A Review of Computer Vision Methods in Network Security

Jiawei Zhao,* Rahat Masood,* Suranga Seneviratne*
* The University of Sydney, Australia
Email: {firstname.lastname}@sydney.edu.au

Abstract— **Network security has become an area of significant importance more than ever as highlighted by the eye-opening numbers of data breaches, attacks on critical infrastructure, and malware/ransomware/cryptojacker attacks that are reported almost every day. Increasingly, we are relying on networked infrastructure and with the advent of IoT, billions of devices will be connected to the internet, providing attackers with more opportunities to exploit. Traditional machine learning methods have been frequently used in the context of network security. However, such methods are more based on statistical features extracted from sources such as binaries, emails, and packet flows.**

On the other hand, recent years witnessed a phenomenal growth in computer vision mainly driven by the advances in the area of convolutional neural networks. At a glance, it is not trivial to see how computer vision methods are related to network security. Nonetheless, there is a significant amount of work that highlighted how methods from computer vision can be applied in network security for detecting attacks or building security solutions. In this paper, we provide a comprehensive survey of such work under three topics; *i) phishing attempt detection, ii) malware detection, and iii) traffic anomaly detection*. Next, we review a set of such commercial products for which public information is available and explore how computer vision methods are effectively used in those products. Finally, we discuss existing research gaps and future research directions, especially focusing on how network security research community and the industry can leverage the exponential growth of computer vision methods to build much secure networked systems.

Keywords—**Network Security, Malware Detection, Image Processing, Computer Vision, Convolutional Neural Networks**

I. INTRODUCTION

We have been witnessing increasing trends in cyber attacks of various kinds. In 2018, there was a 133% increase in the number of breached records compared to the previous year, and it was equivalent to a rate of 291 record breaches per second [1]. It is estimated that in recent years, approximately 50% of emails we receive are spam, and an average user receives 16 malicious spam emails per month [2]. Similar trends can be found for other types of attacks such as Distributed Denial of Service attacks (DDoS), ransomware, cryptojackers, and other malware. As such, network security has become an important aspect not only for large scale enterprises, but also for small and medium-scale institutions as well as for individuals. Security threats to networks are further expected to be increased with the deployments of IoT, where billions of devices are connecting to the internet and providing more opportunities for attackers. Thus, network security has become

highly important more than ever it has been. For example, recent reports from Gartner, Inc. [3] forecast $124 billion global spending in security, which is an 8.7% increase compared to 2018. Regulations such as General Data Protection Regulation (GDPR) [4] and security breach notification laws [5] are likely to further increase such spending over the coming years.

Recently, as a result of breakthroughs in machine learning, computer vision methods are evolving rapidly, in some cases, even surpassing the human-level performance. Such advancements are mostly becoming possible due to the wide availability of large volumes of labelled data and the processing advances in GPUs. Traditionally, though machine learning methods are significantly used in network security and remains well-surveyed [6]–[9], use of computer vision methods for network security doesn't appear as a direct match apart from the obvious cases in authentication such as face, fingerprint, and iris recognition [10]–[13]. Here, we define computer vision methods in related to network security as methods using image representations of network data such as image feature representations (e.g. key point - descriptor methods such as SIFT [14] and SURF [15]), various image hashing methods (e.g. perceptual hashing), and/or use of image-based machine learning methods such as Convolutional Neural Networks (CNNs) on data related to network security.

The trends to utilise advances in computer vision methods for network security is not solely an academic exercise. Recently the information security industry has also shown a keen interest in this area. For example, the latest RSA conference, which is a gathering of security academics and industry professionals, hosted RSAC Early Stage Expo that invited the industry to showcase their innovative projects, aiming to provide higher level of security for consumers and organizations worldwide [16]. Among the 50 identified promising start-ups, two starts-ups *Pixm* and *INKY*, were using state-of-the-art computer vision methods to detect phishing attacks in email inboxes and other channels such as Facebook, LinkedIn, and Instant messages [17, 18]. *INKY* uses computer vision to block brand forgery emails and spear phishing attempts by spotting imposters by pixels. We provide more details on such products in Section VI. We notice that the industry focus to use computer vision techniques for improving security is quite a recent paradigm since most of the companies start introducing such tools from 2017 [19, 20].

In this survey, we explore the research carried out in the intersection between computer vision and network security. This type of a survey is essential as the field of computer

vision is likely to continue advancing rapidly, and there is a unique, yet non-trivial opportunity to leverage such method to build security solutions for networked systems. Our survey is structured as follows. In Section II, we explain the broader network security topics we have chosen and also describe the overall taxonomy of work under the selected topics. In Section III, we survey the research carried out in the domain of *phishing detection* using image matching methods. Section IV presents work where computer vision techniques are used in *malware detection*, while Section V surveys the work in *traffic anomaly detection*. In Section VI, we survey some of the commercial products in this domain for which the public information is available. Section VII discusses the existing gaps in research, challenges to overcome, and potential future research directions and Section VIII concludes the paper.

II. COMPUTER VISION METHODS IN NETWORK SECURITY

As mentioned in the introduction, the applicability of computer vision methods in network security is not straight forward. We started our survey by referring to the network security threats taxonomy compiled by the European Union Agency for Network and Information Security (ENISA) [21]. The report lists 15 types of threats in the likes of malware, web-based attacks, web application attacks, phishing, spam, denial-of-service, and botnets as the most prominent threats in recent times. Next, for each topic we searched academic publication repositories (e.g. Google Scholar, ACM Digital Library, and IEEE Explore) to check whether there are work that applied computer visions methods to detect such threats. During this process, we found several topics for which a body of work existed based on computer vision methods. Those topics included *phishing detection, malware detection, traffic anomaly detection (targeting threats such as DDoS and BotNets), authentication solutions (especially biometric-based ones such as the face, iris, fingerprint recognition), and steganography*. Next, we short-listed network security topics based on their maturity in the field of computer vision as well as how important are they in the current network security threat landscape.

Though authentication solutions are an essential element in network security, we noticed that authentication schemes based on face, iris, or fingerprints, have been significantly investigated and surveyed from the perspective of computer vision. Also, such schemes usually are straightforward applications of computer vision methods. For instance, Jafri et al. [22], Abate et al. [11] and Zhao et al. [10] extensively surveyed face recognition techniques that utilise various computer vision methods. Similarly, other image-based biometric modalities such as iris recognition and fingerprint recognition remain extensively surveyed [12, 13]. Thus, we exclude such topics from our survey. We also investigated steganography and found that it is an interesting application of computer vision methods in network security. There are multiple recent work in steganography that leverage deep generative models [23, 24]. However, steganography is also a well-surveyed topic [25, 26].

Thus, after this analysis, we selected *Phishing Detection, Malware Detection,* and *Traffic Anomaly Detection*, as topics

of this survey. As we discussed later, phishing and malware are by far the two leading root causes behind many data breaches and have caused damages in the scale of billions of dollars. Similarly, the third topic, *traffic anomaly detection* covers topics in the likes of DDoS detection and BotNet detection, which are also becoming a significant threats to internet services. We next define and describe each of these topics with emphasis on the importance of coming up with solutions to detect such threats.

A. Phishing Detection

Phishing is a social engineering attack that manipulates human trust to obtain confidential information about users such as passwords, credit card numbers, usernames, and personal data. The attack is often performed through instant messages or emails where the users are prompted with web pages that appear to be legitimate. The nature of the content of the message or the email will entice the users to enter their user names and passwords, bank account details, or other important personal information. This enables the possibility of a second, more devastating attack using the harvested information.

Phishing attacks are an ongoing problem in the internet. Attackers are using many sophisticated methods such that even the most tech-savvy users struggle to distinguish phishing attempts from legitimate web pages. According to Phishlabs report, phishing attacks grew 40.9% in 2018, while targeting 83.9% of industries that offers financial, email, cloud, payment, and SaaS services [27]. Overall, 98.9% of these attacks target corporate users for credential theft or email scams.

Many high profile network breaches such as iCloud celebrity pictures, Sony hack, Bangladesh Bank heist, and eBay data breach have been traced back to attackers infiltrating the network using a well researched and crafted phishing attacks misguiding employees to enter their credentials. Overall, 90% of the data breaches occurred globally in 2017 were results of phishing attacks. Recent reports also indicated that 75% of the organizations faced at least one phishing attack in 2017, and the average click rate of the targets is an eye-opening 9%. According to FBI's Internet Crime Complaint Center, Global Business Email Compromise (BEC) losses exceeded 12 billion US dollars in 2018. In 2017, these scams cost organizations almost 700 million dollars [28].

It is essential to detect phishing attempts before users click a malicious URL to minimise damages. In this regard, there are two automated phishing attack detection methods: i) blacklist/whitelist-based methods and ii) text similarity-based methods [29, 30]. Blacklist-based methods attempt to keep a list of domain names or exact URLs to known phishing websites and alert users if they are trying to visit those links. However, phishing websites are highly dynamic, and the average lifetime of a phishing web page is only a few hours [31]. Therefore in many cases, zero-hour phishing attacks usually easily bypass blacklist-based methods. Whitelist-based allows users to browse only a set of web pages that are deemed safe, which is not practical in many situations. Text similarity-based methods delve into semantics of the textual content, be it an email or a web page and try to decide whether it is

a phishing attempt or not. Nonetheless, this method is likely to fail in future with the increased use of code obfuscation techniques. As the fundamental characteristic of a phishing page is it's visual similarity to a legitimate target page turning into computer vision method appears as a natural alternative to consider. In Section III, we investigate such solutions that have been proposed for detecting phishing attempts.

B. Malware Detection

Malware is a collective term used for various malicious software variants such as viruses, trojans, ransomware, and spyware. Malware typically consists of a code developed by an attacker to cause damages to data and systems or to gain unauthorised access to a network. Different types of malware follow different propagation methods to increase the reach. For example, some malware are delivered in the form of a link over email and require the users to click on the link while some malware spread through a malicious file download into a system and are self-triggered.

According to a recent report [32], from 2015 to March 2019, 413.73 million new malware programs have been reported, out of which 137.47 million have been reported in 2018 alone. The report also states that on average 279,545 new malicious programs are discovered every day. Malware can cause varying levels of damages from a single computer to national-scale infrastructure. For example, in 2017, Petya ransomware outbreak which started in Ukraine and subsequently spread across Western Europe, North America, and Australia caused losses over $10 billion [33].

Traditional malware analysis is usually based on either static code analysis (signature-based) or dynamic code analysis (behaviour-based) [34, 35]. The signature derived from static analysis, is a short sequence of bytes unique to each known malware, which allows identifying newly encountered malicious files. This type of a detection method is less responsive to new malware and fails when codes are obfuscated. On the other hand, methods based on dynamic analysis (behaviour-based), can detect malware based on its run-time activities. These types of detection methods, however, give high false positive rates where benign programs are falsely classified as malicious programs. Additionally, dynamic code methods are computationally expensive and require a lot of resources for proper execution. To alleviate these problems, recently computer vision approaches have been proposed in the context of malware detection. For example, computer vision methods have been shown to be successful in analysing codes as images when they are obfuscated. Also, computer vision speedups detection process by providing a summarised picture of possible attacks through visual representation such as Treemaps and Thread graphs. In Section IV, we survey such research work.

C. Traffic Anomaly Detection

A network traffic anomaly refers to a deviation from the usual traffic profile in a network. Unlike our first two topics, a traffic anomaly does not directly identify an attack. Rather it indicates that something unusual is happening in the network, which can be caused by a range of attacks such as DDoS, botnet activity, port scanning, or even because of some malware-infected nodes in the network. For instance, in September 2016, Mirai malware launched a DDoS attack on the popular security news web site/blog; *Krebs on Security*, by generating more than 620 Gbps of network traffic [36]. The malware infected over 600,000 IoT devices, turning them into an army of remotely controlled bots. The reason for the use of a large number of IoT devices is to bypass anomaly detection systems which monitor the IP address of incoming requests and filters them if there are too many requests form a limited set of IP addresses. According to recent Kaspersky reports, 50% of all DDoS attack led to a severe disruption of services [37] and there were cases where the attack continued for hundreds of hours.

Traditional network traffic anomaly detection approaches rely on underlying fitted models, which are labelled based on normal traffic behaviour. These approaches detect and characterise unusual patterns in the traffic based on deviations from the corresponding labelled data. Since these models are built on the knowledge of legitimate network traffic behaviour, the traditional detection approaches can identify previously unknown attacks. However, these conventional approaches suffer from high false positive rates. Also, if the set of rules (such as incoming IP addresses, port number, outgoing IP addresses, and protocols), defined within an underlying model are too complicated or too easy, then the accuracy of a detection system is affected. For example, an activity such as directory traversal on a targeted vulnerable server, which complies with a network protocol, easily goes unnoticed as it does not trigger any alert. Thus, efficiency of a detection system depends on how well rule set is implemented and tested on all protocols and scenarios. Another approach to detect traffic anomalies is based on computer vision methods, despite there are limited commonalities shared between network anomaly detection and computer vision tasks, such as image retrieval and object shape recognition. In Section V, we analysed a few traffic anomaly detection methods, that are based on computer vision.

D. Structure of the Survey

We structure our survey with separate sections to the three main topics that we survey. Across the three main topics we broadly categorise the computer vision methods that are being used in to four; *i) image representation-based methods*, *ii) image feature-based methods*, *iii) image hashing-based methods*, *iv) neural network-based methods* as described below and schematically summarised in Figure 1. However, we also highlight that we did not find work in all four sub-categories in all three main topics we survey.

- **Image representation-based methods** have been applied in malware detection and traffic anomaly detection. In these methods, a binary executable, log files, or a data file related to a malware or network traffic is first converted into a visual image and then visual fingerprints are generated based on malware behavioural activities such as API call sequences, I/O request packets (IRPs), resource consumption, network activities etc. These visual fingerprints are represented through properties such as
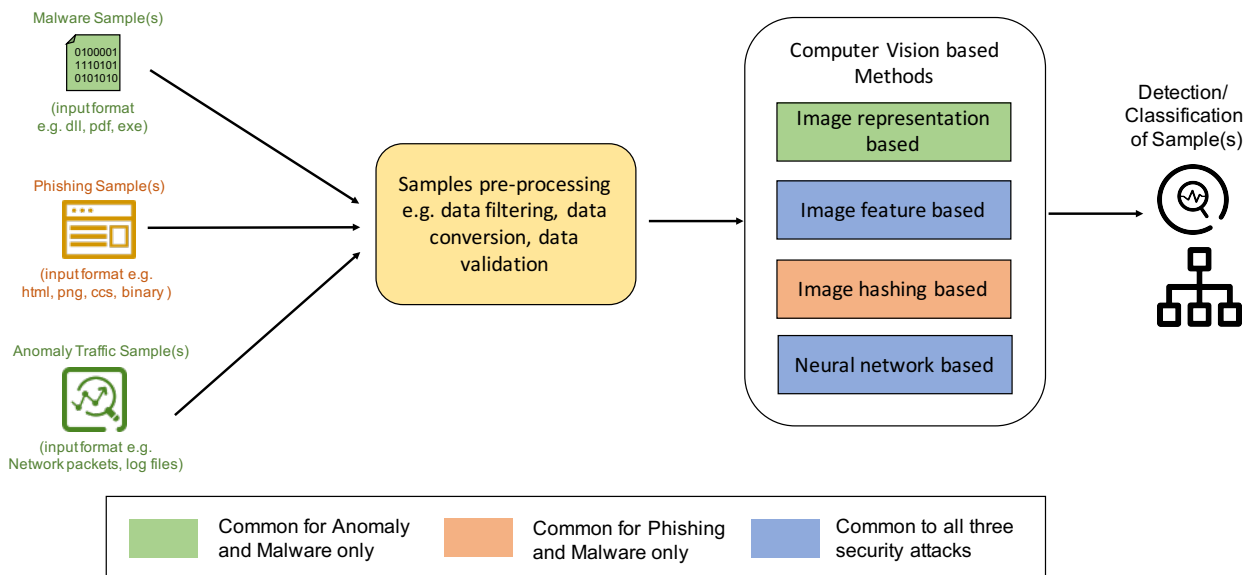
Fig. 1: Taxonomy of Computer Vision Methods in Network Security

x, y coordinates, RGB colours, and intensity level. Unlike the next described feature based methods that compare features, the image representation-based methods compare visual fingerprints to detect maliciousness.

- **Image feature-based methods** generate features from image representations. For example, for phishing websites features are generated from the screenshots of the web pages or images/logos embedded to the web page. Example features include the traditional computer vision features such as Scale Invariant Feature Transform (SIFT) features [14], Speeded Up Robust Features (SURF) [15], and various image histogram features. Similarly, malware binaries or network traffic can be represented as images and features can be extracted.

- **Image hashing-based methods** generate image representations that are efficient during retrieval and comparison tasks. For example, to identify visually similar web pages to a given page, a hash of the web page can be created. When a suspicious web page is found, the two hashes can be compared to see the new page is visually similar to the original page. Similarly, to improve the efficiency of malware detection approaches, local sensitive hashing (LSH) schemes have been introduced that checks the similarity between visual images of malware.

- **Neural network-based methods** accept some form of images as the input, yet do not require to generate features separately. For example, convolutional neural networks are representation learning method that do not require separate feature generation. The network itself learns what features are important to the task at hand.

The rest of our survey is organised as follows. In Section III, we provide detailed analysis on existing computer vision based phishing detection techniques, followed by malware detection and traffic anomaly detection in Section IV and V, respectively. We then discuss industry tools or products in Section VI, that are utilising computer vision techniques for network security attack detection. In Section VII, we discuss open challenges in applying computer vision methods to network security, and also highlight possible future research directions. We conclude our survey in Section VIII.

III. VISUAL SIMILARITY BASED PHISHING DETECTION

Many phishing attempts over various delivery forms such as emails, web, and mobile apps, try to misguide the target users by providing visually similar interfaces to the typical applications and services they use. As such, it makes sense to use computer vision methods in building phishing attempts detection systems. We found a significant body of research fitting into under three sub-topics we discussed in Section II; *i) image feature-based methods*, *ii) image hashing-based approaches*, and *iii) neural network-based approaches*. All the work we surveyed applied some form of image transformation, and as a result, we do not report any work under the sub-topic of *image representation-based methods*.

A. Image feature-based methods

As mentioned before, image feature-based methods first identify an image representation of the possible phishing attempt. For example, the image can be a screenshot of the phishing web page. Then, the image is converted into a feature vector through various methods. Afterwards, different techniques such as k nearest neighbour search, are used to find whether the web page is visually similar to a known legitimate page. The overall pipeline can be evaluated using performance metrics such as accuracy, precision, and recall using a ground-truth established dataset. We show this process in Figure 2.

As such the aspects to investigate under image feature-based methods for phishing detection are; *input data (type of*
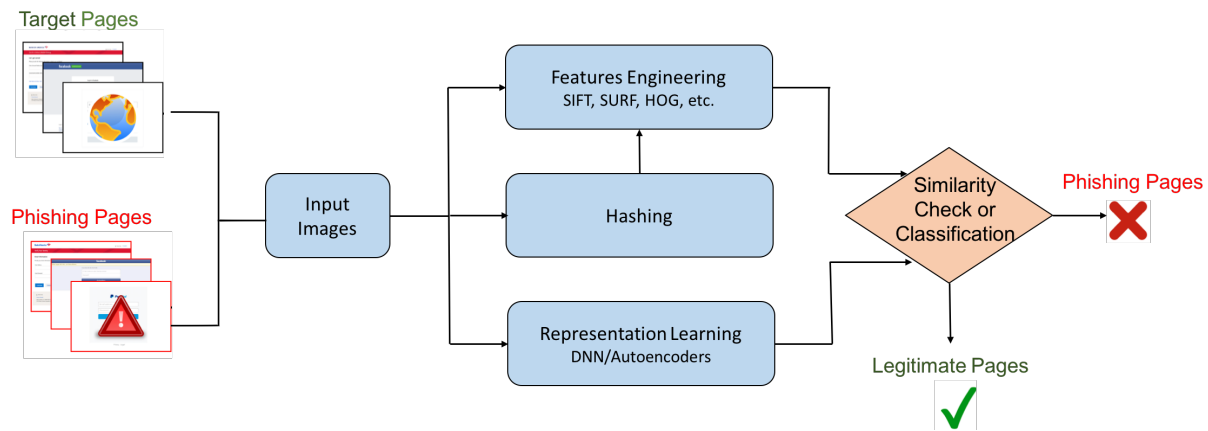
Fig. 2: Process of image feature-based methods

data, how authors collected the data, and how the authors established ground truth), **feature representation**, **similarity calculation**, and **performance levels**. We next present work in this category with particular focus on above aspects.

i) Web page screenshot & image content matching-based approaches

One of the early work in this area is the SiteWatcher proposed by Liu et al. [38, 39]. It is service designed to run on mail servers and flag possible phishing links in both incoming and outgoing mail messages. All mails are monitored to check whether they contain suspicious words associated with a set of protected and sensitive web pages (potential targets for phishing attempts) for a given organization. If an email contains such words, SiteWatcher will consider all the URLs in that mail as suspicious and send those URLS together with the potential target URL to the *visual similarity assessment module* for further analysis.

The *visual similarity assessment module* first identifies a set of *salient blocks* by parsing the DOM (Document Object Model) tree of the web page. Salient blocks for a web page are the clearly distinguishable blocks such as header, footer, sidebars, and any other blocks in the page body that have consistent content either visually or semantically, and considerably different from other blocks [40]. Then, it considers three types of similarity between the blocks; *block-level similarity, layout similarity,* and *style similarity*. A set of features defines each similarity type. For instance, for *block similarity*, *visual similarity assessment module* first decides whether the block contains text of images and based on that calculate features for the block. Text block features include features such as background colour, foreground colour, and anchor colour (colour of the hyperlinks). Examples of image block features are dominant colour of the image, image display, and image display height. The similarity between the two blocks in the phishing web page and the target is calculated as the weighted sum of the individual feature similarities, and the weights are empirically decided. Two blocks are considered to match if their similarity value is higher than an empirically determined threshold. Similar approaches are followed for other similarity

types with different features.

Authors evaluated the SiteWatcher using a ground truth dataset that contained 320 commercial bank home-pages and eight true phishing pages targeting six of the banks' pages based on the reports published by Anti-Phishing Working Group (APWG).[1] Authors considered the six true pages as the query to search for the visually similar pages in the dataset of 320 pages. If the similarity of any of the three metrics (block, layout, and style) were higher than a threshold t, the retrieved web page is reported as a potential phishing web page. Results indicated that for t values greater than 0.7 almost perfect true positive rate and false positive rate can be achieved. For a baseline comparison, authors also evaluated the performance of a pure text features based method and showed that such an approach have a much higher false positive rate.

While the paper has shown some promising results, it has several limitations. First, it identifies the target based on the presence of keywords and only activate the visual similarity assessment module if there are matches with a known set of keywords. Such a method is not scalable and may also miss more sophisticated attempts based on obfuscated code [41]. Similarly, salient block identification is also text based DOM tree analysis which may not produce the optimal results with the use of client-side scripts and code obfuscation. The use of computer vision methods is limited to using manually curated image, colour, and style based features. Finally, the evaluation is limited to a smaller dataset, and it is unclear whether the methodology will scale to the range of thousands and hundreds of thousands of web pages and still able to maintain the same levels of true positive and false positive rates. Nonetheless, SiteWatcher is an important work in this domain as it proposes a useful architecture that allows building more effective and efficient phishing detection systems using computer vision methods. For example, salient block identification can be done entirely using computer vision algorithms and now by leveraging the recent advances of CNNs, we are in a position to design better similarity comparison methods.

Medvet et al. [42] proposed a phishing detection technique that considers the text pieces (including text style), embedded

---

[1] https://apwg.org

images, and overall visual appearance of the page to determine the page similarity. The DOM tree was used to compare the textual features such as content, colour, size and position. 2D Haar Wavelet transformation and colour histogram are the techniques used to analyze the embedded images and the overall image of the web page. Then a single similarity score is calculated and compared with a predefined threshold value to decide whether a web page is a phishing page or not. The authors claim that this method has zero false positive rate.

The evaluation process contains two parts which are signature extraction and signature comparison. The comparison of the two signatures between two web pages requires multiple operations, corresponding to each feature. Authors found that it takes about 11 seconds to compare negative pairs and around 4 seconds to compare positive pairs. Authors also proposed a few optimizations, that significantly reduce the computational cost and make the time of negative comparison reduce to a few milliseconds. This solution is effective to detect phishing web pages that are visually similar to the legitimate ones. Though this work is inspired by AntiPhish [43] and DOMAntiPhish [44], compared to AntiPhish and DOMAntiPhish, it improves in the detection of phishing pages containing many embedded images. The approach is also similar to Liu et al. [38], however has a much higher feature space.

In [45], authors proposed a solution that uses Earth Mover's Distance (EMD) to compare the visual similarity between web pages. In this method, the web page screenshots were first converted to a lower resolution of 100×100 using Lanczos resampling [46] since it can generate sharp images. The features were calculated based on the colours of each pixel where a feature is the colour and the centroid of its position distribution in the image in a reduced colour space of 4,096 colours. Frequency of the colour (i.e. the total number of pixels using that colour in the image) was used as the feature weights to create the full signature of the web page. Finally, authors decided whether a new page is a phishing web page or not based on whether the EMD between two web pages is less than an empirically decided threshold. Authors conducted multiple experiments to evaluate different strategies of choosing the threshold.

To evaluate the proposed solutions, authors established a web page image dataset that contain images of 10,272 legitimate pages and nine phishing pages, which authors collected based on real-world phishing attacks. At the best threshold value, authors were able to identify eight out of nine phishing web pages with only 73 false positives, corresponding to a *true positive rate* of 88.89% and a *false positive rate* of 0.71%. Also, authors compared their approach with a completely text based approach and the above discussed region based approach [38, 39] and showed that new method have lesser *false positive rate*. Work of [47] et al. also used a similar EMD based approach.

This type of work are an improvement compared to Liu et al. [38, 39] and has proposed better methods of capturing the visual similarity as indicated by the results. Nonetheless, going to a lower resolution might lose some information when it comes to larger datasets. Use of only colour features can also be a limitation, and a better approach may be to use both colour and region-based features. Another limitation is that the use of EMD is not justified with results. It is unclear whether a distance metric like L_2 distance which is having lesser computation requirements can provide the same level of performance. Finally, the dataset had only nine cases of phishing attempts. Evaluation in a much larger dataset is required to get a better understanding of the final solution.

Chen et al. [48] also started with the screenshots of target and phishing pages, but proposed to use more advanced methods of image representations. Authors first converted the web page image to grey-scale by averaging the red, green, and blue values of each pixel in the image. Then Harris Laplacian corner detection [49] was used to identify a set of key points in the image that are usually invariant to modifications in the likes of shifting, lighting variation, and colour transformation. To describe each key point, authors use a lightweight version of Contrast Context Histogram (CCH) [50], which captures the contrast distribution in the neighbourhood of a key point. To make CCH computation lightweight, author only considered the grey-scale image of the web page. Authors calculated the similarity between two pages using Euclidean distance-based key point matching. The proposed solution was evaluated in dataset comparable to previously discussed results and achieved similar performance levels.

Corona et al. [51] developed a system called DeltaPhish to detect phishing web pages in compromised websites, which is the first system utilizing websites vulnerability for phishing detection. The overall process of DeltaPhish involves browser automation, feature extraction, HTML-based classification, snapshot-based classification, and fusion classification.

In snapshot-based classification part, the feature representations used to classify snapshots are HOGs (Histogram of Oriented Gradients) and colour histograms. Compared to other feature representation like SIFT (Scale-Invariant Feature Transform), the features used in this paper can achieve better performance for image classification of high inter-class similarity. Authors also do image tiling to gain more spatial information of snapshots. Thus, the features are extracted from the whole snapshot picture, its quarters and sixteens. Based on the feature vectors, the visual similarity is evaluated for input image and homepage. Finally, authors leverage a linear SVM to do the classification. The true positive rate of only snapshot-based classifier is about 82.5% at 1% false positive rate. The processing time of snapshot-based classifier is more than 1.2 seconds averagely, which HOG feature extraction takes most of the time, and colour feature extraction only takes less than 3 ms averagely. The limitation of this method is that it depends on the assumption that within a website, the legitimate pages have a similar visual appearance with the homepage.

In a slightly different approach of using the web page screenshots, Lam et al. [52] started by identifying non-overlapping regions called *layout blocks*. This was done by first using Otsu's threshold methods to convert the screenshot to black and white and identifying the blobs that have the same colour. Then these blobs were used as a mask to identify regions in the original image. The similarity between two pages was calculated as the sum of the differences in width, height and location of each pair of blocks. Then authors

used empirically established thresholds to decide what is the minimum similarity required to decide on a phishing web page. Compared to the EMD algorithm used in [38], this method is faster and can handle non-square image inputs.

Different from the previous papers, GoldPhish proposed by Dunlop et al. [53], used optical character recognition (OCR) to extract text from a screenshot of a web page and used that text to decide whether the page is a phishing page or not. Authors, first take a 1200×400 pixel image of the web page, covert it to black and white, and extract the text of the image using a commercial OCR software. As the image also contains the logos present in the web page, the OCR will extract any text available in logos, which authors claim as an advantage compared to other methods. The extracted text is entered into Google Search as the queries, and related websites were retrieved. The intuition here is that, as the text in logos, and the text in websites are indexed by Goolge, the search will return the corresponding legitimate site (if there is any) as one of its top results.

Testing GoldPhish on 100 active phishing web pages identified from PhishTank and 100 legitimate web pages, showed good results with only two false positives and zero false negatives. Authors also compared their method with text-based methods and heuristics-based methods and showed that GoldPhish achieves better results.

While the results of this method are promising, it has a significant limitation. Use of OCR can be justified only for the websites that have a logo at the top of the web page, and the logo has text in it. For logos, that do not have embedded text, Goldphish might perform worse than logo detection-based approaches. Otherwise, Goldphish is expected to perform the same as a solution that does text matching between web pages without any image processing methods. Also, the comparison was made only with simple methods, and thus it is difficult to compare how GoldPhish will perform compared to other image-based approaches. Nonetheless, the idea of using the Google indexing system as we discussed earlier under [54] is innovative. Due to the limited lifetime of phishing web pages, they are unlikely to be indexed by Google. Thus, the text queries will get hits from the legitimate web page.

In a recent work, Tian et al. [55] also highlighted the necessity of adding visual features to the phishing detection problem to overcome the content obfuscation methods used by attackers. Authors used the open-source OCR engine Tesseract[2] that used the detected words as features in data going into a machine learning classifier.

Hara et al. [56] enhanced phishing detection based on image similarity by detecting phishing sites even without the original registered sites. The comparison of two sites in this method is not only between phishing sites and their victim sites, but also between phishing sites that spoof the same site. The reasons why this method does not need an initial database are: 1) many phishing sites mimicking a same site are visually similar to each other; 2) their system will discover sites on which there is a different image shown from the previous one and use these sites to detect new phishing sites.

By taking the screenshot of web page, Hara et al. utilised an image searching tool called ImgSeek [3], which can output images similar to the input image and their similarity score. Two hundred twenty-four original legitimate sites spoofed by 2262 phishing sites were extracted from 521 legitimate sites by their system. The detection rate is 82.6% while the false positive rate is 18.0%, which can be improved by adding whitelist. Even though this method has a high false positive rate, it can handle many kinds of phishing sites and combine with other methods. Notably, compared to Liu and Medved [38, 42], this system does not rely on an initial set of legitimate pages.

ii) Logo detection-based approaches

Wang et al. [57] proposed a slightly different approach by using the logos embedded in web pages than using the entire screenshots. The method consists of three steps: first, authors determine if the web page contains the given logo. Second, they check if the brand owner has authorised host an IP address to use its logo. Finally, for sites containing unauthorised logos, if the user enters keyboard input on the page, the system will issue a warning.

For the first part, authors use a simple heuristic approach to identify possible logos on the web page under the assumption that logos are usually placed at the top of the web page. Starting from the top of the page, stripes of height $N = 100$ are considered, and continuous blocks of single colour are removed as those parts are unlikely to contain logos. For the remaining blocks, SIFT representation [14], which is another key point and descriptor-based representation similar to above mentioned CCH, is calculated and matched against a local copy of a logo image database. If there is a match, defined by having a Euclidean distance less than a threshold, the brand owner is contacted to check whether the use of logo in the current web page is allowed using two methods; a modified DNS query or certificate base approach where the original logo contains the brand owners public key.

To measure the performance of the solutions, authors first established a dataset of popular logos that covered 166 distinct brands (again taken from 176 distinct instances. The dataset was augmented with the colour inverted versions as it was known to a common technique used by the phishing attempts (i.e. 352 logos). The second dataset was the "brand dataset" that contained screenshots from 23 legitimate websites that are commonly targeted by phishing attacks. For each website, authors collected ten pages, and thus the total size of the *brand dataset* is 230. The third dataset contained the 219 screenshots of known phishing sites taken from PhishTank database.[4] Authors were able to achieve a true positive rate of 90% at 3% false positive rate.

While the use of SIFT might improve the image matching, this work has a significant limitation in detecting logos, which is done based on a straightforward heuristic algorithm. Such an approach is likely to be inefficient in cases where the web site banner does not have large areas of a single colour. In such a

---

[2]https://github.com/tesseract-ocr/

[3]http://www.imgseek.net/

[4]https://www.phishtank.com/

setting, the area the algorithm selects, may not necessarily be logos and might be quite large, resulting in a high latency when it comes to SIFT key point calculation. While it is understandable, that when the authors did the work, object detection methods were not evolved and now this work can be improved by leveraging the recently proposed deep learning-based logo detection methods [58, 59]. Another limitation is the latency. The solution is proposed as a safe browsing solution in contrast to the above discussed email solutions. When it comes to browsing, user experience is highly important, and users will be annoyed by slightest delays. Thus, the total latency (consisting of logo identification, SIFT key point and descriptor generation, and finally matching) in the range of tens of seconds is not acceptable. Nonetheless, such delays may suit well for an email security solution, where there is some lag between the mail is being received by the server and accessed by the user. Also, email server can control access to the email until the scanning is finished.

In [60], Zhou et al. proposed a method that combines logo detection and web page snapshot comparison. The *logo detection and comparison* were based on SURF (Speeded Up Robust Features) [15] descriptors which are faster to calculate than previously discussed SIFT key points and descriptors. For each suspected web page screenshot and each logo image, the authors identified two sets of interest points using the SURF detectors and then generated two sets of SURF descriptors. For each descriptor of the logo image, the descriptor that matches best (above a defined threshold) in the suspected web page is identified using Euclidean distance. Nonetheless, still, there are matches between the logo and irrelevant parts of the web page. To address this, authors proposed to use homography transformation [61] estimated using random sample consensus (RANSAC) [62]. Homography transformation allows projecting the left top and right bottom coordinates of the logo image to the suspected web page image. By doing so, authors then make some decisions based on heuristics such as whether the two points are apart enough to be a logo or which part of the web page the identified logo is located. For *snapshot comparison* authors used the same approach as Fu et al. [45] that was discussed earlier.

Authors collected three types of data to evaluate the performance of their system. The *phishing web page* dataset contained 2,129 phishing web pages that were collected from PhishTank. As *target web page dataset*, authors use the screenshots of the web pages of PayPal, PosteItaliane, eBay, iTunesConnect, and MyAppleID. Finally, as the *irrelevant legitimate web page subset*, authors collected snapshots of 1,367 web pages by discovering URLs via Google search for 26 keyword searches. Overall, the authors were able to achieve 91% true positive rate and 1.46% false positive rate.

While on paper, it can be expected that combination of the ideas of logo comparison and snapshot comparison should improve the detection rates of phishing web pages, the results of this paper do not justify that there is extra benefit. Nonetheless, all the solutions are using different datasets, and it is difficult to establish which methods are superiors. This leads to the following discussions we make on the requirement of establishing benchmark datasets in this domain. Also, as

mentioned above, we emphasise that significant improvements have happened in computer vision space in the last few years and those ideas can be borrowed to improve this type of work (e.g. logo detection and comparison).

PhishZoo [63] used logo matching in a much broader phishing detection systems that consider other similarities between web pages in the likes of URL patterns and HTML text. The image similarity aspect of the solution involved keeping a logo database and matching all the images content of suspected web pages against the stored logos using SIFT key points. Performance evaluation showed that using text features only resulted in very high accuracy (97.6%) but also a very high false positive rate (18.7%), which is not desirable. On the other hand, use of only image matching resulted in a moderate accuracy (82.7%), but a very low false positive rate (2.5%). Overall the combination of images and text resulted in an accuracy of 90.2% and 0.5%.

This work provides an important justification on why image-based methods are required to be incorporated into usual text-based phishing detection systems. Nonetheless, authors highlighted some notable limitations they identified throughout their work. First, the SIFT image matching can identify only near similar matches. For instance, if the logo is rotated significantly, cropped, or combined with other elements, SIFT fails to give a high score. Second, the latency of the system, which was in the range of 7-17 seconds, is mainly contributed by the image matching phase, than the SIFT key point generation (i.e. based on the number of stored logos). As a result, more efficient signature matching methods need to be used in real-world deployments.

In 2014, Fatt et al. [54] proposed a method based on website *Favicons* (i.e website shortcut icon displayed on the address bar of a web browser) to do phishing detection by utilizing the Google Search-by-Image API. Then they used the results of the Google search engine to evaluate the authenticity of the website. Intuition of this approach is that, if the *Favicons* is a copy of a popular website or a popular logo, Google search results will list the target website first as well as some texts describing the logo (e.g. *Best guess for this image: icon paypal*). Then based on the solutions returned by the search query, the authors used four text-based heuristic rules to decide whether the website from where they got the *Favicons* is a phishing site or not. Authors used 500 phishing pages (chosen from PhishTank) and 500 legitimate pages (chosen from Alexa top 500 global websites) to validate the performance of their approach. Authors reported a 97.2% true positive rate and a 2.8% false positive rate.

In contrast to previous methods that keep their own images of the target pages, this method leverages the Google's page indexing system which is an advantage as it eliminates the need of maintaining a database of target web sites. As indicated by results, this method will provide very high accuracy for target pages that are indexed by Google and sufficiently widespread. Nonetheless, the major drawback is that it will work only if the *Favicons* is available on the web page. As a result, this is a method more suitable to be one specific feature of much larger solution covering various other aspects.

B. Image hashing-based approaches

One common aspect of all the work discussed in above Section III-A is that once the feature representation of a suspected page is generated, a $k-NN$ (k nearest neighbour) search is required to identify the target page from a set of potential pages stored in a database. This search can be time-consuming as highlighted by some work [63] and is proportional to the size of the target web page database. This is exacerbated by the higher complexities in calculating EMD and SIFT representation. Though this search process can be optimised by using more efficient data structures such as $k-d$ trees and *vantage-point trees* [64], there is a fundamental requirement of making it more efficient. We next discuss a set of work, that addressed this problem by using *image hashing methods*.

In [65], the phishing websites were detected by analyzing their structural characteristics and page screenshots. The primary computer vision technique used by the authors here was to compute a hash of the web page screenshot and retrieve similar images based on the Hamming Distance. Authors experimented with three hashing methods; *MD5*, *SHA-512*, and *pHash (Perceptual Hashing)*, in which latter is specifically designed for image matching while the former two are more suitable for string/text content matching. Authors first demonstrated the performance benefits of their method using a limited dataset of five phishing pages taken from PhishTank. As expected, *pHash*, which is specifically designed for image matching performed better compared to *MD-5* and *SHA-512* that are more suitable for string or text content matching.

Then, the authors conducted a much large scale study by collecting URLs from Tweets and checking whether there are phishing ULRs or not using the proposed approach. Authors collected 1,829,531 URLs from Tweets. Authors found over 200 phishing pages trying to impersonate top phishing targets. The average processing time to decide on a page was less than 4 seconds, which was an improvement compared to previous work.

This work demonstrates the efficiency of hashing based image matching methods. Nonetheless, there are many other image hashing methods such as *block hashing, wavelet hashing,* and *difference hashing* and more thorough experimentation is required to identify which methods works well for website screenshot matching and provides a good trade-off between accuracy and computational efficiency.

Fotiou et al. designed a model to detect phishing websites based on the perceptual hash of the web pages' screenshots [66]. The whole mechanism contains a web browser extension, a back-end application and a database. The web browser extension is responsible for capturing a screenshot of the currently loaded web page. The back-end application is used to extract and compute information of page screenshots, and the database is used to store all information signatures. In their implementation, authors used the open-source perceptual hash library *(pHash)*, and they evaluated information signatures by applying perceptual hash over a web page's screenshot bytes. And then authors used the hamming distance of the hash value to compare the similarity of the screenshots. Similarity threshold represents the maximum distance between two almost identical hash values of screenshots. If the numerical distance of the hash value of two screenshots is greater than the similarity threshold, it will be considered as a phishing attempt.

The dataset contains 100 phishing pages captured from PhishTank database and 100 legitimate pages obtained from the list of Googles top 100 most visited sites in the United States. As a result, by choosing a good similarity threshold, the authors could achieve a recall of 81% and false positive rate as low as 1%. Even though the dataset is small, the result still shows a real-time performance. While this work is similar to [65], no comparison was provided.

C. Neural networks-based approaches

One of the recent advances in computer vision is CNNs and other related deep learning methods such as auto-encoders that are able to eliminate the need of curating features manually; similar to what was done in above discussed work that were using SIFT or SURF. As such, we next investigate work that leverage the representation learning capabilities of deep learning methods to detect phishing web pages. As phishing detection work spans beyond the point where CNNs became mainstream (i.e. after the success of AlexNet [67] in the ImageNet competition in 2012), we also found some work that used traditional neural networks, and we discuss that work here as well.

Borgolte et al. [68] proposed MEERKAT to detect *de-faced websites*, which is a different but potentially a related problem to phishing web page detection. In a *de-faced* or a *vandalised* website, an attacker replaces some content of the legitimate site, and as a result it's visual appearance is different from what it is supposed to be. First, authors extract 160×160 windows from web page screenshots for both de-faced and legitimate web pages. Then, authors build a deep neural network to learn advanced features from these windows automatically. The structure of this deep neural network is comprised of a stacked auto-encoder and a standard feed-forward neural network. The stacked auto-encoder is utilised to de-noise the images and learn high-level features from the images. The feed-forward neural network with dropouts is leveraged to do the final classification.

Authors showed the feasibility of MEERKAT on a large website defacement dataset that included 10,053,772 damages observed between January 1998 and May 2014, and 2,554,905 legitimate websites. Overall, MEERKAT achieved a true positive rate between 97.422% and 98.816%, a false positive rate between 0.547% and 1.528%, and a Bayesian detection rate between 98.583% and 99.845%.

While the problem that this work addresses is different, the same approach can be used for phishing web page detection. The idea of learning latent representations from data itself, rather than manually identifying and calculating features will improve the efficiency of designing a phishing detection pipeline. The method is right now based on windows; therefore further work is required to understand what is the suitable setting for phishing web pages. For example, the

model can take the whole screenshot as the input, only some key segments, or random patches and such decisions require further analysis.

Adebowale at el. [69] proposed ANFIS (Adaptive Neuro-Fuzzy Inference System), a traditional neural network and fuzzy logic system based on calculated feature representations. Authors created 35 features from web pages that include 22 text-based features to represent the structure of the page, eight features to represent the frame-based properties, and five features to represent the image resources of the web page. However, the image features are not directly related to the visual similarity or the content of the images. For example, *favicon* feature indicates whether the icon is loaded from a domain other than the domain shown in the address bar. As such, this work can't be considered as a working that truly uses computer vision methods for phishing detection.

In 2019, Abdelnabi et al. [70] proposed a new visual similarity-based framework, called WhiteNet, utilising a triplet convolutional neural network for phishing website detection. Authors also released a new dataset called WhitePhish, which can be used to further benchmark phishing detection methods. The dataset contains phishing web pages, legitimate training web pages, and legitimate test web pages. The triplet learning process that was successfully first used in face recognition [71] involves three inputs; an anchor image, a positive image which is a screenshot of the same website as the anchor and a negative image from the different website. The triplet loss that is being optimised during the training process attempts to make the embedding of samples in different classes further apart from the embedding of samples in the same class. Authors were able to achieve a recall of 95.81% at a FP rate of 6.88%. This work indeed is an example of how modern computer vision ideas can be utilised effectively to phishing detection.

D. Summary of visual similarity based phishing detection:

To summarise, in this section, we surveyed research work that used computer vision techniques to detect phishing web pages that are visually similar to target web pages. The motivation of using computer vision methods in this setting is the fact that many malicious websites increasingly use obfuscated codes using client-side scripting languages such as Java-scripts and it is relatively easy to create a website that shows no or minimum similarity to the original website when text features extracted from HTML and Java-script codes are used, yet shows very high visual similarity to the user once the web page is rendered in a browser. This was evidenced by the results; many work reported that image-based methods achieve higher performance compared to text feature-based methods.

We presented these work under three topics; *i) image feature-based methods*, *ii) image hashing-based methods*, and *iii) neural networks-based approaches*. *Image feature-based approaches* create features such as SIFT, SURF, HOG, and CCH from the full screenshot of the web page, segments of the page, or embedded images/logos and measure how similar are those to the features from a target web page. *Image hashing-based methods* improve the speed of the similarity check by creating various hashes from images, and thus converting the time-consuming EMD and L2 distance calculations

to more efficient Hamming Distance calculation. Finally, *neural networks-based approaches* leverage the recent advances in deep learning methods and eliminate the requirement of manually creating the features. We found more work in the category of image feature-based methods, but very recent trends in using deep learning-based methods. We summarise all the work we surveyed in Table I.

We next discuss several challenges and limitations we observed over all the work and worth addressing to make the research contributions more effective and can be deployed in industrial solutions.

- **Datasets:** Almost all the work collected their own dataset using the PhishTank website which is a crowd-sourced phishing verification system where users submit links to suspected phishing pages and other users "vote" if the reported page is a phishing page or not. It is relatively easy to build a sizeable and good quality dataset using PhishTank API that include the links to the phishing web page as well as the target. As a result, the majority of the work used their own dataset. A limited amount of work used data collected from their own institutions (i.e. by harvesting URLs from organizational emails) or crawled social media plot to circumvent the short lifetime issues of phishing web pages. For example, in PhishTank links, there is a delay between someone reporting the link and others voting for it as possible phishing. By the time this flagging happens, the phishing site might be already down, eliminating the possibility of obtaining a screenshot.

- **Performance Comparison:** As a consequence of different work using different datasets, it is not possible to quantitatively compare the performance methods. Analogous to the established datasets for various computer vision tasks such as MNIST [72], CIFAR 10 & 100 [73], ImageNet [74], and COCO [75], a common web page screenshot dataset would enable more re-producible research in this area. We noticed that, there are many new phishing website databases released between 2018 to 2019 such as Phish-IRIS dataset [76], anti-phishing dataset by Universiti Malaysia Sarawak [77], and CIRCL dataset [78] that have the potential to become standard datasets for phishing detection.

- **Deep Neural Networks (DNNs):** We found that several recent work started using deep learning methods on web page screenshots to compute visual similarity, which is not surprising at all given the huge success of deep learning methods in computer vision. However, the current set of work use basic blocks of deep learning, such as MLPs, autoencoders, and CNNs. However, web page images are very diverse and rich in content, and the phishing tactics are becoming more and more advanced. As such more advanced deep learning methods for salient object detection and segmentation [79]–[81], image comparison [82, 83], and neural image hashing [84]–[86] can be leveraged in this area to build more efficient and accurate phishing detection solutions.

TABLE I: Summary of Visual Similarity Based Phishing Detection Approaches

Research Work	Computer Vision Methods	Input Image Structure	Dataset Size (Phishing/ Legitimate)	Summary	Performance Metrics
Web Page Feature Representation Approaches					
Liu et al. [38]	Visual Similarity Assessment based on Key Regions, Page Layouts, and Styles	N/A	8/320	Compares the suspect page to the legitimate page by measuring block-level similarity, layout similarity, and overall style similarity.	Recall - 100% Precision - 91.2%
Medvet et al. [42]	Visual Similarity based on 2D Haar Wavelet Transformation and Colour Histograms	N/A	41/161	Extracts three visual page features which are texts, embedded images, and overall images of the from the snapshot of a web page and then computes the distance between the colour histograms, the positions and the 2D Haar wavelet transformations to gain a similarity score.	Recall - 92.6% False Negative Rate- 7.4% False Positive Rate - 0%
Lam et al. [52]	Layout Block Matching	N/A	6,750/ 312	Converts both the suspicious page and the legitimate page to black-and-white images, divides images into different layout blocks, and then performs layout matching between them to get their similarity score.	Accuracy - 99.6% False Positive Rate - 0.028%
Fu et al. [45]	EMD-based Visual Similarity calculated on Colour and Coordinate Features	100*100	9/10,272	Converts the screenshot of web pages into low resolution images with normalised size, and then uses EMD to assess the visual similarity based on colour and coordinate features.	Recall - 88.8% Precision - 99.7%
Zhang et al. [47]	EMD-based Visual Similarity calculated on Colour and Coordinate Features	100*100	5,711/ 10,272	Transforms web pages to images, generates image signatures and measures the distance of images by EMD, and finally compares the visual similarity with the threshold.	False Negative Rate - 0.88% Classification ratio - 99.9%
Chen et al. [48]	Contrast Context Histogram (CCH)	N/A	-/300	Captures a webpage screenshot and extracts its keypoint feature, then uses CCH descriptors to perform image matching and finally classifies a page based on page similarity degree.	Accuracy - 95 98% False Positive Rate - 1% False Negative Rate - 1%
Corona et al. [51]	Histogram of Oriented Gradients (HOGs) and Colour Histograms	N/A	1,012/ 4,499	Captures a screenshot of a browser, then extracts visual features based on HOGs and colour histograms, and finally performed classification on the similarity between two histograms based on SVM.	Recall - 99%, False Positive Rate- 1% False Negative Rate - 1%
Dunlop et al. [53]	OCR and Google Search API	1,200*400	100/100	Captures a page screenshot, and utilises OCR software to extract its text information, then searches the texts on Google to retrieve results.	Recall - 98% False Positive Rate - 0%
Tian et al. [55]	OCR and Image Hash	N/A	1,224/ 65,7663	Utilises OCR to extract texts from the screenshots of web pages as OCR features, and applies the Image hash (a fuzzy hashing function) to compare the page screenshots with the hash of the real pages to assess the visual similarity.	False Positive Rate - 0.5%, False Negative Rate - 0.05%
Hara et al. [56]	ImgSeek (Wavelet Transform)	N/A	2,262/ 521	Extracts the image shown on the target web page, then uses ImgSeek to search for similar images in the database containing legitimate sites, phishing sites, and undeterminated sites, and checkes the output similarity to perform classification.	Recall - 82.6% False Positive Rate - 18.0%
Logo Feature Representation Approaches					
Wang et al. [57]	SIFT	N/A	219/230	Captures a screenshot, then uses SIFT to match the screenshot against a logo database, and finally checks for keypress and the fact whether the DNS matches with the legitimate brand holder to decide if it is phish.	Recall - 90% False Positive Rate - 3%
Zhou et al. [60]	SURF and EMD	100*100	2,129/ 1,372	Combines and Compares logo detection and web page snapshots using SURF and homography transformation for possible detection of phishing websites.	Recall- 90% False Positive 1.46%
Afroz et al. [63]	SIFT, Fuzzy Hashing, and OCR	N/A	1,000/ 200	Uses SIFT to extract logo features form a web page, and performs image matching based on fuzzy hashing, OCR and SIFT algorithms.	Recall- 90.2% False Positive - 0.5%
Fatt et al. [54]	Google search-by-image API	N/A	500/500	Extracts the favicon of website, then utilises Google search-by-image API to search for relative information, and then determines if the suspect page is a phishing page based on search results.	False Positive Rate - 5.4%, Recall - 97.2%
Hashing Approaches					
White et al. [65]	Perceptual Hashing	N/A	-	Takes a page screenshot, then creates a hash of the image, and finally calculates the Hamming distance between the image hashes of the suspect pages and authentic pages.	
Tian et al. [55]	OCR and Image Hashing	N/A	1,224/ 657,663	Utilises OCR to extract texts from the screenshots of web pages as OCR features, and applies the Image hash (a fuzzy hashing function) to compare the page screenshots with the hash of the real pages to assess the visual similarity.	False Positive Rate - 0.5%, False Negative Rate - 0.05%
Fotiou et al. [66]	Perceptual Hashing	N/A	100/100	Applies perceptual hashing on web page screenshots, then calculates the Hamming distance between the image hashes of the suspect pages and authentic pages.	False Positive Rate - 1%, Recall - 81%
Neural Network Approaches					
Borgolte et al. [68]	Stacked Auto Encoder and MLP	screenshots(1600*900) and window(160*160)	925,817/ 255,490	Captures the snapshots of websites, extracts a specific window from each snapshot, and then uses a sliding window approach on the snapshot to perform defacement detection based on deep neural network.	Recall - 97.42% - 98.82% False Positive Rate - 0.55% - 1.53%
Adebowale et al. [69]	SIFT, Classical Neural Network with Fuzzy Logic	N/A	6,843/ 6,157	Extracts features of images, frames and text, and uses the hybrid features as input to a fuzzy neural network to classify the web page.	Recall - 98.3% False Positive Rate - 3%
Abdelnabi et al. [70]	Triplet CNN	screenshots(224*224)	1195/683	Utilises triplet networks to train all screenshots with random sampling, then fine-tunes the weights of model by iterative training on hard examples, and finally (based on embeddings) performs classification by checking the distance with a threshold distance.	ROC Area-0.9879 Recall -95.81% False Positive Rate - 6.88%

IV. VISUAL SIMILARITY BASED MALWARE DETECTION

More recently, computer vision methods help in alleviating the limitations of traditional malware detection systems such as detection of obfuscated malware or a computationally less expensive detection based on malware behaviour. Computer vision methods provide visual aids in detecting unknown malware to alert on anomalous behaviour patterns quickly. Visual representations of malware patterns also have the advantage of offering a summarised picture of possible attacks, thus speeding up the detection or classification of malware samples. For instance, if the frequency of API calls or usage of system resources is too high in an image representation, then a sample is classified as malware. Moreover, the fact that recent malware such as DuQu [87] and Hammertoss [88] exploited the structure of image files (JPEG and PNG) either to communicate with the Command and Control (C&C) server or to follow commands hidden under images files on a victim's machine. Finally, some malware display a familiar, trusted icon to users to trick them into clicking the malicious program, also motivates malware analysts to utilise computer vision techniques for possible malware detection.

In this section, we survey the existing computer vision-based malware detection or classification approaches and categorise them based on their method of detection, according to the categorization we introduced in Section II. We found a significant amount of work fitting into the categories of *i) image representation-based methods, ii) image feature-based methods, iii) image hashing-based methods*, and *iv) neural network-based methods.*

A. Image representation-based methods

Image representation-based methods visualise behavioural activities of malware as images and then compare them against the benign/normal activities of a system or a network. Malware behaviour refers to what the malware does, exhibits, or causes to its environment during live execution. Such behaviour

aspects could be represented through monitoring changes to operating system resources during malware execution, capturing malware's API call sequence, malware's I/O request packets (IRP), or malware's network activity. It is possible to represent such behaviours as visual fingerprints (e.g. RGB colours, matrix, or intensity levels). For example, a visual fingerprint can be the frequency of calls to each API in RGB colours, where red denotes the most frequently called API (e.g. HTTP outgoing packet send request) and the green denotes the least called API (e.g. invoke anti-virus).

Image representation-based methods for malware detection have several advantages. First, the image representation-based analysis helps in detecting obfuscated malware. For instance, a visual fingerprint of an obfuscated malware can be generated by executing it in a sandbox environment and then tracing the functionalities performed by malware. If the fingerprint matches with existing malware fingerprints, then it is considered malcious, otherwise benign. Similarly, it is possible to detect obfuscated malware, which has an icon disguised as a benign app, by using properties such as RGB pixels colours. In this case, similarity analysis over the sets of images can provide insights into malware sample relationships. Second, the behavioural analysis of malware can help to detect new and unknown malware types. For example, if a new program (malware) communicates very frequently with C&C server, then thread graph can be used to view a temporal behaviour of a new program, showing more often connections to a specific IP address. Similarly, behavioural to colour mapping, may also be used to identify potentially malicious programs by defining high-intensity colours (e.g. dark tones) to system-related activities such as system shut down, system files deletion, copy of system files, sending password files to a remote server etc. Finally, an image-based analysis may help analysts in highlighting current trends (e.g. icons, logos, hidden images, frequently used system functions, sequence/series of functions layout), in which malware authors are attempting to trick users into running malicious programs.

One of the early work in visualizing and analyzing malware behaviour was by Trinius et al. [89]. The proposed work first applied an abstraction method to summarise the reports of the sandbox. Then classical visualization techniques, treemaps and thread graphs are used to visualise the percentages of API calls as well as malware behaviour. The abstraction method groups API calls performing similar functionalities (e.g. all API calls related to file system activity groups to one cluster). These API calls order arguments according to their significance in performing functionality. The purpose of Treemapping is to help in viewing the distribution of operations, e.g. whether the main task lies in the area of network interaction or interaction with other processes. On the other hand, thread graphs help in visualizing the temporal behaviour of a sample.

To validate the effectiveness of this work, a large database of malware samples was collected using honeypots. Then the behaviour of these samples was analysed by executing them for two minutes in a controlled instrumented environment (sandbox). During execution, sandbox records all the system-level activities (e.g. loaded system libraries, outgoing and incoming network connections, accessed or manipulated registry

keys, and reports in XML format). To expand the analysis, the authors also executed the associated data files such as word or adobe acrobat files in a sandbox environment. After the execution, the abstraction method transformed XML reports to a short format; for example, by grouping together API calls having similar functionalities. Treemaps and thread graphs then helped in visualising the abstracted reports. The work claims to be effective against the classification of unknown samples to malware families by analysing over a set of 2,000 malware samples.

The advantage of this work is that it partly circumvents the problems of code obfuscation techniques such as packers or crypters. A packed malware is a modified form of a malware that is compressed using a compression program and is decompressed at runtime using a stub appended with a malware file. Therefore, when a packed malware executes in a sandbox environment, it unpacks itself and then it is possible to study its behaviour. Moreover, this method is capable of detecting malware hidden in data files such as adobe acrobat or Microsoft word files. However, the approach is not fully automated and require human intervention to compare thread graphs or treemaps of benign and malicious files. Another drawback of this approach is the fact that malware may detect the presence of the analysis environment (sandbox) and then behave differently. This approach is not capable of detecting such behaviour.

Following the work of [89], Han et al. [90] also proposed a novel method to detect and classify malware by constructing image matrices from the opcode sequence of executable files. The purpose of the image matrix is to detect malware features conveniently and to calculate similarities between different malware faster than other visualization methods. First, OllyDbg or IDA Pro dissembles binary files to extract binary information, and then opcode sequence is divided into blocks. These blocks are assigned (x,y) coordinates and RGB colours using *Locality Sensitive Hash Function (SimHash)*. The RGB values are then plotted on their corresponding coordinates in a block of dimension 8×8 to get an image matrix. After forming a matrix for the whole image, "selective area matching" is performed using vector angular-based distance measure algorithm [91] to calculate the similarities between image matrices of different images. The image matrices sizes are set to 256×256 pixels. The proposed method was applied to 10 different malware families, and the results were recorded. Results indicate that the image matrices from the same malware family have 0.95 similarity on average, but, those from different families have 0.325 similarity on average. Though this work shows promising results in classifying unknown samples to malware families, it was tested on a small number of malware samples. Also, it is unclear whether the authors collected malware samples or use some public dataset.

Similar to the above approaches, [92] also proposed a malware visual analysis method that generates RGB-coloured pixels on image matrices using the opcode sequences. However, to reduce the computational overheads, authors extracted the opcode sequences only from the blocks that are related to staple behaviours such as functions and API calls. The angular-based distance measurement algorithm was used to calculate

the similarities between images matrices.

In the paper [92], authors performed two sets of experiments, static and dynamic analysis. In a static analysis experiment, 290 malware samples from 16 malware families, were tested. This experiment did not contain any packed or obfuscated malware samples as the purpose was to extract the basic blocks through static analysis using a disassembler. Based on staple behaviour, the basic blocks helped in selecting major blocks, which were used to generate image matrices using opcode sequences. The image matrices were compared against each other. The dynamic analysis experiment consists of 560 malware samples from 14 malware families in which the packed and non-packed malware samples coexist. These samples generated the dynamic traces, and then basic blocks were extracted through repetition filtered technique. Afterwards, the major blocks relating to suspicious behaviours and functions were selected. The proposed method was then applied to these major blocks to generate image matrices and to analyze similarities. For both experiments, the sizes of the generated image matrices were 256×256 pixels. The static analysis shows an accuracy of 98.96% whereas, dynamic analysis shows an accuracy of 97.32%, receptively. The above two approaches [90, 92] suffer from the drawback of disassembly of malware files, which make them ineffective for packers and encrypted malware. Hence, the characteristics of image representation based methods to analyze packed or encrypted malware, are not fully utilised by these methods. Authors also did not discuss any results or an extension of their work for packed and encrypted malware.

In another scheme, Han et. al [93] converted binary files into images and entropy graphs to detect and classify malware variants. The method first converts a binary file into a bitmap image using the "Bitmap Image Converter". It then calculates the entropy values of a greyscale bitmap image to generate the entropy graph. As the final step, the entropy graph similarity calculation algorithm detected unknown samples to a malware family. The method was tested on two sets of data collected from VX Heavens [94]. In the first dataset (consisted of 24 benign binaries and 27 malware binary files from 8 families) determined the threshold to minimise false positives. After setting the threshold to 0.75, the second dataset was tested with a total of 1,000 malware binary files from 50 families, and achieved an average accuracy of 97.9%. The collected malware binary files were Backdoors, Trojans, Viruses, and Worms that are executable files in the Windows operating system.

This work has an advantage over previously proposed methods in terms of detection timing. It has less computational overhead than texture analysis methods, such as GIST [95]. The time to calculate the similarities of all the benign and malware binary files was ranged from 35.29ms to 1.39ms. However, the problem with this work is that it cannot be applied to packed malware because the entropy of packed malware usually are very high and cannot indicate any specific pattern. Moreover, all the three approaches from Han et al. [90, 92, 93] have a limitation of inspecting only a particular aspect of malware patterns and behaviours and did not focus on analyzing unknown processes, system calls and their relations.

Shaid et al. [96] showed another method to visualise and classify malware files. Authors first created a behaviour-to-colour mapping and then generated the behaviour-based image by converting each behaviour of malware to a colour. In order to obtain behavioural data, an API call monitoring utility executed the malware in real-time and captured all user-mode API calls made by a malware sample. API call monitoring utility captures the information such as accessed system files, network connections and protocols used, modified memory content of a remote process. The reason to opt for user-mode API calls is its stability as these calls are rarely changed as compared to the lower-level system call, which differs a lot between different versions of the same operating system. Besides being stable, capturing user-mode API is also crucial in speeding up the process of understanding the workings of malware since user-mode API reflects the exact nature of a particular malware. Afterwards, API calls are grouped and sorted according to their level of maliciousness. Once arranged, colours are assigned to each group such that hot colours represent malicious activity while cold colours represent benign activity. Once the API-to-colour map is ready, a behaviour image of 64×4 pixels size was generated from each malware sample. The experiment was run on 1,101 malware samples belonging to 12 different families and showed an accuracy of between 95.92% and 98.92%.

This work did not focus on detecting malware; instead, their only focus is to classify unknown malware samples to malware families. Moreover, the authors did not explain how they collected malware dataset. The performance of this method was also not discussed as we assume that executing malware on a Virtual Machine and then performing colour mapping might require more time than other detection techniques.

A very recent work [97] proposed SAD THUG (Structural Anomaly Detection for Transmissions of High-value information Using Graphics) to detect malware hidden under the structure of image files. The method consists of two main phases, a training phase, for building a formal model, and a classification phase, to check whether files correspond to that model. Each image file decomposes into the sequence of segments of different length. Most segments start with a header or byte sequence that indicates their type. Often, their payload starts with another header that is needed to interpret the segments payload correctly. This inner header has a significant impact on how the segment is interpreted; for instance, these headers represent the purpose of a segment or even whether the segment should be ignored completely by most decoders. Hence, segments with different inner headers are written and read for different purposes and should thus be assigned different subtypes. Therefore, this method further defines the subtypes of these inner headers to detect malicious purposes hidden under the structure of JPEG or PNG files. Once subtypes are defined, classifiers are trained with benign files that contain the start of a file, end of the file, start of each segment, end of each segment, length of each segment, and subtype of each segment.

To classify the unknown file, the method used two parameters α, and τ, where τ determines the number of times a sequence of segment appears during the training phase and α indicates the reasonable length of each segment. For instance,

τ of 1 means any sequence of segments that have been observed once in training phase are benign, and τ set to high values makes classifier more restrictive. During classification, first, the transition of segments from one to another is checked, if that segment has not been observed during the training phase or observed to follow the previous segment, then the file is considered to be anomalous. After passing this stage, the length of segments is checked if it is under a defined range. If all checks are passed, the file is considered to be benign. The approach used a total of 270,000 JPEG and 33,000 PNG files downloaded from Alexa top 25 popular websites and the third-parties found during the recursively crawling. For JPEG files, this approach achieved a mean true positive ratio of 99.24% while 99.31% for PNG files. On the other hand, the mean true negative ratio was 99.32% and 98.88% for JPEG and PNG, respectively.

While this work is interesting and has shown promising results, it can be further improved. First, the method currently supports only a limited number of image files (i.e. JPEG and PNG) and could be evaded by any other file types that it does not support. Second, its effectiveness is dependent on the training dataset. A training dataset that is not representative enough for all types of benign data observed in the classification phase may increase the false positive ratio.

To detect the repackaged apps and phishing malware, Sun et al. [98] proposed the DroidEagle tool that compares the visual characteristics and components in the XML layouts of Android apps. The tool leverages the app repository analysis, named as RepoEagle, and host-based detection, named as HostEagle. RepoEagle performs large scale detection on app repositories, by extracting the layout tree as the fingerprint, which can indicate a visual structure of an app. The similarity between two layout trees is then measured using Layout Edit Distance (LED). LED measures the similarity between two layout trees by counting the minimum number of operations required to transform from one layout tree to another tree. The operations in layout tree transformation are node deletion, node insertion and node substitution. A small value of LED implies that the two layout trees are very similar, indicating that apps may be impersonated. On the other hand, HostEagle is a lightweight mobile app which conducts a real-time detection on mobile phones. This module extracts a layout tree and calculates the hash value of the layout tree as the visual signature. By comparing the signature of the app in the remote server, it can decide whether this app is a malware or not. Authors crawled and collected 100,096 apps from the Google market and various third-party markets. Results show that RepoEagle found 1,298 visually similar apps from a repository whereas, HostEagle is capable of determining whether a downloaded mobile app is a repackaged app or a phishing malware.

Despite promising results, this method suffers from several limitations. First, all visually similar apps may not be necessarily similar in XML layouts, and it is necessary to consider the similarities in images. Second, app developers are starting to use code encryption methods, thus accessing codes and layout files may not always possible. Third, dependence on specific aspects related to one operating system will not allow making comparisons between heterogeneous app markets, and in such situations, only metadata and image similarity are meaningful.

B. Image feature-based methods

While image representation-based methods are capable of detecting obfuscated/encrypted malware and can also detect/classify malware and its variants through behavioural information, they are less efficient and robust for a large malware dataset [99]. Consider the dynamic analysis of a malware binary, in which binary is executed, provided time to run its program logic, converted into an image based on call sequences, and finally the image is compared for possible malware detection. This whole process takes a few minutes (e.g. 3- 5 minutes) for a single binary sample. Even if the procedure is streamlined for 30 seconds, it will take years to process millions of malware samples such as Symantec Malware Corpus. In this regard, feature-based methods are highly accurate and efficient to detect or classify unknown sample, as compared to image-representation methods [95].

In feature-based methods, a malware sample first converts into an image, and then image-related features are extracted. A given malware sample is generally read as a vector of 8 or 16-bit unsigned integers and then organised into a 2D array. The 2D array can then be visualised as a greyscale image in the range [0,255] (0: black, 255: white). After that, image features such as SIFT, SURF, and Gabor are extracted and fed into a detection system (either a similarity checker or a machine learning classifier) for possible classification. We evaluate each method under this category based on similar aspects mentioned in Section III, *i.e. input data type, computer vision methods used, datasets, and performance*.

Nataraj et al. [100] is one of the early work fitting into this category. The proposed method assumed that greyscale image of the same malware family has similar layout and texture. Authors extracted 320 GIST features and used k-nearest neighbours with Euclidean distance on top of those features to classify unknown malware. Authors evaluated the performance of the solution using malware executables submitted to Anubis analysis system [101] and showed an average accuracy of 98% over a database of 9,458 samples and 25 malware families.

One advantage of this work is that it can detect packed malware as it is based on the assumption that the images of the same malware family when packed with the same packer, appear similar. Authors achieved a classification accuracy of 98.08% when they packed malware of the same family and applied the proposed method on them. Authors also compared their performance with existing methods based on bi-gram executions. They extracted bi-grams distributions from the raw dataset and obtained a classification accuracy of 98%, which is same as their approach. However, the time taken to extract bi-grams are 5 seconds, and the time to classify a malware sample is 56 seconds. In contrast, the proposed approach computed GIST features in 54 milliseconds and classified a sample in 1.4 seconds. The high efficiency of the proposed method is partly explained by the fact that the feature vector length used to characterise a malware image is about 320 whereas the distribution based analysis using the bi-grams needed about 65K elements. However, a major limitation of this approach is

that visual representation only considers images embedded in a resource section of an executable. This limitation allows an attacker to easily bypass detection either by relocating sections in a binary or creating less obvious patterns.

In a later extension [95], Nataraj et al. compared the performance of the initial work [100], named binary-texture analysis, with contemporary dynamic techniques. While the classification accuracy of the binary-texture analysis is comparable with dynamic analysis methods, however, it is 4,000 times faster than dynamic techniques. In addition, the texture-based approach is resilient against packed malware samples. They confirmed that the contemporary packing strategies perform a monotonic transformation which fails to conceal common structures present in the malware binaries, thus easily detectable. Authors used two dynamic analysis techniques. The first one is a commonly used approach called system-call level monitoring which generates a report of all system calls and their arguments performed by malware binary. The second approach is a forensic snapshot which compares features before and after a malware infection.

Authors conducted three different experiments; the first experiment was conducted on 393 malware binaries and showed the classification accuracy of 98% and 95% for dynamic and binary texture analysis, respectively. In the second set of experiments, Malheur dataset [102] was used for comparison that comprises of Reference dataset with 3,131 malware binaries from 24 unique malware families, and an application dataset of roughly 33K binaries that range from malicious, to unknown, to known benign. This experiment showed the classification accuracy of binary texture analysis ranging from 86% to 97%. In the third experiment, VX heavens dataset was used that consists of over 63K malicious binaries from 531 unique families. The binary-texture analysis, in this case, achieved 72% accuracy. The reason for this high accuracy is that packed malware variants that belong to the same family exhibit visual similarity with one another, which can be detectable through binary-texture analysis method. Another reason is that a single malware family is packed with one or two packing software and not more. Finally, there is a fraction of the malware that is unpacked. Thus, all the reasons mentioned above aid binary-texture analysis in detecting malware with high accuracy. However, the limitation is that this method is vulnerable to knowledgeable adversaries who can obfuscate their malicious code to defeat texture analysis.

Kancherla et al. [103] proposed a malware detection method that first converts a binary executable file to an 8-bit 1-dimensional vector and then the 8-bit value converts to the intensity of the pixel. Next, this single-dimensional vector converts into a 2-dimensional vector, known as *byplot*. This byplot (image) has a fixed width based on the size of the file. From this image, the authors extracted 534 features that included 512 Gabor-based, 16 wavelet-based, and six intensity-based features. The purpose of using wavelet and intensity-based features is to examine if light-weight features offer comparable classification accuracy against computationally expensive Gabor-based features. Authors performed three different sets of experiments using only Gabor features, using only intensity and wavelets, and using all features. Results

showed that light-weight intensity and Wavelet-based features produced slightly better classification accuracy as compared to Gabor based features. The resulting accuracy was around 94.32% for the light-weight features, whereas Gabor features showed 93.23% accuracy. These experiments were conducted on offensive computing database containing 25,000 malware samples and from other sources like source forge to collect 12,000 benign samples.

This work is an improvement to [100] as it shows that light-weight features such as intensity-based features and wavelet-based features offer better classification performance when compared to computationally expensive Gabor based features. Few other works have very similar approaches [104]–[106] with the only difference being the final classifier and the dataset on which the solution is evaluated.

Ahmadi et al. [107] emphasised on the convenience to select simple features to design light and efficient malware-detection system. Instead of extracting features directly from a Portable Executable (PE) file, authors calculated a set of features from different categories in hex and assembly view of PE. The feature categories in hex view are n-gram of bytes, the byte-level entropy vector, metadata (e.g. size of the file, address of the first bytes sequence), the histogram of the length of strings, image-features (Haralick features, or the Local Binary Patterns features) whereas, the categories in assembly view are frequency of some specific operation codes, frequency of API calls, frequency of use of the registers, and characteristics of sections in a file. With each category, there is a range of features; for instance, in image-based features, there are a total of 160 features. In order to combine feature categories such that accuracy in classifying/detecting malware increases (i.e. best subset of relevant features), forward step-wise feature selection was used. In forward subset selection, initially, a model contains no features and then gradually augments the feature set by adding more features to the model, one by one. In order to check the performance of each feature category and combination of categories, XGBoost classifier added features from each category into the model one by one. Authors used Microsoft released a dataset containing 21,741 malware samples, where 10,868 samples are used for training, and the rest are used for testing. Results showed that image-based category has a classification accuracy of 95.5% however, with the combination of other features, this accuracy reaches to 99.40%.

The subset of features utilised in [107] is still considered to be expensive and have chances of reducing the efficiency of a malware detection system. Therefore, Zhao et al. [108] proposed an approach that used malicious code homology analysis method (based on texture fingerprint clustering) for possible detection of malware. The approach first extract features from a binary program, create texture fingerprint from the features and then use a clustering algorithm to group the texture fingerprint of malware into families. Finally, the homology analysis matched the unknown sample with one of the templates of malware families stored in the database. To extract features, B2M algorithm converted unknown samples to an image, and then an improved SIFT algorithm was applied to create a multi-scale space. Once space is created, the feature

descriptors of a sample is extracted and fed into DBSCAN for clustering into family libraries. The homology analysis is the same as image recognition, where an unknown sample is compared with the family library to detect its maliciousness and classify into a specific library. The method was tested on 1,789 samples from 11 different malware families, collected from Bigdata Innovators Gathering (BIG) 2015 [109]. Results show the average accuracy of 77.6% against different malware families.

A recent approach by Masouleh et al. [110] classifies windows malware (PE files) using icon-based features. The method used three types of features generated from the original icon pixels: i) Summary statistics features, ii) Histogram of oriented gradient (HOG) features, and iii) Autoencoder (AE) features. Summary statistics features are the mean and standard deviation computed for all the pixels and across all three channels, on each channel separately, and by splitting the icon image into nine different sections. HOG features are captured by sliding a small window over the image and then computing the gradient of the image within that window. To extract AE features, neural-network is trained with a large set of icons. In total, 1,114 features were extracted. Next, hierarchical density-based spatial clustering (HDBSCAN) and K-Nearest Neighbor algorithm were used to cluster the icons with the purpose to reduce the data dimensionality. In order to test the efficacy of the method, a balanced sample of publicly available PE files was obtained from VirusTotal with 1,138 benign and 1,138 malware files. The dataset is split into 80% testing and 20% training and then evaluated using Logistic Regression (L1, L2) and SVM classifiers. Logistic Regression (L1) shows better performance than the other two classifiers with an accuracy of 84.4%.

This work [110] may also fit into the category of neural network-based methods. Still, we have included it here because the neural networks do not play a major role in the methodology to detect malware. The neural network-based classifiers were primarily used to validate the efficacy of the proposed methodology. This work also suffers from several limitations. First, it is not validated on a large number of malware PE samples, and therefore its efficiency cannot be properly determined. Second, extraction of 1,114 features is quite a cumbersome task and time-consuming. Authors did not show the performance of the method with respect to time. Third, the technique is not compared against previously proposed icon-based malware detection technique [111]. Finally, the classification accuracy is quite low as compared to other techniques that were tested on Microsoft Challenge dataset. For instance, Ahmadi et al. [107] showed 95.5% of classification accuracy which is 10% higher than [110]. Furthermore, approaches in neural network-based methods also used Microsoft dataset and achieved accuracy ranging from 95% to 99.8% [112]–[117].

In [118], authors proposed a text-based and content-based image retrieval method to detect camouflaged apps on mobile market places. For each app, four external features are extracted; namely app name, description, icon and screenshot. App name and descriptions are handled by text retrieval systems, while app icon and screenshots are handled by image retrieval system. Authors used unofficial Google API to crawl app info, such as id, name, developer, rating as well as application description, icon and screenshots from the official Android market. Total of 30,625 applications was crawled for the experiment. After crawling, each app was assigned four indexes which are name index, description index, icon index and screenshot index. For the image indexing, colour correlogram algorithm was used. The last step is to detect camouflaged apps by querying the indexed database. For each queried app, the method retrieved apps which have similar user interfaces but are from different developers. Both text-based and content-based systems are used to calculate the cosine similarity scores between the query app and a set of indexed apps. Out of 30,625, this method identified 477 camouflaged apps (approx. 1.56%); however, the manual review showed that 44 apps were falsy flagged camouflaged with 9.22% of the false positive rate.

C. Image hashing-based methods

The work discussed in above Section III-A showed high accuracy for the known malware families; however, the feature extraction and matching (similarity calculation) process still requires a time investment per malware that does not scale well to the high volume of malware samples at run-time. Consider the extraction and similarity calculation of 512 Gabor-based features [103], 1,114 AE features [110], or even 160 Haralick features [107] that require sufficient time to match with the known malware stored in a database, thus highlighting the need of an efficient feature matching algorithms, as highlighted by [119]. A body of work addresses this problem by proposing *image-hashing based methods*, which we discuss next in this subsection.

To improve the efficiency of malware classification, Xiaofang et al. [119] proposed a novel distance (similarity) metrics based on the locality-sensitive hashing (LSH) schemes. The approach classifies malware by calculating the distance metrics on the fingerprints of malware content. First, a binary image is obtained from a malware sample, and then a feature fingerprint of 64*64 dimension is computed using SURF algorithm. Afterwards, a fast fingerprint matching is performed with the LSH to get the topmost visually (structurally) similar variants. Here a euclidean distance is used as a metric to get a matched malware set. The proposed scheme is applied to 25 malware families on a total of 8,410 malware samples. The precision, recall and F-Score is around 80% to 90% for varying dataset sizes. Though this work claims to improve the efficiency of computer-vision based malware detection systems by using LSH scheme, the authors did not provide any results of performance comparison in the paper to prove the claim.

Another body of work used image hashing-based methods to detect mobile malware and greyware. Malisa et al. [120] detected two types of mobile impersonation attacks; *repackaging* and *phishing*. Authors proposed visual impersonation detection system that runs mobile app inside an emulator, explores its user interface dynamically, and extracts app screenshots using GUI crawling techniques. The extracted screenshots are then compared with other apps to detect the case of possible impersonation. If the comparison value is above a

threshold (i.e. near match or exact match), apps are labelled as impersonated apps. The Local Sensitive Hashing (LSH) technique is used to check the similarity between screenshots. Unlike other hashes, LSH maps similar data to the same hash code with high probability, maximizing the probability of a collision for similar inputs. The method was deployed on the Google Cloud Compute platform and tested on 150,000 mobile apps with 4.3 million screenshots. Results indicate the detection of 43,904 impersonating apps with 15% FPR.

The proposed system is capable of detecting impersonated apps that are using evasion techniques such as obfuscation, as the detection purely relies on the visual appearance of the examined applications. Another advantage of this system is that it can efficiently analyze a large number of Android apps before their deployment at Play Store. Moreover, the detection system neither requires human input nor application-specific knowledge as it supports the automated analysis of apps at a larger scale. Nonetheless, this method inherits all the resource limitations of dynamic analysis as each app runs in the emulator.

Besides mobile malware, greyware is another emerging threat to mobile apps. For instance, ad fraud apps could be classified as greyware because of the reason that such apps contain annoying, undesirable or undisclosed behaviours that cannot be classified as malware. In this regard, Andow et al. [111] first identified nine categories of greyware based on installation and runtime behaviour. Later, authors developed lightweight heuristics for triage of the seven categories of greyware on Google Play that leverages text analytics and static program analysis. Among all the greyware categories, imposter apps are detected using computer vision methods. To detect imposters, first, the similarity of apps titles and developers names are computed and then *fuzzy hashing* is used to score the similarity of app icons.

Authors retrieved a list of APKs and app metadata such as user reviews, description, and developer names from the PlayDrone Archive [121] and separate them in four different datasets. The dataset to detect imposters contained 50 most popular apps for each app category in both free and paid subcategories, which a total of 2,500 apps. Results indicate that manual review of 2,500 apps is reduced to 977 with exact title matches, of which 22 app pairs have similar icons. Though fuzzy hashing is suitable for the detection of slightly modified visual images; however, the efficient detection of impersonated apps requires techniques other than hashing (e.g. feature comparisons or deep learning classification methods). Moreover, the proposed imposter detection method was tested only on a small portion of apps, i.e. 2,500; this number is quite low as compared to millions of apps available at PlayStore.

Alike [110, 111], Jiao et al. [122] also proposed a method called ImageStruct that detects repackaged apps using similarity measurement. The method extracted features from app images using pHash and compared the hash codes to detect similar apps. When the similarity of two apps with different signatures exceeds a certain threshold, one of the apps was considered as repackaged. Similarly, Long et al. [123] visualised similarity between malware samples based on their embedded graphical assets. First, images were extracted from

malware using *wrestool* [124] and converted to ,1024 dimensional vector. Then, the vectors were stored using Fast Library for Approximate Nearest Neighbor (FLANN) index. Finally, the top n visually similar images were retrieved by the random tree method using $L2$ distances. The authors claimed that the method has a recall of 83% and precision of 85%. Since the method uses average hashing and FLANN indexing, it is fast; however, the performance is highly depending on the threshold values.

D. Neural network-based methods

Similar to the motivation of moving into neural networks as mentioned in Section III, in malware detection also multiple works proposed to use neural networks, especially deep learning based networks to eliminate the need of manual feature curation. In image-based neural network approaches for malware detection, the malware binaries are converted into greyscale image representation, and deep learning architectures are employed to learn the complex features (image patterns). Most commonly employed deep learning architectures are convolutional neural networks and LSTM (long short-term memory) recurrent neural networks.

Zhang et al. [125] proposed a CNN to compare opcode images of a target with the opcode images generated by known malware sample codes to detect binary malware variants. First, the binary executable is disassembled, and an opcode profile for the binary is built by representing it as a list of opcode sequences of length two and the corresponding frequency. This sequence is converted to an image where each pixel value represents the information gain of the two corresponding opcodes. This image is further enhanced using methods in the likes of histogram normalization, erosion, and dilation. These images are then used to train a simple CNN with two sets of convolution and pooling layers followed by one fully connected layer.

Authors evaluated the classifier using a dataset of 9,168 malware samples belonging to 10 families and 8,640 benign samples. The results showed that image-based malware variant detector could achieve over 90% true positive rate and true negative rate outperforms traditional machine learning-based classifiers such as naive Bayes and k nearest neighbours especially when there is less number of training samples.

Kim et al [112, 113, 126] conducted a series of work on using generative adversarial networks (GANs) for zero-day malware detection. The idea stems from the GANs in computer vision that has found multiple applications in image generation, denoising, and resolution improvements. The key idea in related to zero-day detection is to use the *generator function* of the GAN to generate new malware samples and use the discriminator function to predict whether the given sample is a malware or not. The intuition behind this process is that as the generator builds new malware representations and the discriminator has been trained to identify them, when a new malware that was not part of the training set emerges, the discriminator will be able to detect it.

In this regard, the first work by Kim et al. [126] proposed transferred generative adversarial network (tGAN) that uses a

pre-trained auto-encoder as the generator and a conventional CNN as the discriminator. The auto-encoder learns the characteristics of data using transfer learning method and generates malware images similar to real malware images. On the contrary, the CNN based discriminator tries to distinguish between real and fake malware images. However, this approach suffers from the problem of instability in converging generator and discriminator to Nash equilibrium during the training process. Therefore, Kim et al. extended the idea and proposed two different approaches. In the first approach, Kim et al. [113] proposed transferred deep-convolutional generative adversarial network (tDCGAN) that uses a deep auto-encoder (DAE) as a generator and deep convolutional GAN (DCGAN) as a discriminator. The advantage of using DCGAN is that it produces higher-quality images than original GAN, and the training system is relatively stable. Akin to above method, Kim et al. [112] also proposed another version of GANS called latent semantic controlling generative adversarial networks (LSC-GAN), that generates malware with arbitrary modified features and help detector to learn features with unseen observations.

In all the three approaches [112, 113, 126], authors used a dataset from Kaggle Microsoft Malware Classification Challenge [109], which contains nine malware families. The accuracy of detecting malware in [126] is 96.39% for 90:10 ratio of training-testing data. For [113] and [112], authors achieve the average accuracy of 95.74% and 96.97% respectively. Nevertheless, the major problems in these approaches are the interpretability of the results obtained and the black-box nature of these systems. The difficulties with interpreting and understanding neural networks such as GANs, places a constraint on how these models can be modified. Also, the process of training large scale GANs is not a simple and trivial task. Designers of malware detection systems need an efficient GAN-based security system to defeat the rapidly evolving malware adversarial programs [127].

Ni et al. [128] proposed an algorithm named, Malware Classification using SimHash and CNN (MCSC), that converts SimHash values to greyscale images and then uses CNN to identify malware families. The algorithm combines multiple methods such as cascade hashing, major block selection, and bilinear interpolation to improve accuracy. The algorithm is divided into three parts: the first part extracts opcode sequence-based features and calculates the similarity between sequences using SimHash method. In this step, multiple cascading hash functions strengthen SimHash. Moreover, major block selection chooses opcode sequences that contain representative/informative feature. For example, the opcode sequences containing CALL instructions are considered as a major block because this instruction invokes APIs, library functions etc. In a second step, the SimHash value and bilinear interpolation convert opcode sequences to malware images. Finally, these images train CNN model with the aim to identify their families. Similar to other methods, this algorithm also used Malware Classification Challenge dataset containing 10,805 samples. The classification accuracy is 98.86% on average, which is higher than other algorithms. However, this work can identify only known malware, and no solution was indicated that shows the integration of characteristics of new malware in the detection system.

Akin to above, Yan et al. [114] also proposed a method, MalNet, that uses greyscale images and opcode sequences to detect and classify malware. The method uses CNN for extracting image-based features while Long Short Term Memory (LSTM) is used to learn opcode sequences features. Moreover, in order to overcome the gradient vanishing problem of LST, the truncated backpropagation algorithm is used along with a subsequence selection method to filter out benign opcode sequences. The processing of the MalNet is divided into two stages. The first stage is to preprocess malware sample data; it takes a binary form of a Windows executable file, generates a greyscale image from it, and extracts opcode sequence and metadata features using a disassembler tool. CNN and LTSM then process the greyscale image, and the results are sent to the stacking ensemble to integrate two networks' output and metadata. The experiments were conducted on 21,736 malware samples from Microsoft and 20,650 benign samples. The MalNet achieves detection accuracy of 99.88% and TPR of 99.1% with a false positive rate of 0.1% respectively. In addition, the malware classification accuracy of nine families reaches up to 99.36%.

To address the problem of data imbalance, Cui et al. [129] developed a method that combines the CNN with the Bat algorithm. The Bat algorithm is used to balance the image dataset, and CNN is used to identify and classify the images. In this work, first, samples are converted into greyscale images and then classified using CNN. The size of an image determines the structure (layers) of CNN, whereas, the width of the image could be fixed based on the pre-defined standard given in the literature. For instance, an image file of 100kb to 200kb size should have a fixed image width of 384. The method uses two different types of layers, namely convolution and sub-sampling layer. The former reduces the number of image parameters and preserves the main features such as translation, rotation, or scale invariance; while the latter reduces the dimensions of features, thus weakening the influence of image deformation. Both layer types improve the accuracy of models and avoid overfitting. The malware dataset was taken from Vision Research Lab consisting of 9342 greyscale images of 25 malware families. The achieved average accuracy is 94.5%.

Yue et al. [130] also addresses the problem of class imbalance by proposing a method that incorporates CNN and weighted softmax loss method for effective classification. Authors used vgg-verydeep-19 model with 43 layers and also placed two dropout layers between the three fully connected layers to prevent overfitting. The weighted softmax loss layer is then used as the last layer. The method achieved 98.63% accuracy with the Malimg dataset containing samples from 25 malware families which were highly imbalanced. The problem with this work is that it requires more computation as there is a total of 60 layers which are used for malware classification.

Le et al. [115] proposed a generic image scaling approach, where a raw malware byte code is interpreted as a one-dimensional image and is scaled to a fixed target size. The raw static byte code inputs to a convolutional neural network followed by a recurrent neural network. The approach uses three different models of CNN; i) conventional CNN, ii)

CNN-UniLSTM and iii) CNN-BiLSTM. In the second and third model, a recurrent neural network (RNN) is used to improve the performance of malware classification. Similar to previous works, the authors used the Microsoft Malware Classification Challenge dataset containing 10,868 samples. The results showed that the CNN-BiLSTM has the highest accuracy of 98.20% as compared to the other two models.

This approach is simpler than converting a malware binary file to a 2D image since one doesn't have to make a decision about the height and the width of the image. The image size is fixed to 10000 bytes. However, this work was intended to identify the malware class from nine classes and not to decide if a file is benign or malicious. Akin to above, Kalash et al. [116] also converted malware binaries to images and trained CNN. However, authors used VGG network architecture which is a variant of CNN. The method was tested on two large datasets, Malimg and Microsoft malware, and achieved an accuracy of 98.52% and 99.97% respectively.

Jiawei et al. [131] also used CNN to detect and classify malware residing in IoT devices. The proposed approach uses a rescaling method to standardise the image sizes of malware samples to 64*64 pixels. Such images then serve as input to a CNN that determines whether the analyzed program is malicious or benign. To evaluate the performance, authors used malware captured by the IoTPOT honeypot [132]. The approach was run on 500 malware samples and achieved an average accuracy of 94.0% for two-class classification and 81.8% for three-class classification. (i.e., benign, Mirai, or Gafgyt). Analysis results showed that malware samples have distinct byte patterns in the padded areas and unused data sections. However, the authors admit that their approach is susceptible to complex code obfuscation. In addition, the padded areas and unused data sections can be easily modified by malware developers without affecting the behaviours of malware.

Cao et al. [117] developed a malware identification and classification system based on CNNs. Firstly, the binary files of potentially malicious codes are converted into greyscale images, and then the images are fed into the CNN for malware detection. The proposed system works in an online environment where a browser plug-in checks if a file is malicious or not. If considered malicious, it is sent to the detection server where the CNN performs further analysis, and finally, the database is updated. In another scenario, a user can also upload the malicious file on a website. The method used two popular malware datasets, Vision Research Lab and Microsoft Malware Classification Challenge with the 20 well-known malicious code families. The system achieved an accuracy of 95%.

Similar to above, Akrash et al. [133] classified malware families by converting malware binaries to greyscale images and using a hybrid CNN and Long Short Term Memory (CNN-LSTM). The hybrid model enables the extraction of temporal and spatial features which are later used to identify the malware family. To select a generalised model that is suitable to deploy at run-time, the authors conducted several experiments on Malimg dataset that contains 9,389 malware samples of 25 different families. The accuracy of the method ranges from 95% to 97%. However, this research work did not explain how it is different from other neural-network-based approaches.

A recent approach by Venkatraman et al. [134] used an image-based technique with deep learning architectures to detect and classify obfuscated malware into their respective families. The proposed approach performs an image comparison of different malware families as well as benign datasets to visually demonstrate the significant difference in the behaviour patterns of the malware families. This approach is different from other methods as it uses a self-learning system which is capable of detecting not only the known malware and the variants of known malware but also unknown malware. The architecture composed of three different subsystems, in which one subsystem uses unsupervised learning model and the other two use supervised learning models. First, pre-processing on samples are performed to detect packed and unpacked binaries. After this, image analysis extracts image related features from a pre-trained deep CNN model and then clusters in the image feature space using k-means clustering algorithm. The features are then fed into CNN. The approach was tested on 52,000 malware samples collected privately in previous work as well as commonly used public datasets for benchmarking. Results indicated accuracy ranges from 94.8% to 96.3% on average, whereas F1 score ranged from 90.3% to 91.6%.

A method proposed by Darshan et al. [135] also used a CNN to detect windows-based malware. However, their process to generate an image is quite different from the rest of the methods. It uses the run-time behaviour features (n-grams) of the PE files, selected by a feature selection technique to create images, and the generated images are then fed into the CNN for the detection. In the training phase, behaviour-based feature extractor records the behavioural features (CAT-API features) of the PE file, which is under execution in a Cuckoo sandbox. These features are then processed to derive n-grams. The n-grams are then passed through feature selection techniques such as Chi-Square, Mutual Information, Information Gain, and Relief to choose the best n-grams. The chosen n-grams are passed through the image generator that checks the occurrence of the n-gram present in the final feature set to construct an image for the n-gram file. Based on the presence or absence of the n-gram, the value 255 or 0 is written onto the image (0 represents that the n-gram is absent, and 255 indicates that the n-gram is present). Finally, each input image proceeded through two convolution layers, two sub-sampling layers, and two fully connected layers for CNN training. In the testing phase, the test image was created from the unknown PE files, which was then sent to the trained CNN to measure the detection ability of the trained CNN. Authors used 400 self-generated samples (200 benign, 200 malware) and a Malheur dataset [136] to demonstrate the effectiveness of the proposed CNN-based Windows malware detector. Malheur dataset contains 3,282 benign and 4,151 malware files that included ten different types of malware. A set of experiments were conducted to demonstrate the classification ability of the proposed approach and compared with the chosen six machine learning-based classifiers. The method achieved maximum detection accuracy of 97.968% outperforming machine learning-based classifiers.

To overcome the drawbacks of code obfuscation and high computational costs of analysing malware, Xue et al. [137] recently proposed Malscore, which is based on the probability scoring and machine learning. The purpose of probability scoring is to decide if static analysis needs to concatenate with dynamic analysis. First, greyscale images are generated from raw malware to train a CNN-based classifier. The output from the CNN classifier judges the credibility of each malware sample by comparing the score with the probability threshold. The samples with lower credibility are dynamically analysed by executing the malware in the sandbox and extracting native API call sequences. These API call sequences train n-grams and machine learning-based classifier. To understand this whole process, consider a packed/obfuscated malware that gets lower credibility from the CNN classifier. This sample will then be filtered out for dynamic analysis. In this way, this work improves the robustness for packed/obfuscated malware and also improves the classification accuracy. Similarly, to reduce the cost of detection, probability scoring in this work filters out most malware that get reliable classification results. The experiments were conducted on the dataset from VX Heaven website [94] and after pre-processing, 174,607 malware samples from 63 malware families were used. The result showed that Malscore achieved 98.82% accuracy while reducing the pre-processing and test time by 59.58% and 61.70%, as compared to previous methods.

Though highly efficient, this work did not use dynamic analysis as the primary detection method, but as a complementary method for static analysis to increase the robustness in detecting packed malware. Also, this work did not take into account the concept drift as it did not provide any incremental learning strategies to fortify re-learning ability.

To detect counterfeit mobile apps, Rajasegaran et al. [138] proposed a neural embedding-based approach that combines content and style embeddings generated from pre-trained CNNs. The method relies on the idea that a combination of style and content embeddings achieves better results in detecting visually similar app icons. First, the app icons are fed into a pre-trained VGGNet for the extraction of content and style embeddings. Authors crawled and collected approximately 1.2 million apps from Google Play Store and showed that a combination of content and style embeddings achieved 8% to 12% higher precision and 14% to 18% higher recall, respectively compared to conventional methods. The precision and recall are further improved to 3% to 5% and 6% to 7% by adding text embeddings. Authors identified 7,246 potential counterfeits to the top-10,000 apps, out of which 2,040 may contain malware. Results further indicate that 1,565 of apps ask for at least five dangerous permissions and 1,407 have at least five extra third-party ad libraries.

E. Summary of computer vision based malware detection:

In this section, we surveyed research work that used computer vision methods to detect and classify malware more accurately and efficiently. The main motivation of using computer vision methods is to overcome the limitations of static and dynamic code analysis methods, where the former suffers from the problem of code obfuscation and detecting new malware types, and the latter is computationally expensive. The discussed computer vision-based research work has shown promising results in terms of detection and classification accuracy, and in some cases, offers high accuracy against the detection of obfuscated or encrypted malware. Moreover, some of the research work in this area is independent of the malware execution environment and computationally inexpensive.

In computer vision-based malware detection, the main idea is usually to convert malware binaries or behaviour of malware to an image and then perform analysis on the image using methods such as feature extraction, image comparison, and image hashing. Similar to phishing detection, we categorise the malware detection approaches to *i) image feature-based methods*, *ii) image hashing-based methods*, and *iii) neural networks-based methods*. In addition, we also found some work under the *image representation-based methods* that visualise malware activities such as API calls, resources consumption, file structure decomposition, and frequency of API calls, as images and then compares the unknown sample with abnormal actions using image processing techniques such as bitmap converter, thread graph, treemaps, and colour mapping. We summarise all the work we surveyed in Table II.

Next, we discuss observed challenges and limitations over all the work, with the purpose to make the contribution of computer vision more effective in the area of malware detection, and thereby deploying solutions commercially. Our study on existing computer-vision malware detection methods reveal various findings:

- **Malware Dataset and Families:** We found that most of the work used Microsoft Malware Classification Challenge (BIG 2015) as their dataset. This dataset is 0.5 terabytes in size, consisting of disassembly and bytecode of more than 20K malware samples. Another malware dataset which is most often used is the Malimg dataset. This dataset contains 9,339 malware samples from 25 different malware families. We also found a work proposed by [89] that has set-up honeypots in a controlled network environment for real-time malware detection.

  We also observe that most of the work did not validate their work against multiple datasets. For instance, Silva et al. [110] evaluated their method only on the BIG dataset, whereas, sl et al. [135] used Malheur dataset for evaluation. Nevertheless, there are methods such as Le et al. [115] and Kalash et al. [116], that performed evaluation on more than one datasets. We also notice that most of the work was evaluated on the specific set of malware families provided by the datasets and discussed the extension of their work to new malware variants. There are very few works that perform regressive testing on a broad set of malware families. For example, we found work by Han et al. [93] that has performed their experiments on 50 different malware families. To build a robust malware detection system, it is necessary to collect latest and large datasets, containing several malware families and their corresponding variants, such as EMBER [139]. Moreover, malware datasets should contain samples from a range of platform to check the

TABLE II: Summary of Visual Similarity Based Malware Detection Approaches

Research Work	Computer Vision Methods	Input Image Structure	Dataset Size/Family	Summary	Performance Metrics
Image Representation-based Methods					
Trinus et al. [89]	Treemaps, ThreadGraphs, CWSandbox	(Generic) Size similar to original image	2,000/13	Applies an abstraction method to make malware reports easily understandable and to enable the extraction of unique features for malware detection.	Not Given
Han et al. [90]	SimHash, Vector Angular-based Distance Measure	256×256 pixels of image matrices	-/3	Generates image matrices with (x,y) coordinates and RGB colours and then uses vector angular-based distance measure algorithm to calculate the similarities between image matrices.	Similarity (across same family) - 0.95 Similarity (across diff. family) - 0.32
Han et al. [92]	SimHash, Vector Angular-based Distance Measure	256×256 image	Static- 290/16 Dynamic- 560/14	Extracts the opcode sequences related to staple behaviours and uses angular-based distance measure to calculate the similarities.	Accuracy (Static)- 98.96% Accuracy (Dynamic)- 97.32%
Han et al. [93]	Bitmap Image Convertor, Entropy Graph Generator, Similarity Calculation Algorithm	(Generic) Pixels of original file size	1,000/50	Calculates entropy values of a greyscale bitmap image to generate the entropy graph and further measures similarity of entropy graph for the malware detection.	Accuracy (ROC) - 97.9%
Shaid et al. [96]	API-to-Colour Mapping, API Monitoring Utility, Similarity Measure	64×4 pixel	1,101/12	Utilises API call monitoring utility to obtain behavioural data of an image by mapping to colour, and then groups malicious behaviour.	Accuracy - 95.92% to 98.92%
Chapman et al. [97]	File Structure Decomposition	JPEG and PNG files of variable length	270k JPEG and 33k PNG files	Decomposes each image file into sequence of segments, feeds into the training algorithm to build a formal model, and then classifies unknown samples.	TPR - 99.24% (JPEG), 99.31% (PNG) TNR- 99.32% (JPEG), 98.88% (PNG)
Sun et al. [98]	Layout Edit Distance (LED)	Apps XML Layouts in tree forms	1,00,096 apps	Detects the repackaged and phishing malware in apps by calculating the LED between XML layout of android apps.	Visually Similar Apps: 1298
Image Feature-based Methods					
Nataraj et al. [100]	k-NN with Euclidean Distance	320 GIST descriptors	9,458/25	Converts binary file to a greyscale image and extracts 320 GIST features to classify malware using k-NN.	Accuracy - 98%
Nataraj et al. [95]	k-NN with Euclidean Distance, System Call-level Monitoring, Forensic Screenshots	320 GIST descriptors	393/ 3,131/24 63,000/531	Converts binary file to a greyscale image and extracts 320 GIST features to classify malware using k-NN.	Accuracy (First Dataset) - 98% Accuracy (Second Dataset) - 97% Accuracy (Third Dataset) - 72%
Kancherla et al. [103]	SVM, Gabor Filter, Wavelet Transformation	2-dimensional vector of file size	25,000/6	Converts a binary file to a 2D vector, extracts Gabor, Intensity and Wavelet features, and then uses SVM for malware detection and comparison against each feature set.	Accuracy (Gabor) - 93.23% Accuracy (Lightweight) - 94.32%
Ahmadi et al. [107]	Forward Stepwise Selection Algorithm, XGBoost Classifier	Hex byte as a greyscale pixel	21,741/9	Calculates set of features from different categories of hex and assembly view of PE and applies feature fusion algorithm to generate the most effective concatenation of features categories such as Haralick features and Local Binary Patterns features to classify a malware.	Accuracy (with Img. Feat.) - 95.5% Accuracy (comb. of other feat.)- 99.40%
Zhao et al. [108]	Homology Analysis, B2M, Improved SIFT, Guassian Pyramid, DBSCAN	66 byte feature descriptor	1,789/11	Extracts features from binary sample, creates texture fingerprint from the features, and then applies clustering to group the texture fingerprint of malware into families. Afterwards, the homology analysis was performed to classify the malicious unknown samples.	Accuracy - 77.6%
Masouleh et al. [110]	HOG, AE, Summary Statistics Features from icon pixels, HDBSCAN, K-NN	Icon Pixels of varied length	1,138/	Detects PE malware by extracting HOG, AE and summary statistics features form icon pixels, clusters them, and then trains machine learning classifiers.	Accuracy: 84.4% TPR: 80.2%
Su et al. [118]	Colour Correlogram Algorithm, Cosine Similarity	App Icons and Screenshots	30,625 crawled mobile apps	Detects camouflaged apps on mobile market places using indexed based text and image retrieval systems.	FPR: 9.22%
Image Hashing-based Methods					
Xiaofong et al. [119]	LSH SURF, Euclidean Distance	64*64 image vector	8,410/25	Computes a feature fingerprint and applies LSH to return the top most visually (structurally) similar variants for malware classification.	Precsion/Recall/F-Score - 80% to 90%
Malisa et al. [120]	LSH	Mobile app screenshots	150,000 mobile apps	Detects impersonation attacks by capturing run-time app activities screenshots and measures similarity through LSH.	FPR - 15% Impersonated Apps- 43,904
Andow et al. [111]	Fuzzy Hashing	App icons	2,500 mobile apps	Detects imposter apps through text analytics and static program analysis, where later measures the similarity of app icons through fuzzy hashing.	Manual App Review Reduction: 977 apps out of which 22 app icons are same
Jiao et al. [122]	pHash	App images	Not Given	Detects repackaged apps by comparing hash codes and if the the similarity exceeds a threshold, then flag the app as repackaged app.	Not Given
Long et al. [123]	wrestool, FLANN index, L2 distance	1024 dimensional vector	Not Given	Measures similarity between apps by extracting images, converts images to features, and then uses L2 distance for similarity checking.	Recall - 83% Precision - 85%
Neural Network-based Methods					
Zhang et al. [125]	Opcode Sequence, Histogram, Dilution, Erosion, CNN	2-Dimensional opcode images	17,808/10	Disassembles binary executables into opcodes sequences, converts the opcodes into images and then compares with the known malware sample using CNN.	Accuracy (Max.) - 96.7% TPR (Max.) - 94.3% TNR (Max.) - 99.1%
Kim et al. [126]	GANS, t-SNE Algorithm	Average of all the rows and columns of the data	9,970/9	Utilises transfer learning method to pre-train the generator for generating malware images close to real malware for the possible detection of zero-day attacks.	Accuracy (Avg.) - 96.39%
Kim et al. [113]	DAE, GANS	63*135 pixel image	10,800/9	Devises an architecture by combining several deep learning methods of DAE, GAN, and transfer learning to construct a malware detector that is robust to noise and zero-day attacks.	Accuracy (Avg.) - 96.74%
Kim et al. [112]	GANS, VAE	Image size 256*128 Latent Space Size - 270 dimension	10,867/9	Generates malware from latent space based features and helps the detector to learn features with unseen observations.	Accuracy (avg.) - 96.97%
Ni et al. [128]	SimHash, CNN, Bilinear Interpolation	8*8, 8*16, 16*16, 16*32, 24*32, 28*32	10,805/9	Extracts opcode sequence based features, uses SimHash method to calculate similarity, converts SimHash values to greyscale images and then uses CNN to identify malware families.	Accuracy (Avg.) - 98.86%
Yan et al. [114]	CNN, LSTM, Truccated back-propagation Algorithm	64*64 greyscale image	21,736/9	Uses CNN for extracting image based features and LSTM to learn opcode sequences features for malware detection.	Accracy - 99.88%
Cui et al. [129]	CNN, Bats	Multiple representations depending on original file size	9,342/25	Malicious code is converted into greyscale image and then images are classified using CNN having convolution and subsampling layers.	Accuracy (Avg.) - 94.5%
Yue et al. [130]	CNN and weighted softmax loss	VGG-verydeep-19 CNN model with 60 layers	9,342/25	Proposes a solution to overcome class imbalance problem in malware classification through CNN of 43 and two dropout layers.	Accuracy - 98.63%
Le at al. [115]	CNN, CNN-UniLSTM, CNN-BiLSTM	10000 byte file size	10,868/9	Utilises three different CNNs to improve the performance of malware classification.	Accuracy(CNN) - 95.8% Accuracy (CNN-UniLSTM) - 98.12% Accuracy(CNN-BiLSTM) - 98.20%
Kalash et al. [116]	CNN	320 GIST descriptors	Malimg - 9,339/25 MS - 1,868/9	Converts malware binaries to images and trains CNN for classification.	Accuracy (Malimg) - 98.52% Accuracy (Microsoft) - 99.97%
Jiawei et al. [131]	CNN	64*64 greyscale image	500/4	Utilises CNN to detect and classify IoT DDOS malware residing in IoT devices.	Accuracy (2-class) - 94.0% Accuracy (3-class) - 81.8%
Cao et al. [117]	CNN	2-D array with specified widths 28*28 / 56*56	Vision Lab and Microsoft Datasets/20	Works for online environment where a browser plug-in sends suspected malicious file to detection server running CNN for further analysis.	Accuracy - 95%
Akrash et al. [133]	CNN-LSTM	Greyscale images of varying size	9,389 malware samples	Detects malware families using CNN-LSTM model and tests against various parameters.	Accuracy: 95% to 97%
Venkatraman et al. [134]	CNN, CNN BiLSTM, CNN BiGRU	Greyscale images of varied length	52,000 malware samples collected from private and public sources	Detects and classifies known and unknown malware samples through a hybrid deep learning and image analysis techniques.	Accuracy: 94.8% to 96.3% F1 Score: 90.3% to 91.6%
Darshan et al. [135]	CNN, n-gram features	Greyscale images of varying size selected by feature selector	Self-Generated Dataset: 400 samples Malheur Dataset: 7,433 samples	Detects and classifies known and unknown malware PE files by executing them at run-time, extracting n-gram features, converting features into images and feeding into CNN.	Accuracy: 97.96% F-Score, Precision, Recall: 0.97
Xue at al. [137]	CNN, n-grams	Greyscale images	174607 malware samples	Classifies malware using probability scoring and machine learning models with the purpose to overcome the drawbacks of static and dynamic malware analysis methods i.e. code obfuscation and computation cost.	Accuracy: 98.82%
Rajasegaran et al. [138]	VGGNet	App icons	1.2 million apps from google Play-store	Detects counterfeit apps by using a pre-trained VGGNET that combines content and style embeddings.	Precsion: 8% to 12% Recall: 14% to 18% Counetrfeit Apps: 7,246 from the top-10,000 apps

diversity of malware detection method. An example here is Winnti Windows malware family that was first explored in 2013, however, recently (in 2019) its Linux variant was discovered which dated back to 2015 [140].

- **CNN and GANS:** We noticed that understandably CNNs are an emerging approach in detecting and classifying malware efficiently. We found that several recent works have used different versions of CNNs, either to detect malware presence or to classify malware types. CNN-based methods are fed either with greyscale or coloured malware images and then tested with a number of different input and output layers and various parameters. As a matter of fact, each of the research work in neural network category used a different number of inputs, middle, and output layers and also tuned with different hyper-parameters. There is a need to investigate for a generic neural network model that has a capacity to tune layer and parameters for high malware detection accuracy automatically. Recently, researchers also used GANS to generate new malware training data to detect zero-day attacks [112, 113, 126]. However, the research is still underway as it involves many challenges. For instance, it is still not clear that malware generated through GANS resembles the real malware and how much helpful they are in the detection of zero-day attacks.

- **Real world adoptions:** The computer vision-based malware detection has not fully adopted by the industry, despite a significant amount of work happening. We did not find any company that explicitly mentions the usage of such work in their commercial products. One possible reason for this gap is that computer vision-based malware detection methods are not mature enough to be deployed as commercial products. The proposed methods neither have convincing results nor cover all aspects of malware detection. For instance, most of the methods were tested on a specific dataset and a particular network environment. These methods have not been deployed in open network to verify their robustness, reliability, and performance at large scale. Moreover, there is no standard benchmark to compare these work against each other.

V. VISUAL SIMILARITY BASED TRAFFIC ANOMALY DETECTION

Computer vision methods also alleviate the limitations of traditional traffic anomaly detection approaches that are usually based on statistical features of traffic flows. For instance, a conventional approach that detects short and long-lasting anomalies through variances in the number of bytes (traffic volume) cannot identify low-rate attacks [141]. An example of a low-rate attack is port scanning that does not consume much bandwidth when it tries to access an abnormally large number of ports on a single host [142]. Also, such approaches cannot identify attacks that are statistically close to global traffic behaviour. Finally, these traditional approaches are unable to detect zero-day attacks. As a result, much effort has been put into detecting anomalies through computer vision methods.

In this section, we survey computer vision-based traffic anomaly detection approaches and categorises them based on their method of detection. We found work fitting into the categories of *i) image representation-based methods, ii) image feature-based methods,* and *iii) neural network-based methods.* However, we highlight the volume of work we found for computer vision-based traffic anomaly detection is much less compared to *phishing detection* and *malware detection.* A possible reason for this gap is that traffic anomaly detection is a highly non-trivial application for image processing techniques, as compared to phished web pages or malware binaries that can be directly represented as images We also did not find any work that utilised *image hashing-based methods* to detect traffic anomalies. We further discuss the possible reasons behind this at the end of the Section V-D.

A. Image representation-based methods

Similar to Section IV-A, the image representation-based methods detect traffic anomalies by visualising traffic data or logs as images and then analysing them using computer vision algorithms. The traffic data usually consists of network packets with header and payload sections, where header contains information such as source IP address, destination IP address, source port, destination port, and protocol names. Studies have shown that network traffic can have strong patterns of behaviour over several timescales [143]. Therefore, by passively monitoring packet headers of network traffic at regular intervals and then generating images of traffic usage (which can be correlated with previous images), can help to detect anomalous behaviours. For example, in case of anomalous traffic such as DoS attack, the usage pattern of network at time t_i is different from previous normal states $t_1, t_2, ... t_{i-1}$. These usage patterns at different time intervals can be represented as visual images; for example, images of traffic volume originating from source or destination IP addresses, traffic flow between IP addresses, or intensity of traffic flows. The images are then analysed using computer vision algorithms (e.g. Hough Transformation [144], Canny Edge Detection Algorithm [145], and Video Compression Algorithms [146]), to find peculiarities or further characterise the anomalies into several categories (e.g. random attacks, targeted attacks, multi-source attack, and port scans). Figure 3 shows a generic process of detecting anomalous traffic using image-based representation methods.

Kim et al. [147, 148] is one of the early work in detecting anomalous traffic using image representation methods. Authors proposed an approach, *Netviewer*, that used image processing and compression techniques for detecting and identifying network attacks in real-time. First, samples of network traffic are converted into images. The traffic samples are packet header traces that consist of data such as source/destination addresses, port numbers, and traffic volume in bytes. For instance, an image may be represented as traffic volume in bytes originating from a source IP address, or the traffic between a source and destination IP addresses/port numbers. Second, the network traffic images, from regular time intervals, are then considered as video frames such that a sudden change in a scene (traffic pattern) indicates a possible traffic anomaly. To analyse the data in video frames, various techniques from
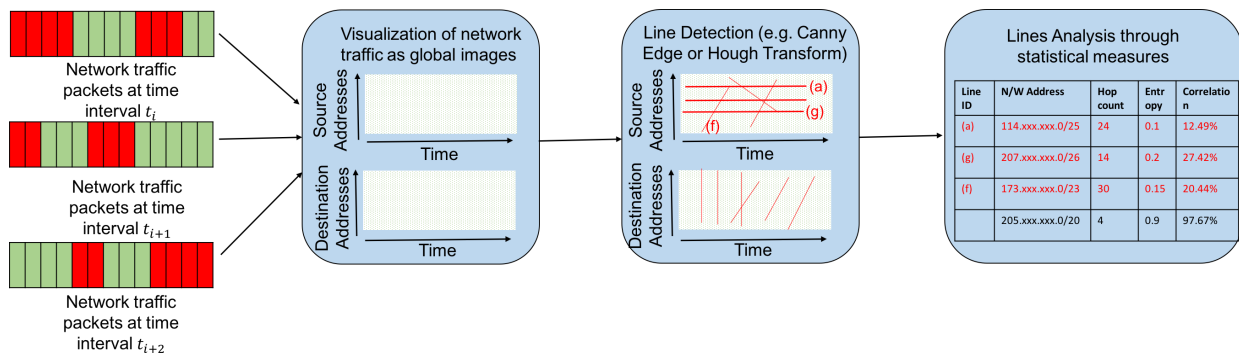
Fig. 3: Generic Process of Anomalous Traffic Detection using image-based representation methods

video compression and image processing such as Discrete Cosine Transform (DCT) coefficients and absolute difference of image pixels are applied. For instance, a mean square error (MSE) is calculated between video frames (which are collected at regular interval) to detect anomalies. Finally, to identify attacks and victims, line or edge detection algorithms are used. For example, consider a representation where the x-axis of an image represents destination IP addresses, and the y-axis represents the source IP addresses. A horizontal line in this image, is an indication that a source is accessing multiple destinations, for example, during worm propagation. A horizontal line may also indicate a host scan of destination machines by a single source with a high probability. Similarly, a vertical line indicates that several sources are accessing a single destination (e.g. a DDoS/Botnet attack emanating from multiple machines towards a single destination). Moreover, Netviewer also estimated the movement of attack patterns using motion prediction algorithms.

Authors used real traffic traces from three major networks to evaluate their approach. The networks were University of Southern California (USC), Korea Research Environment Open NETwork (KREONet2), and Texas A&M University campus [149]. Authors claimed that their method has an average accuracy of 92.3% and a false positive rate of 0.3% when the source IP address was considered in the packet header. There was a 5% reduction in accuracy and 0.45% increment in false positive rate when only the destination address was considered in the analysis. Authors also compared their approach with popular Intrusion Detection System (IDS) Snort and showed that the results of both systems agree with each other. However, both systems differ from their methodology perspective. By investigating payload and packet header, Snort is capable of identifying the source of malicious activities as well as functions performed by those malicious activities. On the other hand, Netviewer investigates packet headers only and reports suspicious IP addresses and the pattern of abnormality in aggregation manner (e.g. by observing the number of packets arriving at a specific destination IP addresses at time t). This makes Snort superior in presenting granularity details of attacks. Similarly, Snort employs qualitative analysis (rule-based), whereas Netviewer is quantitative (aggregation-based). As a result, Snort may miss the identification of some unusual traffic if the rules are not adequately defined. For example,

heavy traffic flows from a single host in a network may not be detected by Snort without an operational rule. In contrast, Netviewer can identify such flows as anomalous because of its aggregation-based method.

In [150], authors used a similar approach to Netviewer [147], with the main difference being using image pixel intensities and discrete cosine transform (DCT) coefficients to detect scene changes in video frames. In addition to other packet header data, an image also represents the intensity of traffic of corresponding IP address through colour and darkness of each pixel. For example, a packet count from a source or destination IP address can be represented as the intensity of the pixel in the image representation. The anomalies are then detected using real-time and postmortem analysis, where the former exploits the variance of pixels intensities of a traffic image and the latter employs the variance of 8-by-8 DCT of the traffic image. Similar to [147], the approach was tested on the traces from three major networks and showed that the real-time analysis has a 0.22% and 2% higher false positive and accuracy rate respectively compared to postmortem analysis based on DCT coefficients.

Besides, the authors also compared their the real-time analysis method with classical detection theory-based *Neyman-Pearson test* and showed the false alarm rate of 0.37% and the detection rate of 74.6%. As compared to the previous approach by Kim et al. [147], the results of this approach are slightly better, which indicates that the real-time analysis (image pixels intensities) on the network traffic is a better performer for the anomaly detection than the postmortem analysis. Moreover, the advantage of using real-time analysis is that it is effective against flood types of attacks. When images are defined as the intensity of traffic flow in the source address, destination address, or destination port, and analyzed as video frames, then the distribution of abnormal traffic flows in some frames are expected to be much different from normal and historical distribution.

Above two approaches [147, 150] detect anomalies by representing traffic flow as images and then using a threshold to determine significant changes from the images of regular network traffic. However, these methods may not be able to detect several other anomalies that do not rely on traffic flow intensities. To this end, Fontugne et al. [142] proposed an approach that is capable of detecting unusual traffic behaviours

by pointing out the abnormal distribution of traffic features. In this approach, network traffic is mapped to snapshots, and then pattern recognition technique such as Hough Transform is used to identify anomalies. Pattern recognition not only allows for unsupervised learning, but also enable the detection of ambiguous short-lived traffic. To detect anomalies, first, time-based sliding windows are adjusted, and then network traces in each sliding window are converted into images using four traffic features (i.e. source address, destination address and ports). Once converted, the Hough transformation is applied to detect lines in snapshots representing unusual and excessive use of traffic features. Once a line is found, network traffic data involved in those lines is retrieved and summarised as an event. An event provides detailed information such as the percentage of use of all protocols, the entropy of each traffic feature, timestamps, and IP addresses, about packets extracted from the whole traffic flow. Moreover, these events are clustered together based on destination or source address to detect the anomaly.

The approach was tested on the traffic traces from the Measurement and Analysis on the Wide Internet (MAWI) archive [151], which were collected in 2004. Overall, it detected 73.8% of anomalies from a dataset containing 2000 anomaly traffic and 630 normal traffic. Moreover, the proposed approach was compared with a statistical approach non-Gaussian multi-timescale models [152] and showed efficient detection of short and long-term anomalous traffic. The proposed method [142] detected 297 more anomalies which were not identified by a statistical method. Also, the method claims to have shorter detection delay as compared to the statistical method (i.e. this approach detects traffic anomalies within three seconds, whereas the statistical method needs at least one minute to detect anomalies). Though the proposed approach is capable of detecting more anomalies than the compared statistical approach, the authors did not discuss false positive anomalies, thus requiring further analysis. In addition, authors applied their method to an old dataset which may not reflect the current characteristics of network traffic. Hence, there is a need to validate this work with the latest network traffic datasets such as CAIDA IPv4, IPv6, or KeroNet. All in all, the approach is capable of detecting volume-based attacks through traffic features.

Jeong et al. [153] proposed a hierarchical approach that performed image analysis in two-tiers to detect network attacks. The first-tier detects random attacks (e.g. DDoS attacks) by analysing global images whereas the second tier detects semi-random attacks (e.g. post scanning attacks) by examining the local traffic images. Global images are created based on the network traffic such that X-axis represents the destination IP addresses and Y-axis represents the source IP addresses. The intensity of the image indicates the traffic flow between the source and destination pairs. Similar to [142], lines in the images are identified using the Hough transform. These lines are an indication of high usage of a port number, IP address or a high traffic flow between specific source and destination addresses. Based on the information provided by the lines in the global image, local images are created that contain more details about the attack. Local images show traffic volume for

each destination host and port number in a network domain. Once created, Connected Component Labelling (CCL) [154] is applied to local images to detect traces of mostly used destination ports (a.k.a post scanning). Authors evaluated their approach on Netflow data of KREONet, which collected network traces from July 20 to July 27, 2003. The method analysed 60,000 to 70,000 events in 3 seconds, thus showing efficiency. However, the authors did not discuss the accuracy of detecting network attacks using the proposed method.

Similar to [153], Kim et al. [155] also represent unusual patterns of traffic as lines on a 2D image space. However, their approach detects traffic anomalies that generate only from a DDoS attack with IP spoofing. The unusual patterns which represent DDoS attack sources are recognised using Canny Edge Detection Algorithm. Afterwards, signature-based pattern extraction algorithm, *pivoted movement*, is applied to prevent DDoS attacks. The pivoted movement algorithm correlates information such as IP address and media access control (MAC) pairings.

The overall proposed scheme [155] works as follow: First, the incoming packets to a victims machine are converted into a 2D image using the network tool, NetSCENE, introduced by the authors. NetSCENE monitors the network packets at regular intervals and analyzes the aggregated data. The Canny Edge Detection algorithm is then applied to an image to identify DDoS attack packets with subnet spoofing. The patterns of DDoS attack appear as edges on NetSCENE where each edge is linked to the local network addresses at which targets (i.e. effected machines under the control of attacker) are located. Once lines are identified, the victims machine alerts the local source router about those network addresses which then applies the pivoted movement algorithm by paring the source IP address with MAC address. The reason to pair the IP address with the MAC is that legitimate packets pair one IP address with one MAC address whereas, the attack packets use different IP addresses with one MAC address. This is often called pivoting the IP addresses. Moreover, the two-step detection process also reduces the risk of false positives and false negatives. Finally, the illegitimate IP addresses are marked with signatures and are filtered out. To test the approach, the authors used two different datasets; IPv4 Routed/24 and IPv6 topology datasets from CAIDA [156, 157]. Results showed an average false positive rate of 1% and an average false negative rate of 0.5% for canny edge detection and pivoted movement algorithm, respectively.

While image representation-based methods offer promising results, there is also a limited amount of work that extracted image-related features from traffic data and then analysed those features for possible traffic anomaly detection. Image-feature based methods help in the discovery of abnormalities in the network that does not depend on traffic volume and would otherwise go undetected by image representation-based methods. For instance, web attacks such as SQL injection can be detected through statistical features (e.g. length of URL, string input types, string matching), which are difficult to detect through image representation-based methods. Moreover, by exploiting the correlation among image related features of network traffic, it is possible to identify the type and

characteristics of anomalies precisely and at the same time, reduce false positive and increase true positive rate. We next discuss image feature-based methods that enable the detection of traffic anomalies in the network.

B. Image feature-based methods

In comparison to image representation-based methods, we did not find a significant amount of work in image feature-based methods for traffic anomaly detection. The most plausible reason being, usually different image features capture different elements of the image such as key-points, colour distribution, and intensity levels and it is not straight forward to decide which feature might be important to identify traffic anomalies with high accuracy. Therefore, such methods may not provide additional benefits over statistical feature-based conventional methods. The two work we found belonging to this category, extracted features from a traffic sample/pattern and determined the presence of an anomaly in network traffic. The extracted features mainly include entropy, contrast, pixel randomness, and local homogeneity of the image.

Early work in this area is by Tan et al. [158] that detected DoS attacks using features of network traffic. The overall approach consists of three steps. In the first step, basic features such as Timestamp, Source IP, Destination IP, Source port, source bytes, average duration of host, average duration of all services, number of ICMP packets, number of packets to all services, and number of SYN flags, are generated from network traffic packets captured at the destination network. In the second step, dimensionality reduction using Principle Component Analysis (PCA) is performed on the generated features. Finally, the selected lower-dimensional features are used in training and testing phase for possible detection of DoS attacks. In a training phase, standard profiles are generated for various types of legitimate/normal traffic records (i.e., TCP, UDP and ICMP traffic). These profiles are built based on the subspace feature set provided in step two. The normal profiles are then converted into two-dimensional feature matrix (image), named Triangle Area Maps (TAM), that reveals the hidden correlations between the features of network traffic. To do so, Multivariate Correlation Analysis (MCA) [159] is applied that selects most discriminative features. The TAM is then stored in a database along with mean and standard deviation of the earth movers distances between individual training samples and the mean of the given training samples. In the testing phase, images (TAM) of individually tested records are generated and compared against the normal images from the Training Phase using the EMD algorithm.

The method was evaluated using two datasets; KDD Cup 99 and ISCX 2012 [160, 161]. Authors also compared their method with then state-of-the-art detection systems; network intrusion detection system based on covariance feature space [162], triangle- area-based nearest neighbours approach [163], DoS attack detection system using TAM-based MCA [159], and Nave Bayes based detection approaches [164]. The overall evaluation showed 99.95% accuracy on KDD Cup datasets, outperforming the techniques in [163] and [162] by 2.06% and 7.8% respectively, and was comparable to the accuracy of [159]. On the ISCX 2012 IDS dataset, the proposed method achieved 90.12% accuracy outperforming NB-based detection approaches. It also showed promising results in terms of computational complexity by analyzing 59,738 traffic records per second. Another advantage of this method is that it processes each traffic sample individually, which helps in prompt detection of an attack and also increases the probability of correctly classifying a sample than group-based detection mechanisms. Group-based detection mechanisms take a certain number of samples together and then classify them as benign or illegitimate [162]. Overall, this work demonstrated to have less detection delay as compared to group-based approaches, yet has equal or better detection precision.

Zou et al. [165] proposed a network anomaly detection method that used image processing to analyse network traffic patterns in Cyber-Physical Systems (CPS). A CPS is an integration of physical and software components that interact with each other in different contexts to perform desired functionalities at different spatial and temporal scales [166]. Examples of CPS include smart grids, medical monitoring systems, and robotics systems. The proposed method relies on the network throughput of all the devices in a CPS network such that a communication image is created by visualising throughput as pixels. The size of the communication image is equal to the number of CPS devices in the network where each image represents the network resource usage of the whole CPS network at a time sample. The method then applies texture feature analysis [167] on images to extract the spatial information about the network resource usage of the CPS devices. More particularly, Grey Level Co-occurrence Matrix (GLCM) was applied to extract five GLCM texture features which were energy, entropy, contrast, IDM (Inverse Difference Moment), and DM (Directional Moment) [168]. These features measure the intensity of pixel pair repetitions, the randomness of pixels, the extent of a pixel and its neighbours, local homogeneity of the communication image, and alignment of it in terms of the angle, respectively. In order to extract more communication features, a temporal analysis of GLCM features was also carried out. Thus, the spatial and temporal features together determine the communication patterns of the CPS network. Afterwards, a k-NN algorithm is used to detect the anomalies.

The proposed method was experimented on four real-world decentralised applications using Common Open Research Emulator (CORE) simulator [169]. CORE is a live emulator that uses existing operating system virtualization techniques to build wired and wireless virtual networks. These networks can be connected to real networks and systems to run various protocols and applications to extend lab test-beds. Authors set up 100 devices on CORE Simulator and run every application for 200 times and capture the network throughput. Authors then simulate the anomaly network traffic by adding different intensity (10%, 20%, 30%) of random network throughput to every device. The overall classification accuracy was 99%; however, the method was not tested on a real dataset or a real environment. Moreover, the time performance of extracting GLCM features has not been discussed in the paper.

To summarise, we only found a limited amount of work under image feature-based methods for traffic anomaly detection. One possible reason may be the difficulty in identifying what features make more sense in traffic anomaly detection, which is potentially solved in the neural-network approaches we next discuss.

C. Neural Network-based methods

Similar to the phishing and malware detection, neural-network based approaches have also been explored for traffic anomaly detection. In this category, images are generated from network traffic and then directly fed into neural networks. The network automatically learns latent representations that can be leveraged to detect traffic anomalies.

Wang [170] demonstrated the idea of representing network flow data as images and utilising a deep neural network to detect anomalies. First, the payload bytes from each TCP session is converted to an image, where each byte represents a pixel after normalising to a scale from 0 to 1. Once an image is formed, the images are fed into Multi-Layer Perceptron (MLP) and Stacked Auto-Encoder (SAE) deep learning models for the identification of anomalous protocols in network traffic. The MLP model represents a supervised learning setting (requires labelled data) while the SAE represents an unsupervised setting (can work without labelled data). In the MLP and SAE models, the nodes that have higher significance (low error cost) in the deep layers are selected as features and trained on the data. Classification results show that SAE performs better on traffic classification than the MLP model.

Both MLP and SAE models were evaluated on 0.3 million TCP flow data records collected from an internal network. These records have 58 different protocol types in total while having 17% of the overall data flows labelled as unknown. Both methods were able to distinguish 54.9% of the unidentified protocols by giving a probability of 0.9 to classify a traffic protocol of each network flow correctly. However, the authors provided limited information on the experimental settings (e.g. dataset used) and results of the methods understandably due to the fact that they represent enterprise data.

In an improvement, Li et al. [171] proposed a CNN based intrusion detection approach while using a graphic conversion approach. The proposed method converts raw packets into an image, and then a CNN is used to learn the features from that graphics. Authors took NSL-KDD dataset [172] as a case study for the evaluation of their approach. In the first step, 41 features, that contain integer or float features, symbolic features and binary features, were extracted from each sample of NS-KDD dataset. The extracted features are then converted into binary vectors with 64 dimensions. Authors used one-hot encoder mapping to convert symbolic features into binary vectors and standard min-max scaling to convert other continuous features into binary. The binary vectors are then converted into 8×8 grey-scale pixel images and input to a CNN model.

Authors used ResNet50 and GoogLeNet as the CNN architectures, where former was trained with 100 epochs, and 256 batch size and latter was trained with 100 epochs and 64 batch size. Authors reported that both the models showed accuracies ranging from 79% to 82% in detecting anomalies in NSL-KDD dataset. Moreover, the authors compared their approach with traditional classifiers such as Random Forest and SVM and showed that CNN performed better. However, the authors did not report the detailed analysis of anomalies (e.g. type and pattern related to each anomaly). Also, the paper did not mention about the treatment of unknown data types/packets by CNN and the effect of such unknown data packets on the performance of the approach.

Akin to above, Wang et al. [173] also proposed an approach based on representation learning where network flow is converted into images and fed into a CNN classifier. Authors used the dataset released by the University of New South Wales (UNSW), referred to as NB-15 [174]. This dataset contains 25,400,44 network communication data records, and the total size of the dataset is 100GB. Compared to the KDD dataset, NB-15 has a large volume of data and more accurate ground truth labelling for anomalous traffic (i.e. attack types). Authors, pre-processed the dataset by performing traffic split, traffic mark, fixed-length interception, and format conversion task. In traffic split, the original traffic file is divided into small traffic files according to the session flow where a session is a communication record of the same IP 5-tuple. The 5-tuple includes the source IP address, source port, protocol type, destination IP address and destination port. Afterwards, the session units are marked (traffic mark) by finding the corresponding entry in a CSV file using IP 5 tuple and the traffic generation time. Once labelled, each byte of session unit data converts into an integer ranging from 0-256 to create a fixed-length interception. The integers then transform into a greyscale image of 784 (28×28) bytes and feed into the CNN that consists of two convolution layers and two pooling layers. The CNN was trained using cross-entropy loss with a learning rate of 0.001 and a batch size of 50 over 1,000 epochs. The authors reported accuracy of 97.3% and an F1 measure of 0.985. However, there was no comparison made with traditional classifiers or with the KDD-dataset. As such, it is difficult to establish where this method stands compared to both classical anomaly detection methods as well as computer vision-based methods.

Though different from traffic anomaly detection methods, Internet traffic classification is another related area where computer vision methods have been applied. Zhou et al. [175] classified botnet encrypted and unencrypted traffic by using CNN-based features self-learning. First, the data is pre-processed by dividing the network traffic into sessions such that the size of a session meets the input requirement of the CNN. The sessions are determined through bidirectional flow standard [176] where not only incoming but also outgoing packets also participate in malicious activity. To form a session, the authors used five tuples; source IP, source port, destination IP, destination port, and the protocol. After pre-processing, the session data was converted into a greyscale image by normalizing each byte of a session to [0, 1] where one byte corresponds to one pixel. The CNN architecture authors used is a standard set up with two blocks of convolution and pooling layers, followed by a fully connected layer of size 1,024 and a final softmax layer of 12 units.

Authors evaluated their solution on the ISCX-Bot-2014 released by The Canadian Institute for Cybersecurity [177]. The dataset includes both botnet and normal traffic; however, the authors used only botnet traffic by filtering out botnet communication IP addresses provided by the dataset owner. Results indicated that increasing the training iterations, improves the performance of a method, where precision and recall of Blackhole and SmokeBot botnets achieve more than 85%. Overall, the technique reached an average accuracy of 99.18%, which is the highest accuracy achieved among all neural network-based anomaly detection methods. However, there is no baseline comparison of this method with others, as different methods were evaluated on different datasets.

Similar to [165] in image feature-based methods, Moore et al. [178] also proposed an intrusion detection system for cyber-physical systems, that maps the controller area network (CAN) data to 2D images. The method extracts network features (e.g. speed and acceleration of a vehicle) of a CAN which are then converted into 2D images. These images are input into CNN for the classification of abnormal traffic. Before image generation, frequency Fourier Transformation is used to perform dimensionality reduction by representing the physical relationships between features. Authors used data from one vehicle network with six different traces. The first five traces were used as labelled data to train a CNN, while the sixth trace was used to test the classification accuracy. Authors reported the accuracy of 93% after ten epochs.

D. Summary of computer vision-based traffic anomaly detection:

In this section, we surveyed existing computer vision-based approaches for detecting network traffic anomalies. The motivation to use computer vision methods, as reported by several papers [142, 153, 158], is to overcome the problem of traditional detection approaches. Most of the conventional approaches are unable to characterise and detect specific types of attacks such low rate attacks, replacement attacks or zero-day attacks. Computer-vision methods help in alleviating such problems by utilising pattern recognition or image processing algorithms such as Hough transformation and Canny Edge Detection, where detection of a single line, edge, or a high-intensity pixel may represent abnormalities.

We categorised the work on computer vision-based methods under three topics; *i) image representation-based methods, ii) image feature-based methods,* and *iii) neural network based-methods.* Image representation-based methods monitor network packet headers at regular intervals, generate images of traffic usage, and then correlate/match current network traffic image with previously generated images. The deviated images can be further analyzed to detect attack type by using pattern recognition algorithms mentioned above. Image feature-based methods extract features such as entropy, contrast, pixel randomness, local homogeneity, and Triangle Area Maps (TAM), from the network traffic packets and then input into a distance metric or a classifier for anomaly detection. Similar to phishing and malware detection, neural network-based methods in anomaly detection also generates images

from network traffic and then fed into a deep neural network. Akin to the above two categories, neural-network is also a recent trend in anomaly detection. We summarise all the work we surveyed in Table III.

We next discuss some challenges and limitations of existing work in computer vision-based traffic anomaly detection that need to be addressed to make the research more effective and deployable in the network industry.

- **Challenges in Computer Vision-based Anomaly Detection Methods:** While computer vision methods have extensively been investigated to detect phishing and malware attacks, we find limited work on the detection of network traffic anomalies. As mentioned earlier, one possible reason could be that network traffic is a non-trivial application for computer vision methods. It is not straightforward to convert network traffic flows into images in a meaningful way, and in some cases, may not provide additional benefits compared to statistical anomaly detection methods. Moreover, network traffic uses several protocols such as TCP, UDP, HTTP, and HTTPS, that have different requirement and configurations. Therefore, a generic computer vision solution is harder to implement. Consider an email service provider that wants to detects anomalies such as self-propagating emails which spread through email virus or worms. The anomaly detection approach, in this case, requires images based on features mainly including several connections from a single host, and email header bytes. On the other hand, a video streaming service provider (e.g. Youtube) is more interested in deploying anomaly detection approaches based on traffic volume/flow to its servers. Thus, a generic computer vision-based solution requires further extensive research and testing. This area may further grow with the advances in unsupervised deep learning methods that is currently an active area of study.

- **Detailed Investigation of Anomalous Traffic:** Though existing computer-vision based anomaly detection methods can detect abnormal behaviour in network traffic with high accuracy, we find that existing work did not provide further insights on different types of anomalies such as botnets, DoS, backdoor, worms, and scanning. The methods generalise all these attack types as anomalous behaviours. However, in reality, each of these attack types has different characteristics which can have different impacts. Thus, there is a need to perform a detailed analysis of the detected anomalies that may highlight specific trends and patterns of varying anomaly types.

- **Traffic Anomaly Datasets:** We found a number of datasets that have been used by existing methods for anomaly detection such as NSL-KDD [160], KERONet [149], CAIDA [156, 157], and ISCX [177]. The majority of work used the NSL-KDD dataset [160]; however, its ability to reflect real-world conditions has been often criticised [179, 180]. The primary concern was the inability to evaluate anomaly detection systems against current and evolving anomalies and network traffic patterns. In this regard, CAIDA [156, 157] and

TABLE III: Traffic Anomaly Detection Approaches

Ref	Techniques Applied	Image Structure	Dataset (Normal/ Anomaly)	Description	Performance Metrics
Image Representation-based Methods					
Kim et al. [147, 148]	DCT, Mean Square Error	4×4 DCT coefficients	3616/729	Visualises and detects anomalies in real-time by passively monitoring packet headers and then converts them into video frames using DCT and absolute differences of pixels.	**Source Address-** Accuracy- 92.3%, FP- 0.3% **Destination Address-** Accuracy - 87.2%, FP- 0.75%
Kim et al. [150]	DCT, pixels intensities	8×8 DCT coefficients	3563/ 782	Visualises and detects anomalies using DCT coefficients and pixel intensities in real-time.	**Real-time analysis** TP-92.8%, FP-0.42% **Postmortem analysis** TP-94.8%, FP- 0.20%
Fontugne et al. [142]	Hough Transform	Image of network packets with time sliding window	630/ 2000	Detects anomalies by converting traffic into images and then applies Hough transform to identify lines representing DoS or port scanning attacks.	Accuracy - 73.8%
Jeong et al. [153]	Hough Transform, CCL	Image of network packets at fixed sampling rate	70,000 events	Detects anomalies using two tier architecture, where first tier detects random attacks such as DoS and second tier detects semi-random attacks such as port scanning.	Not Given
Kim et al. [155]	Canny Edge Detection, Pivoted Movement Algorithm	2-Dimensional Image with upper 24 bit and lower 8 bit IPv4/v6 addresses	CAIDA Ipv4/Routed and IPv6 dataset	Detects DDoS attack with IP subnet spoofing using canny edge detection and pivoted algorithm in a self-developed NetSCENE tool.	FPR: 1% FNR: 0.5%
Image Feature-based Methods					
Tan et al. [158]	EMD, PCA, MCA	2D feature matrix (TAM)	**KDD cup 99:** 97,260/ 3,892,550 **ISCX 2012 IDS:** 2,450,329/ 8720	Detects DoS attacks by extracting image-related network features and then applies EMD algorithm on features.	**KDD cup 99:** Accuracy - 99.95%, FP-1.93% **ISCX 2012 IDS:** Accuracy - 90.12%, FP-7.92%
Zou et al. [165]	GLCM texture features, k-NN, Euclidean Distance	Size of the image is equal to number of CPS devices in the system	100 CPS device traffic	Detects anomalies in a CPS by conducting a spatial and temporal analysis of network resources and using textural features.	Accuracy - 99%
Neural Networks-based Methods					
Wang et al. [170]	ANN, SAE	1 byte = 1 pixel	0.3 million data flow records	Utilises a deep neural network technique to extract and select features significant for detecting unknown and anomalous network protocols in a network traffic.	Unknown Protocol Identification rate - 54.9%
Li et al. [171]	ResNet50, GoogLeNet as CNN	8×8 input image	**NSL-KDD:** **Test$^+$:** 9711/12833 **Test^{-21}:** 2152/9698	Detects intrusions by converting raw packets to images and then input into CNN for automatic feature learning.	**ResNet 50:** **Test$^+$:** Accuracy - 79.14% **Test^{-21}:** Accuracy - 81.57% **GoogleNet:** **Test$^+$:** Accuracy - 77.04% **Test^{-21}:** Accuracy - 81.84%
Wang et al. [173]	CNN	78*78 image (784 bytes image)	100GB (25,400,44 network communication data records)	Converts network flow to images and then feeds into CNN for malicious network identification using representation learning.	Accuracy - 97.3%, F1 - 98.5%, Precision - 98.65%, Recall - 98.4%
Zhou et al. [175]	CNN	1 byte = 1 pixel	7.93GB network flow data	Classifies botnet encrypted and unencrypted network traffic by using CNN-based features self-learning and images as input.	Accuracy - 99.18%
Moore et al. [178]	CNN	18x25 pixel image	6 network traces from one vehicle network	Detects cyber attacks on a CAN bus network by exploiting the features related to physics such as frequency domain.	Accuracy - 93%

WIDE [151] datasets are being periodically updated as compared to KDD dataset. Though, it is always good to have a range of datasets for testing and validation purposes; a benchmark is helpful in setting as a standard. In addition to the standard dataset, continuously evolving datasets with properties such as modifiable, extensible, reproducible, are also required to include new network trends. Nevertheless, we do not find any dataset that serves the purposes stated above.

- **Convolutional Neural Networks (CNN):** We found that the use of CNNs is an emerging trend for detecting anomalous traffic. Several recent work utilised CNNs for network anomaly detection where the image is taken as an input. While CNN is more common neural network model among several works, we also find SAE being used by one method [170]. This direction is expected to grow with further advancements in CNN-based methods.

VI. PROTOTYPE AND COMMERCIAL SOLUTIONS

As we mentioned in the introduction, the use of computer vision methods in network security is not purely academic interest. There is an increasing trend from the industry to use such methods to overcome the limitations of the exiting solutions and obtain a competitive advantage. In this section, we survey computer vision-based tools/software that are offering network security-based services. Tools we study are either developed by individuals working under the capacity of developers or researchers (prototype solutions), or by companies who are trying to improve their existing network security products (commercial solutions). We highlight that

this product survey is carried out based on publicly available information. Thus, indeed this list is not comprehensive as many companies might not release the inner workings of their products publicly due to intellectual property reasons. Also, in many cases, the available public information is limited to provide an in-depth analysis of the tools and products. Our search resulted in 12 such tools/products as enumerated below.

i) Blazar [181] is a tool offered by Endgame Research [182] to detect Homoglyph attacks using a Siamese Neural Network. Homoglyph attacks are a type of phishing attacks where the users are prompted to with a ULR that appear visually similar to a legitimate website (e.g. `www.google.com` vs. `www.g0ogle.com`). The key idea of Blazar is to convert text URLs to images and feed them to a Siamese Neural Network that can capture visual similarities between characters. This approach works better compared to edit distance based text similarity methods, mainly because of extended character sets. The model was trained using a synthesised dataset that contained pairs of images of popular domain names and modified versions of them that appear visually similar using various extended characters. Blazar also implements a KD tree-based indexing and lookup system to improve performance [183]. Blazar is released in GitHub as an open-source tool.

ii) SpeedGrapher [184] is another tool by Endgame Research that detects macro-enabled document-based phishing attacks. In such attacks, the attackers trick victims by requesting them to open malware-embedded documents. The company already had a labelled dataset of macro-based documents that are known to be malware and to develop SpeedGrapher they

generated screenshots of such documents using the preview from the Word Interop class [185]. From the screenshots, a range of features such as prominent colours, blur/blank areas, embedded characters, and icons (using YOLO CNN). Finally, a classification algorithm, Random Forest, is used to detect malicious macros in a document. According to Endgame Research, this is a product that is still under active development and as such no code is released to date.

iii) Lookout Phishing AI [186] is an internet monitoring tool that identifies early signals of a phishing attack and alerts its users of the potential harm. This tool incorporates a machine learning engine that continuously scans the URLs in a browser and watches the behaviour of the site. Phishing AI claims to use computer vision methods in their website behaviour analysis [186]. However, the exact details are not provided on the company website. Recently Phishing AI was able to detect a large scale phishing campaign targeting non-governmental organizations such as UNICEF [187].

iv) Phish.ai [188] is a solution that can detect zero-day phishing attacks. Similar to many solutions discussed in Section III, Phish.ai keep a visual signature database extracted from the screenshots of top websites and brands. The API accepts the links as the input and obtain a screenshot of the website and compare its visual signature against ones stored in the database to check whether the new site is trying to impersonate an existing, legitimate site. While it might be interesting to know what types of visual signatures are used in this product, such information is not disclosed by the company.

v) MalSee [189] is a malware detection and classification system that is developed by researchers at Mayachitra Inc. It is similar to the ideas presented in the early work by Nataraj et al. [100], where the malware binaries are represented as images. The tool is offered as a web-based API and claims to provide ∼1,000x speedup detection of harmful computer viruses compared to existing methods.

vi) Anti-Pixm [18] by Pixm Inc. is another phishing detection tool that uses deep learning and computer vision to analyze the web pages visually. Compared to other similar products, Anti-Pixm claims to provide real-time, end-device based protection against stealth mode phishing attacks (i.e. the phishing attacks against cloud-based email security tools). When a user clicks a link received in an email, Anti-Pixm browser extension quickly grabs a screenshot in the background, compares its visual signature with common targets, and blocks the URL within less than a second. The tool claims to find phishing attacks before they are being identified by major cloud-based phishing solutions that maintain blacklists. The exact vision-based deep learning methods used by Pixm is not publicly available.

vii) Ironscales's IronShield [190] is a cloud-based solution to defend against real-time malware, credential threats, account takeovers, and phishing websites. One specific aspect of the solution is to validate the legitimacy of login blocks using computer vision methods.

viii) INKY [17] is a cloud-based email security solution that blocks spam, malware, and phishing attacks. One of the main features the solution provides is detecting peculiarities and brand forgeries in email graphics and icons. INKY offers seamless integration with Office 365, Microsoft Exchange Solutions, Microsoft Security, and Google G Suit so that phishing emails can be analysed and contained in realtime.

Cyberfish [191], **GreatHorn [192]**, **Area1 [193]**, and **Zero-Fox [194]** are similar phishing email detection solutions that use visual signatures. No significant information is available about these solutions apart from the fact that they are using computer vision methods to match visually similar web pages.

One key observation regarding commercial tools is that most of the solutions are developed in recent years despite there are published academic work since the mid-2000s as our survey found. One possible reason for this can be the exponential growth in deep learning-driven computer vision methods happened in the last few years making computer vision accessible for other domain experts. Also, such attempts are further exacerbated by the ubiquitous availability of deep learning and computer vision libraries such as Tensorflow, PyTorch, OpenCV, and SciKit Learn that enables rapid prototyping and testing without any specialised knowledge. Thus, we believe that in future, there will be more of these types of security solutions as well as research work.

Another interesting observation is that apart from Malsee [189] all the rest of the products we found are built to detect phishing attacks. However, when it comes to academic work, both phishing detection and malware detection had a significant body of work that used computer vision methods. Furthermore, we did not find any commercial or prototype solution that used computer vision in anomaly detection. While there is no conclusive evidence on why this is the case, we believe that from an applied research point of view, using computer vision techniques for phishing detection is conceptually straightforward compared to malware detection or traffic anomaly detection. As computer vision methods further advances, we believe the information security community will start adapting those ideas in malware and traffic anomaly detection as well. Finally, only very few tools were available as open-source tools. One thing that can be done to further enhance the research in this domain is to release the code and the data together with the academic publication.

VII. Open Issues and Research Challenges

Computer vision methods for network security is an evolving research area, and it is likely to grow further with the exponential advances in deep learning-driven computer vision methods. We also provided the evidence that the information security industry has also started to adapt of build security products or features adapting the ideas from computer vision, mainly to address the limitations of existing solutions and to obtain a competitive advantage. As a growing field, there are several exciting research directions and challenges to solve, to build more robust security solutions using computer vision methods. We next discuss such opportunities and challenges.

TABLE IV: Computer Vision Based Network Security Products and Prototypes

Product Name	Security Service	Deployment Level	Open Source?	Description
Blazar [181]	Homoglyph Attacks (Domain/URLs and File Names) Detection	Browser and Network Infrastructure	Yes	Utilises Siamese Convolutional Neural Network and Euclidean Distance to measure similarity between original and malicious URLs/file names.
SpeedGrapher [184]	Detects macro-enabled document based phishing	Desktop for Documents e.g. Microsoft Word	No	Utilises various computer vision techniques such as Blur Detection, Blank Detection, Optical Character Recognition (OCR), and Icon Detection to extract features from a document.
Lookout Phishing AI [186]	Detect Phishing Websites	Browser	No	Utilises computer vision techniques to identify the phishing sites by analyzing the use of logos and graphics.
Phish.ai [188]	Detects zero-day phishing attacks	Customized API and web Browser	Yes	Creates a unique and up-to-date computer-vision based database of legitimate websites and afterwards compares a screenshot of suspicious website with the database.
MalSee [189]	Malware Detection	Web-accessible service	No	Utilises pattern recognition from image processing methods to detect and classify a malware.
Pixm [18]	Detect Phishing Websites	Browser	No	Uses computer vision techniques to compare screenshots of websites for possible phishing detection.
Ironscales' IronShield [190]	Provide multiple security services such as phishing detection, zero-day malware detection, Malicious Email Detection and credential thefts	Cloud level Network Infrastructure, Browser, and Email	No	Utilises computer vision to detect in real-time visual deviations and determines whether or not a login page is legitimate.
INKY [17]	Block spam, malware and phishing attacks	Cloud based network infrastructure, and Email, Browser	No	Uses computer vision to spot differences in email graphics and iconography to detect forgeries.
Cyberfish [191]	Detect phishing attacks in emails and web pages	Cloud and On-Premise Email and Browser	No	Analyzes visual representation of emails and web pages to detect attacks in real-time.
GreatHorn [192]	Credential Theft Detection	Email	No	Utilises CNN and image representation-based computer vision methods to detect credential thefts via email.
Area1 [193]	Phishing detection on multiple platforms	Email, Web and Network	No	Utilises AI and computer vision methods to provide effective phishing protection.
ZeroFox [194]	Risk identification	Web/Domain, Social Media Accounts	No	Utilises computer vision techniques like OCR, logo and face detection for risk identification in multiple platforms such as social media, web and domain.

A. Hybrid Solutions

Almost all the work we discussed, used computer vision methods and compared them against traditional methods. We find only a very limited literature [92, 107, 111] that propose a hybrid approach of combining conventional and computer vision methods. Nonetheless, it is in favour of security defenders to deploy both types of methods to overcome the limitations of each other. For example, Fatt et al. [54] used *Favicons* for phishing detection. However, this work does not apply to web pages that do not have *Favicons*. In such a case, traditional text or network-based methods can be merged into this approach. PhishZoo [63] combined image and text-based methods and showed an overall accuracy of 90.2% in phishing detection.

We find similar examples in malware detection where a single type of method does not holistically cover all aspects of detection and classification. For instance, the method proposed by Nataraj et al. [100] focused only on the visual representations of a resource section in a malware executable file, which allows an adversary to bypass a detection system if malware exists in other parts of a file. In this case, the use of static-analysis methods (e.g. signature-based) may help in detecting malware that is relocated to different sections of a file. Similarly, the resource limitation in running vision-based dynamic analysis on mobile devices can be addressed by first performing static analysis followed by methods such as [120]. Thus, we see the need for more hybrid security solutions that combine traditional techniques with computer vision methods for better reliability and accuracy.

B. Adversarial Example Attacks

Several works pointed out that computer vision systems, such as object detection and classification solutions, are vulnerable to adversarial examples. That is, an attacker can craft inputs that can trigger misclassifications, *evasion attacks* [195]–[197] or feed the training process with adversarial data so that the model is biased towards a direction set by the attacker, *poisoning attacks* [198]–[200]. Such attacks have been demonstrated extensively for traditional computer vision or machine learning settings. However, recently multiple works showed such vulnerabilities exist in security solutions.

Example work for *evasion attacks* include Al-Dujaili et al. [201] and Grosse et al. [202] that demonstrated the generation of functional malware that can act as adversarial examples for model-based malware classifiers. Hu et al. [203] proposed MalGAN, a GAN-based generative model capable of generating malware samples that can bypass classifiers such as random forest and logistic regression. Rosenberg et al. [204] is another comparable work.

Recent work [205] also showed how malware traffic could be synthesised to bypass intrusion detection systems in practical scenarios. Authors tried to mimic the behaviour of Facebook chat traffic by feeding its network flow parameters to GAN for a predefined number of epochs. The trained GAN generator then communicates its output parameters to the command and control server so that malware can change its traffic accordingly misguiding middleboxes to think malware traffic as Facebook chat traffic.

Biggio et al. [206] demonstrated *poisoning attacks* on spam filtering and intrusion detection scenarios where an adversary is capable of controlling a subset of samples that are used to train or update a classifier. The attacker carefully constructs these samples such that the classifier can be misguided. For instance, in spam filtering attack, adversaries can modify spam emails by adding some non-suspicious words which are likely

to appear in legitimate emails whereas, in intrusion detection, adversaries may inject poisoning pattern in the network that matches with legitimate activities.

As vision-based network security solutions become more and more mainstream, attackers will inevitably be exploring adversarial attacks. These attacks are exacerbated by the easy access to software libraries that can generate such attacks (e.g. [207]). As our survey indicated, at this stage, the research focus is more on getting working systems with the required levels of accuracy. None of the work we presented in Section III, IV, and V looked into the attacks against their systems. Assessing the adversarial robustness of classifiers and coming up with defence strategies is an active area of research in computer vision and machine learning [208, 209]. We believe that in the next phase of research in vision-based network security must look into the robustness of classifiers by exploring the defensive mechanisms proposed in computer visions systems such as adversarial training, model distillation, feature squeezing, gradient hiding, and blocking transferability [210] and adapt those to network security solutions.

C. One-shot and Few-shot Learning

Most of the methods that we discussed in the above sections used deep learning models such as CNN that are well known for their high training data requirements. In object classification or image recognition, it is easy to collect a large volume of data. However, collecting an equivalent amount of data from security events is somewhat tricky. For instance, in the case of malware and traffic anomalies, it is not practical to obtain large volumes of samples as these events are rare and often go unnoticed. Moreover, malware samples and security incident data (e.g. logs and configurations) are usually not publicly released for safety reasons as well as to protect corporate secrets. Finally, to build labelled datasets for security problems, specialised domain experts are required compared to image labelling or voice transcribing tasks. Although crowdsourcing has been tried (e.g. PhishTank), the ensuring quality remains challenging. Hence it is necessary to explore solutions that require lesser volumes of labelled data.

A possible solution to this problem is the ideas of one/few-shot learning [211]–[218] and unsupervised feature learning/self-supervised learning [219, 220] where models learn from minimal or no labelled data. This is an active area of research in computer vision. However, its application in vision-based security solutions remains mostly unexplored. There have been recent attempts in classifying malware through few/one-shot learning models [221]–[223]. Tran et al. [222] proposed a few-shot learning approach to malware classification by first converting malware binaries into greyscale images and then adapting meta-learning models (i.e. matching networks and prototypical networks). The results showed an average accuracy of 92.4% with a prototypical network for1-shot learning and 95.3% accuracy with 4-shot learning. Similarly, Hsiao et al. [221] used one-shot Siamese neural network to classify unknown malware images. Atapour-Abarghouei et al. [223] also proposed a ransomware classification approach using one-shot learning through data augmentation.

Nevertheless, few-shot learning is still an emerging area in cybersecurity, and extensive research must be conducted to understand how such methods can be effectively utilised.

D. Open Set for Classification

Majority of the work we discussed operates under a closed set assumption. For example, a malware classifier will classify a given sample as one of the known malware families or as benign. If a sample from a new malware family appears, the classifier will still make a similar decision. Similarly, the same behaviour will happen for a new benign sample as well. In fact, in such settings, there is a possibility that the classifier is making very confident yet highly inaccurate decision [224, 225] as the classifier is not trained to handle *unknown* (i.e. data samples that are out of the distribution of the training set) data samples [226].

While classifiers that can't handle unknown inputs are sufficient for prototyping, in real-world deployments, often there will be inputs that are entirely outside the distributions of the training data the classifier has seen. This situation may not be critical in some applications like photo tagging. However, when it comes to network security, the ramifications of such solutions can be disastrous. This problem also becomes challenging because of the incomplete knowledge of the world during training (i.e. only the known classes are accessible). Thus, steps must be included to make sure the model knows its limitation on the closed set operation. None of the work we surveyed addressed this challenge, and we believe as the vision-based network security solutions become more mainstream. this is an area that should show growth.

An emerging solution to correctly identify known and unknown samples is through *open-set recognition*. Open set recognition describes a scenario where new classes that were not seen in the training phase appear in testing and requires the classifiers to not only accurately classify known classes but also effectively deal with unknown ones. In the most basic sense, classes utilised in testing are not present in training [224]. This is currently an active area of research in computer vision [227]–[232]. For instance, Dhamija et al. [227, 228] introduced a novel loss function and regularisation methods to improve the handling of background and unknown inputs. Oza et al. [229] proposed an approach based class conditioned auto-encoders and using the reconstruction loss for open set recognition. Similarly, Liu et al. [230] introduced Open Long-Tailed Recognition (OLTR) that handles imbalanced classification, few-shot learning, and open-set recognition.

Some recent work in network security has started looking into the open set classification. Cruz et al. [233] detected intrusions in the network by considering both closed set and open set recognition problem. Authors proposed a fine-grained recognition approach that considers all intrusion detection problem as a recognition problem and categorically assigns intrusion detection as a closed set problem and intrusion recognition as an open set problem. Authors used a loosely open set dataset which is labelled according to individual intrusion types (e.g. (sendmail, snmp guess), rather than more general attack categories (e.g., DoS). To validate the approach, two

different classifier types are employed - Gaussian RBF kernel SVMs, which are not theoretically guaranteed to bound open space risk, and W-SVMs, which are theoretically guaranteed to bound open space risk. Results showed that W-SVM offers better performance in classifying unknown classes.

Similarly, Henrydoss et al. [234] proposed an approach that recognises unknown intrusion classes by understanding the characteristics and also performs incremental learning for updating the classifiers using feedback from the classifier itself after successfully identifying the unknown classes during query time. Authors used a multi-class classifier that is proposed explicitly for open set recognition setting and utilised Extreme Value Machine (EVM) to classify intrusion detection data. The EVM has a capability to perform kernel-free, nonlinear, variable bandwidth outlier detection in combination with incremental learning. We recommend interested readers to look into a survey paper from Rudd et al. [235] for open set recognition approaches in intrusion detection. Incorporating open-set recognition techniques to network security solutions is vital, and we are still at the early stages of this research direction. We believe this is as an open yet challenging area that needs a significant amount of further work.

E. Online Learning and Real-time Predictions

Many of the work that we surveyed were operating on offline mode where a data collection phase followed by a training phase. In the real-world, often network security systems must deal with real-time data, rapid fluctuations, and trends. Also, models need to be re-trained frequently with the availability of new data. As such, when models are deployed, there must be mechanisms to update them at short intervals using methods such as online learning [236, 237]. However, there are many challenges associated with such setups. Catastrophic forgetting [238, 239] is one of the challenges where the learning model might forget its prior knowledge when fine-tuned with new data. This is widely believed to be a serious problem for neural networks [240]. Also, in online learning, designs must consider poisoning attacks as well.

Another aspect of handling real-time data, especially related to traffic anomaly detection is how to do traffic flow sampling accurately. In large networks, it is not practical to analyse every traffic flow. As a result, it might be essential to explore the work proposed in traffic sampling [241]–[244] and integrate them into the training pipeline.

VIII. Conclusion

Network security threats are becoming increasingly common as well as sophisticated to a point where not only traditional security solutions but also machine learning-based security solutions are becoming less effective. As an alternative, academic research and commercial solutions providers are exploring the feasibility of using methods from computer vision in combination with machine learning models. We surveyed such work under three broad topics; *phishing detection*, *malware detection*, and *traffic anomaly detection*. Overall our survey showed that there are distinct advantages of using computer vision methods, especially when it comes to detecting

zero-day attacks as well as developing more scalable and accurate phishing detection systems. Also, our survey found that in order to stem further research in this area, it is necessary to establish much larger and more recent security datasets that can be conveniently used to establish benchmarks and compare different solutions. Finally, we discussed the potential research directions in this domain such as hybrid solutions combining traditional methods with computer vision methods, building solutions that are resilient of adversarial attacks, few-shot learning to address the challenges in obtaining large samples of data, and open set classification which is required when it comes to real world deployments.

References

[1] L. Strain, "2018: The year of the data breach tsunami," 2019. [Online]. Available: https://blog.malwarebytes.com/101/2018/12/2018-the-year-of-the-data-breach-tsunami/

[2] "Internet security threat report," 2018. [Online]. Available: https://www.symantec.com/content/dam/symantec/docs/reports/istr-23-2018-en.pdf

[3] "Gartner forecasts worldwide information security spending to exceed $124 billion in 2019," 2018. [Online]. Available: https://www.gartner.com/en/newsroom/press-releases/2018-08-15-gartner-forecasts-worldwide-information-security-spending-to-exceed-124-billion-in-2019

[4] "2018 reform of EU data protection rules," 2018. [Online]. Available: https://ec.europa.eu/commission/priorities/justiceandfundamental-rights/dataprotection/2018reformeudataprotectionrules_en

[5] "Mandatory data breach notification," 2017. [Online]. Available: https://www.oaic.gov.au/media-and-speeches/statements/mandatory-data-breach-notification

[6] A. L. Buczak and E. Guven, "A survey of data mining and machine learning methods for cyber security intrusion detection," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 2, pp. 1153–1176, 2015.

[7] J. Gardiner and S. Nagaraja, "On the security of machine learning in malware C&C detection: A survey," *ACM Computing Surveys (CSUR)*, vol. 49, no. 3, p. 59, 2016.

[8] Y. Xin, L. Kong, Z. Liu, Y. Chen, Y. Li, H. Zhu, M. Gao, H. Hou, and C. Wang, "Machine learning and deep learning methods for cybersecurity," *IEEE Access*, vol. 6, pp. 35 365–35 381, 2018.

[9] R. Sommer and V. Paxson, "Outside the closed world: On using machine learning for network intrusion detection," in *2010 IEEE symposium on security and privacy*. IEEE, 2010, pp. 305–316.

[10] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld, "Face recognition: A literature survey," *ACM computing surveys (CSUR)*, vol. 35, no. 4, pp. 399–458, 2003.

[11] A. F. Abate, M. Nappi, D. Riccio, and G. Sabatino, "2D and 3D face recognition: A survey," *Pattern recognition letters*, vol. 28, no. 14, pp. 1885–1906, 2007.

[12] J. Daugman, "How iris recognition works," in *The essential guide to image processing*. Elsevier, 2009, pp. 715–739.

[13] ——, "New methods in iris recognition," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 37, no. 5, pp. 1167–1175, 2007.

[14] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, no. 2, pp. 91–110, 2004.

[15] H. Bay, T. Tuytelaars, and L. Van Gool, "Surf: Speeded up robust features," in *European conference on computer vision*. Springer, 2006, pp. 404–417.

[16] A. Fiscutean, "12 of the hottest startups at the RSA conference 2019," 2019. [Online]. Available: https://www.csoonline.com/article/3342162/12-of-the-hottest-startups-at-the-rsa-conference-2019.html

[17] "INKY: The next generation of email protection," 2019. [Online]. Available: https://inky.com/solutions/

[18] "Worlds first phishing protection using computer vision," 2019. [Online]. Available: https://pixm.net

[19] L. Ackerson, "INKY announces additional $6 million in funding led by clearsky security to extend technology availability and strategic partnerships," 2019. [Online]. Available:

https://www.prnewswire.com/news-releases/inky-announces-additional-6-million-in-funding-led-by-clearsky-security-to-extend-technology-availability-and-strategic-partnerships-300957916.html

[20] "Pixm takes on phishing attacks with deep learning using Apache MXNet on AWS," 2018. [Online]. Available: https://www.wired.com/brandlab/2018/08/phishing-attacks-with-deep-learning-using-apache-mxnet-on-aws/

[21] "ENISA threat landscape report 2018," 2019. [Online]. Available: https://www.enisa.europa.eu/publications/enisa-threat-landscape-report-2018

[22] R. Jafri and H. R. Arabnia, "A survey of face recognition techniques." *Jips*, vol. 5, no. 2, pp. 41–68, 2009.

[23] J. Hayes and G. Danezis, "Generating steganographic images via adversarial training," in *Advances in Neural Information Processing Systems*, 2017, pp. 1954–1963.

[24] D. Volkhonskiy, I. Nazarov, B. Borisenko, and E. Burnaev, "Steganographic generative adversarial networks," *arXiv preprint arXiv:1703.05502*, 2017.

[25] B. Li, J. He, J. Huang, and Y. Q. Shi, "A survey on image steganography and steganalysis," *Journal of Information Hiding and Multimedia Signal Processing*, vol. 2, no. 2, pp. 142–172, 2011.

[26] A. Cheddad, J. Condell, K. Curran, and P. Mc Kevitt, "Digital image steganography: Survey and analysis of current methods," *Signal processing*, vol. 90, no. 3, pp. 727–752, 2010.

[27] "2019 phishing trends and intelligence report," 2019. [Online]. Available: https://info.phishlabs.com/hubfs/2019%20PTI%20Report/2019%20Phishing%20Trends%20and%20Intelligence%20Report.pdf

[28] "Business emails compromise the 12 billion dollar scam," 2018. [Online]. Available: https://www.ic3.gov/media/2018/180712.aspx

[29] M. Khonji, Y. Iraqi, and A. Jones, "Phishing detection: a literature survey," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 4, pp. 2091–2121, 2013.

[30] A. Aleroud and L. Zhou, "Phishing environments, techniques, and countermeasures: A survey," *Computers & Security*, vol. 68, pp. 160–196, 2017.

[31] D. Palmer, "Phishing attacks: How hunting down fake websites is making life harder for hackers," 2018. [Online]. Available: https://www.zdnet.com/article/phishing-attacks-how-hunting-down-fake-websites-is-making-life-harder-for-hackers/

[32] "The AV-TEST security report 2017/2018," 2018. [Online]. Available: https://www.av-test.org/fileadmin/pdf/security_report/AV-TEST_Security_Report_20172018.pdf

[33] "Petya ransomware outbreak goes global," 2017. [Online]. Available: https://krebsonsecurity.com/2017/06/petya-ransomware-outbreak-goes-global/

[34] E. Gandotra, D. Bansal, and S. Sofat, "Malware analysis and classification: A survey," *Journal of Information Security*, vol. 5, no. 02, p. 56, 2014.

[35] Y. Ye, T. Li, D. Adjeroh, and S. S. Iyengar, "A survey on malware detection using data mining techniques," *ACM Computing Surveys (CSUR)*, vol. 50, no. 3, p. 41, 2017.

[36] Sentryo, "The mirai IoT botnet, a publically available turn-key threat," 2016. [Online]. Available: https://www.sentryo.net/the-mirai-iot-botnet-a-publically-available-turn-key-threat-2/

[37] E. B. Oleg Kupreev and A. Gutnikov, "DDoS attacks in Q4 2018," 2019. [Online]. Available: https://securelist.com/ddos-report-q1-2019/90792/

[38] W. Liu, X. Deng, G. Huang, and A. Y. Fu, "An antiphishing strategy based on visual similarity assessment," *IEEE Internet Computing*, vol. 10, pp. 58–65, 2006.

[39] L. Wenyin, G. Huang, L. Xiaoyue, X. Deng, and Z. Min, "Phishing web page detection," in *Eighth International Conference on Document Analysis and Recognition (ICDAR'05)*. IEEE, 2005, pp. 560–564.

[40] Y. Liu, W. Liu, and C. Jiang, "User interest detection on web pages for building personalized information agent," in *International Conference on Web-Age Information Management*. Springer, 2004, pp. 280–290.

[41] E. Kirda and C. Kruegel, "Protecting users against phishing attacks," *The Computer Journal*, vol. 49, no. 5, pp. 554–561, 2006.

[42] E. Medvet, E. Kirda, and C. Kruegel, "Visual-similarity-based phishing detection," in *4th International Conference on Security and Privacy in Communication Netowrks*, 2008, pp. 221–226.

[43] E. Kirda and C. Kruegel, "Protecting users against phishing attacks with AntiPhish," in *29th Annual International Computer Software and Applications Conference*, vol. 1, July 2005, pp. 517–524 Vol. 2.

[44] A. P. E. Rosiello, E. Kirda, . Kruegel, and F. Ferrandi, "A layout-similarity-based approach for detecting phishing pages," in *2007 Third International Conference on Security and Privacy in Communications Networks and the Workshops*, Sep. 2007, pp. 454–463.

[45] A. Y. Fu, L. Wenyin, and X. Deng, "Detecting phishing web pages with visual similarity assessment based on earth mover's distance (EMD)," *IEEE Transactions on Dependable and Secure Computing*, vol. 3, no. 4, pp. 301–311, 2006.

[46] K. Turkowski, "Filters for common resampling tasks," in *Graphics gems*. Academic Press Professional, Inc., 1990, pp. 147–165.

[47] H. Zhang, G. Liu, T. W. S. Chow, and W. Liu, "Textual and visual content-based anti-phishing: A bayesian approach," *IEEE Transactions on Neural Networks*, vol. 22, no. 10, pp. 1532–1546, Oct 2011.

[48] K. Chen, J. Chen, C. Huang, and C. Chen, "Fighting phishing with discriminative keypoint features," *IEEE Internet Computing*, vol. 13, no. 3, pp. 56–63, 2009.

[49] K. Mikolajczyk and C. Schmid, "Indexing based on scale invariant interest points," in *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, vol. 1. IEEE, 2001, pp. 525–531.

[50] C. R. Huang, C. S. Chen, and P. C. Chung, "Contrast context histogram-a discriminating local descriptor for image matching," in *18th International Conference on Pattern Recognition (ICPR'06)*, vol. 4. IEEE, 2006, pp. 53–56.

[51] I. Corona, B. Biggio, M. Contini, L. Piras, R. Corda, M. Mereu, G. Mureddu, D. Ariu, and F. Roli, "Deltaphish: Detecting phishing webpages in compromised websites," *CoRR*, vol. abs/1707.00317, 2017. [Online]. Available: http://arxiv.org/abs/1707.00317

[52] I. F. Lam, W. C. Xiao, S. C. Wang, and K. T. Chen, "Counteracting phishing page polymorphism: An image layout analysis approach," in *International Conference on Information Security and Assurance*. Springer, 2009, pp. 270–279.

[53] M. Dunlop, S. Groat, and D. Shelly, "GoldPhish: Using images for content-based phishing analysis," in *Fifth International Conference on Internet Monitoring and Protection*, 2010, pp. 123–128.

[54] J. C. S. Fatt, C. K. Leng, and S. S. Nah, "Phishdentity: Leverage website favicon to offset polymorphic phishing website," in *2014 Ninth International Conference on Availability, Reliability and Security*, Sep. 2014, pp. 114–119.

[55] K. Tian, S. T. Jan, H. Hu, D. Yao, and G. Wang, "Needle in a haystack: Tracking down elite phishing domains in the wild," in *Proceedings of the Internet Measurement Conference 2018*, 2018, pp. 429–442.

[56] M. Hara, A. Yamada, and Y. Miyake, "Visual similarity-based phishing detection without victim site information," in *2009 IEEE Symposium on Computational Intelligence in Cyber Security*, March 2009, pp. 30–36.

[57] G. Wang, H. Liu, S. Becerra, K. Wang, S. Belongie, H. Shacham, and S. Savage, "Verilogo: Proactive phishing detection via logo recognition," in *Technical Report t CS2011- 0969, University of Califonia, San Diego*, 2011.

[58] I. Fehérvári and S. Appalaraju, "Scalable logo recognition using proxies," in *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2019, pp. 715–725.

[59] Y. Wang, W. Yang, and H. Zhang, "Deep learning single logo recognition with data enhancement by shape context," in *2018 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2018, pp. 1–7.

[60] Y. Zhou, Y. Zhang, J. Xiao, Y. Wang, and W. Lin, "Visual similarity based anti-phishing with the combination of local and global features," in *IEEE 13th International Conference on Trust, Security and Privacy in Computing and Communications*, 2014, pp. 189–196.

[61] O. Chum, T. Pajdla, and P. Sturm, "The geometric error for homographies," *Computer Vision and Image Understanding*, vol. 97, no. 1, pp. 86–102, 2005.

[62] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.

[63] S. Afroz and R. Greenstadt, "PhishZoo: Detecting phishing websites by looking at them," in *IEEE Fifth International Conference on Semantic Computing*, 2011, pp. 368–375.

[64] N. Kumar, L. Zhang, and S. Nayar, "What is a good nearest neighbors algorithm for finding similar patches in images?" in *European conference on computer vision*. Springer, 2008, pp. 364–378.

[65] J. S. White, J. N. Matthews, and J. L. Stacy, "A method for the automated detection phishing websites through both site characteristics and image analysis," in *Cyber Sensing 2012*, vol. 8408. International Society for Optics and Photonics, 2012, p. 84080B.

[66] N. Fotiou, G. F. Marias, and G. C. Polyzos, "Fighting phishing the information-centric way," in *2012 5th International Conference on New Technologies, Mobility and Security (NTMS)*, May 2012, pp. 1–5.

[67] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.

[68] K. Borgolte, C. Kruegel, and G. Vigna, "Meerkat: Detecting website defacements through image-based object recognition," in *24th {USENIX} Security Symposium ({USENIX} Security 15)*, 2015, pp. 595–610.

[69] M. Adebowale, K. Lwin, E. SAnchez, and M. Hossain, "Intelligent web-phishing detection and protection scheme using integrated features of images, frames and text," *Expert Systems with Applications*, vol. 115, pp. 300 – 313, 2018.

[70] S. Abdelnabi, K. Krombholz, and M. Fritz, "WhiteNet: Phishing website detection by visual whitelists," *arXiv preprint arXiv:1909.00300*, 2019.

[71] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 815–823.

[72] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.

[73] A. Krizhevsky and G. Hinton, "Learning multiple layers of features from tiny images," Citeseer, Tech. Rep., 2009.

[74] J. Deng, W. Dong, R. Socher, L. J. Li, K. Li, and L. Fei Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.

[75] T. Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European conference on computer vision*. Springer, 2014, pp. 740–755.

[76] F. C. Dalgic, A. S. Bozkir, and M. Aydos, "Phish-IRIS: A new approach for vision based brand prediction of phishing web pages via compact visual descriptors," *CoRR*, vol. abs/1905.07767, 2019. [Online]. Available: http://arxiv.org/abs/1905.07767

[77] K. L. Chiew, E. H. Chang, C. L. Tan, J. Abdullah, K. Sheng, and C. Yong, "Building standard offline anti-phishing dataset for benchmarking," *International Journal of Engineering & Technology*, vol. 7, no. 4.31, pp. 7–14, 2018. [Online]. Available: https://www.sciencepubco.com/index.php/ijet/article/view/23333

[78] V. Falconieri, "Open dataset of phishing and tor hidden services screen-captures," *arXiv preprint arXiv:1908.02449*, 2019.

[79] G. Li and Y. Yu, "Deep contrast learning for salient object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 478–487.

[80] N. Liu and J. Han, "Dhsnet: Deep hierarchical saliency network for salient object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 678–686.

[81] X. Li, L. Zhao, L. Wei, M.-H. Yang, F. Wu, Y. Zhuang, H. Ling, and J. Wang, "Deepsaliency: Multi-task deep neural network model for salient object detection," *IEEE Transactions on Image Processing*, vol. 25, no. 8, pp. 3919–3930, 2016.

[82] E. Hoffer and N. Ailon, "Deep metric learning using triplet network," in *International Workshop on Similarity-Based Pattern Recognition*. Springer, 2015, pp. 84–92.

[83] J. Wan, D. Wang, S. C. H. Hoi, P. Wu, J. Zhu, Y. Zhang, and J. Li, "Deep learning for content-based image retrieval: A comprehensive study," in *Proceedings of the 22nd ACM international conference on Multimedia*. ACM, 2014, pp. 157–166.

[84] V. Erin Liong, J. Lu, G. Wang, P. Moulin, and J. Zhou, "Deep hashing for compact binary codes learning," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 2475–2483.

[85] K. Lin, H.-F. Yang, J.-H. Hsiao, and C.-S. Chen, "Deep learning of binary hash codes for fast image retrieval," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2015, pp. 27–35.

[86] H. Lai, Y. Pan, Y. Liu, and S. Yan, "Simultaneous feature learning and hash coding with deep neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3270–3278.

[87] "The Duqu 2.0," 2015. [Online]. Available: https://media.kasperskycontenthub.com/wp-content/uploads/sites/43/2018/03/07205202/The_Mystery_of_Duqu_2_0_a_sophisticated_cyberespionage_actor_returns.pdf

[88] "HAMMERTOSS: stealthy tactics define a Russian cyber threat group," *FireEye, Milpitas*, 2015.

[89] P. Trinius, T. Holz, J. Göbel, and F. C. Freiling, "Visual analysis of malware behavior using treemaps and thread graphs," in *2009 6th International Workshop on Visualization for Cyber Security*. IEEE, 2009, pp. 33–38.

[90] K. Han, J. H. Lim, and E. G. Im, "Malware analysis method using visualization of binary files," in *Proceedings of the 2013 Research in Adaptive and Convergent Systems*. ACM, 2013, pp. 317–321.

[91] D. Androutsos, K. N. Plataniotis, and A. N. Venetsanopoulos, "A novel vector-based approach to color image retrieval using a vector angular-based distance measure," *Computer Vision and Image Understanding*, vol. 75, no. 1-2, pp. 46–58, 1999.

[92] K. Han, B. Kang, and E. G. Im, "Malware analysis using visualized image matrices," *The Scientific World Journal*, vol. 2014, 2014.

[93] K. S. Han, J. H. Lim, B. Kang, and E. G. Im, "Malware analysis using visualized images and entropy graphs," *International Journal of Information Security*, vol. 14, no. 1, pp. 1–14, 2015.

[94] V. Bontchev, "Are good computer viruses still a bad idea?" [Online]. Available: vxheavens.com

[95] L. Nataraj, V. Yegneswaran, P. Porras, and J. Zhang, "A comparative assessment of malware classification using binary texture analysis and dynamic analysis," in *Proceedings of the 4th ACM Workshop on Security and Artificial Intelligence*. ACM, 2011, pp. 21–30.

[96] S. Z. M. Shaid and M. A. Maarof, "Malware behavior image for malware variant identification," in *2014 International Symposium on Biometrics and Security Technologies*. IEEE, 2014, pp. 238–243.

[97] J. P. Chapman, "{SAD}{THUG}: Structural anomaly detection for transmissions of high-value information using graphics," in *27th {USENIX} Security Symposium*, 2018, pp. 1147–1164.

[98] M. Sun, M. Li, and J. Lui, "DroidEagle: Seamless detection of visually similar android apps," in *Proceedings of the 8th ACM Conference on Security & Privacy in Wireless and Mobile Networks*, 2015, p. 9.

[99] P. Li, L. Liu, D. Gao, and M. K. Reiter, "On challenges in evaluating malware clustering," in *International Workshop on Recent Advances in Intrusion Detection*. Springer, 2010, pp. 238–255.

[100] L. Nataraj, S. Karthikeyan, G. Jacob, and B. S. Manjunath, "Malware images: Visualization and automatic classification," in *8th International Symposium on Visualization for Cyber Security*, ser. VizSec '11, 2011, pp. 4:1–4:7.

[101] "Anubis: Analyzing unknown binaries," 2011. [Online]. Available: https://seclab.cs.ucsb.edu/academic/projects/projects/anubis/

[102] K. Rieck, "Malheur dataset," 2016. [Online]. Available: https://www.sec.cs.tu-bs.de/data/malheur/

[103] K. Kancherla and S. Mukkamala, "Image visualization based malware detection," in *2013 IEEE Symposium on Computational Intelligence in Cyber Security (CICS)*. IEEE, 2013, pp. 40–44.

[104] A. Makandar and A. Patrot, "Malware analysis and classification using artificial neural network," in *2015 International conference on trends in automation, communications and computing technology (I-TACT-15)*. IEEE, 2015, pp. 1–6.

[105] ——, "Malware class recognition using image processing techniques," in *2017 International Conference on Data Management, Analytics and Innovation (ICDMAI)*. IEEE, 2017, pp. 76–80.

[106] ——, "Trojan malware image pattern classification," in *Proceedings of International Conference on Cognition and Recognition*. Springer, 2018, pp. 253–262.

[107] M. Ahmadi, D. Ulyanov, S. Semenov, M. Trofimov, and G. Giacinto, "Novel feature extraction, selection and fusion for effective malware family classification," in *Proceedings of the sixth ACM conference on data and application security and privacy*. ACM, 2016, pp. 183–194.

[108] Z. Xiaolin, Z. Yiman, L. Xuhui, and C. Quanbao, "Research on malicious code homology analysis method based on texture fingerprint clustering," in *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*. IEEE, 2018, pp. 1914–1921.

[109] "Big 2015: Bigdata innovators gathering," 2015. [Online]. Available: http://www.www2015.it/big-2015/

[110] P. Silva, S. Akhavan-Masouleh, and L. Li, "Improving malware detection accuracy by extracting icon information," in *2018 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*. IEEE, 2018, pp. 408–411.

[111] B. Andow, A. Nadkarni, B. Bassett, W. Enck, and T. Xie, "A study of grayware on google play," in *2016 IEEE Security and Privacy Workshops (SPW)*. IEEE, 2016, pp. 224–233.

[112] J. Y. Kim and S. B. Cho, "Detecting intrusive malware with a hybrid generative deep learning model," in *International Conference on Intelligent Data Engineering and Automated Learning*. Springer, 2018, pp. 499–507.

[113] J. Y. Kim, S. J. Bu, and S. B. Cho, "Zero-day malware detection using transferred generative adversarial networks based on deep autoencoders," *Information Sciences*, vol. 460, pp. 83–102, 2018.

[114] J. Yan, Y. Qi, and Q. Rao, "Detecting malware with an ensemble method based on deep neural network," *Security and Communication Networks*, vol. 2018, 2018.

[115] Q. Le, O. Boydell, B. Mac Namee, and M. Scanlon, "Deep learning at the shallow end: Malware classification for non-domain experts," *Digital Investigation*, vol. 26, pp. S118–S126, 2018.

[116] M. Kalash, M. Rochan, N. Mohammed, N. D. Bruce, Y. Wang, and F. Iqbal, "Malware classification with deep convolutional neural networks," in *2018 9th IFIP International Conference on New Technologies, Mobility and Security (NTMS)*. IEEE, 2018, pp. 1–5.

[117] D. Cao, X. Zhang, Z. Ning, J. Zhao, F. Xue, and Y. Yang, "An efficient malicious code detection system based on convolutional neural networks," in *Proceedings of the 2018 2nd International Conference on Computer Science and Artificial Intelligence*. ACM, 2018, pp. 86–89.

[118] S. M. Kywe, Y. Li, R. H. Deng, and J. Hong, "Detecting camouflaged applications on mobile application markets," in *Information Security and Cryptology - ICISC 2014*, J. Lee and J. Kim, Eds. Cham: Springer International Publishing, 2015, pp. 241–254.

[119] B. Xiaofang, C. Li, H. Weihua, and W. Qu, "Malware variant detection using similarity search over content fingerprint," in *The 26th Chinese Control and Decision Conference*. IEEE, 2014, pp. 5334–5339.

[120] L. Malisa, K. Kostiainen, M. Och, and S. Capkun, "Mobile application impersonation detection using dynamic user interface extraction," in *European Symposium on Research in Computer Security*, 2016.

[121] "Playdrone archive." [Online]. Available: https://archive.org/details/android apps

[122] S. Jiao, Y. Cheng, L. Ying, P. Su, and D. Feng, "A rapid and scalable method for android application repackaging detection," in *Information Security Practice and Experience*, J. Lopez and Y. Wu, Eds. Cham: Springer International Publishing, 2015, pp. 349–364.

[123] A. Long, J. Saxe, and R. Gove, "Detecting malware samples with similar image sets," in *Eleventh Workshop on Visualization for Cyber Security*, ser. VizSec '14, 2014, pp. 88–95.

[124] O. Liljeblad, "icoutils - introduction," 2013. [Online]. Available: http://www.nongnu.org/icoutils/

[125] J. Zhang, Z. Qin, H. Yin, L. Ou, and Y. Hu, "IRMD: malware variant detection using opcode image recognition," in *2016 IEEE 22nd International Conference on Parallel and Distributed Systems (ICPADS)*. IEEE, 2016, pp. 1175–1180.

[126] J. Y. Kim, S. J. Bu, and S. B. Cho, "Malware detection using deep transferred generative adversarial networks," in *International Conference on Neural Information Processing*. Springer, 2017, pp. 556–564.

[127] C. Yinka Banjo and O. A. Ugot, "A review of generative adversarial networks and its application in cybersecurity," *Artificial Intelligence Review*, pp. 1–16, 2019.

[128] S. Ni, Q. Qian, and R. Zhang, "Malware identification using visualization images and deep learning," *Computers & Security*, vol. 77, pp. 871–885, 2018.

[129] Z. Cui, F. Xue, X. Cai, Y. Cao, G. g. Wang, and J. Chen, "Detection of malicious code variants based on deep learning," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 7, pp. 3187–3196, 2018.

[130] S. Yue, "Imbalanced malware images classification: a CNN based approach," *CoRR*, vol. abs/1708.08042, 2017.

[131] J. Su, V. D. Vasconcellos, S. Prasad, S. Daniele, Y. Feng, and K. Sakurai, "Lightweight classification of iot malware based on image recognition," in *2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC)*, vol. 2, 2018, pp. 664–669.

[132] Y. M. P. Pa, S. Suzuki, K. Yoshioka, T. Matsumoto, T. Kasama, and C. Rossow, "IoTPOT: A novel honeypot for revealing current IoT threats," *Journal of Information Processing*, vol. 24, no. 3, pp. 522–533, 2016.

[133] S. Akarsh, P. Poornachandran, V. K. Menon, and K. Soman, "A detailed investigation and analysis of deep learning architectures and visualization techniques for malware family identification," in *Cybersecurity and Secure Information Systems*. Springer, 2019, pp. 241–286.

[134] S. Venkatraman, M. Alazab, and R. Vinayakumar, "A hybrid deep learning image-based analysis for effective malware detection," *Journal of Information Security and Applications*, vol. 47, pp. 377–389, 2019.

[135] S. D. SL and C. Jaidhar, "Windows malware detector using convolutional neural network based on visualization images," *IEEE Transactions on Emerging Topics in Computing*, 2019.

[136] K. Rieck, P. Trinius, C. Willems, and T. Holz, "Automatic analysis of malware behavior using machine learning," *Journal of Computer Security*, vol. 19, no. 4, pp. 639–668, 2011.

[137] D. Xue, J. Li, T. Lv, W. Wu, and J. Wang, "Malware classification using probability scoring and machine learning," *IEEE Access*, vol. 7, pp. 91 641–91 656, 2019.

[138] J. Rajasegaran, N. Karunanayake, A. Gunathillake, S. Seneviratne, and G. Jourjon, "A multi-modal neural embeddings approach for detecting mobile counterfeit apps," in *The World Wide Web Conference*. ACM, 2019, pp. 3165–3171.

[139] "Ember," 2019. [Online]. Available: https://github.com/endgameinc/ember

[140] INQUIRER, "Security boffins uncover linux variant of winnti malware," 2019. [Online]. Available: https://www.theinquirer.net/inquirer/news/3076093/winnti-linux-malware-uncovered

[141] A. Kuzmanovic and E. W. Knightly, "Low-rate tcp-targeted denial of service attacks: the shrew vs. the mice and elephants," in *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*. ACM, 2003, pp. 75–86.

[142] R. Fontugne, T. Hirotsu, and K. Fukuda, "An image processing approach to traffic anomaly detection," in *4th Asian Conference on Internet Engineering*, ser. AINTEC '08, 2008, pp. 17–26.

[143] P. Barford, J. Kline, D. Plonka, and A. Ron, "A signal analysis of network traffic anomalies," in *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment*. ACM, 2002, pp. 71–82.

[144] R. O. Duda and P. E. Hart, "Use of the hough transformation to detect lines and curves in pictures," Sri International Menlo Park Ca Artificial Intelligence Center, Tech. Rep., 1971.

[145] J. Canny, "A computational approach to edge detection," *IEEE Transactions on pattern analysis and machine intelligence*, no. 6, pp. 679–698, 1986.

[146] T. Wiegand, G. J. Sullivan, G. Bjontegaard, and A. Luthra, "Overview of the H. 264/AVC video coding standard," *IEEE Transactions on circuits and systems for video technology*, vol. 13, no. 7, pp. 560–576, 2003.

[147] S. S. Kim and A. L. N. Reddy, "A study of analyzing network traffic as images in real-time," in *IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies.*, vol. 3, 2005, pp. 2056–2067.

[148] ——, "Modeling network traffic as images," in *IEEE International Conference on Communications*, vol. 1, 2005, pp. 168–172.

[149] A. Hussain, J. Heidemann, J. Heidemann, and C. Papadopoulos, "A framework for classifying denial of service attacks," in *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*. ACM, 2003, pp. 99–110.

[150] S. S. Kim and A. L. N. Reddy, "Image-based anomaly detection technique: Algorithm, implementation and effectiveness," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 10, pp. 1942–1954, 2006.

[151] C. Sony and K. Cho, "Traffic data repository at the WIDE project," in *Proceedings of USENIX 2000 Annual Technical Conference: FREENIX Track*, 2000, pp. 263–270.

[152] G. Dewaele, K. Fukuda, P. Borgnat, P. Abry, and K. Cho, "Extracting hidden anomalies using sketch and non gaussian multiresolution statistical detection procedures," in *Proceedings of the 2007 workshop on Large scale attack defense*. ACM, 2007, pp. 145–152.

[153] C. Y. Jeong, B. H. Chang, and J. C. Na, "A hierarchical approach to traffic anomaly detection using image processing technique," in *6th International Conference on Networked Computing and Advanced Information Management*, 2010, pp. 592–594.

[154] L. He, X. Ren, Q. Gao, X. Zhao, B. Yao, and Y. Chao, "The connected-component labeling problem: A review of state-of-the-art algorithms," *Pattern Recognition*, vol. 70, pp. 25–43, 2017.

[155] T. H. Kim, D. S. Kim, and H. Y. Jung, "Defending against DDoS attacks under ip spoofing using image processing approach," *IEICE Transactions on Communications*, vol. 99, no. 7, pp. 1511–1522, 2016.

[156] "The IPv4 routed /24 topology dataset." [Online]. Available: https://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml

[157] "The IPv6 topology dataset." [Online]. Available: https://www.caida.org/data/active/ipv6_allpref_topology_dataset.xml

[158] Z. Tan, A. Jamdagni, X. He, P. Nanda, R. P. Liu, and J. Hu, "Detection of denial-of-service attacks based on computer vision techniques," *IEEE Transactions on Computers*, vol. 64, no. 9, pp. 2519–2533, 2015.

[159] Z. Tan, A. Jamdagni, X. He, P. Nanda, and R. P. Liu, "A system for denial-of-service attack detection based on multivariate correlation analysis," *IEEE transactions on parallel and distributed systems*, vol. 25, no. 2, pp. 447–456, 2013.

[160] "KDD cup 1999 data," 1999. [Online]. Available: http://kdd.ics.uci.edu/databases/kddcup99/kddcup99.html

[161] A. Shiravi, H. Shiravi, M. Tavallaee, and A. A. Ghorbani, "Toward developing a systematic approach to generate benchmark datasets for intrusion detection," *computers & security*, vol. 31, no. 3, pp. 357–374, 2012.

[162] S. Jin, D. S. Yeung, and X. Wang, "Network intrusion detection in covariance feature space," *Pattern Recognition*, vol. 40, no. 8, pp. 2185–2197, 2007.

[163] C. F. Tsai and C. Y. Lin, "A triangle area based nearest neighbors approach to intrusion detection," *Pattern recognition*, vol. 43, no. 1, pp. 222–229, 2010.

[164] G. Kumar and K. Kumar, "Design of an evolutionary approach for intrusion detection," *The Scientific World Journal*, vol. 2013, 2013.

[165] M. Zou, C. Wang, F. Li, and W. Song, "Network phenotyping for network traffic classification and anomaly detection," *CoRR*, vol. abs/1803.01528, 2018.

[166] "Cyber-physical systems (CPS)," 2010. [Online]. Available: https://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf11516

[167] A. Materka and M. Strzelecki, "Texture analysis methods–a review," *Technical university of lodz, institute of electronics, COST B11 report, Brussels*, pp. 9–11, 1998.

[168] A. Baraldi and F. Parmiggiani, "An investigation of the textural characteristics associated with gray level cooccurrence matrix statistical parameters," *IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING*, vol. 33, no. 2, 1995.

[169] J. Ahrenholz, "Comparison of core network emulation platforms," in *Military Communications Conference*. IEEE, 2010, pp. 166–171.

[170] Z. Wang, "The applications of deep learning on traffic identification," *BlackHat USA*, vol. 24, 2015.

[171] Z. Li, Z. Qin, K. Huang, X. Yang, and S. Ye, "Intrusion detection using convolutional neural networks for representation learning," in *Neural Information Processing*, D. Liu, S. Xie, Y. Li, D. Zhao, and E.-S. M. El-Alfy, Eds., 2017, pp. 858–866.

[172] P. Aggarwal and S. K. Sharma, "Analysis of KDD dataset attributes-class wise for intrusion detection," *Procedia Computer Science*, vol. 57, pp. 842–851, 2015.

[173] Y. Wang, J. An, and W. Huang, "Using cnn-based representation learning method for malicious traffic identification," in *2018 IEEE/ACIS 17th International Conference on Computer and Information Science (ICIS)*. IEEE, 2018, pp. 400–404.

[174] "The UNSW-NB15 dataset description," 2016. [Online]. Available: https://www.unsw.adfa.edu.au/australian-centre-for-cyber-security/cybersecurity/ADFA-NB15-Datasets/

[175] Z. Zhou, L. Yao, J. Li, B. Hu, C. Wang, and Z. Wang, "Classification of botnet families based on features self-learning under network traffic censorship," in *Third International Conference on Security of Smart Cities, Industrial Control System and Communications*, 2018, pp. 1–7.

[176] B. Anderson and D. McGrew, "Identifying encrypted malware traffic with contextual flow data," in *Proceedings of the 2016 ACM workshop on artificial intelligence and security*. ACM, 2016, pp. 35–46.

[177] "Botnet dataset." [Online]. Available: https://www.unb.ca/cic/datasets/botnet.html

[178] M. R. Moore and J. M. Vann, "Anomaly detection of cyber physical network data using 2D images," in *2019 IEEE International Conference on Consumer Electronics (ICCE)*. IEEE, 2019, pp. 1–5.

[179] J. McHugh, "The 1998 lincoln laboratory ids evaluation," in *International Workshop on Recent Advances in Intrusion Detection*. Springer, 2000, pp. 145–161.

[180] C. Brown, A. Cowperthwaite, A. Hijazi, and A. Somayaji, "Analysis of the 1999 darpa/lincoln laboratory ids evaluation data with netadhict," in *2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications*. IEEE, 2009, pp. 1–7.

[181] D. Grant, "Detecting phishing with computer vision: Part 1, blazar," 2018. [Online]. Available: https://www.endgame.com/blog/technical-blog/detecting-phishing-computer-vision-part-1-blazar

[182] "Endgame," 2018. [Online]. Available: https://www.endgame.com

[183] J. Woodbridge, H. S. Anderson, A. Ahuja, and D. Grant, "Detecting homoglyph attacks with a siamese neural network," in *2018 IEEE Security and Privacy Workshops (SPW)*. IEEE, 2018, pp. 22–28.

[184] D. G. Bill Finlayson, "Detecting phishing with computer vision: Part 2, speedgrapher," 2018. [Online]. Available: https://www.endgame.com/blog/technical-blog/detecting-phishing-computer-vision-part-2-speedgrapher

[185] "Microsoft.Office.Interop.Word.Namespace." [Online]. Available: https://docs.microsoft.com/en-us/dotnet/api/microsoft.office.interop.word?view=word-pia

[186] J. Richards, "What is lookout phishing AI?" 2019. [Online]. Available: https://blog.lookout.com/lookout-phishing-ai

[187] "Phishing attack targeting united nations and humanitarian organizations discovered by lookout phishing AI." [Online]. Available: https://blog.lookout.com/lookout-phishing-ai-discovers-phishing-attack-targeting-humanitarian-organizations

[188] "Next-generation anti-phishing platform powered by AI & computer vision," 2019. [Online]. Available: https://www.phish.ai

[189] MayaChitra, "Mayachitra is developing advanced cybersecurity tools using deep learning and pattern recognition methods," 2019. [Online]. Available: https://mayachitra.com/cyber-security/

[190] IronScales, "Ironscales raises $15 million to defeat phishing attacks with AI," 2019. [Online]. Available: https://venturebeat.com/2019/06/17/ironscales-raises-15-million-to-defeat-phishing-attacks-with-ai/

[191] "Total anti-phishing protection," 2019. [Online]. Available: https://cyberfish.io

[192] 406Venture, "Greathorn prevents Microsoft Office 365 and Google G Suite credential theft phishing using advanced computer vision," 2019. [Online]. Available: https://www.406ventures.com/news/2418-greathorn-prevents-microsoft-office-365-and-google-g-suite-credential-theft-phishing-using-advanced-computer-vision

[193] O. Falkowitz, "Seeing through normal with computer vision," 2017. [Online]. Available: https://www.area1security.com/blog/seeing-normal-computer-vision/

[194] M. Price, "Artificial intelligence is key to digital risk protection," 2019. [Online]. Available: https://www.zerofox.com/blog/artificial-intelligence-computer-vision/

[195] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv preprint arXiv:1412.6572*, 2014.

[196] S. M. Moosavi Dezfooli, A. Fawzi, and P. Frossard, "Deepfool: a simple and accurate method to fool deep neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2574–2582.

[197] A. Kurakin, I. Goodfellow, and S. Bengio, "Adversarial machine learning at scale," *arXiv preprint arXiv:1611.01236*, 2016.

[198] X. Chen, C. Liu, B. Li, K. Lu, and D. Song, "Targeted backdoor attacks on deep learning systems using data poisoning," *arXiv preprint arXiv:1712.05526*, 2017.

[199] L. Muñoz González, B. Biggio, A. Demontis, A. Paudice, V. Wongrassamee, E. C. Lupu, and F. Roli, "Towards poisoning of deep learning algorithms with back-gradient optimization," in *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, 2017.

[200] A. Shafahi, W. R. Huang, M. Najibi, O. Suciu, C. Studer, T. Dumitras, and T. Goldstein, "Poison frogs! targeted clean-label poisoning attacks on neural networks," in *Advances in Neural Information Processing Systems*, 2018, pp. 6103–6113.

[201] A. Al Dujaili, A. Huang, E. Hemberg, and U.-M. OReilly, "Adversarial deep learning for robust detection of binary encoded malware," in *2018 IEEE Security and Privacy Workshops (SPW)*. IEEE, 2018, pp. 76–82.

[202] K. Grosse, N. Papernot, P. Manoharan, M. Backes, and P. McDaniel, "Adversarial perturbations against deep neural networks for malware classification," *arXiv preprint arXiv:1606.04435*, 2016.

[203] W. Hu and Y. Tan, "Generating adversarial malware examples for black-box attacks based on GAN," *arXiv preprint arXiv:1702.05983*, 2017.

[204] I. Rosenberg, A. Shabtai, Y. Elovici, and L. Rokach, "Query-efficient GAN based black-box attack against sequence based machine and deep learning classifiers," *arXiv preprint arXiv:1804.08778*, 2018.

[205] M. Rigaki and S. Garcia, "Bringing a GAN to a knife-fight: Adapting malware communication to avoid detection," in *2018 IEEE Security and Privacy Workshops (SPW)*. IEEE, 2018, pp. 70–75.

[206] B. Biggio, I. Corona, G. Fumera, G. Giacinto, and F. Roli, "Bagging classifiers for fighting poisoning attacks in adversarial classification tasks," in *International workshop on multiple classifier systems*. Springer, 2011, pp. 350–359.

[207] N. Papernot, N. Carlini, I. Goodfellow, R. Feinman, F. Faghri, A. Matyasko, K. Hambardzumyan, Y.-L. Juang, A. Kurakin, and R. Sheatsley, "cleverhans v2. 0.0: an adversarial machine learning library," *arXiv preprint arXiv:1610.00768*, 2016.

[208] H. Xu, Y. Ma, H. Liu, D. Deb, H. Liu, J. Tang, and A. Jain, "Adversarial attacks and defenses in images, graphs and text: A review," *arXiv preprint arXiv:1909.08072*, 2019.

[209] N. Akhtar and A. Mian, "Threat of adversarial attacks on deep learning in computer vision: A survey," *IEEE Access*, vol. 6, pp. 14 410–14 430, 2018.

[210] A. Chakraborty, M. Alam, V. Dey, A. Chattopadhyay, and D. Mukhopadhyay, "Adversarial attacks and defences: A survey," *arXiv preprint arXiv:1810.00069*, 2018.

[211] O. Vinyals, C. Blundell, T. Lillicrap, and D. Wierstra, "Matching networks for one shot learning," in *Advances in neural information processing systems*, 2016, pp. 3630–3638.

[212] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Advances in Neural Information Processing Systems*, 2017, pp. 4077–4087.

[213] S. Qiao, C. Liu, W. Shen, and A. L. Yuille, "Few-shot image recognition by predicting parameters from activations," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7229–7238.

[214] S. Gidaris and N. Komodakis, "Dynamic few-shot visual learning without forgetting," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4367–4375.

[215] M. Woodward and C. Finn, "Active one-shot learning," *arXiv preprint arXiv:1702.06559*, 2017.

[216] R. Socher, M. Ganjoo, C. D. Manning, and A. Ng, "Zero-shot learning through cross-modal transfer," in *Advances in neural information processing systems*, 2013, pp. 935–943.

[217] L. Fei Fei, R. Fergus, and P. Perona, "One-shot learning of object categories," *IEEE transactions on pattern analysis and machine intelligence*, vol. 28, no. 4, pp. 594–611, 2006.

[218] Y. Xian, B. Schiele, and Z. Akata, "Zero-shot learning-the good, the bad and the ugly," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 4582–4591.

[219] O. J. Hénaff, A. Razavi, C. Doersch, S. Eslami, and A. V. D. Oord, "Data-efficient image recognition with contrastive predictive coding," *arXiv preprint arXiv:1905.09272*, 2019.

[220] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," *arXiv preprint arXiv:1911.05722*, 2019.

[221] S. C. Hsiao, D. Y. Kao, Z. Y. Liu, and R. Tso, "Malware image classification using one-shot learning with siamese networks," *Procedia Computer Science*, vol. 159, pp. 1863–1871, 2019.

[222] T. K. Tran, H. Sato, and M. Kubo, "Image-based unknown malware classification with few-shot learning models," in *2019 Seventh International Symposium on Computing and Networking Workshops (CANDARW)*. IEEE, 2019, pp. 401–407.

[223] A. Atapour Abarghouei, S. Bonner, and A. S. McGough, "A kings ransom for encryption: Ransomware classification using augmented one-shot learning and bayesian approximation," *arXiv preprint arXiv:1908.06750*, 2019.

[224] W. J. Scheirer, A. De Rezende Rocha, A. Sapkota, and T. E. Boult, "Toward open set recognition," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 7, pp. 1757–1772, 2012.

[225] C. Geng, S. j. Huang, and S. Chen, "Recent advances in open set recognition: A survey," *arXiv preprint arXiv:1811.08581*, 2018.

[226] T. Boult, S. Cruz, A. Dhamija, M. Gunther, J. Henrydoss, and W. Scheirer, "Learning and the unknown: Surveying steps toward open world recognition," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 9801–9807.

[227] A. R. Dhamija, M. Günther, and T. Boult, "Reducing network agnostophobia," in *Advances in Neural Information Processing Systems*, 2018, pp. 9157–9168.

[228] A. Raj Dhamija, M. Gunther, and T. E. Boult, "Improving deep network robustness to unknown inputs with objectosphere," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 87–90.

[229] P. Oza and V. M. Patel, "C2AE: Class conditioned auto-encoder for open-set recognition," *arXiv preprint arXiv:1904.01198*, 2019.

[230] Z. Liu, Z. Miao, X. Zhan, J. Wang, B. Gong, and S. X. Yu, "Large-scale long-tailed recognition in an open world," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

[231] R. Yoshihashi, W. Shao, R. Kawakami, S. You, M. Iida, and T. Naemura, "Classification-reconstruction learning for open-set recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4016–4025.

[232] M. Hein, M. Andriushchenko, and J. Bitterwolf, "Why ReLU networks yield high-confidence predictions far away from the training data and how to mitigate the problem," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 41–50.

[233] S. Cruz, C. Coleman, E. M. Rudd, and T. E. Boult, "Open set intrusion recognition for fine-grained attack categorization," in *2017 IEEE International Symposium on Technologies for Homeland Security*.

[234] J. Henrydoss, S. Cruz, E. M. Rudd, and T. E. Boult, "Incremental open set intrusion recognition using extreme value machine," in *2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*. IEEE, 2017, pp. 1089–1093.

[235] E. M. Rudd, A. Rozsa, M. Günther, and T. E. Boult, "A survey of stealth malware attacks, mitigation measures, and steps toward autonomous open world solutions," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 2, pp. 1145–1172, 2016.

[236] L. Bottou and Y. L. Cun, "Large scale online learning," in *Advances in neural information processing systems*, 2004, pp. 217–224.

[237] D. Sahoo, Q. Pham, J. Lu, and S. C. Hoi, "Online deep learning: Learning deep neural networks on the fly," *arXiv preprint arXiv:1711.03705*, 2017.

[238] M. McCloskey and N. J. Cohen, "Catastrophic interference in connectionist networks: The sequential learning problem," in *Psychology of learning and motivation*. Elsevier, 1989, vol. 24, pp. 109–165.

[239] R. Ratcliff, "Connectionist models of recognition memory: constraints imposed by learning and forgetting functions." *Psychological review*, vol. 97, no. 2, p. 285, 1990.

[240] I. J. Goodfellow, M. Mirza, D. Xiao, A. Courville, and Y. Bengio, "An empirical investigation of catastrophic forgetting in gradient-based neural networks," *arXiv preprint arXiv:1312.6211*, 2013.

[241] P. Tune and D. Veitch, "Sampling vs sketching: An information theoretic comparison," in *2011 Proceedings IEEE INFOCOM*. IEEE, 2011, pp. 2105–2113.

[242] Z. Liu, A. Manousis, G. Vorsanger, V. Sekar, and V. Braverman, "One sketch to rule them all: Rethinking network flow monitoring with univmon," in *Proceedings of the 2016 ACM SIGCOMM Conference*. ACM, 2016, pp. 101–114.

[243] T. Yang, J. Jiang, P. Liu, Q. Huang, J. Gong, Y. Zhou, R. Miao, X. Li, and S. Uhlig, "Elastic sketch: Adaptive and fast network-wide measurements," in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, 2018, pp. 561–575.

[244] P. Tune and D. Veitch, "OFSS: Skampling for the flow size distribution," in *Proceedings of the 2014 Conference on Internet Measurement Conference*. Citeseer, 2014, pp. 235–240.

Jiawei Zhao Jiawei Zhao received her Bachelors Degree in Electrical Engineering (First Class Hons.) jointly from The University of Sydney, Australia, and The Harbin Institute of Technology, China, in 2017. She is currently pursuing a PhD degree at The School of Computer Science in the University of Sydney, Australia. Her research interests include cybersecurity, deep learning, fraud detection, and security of machine learning.

Rahat Masood Rahat Masood is currently a research fellow at The University of Sydney. She received her PhD in 2019 from the University of New South Wales, Australia, in collaboration with information security and privacy group at Data61-CSIRO. Her research area focuses on security and privacy of mobile and web platforms, authentication and authorization, biometric security, and critical infrastructure protection. She was also a visiting scholar at the Sandia National Laboratories (SNL), New Mexico, and Cyber Security Policy and Research Institute (CSPRI) at The George Washington University.

Suranga Seneviratne Suranga Seneviratne is a Lecturer in Security at the School of Computer Science, The University of Sydney. He received his PhD from University of New South Wales, Australia in 2015. His current research interests include privacy and security in mobile systems, AI applications in security, and behaviour biometrics. Before moving into research, he worked nearly six years in the telecommunications industry in core network planning and operations. He received his bachelor degree from University of Moratuwa, Sri Lanka in 2005.