

Detecting Phishing Sites Using ChatGPT

Takashi Koide
takashi.koide@global.ntt
NTT Security Japan (KK)
Tokyo, Japan

Naoki Fukushi
naoki.fukushi@global.ntt
NTT Security Japan (KK)
Tokyo, Japan

Hiroki Nakano
hi.nakano.sec@gmail.com
NTT Security Japan (KK)
Tokyo, Japan

Daiki Chiba
daiki.chiba@ieee.org
NTT Security Japan (KK)
Tokyo, Japan

ABSTRACT

The rise of large language models (LLMs) has had a significant impact on various domains, including natural language processing and artificial intelligence. While LLMs such as ChatGPT have been extensively researched for tasks such as code generation and text synthesis, their application in detecting malicious web content, particularly phishing sites, has been largely unexplored. To combat the rising tide of automated cyber attacks facilitated by LLMs, it is imperative to automate the detection of malicious web content, which requires approaches that leverage the power of LLMs to analyze and classify phishing sites.

In this paper, we propose a novel method that utilizes ChatGPT to detect phishing sites. Our approach involves leveraging a web crawler to gather information from websites and generate prompts based on this collected data. This approach enables us to detect various phishing sites without the need for fine-tuning machine learning models and identify social engineering techniques from the context of entire websites and URLs. To evaluate the performance of our proposed method, we conducted experiments using a dataset. The experimental results using GPT-4 demonstrated promising performance, with a precision of 98.3% and a recall of 98.4%. Comparative analysis between GPT-3.5 and GPT-4 revealed an enhancement in the latter's capability to reduce false negatives. These findings not only highlight the potential of LLMs in efficiently identifying phishing sites but also have significant implications for enhancing cybersecurity measures and protecting users from the dangers of online fraudulent activities.

KEYWORDS

Phishing Sites, Social Engineering, and ChatGPT.

1 INTRODUCTION

The emergence of large language models (LLMs) [15, 30, 31, 35] has revolutionized various domains, including natural language processing and artificial intelligence. These models have gained significant popularity due to their remarkable performance across a wide range of tasks. Among them, ChatGPT [5] has consistently achieved state-of-the-art results, demonstrating its ability to generate coherent and contextually relevant responses in conversational settings. While extensive research has focused on exploring the capabilities of LLMs for tasks such as code generation and text synthesis, little attention has been paid to leveraging these models for the analysis and detection of malicious web content, particularly phishing sites. Phishing sites, which employ social engineering techniques [28, 41, 45], pose a severe threat to Internet users by masquerading as legitimate platforms and tricking them into revealing sensitive information or causing financial harm. To effectively combat the growing threat of automated cyber-attacks facilitated by LLMs, it is important to automate the detection of malicious

web content. This requires the exploration of methodologies that leverage the power of LLMs to analyze and classify phishing sites.

In this research, we propose a novel method that utilizes ChatGPT to detect phishing sites. Our approach involves leveraging a web crawler to gather information from websites and generate prompts [23, 43] based on this collected data. These prompts are then presented to ChatGPT, which determines whether a given website is a phishing site or not. The integration of web crawling and ChatGPT's contextual understanding enables us to make informed decisions concerning the legitimacy or suspiciousness of websites. By employing ChatGPT, we can (1) detect various phishing sites without the need for fine-tuning machine learning models and (2) identify social engineering (SE) techniques from the context of entire websites and URLs. To the best of our knowledge, this is the first study to analyze the ability of LLMs to identify phishing sites. To evaluate the performance of our proposed method, we conducted experiments using a carefully curated dataset for phishing site detection. The experimental results using GPT-4 [30] showed promising performance, with a precision of 98.3% and a recall of 98.4%. Moreover, a comparative analysis between GPT-3.5 and GPT-4 revealed a significant improvement in the latter's capabilities, particularly in terms of reducing false negatives. GPT-4 outperformed GPT-3.5 in its ability to determine the suspiciousness of domain names, identify SE techniques from the website content, and provide comprehensive phishing detection by considering multiple factors. The results of this study emphasize the potential of LLMs in efficiently detecting phishing sites, particularly in uncovering SE techniques aimed at psychologically manipulating users. These findings have significant implications for enhancing automated cybersecurity measures and mitigating the risks of online fraudulent activities faced by users.

In summary, we make the following contributions:

- We propose a novel method for detecting phishing sites using ChatGPT. By utilizing web crawling techniques to gather information from websites and generate prompts, our approach enables ChatGPT to classify websites as phishing or not.
- We present an experimental evaluation of our proposed method using a dataset designed for phishing site detection. Our results show promising performance with a precision of 98.3% and a recall of 98.4%, demonstrating the effectiveness of LLMs in identifying phishing sites.
- We performed a comparative analysis between GPT-3.5 and GPT-4 and found notable improvements in GPT-4's ability to identify phishing sites, particularly in minimizing false negatives. It excelled at identifying suspicious domains, detecting SE techniques in web content, and providing comprehensive phishing detection by considering multiple factors.

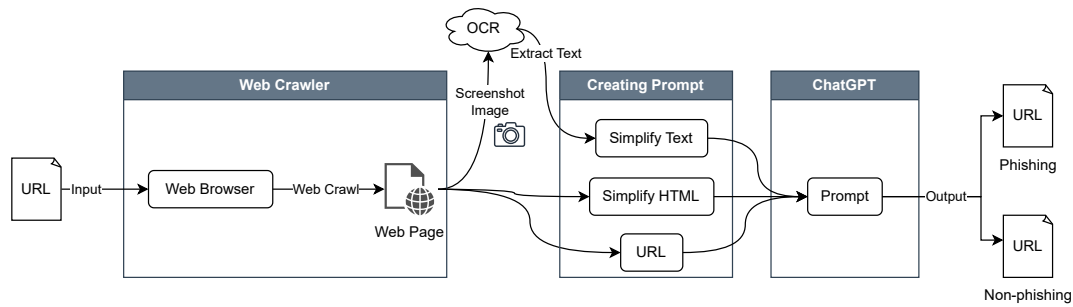


Figure 1: Overview of Proposed Method.

2 BACKGROUND

Phishing sites are fraudulent websites that aim to steal personal information, money, or cause malware infections by psychologically luring users. These sites use email, short message service (SMS), and web advertisements to attract users and redirect them to phishing sites by having them click on malicious links [37, 38]. There are two key components that constitute phishing sites, which either have one or both of these elements: (1) Imitation of domain names and the use of official logos and branding to deceive users into believing they are interacting with legitimate services, and (2) Employing SE techniques to deceive users and manipulate their actions. By impersonating legitimate services, phishing sites deceive users and trust them [24, 25]. These websites often imitate legitimate platforms such as online banking [13], e-commerce sites [14], and social [39] media. In addition, they can create a sense of urgency or fear by displaying fake malware infection warnings or account problems, or generate interest by displaying fake rewards [22]. Through these SE techniques, users are misled into providing sensitive information such as login credentials and credit card numbers. Furthermore, various user actions may be induced, such as sending cryptocurrency, calling fake technical support centers [26], or downloading apps and executable files [28].

Previous studies have been conducted to understand the characteristics, techniques, and fundamental mechanisms used by attackers in phishing sites. Researchers analyze the design, structure, and content of these sites, identify common patterns, and develop methods for detecting phishing sites. These studies can contribute to improving security practices, educating users about potential threats, and devising effective strategies to mitigate online fraud and associated risks. For example, studies exist on identifying websites that abuse legitimate branding based on their appearance [11, 24, 25], as well as on discovering brand information in domain names and URLs [12, 20, 29]. Some studies have also been conducted on identifying phishing sites based on information contained in certificates and domain names [16, 19], and on capturing the context of social engineering in technical support scams, fake infection warnings, and fake rewards [18, 21, 22].

Although various methods have been employed to detect phishing sites, there are two main problems. One is the need for learning targeted brands and modifying algorithms depending on the phishing sites. For example, it is necessary to learn in advance the logo images that phishing sites abuse, or to create rules to detect them

according to domain squatting techniques. The other is the inability to analyze the context of psychological manipulation induced by social engineering in detail. While some studies have been conducted on keyword matching and deep learning-based analysis [18, 44], no attempt has been made to automatically analyze and understand the various contexts of psychological manipulation by examining the entire content of a website.

3 METHOD

We propose a method for detecting phishing sites by using ChatGPT, to analyze the contextual information of websites and their respective URLs. Our proposed method takes advantage of ChatGPT’s highly accurate contextual understanding to precisely identify textual representations involving SE techniques, as well as inconsistencies between brands deceived by the websites and domain names. By employing ChatGPT, which is trained on extensive text data, we can detect various phishing sites that disguise themselves as legitimate services without learning from collected phishing site data. In this manner, the proposed method addresses two problems present in previous studies. An overview of the proposed method is illustrated in Figure 1. Our method uses a web crawler to access the input URL and obtain information from the visited website, such as screenshot images, HTML, and URLs. This information is used to create a prompt for input into ChatGPT, which then determines whether the site is a phishing site or a non-phishing site.

3.1 Web Crawling

We implement a web crawler that automates Google Chrome with Chrome DevTools Protocol [4] to visit websites and collect information. Given an input URL, the web crawler retrieves the URL of the reached web page (the final destination after any redirects), the HTML after JavaScript execution, and captures a screenshot image. The reason for obtaining the HTML after JavaScript execution (Browser-rendered HTML) is that some phishing sites employ obfuscated JavaScript to generate DOM elements as a means to evade analysis. As a result, it becomes challenging to analyze any traces of phishing solely based on the HTML before JavaScript execution. The web crawler is configured to emulate two different environments: Windows + Chrome, and iPhone + Safari. These configurations included specifying the *UserAgent* and browser size for each environment.

3.2 Prompt

We describe the process of generating a prompt for detecting phishing sites. The template for the prompt is shown in Prompt 1. The purpose of this prompt is to provide website information to ChatGPT and determine whether it is a phishing site or not. We create the template based on the Chain of Thought (CoT) prompting technique [23, 42]. CoT is a prompting technique that encourages LLMs to explain their reasoning. CoT has been shown to be effective in improving performance on a variety of reasoning tasks, such as arithmetic and symbolic reasoning. The task of phishing site detection is divided into four subtasks, which facilitate the execution of specific reasoning processes. These subtasks are as follows:

- (1) Analyze whether the website contains SE techniques that deceive or attract users. The prompt provides typical SE techniques commonly used by phishing sites, such as cash prizes, fake malware infection warnings, account problems, and postal parcel issues. The LLM is instructed that the presence of these phrases without context in a non-login state indicates a high likelihood of a phishing site.
- (2) Extract the brand name of the website. Phishing sites may be created by copying resources such as HTML and images from legitimate sites, making it difficult to determine their authenticity based on HTML alone. Therefore, the LLM is prompted to confirm whether the URL corresponds to the legitimate site.
- (3) Determine if the site is a phishing site or not and explicitly state the rationale for the decision. Providing detailed descriptions improves the accuracy of responses and makes it easier for humans to analyze them.
- (4) Generate output in JSON format. If the website uses SE techniques, the LLM is expected to return *phishing* as true. If the brand name can be identified and its domain name differs from the legitimate one, *suspicious_domain* is expected to be true. Additionally, the LLM should produce a *phishing_score* ranging from 0 to 10.

The prompt is provided with website information in the form of HTML, URL, and text extracted from screenshot images using optical character recognition (OCR). Some phishing sites evade HTML-based phishing site analysis by displaying brand information (logos and headings) in image files or canvas elements without including specific brand names or text in the HTML. To input text that cannot be extracted from HTML into the prompt, we employ OCR for text extraction.

Some websites may exceed the token limit of ChatGPT when the entered data is large. Tokens in ChatGPT refer to individual units of text that represent characters or words and are used for processing and analysis by LLMs. The GPT-3.5 model has a token limit of 4,096, whereas GPT-4 (8K) has a limit of 8,000 tokens, and GPT-4 (32k) has a limit of 32,000 tokens. In fact, the median token count for HTML in our dataset (see Section 4) is 30,398 (with a median of 74,937 for non-phishing pages), indicating that a large number of web pages exceed this token limit. To compare the performance of the GPT-3.5 model and GPT-4 model, we simplify the HTML and OCR-extracted text to fit within 4,096 tokens in this study. Table 1 presents the maximum token counts for each data category. This prompt template consumes 362 tokens, with a maximum limit of

Prompt Template 1

You are a web programmer and security expert tasked with examining a web page to determine if it is a phishing site or a legitimate site. To complete this task, follow these sub-tasks:

1. Analyze the HTML, URL, and OCR-extracted text for any SE techniques often used in phishing attacks. Point out any suspicious elements found in the HTML, URL, or text.
2. Identify the brand name. If the HTML appears to resemble a legitimate web page, verify if the URL matches the legitimate domain name associated with the brand, if known.
3. State your conclusion on whether the site is a phishing site or a legitimate one, and explain your reasoning. If there is insufficient evidence to make a determination, answer "unknown".
4. Submit your findings as JSON-formatted output with the following keys:
 - *phishing_score*: int (indicates phishing risk on a scale of 0 to 10)
 - *brands*: str (identified brand name or None if not applicable)
 - *phishing*: boolean (whether the site is a phishing site or a legitimate site)
 - *suspicious_domain*: boolean (whether the domain name is suspected to be not legitimate)

Limitations:

- The HTML may be shortened and simplified.
- The OCR-extracted text may not always be accurate.

Examples of social engineering techniques:

- Alerting the user to a problem with their account
- Offering unexpected rewards
- Informing the user of a missing package or additional payment required
- Displaying fake security warnings

URL:

{URL}

HTML:

...

{Browser-rendered HTML}

...

Text extracted using OCR:

...

{OCR-extracted text}

...

Table 1: Token Counts for Prompt Template, HTML, OCR-extracted Text, and URL

Data	Maximum Token Count
Prompt Template	362 (fix)
HTML	3,000
OCR-extracted Text	500
URL	234
Total	4,096

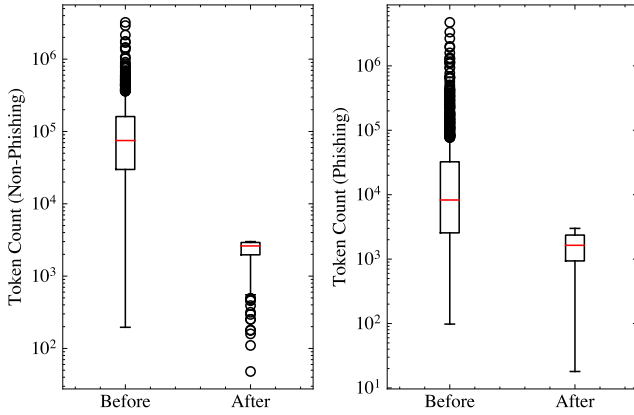


Figure 2: Number of Tokens Before and After Simplified HTML

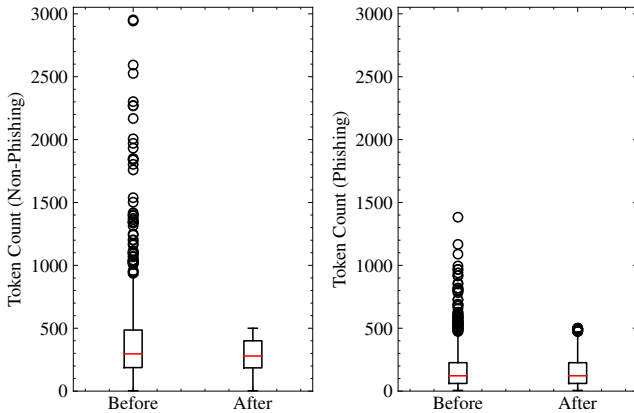


Figure 3: Number of Tokens Before and After Simplified OCR-extracted Text

3,000 tokens for HTML and 500 tokens for OCR-extracted text. Since all generated prompts in this dataset fit within the 4,096-token limit, we did not simplify URLs. The median token count for URLs in the dataset is 9, with the longest URL consisting of 2,209 tokens. In the following sections, we describe the specific processes to simplify the HTML and OCR-extracted text.

Algorithm 1 Simplifying HTML

```

Require: inputHTML: Input HTML
Ensure: processedHTML: Processed HTML
1: function SIMPLIFYHTML(inputHTML)
2:   Remove style, script, and comment elements from
   inputHTML
3:   processedHTML ← Result after removal
4:   if lengthToken(processedHTML) < 3000 then
5:     return processedHTML
6:   end if
7:   Unwrap elements except for important tags
8:   Remove elements without text content
9:   Shorten href in a tags and src in img tags
10:  processedHTML ← Result after removal
11:  if lengthToken(processedHTML) < 3000 then
12:    return processedHTML
13:  end if
14:  while lengthToken(processedHTML) > 3000 do
15:    Remove an HTML element from the midpoint of the
    processedHTML
16:    processedHTML ← Result after removal
17:  end while
18:  return processedHTML
19: end function

```

3.3 Simplifying HTML

The proposed method simplifies the HTML while preserving essential information for phishing detection and brand identification. Elements located at the top of the HTML, such as the title and the meta description element, primarily contain cues for determining the services provided by the website. Additionally, the form element used for inputting login information and text displayed on the page are crucial for analyzing the role of the website. To achieve maximum retention of such information while minimizing its length, we follow the steps outlined in Algorithm 1.

The proposed method removes HTML elements, such as style, script, and comment tags, that have low relevance and a high number of tokens. Subsequently, we unwrap HTML elements other than the important tags listed as head, title, meta, body, h1, h2, h3, h4, h5, h6, p, strong, a, img, hr, table, tbody, tr, th, td, ol, ul, li, ruby, and label. Unimportant elements are eliminated, and their child elements are incorporated into their parent elements. Then, our method removes HTML elements that do not contain text enclosed within tags. We also shorten the src attribute img elements encoded in base64 and the href attribute of a elements containing lengthy URLs. The proposed method repetitively removes intermediate HTML elements until the token count falls below 3,000.

3.4 Simplifying OCR-extracted Text

Some websites may have a large amount of text on the screen, and OCR-extracted text as well as HTML may exceed the token limit. The proposed method also simplifies OCR-extracted text.

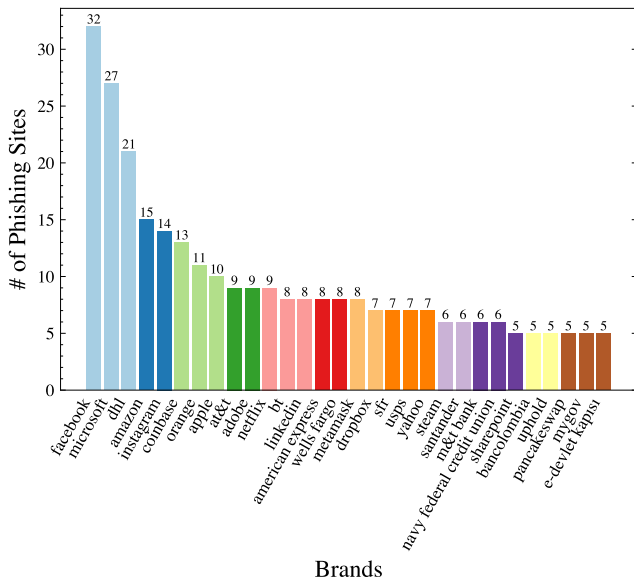


Figure 4: Top 30 Brands Targeted by Phishing Sites.

Specifically, it examines the font size of the identified text from OCR and removes sentences starting from the smallest font size. This process is repeated until the number of tokens in the text falls below 500 tokens.

4 DATASET

We describe how we created the dataset used in our experiment. The dataset includes a total of 1,000 phishing sites and an equal number of non-phishing sites. To gather these websites, we conducted web crawling starting from the collected seed URLs.

4.1 Phishing Sites

To collect phishing sites, we used OpenPhish [8] and PhishTank [9], which are phishing intelligence sources, and CrowdCanary [27], a method for extracting phishing-related posts from Twitter using machine learning models. By using keyword-based filters such as “phishing” and “scam,” CrowdCanary extracts URLs from posts associated with reports of phishing encountered by victims and observed by experts. We collected URLs from March to April 2023 and used them as seed URLs. This enabled us to obtain a wide range of phishing sites that are not limited to incoming channels such as email, SMS, and web advertisements. Subsequently, we accessed the seed URLs using the web crawler described in Section 3.1. To ensure the integrity of our dataset, we manually excluded websites with incomplete rendering and image loading. We conducted a thorough examination of the screenshot images of websites, focusing on two factors:

- Websites that display logos or brand names associated with well-known services or that visually resemble legitimate sites.

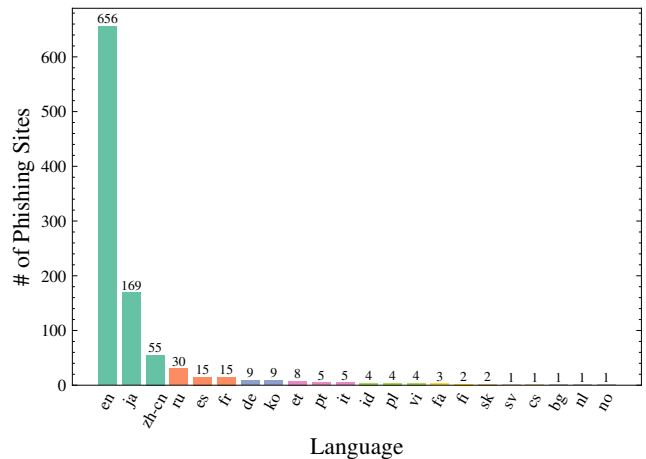


Figure 5: Language Distribution of Phishing Sites.

- Websites employing SE techniques, such as presenting false information (e.g., fake rewards, virus infection alerts, account issues), to generate a sense of urgency or interest.

From the gathered website information, we excluded those meeting the following conditions:

- Websites that contain R18 (adult, gambling, dating) content due to potential conflicts with ChatGPT’s content filter.
- In cases where similar screenshots were found among the extracted phishing sites, all but one were excluded. The criteria for determining similarity included the Fully Qualified Domain Name (FQDN), the page title, the website’s appearance, and the text content. However, if websites have the same appearance but different languages, we keep them.

As a result of the above analysis, we obtained a dataset of 1,000 phishing sites with unique 1,000 FQDNs. These phishing sites targeted a total of 147 legitimate service brands (see Appendix B). We identified 32 phishing sites that did not contain any brand information. Figure 4 shows the top 30 brands used by phishing sites in the dataset. In addition, we extracted text from the screenshot images using OCR. In this paper, we used Azure Cognitive Service [2] to streamline the OCR process for multilingual websites. We identified the languages of the OCR-extracted text using langdetect [7], a language detection library, revealing that the 1,000 phishing sites were distributed across 22 different languages. Figure 5 illustrates the number of phishing sites in each language.

4.2 Non-phishing Sites

We collected an equal number of websites as phishing sites for non-phishing sites. The seed URLs consisted of legitimate websites for the 153 brands targeted by phishing sites, as well as the top 2k domain names from the Tranco list [33]. From the legitimate sites of the 153 brands, we extracted the URLs of their homepages and, if available, the login pages (a total of 236 URLs with 196 FQDNs). We accessed the seed URLs using a web crawler. After crawling the Tranco top 2k, we successfully accessed 1,661 URLs. The Tranco top sites include certain categories such as adult content, illegal downloading/streaming, and gambling. Since these categories can trigger

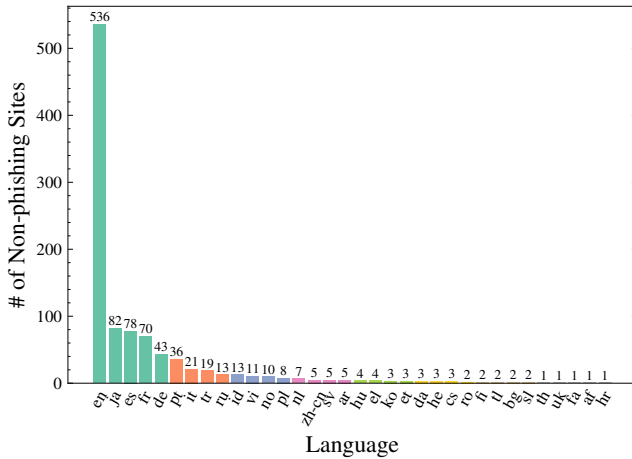


Figure 6: Language Distribution of Non-phishing Sites.

content filters in ChatGPT or lead to phishing sites through malicious advertisements, we excluded such websites. The breakdown of excluded sites is as follows: 28 porn, 3 gambling, and 9 illegal sites. The remaining 764 websites (out of 1,000 non-phishing sites) were randomly selected from the crawled Tranco top sites. The non-phishing sites are distributed across 34 languages. Figure 6 presents the distribution of non-phishing sites across different languages.

4.3 Simplifying HTML and OCR-Extracted Text

We simplified the collected HTML and OCR-extracted text of phishing sites and non-phishing sites using the process explained in Section 3.3 and 3.4. Figure 2 shows the comparison of token counts for the HTML content before and after simplification. The median token count for non-phishing sites was 74,937 before simplification, while it was 8,275 for phishing sites. In general, non-phishing sites have a higher token count. One reason for this difference is that non-phishing sites often implement a variety of sophisticated features through complex JavaScript code or use large platforms, while phishing sites have only the minimum functionality necessary to deceive users or steal sensitive information. Of the 2,000 sites in the dataset, 980 exceed the current ChatGPT model’s maximum token limit of 32,000. Therefore, simplification is required for website inspections under the existing ChatGPT model. Figure 3 compares the token counts of OCR-extracted text before and after simplification. The median token count for non-phishing sites was 296, whereas for phishing sites, it was 122. Similar to the HTML comparison, non-phishing sites tend to have a higher token count in this case as well.

5 EVALUATION

To assess the detection accuracy of the proposed method for phishing sites, we conducted an evaluation experiment using the Azure OpenAI API [3] with our own dataset. We used two models: gpt-4-0314 (referred to as GPT-4) and gpt-3.5-turbo-0301 (referred to as GPT-3.5). To compare their detection accuracy, we analyzed the responses from both models for each URL in the dataset using the same prompt. The dataset and experimental results cannot be made

Table 2: Confusion Matrix for GPT-3.5

		Predicted	
		Phishing	Non-phishing
Actual	Phishing	TP 867	FN 133
	Non-phishing	FP 15	TN 985

Table 3: Confusion Matrix for GPT-4

		Predicted	
		Phishing	Non-phishing
Actual	Phishing	TP 984	FN 16
	Non-phishing	FP 17	TN 983

Table 4: Performance Metrics Comparison

	GPT-3.5	GPT-4
Precision	98.3%	98.3%
Recall	86.7%	98.4%
Accuracy	92.6%	98.4%
F-measure	92.1%	98.4%

publicly available due to the inclusion of commercially licensed data and copyright issues, but we can provide them to researchers upon request.

5.1 Summary of Result

To extract detection results from the API responses, we defined a response as phishing if either the value of the *phishing* key or the *suspicious_domain* key in the JSON response was true. Conversely, a response was labeled as non-phishing if both keys were false. Due to probabilistic fluctuations in ChatGPT’s output, some responses did not contain machine-readable JSON text, such as using a single quote instead of a double quote. Among the responses, 2 (0.2%) for GPT-4 and 13 (1.3%) for GPT-3.5 resulted in JSON parse errors. However, since all responses included the corresponding answers for each key, we heuristically parsed and aggregated them.

Table 2 and Table 3 present the phishing detection results of GPT-3.5 and GPT-4. GPT-3.5 had 15 false positives (FPs), while GPT-4 had 17 FPs, indicating similar accuracy for non-phishing sites. On the other hand, GPT-3.5 had a higher number of false negatives (FNs) with 133 websites, whereas GPT-4 improved with only 16 websites. These results indicate that GPT-4 demonstrates an enhanced predictive accuracy for identifying actual phishing sites compared to GPT-3.5. Out of the 867 true positives (TPs) for GPT-3.5, 759 responses marked *phishing* as true, 796 marked *suspicious_domain* as true, and 688 had both marked as true. For GPT-4, out of the 984 TPs, 906 marked *phishing* as true, 983 marked *suspicious_domain* as true, and 905 had both marked as true.

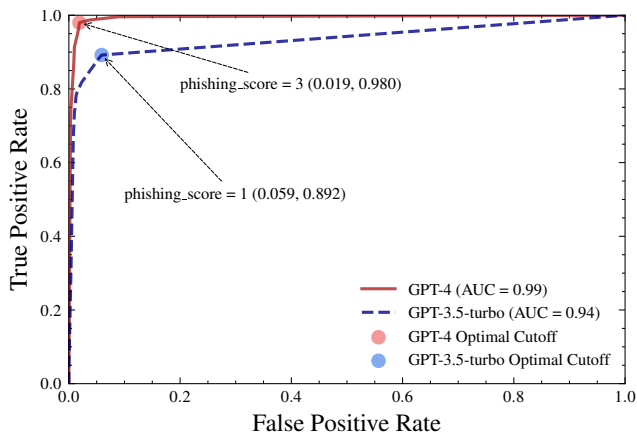


Figure 7: ROC Curve for the `phishing_score`.

Table 4 shows the results of the performance metrics comparison. Precision ($= \frac{TP}{TP+FP}$), Recall ($= \frac{TP}{TP+FN}$), Accuracy ($= \frac{TP+TN}{TP+TN+FP+FN}$), and F-measure ($= 2 \times \frac{Precision \times Recall}{Precision+Recall}$) were used as performance metrics. The capability of non-phishing site classification (FP and TN) was comparable between GPT-3.5 and GPT-4, resulting in similar precision values. However, GPT-4 outperforms GPT-3.5 by approximately 5.8% to 11.7% in the other metrics. Despite the inclusion of non-English local brands in the dataset, both GPT-3.5 and GPT-4 achieved accurate non-phishing site classification. On the other hand, for phishing site classification, GPT-4 significantly improved with a recall of 98.4% compared to GPT-3.5's 86.7%. Based on these results, the subsequent sections will provide a detailed analysis of the detection capability of the proposed method.

Additionally, we can classify the responses by selecting an appropriate threshold for the `phishing_score` value. By incrementing the threshold by 1, we calculated the true positive rate and false positive rate and plotted the receiver-operating characteristics (ROC) curve, as illustrated in Figure 7. The optimal cut-off values for `phishing_score`, calculated using Youden's J statistic ($J = \frac{TP}{TP+FN} + \frac{TN}{TN+FP} - 1$), were 3 for GPT-4 and 1 for GPT-3.5. The area under the curve (AUC) was 0.99 for GPT-4 and 0.94 for GPT-3.5. The detection accuracy based on the threshold setting for `phishing_score` also demonstrated the superior performance of GPT-4 compared to GPT-3.5. The corresponding performance metrics for this case can be found in Appendix A.

5.2 Examples of ChatGPT's Responses

In this section, we explain specific examples of ChatGPT's response when it correctly detected phishing sites.

Phishing Example A

Figure 8 displays a screenshot image of the phishing site ([https://securitypages.start\[.\]page/](https://securitypages.start[.]page/)), and Response 1 represents the response generated by ChatGPT. This website impersonates Facebook and attempts to create a sense of urgency regarding account issues for the user. ChatGPT correctly identified this website as being associated with the "Meta Facebook" brand and determines that the domain name is not legitimate Facebook. Furthermore, ChatGPT recognized the presence of SE techniques and highlights

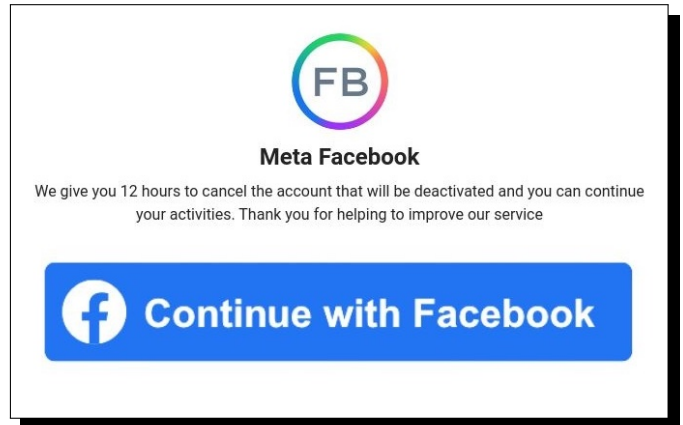


Figure 8: An Example of Facebook Phishing Site.

the existence of suspicious links within the HTML. ChatGPT set the `phishing_score` to 9, marking both `phishing` and `suspicious_domain` as true. Below, we have defanged some of the URLs included in the response.

Response 1

1. Suspicious elements:
 - The URL contains a subdomain "securitypages", which is unrelated to the legitimate Facebook domain (www.facebook[.]com).
 - The extracted text alerts the user about having 12 hours to cancel the account. This is a social engineering technique often used in phishing attacks.
 - Furthermore, the link provided redirects to a suspicious, non-legitimate URL: "http://centrenotificationvice.s.duckdns[.]org/cntre.html".

2. The brand name identified is Meta Facebook.

3. The site appears to be a phishing site because of the suspicious URL and social engineering tactics in the text.

4. JSON-formatted output:

```

---
{
  "phishing_score": 9,
  "brands": "Meta Facebook",
  "phishing": true,
  "suspicious_domain": true
}
---
```

Phishing Example B

Figure 9 illustrates a phishing site ([https://siges.com\[.\]br/service/](https://siges.com[.]br/service/)) impersonating DHL, attempting to deceive users into providing their financial information under the pretext of a

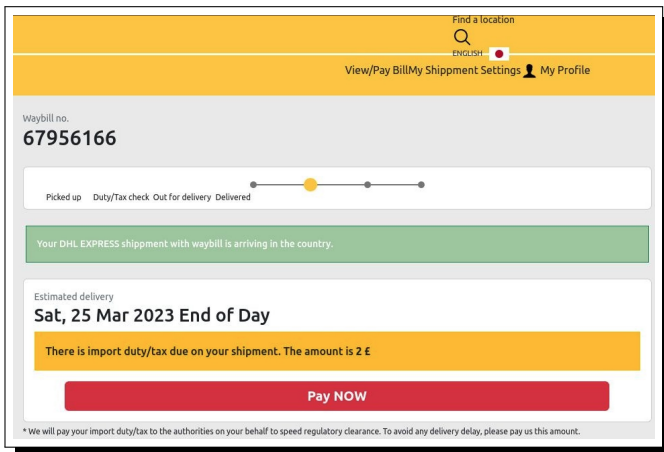


Figure 9: An Example of DHL Phishing Site.

payment requirement for receiving a package. ChatGPT accurately recognized the brand name as DHL and correctly identified the SE technique of requesting payment for import duty/tax expenses. It also determined that the domain name does not belong to the legitimate DHL website. Furthermore, by analyzing the HTML structure, ChatGPT provided detailed insights into the structure of the phishing site, which attempts to steal credit card information such as card numbers and CVC numbers.

Response 2

Based on my analysis of the HTML, URL, and OCR-extracted text, here are my findings:

1. Suspicious elements in the HTML and OCR-extracted text:

a. The site prompts the user to pay an import duty/tax amounting to 2€ on their shipment, which could be a social engineering technique to trick the user into providing their payment information.

b. The HTML structure has a card data input form where the user is expected to enter their card details, such as card number, expiry date, CVC number, and cardholder's name. This could be an attempt to steal sensitive card information.

2. The brand name identified is "DHL EXPRESS".

3. The URL (<https://siges.com/br/service/>) does not match the legitimate domain name associated with the DHL brand, which is "dhl.com". This mismatch is a red flag as it is not coming from the legitimate source.

Considering my findings, I conclude that this site is a phishing site attempting to steal user payment information by posing as a DHL EXPRESS page.

Here's the JSON-formatted output of my findings:

```
{
  "phishing_score": 9,
  "brands": "DHL EXPRESS",
  "phishing": true,
  "suspicious_domain": true
}
```

5.3 Evidence for Identifying Phishing Sites

We analyze the evidence used by GPT-4 to identify phishing sites. We provide a detailed explanation of each of the seven types of evidence used to detect phishing sites, including the characteristics of the corresponding SE techniques and how ChatGPT identified them.

Domain Name is Not Legitimate The most crucial evidence for identifying phishing sites is when the inspected website impersonates a legitimate brand but has a domain name that does not match the official domain name. The prompt of our proposed method instructs ChatGPT to extract the brand name from the HTML or OCR-extracted text and compare it with the legitimate domain name. In many cases, ChatGPT accurately detected phishing sites by identifying inconsistencies between the extracted brand and its official domain name. ChatGPT can also detect if the domain name included in the provided URL is a fake domain name (domain squatting [20, 40]) attempting to deceive users. For instance, ChatGPT correctly recognized that `disc0rd[.]pro`, which displayed a download page for Discord's installer, was not the legitimate domain name `discord[.]com`. It also identified that `amazon761[.]com`, which displayed the installation page for Amazon's mobile app, did not match `amazon[.]com`. Moreover, ChatGPT was able to discern that `aib[.]monster`, which showed the login page for Allied Irish Bank, although having the same second-level domain as the legitimate one (`aib[.]ie`), was not legitimate due to the different top-level domain (TLD). The explanation is as follows: *The URL provided in the question is "http://aib[.]monster/", but the actual legitimate domain for AIB is "https://www.aib[.]ie/".*

Fake Virus Infection Warning We also observed the effectiveness of ChatGPT in identifying fake malware infection warnings, a commonly used social engineering attack on the web. For instance, the website `landiingpages[.]beauty`, which displayed a screen resembling Microsoft Windows Defender's virus scan, was analyzed to display fake security warnings such as *このPCへのアクセスはセキュリティ上の理由でブロックされています (Access to this PC is blocked for security reasons)* and *脅威を発見 (Threat found) - Trojan Spyware App: Ads.financetrack(1).exe*. ChatGPT determined that these warnings aim to make the user believe their computer is in danger. Furthermore, the website `plfkwyacu.duckdns[.]jorg`, which displayed the logo of mobile carrier SoftBank, was recognized for a series of social engineering tactics. After a fake security warning stating *マルウェアが検出されました (Malware detected)*, the user was urged to install the *SoftBankセキュリティ無料版アプリ (SoftBank Security App Free Edition)*.

Fake Account Issue Phishing sites often display warnings claiming that there are issues with the user’s account, aiming to steal their login credentials. An example of this is the website `m4agence.web[.]app`, which ChatGPT analyzed as a phishing site targeting the bank Société Générale. ChatGPT identified that it employed a SE technique as follows: *The text extracted using OCR alerts the user to a problem with their account, which can be a typical technique of phishing attacks.* Furthermore, ChatGPT explained the presence of suspicious HTML elements on this website as follows: *The HTML contains a link to a different domain than the one specified in the visible text: "https://dev-sgwebnetauth.pantheonsite[.]io/wp-content/sysconnect" instead of "www.societegenerale[.]fr/synchronisation-agence."* Thus, ChatGPT has a high capability of HTML analysis, such as detecting inconsistencies between the text of the `a` element and the link specified in the `href`.

Urgent Payment Request As shown in Section 5.2, phishing sites that impersonate courier services and generate alerts about package issues, while creating a sense of urgency for payment, are targeting brands worldwide. An example is `royalmail-online[.]com`, which disguised as Royal Mail and requested additional payment due to the overweight of the shipment. ChatGPT identified this as a phishing site based on the mismatch in the domain name and the presence of phrases such as *Interrupted delivery* and *Additional charges of 0.76€*, indicating the employment of SE techniques. In the case of a phishing site (`www.caaarem[.]mx`) impersonating UPS, ChatGPT detected the suspicious phrase *790,45 HUF fizetés: Szállítási költség (Pay 790,45 HUF: Shipping cost)*. Additionally, ChatGPT provided an explanation stating, *The request for the user’s credit card information, including card number, expiration date, and security code.* By analyzing the HTML, ChatGPT gained a detailed understanding of the elements prompting user input. Moreover, `tlmcjohsvz.duckdns[.]org` was identified as a phishing site demanding payment for outstanding tax obligations. ChatGPT provided the following statement as evidence: *The HTML contains alerting text that seems like a phishing attempt, such as “差押最終通知” (Final Attachment Notice), informing the user about an unpaid tax and urging them to pay using specific methods.*

Fake Login Error As demonstrated in Section 5.2, some phishing sites display misleading error messages, such as indicating that the login credentials are wrong even when no input has been provided or alerting users that their credit card number is invalid. These SE techniques aim to deceive users and extract sensitive information. For instance, the website `www.interceptionbookingconfirmation[.]com`, pretending to be Booking.com, prompted users to input their card information within 24 hours, accompanied by an alert stating, *Your credit card was marked as invalid.* ChatGPT explained as follows: *The presence of a "Wrong Credential" error message in the HTML text is a common social engineering technique used to trick users into re-entering their login credentials.* Similarly, ChatGPT identified the website `mail.sikkimrajshree[.]in`, which impersonated the logistics company SF Express. Based on the displayed text on the website, ChatGPT categorized it as a phishing site, providing the following statement as evidence: *The HTML content contains an error message ‘无效的用户名或密码’ which translates to ‘Invalid username or password’, indicating that the user’s account*

might have a problem, which can be a social engineering technique used in phishing attacks.

Fake Reward Phishing sites often employ deceptive tactics to capture users’ interest, and one common initial step is to offer fake rewards, such as monetary prizes, cryptocurrencies, gift cards, or popular products like smartphones. For instance, the website `www.mobilegoodies4you[.]com` impersonates SFR, a telecommunication company, and displayed a fraudulent cash prize. ChatGPT identified the phrase *Votre adresse IP a été tirée au sort et vous avez une chance de gagner un 460,00 € en espèces! (Your IP address has been entered into a draw for a chance to win €460.00 in cash!)* within the website as an example of an unexpected reward. Similarly, ChatGPT identified the phishing site `allesettlemenie[.]top`, which impersonates the e-commerce brand Allegro, as using SE techniques to captivate users’ interest such as *Sign up or log in today to shop and earn cash! Withdrawal to bank account (limited to one mobile number)*. In the case of a phishing site (`me7q1.vetedeb[.]xyz`) targeting Sberbank, users are prompted to participate in a survey in exchange for cash. ChatGPT recognized the use of a cash reward to attract users’ interest, as indicated by the statement: *The text contains social engineering techniques such as enticing users with rewards (‘П О Л И У Ч И Т Ь В О З Н А Г Р А Ж Д Е Н И Е Д О 600 000 Р У Б ’), which means “get reward up to 600,000 RUB.”*

SMS Authentication Request Phishing sites not only target users’ login credentials but also attempt to obtain SMS authentication codes from them. For instance, the website `demo.crustncakes[.]com` impersonates the payment service BenefitPay and deceived users into inputting an SMS code for the purpose of verifying a transaction. ChatGPT explained that this technique as follows: *The OCR-extracted text informs about an expired SMS code that needs to be re-entered, which can be considered as a social engineering technique to alert users.*

5.4 False Positives (GPT-4)

GPT-4 identified phishing sites with high accuracy, however, it contained a small number of false positives (17). We analyzed the causes of these false positives by categorizing them into seven factors based on the responses generated by ChatGPT.

Misidentifying SE Techniques As discussed in Section 5.3, ChatGPT demonstrated precise identification of various SE techniques commonly used in phishing sites. However, it falsely flagged some non-phishing sites as phishing. In our experiment, GPT-4 misidentified `cfspart.impots.gouv[.]fr` as containing SE techniques in HTML elements. This element, initially hidden using the `display:none` property, triggers an alert message upon receiving incorrect input in the form. GPT-4 recognized `impots.gouv[.]fr` as a legitimate domain name, however, it misidentified `cfspart.impots.gouv[.]fr` as a different domain name because of the existence of the subdomain “cfspart”. Although there were other cases where GPT-4 reported that some legitimate sites contained elements using SE techniques, they were all classified as non-phishing because GPT-4 correctly verified that they matched legitimate domain names.

Multiple Domain Names of Legitimate Brands ChatGPT sometimes misclassifies legitimate brands that operate multiple domains as phishing sites. This occurs when ChatGPT can identify one or

more of these domain names but does not have the knowledge of the specific domain name that is being checked. In our experiment, GPT-4 produced four false positives for the following domain names: `m.botw[.]com` (an alternate domain of `bankofthewest[.]com`), `aplusfcu.balancepro[.]org` (`aplusfcu[.]org`), `hb.redlink.com[.]ar` (`bancode1apampa.com[.]ar`), and `cloudflare[.]net` (`cloudflare[.]com`). Additionally, it falsely flagged the law firm Pinsent Masons' website `www.aboutcookies[.]org` as phishing. This false positive occurred because the extracted brand name Pinsent Masons was not present in the domain name.

Global Brands with Local Domain Names False positives can occur when global brands use local domain names specific to different countries. GPT-4 produced false positives for three domain names due to their mismatch with the most recognized legitimate domain names: `www.aliexpress[.]us`, `sube.garantibva.com[.]tr`, and `www.sicredi.com[.]br`.

Numerous Subdomains ChatGPT sometimes fails to recognize a legitimate domain name if it has multiple subdomains, even if the domain name is under a legitimate domain name. The following domain names were false positives in this category: `cloud.walletconnect[.]com`, `entry11.bk.mufg[.]jp`, and `www.my.commbank.com[.]au`.

Non-English Websites While ChatGPT demonstrated high accuracy in identifying SE techniques across a wide range of languages, false positives can occur due to insufficient knowledge of non-English websites. The following two domain names were falsely flagged as they were not well-known: `justhost[.]ru` and `www.jb51[.]net`.

Long URL Query String ChatGPT correctly recognized the website `auth.talktalk.co[.]uk` as legitimate and did not detect any SE techniques within the web page content. However, it falsely flagged this site as phishing due to an extremely long URL query string (360 characters) in the URL path.

5.5 False Negatives (GPT-4)

GPT-4 demonstrated an improvement in reducing false negatives compared to GPT-3.5, achieving an 11.7% reduction. However, there are cases where GPT-4 erroneously classifies certain phishing sites as non-phishing (16 cases), necessitating a thorough analysis of the underlying causes. Out of the 16 FNs, except for 2 cases, GPT-4 correctly identified the brand names of the remaining websites. Among these FNs, 8 websites lacked any descriptive text other than JSON-formatted content in their responses, making it impossible to analyze the basis of their classification. There were 4 FNs where GPT-3.5 correctly identified the phishing sites, but GPT-4 did not.

Falsely Identified Domain Names as Legitimate Despite correctly identifying the targeted brand names for the following 5 phishing sites, GPT-4 erroneously classified the domain names as legitimate. For instance, GPT-4 misclassified the phishing site `www.phototan-push[.]de` (disguised as `commerzbank[.]com`) by stating *there's no suspicion surrounding the domain name*, thus incorrectly labeling it as a non-phishing site. The phishing site `www.gmenduroracingteam[.]it` was identified as non-phishing by GPT-4, because the legitimate URL of the brand (`https://www.mooney[.]it/`) that this phishing site deceived was included as

a link in the HTML, even though the actual domain name is different. The following three sites were identified as false negatives by GPT-4, but correctly identified by GPT-3.5. A website impersonating `www.gov[.]uk` (`claim.redundancy-payments.org[.]uk`) was classified as phishing by GPT-3.5 due to its SE technique of requesting the national insurance numbers and bank details. In contrast, GPT-4 accurately identified the brand as "GOV.UK" but indicated *it is likely that the site is a legitimate one*. GPT-4 also suggested *it is always better to verify the authenticity of the URL by visiting the official government website*. `etmmetaverse[.]com`, a phishing site for Office 365, was correctly classified as phishing by GPT-3.5 based on the mismatch in the domain name and links to other pages. However, GPT-4 offered a neutral explanation, stating *there is not enough evidence supporting that this is a phishing site*, despite assigning a `phishing_score` of 5 and labeling both `phishing` and `suspicious_domain` as `unknown`. Furthermore, both GPT-3.5 and GPT-4 correctly identified the brand SI-DEP (France's national Covid-19 screening system) for the phishing site `ca8567c7fa0141658f08b0dabe13d5ee.v1.radwarecloud.net`. While GPT-3.5 correctly pointed out *The domain name "radwarecloud.net" does not appear to be associated with the brand name "SI-DEP"*, GPT-4 failed to detect the phishing attempt.

Failure to Identify Domain Squatting The website `www.americanexpressseguros[.]com` was a phishing site for American Express, offering insurance solicitations. Despite containing the phrase *Recibe hasta \$2,000.00 M.N de bonificación al contratar (Receive up to \$2,000.00 M.N of bonus when contracting)*, GPT-4 failed to recognize the SE technique employed. Although it was a case of domain squatting, GPT-4 incorrectly identified it as a legitimate domain.

Failure to Identify SE Techniques The phishing site `phpstack-197144-1061735.cloudwaysapps[.]com` presented an insurance survey, without specifically targeting any particular brand. Since it lacks prominent elements such as rewards or a sense of urgency, it was mistakenly classified as a legitimate site. Similarly, the e-commerce site `lojanewgeneration[.]com` attempted to attract users with the phrase *daily offers*, yet GPT-4 failed to identify the SE technique employed, leading to its classification as a legitimate site.

5.6 Comparative Analysis of Phishing Site Detection: GPT-4 vs. GPT-3.5

We conducted a comparative analysis of the capabilities of GPT-4 and GPT-3.5 in accurately identifying phishing sites. Specifically, we thoroughly analyzed the phishing sites that were successfully detected by GPT-4 but failed by GPT-3.5. Among the 133 FNs of GPT-3.5, a total of 121 phishing sites were correctly identified by GPT-4. Upon analyzing these 121 phishing sites, we discovered three abilities in which GPT-4 exhibited superior performance compared to GPT-3.5:

- Ability to determine the authenticity of domain names.
- Ability to recognize SE techniques.
- Ability to detect phishing sites comprehensively by using multiple factors.

Ability to Determine the Suspiciousness of Domain Names GPT-4 and GPT-3 showed differences in their ability to assess

the suspiciousness of domain names. This involves determining whether the domain name matches a legitimate domain name or has suspicious characteristics. For instance, when analyzing the OpenAI phishing site `openai-gpt-4[.]com`, GPT-3 classified it as non-phishing due to the presence of the term “openai.” In contrast, GPT-4 correctly identified it as phishing since the actual domain name, `openai[.]com`, differs from the suspicious one. In the case of a phishing site (`krakken-logi.mystrikingly[.]com`) impersonating the cryptocurrency exchange Kraken, both GPT-3 and GPT-4 successfully recognized the brand name. However, GPT-4 identified the misspelling of “krakken-logi” as suspicious, whereas GPT-3 erroneously classified it as a legitimate domain name based on the presence of “krakken”. Furthermore, both GPT-3.5 and GPT-4 accurately identified the phishing site `firstcitiznbc[.]com` as distinct from the legitimate domain name `firstcitizens[.]com`. However, GPT-3.5 considered it to be a variable of the legitimate domain name and not a phishing site. Another example is the phishing site `correos-es.firebaseio[.]com`, which impersonates Spain’s national postal service, Correos. While GPT-3.5 classified it as a legitimate domain name, GPT-4 correctly recognized it as a suspicious domain due to its hosting on Google Firebase.

Ability to Identify SE Techniques GPT-4 successfully identified SE techniques that were missed by GPT-3.5 in some phishing sites. For example, in the case of `verify.vodafone-uk[.]com`, which GPT-3.5 explained that *The page warns users about the potential termination of their phone number*. However, GPT-4 recognized the presence of SE techniques by explaining that the websites warned users about the potential termination of their phone number. Additionally, GPT-4 noted that *The domain name in the URL is "verify.vodafone-uk.com", which seems odd with the hyphen between "vodafone" and "uk"*. In the phishing site (`aktiverer-bnkid.web[.]app`) impersonating Sweden’s ID system, BankID, GPT-4 identified the SE technique by highlighting that it requested personal information such as the user’s personal number, phone number, and user-ID. In contrast, GPT-3 incorrectly states that the website *does not contain any obvious social engineering techniques*.

Ability to Detect Phishing Sites Comprehensively by Using Multiple Factors In some cases, ChatGPT may have incorrectly determined that legitimate text on the website was suspicious. For example, it identified legitimate error messages (e.g., “wrong password”) displayed after submitting a form as indicative of social engineering. It also flagged the inclusion of the year 2023 in a copyright notice as suspicious. This is because ChatGPT’s training data only goes back to September 2021, and it lacks knowledge of the current year, which could lead to inaccurate results. However, GPT-4, while occasionally providing irrelevant observations, tends to make overall accurate determinations by combining them with more reliable information. On the other hand, GPT-3.5 may classify a phishing site as benign even though it correctly recognizes messages using SE techniques in the HTML, or the domain name differs from the legitimate one. The following are some examples of non-phishing sites. For instance, the website `b2meguy.com` displayed an attention-grabbing statement at the top of the website: *Obtenga acceso instantáneo al software de Quantum AI, y gane de 3200€ al día! (Get instant access to Quantum AI’s software, and earn from 3200€ per day!)*. Both GPT-3.5 and GPT-4 identified this statement as suspicious, however, while GPT-3.5 did not classify it as a phishing site,

GPT-4 correctly identified it as phishing based on the explanation that *it presents unrealistic promises of gains*. Another example is the phishing site `verification-appeal-code.firebaseio[.]com`, which pretended to be Facebook and asked users to respond to claims of intellectual property infringement. Both GPT-3.5 and GPT-4 detected the request for sensitive information from visitors, however, GPT-3.5 incorrectly classified it as a legitimate site.

6 LIMITATIONS

ChatGPT provides flexible and varied responses because its output is determined probabilistically. However, this also means that detection results can change depending on the experiment. In this study, we used default parameters for ChatGPT and prepared a dataset consisting of various websites to statistically evaluate the detection accuracy. In order to obtain reliable results for the classification of phishing sites, some parameters such as *Top_k* and *Temperature* may require adjustments for specific purposes.

For the dataset used in this experiment, we excluded phishing sites and non-phishing sites containing R-18 content to eliminate the influence of ChatGPT’s content filter. To examine websites with adult or aggressive content, it will be necessary to modify or select an appropriate content filter for ChatGPT.

Methods that use LLMs, including the proposed approach, may be susceptible to prompt injection attacks [17, 32], where the original prompt is overwritten and malicious content is inserted into the response. Our proposed method employs simple countermeasures by simplifying the HTML and clarifying text sections through the placement of triple backticks. However, phishing sites may use prompt injections in the future to avoid analysis by LLMs. More advanced defenses against prompt injection attacks will be needed.

Because ChatGPT has been trained on data up to September 2021, it may not be able to accurately classify phishing and non-phishing sites associated with services created after that date. The classification ability could potentially be improved by externally referencing a list of domain names corresponding to brands or by using a fine-tuned LLM.

7 RELATED WORK

An effective method for detecting phishing sites is to compare the appearance of a website with that of a legitimate one. Previous studies such as Abdelnabi et al. [11], Lin et al. [24], and Liu et al. [25] detect the abuse of logo images by identifying and comparing them with legitimate ones, or extract features from the overall appearance of web pages to analyze their similarity. While these deep learning-based methods are effective in detecting phishing sites created by copying HTML or reusing logos, they cannot detect phishing sites that use their own logos or do not impersonate brands. Our proposed method aims to identify phishing sites that use SE techniques by analyzing the context of the web page, allowing for the detection of different types of phishing sites, not limited to the misuse of branding.

Methods have been proposed to detect domain squatting, where attackers obtain domain names similar to legitimate service domain names. Nikiforakis et al. [29], Quinkert et al. [34], Kintis et al. [20], Agten et al. [12], and dnstwist [1] have proposed techniques that use rule-based approaches or machine learning models to identify

malicious domain names by considering various factors such as minor variations, the inclusion of brand names, and changes in top-level domains. ChatGPT showed high capabilities against domain squatting, as demonstrated in our evaluation experiments, by detecting minor character differences and fake domain names containing brand names. There are also methods that use TLS certificates to detect phishing sites. Kim et al. [19], Bijmans et al. [13], and Drichel et al. [16] have proposed such approaches that use machine learning algorithms to analyze features derived from TLS certificates or certificate transparency logs to detect phishing sites.

In recent years, there have been attempts to apply LLMs to cybersecurity. One such example is VirusTotal Code Insight [10], which is based on Google’s security-focused LLM, Sec-PaLM. This model can explain the purpose and functionality of malware code in natural language, allowing security experts to gain a general understanding of the code’s intentions. An article [6] analyzed phishing URLs using ChatGPT with a simple prompt. The experiment showed a detection rate of 87.2% and a false positive rate of 23.2%, indicating that while it can block phishing sites, it can also mistakenly block legitimate sites. In contrast, the method proposed in this paper includes not only the URL, but also HTML, text extracted from screenshots using OCR, and carefully designed prompts to achieve superior detection accuracy. Roy et al. [36] explored the ability of ChatGPT to automatically generate phishing sites. Their findings suggest that these generated phishing sites can mimic popular brands and employ various evasion tactics to evade anti-phishing systems. To counteract such automated phishing campaigns, it becomes important to employ security measures using LLMs, such as our proposed method.

8 CONCLUSION

In this study, we have presented a novel method for detecting phishing sites using ChatGPT, a state-of-the-art large language model. Our approach combines web crawling techniques with ChatGPT’s contextual understanding to classify websites as phishing sites or not. Through detailed experiments, our method achieved remarkable performance with a precision of 98.3% and a recall of 98.4%, demonstrating the potential of LLMs to efficiently detect phishing sites and uncover social engineering techniques. Furthermore, our comparative analysis between GPT-3.5 and GPT-4 revealed significant improvements in GPT-4’s ability to minimize false negatives. GPT-4 excelled at identifying suspicious domains, detecting social engineering techniques in web content, and considering multiple factors in its evaluation. These findings have important implications for enhancing automated cybersecurity measures and mitigating the risks of online fraudulent activities faced by users. This research provides new directions for exploring the application of LLMs in various cybersecurity domains and further optimizing their capabilities for analyzing and detecting malicious web content.

REFERENCES

[1] [n. d.]. GitHub - elceef/dnstwist. <https://github.com/elceef/dnstwist>
[2] 2023. Azure Cognitive Service for Vision with OCR and AI | Microsoft Azure. <https://azure.microsoft.com/en-us/products/cognitive-services/vision-services>
[3] 2023. Azure OpenAI Service – Advanced Language Models | Microsoft Azure. <https://azure.microsoft.com/en-us/products/cognitive-services/openai-service>

[4] 2023. Chrome DevTools Protocol. <https://chromedevtools.github.io/devtools-protocol/>.
[5] 2023. Introducing ChatGPT. <https://openai.com/blog/chatgpt>
[6] 2023. Investigating ChatGPT phishing detection capabilities | Securelist. <https://securelist.com/chatgpt-anti-phishing/109590/>
[7] 2023. Mimino666/langdetect: Port of Google’s language-detection library to Python. <https://github.com/Mimino666/langdetect>
[8] 2023. OpenPhish. <https://openphish.com/>
[9] 2023. PhishTank. <https://phishtank.org/>
[10] 2023. VirusTotal Code Insight. <https://blog.virustotal.com/2023/04/introducing-virustotal-code-insight.html>
[11] Sahar Abdelnabi, Katharina Krombholz, and Mario Fritz. 2020. VisualPhishNet: Zero-Day Phishing Website Detection by Visual Similarity. In *CCS ’20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*, Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna (Eds.). ACM, 1681–1698.
[12] Pieter Agten, Wouter Joosen, Frank Piessens, and Nick Nikiforakis. 2015. Seven Months’ Worth of Mistakes: A Longitudinal Study of Typosquatting Abuse. In *22nd Annual Network and Distributed System Security Symposium, NDSS 2015, San Diego, California, USA, February 8-11, 2015*. The Internet Society.
[13] Hugo L. J. Bijmans, Tim M. Booij, Anneke Schwedersky, Aria Nedgabat, and Rolf van Wegberg. 2021. Catching Phishers By Their Bait: Investigating the Dutch Phishing Landscape through Phishing Kit Detection. In *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, Michael Bailey and Rachel Greenstadt (Eds.). USENIX Association, 3757–3774. <https://www.usenix.org/conference/usenixsecurity21/presentation/bijmans>
[14] M. Bitaab, H. Cho, A. Oest, Z. Lyu, W. Wang, J. Abraham, R. Wang, T. Bao, Y. Shoshitaishvili, and A. Doupe. 2023. Beyond Phish: Toward Detecting Fraudulent e-Commerce Websites at Scale. In *2023 IEEE Symposium on Security and Privacy (SP) (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 2566–2583. <https://doi.org/10.1109/SP46215.2023.00152>
[15] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. *CoRR abs/2005.14165* (2020). arXiv:2005.14165 <https://arxiv.org/abs/2005.14165>
[16] Arthur Drichel, Vincent Drury, Justus von Brandt, and Ulrike Meyer. 2021. Finding Phish in a Haystack: A Pipeline for Phishing Classification on Certificate Transparency Logs. In *ARES 2021: The 16th International Conference on Availability, Reliability and Security, Vienna, Austria, August 17-20, 2021*, Delphine Reinhardt and Tilo Müller (Eds.). ACM, 59:1–59:12.
[17] Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. 2023. Not what you’ve signed up for: Compromising Real-World LLM-Integrated Applications with Indirect Prompt Injection. arXiv:2302.12173 [cs.CR]
[18] Amin Kharraz, William K. Robertson, and Engin Kirda. 2018. Surveillance: Automatically Detecting Online Survey Scams. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA*. IEEE Computer Society, 70–86. <https://doi.org/10.1109/SP.2018.00044>
[19] Doowon Kim, Haehyun Cho, Yonghwi Kwon, Adam Doupe, Soeul Son, Gail-Joon Ahn, and Tudor Dumitras. 2021. Security Analysis on Practices of Certificate Authorities in the HTTPS Phishing Ecosystem. In *ASIA CCS ’21: ACM Asia Conference on Computer and Communications Security, Virtual Event, Hong Kong, June 7-11, 2021*, Jiannong Cao, Man Ho Au, Zhiqiang Lin, and Moti Yung (Eds.). ACM, 407–420.
[20] Panagiotis Kintis, Najmeh Miramirkhani, Charles Lever, Yizheng Chen, Rosa Romero Gómez, Nikolaos Pitropakis, Nick Nikiforakis, and Manos Antonakakis. 2017. Hiding in Plain Sight: A Longitudinal Study of Combosquatting Abuse. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, Bhavani Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM, 569–586.
[21] Takashi Koide, Daiki Chiba, and Mitsuaki Akiyama. 2020. To Get Lost is to Learn the Way: Automatically Collecting Multi-step Social Engineering Attacks on the Web. In *ASIA CCS ’20: The 15th ACM Asia Conference on Computer and Communications Security, Taipei, Taiwan, October 5-9, 2020*, Hung-Min Sun, Shih-Pyng Shieh, Guofei Gu, and Giuseppe Ateniese (Eds.). ACM, 394–408. <https://doi.org/10.1145/3320269.3384714>
[22] Takashi Koide, Daiki Chiba, Mitsuaki Akiyama, Katsunari Yoshioka, and Tsutomu Matsumoto. 2020. It Never Rains but It Pours: Analyzing and Detecting Fake Removal Information Advertisement Sites. In *The 17th Conference on Detection of Intrusions and Malware & Vulnerability Assessment, DIMVA 2020*. 171–191.
[23] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large Language Models are Zero-Shot Reasoners. In *NeurIPS*.

http://papers.nips.cc/paper_files/paper/2022/hash/8bb0d291acd4acf06ef112099c16f326-Abstract-Conference.html

[24] Yun Lin, Ruofan Liu, Dinil Mon Divakaran, Jun Yang Ng, Qing Zhou Chan, Yiwen Lu, Yuxuan Si, Fan Zhang, and Jin Song Dong. 2021. Phishpedia: A Hybrid Deep Learning Based Approach to Visually Identify Phishing Webpages. In *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, Michael Bailey and Rachel Greenstadt (Eds.). USENIX Association, 3793–3810.

[25] Ruofan Liu, Yun Lin, Xianglin Yang, Siang Hwee Ng, Dinil Mon Divakaran, and Jin Song Dong. 2022. Inferring Phishing Intention via Webpage Appearance and Dynamics: A Deep Vision Based Approach. In *31st USENIX Security Symposium, USENIX Security 2022, Boston, MA, USA, August 10-12, 2022*, Kevin R. B. Butler and Kurt Thomas (Eds.). USENIX Association, 1633–1650.

[26] Najmeh Miramirkhani, Oleksii Starov, and Nick Nikiforakis. 2017. Dial One for Scam: A Large-Scale Analysis of Technical Support Scams. In *24th Annual Network and Distributed System Security Symposium, NDSS*.

[27] Hiroki Nakano, Daiki Chiba, Takashi Koide, Naoki Fukushi, Takeshi Yagi, Takeo Hariu, Katsunari Yoshioka, and Tsutomu Matsumoto. 2023. Canary in Twitter Mine: Collecting Phishing Reports from Experts and Non-experts. arXiv:2303.15847 [cs.CR]

[28] Terry Nelms, Roberto Perdisci, Manos Antonakakis, and Mustaque Ahamad. 2016. Towards Measuring and Mitigating Social Engineering Software Download Attacks. In *25th USENIX Security Symposium, USENIX Security 16*. 773–789.

[29] Nick Nikiforakis, Steven Van Acker, Wannes Meert, Lieven Desmet, Frank Piessens, and Wouter Joosen. 2013. Bitsquatting: exploiting bit-flips for fun, or profit?. In *22nd International World Wide Web Conference, WWW ’13, Rio de Janeiro, Brazil, May 13-17, 2013*, Daniel Schwabe, Virgílio A. F. Almeida, Hartmut Glaser, Ricardo Baeza-Yates, and Sue B. Moon (Eds.), 989–998.

[30] OpenAI. 2023. GPT-4 Technical Report. CoRR abs/2303.08774 (2023). <https://doi.org/10.48550/arXiv.2303.08774> arXiv:2303.08774

[31] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *NeurIPS*. http://papers.nips.cc/paper_files/paper/2022/hash/b1efde53be364a73914f58805a001731-Abstract-Conference.html

[32] Fábio Perez and Ian Ribeiro. 2022. Ignore Previous Prompt: Attack Techniques For Language Models. arXiv:2211.09527 [cs.CL]

[33] Victor Le Pochat, Tom van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczynski, and Wouter Joosen. 2019. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *26th Annual Network and Distributed System Security Symposium, NDSS 2019, San Diego, California, USA, February 24-27, 2019*. The Internet Society. <https://www.ndss-symposium.org/ndss-paper/tranco-a-research-oriented-top-sites-ranking-hardened-against-manipulation/>

[34] Florian Quinkert, Martin Degeling, Jim Blythe, and Thorsten Holz. 2020. Be the Phisher - Understanding Users’ Perception of Malicious Domains. In *ASIA CCS ’20: The 15th ACM Asia Conference on Computer and Communications Security, Taipei, Taiwan, October 5-9, 2020*, Hung-Min Sun, Shih-Pyng Shieh, Guofei Gu, and Giuseppe Ateniese (Eds.). ACM, 263–276.

[35] Alec Radford and Karthik Narasimhan. 2018. Improving Language Understanding by Generative Pre-Training.

[36] Sayak Saha Roy, Krishna Vamsi Naragam, and Shirin Nilizadeh. 2023. Generating Phishing Attacks using ChatGPT. arXiv:2305.05133 [cs.CR]

[37] Janos Szurdi, Meng Luo, Brian Kondracki, Nick Nikiforakis, and Nicolas Christin. 2021. Where are you taking me? Understanding Abusive Traffic Distribution Systems. In *WWW ’21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021*, Jure Leskovec, Marko Grobelnik, Marc Najork, Jie Tang, and Leila Zia (Eds.). ACM / IW3C2, 3613–3624.

[38] Siyuan Tang, Xianghang Mi, Ying Li, XiaoFeng Wang, and Kai Chen. 2022. Clues in Tweets: Twitter-Guided Discovery and Analysis of SMS Spam. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS 2022, Los Angeles, CA, USA, November 7-11, 2022*, Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi (Eds.). ACM, 2751–2764.

[39] Kurt Thomas, Frank Li, Ali Zand, Jacob Barrett, Juri Ranieri, Luca Invernizzi, Yarik Markov, Oxana Comanescu, Vijay Eranti, Angelika Moscicki, Daniel Margolis, Vern Paxson, and Elie Bursztein. 2017. Data Breaches, Phishing, or Malware?: Understanding the Risks of Stolen Credentials. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, Bhavani Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM, 1421–1434. <https://doi.org/10.1145/3133956.3134067>

[40] Ke Tian, Steve T. K. Jan, Hang Hu, Danfeng Yao, and Gang Wang. 2018. Needle in a Haystack: Tracking Down Elite Phishing Domains in the Wild. In *Proceedings of the Internet Measurement Conference 2018, IMC 2018, Boston, MA, USA, October 31 - November 02, 2018*. ACM, 429–442. <https://dl.acm.org/citation.cfm?id=3278569>

[41] Phani Vadrevu and Roberto Perdisci. 2019. What You See is NOT What You Get: Discovering and Tracking Social Engineering Attack Campaigns. In *Proceedings of the Internet Measurement Conference, IMC 2019*. 308–321.

[42] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. arXiv:2201.11903 [cs.CL]

[43] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *NeurIPS*. http://papers.nips.cc/paper_files/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html

[44] Peng Yang, Guangzhen Zhao, and Pengjie Zeng. 2019. Phishing Website Detection Based on Multidimensional Features Driven by Deep Learning. *IEEE Access* 7 (2019), 15196–15209.

[45] Zheng Yang, Joey Allen, Matthew Landen, Roberto Perdisci, and Wenke Lee. 2023. TRIDENT: Towards Detecting and Mitigating Web-based Social Engineering Attacks. In *32st USENIX Security Symposium, USENIX Security 2023, 2023*. USENIX Association, 1681–1698.

A DETECTION ACCURACY USING PHISHING SCORE

Table 5: Performance Metrics Comparison Using `phishing_score`.

	GPT-3.5	GPT-4
Precision	0.938	0.981
Recall	0.892	0.980
Accuracy	0.916	0.981
F-measure	0.914	0.980

Table 5 shows the results of calculating the detection accuracy metrics - Precision, Recall, Accuracy, and F-1 measure - using the optimal threshold values for `phishing_score` in GPT-3.5 and GPT-4.

B PHISHING BRANDS

Table 6 is a list of 147 brands that were targeted by phishing sites in our dataset.

Table 6: Brands Targeted by Phishing Sites.

A+ Federal Credit Union	AEON CARD	Alpha Web
Amazon	American Express	Apple
Ardoiz	BECU Online Banking	BNL
BRImo	Banca Sella	Banco Desio
Banco La Pampa	Bancoagrícola	Bancolumbia
Bank of America	Barclays	Battle.net
Bendigo Bank	Bradesco	CAJA
CIMB Ni	Chunghwa Post Co.	Citizens Bank
CodeSquare	Coin Wallet	Colissimo
Commonwealth Bank	Correos	Credomatic
Crypto.com	DHL	DKB
Deutsche Bank	Deutsche Post DHL Group	Disney
Docomo	DocuSign	Dropbox
EATA	EClick PORTAL	ETC
EXPRESSPAY	Ebay	Emirates Post
Entrust	Facebook	Fifth Third Momentum Banking
First Citizens Bank	GLOBAL PASS	GO online
GOV.UK	Garanti BBVA	Garena
Gazprom	Gemini	GitLab
Google	Gruppo BNP Paribas	IBC
ING	ImToken	Impots.gouv.fr
Instagram	Involta	JUHACHI-SHINWA BANK
JUNO	Kakao	Kusainon
La Banque Postale	Livelo	METAMASK
MIR VISA	M&T Bank	Ma Banque
Mashreq Online Banking	MasterCard	Mercari
Microsoft	Minnesota Unemployment Insurance Benefits System	Mitsubishi UFJ Bank
Mitsui Sumitomo Card	Mygrow	NAB Internet Banking
NAVY FEDERAL Credit Union	NETFLIX	NLB Banka
Nexi	Nordea	ORLEN
OakPay	OneDrive	Orange
OurTime	Outlook.com	PNC Online Banking
PRESTIA	PancakeSwap	PayPal
PayPay	Post canada	QUOTY
Qatar Post	RENNER	REVOLUT INTERNATIONAL BANK
Ronin Wallet	SAISONCARD	SBJ Bank
SDCCU	SFR	SMBC
SMBC Trust Bank	SPANKKI	STEAM
Santander	Satang	Schwab Safe
Scotiabank	Security Bank	SharePoint
Sicredi	Siemens	Slovenská POŠTA
Spotify	Swiss	Swiss Post
Swisscom	T-Mobile	TD Ameritrade
TESCO Bank	THEWEST	TalkTalk
Theta	Transcash	Trust Wallet
USAA	USPS Tracking	Ubisoft
UniCredit	Uphold	UPS
VALVE	VISA	Verizon
Vietcombank	WELLS FARGO	WalletConnect
WeTransfer	Weebly	Yahoo!