It Never Rains but It Pours: Analyzing and Detecting Fake Removal Information Advertisement Sites

Takashi Koide[1,2(✉)], Daiki Chiba[1], Mitsuaki Akiyama[1], Katsunari Yoshioka[2], and Tsutomu Matsumoto[2]

[1] NTT Secure Platform Laboratories, Musashino, Japan
takashi.koide.fk@hco.ntt.co.jp
{daiki.chiba,akiyama}@ieee.org
[2] Yokohama National University, Yokohama, Japan
{yoshioka,tsutomu}@ynu.ac.jp

Abstract. Fake antivirus (AV) software is a serious threat on the Internet to make users install malware and expose their personal information. Fake removal information advertisement (FRAD) sites, which introduce fake removal information for cyber threats, have emerged as platforms for distributing fake AV software. Although FRAD sites seriously threaten users who have been suffering from cyber threats and need information for removing them, little attention has been given to revealing these sites. In this paper, we propose a system to automatically crawl the web and identify FRAD sites. To shed light on the pervasiveness of this type of attack, we performed a comprehensive analysis of both passively and actively collected data. Our system collected 2,913 FRAD sites in 31 languages, which have 73.5 million visits per month in total. We show that FRAD sites occupy search results when users search for cyber threats, thus preventing the users from obtaining the correct information.

Keywords: Fake AV software · Social engineering attacks

1 Introduction

Antivirus (AV) software is one of the basic defense strategies for protecting users' devices. The major AV software market was valued at 3,770 million USD in 2018 [12], and attackers focus on the needs of such pervasive AV software to gain financial benefits. Specifically, *fake AV software*, which are rogue applications disguised as legitimate AV software, is used to manipulate users' devices and steal money or sensitive information [2,18]. For example, once fake AV software is installed, the software displays fake virus scan results to get users to purchase additional licenses [4,23].

Fake AV software is a traditional cyber threat that can effectively spread malware and unwanted software on the web [11,22]. To infect users and gain more profit, attackers take advantage of online advertisements that target many people

C. Maurice et al. (Eds.): DIMVA 2020, LNCS 12223, pp. 171–191, 2020.
https://doi.org/10.1007/978-3-030-52683-2_9

to distribute fake AV software [26]. The web pages served by these advertisements typically show fake virus infection alerts or messages claiming the necessity of installing their software. These web pages also attract users with promises of speeding up their machines [24]. Attackers use such social engineering techniques that exploit users' psychological vulnerabilities to lure users to download fake AV software. These web pages are known to be major distribution paths for fake AV software [7,15,27].

In this paper, we focus on new techniques that psychologically encourage users to install fake AV software from the web. Attackers create web pages that introduce fake information for handling specific cyber threats, such as malware infection or visits to malicious web pages, and suggest fake AV software. We call these web pages *fake removal information advertisement (FRAD) sites*, which target users who have already suffered from security problems and which make them victims of another one. For example, users who notice their malware infection try to search for removal information using the malware detection names given by virus scanners, and they reach the FRAD sites from search results. Believing the FRAD information, the users follow the instructions and inadvertently install the suggested fake AV software. Although it is well known that attackers induce users to install fake AV software using scaring or attracting messages–such as fake infection alerts or promises to speed up their machines–little attention has been given to analyzing the FRAD sites.

Here, we propose a system that automatically crawls the web pages and detects FRAD sites. Using the linguistic and visual features of the web pages, we accurately identify FRAD sites with 98.8% true positives and only 3.3% false positives. We used our system for a large-scale collection of FRAD sites and found 2,913 distinct domain names of FRAD sites written in 31 languages. The total user accesses to these FRAD sites was 73.5 million visits per month. We observed that these FRAD sites are not adequately reported by existing blacklists.

To reveal the ecosystem of FRAD sites, we performed a measurement study using both passively collected statistical data on user accesses and actively crawled data. We first investigated the incoming traffic to FRAD sites to determine what types of user behaviors are at risk of reaching FRAD sites. We found that many users not only accessed these sites from search engines directly but also reached FRAD sites from videos or messages posted on social media by attackers' accounts. To determine what kinds of attacks users encounter from FRAD sites, we then analyzed the transferred web pages and downloaded files from the FRAD sites. We confirmed that the FRAD sites led to 76 fake AV software families by directly distributing installers and luring users to payment and distribution sites. Also, we investigated search results for the names of specific cyber threats, and we found that 82.6% of the top 10 search results were occupied by FRAD sites. In other words, search results for information concerning cyber threats are poisoned by FRAD sites, making it difficult for users to obtain correct removal information. To the best of our knowledge, this is the first study that has revealed the prevalence and ecosystem of FRAD sites.

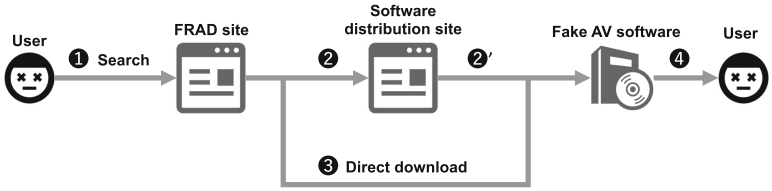

Fig. 1. Overview of fake AV software distribution via FRAD sites. Users that require removal information for cyber threats access FRAD sites via a web search (e.g., search engines or social media) (❶). They click on download buttons on the FRAD sites and are navigated to software distribution sites (❷). They download fake AV software from these sites (❷') or from the FRAD sites (❸) directly. Then, they make the damage even worse by installing the fake AV software (❹).

In summary, our contributions are as follows:

- We propose a system to crawl the web and detect FRAD sites automatically. By extracting linguistic and visual features from crawled web pages, our system detected FRAD sites with 98.8% true positives and 3.3% false positives.
- We performed a large-scale collection of FRAD sites on the web by leveraging a search engine, which is the most common channel used to reach FRAD sites. Using our system, we discovered 2,913 domain names of FRAD sites written in 31 languages. We found that attackers widely deploy FRAD sites targeting users in various countries to increase the number of page views.
- We conducted a comprehensive measurement study using both passively collected statistics data and actively crawled data to reveal the ecosystem of FRAD sites. Our measurement study also clarified the typical incoming channels employed by users to reach FRAD sites and the types of potential threats directed from the FRAD sites. We also found that it is difficult for users who need removal information for specific cyber threats to reach correct information, because most of the search results concerning cyber threats are poisoned by the FRAD sites.

2 Background

We first consider an attack technique for distributing fake AV software via FRAD sites. The purpose of the FRAD sites is to deceive users who need ways to deal with cyber threats, i.e., malicious acts that damage the users' devices and steal their sensitive information. Examples of cyber threats include malware infection, fraudulent popup messages, and malicious browser extensions. Attackers post multiple entries on FRAD sites that introduce fake threat removal guides, using the names of specific cyber threats, such as malware detection names or the domain names of malicious sites. For instance, there can be more than 15k entries in a single FRAD site, and dozens of new entries are added to the FRAD

site every day. When users notice that they have security issues by looking at the results from legitimate virus scanners or from suspicious alert messages on web pages, they search for information to remove them. Users who reach FRAD sites and are deceived by false information install fake AV software, which makes matters worse. We focus on such scams on the web in this paper.

Figure 1 shows an overview of the distribution of fake AV software via FRAD sites. First, users who have security problems reach FRAD sites by searching for the specific names of cyber threats they want to remove (❶). Attackers leverage search engine optimization (SEO) techniques that target specific names of cyber threats to increase the web traffic to FRAD sites. Attackers also post fake videos on YouTube that introduce ways to remove the threats, and they post similar articles on Facebook and other social media to lure users to click on links to FRAD sites. Forum and community sites where anyone can post messages are also used by the attackers in the same manner. Thus, users not only visit FRAD sites from results provided by search engines but also reach FRAD sites through social-media postings and other web pages hit by the search results. The FRAD sites contain detailed fake removal guides for individual threats as well as large buttons or banners to direct users to fake AV software. The FRAD sites usually display the logos of famous security vendors or third-party organizations (e.g., software certification companies) to make them look as if they are legitimate web pages. Users who click on the buttons or banners are navigated to software distribution sites (❷). Most of the software distribution sites use domain names containing the names of the fake AV software and disguise themselves as official sites for legitimate AV software by displaying product information and purchase menus. These sites are also reachable through search engines and even provide customer support such as web chats or toll-free calls. On these web pages, users follow the payment and download instructions and then obtain fake AV software installers (❷'). These installers can also be downloaded from the FRAD sites directly (❸). Users install the fake AV software and thus become victims of other cyber threats (❹).

Some social engineering techniques are already known, such as threatening users using fake infection alerts or attracting them by the prospect of improving computer performance. However, it has not been clarified whether attackers use techniques for distributing fake AV software that exploit the weaknesses of users who have already suffered from cyber threats.

3 Method

In this section, we introduce our system for collecting and detecting FRAD sites on the Internet automatically. The system consists of two steps: web crawling and classification.

3.1 Web Crawling

The implementation of a web crawler that collects and stores browser-level information from web pages is the first step in our system. The requirement of the

Table 1. List of terms for each category; used to check the term's frequency in the title, URL paths, domain names, and text content of a web page.

Category	Example terms
way	"how to", "guide", "solution", "tips", "report", "instruction"
removal	"remove", "get rid of", "uninstall", "delete", "fix", "clean", "kill", "block", "repair", "anti", "entfernen", "eliminar", "verwijderen", "deinstallieren", "desinstalar", "supprimer", "remuovere", "usunac"
problem	"virus", "malware", "spyware", "trojan", "backdoor", "adware", "threat", "infection", "ransom", "error", "pop up", "redirect"
device	"computer", "pc", "windows", "mac", "browser"

crawler is to extract linguistic and image features from a web page rendered by a web browser and to compose a feature vector for the result. To analyze the FRAD sites in detail, we also need to capture the network traffic to and perform browser interactions on the web page. To achieve this, we designed and implemented the crawler using Scrapy[1], which is a web crawling framework for Python, in order to develop functions for monitoring and managing logged data. We used Selenium[2] as the middleware for Scrapy to automate a real web browser. We used Google Chrome as the default web browser for the crawler. To monitor network traffic in detail, we used Chrome DevTools API[3]. This is necessary, because we collect network-level information such as HTTP requests and responses that Selenium API does not handle directly. The collected information–such as screenshots, HTML source codes, and network traffic–are stored to MongoDB. We use those kinds of information for the next step, classification.

3.2 Classification

In the second step, our system extracts features from the information collected from the web pages and identifies FRAD sites using a supervised machine learning approach. In particular, the system analyzes term frequencies in web pages and URLs, the presence of logo images on screenshots, and HTML structures, such as the number of tags, and combines them into a feature vector. We explain the detail of each feature below.

Term Frequencies. To capture the linguistic characteristics of FRAD sites, frequencies of terms are used as a feature. To improve the SEO ranking and ensure an easy web page topic for users to understand, FRAD sites use terms meaning for the removal of cyber threats in the titles, URL paths, domain names, and text content of their web pages. Examples

[1] https://scrapy.org/.
[2] https://selenium.dev/.
[3] https://developer.chrome.com/extensions/devtools.

of such titles are "*Remove Trojan.Zerocleare (Virus Removal Guide)*" and "*Remove Magiballs.com (Free Guide).*" The URL paths include forms such as "`/2019/12/27/how-to-remove-my-login-hub-virus-removal-guide/`" and "`/uninstall-nvux-xyz-from-windows-7-8-8-1-10`." Examples of domain names are `uninstallmalwarefrompc[.]example` and `virusremovalguide[.]example`. The text content of the web page is written with a summary of the cyber threat and specific removal information for it.

Our key insight is that the FRAD sites must include a phrase composed of the following four categories of terms: *way*, *removal*, *problem*, and *device*. Table 1 shows a list of example terms. As the feature vector, we use the number of occurrences of each term category in the following four fields: the title, URL path, domain name, and text content. The terms in the four categories are intended to capture phrases such as "*how to remove Trojan.Zerocleare virus from my PC.*" Because the FRAD sites are created in many languages, we leverage machine translation services such as Cloud Translation API[4] and Amazon Translate[5]. We translate the title and text content of the crawled web pages into English and then calculate the frequencies of the terms.

To create the list of terms, we extracted all terms that match each category from the title, URL paths, domain names, and text content of 300 FRAD sites that were randomly selected from our created dataset, as discussed below in Sect. 4. Some domain names include non-English terms in the *removal* category, such as "entfernen" in German and "eliminar" in Spanish. Because these domain names are difficult to translate, we manually obtained such terms as much as possible. To this end, we separated the domain names by "." or "-" and used word segmentation[6] and then searched for the meaning of each extracted word.

Logo Images. We next consider features that specify logo images on the FRAD sites. The FRAD sites include download buttons and software packages that may be shared among multiple FRAD sites. The FRAD sites also display logos of security vendors, operating system (OS) vendors or software certification companies in order to pretend to be legitimate sites. These logos are copied from vendors' sites or used as image files modified from the original images. To find such visual characteristics, our system uses an image matching approach on the basis of our logo image database. Specifically, the system extracts images from `img` tags and crops images for which the area matches `a` or `button` tag elements from screenshots. It calculates the perceptual hash[7] of these images and compares them to the image database. If the target image is more than 85% similar to the image in the database, the system determines it to be a logo image. Three types of images are stored in the database: logos of security vendors or software certification company (19 images), package images of fake AV software (33 images), and images of the download buttons (56 images). We extracted images

[4] https://cloud.google.com/translate/.
[5] https://aws.amazon.com/translate/.
[6] http://www.grantjenks.com/docs/wordsegment/.
[7] https://github.com/JohannesBuchner/imagehash.

belonging to the three types from the 300 FRAD sites used in the above. Our system counts the number of images that match each type to create feature vectors.

HTML Structure. Here, we explain the features extracted from the HTML structure that we use for identifying FRAD sites. As with previous works that identify specific types of malicious web pages [7,17], the numbers of `a` and `iframe` tags are important indicators of FRAD sites. Also, FRAD sites often re-use web page templates so that they have similar structures of HTML source codes. In other words, the frequency of HTML tags and combinations of those numbers characterize FRAD sites. To find such features, the system counts the number of appearances of HTML tags. The HTML tags to be counted are the top 30 tags frequently used in the 300 FRAD sites mentioned above.

4 Data Collection

We explain the method used to collect FRAD sites in the wild in order to make the dataset employed to evaluate our classification model. We first collected the names of cyber threats. Then, we searched for and gathered candidates of FRAD sites using the names of those cyber threats. Finally, we manually created a labeled dataset for our evaluation experiment.

4.1 Collecting Cyber Threats

We collected the names of cyber threats to make search queries to find candidate FRAD sites. As described in Sect. 2, FRAD sites prepare many entries that introduce ways of removing specific cyber threats such as malware detection names and malicious domain names. To collect such names efficiently, we crawled the database pages of security vendors (e.g., Symantec Security Center[8]) and a security community site (e.g., malwaretips[.]com) in October 2019. We collected 806 names of threats, including 500 malware detection names, 200 malicious domain names, and 106 popup messages.

4.2 Web Search

We created search queries using the collected names of cyber threats and gathered the URLs of web pages using a search engine. To collect FRAD sites efficiently, we added "how to remove" to the name of the cyber threat to create the search query, instead of searching only for the name of the threat. We found that we can collect more FRAD sites by searching with "how to remove" in our experiment described in Sect. 6.3. To collect search results systematically, we used Microsoft Bing Web Search API[9] and gathered 34k URLs. We chose one URL for each domain name from among the gathered URLs. As a result, we extracted 4,188 URLs with 4,188 unique domain names to crawl.

[8] https://www.symantec.com/security-center/a-z.
[9] https://azure.microsoft.com/en-us/services/cognitive-services/.

4.3 Creating the Dataset

We crawled 4,188 web pages using our system and created a labeled dataset. Since there is no existing URL blacklist that accurately identifies FRAD sites, we manually labeled them by analyzing the crawled web pages and actually accessed them as necessary. To efficiently conduct this process, we created a web application that displays screenshots and buttons to choose labels (FRAD and non-FRAD sites). This application extracts information about the crawled web pages from our MongoDB database and generates the web pages for labeling. We implemented it using Node.js and the Express[10] framework. We labeled web pages as FRAD sites if they satisfied following heuristic rules. If not, we labeled the web pages as non-FRAD sites.

i. We check whether a web page introduces a removal guide for a specific cyber threat. If so, we check rule ii.
ii. We check whether the web page has visual characteristics specific to FRAD sites, as described in Sect. 3.2. Specifically, we check whether the web page has an image of a fake AV software package or a logo of a security vendor or a software certification company. We also check screenshots of the removal instructions or download buttons, which are often shared with multiple FRAD sites. If the web page has these characteristics, we identify it as an FRAD site. If not, we further check rule iii.
iii. We confirm that clicking a download button on the web page triggers a download of a fake AV software installer or initiates a web transition to a distribution or payment site for fake AV software. We performed this process by manually accessing the web page and clicking the download button.

From the 15-h labeling process, we obtained 804 web pages of FRAD sites with 804 unique domain names. To create a dataset, we randomly selected 800 web pages from these FRAD sites. We also randomly selected 800 web pages from non-FRAD sites, which are the web pages remaining after excluding the 804 web pages of FRAD sites. Since we collected the non-FRAD sites using the same search queries as for the FRAD sites, they often introduce removal information for cyber threats, details of malware, or introductions to legitimate AV software, just as FRAD sites do. Thus, it is a challenging task to identify FRAD sites accurately from these similar web pages.

5 Evaluation

We next evaluated the detection capability of our system in terms of its capability to classify web pages accurately as FRAD sites or non-FRAD sites. We also conducted an experiment to discover unknown FRAD sites in the wild using the trained classification model.

[10] https://expressjs.com/.

5.1 Detection Accuracy

We first evaluated the detection accuracy of our system using the balanced dataset including 800 FRAD sites and 800 non-FRAD sites. We used a random forest classifier as the machine learning algorithm for two-class classification, because we can easily tune it due to the small number of hyper parameters to be considered. We conducted a 10-fold cross validation to determine how accurately our system performed classifications. We found that our system classified web pages with a 98.8% true positive (TP) rate ($= \frac{TP}{TP+FN}$), where FN = false negative, a 3.3% false positive (FP) rate ($= \frac{FP}{FP+TN}$), and with 96.8% precision ($= \frac{TP}{TP+FP}$). The system identified 26 non-FRAD sites as FRAD sites (FPs). Examples include articles from security vendors that introduce malware information, ranking web pages for legitimate AV software, and blog entries that describe correct removal instructions. Five FPs were security vendors' web pages that often appear in search results when searching for removal information for cyber threats. We can therefore reduce FPs by placing the domain names of major security vendors on a whitelist. Examples of false negatives include web pages with domain names that do not include words such as "remove" or "malware." Other false negatives do not contain visual features such as images of fake AV software packages or logos of security vendors.

5.2 Detecting Unknown FRAD Sites

To collect unknown FRAD sites that have not been found in Sect. 5.1, we conducted additional data collection and detection using our classification model, which has high detection accuracy.

Additional Data Collection. We first describe additional data collection to find more FRAD sites in the wild, such as non-English FRAD sites and FRAD sites with content copied from other sites. In the process of creating the dataset described in Sect. 4, we found many FRAD sites written in various languages. Some of them were translated automatically according to the browser's language setting when the web pages were loaded. Some web pages were also written in multiple languages to enable users to switch languages. In addition, we found FRAD sites dedicated to certain languages. In such cases, the domain names contain words in those languages (e.g., "entfernen" in entfernen-spyware[.]example and "eliminar" in eliminarvirus[.]example), as described in Sect. 3.2. We also found that FRAD sites are often copied from other FRAD sites and from legitimate sites that introduce specific malware removal information. These FRAD sites not only use the names of cyber threats extracted from legitimate sites but also copy page titles or entire articles from them. To find such FRAD sites, we collected page titles from legitimate sites (malwaretips[.]com and malwarefixes[.]com) and from the 804 FRAD sites we labeled, which include non-English sites, and we searched for the titles using Bing API. Although it is difficult to create search queries in multiple languages to collect non-English FRAD sites,

we can gather them efficiently in this way. We gathered 16k page titles from these web pages and collected 836,731 URLs (111,161 domain names) from these search. We extracted up to three URLs from each domain name and crawled them (120,577 URLs) using our system.

Detection Result. As a result of the classification of additionally crawled web pages, we identified 6,130 URLs as FRAD sites. To find FPs, we manually checked web pages classified as positive in the same way as described in Sect. 4.3. Examples of FPs include the following. Some technical-support scam [14,21] sites were falsely identified as FRAD sites, because they offered support for malware removal and displayed noticeable phone numbers and web-chat support. These FPs are not FRAD sites, however, because they did not lead users to fake AV software but instead are actually malicious web pages themselves, which are listed in VirusTotal[11]. Moreover, our system falsely detected pirate web pages that introduce free downloads of fake AV software. Although such fake AV software is useless and not very well-known, some web pages illegally offered such software. Other FPs include software review and download sites, which distribute fake AV software as well as legitimate software. We also found FPs similar to those described in Sect. 5.1. By excluding these FPs, we finally determined 5,780 URLs (2,109 domain names) as FRAD sites. The precision of this classification result was 94.3%. Although this precision is somewhat less than the results obtained in Sect. 5.1, we accurately identified FRAD sites. The reason for this decrease in detection capability is that we changed the search queries from "how to remove" and the name of threats (used in Sect. 4.2) to page titles of known FRAD sites, so that the types of web pages in the search results were somewhat changed.

Summary of Collected FRAD Sites. Overall, in this paper we have identified 2,913 domain names, including the newly discovered 2,109 domain names, to be FRAD sites. To confirm the FRAD sites already reported by security vendors, we searched for all 2,913 domain names in VirusTotal. Of the total, 32.7% (952 domain names) of the domain names had URLs that had already been detected by one or more vendors. We also found 21.5% (626/2,913) of the domain names had URLs that are sources of detected files. Although some FRAD sites have been detected by a small number of security vendors, most of the FRAD sites we found in this paper have been unreported to date. These FRAD sites are less likely to be filtered out from search results, even if they were reported as malicious. Thus, most of these FRAD sites remain easily accessible to users and remain threatening to them.

[11] https://www.virustotal.com.

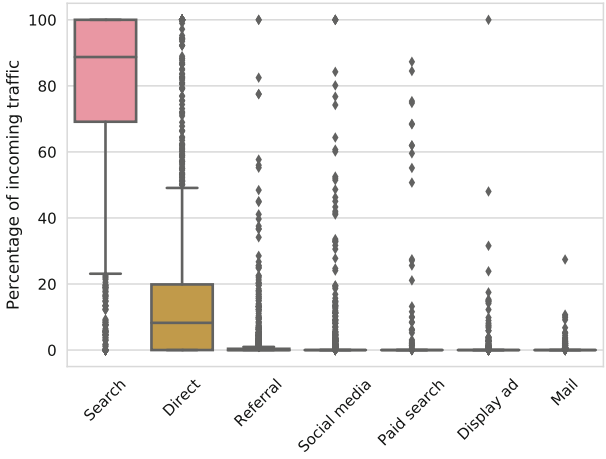

Fig. 2. Percentage of incoming traffic to FRAD sites from each channel.

6 Measurement Study

We measured the ecosystem and risk of FRAD sites using both passively collected statistical data of user accesses and actively crawled data. In the experiment described above, we found FRAD sites using our system and simply checked the detection status for each of them on VirusTotal. Here, we analyze deeply the 2,913 domain names of FRAD sites that we found in Sect. 5 in terms of incoming traffic to those FRAD sites, the distribution of fake AV software from those sites, and poisoned search results that are occupied by FRAD sites.

6.1 Incoming Traffic to FRAD Sites

To find out what browsing behaviors of users are at risk of reaching FRAD sites, we analyzed the incoming channels (i.e., ❶ in Fig. 1 in Sect. 2) of the FRAD sites that we found in Sect. 5. To this end, we need data on the history of user accesses to and traffic volumes of those web pages. Thus, we leveraged the statistical data provided by SimilarWeb[12], which passively observes hundreds of millions of global devices and covers over 220 countries and territories. Using this approach, we collected statistical data from October to December in 2019 that we used in the measurement studies described below.

Overview of Incoming Traffic. We first show an overview of seven types of incoming traffic to FRAD sites. We investigated 1,451 domain names of FRAD sites for which data are available in SimilarWeb (out of 2,913 domain names of the FRAD sites we discovered in this paper). Note that statistical data of web pages with few user accesses are not provided. These FRAD sites have

[12] https://www.similarweb.com/.

Table 2. Search queries used by the users to reach FRAD sites.

Category	Search query	#
Cyber threats	how to <remove><threat>	576
	<remove><threat>	438
	<threat>	849
	is <threat> safe ?	27
	what is <threat>	113
	<error>	140
Download	download <software>	421
	crack <software>	101
Fake AV software	<fake AV software>	66
Other	<other>	1,802
Total		4,510

73.5 million visits per month in total. Figure 2 shows the percentage of traffic to the FRAD sites from each incoming channel. The channels consist of seven labels: *Search* (accessed from a search engine), *Direct* (directly accessed by entering URLs in a web browsers), *Referral* (accessed from other web pages), *Social media* (accessed from Social Media), *Paid search* (accessed from keyword advertisements on search engines), *Display ad* (accessed from advertisements on web pages), and *Mail* (accessed from hyperlinks on email). Note that the incoming traffic measured as Mail comes only from web mail. Incoming traffic from email client software or other applications is measured as Direct. The mean values of Search, Direct, Referral, and Social media were 76.7%, 16.5%, 1.7%, and 1.7%, respectively. The value for each of the other three channels is less than 0.6%. Paid search, Display ad, and Mail have few data for further investigation. Also, we only know the amount of incoming traffic that we have shown here from the data of Direct. Therefore, in the following, we analyzed the detail of three channels: Search, Referral, and Social media.

Search. To find out how users reached FRAD sites via search engines, we investigated the statistics of the search queries. We extracted the top 10 English search queries (4,510 unique queries in total) for each FRAD site and categorized them. Table 2 shows the categories and the number of search queries. We found that 47.5% (2,143/4,510) of the search queries were related to the names of specific cyber threats. They included malware detection names (e.g., trojan:win32/bearfoos.a!ml), malicious domain names, and alert dialog messages (e.g., "your computer is infected with dangerous viruses"). Among them, 12.8% (576/4,510) are search queries combining "how to" with words meaning removal (e.g., "remove", "delete") and the names of cyber threats. We found that 9.7% (438/4,510) of the search queries combined words meaning removal with the names of cyber threats. Users also searched for the names of cyber threats alone

Table 3. Top 10 social media that led to FRAD sites.

Social media	# of FRAD sites
Youtube	160
Facebook	111
Reddit	58
Quora	35
Pinterest	22
Pocket	9
Twitter	7
Linkedin	6
Instagram	5

Table 4. Top 10 categories of referral web pages to FRAD sites.

Category of referral web pages	#
Computers electronics and technology	517
Games	29
News and media	25
Science and education	22
Business and consumer services	20
Arts and entertainment	19
Hobbies and leisure	8
Adult	8
Reference materials	7
E-commerce and shopping	6

(18.8%, 849/4,510) of for software or OS error messages (e.g., "MSVCP140.dll missing"). Thus, many users reach FRAD sites by searching for cyber threats and corresponding removal guides. The names of fake AV software were also used as search queries to reach FRAD sites (66/1,802). We found that 11.6% (522/4,510) of the search queries were used to search for downloads of software such as office software or video games and guides of cracking them. Forty percent (1,802/4,510) of the search queries were not included in these categories.

Social Media. We also analyzed incoming traffic from social media. We investigated 167 FRAD sites for which statistical data for queries incoming from social media is available from SimilarWeb. Table 3 shows the top 10 social media that led users to FRAD sites and the number of FRAD sites to which users were redirected from each type of social media. Users visited 95.8% (160/167) of FRAD sites from YouTube and 66.5% (111/167) of those from Facebook. Attackers create social-media accounts for these FRAD sites and post videos or messages to lure users to FRAD sites. These accounts pretended to be official accounts that use the web-site names or domain names of FRAD sites. They introduce removal information for cyber threats in the same way as entries for FRAD sites, and they put hyperlinks leading to FRAD sites in the description of their videos and messages. We found that some accounts post such instruction videos on YouTube several times a day. These videos got as many as 700k views. We also found that attackers created such accounts across multiple social media. In summary, attackers not only optimize search results to lead users directly to FRAD sites, but also they use various social media to increase user accesses to FRAD sites.

Referrals. In addition, we investigated referral traffic that leads users to FRAD sites. In other words, we analyzed the incoming traffic to FRAD sites when users

accessed them from other web pages, excluding search engines and social media. We found that users visited 891 web pages belonging to various categories before reaching FRAD sites. Table 4 shows the top 10 SimilarWeb categories of these referral web pages. The most common category of referral web page is Computers Electronics and Technology, which includes forum and community sites such as social.technet.microsoft[.]com., ubuntuforums[.]org, and discussions.apple[.]com. In most cases, attackers abuse these sites, where anyone can post messages, to impersonate good users who introduce removal information for cyber threats with URLs of FRAD sites. The web pages categorized as Games (e.g., steamcommunity[.]com) were used in the same manner. Attackers also posted FRAD sites' URLs in comment sections in articles in News and Media and other categories. In short, attackers leverage popular web pages where they can post comments and hyperlinks to lure users to visit FRAD sites.

6.2 Downloads and Page Transitions from FRAD Sites

To identify threats that occur when users access FRAD sites, we performed an additional crawling experiment. While we simply found FRAD sites using our system in Sect. 5, and we investigated users' incoming traffic to them in Sect. 6.1, the malicious activity derived from them was not revealed by these experiments. Therefore, we actively crawled the FRAD sites and collected installers of fake AV software and their respective distribution sites. To this end, we added a function to the crawler of our system to enable it to detect a download button on an FRAD site and click it. Then we analyzed the downloaded files and transferred the web pages from those FRAD sites.

Collecting File Downloads and Web-Page Transitions. We first describe the details of the new function that enables our crawler to interact with the FRAD sites. The crawler crops images with areas that match the `a` tag and `img` tag elements of FRAD sites. If the crawler finds a "download" string in the images using optical character recognition, it clicks on that area. We used two types of UserAgent with different OS (Windows 10 and macOS v10.14). This is because FRAD sites change the fake AV software to be distributed according to the UserAgent's OS, typically Windows or Mac. To collect the URLs of FRAD sites to crawl, we searched for the 2,913 domain names of FRAD sites using Bing API and selected up to three URLs based on the search results for each domain name. The reason for this is that web pages of FRAD sites with the same domain names can lead to different destinations (e.g., different software distribution sites) depending upon their URLs. To find more fake AV software, we collected 8,099 URLs and crawled them twice with two types of UserAgent. As a result, the crawler downloaded 4,548 files with 594 unique MD5 hash values and reached 136 domain names (630 URLs) of web pages from FRAD sites. In the following, we investigated the downloads of fake AV software originating from the FRAD sites (i.e., ❸ in Fig. 1 in Sect. 2), web pages transferred from those sites (i.e., ❷ in Fig. 1), and redirectors that relayed these downloads and web page transitions.

Fake AV Software Downloaded from FRAD Sites. We analyzed the files that our crawler downloaded (see ❸ in Fig. 1) to identify the installers of fake AV software. First, we checked 594 files with unique MD5 hash values on VirusTotal and found that 89 of those files had been detected. To specify fake AV software families from the detected files, we manually analyzed and searched them using their filenames and metadata (e.g., product name, legal copyright, and file description) read by ExifTool[13]. We examined whether the 89 files were related to malware removal, registry fix, or speed up based on the above information and on the software distribution sites that we obtained from the search results. We classified 84 files into 58 unique fake AV software families with different software names. All 58 fake AV software families have software distribution sites reachable from search engines. The software distribution sites profess to be official sites for these fake AV software families. For example, these sites show download and purchase menus and provide customer support such as web chats or toll-free calls. The remaining five detected files were not fake AV software but instead were malware that pretend to be installers of legitimate software, such as music-production software and video games.

To find more fake AV software from the 505 undetected files, we compared their filenames and metadata with those of the classified 58 fake AV software families. As a result of determining files with the same strings as the fake AV software, we additionally found 189 files to be fake AV software. Overall, we found 278 files (31 `dmg` files and 247 `exe` files) of the 58 fake AV software families.

Web Pages Transferred from FRAD Sites. We also analyzed the web pages of 136 domain names that our crawler reached after clicking on download buttons (see ❷ in Fig. 1). In the above measurements, we investigated fake AV software directly downloaded from FRAD sites. However, FRAD sites also navigate users to software distribution sites that lure them to purchase and download fake AV software. To find such web pages, we analyzed the crawled data (e.g., screenshots of web pages) and manually classified the malicious web pages. We first checked the 136 domain names on VirusTotal and found that 57 domain names were detected. We then specified the web pages that offered license purchases of known fake AV software or were related to malware removal, registry fixes, and speed-up from the web pages of the 57 detected domain names. We found that 34 domain names were related to distributions of fake AV software, including six domain names of payment sites and 27 domain names of software distribution sites. The payment sites required inputting credit card numbers and personal information to purchase fake AV software. Out of the 27 domain names, we found that 18 domain names were distribution sites for 18 new fake AV software families in addition to the measurements described above, where we found 58 fake AV software families. Thus, we found 76 fake AV software families in total. The detected domain names also included five domain names of FRAD sites that we found in Sect. 5. That is, users may be transferred from one FRAD site to another. We also found malicious web pages that distribute malicious Chrome

[13] https://exiftool.org/.

Table 5. The percentage of FRAD sites included in search results.

	Threat name <threat name>	remove <threat name>	how to remove <threat name>
Malware	69.4%	87.9%	87.9%
Domain name	88.5%	93.5%	88.0%
Extension	36.1%	85.1%	87.2%
Total	70.6%	89.7%	87.8%

extensions. We found 14 domain names associated with such threats and four domain names related to distributions of other types of malware.

Redirectors. To reveal the network infrastructure related to the distribution of fake AV software, we investigated the redirectors that relayed the above fake AV software downloads and web page transitions. We analyzed the network traffic that our crawler captured and extracted redirectors for which the effective second-level domains (e2LD; e.g., example.com is a e2LD of www.example.com) are different from those of the source web pages (i.e., the FRAD sites) and destination web pages. We found 169 domain names (38 e2LD names) as redirectors of 1,048 URL redirections associated with fake AV software downloads and web transitions to software distribution sites. Nine of these domain names were known advertising domain names listed in EasyList[14]. In addition, we found a small number of redirectors that were involved in many fake AV software distributions. For example, we found that 76.4% of the URL redirections were associated with just two domain names: safecart[.]com and revenuewire[.]net. These two redirectors navigated to 17 and 14 fake AV software families, respectively. The domain name safecart[.]com not only is a redirector but also is a payment web page that prompts users for their credit card numbers. Some redirectors, such as reimageplus[.]com and paretologic[.]com, which are software distribution sites, navigated to other software distribution sites.

6.3 Search Poisoning

We conducted a further measurement experiment to analyze the percentage of FRAD sites in the search results. In Sect. 6.1, we used statistical data to investigate search queries that users used to reach FRAD sites. Then, we determined the risk of users reaching these FRAD sites by actually searching with those search queries and analyzing the search results. When users search for specific names of cyber threats to find removal information, many FRAD sites prominently show up in search results. To confirm these poisoned search results, we investigated 150 search queries, combining 50 cyber threats and three search patterns. The three search patterns are those that users frequently use, as found

[14] https://easylist.to/.

in the measurements in Sect. 6.1: "how to remove" and the name of a cyber threat, "remove" and the name of a cyber threat, and only the name of a cyber threat. We extracted the latest names of cyber threats from public lists: 20 malware detection names from Symantec Security Center and 20 malicious domain names from malwaretips[.]com. Also, we randomly chose 10 malicious browser extensions out of 14 browser extensions that we found in Sect. 6.2. We investigated the top 10 search results for each search query, which are the top result pages from popular search engines such as Google and Bing.

We collected 1,461 web pages from the top 10 search results for each of the 150 search queries in total. By matching the 2,913 domain names of the FRAD sites collected in Sect. 5.2, we found that 1,207 web pages (82.6%) were FRAD sites. Table 5 shows the percentages of FRAD sites included in the search results for each search query and the names of the cyber threats. When we searched for the names of cyber threats with "how to remove" or "remove," the percentages of FRAD sites were 87.8% and 89.7%, respectively. The FRAD sites were also included at a high rate in the results of searching only for the names of cyber threats. In particular, 88.5% of search results for the domain names were FRAD sites. Search results for malicious browser extensions did not include many FRAD sites (36.1%), but there was less useful information available for users to use to remove the threats or determine whether they are malicious. We also found 22 YouTube web pages as search results, with videos and descriptions that introduced FRAD sites. We found that 26.7% (40/150) of the search queries returned search results for which the top 10 web pages were all FRAD sites. In summary, we found that most of the search results were occupied by FRAD sites when users searched for removal information for cyber threats, making it difficult for users to reach correct information.

7 Discussion

Ethical Considerations. We followed research ethics principles and best practices to conduct this study [3]. We analyzed users' behavior to visit FRAD sites using anonymized statistical data on user accesses for this study. We purchased a license to access data that is legally collected based on SimilarWeb's privacy policy. The information extracted from the web pages we crawled is publicly available data. To reduce server load, our experiment that interacted with download buttons was performed only once for each web page that we identified as an FRAD site.

Limitation. Although our system can accurately identify FRAD sites, there are some limitations. Since our system is specialized for collecting and detecting FRAD sites, which are the important platforms used by attackers to distribute fake AV software, detecting software distribution sites is out of scope for this paper. We identified software distribution sites that pretended to be official sites for legitimate AV software on the basis of detection results from VirusTotal and manual analysis. We showed that we can visit various software distribution sites from FRAD sites by clicking on the FRAD sites. We also found that

these software distribution sites share common network infrastructures, such as ad networks and redirectors. Thus, further analyses focusing on the web pages arriving from the FRAD sites collected by our system should support efficient collections of software distribution sites.

We then discussed a technique that can be used to evade our classification of FRAD sites. Developers of FRAD sites employ phrases related to the removal information for threats in domain names, URLs, titles, and text contents. This is because they use the topic of the web pages to attract or persuade users. They also place logos of trusted companies to disguise FRAD sites as legitimate sites. A possible evasion technique would be to remove these characteristics that psychologically affect users. However, this also would reduce the interest of users and the usefulness of the FRAD sites to the attackers. In addition, excluding phrases related to malware removal lowers the SEO rankings of FRAD sites and user accesses. Since our system relies on these characteristics to identify FRAD sites, we can accurately detect high-risk FRAD sites that strongly affect the users' psychology.

Since our collection of FRAD sites depends on search engine results, we have not collected all FRAD sites on the Internet. To efficiently collect FRAD sites, we used the names of the cyber threats that are mainly used by attackers to lure users and leverage search engines, which are the most common channel to lead a user to FRAD sites. As a result, our analysis found that FRAD sites are created in many languages and have a large amount of user access. Our system is useful for continuously collecting FRAD sites to create URL blacklists and for analyzing trends for this type of attack.

8 Related Work

We have reviewed related work that investigated the distribution infrastructure for fake AV software and the social engineering techniques attackers use to trick users. Using a combination of unsupervised, graph-based clustering, Cova et al. analyzed the network infrastructure (e.g., domain registration information and IP addresses) of fake AV software distributions to reveal their ecosystem and attack campaigns [2]. Although they investigated the relationship of servers hosting fake AV software, they did not discuss how users access these web pages. Rajab et al. conducted a measurement study that discovered web pages related to the distribution of fake AV software from data collected by Google [18]. They showed the prevalence of fake AV software in malware distributions on the web. Stone-Gross et al. proposed an economic model and estimated attackers' revenue by analyzing back-end servers that attackers used to support fake AV software businesses [23]. They identified the incoming channels that users employ to reach distribution sites, such as landing pages that exploit browsers to redirect users. They also described the social engineering techniques used to install fake AV software using web pages that display fake infection alerts. Although these studies analyzed the infrastructure and traditional distribution techniques for fake AV software–such as drive-by downloads and fake infection alerts–new distribution tactics using FRAD sites have not been revealed. There is also related work

that describes case studies of fake AV software distribution from social engineering aspects [1,4–6,8–10,13–16,19–21,25,27–29]. In most studies, they analyzed fake infection alerts via advertisements that threaten or attract users to install fake AV software. However, no previous study has focused on the FRAD sites or analyzed attackers' techniques that exploit the psychological weakness of users who are suffering security problems.

9 Conclusion

We have proposed a system to crawl the web and automatically identify FRAD sites that introduce fake removal information for cyber threats and lure users to fake AV software. Using the proposed system, the first comprehensive measurement study was conducted to disclose the ecosystem of distributing fake AV software via FRAD sites. We have analyzed both passively collected statistical data on user accesses and actively crawled data to clarify users' risky behavior that leads them to reach FRAD sites and which exposes them to attacks navigated from FRAD sites. Our findings emphasize that it is very difficult for users who are suffering from cyber threats to reach correct removal information, because search results related to the specific cyber threats are poisoned by FRAD sites. Our system is useful for search engine providers and security vendors for excluding and blocking FRAD sites.

References

1. Caballero, J., Grier, C., Kreibich, C., Paxson, V.: Measuring pay-per-install: the commoditization of malware distribution. In: USENIX Security Symposium (2011)
2. Cova, M., Leita, C., Thonnard, O., Keromytis, A.D., Dacier, M.: An analysis of rogue AV campaigns. In: Recent Advances in Intrusion Detection, RAID 2010 (2010)
3. Dittrich, D., Kenneally, E.: The Menlo report: ethical principles guiding information and communication technology research. Technical report, U.S. Department of Homeland Security (2012)
4. Grier, C., et al.: Manufacturing compromise: the emergence of exploit-as-a-service. In: The ACM Conference on Computer and Communications Security, CCS 2012, pp. 821–832 (2012)
5. Invernizzi, L., Comparetti, P.M.: Evilseed: A guided approach to finding malicious web pages. In: IEEE Symposium on Security and Privacy, SP 2012 (2012)
6. John, J.P., Yu, F., Xie, Y., Krishnamurthy, A., Abadi, M.: deSEO: Combating search-result poisoning. In: 20th USENIX Security Symposium (2011)
7. Kharraz, A., Robertson, W.K., Kirda, E.: Surveylance: automatically detecting online survey scams. In: IEEE Symposium on Security and Privacy, SP 2018 (2018)
8. Kwon, Y., Saltaformaggio, B., Kim, I.L., Lee, K.H., Zhang, X., Xu, D.: Self destructing exploit executions via input perturbation. In: 24th Annual Network and Distributed System Security Symposium, NDSS (2017)
9. Li, Z., Zhang, K., Xie, Y., Yu, F., Wang, X.: Knowing your enemy: understanding and detecting malicious web advertising. In: The ACM Conference on Computer and Communications Security, CCS 2012, pp. 674–686 (2012)

10. Lu, L., Perdisci, R., Lee, W.: SURF: detecting and measuring search poisoning. In: Proceedings of the 18th ACM Conference on Computer and Communications Security, CCS 2011, pp. 467–476 (2011)
11. Malwarebytes Corporation: 2020 State of Malware Report (2020). https://resources.malwarebytes.com/files/2020/02/2020_State-of-Malware-Report.pdf
12. MarketWatch Inc: Global Antivirus Software Market Report 2019 and Future Opportunity Assessment 2024 (2019). https://www.marketwatch.com/press-release/global-antivirus-software-market-report-2019-and-future-opportunity-assessment-2024-2019-09-30
13. Mekky, H., Torres, R., Zhang, Z., Saha, S., Nucci, A.: Detecting malicious HTTP redirections using trees of user browsing activity. In: 2014 IEEE Conference on Computer Communications, INFOCOM 2014, pp. 1159–1167 (2014)
14. Miramirkhani, N., Starov, O., Nikiforakis, N.: Dial one for scam: a large-scale analysis of technical support scams. In: 24th Annual Network and Distributed System Security Symposium, NDSS (2017)
15. Nelms, T., Perdisci, R., Antonakakis, M., Ahamad, M.: Towards measuring and mitigating social engineering software download attacks. In: 25th USENIX Security Symposium, USENIX Security 16, pp. 773–789 (2016)
16. Nikiforakis, N., et al.: Stranger danger: exploring the ecosystem of ad-based URL shortening services. In: World Wide Web Conference, WWW (2014)
17. Rafique, M.Z., van Goethem, T., Joosen, W., Huygens, C., Nikiforakis, N.: It's free for a reason: exploring the ecosystem of free live streaming services. In: 23rd Annual Network and Distributed System Security Symposium NDSS (2016)
18. Rajab, M.A., Ballard, L., Mavrommatis, P., Provos, N., Zhao, X.: The nocebo effect on the web: an analysis of fake anti-virus distribution. In: 3rd USENIX Workshop on Large-Scale Exploits and Emergent Threats LEET 2010 (2010)
19. Sharif, M., Urakawa, J., Christin, N., Kubota, A., Yamada, A.: Predicting impending exposure to malicious content from user behavior. In: The 2018 ACM SIGSAC Conference on Computer and Communications Security CCS 2018 (2018)
20. Shen, Y., Mariconti, E., Vervier, P., Stringhini, G.: Tiresias: predicting security events through deep learning. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security CCS 2018, pp. 592–605 (2018)
21. Srinivasan, B., et al.: Exposing search and advertisement abuse tactics and infrastructure of technical support scammers. In: Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW 2018, pp. 319–328 (2018)
22. Stein, E.: Hong Kong Based Malvertiser brokers traffic to fake antivirus scams (2019). https://blog.confiant.com/hong-kong-based-malvertiser-brokers-traffic-to-fake-antivirus-scams-over-100-million-ads-300e251eff06
23. Stone-Gross, B., Abman, R., Kemmerer, R.A., Kruegel, C., Steigerwald, D.G.: The underground economy of fake antivirus software. In: 10th Annual Workshop on the Economics of Information Security WEIS 2011 (2011)
24. Thomas, K., et al.: Investigating commercial pay-per-install and the distribution of unwanted software. In: 25th USENIX Security Symposium USENIX Security 16, pp. 721–739 (2016)
25. Vadrevu, P., Liu, J., Li, B., Rahbarinia, B., Lee, K.H., Perdisci, R.: Enabling reconstruction of attacks on users via efficient browsing snapshots. In: 24th Annual Network and Distributed System Security Symposium NDSS 2017 (2017)
26. Vadrevu, P., Perdisci, R.: What you see is NOT what you get: discovering and tracking social engineering attack campaigns. Proc. Internet Meas. Conf. IMC **2019**, 308–321 (2019)

27. Vissers, T., Joosen, W., Nikiforakis, N.: Parking sensors: analyzing and detecting parked domains. In: Network and Distributed System Security Symposium NDSS (2015)
28. Zhang, J., Yang, C., Xu, Z., Gu, G.: Poisonamplier: a guided approach of discovering compromised websites through reversing search poisoning attacks. In: Research in Attacks, Intrusions, and Defenses, RAID 2012 (2012)
29. Zhang, M., Meng, W., Lee, S., Lee, B., Xing, X.: All your clicks belong to me: investigating click interception on the web. In: 28th USENIX Security Symposium, USENIX Security 2019, pp. 941–957 (2019)