discovered that CrowdCanary was able to identify about 80% of the common URLs more rapidly than the other two feeds, demonstrating its superior detection speed compared to existing technologies.
- We conducted an analysis of the domain names and hosting providers that attackers typically use to deploy phishing sites, using the collected phishing URLs. We found that the phishing sites in our study have a bias toward certain top level domains that are generally regarded as malicious, and that the hosting providers to which they are deployed are biased toward a small number of IP addresses, many of which are controlled by organizations in the United States.

2. Motivating Examples

In this section, we discuss examples of user-reported phishing attacks and the challenges of extracting URLs and domain names related to phishing attacks.

2.1 Reports on Phishing Message

With the increased usage of social media platforms and smartphones, people post phishing emails and SMSs content they discover or encounter [23]. Figures 1 (1), (2), and (3) show reports of phishing attacks posted by users on Twitter, which we refer to as cases (1), (2), and (3), respectively. These are examples where Twitter users discover or encounter a phishing email or SMS and share that information along with the tweet's text or a screenshot taken with their smartphone.

In case (1), a user discovered Google phishing emails. He/she used hashtags and mentions to alert Twitter users to the email title, the sender's email address, and the phishing URL. It's relatively easy for us to collect reports and extract information if the report includes alerting hashtags or mentions the company's official account, and if the threat intelligence is in the body of the tweet. In case (2), a user clicks on a URL in a phishing email, understands that he/she has arrived at a phishing site, and shares a screenshot of the email and his/her browser. You can find the URL and domain name related to the phishing site in the information. In case (3), a user shares a phishing SMS he/she received to get feedback because he/she are unsure if the information is real or fake.

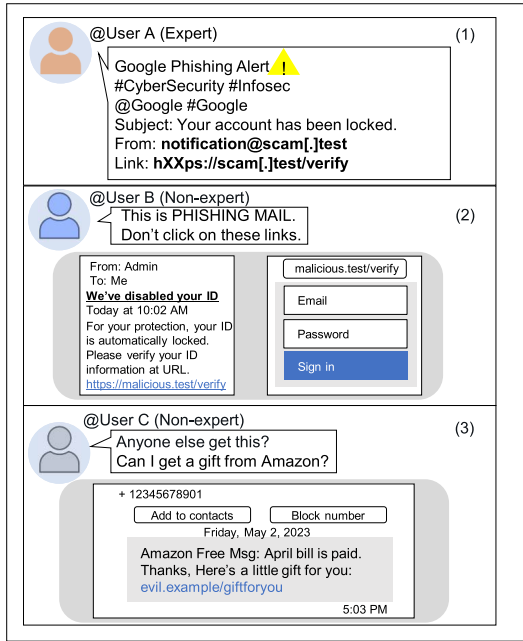


Fig. 1 Reports on phishing messages.

In addition to the URL in the SMS, the text of the tweet and SMS contains the company string “Amazon,” which was abused in the phishing attack. Compared to case (2), this case lacks keywords such as “PHISHING”. Therefore, to collect such phishing reports, we need to monitor Twitter at the right time and with the relevant keywords. Specifically, we need a system that can extract the keyword “Amazon” when phishing attacks with context related to “Amazon” are prevalent and promptly collect phishing reports from Twitter using that keyword. We will have important information about phishing attacks if we can extract URLs, domain names, and exploited company brand names as character strings from collected reports. Because this information is based on live phishing attacks that bypassed existing countermeasure technologies and reached end users, it is valuable to consider better countermeasure technologies to detect and prevent phishing attacks before they reach users.

2.2 Challenges

Collecting phishing-related posts from users and extracting only phishing-related information from them presents three challenges.

Collection of posts from various users on Twitter. There are a lot of tweets on Twitter, including phishing reports from security experts and non-experts. To examine them realistically, we need to collect the tweets as narrowly as possible. However, keywords commonly used by security experts in their reports, such as “#phishing”, are not always included in the reports of security non-experts. Therefore, we need to dynamically determine keywords to include in phishing reports and collect tweets at the right time to collect reports from a wide range of users.

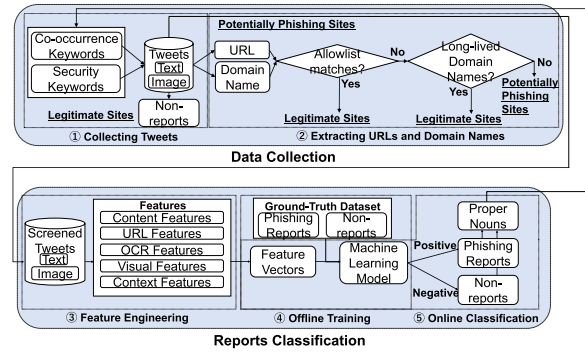


Fig. 2 Overview of CrowdCanary.

Extraction of information from collected user posts.

Phishing reports from non-experts are often presented in more diverse formats than those used by security experts. For example, phishing-related information may only be included in the image of a tweet, not in the body of the tweet. Without human intervention, it is difficult to determine whether the tweet is a report related to sharing information about phishing attacks from texts and images. Since we cannot manually analyze all tweets, we need a mechanical way to extract information from both the texts and images of a large set of tweets.

Validation of extracted information. It is necessary to extract only information about URLs and domain names related to phishing attacks from user reports. Some of the information we collect may be user-generated misinformation about legitimate sites or entirely unrelated to phishing attacks. As a result, we need to confirm the accuracy of the information extracted from the texts and images of the collected reports.

3. Proposed System: Data Collection

We propose CrowdCanary, a system that collects large-scale reports of phishing attacks in English and Japanese from Twitter users, including experts and non-experts, and allows for structured and accurate extraction of phishing information. We selected English and Japanese as the languages for our analysis because they are the top two languages used by Twitter users and thus likely to share information on phishing attacks using those languages [26]. Figure 2 shows an overview of CrowdCanary. CrowdCanary has two core components, *Data Collection* and *Reports Classification*. In this section, we describe the first component of CrowdCanary, *Data Collection*. This component takes keywords as input for searching tweets, collects data for *Report Classification*, and outputs them at one-hour intervals. The one-hour collection interval is a customizable system parameter. This component is designed to collect a wide range of tweets related to phishing attacks from different users. In addition, this component extracts information about URLs and domain names that are candidates for phishing sites from the collected tweets, and excludes information that is in a notationally invalid form or related to legitimate sites. This component consists of the following steps: *Collecting Tweets*

and *Extracting URLs and Domain Names*.

3.1 Collecting Tweets

In this step, we collect tweets using two types of keywords, *Security Keywords*, which are often used to share security information, and *Co-occurrence Keywords*, which co-occur with *Security Keywords* only at certain times. We use the Twitter Search API [27] as a means of collecting tweets. Otherwise, equivalent analysis can be performed using a stream of tweets as input, such as the firehose API [28] or the Decahose API [29] (10% random sampling of the firehose API). Since a large number of users on Twitter routinely post tweets that are unrelated to phishing reports, we considered that a search approach using appropriate keywords would be more efficient in collecting candidate reports of phishing attacks than an analysis of all such tweets or a random sampling of tweets.

Security Keywords. *Security Keywords* in this paper refers to keywords that are regularly posted on Twitter for cybersecurity-related information. *Security Keywords* allows us to collect tweets from security experts and tweets from non-security experts sharing phishing attacks they have discovered. Specifically, we select multiple keywords from two perspectives: related to the attack type (e.g., phishing) and information sharing (e.g., #infosec). The keyword defined as attack type (e.g., phishing) is sometimes used as a hashtag (e.g., #phishing), which is also included in the search. Finally, we selected the 20 security keywords in Table 1 for the following experiments. Based on previous researches [18], [19] and our preliminary study, we selected keywords most likely to be shared on Twitter for information about phishing sites. We also selected the same number of *Security Keywords* in Japanese as those translated from English.

In our preliminary study, we collected and analyzed 100,000 tweets using these common keywords (e.g., “attack” and “email,”) and found that more than 95% of the tweets were unrelated to phishing attacks. On the other hand, we also found that most tweets related to phishing attacks contained 20 selected *Security Keywords*. Specifically, 4,921 tweets, or 4.92% of the 100,000 tweets mentioned above, contained information about phishing attacks that had one of *Security Keywords*. Therefore, *Security keywords* selected in this study are reasonable for collecting and analyzing as many reports of phishing attacks as possible from many tweets on Twitter while reducing the number of false positives.

Co-occurrence Keywords. *Co-occurrence Keywords* in this paper are not directly security-related keywords, but keywords (e.g., Amazon and ATT) that co-occur with *Security*

Keywords at certain times and are included in non-expert tweets. The purpose of designing Co-occurrence Keywords is to collect as many phishing report attacks as possible that would otherwise be missed by Security Keywords. Specifically, *Co-occurrence Keywords* are extracted using the following procedure. First, we consider the tweets collected during the last period when the system is running as the *Co-occurrence Keywords* extraction target. The strength of association (*SoA*) is then calculated using the idea of pointwise mutual information (*PMI*). We define $P(X)$ and $P(Y)$ as the probability of the occurrence of a proper noun X and a type of tweet Y , respectively, in a given tweet. The probability that X and Y co-occur is $P(X, Y)$. In this case, *PMI* is represented by the following:

$$PMI(X, Y) = \log \left(\frac{P(X, Y)}{P(X)P(Y)} \right) \quad (1)$$

Next, we use positive pointwise mutual information (*PPMI*) as in the following equation to avoid the case where *PMI* goes to negative infinity (i.e., where $P(X, Y) = 0$).

$$PPMI(X, Y) = \max(0, PMI(X, Y)) \quad (2)$$

If X and Y do not occur at all in a single tweet, the *PPMI* will be 0. If X and Y are likely to occur in a single tweet, the *PPMI* will be positive or negative. Then, given a pair of proper nouns W in a tweet and a binary label L in the tweet (i.e., a phishing report or non-report), the *SoA* is given by the following equation:

$$SoA(W, L) = PPMI(W, L) - PPMI(W, \neg L) \quad (3)$$

If W appears only in phishing reports or non-reports, $PPMI(W, \neg L)$ is zero, then *SoA* is equal to *PPMI* ($SoA(W, L) = PPMI(W, L)$). Furthermore, W , which appears frequently in phishing and non-reports, has $PPMI(W, L)$ and $PPMI(W, \neg L)$ almost equal. As a result, $SoA(W, L)$ takes on a value close to zero. In other words, given a proper noun in a tweet for a given time period and a binary label of a phishing report or not, it is possible to extract keywords that are frequently found only in the user’s report for that time period. Since the common duration of the same phishing attack is 21 hours [30], we calculate the *PMI* for tweets within the previous 21 hours in our study. For the proper noun extraction task, we use the English model [31] and the Japanese model [32], which have been pre-trained on a large amount of data and confirmed to be highly accurate for this task. We evaluated whether we could extract as many brand names (e.g., Amazon, ATT, Microsoft 365) as possible from the aforementioned 100,000 tweets, and finally set the *SoA* threshold to 4. Then, the top 10 keywords that exceed the threshold are selected as *Co-occurrence Keywords*. The default state is no *Co-occurrence Keywords*, and *Co-occurrence Keywords* will be selected each time this step is performed.

3.2 Extracting URLs and Domain Names

This step extracts URLs and domain names potentially asso-

Table 1 Selected security keywords (English).

Keywords Related to Security Threats	Cyber Attack, Fake Site, Fraud, Scam, Malicious Site, Phishing, Opendir, Spam, Social Engineering, Smishing
Keywords with Frequent Shared Security Threats	#CyberCrime, #CyberSecurity, #CyberThreat, #IdentityTheft, #InformationSecurity, #InfoSec, #EmailSecurity, #ThreatHunting, #Threat, #Security

ciated with phishing attacks from the collected tweets. The extraction targets include both the texts and images contained in the tweets.

Image Analysis. We extract URLs and domain names from the images in the collected tweets by identifying the body area of the SMS or email. Specifically, we used YOLOv5 [33] as in the previous study [20], to identify body text areas in email or SMS screenshots, annotated with 3,000 images in the dataset described in Sect. 4.3. For the 3,000 images used for training, we analyzed valid thresholds with confidence scores ranging from 0.0 to 1.0 for the body text areas extracted by YOLOv5. As a result, all areas with a confidence score of 0.8 or higher corresponded to the body text area in the image. Therefore, in this study, if YOLOv5 extracts an area with a confidence score of 0.8 or higher, it is considered to be the body text area. Then, we use Tesseract [34] to extract character strings from the body text areas in both English and Japanese. If the body text area is not identified, we apply Tesseract to the entire image. We extracted text from English tweets using models pre-trained in English, while we extracted text from Japanese tweets using models pre-trained in both English and Japanese. This is because Japanese phishing emails/SMSs also contain English words.

Text Analysis. Next, we extract URLs and domain names from the text of images and tweets. Our study focuses on URLs and domain names that non-experts are likely to post as phishing attack information. Using regular expressions, we retrieved only the matches of URLs and domain names as candidate phishing sites from both the text of tweets related to the reports and the text derived from images. In particular, if there are defanged strings (e.g., example.com to example[.]com and http to hXXp) in a text, we refang the text (e.g., example[.]com to example.com and hXXp to http) and extract the URL and domain name matched by the regular expression.

Screening Phishing-related URLs and Domain Names. Finally, we exclude URLs and domain names that are incorrectly formatted or related to legitimate sites. Specifically, we check that it conforms to the format specified by RFC 3986 [35] and RFC 1035 [36]. If the URL or domain name that passed format validation is not included in both the image and the text, the tweet will be excluded from further analysis.

Then, we also exclude as legitimate sites any domain name in the top 10,000 on the Tranco list [37] and that does not match the shortened URL list [38]. Existing research [30] has shown that the registration of a domain name and the execution of a phishing attack can occur within a few days or tens of days at most. Therefore, we obtain domain name information from WHOIS and eliminate legitimate sites registered more than 365 days ago. CrowdCanary focuses on fresher domain names to detect newer phishing attacks, thus phishing sites that are more than one year old are excluded from our study. Furthermore, we exclude from our analysis false reports due to attackers sharing obviously legitimate sites or users mistakenly sharing legitimate sites. We out-

put any tweets with at least one or more domain names that remain after the screening as *screened tweets*.

4. Proposed System: Reports Classification

We describe the second component of CrowdCanary, *Reports Classification*, in this section. For the screened tweets obtained in the first component, we extract features in the tweets. Using supervised learning, we train a classifier to identify highly relevant reports of phishing attacks with high accuracy. From the created features, we select some features for training to achieve highly accurate and efficient classification. This step includes the following steps: *Feature Engineering, Training and Classification, and Evaluation of Classification Accuracy*.

4.1 Feature Engineering

We extract features from the screened tweets that help us identify user reports. This component classifies a single tweet as either a phishing report or a non-report. Specifically, we generated vectorizable features from Twitter user information, tweet body text, and images. Then, we selected helpful features from the generated features that improve the classification accuracy of phishing reports and non-reports using Boruta SHAP [41]. Boruta SHAP is a method that uses Shapley values for feature selection in Boruta, allowing for more accurate calculation of feature contributions and increasing the robustness of the Boruta algorithm [42]. Finally, we use the five types of features shown in Table 2: *Content Features, URL Features, OCR Features, Visual Features and Context Features*.

Content Features. From the content of the tweets collected in the previous component, we extract features relevant to identifying sharing related to phishing attacks, focusing mainly on the text. Our idea is straightforward: identify the actual content of the user’s tweet. We extract five features from the information in a user’s tweet. Specifically, we designed the following six types: number of characters (No. 1),

Table 2 List of features.

Feature Type	No.	Features Name	Vector Type	Dimensions
Content	1	# of characters	Integer	1
	2	# of words	Integer	1
	3	# of hashtags	Integer	1
	4	# of images	Integer	1
	5	Defanged type	Category	9
URL	6	Total # of characters	Integer	1
	7	# of characters in domain name	Integer	1
	8	# of digits	Integer	1
	9	Top level domain	Category	10
OCR	10	Number of characters	Integer	1
	11	# of words	Integer	1
	12	# of symbols	Integer	1
	13	# of digits	Integer	1
Visual	14	EfficientNet Vector [39]	Embedding	16
Context	15	BERT Vector [40]	Embedding	58
Total				104

number of words (No. 2), number of hashtags (No. 3), number of images (No. 4), and defanged type (No. 5).

Features No. 1 to No. 4 are each a vector of integer values obtained from tweets. Defanged type (No. 5) is a 9-dimensional feature vector with the one-hot encoding of 9 types of defanged types (“example.com”, “example[.]com”, “example(.).com”, “example{.}.com”, “example\com”, “hxxp://example.com”, “hXXp://example.com”, “http[:]//example.com” and “http://example.com[/]”). We believe that the number of characters and words in a warning-only post is relatively small. In addition, when users post reports, they often include numerous screenshots of emails and SMSs, and these features can efficiently identify user reports. Related studies [18], [43] have shown that these similar features can effectively determine whether a string contains warning information.

URL Features. We extract phishing site-specific features from the URLs contained in the texts and images of the screened tweets. Phishing sites often include characteristic strings in the domain name or path portion of the URL (e.g., abuse of subdomain names and long domain names) compared to legitimate sites [4]. It is possible to classify whether URLs are associated with phishing attacks by capturing the differences between the strings in the URLs of phishing sites and legitimate sites. Specifically, we designed the following four types: total number of characters (No. 6), number of characters in the domain name (No. 7), number of digits (No. 8) and top level domain (TLD) (No. 9).

No. 6 to No. 8 are the respective vectors of integer values calculated from the URLs (domain names) contained in the texts or images of the tweets. We conducted a preliminary survey of the TLDs in the ground-truth dataset (Sect. 4.3) and found 841 different TLDs. We investigated whether TLDs contribute to the identification of phishing sites using Boruta SHAP and identified 10 TLDs (“com”, “org”, “top”, “info”, “xyz”, “online”, “net”, “shop”, “cn” and “vip”) as important. TLD (No. 9) is a 10-dimensional feature vector with the one-hot encoding of 10 types of TLD, as mentioned above. For example, the fully qualified domain names (FQDNs) of phishing sites have more characters than those of legitimate sites, indicating subdomain abuse (e.g., login.security.account.example.com). In addition, Spamhaus reports that in 2023, TLDs such as “cn” and “top” have many cases of abuse [44] and may not be reviewed by registrars. As a result, TLDs abused by phishing sites tend to cluster in the same TLD.

OCR Features. We use Tesseract [34] to extract texts from the images in screened tweets. Reports of phishing attacks shared by people in images are typically screenshots of people’s smartphones, significantly different from other images commonly posted on Twitter. We can determine if the images in the tweets are related to the report of a phishing attack by performing OCR on the images and capturing differences in the extracted strings. If there is no image in a tweet, all OCR features are set to 0. If a tweet has multiple images, split it, create OCR features for each image, and classify all split tweets using the same other features.

Specifically, we designed the following four types: number of characters (No. 10), number of words (No. 11), number of symbols (e.g., !, ? and &) (No. 12) and number of digits (No. 13). No. 10 to No. 13 are the respective integer vectors calculated from the texts extracted by applying OCR to the tweet images. In addition to the URL and domain name, the image that the user shares as a phishing report includes the email or SMS text. In other words, texts and words that deceive users into clicking on URLs are also included in the extracted strings. Phishing SMSs and emails that deceive people have a predetermined amount of characters in a similar context (e.g., Your account has been suspended! Verify now [URL]), and hence the features differ significantly from strings extracted from other images.

Visual Features. We construct a fixed dimensional feature vector if the tweets obtained in the previous component contain images. Then, if there is no image in a tweet, the visual features vectors are set to 0. If a tweet has multiple images, split it, create visual features for each image, and classify all split tweets using the same other features. This feature captures the similarity in appearance of common phishing emails and SMSs.

Specifically, because emails, SMSs, and browser screenshots are usually images with a specific appearance, this feature is useful for classifying such images from other images. These images are essential for distinguishing phishing reports from non-reports, as they are included when users post information in the form of images. We use EfficientNet [39] as our visual feature generation model. We selected EfficientNet as the model for generating visual features since it is one of the state-of-the-art methods in image classification [45], [46]. We fine-tuned the model pre-trained on ImageNet (EfficientNet model) in English and Japanese with images related to the report (e.g., phishing email images and SMS phishing images) and images unrelated to the report (e.g., food images and landscape images). We successfully improved the feature generation to decide whether or not to include images related to the report.

We generate a 1,280-dimensional image feature vector from tweets using a retrained model. Then, we compressed the dimensions to achieve a cumulative contribution rate of 99% using TruncatedSVD [47], and the result was 16 dimensions for both English and Japanese. Here, we employ a fixed-dimensional vector, a compressed version of the vector created by the optimized EfficientNet model (No. 14).

Context Features. The contextual information from the tweet sentences obtained in the previous component is represented as a fixed-dimensional feature vector. When people share reports of phishing attacks, they often include alarming and angry statements, and are usually in a specific context. We cannot adequately capture these contexts based on the number of characters or words in a tweet. To this end, we use vectors created by a model trained on a large amount of text to capture the context of a tweet’s text.

Specifically, we use BERT [40] as the context feature generation model. BERT and BERT-based methods are state-of-the-art for several natural language processing

tasks [48]–[50]. We fine-tuned the sentences of tweets related to reports in both English and Japanese using the ground-truth dataset (Sect. 4.3). We optimized feature generation for a pre-trained model with many words to determine whether a tweet is related to user reports or not. In certain scenarios, a user who receives a phishing attack alerts, suspects, or incites the attacker. As a result, the contextual characteristics are different from other people’s daily posts.

We create a 768-dimensional context feature vector from tweets using a retrained model. Then we also compressed the dimensions to achieve a cumulative contribution rate of 99% using TruncatedSVD [47], and the result was 58 dimensions for both English and Japanese. Here, we use a fixed-dimensional vector, a compressed version of the vector generated by the optimized BERT model (No. 15).

4.2 Training and Classification

Using the many features we have created so far, we train a model for binary classification of whether a tweet is a report of a phishing attack or not.

Method. Given labeled positive or negative training data, a supervised learning model can be trained that uses the characteristics of each tweet to predict the binary value of tweets associated with phishing reports or non-reports. We then aim to predict with a high degree of accuracy whether new tweets are similar to previous phishing reports or non-reports. We compared and evaluated eight commonly used supervised learning algorithms: Random Forest, Neural Network, Decision Tree, Support Vector Machine, Logistic Regression, Naïve Bayes, Gradient Boosting, and Stochastic Gradient Descent. To account for the influence of some algorithms on accuracy loss, all feature vectors were preprocessed to set the mean to 0 and the variance to 1. Here, we train and evaluate using a ground-truth dataset labeled with phishing or non-phishing reports, which will be explained later in Sect. 4.3.

Results. We adopted Random Forest as the training and classification algorithm for the following three reasons. (1) Random Forest showed the best binary classification accuracy for the ground-truth data among the eight algorithms. (2) Random Forest performed consistently well with stable speed in both the training and inference phases for large amounts of data. (3) The importance of the features in the Random Forest was distributed among *Content Features*, *URL Features*, *OCR Features*, *Visual Features*, and *Context Features*, thus the classifier does not depend on any particular feature in its decision. We perform a classification accuracy evaluation on the ground-truth datasets (Sect. 4.3) and, in the live operation using CrowdCanary (Sect. 5), a model trained with the Random Forest algorithm, to perform the binary classification of phishing reports and non-reports.

4.3 Evaluation of Classification Accuracy

Before taking measurements with CrowdCanary in live operation, we evaluated the classification accuracy of phishing

Table 3 Ground-truth dataset for evaluating the accuracy of machine learning models.

Language	Collected Time	Label	# of Tweets
English	May. 1, 2021 – Jul. 19, 2021 (80 days)	Phishing Reports	5,000
		Non-Reports	15,000
Japanese	May. 1, 2021 – Jul. 19, 2021 (80 days)	Phishing Reports	5,000
		Non-Reports	15,000

Table 4 Classification accuracy evaluation results.

Language	Features	Accuracy	TPR	TNR	Precision	F-measure
English	All Features	0.957	0.952	0.962	0.962	0.957
	Content+URL+OCR	0.838	0.829	0.847	0.845	0.837
Japanese	All Features	0.949	0.948	0.960	0.951	0.943
	Content+URL+OCR	0.798	0.754	0.843	0.827	0.789

reports and non-reports in CrowdCanary.

Ground-truth Datasets. Table 3 shows the dataset used for the evaluation. First, we used the 20 English keywords from Table 1 and the 20 translated Japanese keywords. Then, we searched on Twitter using the keywords for 80 days from May 1, 2021 – Jul. 19, 2021, and collected 1,543,245 and 1,023,368 tweets in English and Japanese, respectively. Existing studies or publicly available datasets do not provide ground-truth datasets for the correct answers to phishing reports and non-reports, which are our research goals. As a result, we have to annotate them ourselves. Therefore, we randomly sampled the collected tweets and manually labeled them with a binary value of either phishing reports or non-reports. We excluded from our annotations tweets that do not have a URL or domain name in the text or image of the tweet. We then accessed the URLs and domain names in the text and images of the collected tweets from the experimental environment, examined the collected web content, and performed a similarity analysis with legitimate sites. Four security engineers conducted this annotation, and we labeled each of the tweets that we all agreed were reports of phishing attacks and non-reports. As a result of the annotations, we labeled the tweets as “phishing reports” when we determined they were related to phishing attacks and “non-reports” when they were not. Finally, we created 5,000 “phishing reports” and 15,000 “non-reports” in English and Japanese, respectively. To account for the effect of temporal bias, we split the training and testing data 7:3 in time order for the evaluation experiment.

Evaluation Results. The evaluation results are shown in Table 4. When combining all features (Content+URL+OCR+Visual+Context) for the English case, Accuracy was 0.957, True Positive Rate (TPR) was 0.952, True Negative Rate (TNR) was 0.962, Precision was 0.962, and F-measure was 0.957. The results show that the accuracy is sufficient to classify phishing reports from the large volume of tweets collected. We also found that it is difficult to detect user reports of phishing attacks with high accuracy using only simple features generated from meta information on Twitter. The same result is obtained for the Japanese case. We conclude that feature vectors with information em-

bedded in a fixed dimension, pre-trained on many languages and images, significantly improve classification accuracy. To summarize, in subsequent evaluations for Sect. 5, we will use a machine learning model trained by combining five types of features: Content+URL+OCR+Visual+Context.

5. Evaluating User Reports in the Wild

We used CrowdCanary, which was confirmed to detect user reports with high accuracy in Sect. 4.3, to classify unknown tweets in the wild. We then performed a comparative evaluation with two existing systems [20], [24] that collect and publish malicious URLs and domain names from Twitter.

5.1 Operating Environment

We operated the proposed system in a virtual machine (VM) on Azure. Specifically, we used the Linux OS Ubuntu 20.04 on a Standard D32as v4 (32 vCPU, 128 GB Memory) VM. We used twscrape [51], an open source tool, as the means of data retrieval from Twitter, and scikit-learn [52] for the analysis process related to machine learning. We employed luigi [53], a Python-based pipeline framework, to ensure that each task can be efficiently scheduled and processed in all sub-steps. With the operational environment described above, the proposed system operated without error during all the experimental periods in this study.

5.2 Datasets for Evaluation

A summary of the datasets for CrowdCanary and the two existing systems for comparison is shown in Table 5. These two existing systems collect information from Twitter, but the information they collect is not limited to phishing attacks. Although CrowdCanary focuses specifically on phishing attacks, we demonstrate that the quantity and quality of information collected by CrowdCanary outperforms the two existing systems. While CrowdCanary is a newly implemented system that works perfectly on the current version of Twitter, the existing systems rely heavily on older Twitter APIs and are unable to analyze the latest tweets. Therefore, we used datasets [24], [54] from when these systems were publicly available for our evaluation.

Proposed System (CrowdCanary). We ran CrowdCanary continuously every hour for three months, from Nov. 1, 2022 – Jan. 31, 2023. We set the *Security Keywords* to

20 English and 20 Japanese words in Table 1, and the initial state of the *Co-occurrence Keywords* to none. CrowdCanary selected new *Co-occurrence Keywords* every hour from the collected user reports. During the two-month experiments, we collected 18,765,699 tweets, screened 324,589 tweets, and identified 38,935 phishing reports. For domain names included in user reports, we considered them to be URLs by appending the protocol “https” to the domain name. Finally, we merged these URLs with the extracted URLs to obtain 35,432 unique URLs extracted by CrowdCanary.

Existing System (SpamHunter). We selected the dataset of the previous study [20] as our existing system for comparison. Their “SpamHunter” system collects tweets with SMS-related keywords, performs image analysis, and extracts phishing-related URLs. SpamHunter comes closest to our motivation in terms of the information we want to collect, however their method of collecting tweets is very limited. This is because SpamHunter only analyzes tweets when the keyword “sms” is included in the body of the tweet, which results in a large number of non-expert tweets that should be analyzed being missed. They published the collected URLs [54], and obtained 15,553 threats from Jan. 1, 2018 – Aug. 31, 2022. In addition, we added “https” to threats that lacked protocol information, excluded URLs with formatting deficiencies, and finally prepared 15,269 detected URLs.

Existing System (Twitter IOC Hunter). Next, we selected the existing system [24] for comparison because it extracts cybersecurity-related information (e.g., malicious URLs, IP addresses, etc.) from Twitter and allows us to obtain data for a specified time period through its API. We obtained 10,092 threats using the API of Twitter IOC Hunter [24] from Aug. 1, 2021 – Jul. 31, 2022. Similar to SpamHunter, we added “https” to threats that lacked protocol information, excluded URLs with formatting deficiencies, and finally prepared 9,344 detected URLs.

5.3 Comparison of Maliciousness using VirusTotal

We analyzed how VirusTotal (VT) [14] flags the URLs detected by CrowdCanary and the two existing systems [20], [24]. When we request VirusTotal to scan a URL, it evaluates the maliciousness of about 70 different types of anti-virus software and returns the results to us. Several studies [18], [55]–[58] used VirusTotal as a metric for evaluation. Then it is appropriate for our study to evaluate how much of the information collected from Twitter are actually malicious URLs.

VirusTotal provides five types of results for scanned URLs: malicious, suspicious, harmless, undetected and timeout. Because CrowdCanary immediately collects/ outputs phishing attacks shared by Twitter users, sometimes VirusTotal does not detect them even though the URLs are malicious. We then requested scans and obtained results at least one week after detection in CrowdCanary. Since the URLs of the existing systems had already mainly been analyzed by VirusTotal, we obtained the results of the scans. If VirusTotal had no previous scan results, we requested

Table 5 Overview of datasets for evaluation.

System	Period	Datasets	#
CrowdCanary	Nov. 1, 2022 – Jan. 31, 2023 (3 months)	Collected Tweets	18,765,699
		Screened Tweets	324,589
		Phishing Reports	38,935
		Detected Threats	42,987
		Detected URLs	35,432
SpamHunter	Jan. 1, 2018 – Aug. 31, 2022 (56 months)	Detected Threats	15,553
		Detected URLs	15,269
Twitter IOC Hunter	Aug. 1, 2021 – Jul. 31, 2022 (12 months)	Detected Threats	10,092
		Detected URLs	9,344

Table 6 Overview of comparison results between CrowdCanary and existing systems.

System	VT \geq 1	VT \geq 5	Total	VT \geq 1/day	VT \geq 5/day
CrowdCanary (I+T, E+J)	31,960 (90.2%)	15,768 (44.5%)	35,432 (100.0%)	347	171
CrowdCanary (I, E+J)	17,633 (84.2%)	7,267 (37.8%)	19,205 (100.0%)	187	83.2
CrowdCanary (T, E+J)	15,260 (88.6%)	8,452 (49.1%)	17,231 (100.0%)	164	88.5
CrowdCanary (I+T, E)	17,126 (91.2%)	7,558 (40.2%)	18,779 (100.0%)	186	82.2
CrowdCanary (I+T, J)	14,834 (89.1%)	8,210 (49.3%)	16,653 (100.0%)	161	89.2
CrowdCanary (I, J)	7,528 (83.4%)	4,107 (45.5%)	9,026 (100.0%)	81.8	44.6
CrowdCanary (T, J)	7,081 (87.5%)	4,014 (49.6%)	8,093 (100.0%)	77.0	43.6
CrowdCanary (I, E)	9,048 (88.9%)	3,583 (35.2%)	10,178 (100.0%)	98.3	38.9
CrowdCanary (T, E)	8,286 (90.8%)	4,244 (46.5%)	9,126 (100.0%)	90.1	46.1
SpamHunter (I, E) [20]	8,266 (59.8%)	1,718 (10.9%)	15,269 (100.0%)	4.85	1.01
Twitter IOC Hunter (T, E) [24]	5,228 (56.0%)	2,172 (23.2%)	9,344 (100.0%)	14.3	5.95

I: Image analysis, T: Text analysis, E: English analysis, J: Japanese analysis

a scan and obtained the scan results. VirusTotal has also seen cases of false positives from anti-virus vendors [56]; therefore, URLs identified as malicious/suspicious by one anti-virus vendor are not necessarily phishing URLs. As a result, in our study, we compared CrowdCanary and the two existing systems in terms of the number of URLs flagged as malicious/suspicious by at least one and as many as five anti-virus vendors in VirusTotal. The comparison results are shown in Table 6. We conducted the analysis with images and text as the threat information extraction targets, and with English and Japanese as the tweet collection languages. Focusing on URLs that were flagged as positive by five or more antiviruses in VirusTotal, 15,768 (44.5%) were positive for CrowdCanary (Image+Text), 7,267 (37.8%) were positive for CrowdCanary (Only Image), 8,452 (49.1%) were positive for CrowdCanary (Only Text), 1,718 (10.9%) were positive for SpamHunter and 2,172 (23.2%) were positive for Twitter IOC Hunter. We confirmed that CrowdCanary was superior to the proposed and existing systems in terms of both the absolute number and detection rate of URLs later detected by multiple antiviruses in VirusTotal. Early detection of URLs that will later be detected by VirusTotal is important for the future development of countermeasure technology. SpamHunter is a system that extracts information from tweet images shared in English, and Twitter IOC Hunter is a system that extracts threats from tweet texts shared in English. Due to the different experimental periods of the proposed system and the two existing systems, we compared the average per day of URLs detected by VirusTotal. In this case as well, the results showed that CrowdCanary was superior to the existing systems. When CrowdCanary’s threat information extraction is limited to images and English (equivalent to SpamHunter’s analysis target with CrowdCanary (I, E) in Table 6), CrowdCanary extracted 20 times and

39 times the amount of malicious URLs for VT \geq 1/day and VT \geq 5/day, respectively, compared to SpamHunter. Additionally, when CrowdCanary’s threat information extraction is limited to texts and English (equivalent to Twitter IOC Hunter’s analysis target with CrowdCanary (T, E) in Table 6), CrowdCanary extracted 6 times and 8 times the amount of malicious URLs for VT \geq 1/day and VT \geq 5/day, respectively, compared to Twitter IOC Hunter.

In addition, we manually investigated the remaining 3,472 (35,432 – 31,960) URLs that VirusTotal did not detect during the experimental period. We identified malicious URLs that could be identified as phishing sites based on the content of tweets, website content, screenshots, WHOIS information, etc. As in Sect. 4.3, this investigation was conducted by four security engineers and took a total of 30 hours to check for undetected URLs in VirusTotal. As a result, we found that 2,635 (7.44%) URLs were truly phishing sites (false negatives by VirusTotal). Most of these URLs were used for redirects under the domain names of duckdns.org, which abused the dynamic DNS provider, and cutt.ly, which abused the URL shortening service and made it difficult to determine the maliciousness of the URLs mechanically. On the other hand, 482 (1.36%) URLs were incorrect information due to OCR misidentification (e.g., misidentifying “1” as “l”), 160 (0.45%) URLs were not phishing site URLs included in the user’s report (e.g., minor legitimate sites that users cannot accurately determine whether they are phishing or not), and 195 (0.56%) URLs were misclassified by the machine learning model (e.g., legitimate SMSs or emails). The next Sect. 6 analyzes a total of 34,595 (31,960 + 2,635) URLs detected by VT or manually identified as malicious URLs.

6. Comparison of Experts and Non-Experts

We analyze the reports collected by CrowdCanary with a focus on the characteristics of the users (i.e., security experts or non-experts). In this section, we use 34,595 URLs (32,813 phishing reports) containing malicious information related to phishing attacks identified by VirusTotal and manual investigation in Sect. 5.3.


6.1 Analysis of Users who Shared Reports

Of the 32,813 phishing reports, the number of unique users was 9,025. We identify the users who shared these reports as experts or non-experts. Specifically, users who satisfy either of the following two conditions are considered experts, and users who satisfy neither of the two conditions are considered non-experts. (1) The user has security-related keywords (e.g., phishing, threat hunter) in their Twitter profile. (2) The user has posted more than half of their last 10 tweets related to cybersecurity.

As a result, we categorized users in the method described above, resulting in 25 users (2.77%) as experts and 9,000 users (97.23%) as non-experts, as shown in Table 7. We reviewed the results as a manual and verified that they

Table 7 User categorization results.

User Type	# Users	# Reports	# Shared			
			min	median	mean	max
Expert	25	15,263	1	280	610	3,900
Non-expert	9,000	17,577	1	1	1.95	73


**Fig. 3** Correlation between users and number of times reports were shared.

were categorized as intended. We found that experts share phishing reports an average of 610 times, while non-experts share phishing reports an average of 1.95 times. In particular, we confirmed that many expert shares appeared to be mechanical, with some accounts only posting phishing attack threats up to 3,900 times during the experimental period. Most non-experts shared phishing emails and SMS messages they received only a few times. However, in rare cases, we found some non-experts who shared phishing emails and SMS messages they received 73 times during the experimental period.

Additionally, Fig. 3 shows the correlation between users and the number of times reports are shared. The x-axis represents the number of times a user shared a report, the blue bar on the y-axis represents the number of reports based on the number of times the report was shared, and the red line on the y-axis represents the cumulative distribution function (CDF) value of the reports. From Fig. 3, users who shared only one report accounted for 53.1% of the total, while users who shared two reports accounted for 78.6% of the total. In other words, if we collect information from Twitter limited to accounts of users who frequently share, as in existing studies [17], [18], we would miss phishing reports from numerous users. We demonstrated that CrowdCanary can collect not only the limited information shared by security experts, but also information posted by a large number of users, including reports of phishing attacks by non-experts.

6.2 Analysis of the Detected URLs' Characteristics

We analyzed the value of the URLs included in the phishing reports. Specifically, we analyzed the number of times each URL was shared as a phishing report. The correlation between user types (i.e., experts or non-experts) and the number of times reports containing that URL were shared is shown in Fig. 4. The x-axis represents the number of times a URL has been shared, the green and orange bars on the y-axis represent the number of URLs found that match, and the red line represents the CDF value of the unique URLs.

**Fig. 4** Correlation between users types and number of times URLs were shared.**Table 8** Comparison of URLs characteristics.

User Type	# URLs	# Shortened URLs	# Services
Expert	16,778	102 (0.61%)	7
Non-expert	18,654	2,896 (15.5%)	13

Table 9 Comparison of FQDNs characteristics.

User Type	# FQDNs	# Dynamic DNSs	# Providers
Expert	6,530	668 (10.2%)	1
Non-expert	8,699	3,612 (41.5%)	3

From Fig. 4, URLs extracted from phishing reports shared only once by users accounted for 77.5% of the total, while URLs extracted from user reports shared twice by users accounted for 90.8%. As shown in our results, we found that extracting information from the tweets of a fixed set of users with a limited observation target would miss the majority of high-value malicious URLs that are shared a few times at most.

We then analyze the characteristics of the URLs and FQDNs shared by security experts and non-experts. The results are shown in Tables 8 and 9. The unique URLs included in the expert and non-expert reports were 16,778 and 18,654, respectively. Attackers sometimes use redirects from the landing URL to the phishing site where they ultimately want to direct the user [30]. Specifically, we investigated how many URLs exploited the dynamic DNS providers [59] and URL shortening services [38] used to redirect phishing attacks. Among dynamic DNS providers, duckdns.org was found to be abused 99.3% in total, and among URL shortening services, cutt.ly and bit.ly were abused 70.5% in total. Because these services and providers are free, can generate a large number of URLs, and have no countermeasures to exploit for phishing attacks, it is believed that attackers use them to evade detection (i.e., spam emails and SMSs detection) of phishing sites they have created. Many of the threats shared by non-experts are URLs that are actually spread in phishing e-mails and SMSs. These URLs can be used as a starting point for analyzing the full picture of attacks, or as intelligence for block lists that automatically detect spam e-mails and SMSs. Many experts share URLs after redirects, which sometimes cannot be analyzed because they are inaccessible without the proper referrer [60], and are not suitable information to prevent the spread of phishing emails and SMSs.

Table 10 Comparison of report sharing methods.

User Type	# URLs in Images	# URLs in Texts	#Hashtags		#Mentions	
			median	mean	median	mean
Expert	1,523 (10.0%)	13,740 (90.0%)	4	3.83	0	0.21
Non-expert	16,659 (94.8%)	918 (5.22%)	0	0.73	0	0.15

6.3 Analysis of Report Sharing Methods

We analyze the differences in the way experts and non-experts share information. First, we compared experts and non-experts on how users share information about phishing attacks. The results are shown in Table 10. We found a significant difference in how information was shared: 90% of expert reports included URL information in the text of their tweets. In contrast, 95% of non-expert reports included URL information in the images of their tweets. Experts identify threats through their own investigation rather than by encountering them, and they often share the information in a formatted text (in the text of a tweet). On the other hand, non-experts often store the phishing attacks they encounter it (receiving an email or SMS, or reaching the site with a browser) as screenshots from their smartphones, etc., and attach the images directly to their tweets and share them. Although it is difficult to collect a large number of these reports from non-experts and extract information properly, CrowdCanary was able to extract as many threats as experts and more, as shown in Fig. 4.

We also found significant differences in features between experts and non-experts in the context of the text when sharing reports. The median and the mean number of hashtags and mentions in the phishing reports of experts and non-experts are shown in Table 10. Hashtags are referred to as “#phishing” and are primarily used by users on Twitter to share information. People looking for information can find tweets containing the hashtag relatively easily using the search function. In this case, the expert report shows an average of 3.83 hashtags in the tweets, while the non-expert report shows an average of only 0.73 hashtags. As a result, collecting non-expert reports with appropriate keywords is more difficult than collecting expert reports shared using fixed hashtags. Similarly, we examined user reports that included mentions that could be posted to a specific user account on Twitter and found no significant differences between experts and non-experts.

Finally, we discuss query keywords that were useful in collecting phishing reports. The top 10 keywords that resulted in the collection of expert and non-expert reports are listed in Table 11. Among the top 10 keywords for experts, 8 were hashtagged and 4 were security (as defined in Sect. 3.1) keyword types. In particular, we found that a large number of experts shared their information using the hashtags “#infosec” and “#cybersecurity”, which are not commonly used by non-experts. On the other hand, only 3 of the top 10 non-expert keywords were hashtagged. Although “#phishing” was sometimes the most effective keyword for

Table 11 Top 10 keywords collected phishing reports.

Rank	Keywords (Expert)	Type	#	Keywords (Non-expert)	Type	#
1	#phishing (E+J)	S	3,982	#phishing (E+J)	S	1,545
2	#scam (E+J)	S	2,733	National Tax Agency (J)	C	1,175
3	#phishingmail (J)	C	1,406	Fraud (J)	S	1,035
4	#infosec (E)	S	1,283	#Amazon (E+J)	C	889
5	#cybersecurity (E)	S	1,280	#scam (E+J)	C	820
6	#Amazon (E+J)	C	1,119	Softbank (J)	C	712
7	#phishingsite (J)	C	1,079	Docomo (J)	C	688
8	National Tax Agency (J)	C	1,022	American Express (E+J)	C	653
9	SMBC (J)	C	894	Google (E+J)	C	512
10	#bank (E)	C	822	Please retweet (J)	C	488

E: English only, J: Japanese only, E+J: Contains identical semantic words in both languages

S: Security Keywords, C: Co-occurrence Keywords

collecting phishing reports, as it was for experts, many of the non-experts shared reports using the name of the company brand that was exploited in the phishing attack. However, simply searching for a company’s brand name will return a number of irrelevant tweets. Therefore, either a search using appropriate keywords at the right time, as in this study, or a highly accurate detection mechanism from among the tweets continuously collected by company brand name is required.

7. Analyzing Phishing Attacks in User Reports

To deepen our understanding of phishing attack reports shared on Twitter, we evaluate the effectiveness of information regarding countermeasure techniques and analyze the actual phishing attack infrastructure.

7.1 Analysis of Common URLs with Existing Data Feeds


We collected two types of datasets for comparative evaluation, OpenPhish [13] and PhishTank [12], both specialized for phishing attacks. OpenPhish is an open feed of large-scale data on phishing, and various existing countermeasure technologies reference the OpenPhish dataset. PhishTank is a crowdsourcing service that stores phishing data from users via URL submission and phishing verification. PhishTank determines whether a URL submitted by one user is phishing or not depending on the criteria of other users, and if the URL exceeds a specified PhishTank criterion, it is classified as a phishing site. In this comparative evaluation, we used only data of PhishTank labeled as phishing sites.

These data feeds are widely used in existing researches [30], [43], [57], [61] for the evaluation of phishing attacks as open threat intelligence that anyone can use. The two data feeds are explicitly indicated as providing information to APWG [62] and national CSIRTs; thus, when phishing URLs are published in the data feeds, the CSIRT in each country moves to handle takedowns of them. We continuously collected the latest data feeds from OpenPhish and PhishTank hourly during the same three-month period of Nov. 1, 2022 – Jan. 31, 2023. As a consequence, we collected 82,963 and 28,164 URLs from OpenPhish and PhishTank, respectively.

We evaluated the ratio of common phishing URLs and the latency of the same URLs across two data feeds using

Table 12 Comparative dataset of phishing URLs.



Datasets	#
CrowdCanary's Phishing URLs	34,595
OpenPhish's Phishing URLs [13]	82,963
PhishTank's Phishing URLs [12]	28,164
CrowdCanary \cap OpenPhish	11,589
CrowdCanary \cap PhishTank	9,213
OpenPhish \cap PhishTank	12,748
CrowdCanary \cap OpenPhish \cap PhishTank	4,620

**Fig. 5** Latency comparison of phishing URLs in CrowdCanary and OpenPhish.

CrowdCanary. The target time for evaluation is the posting time of the extracted user reports on Twitter for CrowdCanary, the `discover_time` in the data available in the API for OpenPhish, and the `verification_time` in the data available in the API for PhishTank.

Ratio of Common URLs. As shown in Table 12, CrowdCanary's phishing URLs and OpenPhish URLs have 11,589 URLs in common, CrowdCanary's phishing URLs and PhishTank URLs have 9,213 URLs in common, OpenPhish and PhishTank URLs have 12,748 URLs in common, and all three types of data have 4,620 URLs in common. We discovered that less than half of the phishing URLs in each dataset had anything in common, and that the observed targets differed greatly across them. Especially, among the phishing URLs extracted independently by CrowdCanary, only 13.4% of the total URLs were listed in both OpenPhish and PhishTank. This comparison revealed that the phishing URLs extracted by CrowdCanary had a large amount of unique information that was not present in other data feeds. In other words, reports of phishing attacks by users on Twitter are worth analyzing and extracting, and may be utilized as a new data feed for countermeasure techniques in the future.

Latency Comparison with OpenPhish. We compared and evaluated 11,589 phishing URLs common to CrowdCanary and OpenPhish. A summary of the results is shown in Fig. 5. The x-axis is the difference in latency in days, and the y-axis is the number of relevant phishing URLs. The blue bars represent the number of phishing URLs collected faster by CrowdCanary, whereas the orange bars represent the number of phishing URLs collected faster by OpenPhish. The number of phishing URLs that were collected faster by CrowdCanary was 9,132, which was 78.8% of the total common phishing URLs. From Fig. 5, most latency differences were less than one day. Because phishing sites have a short survival time [30], it is possible to improve existing countermeasure techniques with information from user reports, as the

**Fig. 6** Latency comparison of phishing URLs in CrowdCanary and PhishTank.**Fig. 7** Distribution of top level domains.

majority of common phishing URLs were collected earlier by CrowdCanary.

Latency Comparison with PhishTank. We compared and evaluated 9,213 phishing URLs common to CrowdCanary and PhishTank. A summary of the results is shown in Fig. 6. The x-axis is the difference in latency in days, whereas the y-axis is the number of relevant phishing URLs. The blue bars represent the number of phishing URLs collected faster by CrowdCanary, whereas the green bars represent the number of phishing URLs collected faster by PhishTank. The number of phishing URLs collected faster by CrowdCanary was 7,853, which was 85.3% of the total common phishing URLs. From Fig. 6, we found that Twitter users more often report information about phishing attacks earlier than security experts who use PhishTank routinely. These results mean more users can be protected from phishing attacks by using CrowdCanary information as a countermeasure than by referring to PhishTank to detect phishing attacks.

7.2 Analysis of Phishing Infrastructure

We analyzed the structure of website infrastructure commonly used by attackers for phishing sites by extracting URL and domain name information using CrowdCanary. We focused on what trends exist in the web resources (e.g., domain names and web servers) that attackers use to deploy phishing sites.

Distribution of Top Level Domains. We aggregated the top level domains (TLDs) of 34,595 phishing URLs detected by CrowdCanary. The top 20 most frequently occurring TLDs and the number of URLs are shown in Fig. 7. In total, we found 279 TLDs. As shown, the TLD "org", which was found in the largest number this time, is often used as a domain name for organizations and associations. In this case study, 8,591 URLs (98.5%) with a top level domain of "org" were found to be exploited by "duckdns.org", a dynamic




Fig. 8 Distribution of IP address locations.

DNS provider described in Sect. 6.2. To address this issue, it is essential to collaborate with the operators of these URLs. The top level domain “com” is the most commonly registered domain name and was found to be frequently abused by numerous phishing sites in this survey. The following TLDs, “top”, “cn”, “shop”, “cc”, “icu”, and “xyz”, were reported to have a high rate of abuse [44], and were shared on Twitter as many phishing attack reports as well. There is a clear trend in the TLDs that are exploited in phishing attacks, and this information is valuable in determining maliciousness.

Distribution of IP Addresses Locations. We analyzed the location information of IP addresses by querying the GeoIP2 database [63]. Out of 34,595 phishing URLs, we were able to analyze 31,385 URLs for which we were able to resolve names. Figure 8 shows the top 20 country codes that occur most frequently, along with their corresponding number of URLs. In total, there are 63 different country codes, and we found that 68.2% of IP addresses were located in the United States. Although CrowdCanary targeted phishing sites shared between English and Japanese languages for analysis, the top-ranked English-speaking countries were predominantly biased toward the United States. In other words, many phishing site web servers targeting Japanese people were also located in the United States.

Distribution of Hosting Providers. We analyzed the hosting providers that manage the aforementioned range of IP addresses. The top 20 most frequently occurring hosting providers with country codes and the number of URLs are shown in Fig. 9. Overall, there were 457 different hosting providers, and 15 of the top 20 were managed by United States organizations. Instead of exploiting only specific hosting providers, attackers are deploying phishing sites across a wide range of hosting providers. It is possible to reduce potential victimization by collaborating with these higher-ranking providers and guiding them towards taking down the malicious sites.

Distribution of Frequent IP Addresses and Hosting Providers. We analyzed the number of IP addresses linked to unique URLs and their relevance to hosting providers. We believe that if an attacker deploys a phishing attack using the same IP address, we can find similar attacks when we detect a single URL, even if the domain names are different. The hosting providers linked to the top 20 IP addresses and the number of corresponding URLs are shown in Fig. 10. We found that the top 20 IP addresses were managed by eight hosting providers. In particular, 1,791 and 1,400 phishing URLs were linked to a single IP address administered by




Fig. 9 Distribution of hosting providers.




Fig. 10 Distribution of frequent IP addresses and hosting providers.

“INTERNAP-BLK (US)” and “LG Uplus Corp (KR)”. The three IP addresses (orange bars) managed by “LG Uplus Corp (KR)” have a total of 2,333 phishing URLs linked to them, indicating that the operators’ anti-phishing site measures are insufficient. These results revealed that attackers deployed phishing sites using a variety of domain names, but with a bias toward specific IP addresses. As a criterion for determining whether a site is phishing or not, information such as IP addresses close in range to those already abused, or IP addresses managed by the same hosting provider, may be useful.

8. Discussion

We describes the potential for using CrowdCanary output information to defend against phishing attacks, the limitations of CrowdCanary, and ethical considerations of the experimental design.

8.1 Utilizing the Intelligence Collected for Phishing Attack Defense

We have demonstrated that CrowdCanary can collect threat intelligence on a large number of phishing attacks with greater accuracy than existing technologies. How can this collected intelligence be applied to actual defensive strategies? We believe that intelligence can be used from two main perspectives.

First, the phishing information collected can add to the intelligence in the block lists. It has been reported that the spread of phishing attacks does not end with the first wave of attacks; the second and third waves of attacks are sometimes spread using the same domain names [64]. By extracting information about the attack as early as possible, such as during the first wave, and feeding it into blocklists (e.g., email spam filters), it may be possible to protect users who may become victims of the second and third waves. It was also reported that among users who receive phishing emails, the average time difference between the timing of the first user to click on the URL and the last user to click on the URL is 21 hours [30]. By sharing information with the browser vendor's block list during this time difference, the browser can warn the user and protect them from phishing attacks if they visit the same URL.

Second, the characteristics of phishing attacks contained in the collected information can be analyzed and used as countermeasure information for similar attacks that may occur in the future. It has been reported that phishing sites change domain names frequently, but may continue to be hosted at a particular IP address [65]. For example, using passive DNS (e.g., Farsight DNSDB [66]), it is possible to detect attacks early using CrowdCanary intelligence if the A record of a newly appearing domain name is linked to the same IP address as a domain name that has been exploited for phishing attacks in the past. In addition, phishing sites created using phishing toolkits often have the identical HTML source, images on the site, and scripts [67]. This information can be useful in techniques such as content-based phishing site detection [68]. In addition, information about phishing attacks received by many users can be used to understand trends in company brands being exploited in attacks and to keep an eye on companies and industries that attackers will be targeting in the future.

8.2 Limitation

Our study has three limitations.

First, our study does not focus on extracting reports only from information about the final destination of phishing sites that involve user interaction or redirection. For example, some users may only share a screenshot on Twitter with the URL of the final destination after the entry or redirect occurs. CrowdCanary cannot properly extract reports in this case because it has no information about the user's input or redirection behavior on the browser. In particular,

CrowdCanary is a system that collects URLs that are the seeds of phishing attacks. CrowdCanary does not focus on phishing attacks that do not redirect without an acceptable referrer or can only be reached by clicking. These attacks can be handled by crawling URLs extracted by CrowdCanary as seeds in existing researches [55], [61].

Second, the features designed in this paper are chosen to be invariant with respect to user reporting of phishing attacks. However, the system's accuracy will inevitably decrease over time, and the system will need to be relearned each time. For example, when trends in the appearance of shared images change, or when phishing sites that look completely different become trending, it is necessary to re-annotate and relearn the classification model. Since we do not believe that the way users share information on Twitter itself will change significantly, and accuracy will not drop significantly immediately, evaluating how much accuracy will decrease as trends change is one of the issues we will address in the future.

Finally, depending on recent Twitter specification changes [69], equivalent information may no longer be available via the API. Since CrowdCanary is a system that extracts phishing attacks based on user reports, if the number of users using Twitter decreases (i.e., fewer users share phishing reports), the number of candidate phishing reports will decrease as tweets to be analyzed. Therefore, both the quantity and quality of threat information related to phishing attacks extracted by the proposed system will inevitably decline. In such cases, any social networking service that allows users to post photos and text, as popular as Twitter, can be used as a source of threat intelligence in the same way. In addition, CrowdCanary is adaptable to changes in Twitter because it was designed based on the characteristics of users sharing information about phishing attacks, rather than using Twitter-dependent features.

8.3 Ethical Consideration

We took into account the ethical considerations of collecting data from Twitter on a large scale. Although the collection and analysis targets contain massive amounts of information about Twitter accounts, the content of their tweets is public. In other words, since both expert and non-expert reports are shared with other users in public webspace for the purpose of alerting them, our experiment did not violate their intended use. Then, we believe there is no ethical issue because we did not take any actions that directly harmed users (e.g., actions on victims' email addresses or Twitter accounts).

We used common open source tools to collect data from Twitter at scale and send requests accordingly. We conducted the experiments according to the best practices of related research on Twitter's usage guidelines, minimizing the influence on the platform. In this experiment, we sent only 40 requests to Twitter (20 Security Keywords + 20 Co-occurrence keywords) per hour in English and Japanese. Therefore, we believe that the availability of the platform was unaffected.

9. Related Work

We describe the related research on identifying malicious tweets and generating threat intelligence from Twitter.

Identification of “Malicious” Tweets. Numerous studies [70]–[74] have analyzed phishing attacks that direct users to external malicious sites from Twitter. Gao et al. proposed a system that can detect malicious posts in real time using features common to Twitter and Facebook, such as user connections and the number of characters in a post [75]. *These studies analyze only malicious tweets (i.e., those distributed by attackers with malicious intent). However, our study extracts benign tweets (i.e., shared by users with good intentions), and the information in the benign tweets, such as URLs or domain names, is phishing information; thus, the analysis targets are completely different.*

Threat Intelligence Extraction from Twitter. Research on threat intelligence generation using Twitter information has been conducted from various perspectives [16]–[19], [76]. Shin et al. proposed a system to extract four types of information from a text on Twitter and external blogs: URLs, domain names, IP addresses, and hash values related to cyberattacks [18]. It has been demonstrated that the proposed system can detect threats, especially malware-related threats, earlier than other threat intelligence systems. Roy et al. focused on defanging and phishing attack-related hashtag strings, extracted information about phishing attacks from Twitter, and analyzed the characteristics of the accounts posting information [43]. It has been shown that information that interacts with other accounts, such as replies and retweets to the information posted on Twitter, is reflected more quickly in the block list. *Unlike our studies, tweets were collected from security experts by account names or limited keywords; thus, only limited information on Twitter was analyzed.*

10. Conclusion

This paper proposed CrowdCanary, a system that harvests phishing information from tweets of users who have discovered or encountered phishing attacks. The results suggest that reports from infrequent contributors (i.e., non-experts) contain a lot of valuable information for countering phishing attacks that is not included in the information posted by security experts. In addition, we identified tendencies in the domain names and hosting providers to which phishing sites were actually deployed, and indicated characteristics that are useful for detecting new phishing sites. Since this research showed the usefulness of information about new observation points on Twitter, we are ready to operate CrowdCanary in the future and to provide the data obtained to the national CSIRTs. We hope that the findings of this paper will be useful for future researches and countermeasure developments. We plan to share anonymized sample datasets with interested researchers upon request at <https://crowdc Canary.github.io/>.

References

- [1] M. Liu, Y. Zhang, B. Liu, Z. Li, H. Duan, and D. Sun, “Detecting and characterizing SMS spearphishing attacks,” Proc. 37th Annual Computer Security Applications Conference (ACSAC), pp.930–943, Association for Computing Machinery, 2021.
- [2] B. Reeves, N. Scaife, D. Tian, L. Blue, P. Traynor, and K.R.B. Butler, “Sending out an SMS: Characterizing the security of the SMS ecosystem with public gateways,” Proc. 37th IEEE Symposium on Security and Privacy (SP), pp.339–356, IEEE, Dec. 2016.
- [3] SafetyDetectives, “11 facts + stats on smishing (sms phishing) in 2021,” 2021. <https://www.safetydetectives.com/blog/what-is-smishing-sms-phishing-facts/>
- [4] B. Srinivasan, P. Gupta, M. Antonakakis, and M. Ahamad, “Understanding cross-channel abuse with SMS-spam support infrastructure attribution,” Proc. 21th European Symposium on Research in Computer Security (ESORICS), pp.3–26, Springer, Sept. 28–30, 2016.
- [5] “Smishing reports increase nearly 700% in first six months of this year,” 2021. <https://news.sky.com/story/smishing-reports-increase-nearly-700-in-first-six-months-of-this-year-12407504>
- [6] K. Thomas, F. Li, A. Zand, J. Barrett, J. Ranieri, L. Invernizzi, Y. Markov, O. Comanescu, V. Eranti, A. Moscicki, D. Margolis, V. Paxson, and E. Bursztein, “Data breaches, phishing, or malware? Understanding the risks of stolen credentials,” Proc. Conference on Computer and Communications Security (CCS), pp.1421–1434, Oct. 30–Nov. 3, 2017.
- [7] Y. Lin, R. Liu, D.M. Divakaran, J.Y. Ng, Q.Z. Chan, Y. Lu, Y. Si, F. Zhang, and J.S. Dong, “Phishpedia: A hybrid deep learning based approach to visually identify phishing webpages,” Proc. 30th USENIX Security Symposium (USENIX Security), pp.3793–3810, Aug. 11–13, 2021.
- [8] G. Ho, A. Cidon, L. Gavish, M. Schweighauser, V. Paxson, S. Savage, G.M. Voelker, and D.A. Wagner, “Detecting and characterizing lateral phishing at scale,” Proc. 28th USENIX Security Symposium (USENIX Security), pp.1273–1290, Aug. 14–16, 2019.
- [9] D. Kim, H. Cho, Y. Kwon, A. Doupé, S. Son, G. Ahn, and T. Dumitras, “Security analysis on practices of certificate authorities in the https phishing ecosystem,” Proc. 2021 ACM Asia Conference on Computer and Communications Security (ASIACCS), pp.407–420, June 2021.
- [10] Google, “Google safe browsing,” 2022. <https://safebrowsing.google.com/>
- [11] Microsoft, “Microsoft defender smartscreen,” 2022. <https://docs.microsoft.com/en-us/windows/security/threat-protection/microsoft-defender-smartscreen/microsoft-defender-smartscreen-overview>
- [12] PhishTank, “PhishTank,” 2022. <https://www.phishtank.com/>
- [13] OpenPhish, “OpenPhish,” 2022. <https://openphish.com>
- [14] VirusTotal, “VirusTotal,” 2022. <https://www.virustotal.com/>
- [15] SecurityTrails, “urlscan.io,” 2022. <https://urlscan.io/>
- [16] F. Alves, A. Andongabo, I. Gashi, P.M. Ferreira, and A. Bessani, “Follow the blue bird: A study on threat data published on twitter,” Proc. 25th European Symposium on Research in Computer Security (ESORICS), pp.217–236, Springer, Sept. 14–18, 2020.
- [17] C. Sabottke, O. Suciu, and T. Dumitras, “Vulnerability disclosure in the age of social media: Exploiting twitter for predicting real-world exploits,” Proc. 24th USENIX Security Symposium (USENIX Security), pp.1041–1056, USENIX Association, Aug. 12–14, 2015.
- [18] H. Shin, W. Shim, S. Kim, S. Lee, Y.G. Kang, and Y.H. Hwang, “#twiti: Social listening for threat intelligence,” Proc. Web Conference 2021 (WWW), pp.92–104, ACM, April 12–16, 2021.
- [19] H. Shin, W. Shim, J. Moon, J.W. Seo, S. Lee, and Y.H. Hwang, “Cybersecurity event detection with new and re-emerging words,” Proc. 15th Asia Conference on Computer and Communications Security (ASIACCS), pp.665–678, ACM, Oct. 5–9, 2020.
- [20] S. Tang, X. Mi, Y. Li, X. Wang, and K. Chen, “Clues in tweets: Twitter-guided discovery and analysis of SMS spam,” Proc. Confer-

- ence on Computer and Communications Security (CCS), pp.2751–2764, Nov. 7–11, 2022.
- [21] R. Saeki, L. Kitayama, J. Koga, M. Shimizu, and K. Oida, “Smishing strategy dynamics and evolving botnet activities in Japan,” *IEEE Access*, vol.10, pp.114869–114884, 2022.
- [22] NIST, “National vulnerability database,” 2021. <https://nvd.nist.gov/>
- [23] WeLiveSecurity, “Why do we fall for sms phishing scams so easily? | welivesecurity,” 2021. <https://www.welivesecurity.com/2021/01/22/why-do-we-fall-sms-phishing-scams-so-easily/>
- [24] T.I. Hunter, “Twitter ioc hunter,” 2022. <http://tweettioc.com/>
- [25] H. Nakano, D. Chiba, T. Koide, N. Fukushi, T. Yagi, T. Hariu, K. Yoshioka, and T. Matsumoto, “Canary in twitter mine: Collecting phishing reports from experts and non-experts,” *Proc. 18th International Conference on Availability, Reliability and Security, ARES 2023*, Benevento, Italy, 29 Aug. – 1 Sept., pp.6:1–6:12, ACM, 2023.
- [26] Statista, “Twitter: Most-used languages 2013 | statista,” 2023. <https://www.statista.com/statistics/267129/most-used-languages-on-twitter/>
- [27] Twitter, “Search api | twitter api | docs | twitter developer platform,” 2023. <https://developer.twitter.com/en/docs/twitter-api/enterprise/search-api/overview>
- [28] Twitter, “Compliance firehose api | twitter api | docs | twitter developer platform,” 2023. <https://developer.twitter.com/en/docs/twitter-api/enterprise/compliance-firehose-api/overview>
- [29] Twitter, “Decahose api | twitter api | docs | twitter developer platform,” 2023. <https://developer.twitter.com/en/docs/twitter-api/enterprise/decahose-api/overview/decahose>
- [30] A. Oest, P. Zhang, B. Wardman, E. Nunes, J. Burgis, A. Zand, K. Thomas, A. Doupé, and G.J. Ahn, “Sunrise to sunset: Analyzing the end-to-end life cycle and effectiveness of phishing attacks at scale,” *Proc. 29th USENIX Security Symposium (USENIX Security)*, pp.361–377, USENIX Association, Aug. 12–14, 2020.
- [31] A. Akbik, D. Blythe, and R. Vollgraf, “Contextual string embeddings for sequence labeling,” *Proc. 27th International Conference on Computational Linguistics (COLING)*, pp.1638–1649, Aug. 20–26, 2018.
- [32] M. Labs, “megagonlabs/transformers-ud-japanese-electra-base-ginza · hugging face,” 2021. <https://huggingface.co/megagonlabs/transformers-ud-japanese-electra-base-ginza>
- [33] G. Jocher et al., “ultralytics/yolov5: v3.1 - Bug Fixes and Performance Improvements,” 2020.
- [34] “Tesseract ocr,” 2022. <https://github.com/tesseract-ocr/tesseract>
- [35] “RFC 3986,” 2005. <https://datatracker.ietf.org/doc/html/rfc3986>
- [36] “RFC 1035,” 1987. <https://datatracker.ietf.org/doc/html/rfc1035>
- [37] V. Le Pochat, T. Van Goethem, S. Tajalizadehkhoob, M. Korczyński, and W. Joosen, “Tranco: A research-oriented top sites ranking hardened against manipulation,” *Proc. 26th Network and Distributed System Security Symposium (NDSS)*, The Internet Society, Feb. 24–27, 2019.
- [38] PeterDaveHello, “Url shorteners,” 2022. <https://github.com/PeterDaveHello/url-shorteners>
- [39] M. Tan and Q. Le, “EfficientNet: Rethinking model scaling for convolutional neural networks,” *Proc. 36th International Conference on Machine Learning (ICML)*, pp.6105–6114, PMLR, June 9–15, 2019.
- [40] J. Devlin, M.W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” *Proc. 2019 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, pp.4171–4186, Association for Computational Linguistics, June 3–5, 2019.
- [41] E. Keany, “BorutaShap: A wrapper feature selection method which combines the Boruta feature selection algorithm with Shapley values,” 2020. <https://zenodo.org/badge/latestdoi/255354538>
- [42] M.B. Kursu and W.R. Rudnicki, “Feature selection with the Boruta package,” *Journal of Statistical Software*, vol.36, no.11, p.1–13, 2010.
- [43] S.S. Roy, U. Karanjit, and S. Nilzadeh, “Evaluating the effectiveness of phishing reports on twitter,” *Proc. APWG Symposium on Electronic Crime Research (eCrime)*, Dec. 1–3, 2021.
- [44] “The top 10 most abused tlds,” 2022. <https://www.spamhaus.org/statistics/tlds/>
- [45] H. Alhichri, A.S. Alswayed, Y. Bazi, N. Ammour, and N.A. Alajlan, “Classification of remote sensing images using EfficientNet-B3 CNN model with attention,” *IEEE Access*, vol.9, pp.14078–14094, 2021.
- [46] G. Marques, D. Agarwal, and I. de la Torre Díez, “Automated medical diagnosis of COVID-19 through efficientnet convolutional neural network,” *Applied Soft Computing*, vol.96, 106691, 2020.
- [47] P.C. Hansen, “The truncatedSVD as a method for regularization,” *BIT Numerical Mathematics*, vol.27, pp.534–553, 1987.
- [48] A. Adhikari, A. Ram, R. Tang, and J. Lin, “DocBERT: BERT for document classification,” 2019.
- [49] F. Feng, Y. Yang, D. Cer, N. Arivazhagan, and W. Wang, “Language-agnostic BERT sentence embedding,” *Proc. 60th Annual Meeting of the Association for Computational Linguistics (ACL)*, Dublin, Ireland, pp.878–891, Association for Computational Linguistics, May 22–27, 2022.
- [50] R. Nogueira and K. Cho, “Passage re-ranking with BERT,” *CoRR*, abs/1901.04085, 2019.
- [51] vladkens, “Twitter api scrapper with authorization support,” 2023. <https://github.com/vladkens/twscrape>
- [52] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol.12, pp.2825–2830, 2011.
- [53] Spotify, “Luigi is a python module that helps you build complex pipelines of batch jobs.,” 2023. <https://github.com/spotify/luigi>
- [54] opmusic, “Spamhunter_dataset,” 2023. https://github.com/opmusic/SpamHunter_dataset/blob/main/sms_spam_urls/tweet_sms_url_latest.txt
- [55] T. Koide, D. Chiba, and M. Akiyama, “To get lost is to learn the way: Automatically collecting multi-step social engineering attacks on the web,” *Proc. 15th ACM Asia Conference on Computer and Communications Security (ASIACCS)*, pp.394–408, ACM, Oct. 5–9, 2020.
- [56] P. Peng, L. Yang, L. Song, and G. Wang, “Opening the blackbox of VirusTotal: Analyzing online phishing scan engines,” *Proc. Internet Measurement Conference (IMC)*, pp.478–485, Association for Computing Machinery, Oct. 21–23, 2019.
- [57] K. Tian, S.T.K. Jan, H. Hu, D. Yao, and G. Wang, “Needle in a haystack: Tracking down elite phishing domains in the wild,” *Proc. Internet Measurement Conference (IMC)*, pp.429–442, Association for Computing Machinery, Oct. 31–Nov. 2, 2018.
- [58] Z. Zhu and T. Dumitras, “ChainSmith: Automatically learning the semantics of malicious campaigns by mining threat intelligence reports,” *Proc. 3rd IEEE European Symposium on Security and Privacy (EuroSP)*, pp.458–472, IEEE, April 24–26, 2018.
- [59] dynamic.domains, “25 dynamic DNS (DDNS) providers - dynamic.domains,” 2022. <https://dynamic.domains/dynamic-dns/providers-list/default.aspx>
- [60] P. Zhang, A. Oest, H. Cho, Z. Sun, R.C. Johnson, B. Wardman, S. Sarker, A. Kapravelos, T. Bao, R. Wang, Y. Shoshitaishvili, A. Doupé, and G.-J. Ahn, “CrawlPhish: Large-scale analysis of client-side cloaking techniques in phishing,” *Proc. 42nd IEEE Symposium on Security and Privacy (SP)*, pp.1109–1124, IEEE, May 24–27, 2021.
- [61] T. Nelms, R. Perdisci, M. Antonakakis, and M. Ahamad, “WebWitness: Investigating, categorizing, and mitigating malware download paths,” *Proc. 24th USENIX Security Symposium (USENIX Security)*, pp.1025–1040, USENIX Association, Aug. 12–14, 2015.
- [62] “Unifying the global response to cybercrime,” 2022. <https://apwg.org/>
- [63] MaxMind, “Geoip2 databases,” 2023. <https://www.maxmind.com/en/geoip2-databases>

- [64] C. Yang, R. Harkreader, J. Zhang, S. Shin, and G. Gu, "Analyzing spammers' social networks for fun and profit: A case study of cyber criminal ecosystem on twitter," Proc. 21st International Conference on World Wide Web (WWW), pp.71–80, The International World Wide Web Conference Committee, April 16–20, 2012.
- [65] Q. Cui, G.-V. Jourdan, G.V. Bochmann, R. Couturier, and I.-V. Onut, "Tracking phishing attacks over time," Proc. 26th International Conference on World Wide Web (WWW), pp.667–676, The International World Wide Web Conference Committee, April 3–7, 2017.
- [66] Farsight Security Inc, "Dnsdb," <https://www.dnsdb.info/>, 2020.
- [67] H. Bijmans, T. Booij, A. Schwedersky, A. Nedgabat, and R. van Wegberg, "Catching phishers by their bait: Investigating the dutch phishing landscape through phishing kit detection," Proc. 30th USENIX Security Symposium (USENIX Security), pp.3757–3774, USENIX Association, Aug. 11–13, 2021.
- [68] G. Xiang, J. Hong, C.P. Rose, and L. Cranor, "CANTINA+: A feature-rich machine learning framework for detecting phishing web sites," ACM Trans. Information and System Security (TISSEC), vol.14, no.2, pp.1–28, 2011.
- [69] "Twitter dev," 2023. <https://twitter.com/TwitterDev/status/1615405842735714304>
- [70] S. Lee and J. Kim, "Warningbird: Detecting suspicious URLs in twitter stream," Proc. 19th Network and Distributed System Security Symposium (NDSS), The Internet Society, Feb. 5–8, 2012.
- [71] A. Aggarwal, A. Rajadesingan, and P. Kumaraguru, "PhishAri: Automatic realtime phishing detection on twitter," Proc. APWG Symposium on Electronic Crime Research (eCrime), Oct. 23–24, 2012.
- [72] S. Gupta, A. Khattar, A. Gogia, P. Kumaraguru, and T. Chakraborty, "Collective classification of spam campaigners on twitter: A hierarchical meta-path based approach," Proc. 27th International Conference on World Wide Web (WWW), pp.529–538, The International World Wide Web Conference Committee, April 23–27, 2018.
- [73] K. Thomas, C. Grier, D. Song, and V. Paxson, "Suspended accounts in retrospect: An analysis of Twitter spam," Proc. Internet Measurement Conference (IMC), pp.243–258, Association for Computing Machinery, Nov. 2–4, 2011.
- [74] H. Nakano, D. Chiba, T. Koide, and M. Akiyama, "Detecting event-synced navigation attacks across user-generated content platforms," Proc. IEEE 45th Annual Computers, Software, and Applications Conference (COMPSAC), pp.704–713, IEEE, July 12–16, 2021.
- [75] H. Gao, Y. Chen, K. Lee, D. Palsetia, and A.N. Choudhary, "Towards online spam filtering in social networks," Proc. 19th Network and Distributed System Security Symposium (NDSS), The Internet Society, Feb. 5–8, 2012.
- [76] R.P. Khandpur, T. Ji, S. Jan, G. Wang, C.-T. Lu, and N. Ramakrishnan, "Crowdsourcing cybersecurity: Cyber attack detection using social media," Proc. 2017 ACM on Conference on Information and Knowledge Management (CIKM), pp.1049–1057, Nov. 6–10, 2017.



Hiroki Nakano received his B.S. and M.S. degrees in computer science from Yokohama National University, Japan, in 2016 and 2018. Since joining Nippon Telegraph and Telephone Corporation (NTT) in 2018, he has been engaged in research and development on cybersecurity. He is currently an engineer at NTT Security (Japan) KK, Tokyo, Japan, and a Ph.D. candidate at Yokohama National University, Kanagawa, Japan. He is a member of IPSJ.



Daiki Chiba received the B.E., M.E., and Ph.D. degrees in computer science from Waseda University, in 2011, 2013, and 2017, respectively. Since 2013, he has been with Nippon Telegraph and Telephone Corporation (NTT), where he has been engaged in research on cyber security through data analysis. He is currently a Senior Manager at NTT Security (Japan) KK, Tokyo, Japan. He is a member of IEEE and IEICE.



Takashi Koide received the B.S., M.S., and Ph.D. degrees in Informatics from Yokohama National University in 2014, 2016, and 2021. He is currently a researcher at NTT Security (Japan) KK, Tokyo, Japan. His research interests include network and Web security. He won the Research Award from the IEICE Technical Committee on Information and Communication System Security in 2018 and IPSJ Outstanding Paper Award in 2022.



Naoki Fukushi received his B.E. and M.E. degrees in Computer Science from Waseda University in 2018 and 2020. His research interest covers Cyber Security. He is now with NTT Security (Japan) KK, Tokyo, Japan.




Takeshi Yagi received the B.E. degree in electrical and electronic engineering and the M.E. degree in science and technology from Chiba University, Japan, in 2000 and 2002, respectively, and the Ph.D. degree in information science and technology from Osaka University, Osaka, Japan, in 2013. He joined Nippon Telegraph and Telephone Corporation (NTT), in 2002. He is currently the Director of the Research and Development Planning Department, NTT. His research interests include cybersecurity monitoring and security intelligence. He is a member of the IEEE, IPSJ, and IEICE.



Takeo Hariu received his M.S. degree in electro-communications from the University of Electro-Communications. He is currently a Senior Vice President at NTT Security (Japan) KK, Tokyo, Japan.



Katsunari Yoshioka is a Professor at Yokohama National University since 2011. His research interests cover wide area of system security and network security including malware analysis and IoT security. He received the commendation for science and technology by the minister of MEXT, Japan in 2009, the award for contribution to Industry-Academia-Government Collaboration by the minister of MIC, Japan in 2016, and the Culture of Information Security Award in 2017.



Tsutomu Matsumoto is a professor of Faculty of Environment and Information Sciences, Yokohama National University. He is also the Director of Cyber Physical Security Research Center (CPSEC) at National Institute of Advanced Industrial Science and Technology (AIST). He received Doctor of Engineering from the University of Tokyo in 1986. Starting from Cryptography in the early 80's, he has opened up the field of security measuring for logical and physical security mechanisms. Currently he is interested in research and education of Embedded Security Systems such as IoT Devices, Cryptographic Hardware, In-vehicle Networks, Instrumentation and Control Security, Tamper Resistance, Biometrics, Artifact-metrics, and Countermeasure against Cyber-Physical Attacks.

He is serving as the chair of the Japanese National Body for ISO/TC68 (Financial Services) and the Cryptography Research and Evaluation Committees (CRYPTREC) and as an associate member of the Science Council of Japan (SCJ). He was a director of the International Association for Cryptologic Research (IACR) and the chair of the IEICE Technical Committees on Information Security, Biometrics, and Hardware Security. He received the IEICE Achievement Award, the DoCoMo Mobile Science Award, the Culture of Information Security Award, the MEXT Prize for Science and Technology, and the Fuji Sankei Business Eye Award.