ShamFinder: An Automated Framework for Detecting IDN Homographs

Hiroaki Suzuki
Waseda University
Tokyo, Japan
sz0213hk@nsl.cs.waseda.ac.jp

Daiki Chiba
NTT Secure Platform Laboratories
Tokyo, Japan
daiki.chiba@ieee.org

Yoshiro Yoneya
Japan Registry Services
Tokyo, Japan
yone@jprs.co.jp

Tatsuya Mori
Waseda University/NICT/RIKEN AIP
Tokyo, Japan
mori@nsl.cs.waseda.ac.jp

Shigeki Goto
Waseda University
Tokyo, Japan
sgoto@waseda.jp

ABSTRACT

The internationalized domain name (IDN) is a mechanism that enables us to use Unicode characters in domain names. The set of Unicode characters contains several pairs of characters that are visually identical with each other; e.g., the Latin character ‘a’ (U+0061) and Cyrillic character ‘a’ (U+0430). Visually identical characters such as these are generally known as *homoglyphs*. IDN *homograph attacks*, which are widely known, abuse Unicode homoglyphs to create lookalike URLs. Although the threat posed by IDN homograph attacks is not new, the recent rise of IDN adoption in both domain name registries and web browsers has resulted in the threat of these attacks becoming increasingly widespread, leading to large-scale phishing attacks such as those targeting cryptocurrency exchange companies. In this work, we developed a framework named “ShamFinder,” which is an automated scheme to detect IDN homographs. Our key contribution is the automatic construction of a homoglyph database, which can be used for direct countermeasures against the attack and to inform users about the context of an IDN homograph. Using the ShamFinder framework, we perform a large-scale measurement study that aims to understand the IDN homographs that exist in the wild. On the basis of our approach, we provide insights into an effective countermeasure against the threats caused by the IDN homograph attack.

CCS CONCEPTS

• Security and privacy → Network security.

KEYWORDS

DNS, IDN homograph, Unicode, Homoglyph

ACM Reference Format:

Hiroaki Suzuki, Daiki Chiba, Yoshiro Yoneya, Tatsuya Mori, and Shigeki Goto. 2019. ShamFinder: An Automated Framework for Detecting IDN

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IMC '19, October 21–23, 2019, Amsterdam, Netherlands

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6948-0/19/10...\$15.00

<https://doi.org/10.1145/3355369.3355587>

Homographs. In *Internet Measurement Conference (IMC '19)*, October 21–23, 2019, Amsterdam, Netherlands. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3355369.3355587>

1 INTRODUCTION

Internationalized domain name (IDN) is a mechanism that allows us to use various non-English characters such as Arabic, Chinese, Cyrillic, Hangul, Hebrew, Hiragana, or Tamil. IDN was first proposed by Dürst in 1996 as an Internet Draft (I-D) [19]. Subsequently, a system known as Internationalizing Domain Names in Applications (IDNA) was adopted as an Internet standard [11]. Currently, the IDNA system is widely deployed in various domains including hundreds of top-level domains (TLDs). In addition, the majority of modern web browsers are capable of accommodating IDNs.

Character sets permitted to be used as IDNs contain several pairs of characters that are visually similar with each other. These characters are known as *homoglyphs*. The existence of homoglyphs enables an attacker to create a spoofing domain name. For instance, by using a Unicode character ‘é’, which is a Latin lowercase letter e with an acute accent (U+00E9), an attacker can create a spoofing domain name, “facebook.com,” which is visually similar to the original domain name “facebook.com.” The domain spoofing attack exploiting Unicode homoglyphs is known as “IDN homograph attack” and has been used for malicious purposes such as phishing attacks. IDN homograph attacks are not a new problem. In 2002, Gabrilovich and Gontmakher [12] demonstrated that they successfully registered an IDN homograph using the two Russian letters ‘c’ and ‘o’.

As the adoption rate of IDN was not high in the past, an IDN homograph has been recognized as a proof-of-concept attack. However, the recent rise in the number of IDN registrations and the adoption of an IDN in many TLDs together with the adoption of IDNs in modern browsers have resulted in the threat of IDN becoming realistic and has attracted interest from researchers [21, 46] as well as from attackers. In May 2018, Binance, which is a cryptocurrency exchange company, reported that their primary domain name binance.com was the victim of an IDN homograph attack [3]. We note that as this incident implied, the targets of IDN homograph attacks are not only browsers, but also email clients, where a victim could click a malicious URL composed of an IDN homograph.

A straightforward and effective countermeasure against the threat of an IDN homograph attack is to identify possible IDN homographs. The key technical challenge here is to automate the

process of detecting homoglyphs that could be abused for creating an IDN homograph. As of May 2019, of the 137,928 characters included in Unicode 12.0.0 [39], 123,006 characters can be used for IDN, following the specification of *IDNA2008* [8]. Furthermore, the number of IDNs registered has continued to increase. According to the IDN World Report [18], the estimated number of IDNs registered was 2.0 million in 2009 and this number increased to 7.5 million IDNs in December 2017.

In this work, we developed a generic framework named “*ShamFinder*,” which aims to identify IDN homographs in a scalable manner. The key technical contribution of *ShamFinder* is to build a new homoglyph database named *SimChar*, which can be maintained without requiring time-consuming manual effort. Unlike previous approaches for detecting IDN homographs [21, 30], the notable advantage of *ShamFinder* is that it can pinpoint the differential characters; thus, it can be used for direct countermeasures such as building a blacklist of the confusable characters or highlighting the anomalous characters to inform the user of the potential risk of an IDN homograph attack. We note that our homoglyph database covers a wide range of homoglyphs that have not been listed in the existing database maintained by the Unicode consortium [36].

Using the *ShamFinder* framework, we attempt to understand the IDN homographs registered in the wild. In our study, we investigated the way in which the registered IDN homographs are abused by collecting IDNs from the world’s most popular TLD, .com. In addition, using *ShamFinder* as a building block, we discuss a proof-of-concept system that aims to mitigate the threats posed by an IDN homograph attack.

The main contributions of this work are summarized as follows:

- We developed a framework named *ShamFinder*, which aims to identify IDN homographs in an automated manner.
- We built a new homoglyph database named *SimChar*, which can be automatically updated and can be used for other security applications such as detecting plagiarism that exploits homoglyphs.
- Using the *ShamFinder* framework, we performed a large-scale measurement study on how IDNs are used or abused in the wild. The measurement study demonstrated that our framework efficiently extracted IDN homographs, which contained malicious ones.
- Based on the *ShamFinder* framework, we propose a practical countermeasure against the generic threat of IDN homograph attacks.

The remainder of the paper is organized as follows: Section 2 presents an overview of IDN and IDN homograph attacks. In Section 3, we introduce the *ShamFinder* framework. Section 4 contains an evaluation of the performance of the *ShamFinder* framework from the viewpoints of human perception and computational costs. In Sections 5 and 6, we present our data sources and findings derived from the large-scale measurement of IDN in the wild, using the *ShamFinder* framework. Section 7 discusses the limitations of our work and effective countermeasures against the threats posed by IDN homograph attacks. In Section 8, we review related work in comparison with ours. We conclude our work in Section 9.

2 BACKGROUND

This section first presents an overview of IDNs. We then provide an overview of IDN homograph attacks and recent studies on the threats posed by these attacks.

2.1 IDN and Permitted Unicode Characters

Since the initial proposal of IDN in 1996, its protocol specification has been standardized. In 2003, the Internet Corporation for Assigned Names and Numbers (ICANN) and top IDN registries such as .cn, .info, .jp, .org, and .tw have published a guideline for the implementation/operation of IDN [27]. The guideline requires TLD registries to employ an “inclusion-based” approach, i.e., in each TLD, only code points that are permitted by the TLD can be used for IDN. Each TLD employs language-specific registration and administration rules, which are publicly available as IDN tables [28]. The tables are maintained by the Internet Assigned Number Authority (IANA).

This restriction introduced by the inclusion-based approach is expected to thwart the threats of IDN homograph attacks because the set of characters that can be used for IDN are limited with the tables. For instance, the JP domain, which is the country code top-level domain (ccTLD) for Japan, limits the permitted character sets for IDN to LDH, which consists of case-insensitive English letters, digits, and hyphens (Letter-Digit-Hypen), Hiragana, Katakana, and a subset of CJK unified ideographs (character set used in Chinese, Japanese, and Korean). Therefore, it is not possible to register Latin-based IDN homographs with names such as “ác.jp” because the permitted characters for IDN of the JP domain do not contain a homoglyph of LDH.

However, as we shall present later, among the characters permitted for each TLD such as .com, there are many homoglyphs, indicating that an attacker can leverage such homoglyphs to execute an IDN homograph attack. We note that an attacker can create an IDN homograph of a non-Latin IDN homograph. One of the key contributions of our work is to automatically build a comprehensive list of homoglyphs, which could be potentially abused for IDN homograph attack.

Although the IDN extension allows us to use non-Latin characters for domain names, we need to use LDH at the protocol level for backward compatibility reasons. Therefore, we need a mechanism that transcodes a domain name consisting of Unicode characters into one with LDH characters. In this regard, Punycode is a character encoding scheme for transcoding a Unicode string to a string with LDH. The specification of Punycode is defined in RFC 3492 [6]. When using a string transcoded by Punycode for IDN, we add the prefix “xn--” to the beginning of the transcoded string. For instance, the string “阿里巴巴” is represented as “tsta8290bfzd” by the Punycode transcoding, and the corresponding IDN is “xn--tsta8290bfzd.com”.

Finally, we note that each web browser implements the processing of IDN in a different way [23, 35]. As we explain below, the way IDN is displayed in the address bar could increase or decrease the threat of an IDN homograph attack. Thus, the implementation largely affects the way users react to the IDN homograph presented in a browser. In Section 7, we discuss a proof-of-concept implementation of IDN processing on a browser to enable users to become

knowledgeable of the existence of a possible IDN homograph attack without sacrificing the usability of IDN for them.

2.2 IDN Homograph Attack

As mentioned in the previous section, the history of IDN homograph attacks can be traced back to the early 2000s. As Gabrilovich and Gontmakher [12] reported in 2002, numerous English domain names can be *homographed* by leveraging non-Latin letters.

Despite the fact that threats of IDN homograph attacks were pointed out earlier, effective and usable countermeasures against these threats have not been developed. We conjecture that the reason behind abandoning the threats is that IDN has not been widely deployed in the world and there have been few web clients that can correctly process IDN. However, the situation has changed because popular web browsers today have developed the ability to handle IDNs. In addition, according to the IDN World Report [18], 7.5 million of IDNs have been registered by December 2017. These observations imply that the threat of IDN homograph attacks have become real. In fact, as mentioned in Section 1, the cryptocurrency exchange company Binance was the victim of an IDN homograph attack.

As countermeasures against IDN homograph attacks, many browser vendors have updated the implementation of displaying IDN in the address bar after the threat of an IDN homolog attack was widely publicized by a blog post on the web [46] in April 2017. Specifically, Firefox and Chrome have changed their implementations as follows: when characters originating from multiple scripts (character sets) are mixed in a character string constituting an IDN, the IDN is displayed in the form of Punycode instead of Unicode [23, 35]. For instance, if a Latin-script-based domain name comprises non-English scripts such as Latin scripts, Cyrillic scripts, or Greek scripts, the domain name is displayed in the form of Punycode; i.e., for “facé-book”, its Punycode, xn--facbook-dya is displayed in the address bar.

Although this update can be expected to mitigate the threats of IDN homographs to some extent, it is likely to impair the usability because Punycode is not a human-friendly representation. As the human-readable domain name provides hints as to the authenticity of the website, masking the original domain name may leave users less knowledgeable. Although the aforementioned countermeasure by the browsers becomes a temporary countermeasure against IDN homograph attacks, if it is compulsorily displayed in Punycode, it is problematic in that it becomes difficult to understand the cause of the threat. That is, because the user does not notice that the domain name entered in the browser is a homograph attack, the user risks visiting the site with the same domain name again.

We also note that, in the above implementations, even in the case of an IDN composed of multiple scripts, if the domain name comprises both Latin script and a CJK ideograph, it will be displayed with Unicode. Furthermore, an attacker can create an IDN by not only combining Latin script, Cyrillic script, or Greek script but also by combining characters belonging to the set of CJK ideographs. We refer to such a homograph as a non-Latin homograph. For instance, the string “工業大学” (meaning an institute of technology in English) has the homograph, “工業大学”, where ‘工’ is a CJK Unified Ideograph (U+5DE5) and ‘厶’ is a Katakana Letter (U+30A8).

Figure 1: Overview of the ShamFinder Framework.

Current web browsers do not have a way to identify non-Latin IDN homographs such as this.

3 SHAMFINDER FRAMEWORK

In this section, we first provide a high-level overview of the *ShamFinder* framework. Next, we present several Unicode character sets and those used for IDN. We note that precise understanding of these character sets is essential in extracting Unicode homoglyphs that could be abused for an IDN homograph attack. We then describe the approach we followed to build the homoglyph database, which plays a key role in the *ShamFinder* framework. Finally, we describe the characteristics of the homoglyph database.

3.1 High-level Overview

Figure 1 presents a high-level overview of the *ShamFinder* framework.

Step 1: First, we collect registered/active domain names for each TLD. To this end, we can either make use of the DNS zone file for each TLD or publicly available/commercial domain name lists such as [7]. We introduce the datasets we used for our analysis in Section 5.

Step 2: Next, we extract IDNs from the collected domain names by searching for those starting with the prefix “xn--”.

Step 3: To find IDN homographs, we leverage a list of popular domain names as reference. As representative reference we can leverage a website ranking lists [29, 31] such as Alexa Top Sites [2] or Majestic Million [22]. Next, we leverage the database of homoglyphs to identify potential IDN homographs; as we show in the next subsection, our contribution is to present a way of automatically building such a database.

Figure 2 and Algorithm 1 show the IDN detection scheme. We check the length (number of characters) of each domain name listed in the reference domain names list and extract the IDNs with the same number of characters. For each pair consisting of a reference domain name and sampled IDN, we check their letters one by one to determine whether they correspond. If two corresponding letters match each other, we proceed to the next pair of letters. If the letters do not match, we check whether the pair is listed in the homoglyph database, which we present in the next subsection. If they are listed, we proceed to the next pair of letters and repeat the same process. If we find letters that do not match, we conclude that the IDN is not an IDN homograph of the reference domain name. The computational complexity of the algorithm is $|N||M||L|$ where $|N|$, $|M|$, and L are

Figure 2: IDN homograph detection scheme for a given TLD. Left: a domain name is detected as an IDN homograph. Right: a domain name is not detected as an IDN homograph. In this example, the TLD part has been removed.

Algorithm 1 Homograph detection algorithm

```

 $\Omega \leftarrow$  A set of IDNs for a given TLD (TLD part removed)
 $R \leftarrow$  A set of reference domain names for a given TLD (TLD part removed)
 $\triangleright \text{len}(x)$  returns the length of a string  $x$ .
 $\triangleright x[i]$  represents the  $i$ -th character of string  $x$ .
for  $r \in R$  do
  Let  $\omega \subset \Omega$  be a set s.t.  $\text{len}(r) = \text{len}(x) \forall x \in \omega$ 
  for  $x \in \omega$  do
    for  $i = 1; i \leq \text{len}(x); i \rightarrow i + 1$  do
      if  $r[i] = x[i]$  then next
      else if  $r[i]$  and  $x[i]$  are listed as a pair in the SimChar database then
        next
      else
        Mark  $x$  as not being a homograph of  $r$  and skip the loops for the
        next  $x$ 
    Mark  $x$  as a homograph of  $r$ 

```

the number of reference websites, number of IDNs, and number of characters contained in a domain name, respectively. Although this is a naïve approach, the actual calculation cost has been reduced by restricting the computation of matching to those pairs of strings with the same length. The evaluation of the time needed for the computation appears in Section 4.

3.2 Unicode Characters Sets and IDN

Our primary goal is to compile a database that lists pairs of visually identical Unicode characters (homoglyphs) that are permitted to be used for IDN. We explain how we compile this database by beginning with a description of several Unicode character sets. Figure 3 summarizes the contamination and overlap of the Unicode character sets. The root set is the characters contained in Unicode 12.0.0 [39]. The set contains a total of 137,928 characters, covering 150 scripts, including modern/historic characters, signs, and symbols such as Emoticons. Of the character sets defined in Unicode 12.0, the latest set of characters permitted for the use in IDN is defined in the Internet draft, named draft-faltstrom-unicode12-00 (“IDNA2008 and Unicode 12.0.0”) [8]. The number of Unicode characters contained in the IDNA2008 draft is 123,006; these characters are listed in the section, “Code points in Unicode Character Database (UCD),” of the draft with the property of “PROTOCOL VALID (PVALID),” which indicates that the code points with the property value are permitted for general use in IDNs [10].

In the document named Unicode Technical Standard #39 (UNICODE SECURITY MECHANISMS), a database named “confusables.txt”

Figure 3: Contamination and overlap of character sets. UC denotes the recommended confusable mapping for IDN, “confusables.txt.”

Table 1: Number of characters contained in each set shown in Figure 3; IDNA2008 draft (IDNA), Unicode confusables.txt (UC), and SimChar.

Sets	# characters	# homoglyph pairs
IDNA	123,006	n/a
UC	9,605	6,296
$UC \cap IDNA$	980	627
<i>SimChar</i>	12,686	13,208
$SimChar \cap UC$	233	127
$SimChar \cup (UC \cap IDNA)$	13,210	13,708

is provided. This text file compiles the confusable mapping for IDN. In this work, we refer to the database as UC for brevity. The UC database lists visually *confusable* characters and provides a mapping for visual confusables for use in detecting security problems such as an IDN homograph attack. Although UC covers a wide range of homoglyphs that could be abused for IDN homograph attacks, our empirical observations revealed that a non-negligible number of homoglyphs are *not* contained in UC as shown in Table 1. This observation motivated us to build a new homoglyph database, *SimChar*, which is described in the next subsection. We note that although UC has been manually maintained, we can build *SimChar* in an automated way, implying that it can discover new homoglyphs from newly registered Unicode characters in future. Furthermore, as explained in Section 4, homoglyphs contained in *SimChar* are more confusing than those contained in UC.

We note that UC covers several characters that are *not* contained in the IDNA2008 draft. Of the characters defined in the IDNA2008 draft, 980 characters are listed in UC; i.e., these 980 characters are potentially abused for IDN homograph attacks. Our contribution is to build a complementary database named *SimChar*, which is compiled of a set of characters that have at least one homoglyph character from the IDNA2008 draft character set. The new set has 13,210 characters that are also included in UC. It adds 3,605 characters that have not been listed in UC. Moreover, as seen in Table 3, *SimChar* adds 316 homoglyphs of Basic Latin characters that are *not* listed in UC. Table 1 summarizes the number of characters contained in the character sets shown in Figure 3. We note that the *ShamFinder* framework makes use of the union of two sets UC and

Figure 4: Contamination and overlap of character sets.

SimChar to find IDN homographs. We also note that a character could be the homoglyph of several other characters. We count such pairs as “Homoglyph pairs.” Homoglyphs contained in *SimChar* are built from a set of characters contained in IDNA. We notice that although the number of characters contained in UC is roughly 10K, if we consider the number of IDNA-permitted characters, the size becomes much smaller by a factor of 10. The details of *SimChar* will be shown later.

3.3 Building Homoglyph Database

As shown in Fig. 2, we use UC and *SimChar* as the components of the homoglyph database we used to detect IDN homographs. The key idea of *SimChar* is to extract homoglyphs by computing the similarity between the glyphs of corresponding characters. We first need to represent each code point as a visual image (glyph). To this end, we can make use of various Unicode fonts such as those listed in [38]. In this work, we adopt GNU Unifont Glyphs [40], which covers the entire collection of characters contained in the Unicode Basic Multilingual Plane (BMP) as well as several other characters of the Supplemental Multilingual Plane (SMP). Although BMP contains characters for almost all modern languages and a large number of symbols, SMP contains historic characters and signs as well as the symbols used in various fields such as Emoticons. Even though the choice of a font may affect the detected homoglyphs, the following procedure can easily be extended to other font sets. We aim to evaluate other fonts in future work.

Figure 4 presents the relationship between the character sets. Of the characters contained in the IDNA2008 draft, the latest version of Unifont (Unifont12 for short) covers 52,457 characters. Several IDN-permitted characters are *not* covered by Unifont12. However, as Unifont provides much larger coverage than other proprietary Unicode fonts such as Microsoft JhengHei, we deem the choice to be reasonable. In fact, of the 2,990 IDN-permitted characters in UC, 2,877 characters are covered by Unifont12. Table 2 summarizes the number of characters contained in the character sets shown in Figure 4. In the following, we denote UC and *SimChar* as those with the union sets of Unifont12 for brevity.

Next, we attempt to identify homoglyphs by testing their similarity as images. The structural similarity index measure (SSIM) is a widely used metric to quantify the degradation of image quality caused by processing methods such as data compression or by losses

¹Note that *SimChar* is composed using the union set of IDNA and Unifont12. Therefore, $|SimChar| = |SimChar \cap Unifont12|$.

Table 2: Number of characters contained in the Unifont12 font set (Figure 4).

Sets	# Chars	# Pairs
$IDNA \cap Unifont12$	52,457	n/a
$UC \cap Unifont12$	5,080	3,696
$SimChar \cap Unifont12^1$	12,686	13,208

Figure 5: Examples of Unifont glyph images of Unicode characters. Top: Basic Latin characters and their homoglyphs. Bottom: CJK Unified Ideographs, Hangul syllables, and Oriya (Indo-Aryan language spoken in the Indian state of Odisha), and their homoglyphs.

in data transmission [16, 45]. Thus, it can also quantify the similarity between a pair of images. However, because our objective is not assessing the perceptual metric that quantifies image quality degradation, we directly count the number of different pixels between two images. Let $I(i, j) \in \{0, 1\}$ ($0 \leq i \leq N-1$, $0 \leq j \leq N-1$) be a square image having $N \times N$ pixels, where each pixel is represented as a binary digit. Our metric, Δ is computed as

$$\Delta = \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} |I_1(i, j) - I_2(i, j)|.$$

When $\Delta = 0$, it indicates that two images are completely identical.

We note that Δ can be associated with the peak signal-to-noise ratio (PSNR), which is another widely used metric aimed at quantifying the reproducibility of images [16, 45]. In our model, $I(i, j)$ is represented as a binary bit. Therefore, the mean square error (MSE) is computed as

$$MSE = \frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \{I_1(i, j) - I_2(i, j)\}^2 = \frac{\Delta}{N^2}.$$

Using the MSE, the PSNR is computed as

$$PSNR = 10 \log_{10} \left(\frac{1}{MSE} \right) = 20 \log_{10} N - 10 \log_{10} \Delta.$$

In the following, we show the processes we employed to construct the *SimChar* database.

Step I For the 52,457 characters in the intersection of the IDNA2008 draft and Unifont12, we represent the characters as bitmap images of 32×32 pixels, using the Unifont glyphs. Note that the original size of Unifont11 is 8×8 pixels for Latin characters and 16×16 for other characters. Figure 5 presents the example of the generated Unifont glyph images where we intentionally chose visually similar pairs.

Step II For all the pairs in the pairwise combinations of the 52,457 characters, we compute the metric Δ . If Δ is less than

Table 3: Number of homoglyphs of Latin letters (lowercase) contained in *SimChar* and $UC \cap IDNA$.

<i>SimChar</i>						$UC \cap IDNA$					
	#		#		#		#		#		#
‘o’	40	‘s’	14	‘f’	8	‘o’	34	‘c’	4	‘p’	3
‘e’	26	‘r’	14	‘m’	8	‘l’	12	‘d’	4	‘x’	3
‘n’	24	‘a’	14	‘g’	7	‘y’	10	‘g’	4	‘j’	2
‘w’	20	‘k’	13	‘j’	7	‘i’	9	‘f’	4	‘n’	2
‘c’	19	‘t’	13	‘p’	7	‘u’	9	‘a’	3	‘z’	2
‘l’	18	‘z’	12	‘x’	6	‘w’	8	‘b’	3		
‘u’	18	‘d’	10	‘q’	2	‘v’	6	‘e’	3		
‘h’	17	‘y’	9	‘v’	1	‘s’	5	‘h’	3		
‘i’	16	‘b’	8			‘r’	5	‘q’	3		
Total		351				Total		141			

or equal to a threshold θ , the two characters are identified as homoglyphs. In this work, we empirically derived a conservative threshold as $\theta = 4$; i.e., a pair of characters are detected as homoglyphs if $\Delta \leq 4$. Figure 6 shows examples of Unicode characters with various values of Δ . Although $\Delta = 4$ would not indicate obvious *false positives* (i.e., those that should not be detected as homoglyphs), we can observe several *false negatives* (i.e., those that could be detected as homoglyphs) among characters with $\Delta \geq 5$. In Section 4, we consider an evaluation of the validity of the threshold by presenting a human study.

Step III Finally, from the extracted pairs, we eliminate *sparse* characters that contain fewer than 10 black pixels. The threshold was empirically derived as a result of careful manual effort. In most cases, these characters are used for punctuation, spacing/nonspacing, or combining in various languages. Figure 7 presents examples of the eliminated characters.

After performing the four processes described above, we obtained a set of 12,636 characters. The set constitutes 13,126 pairs, which we named *SimChar*. As shown in Table 1, the size of the intersection of *SimChar* and UC is fairly small, indicating that *SimChar* successfully adds new homoglyphs that have not been covered by UC. We also note that there are several characters that are not covered by *SimChar*, but are covered by UC. Thus, the two character sets can be used complementary to identify potential IDN homograph attacks.

3.4 Characteristics of *SimChar*

Homoglyphs of Latin Letters As the majority of popular websites make use of the 26 Latin letters to construct their primary domain names, it is essential to study the extent to which our homoglyph database covers the homoglyphs of Latin letters. Table 3 lists the results. We first notice that *SimChar* successfully extracted new homoglyphs that have not been contained in UC. For instance, whereas the intersection of IDNA2008 and UC contains only three homoglyphs for the Basic Latin lowercase letter ‘e’, *SimChar* contains 26 homoglyphs of ‘e’ as shown in Figure 6. We also notice that several characters have many homoglyphs. In total, *SimChar* contains 351 homoglyphs of Latin letters, whereas UC contains 141 of these homoglyphs. In the *SimChar* dataset, the Basic Latin

Table 4: Top-5 Unicode blocks contained in *SimChar* (left) and $UC \cap IDNA$ (right). CJK, CDM, Hangul, and CA are abbreviations of the CJK Unified Ideographs, Combining Diacritical Marks, Hangul Syllables, and Unified Canadian Aboriginal Syllabics, respectively.

<i>SimChar</i>		$UC \cap IDNA$	
Block	#chars	Block	#chars
Hangul	8,787	CJK	91
CJK	395	CDM	56
CA	387	Arabic	52
Vai	134	Cyrillic	40
Arabic	107	Thai	36

lowercase letter ‘o’ has 40 characters that are visually similar to it, indicating that the character is “vulnerable” to an IDN homograph attack. We note that the intersection of the sets of homoglyphs for ‘o’ for *SimChar* and UC contains 5 characters, implying that they cover different sets of homoglyphs of ‘o’; i.e., the majority of homoglyphs of ‘o’ listed in *SimChar* were accented characters of ‘o’, whereas the majority of homoglyphs of ‘o’ listed in UC were characters of which the appearance resembles a circle.

Unicode Blocks In Unicode, a block is a contiguous range of code points. A block consists of hundreds to tens of thousands of characters. The characters contained in a block are typically associated with the writing systems in which the characters are used; e.g., the Basic Latin block consists of all the characters and control codes of the ASCII character set. The majority of the blocks are classified into two planes: the Basic Multilingual Plane (BMP) and Supplementary Multilingual Plane (SMP). In the BMP, the largest block is the CJK Unified Ideograph, the characters of which are used in the Chinese, Japanese, and Korean languages, and it contains more than 20 K of Chinese characters.

Table 4 compares UC and *SimChar* with respect to their top-5 blocks. Although the two scripts, CJK Unified Ideographs and Arabic are commonly found, the breakdown of these scripts differ from each other, indicating that the coverage of UC and *SimChar* is different. Our contribution is to automatically build *SimChar*, which can complement the manually compiled list of homoglyphs, i.e., UC. We note that the .com TLD is allowed to use characters from either of these blocks for IDN.

4 PERFORMANCE EVALUATION

This section presents our evaluation of the performance of the *ShamFinder* framework from the viewpoints of (1) human perception and (2) computational cost.

4.1 Human Perception

We evaluated the human perception of the homoglyphs listed in our *SimChar* database; i.e., to determine whether humans perceive their homoglyphs as confusing. To this end, we employed a series of human study experiments using a crowd sourcing platform, Amazon Mechanical Turk (MTurk in short). We designed two types of experiments. In our first experiment, we studied the effect of the threshold Δ , which was introduced in Section 3.2, on the extent to

Figure 6: Basic Latin lowercase letter ‘e’ and characters under different values of the threshold Δ . In this work, characters with $\Delta \leq 4$ are detected as homographs.

Figure 7: Examples of sparse characters.

Figure 8: Screenshot of the MTurk task.

which *SimChar* homographs could be confused, i.e., their “confusability.” This experiment is intended to demonstrate the validity of the threshold we determined for detecting homographs, i.e., $\Delta = 4$. Next, we compare the confusability of *SimChar* and UC, with the baseline of random pairs of characters.

Experimental Setup We created a crowd sourcing task that asks a participant whether pairs of two characters, which may contain homographs, are confusing or distinct. Before performing the large-scale experiment, we carefully designed our experiment by conducting a series of pilot study trials that enabled us to adjust the wording of questions and answers. Several trials of the pilot study allowed us to obtain useful feedback from coworkers and participants, and we ultimately worded the question as “*There are two characters shown in the image. Are they distinct or confusing?*” In terms of the answer, the following words were selected as the options for the five-level Likert scale score, “1: *very distinct*,” “2: *distinct*,” “3: *neutral*,” “4: *confusing*,” and “5: *very confusing*.” In this work, we refer to this score as the “confusability score.”

Figure 8 presents a screenshot of an assignment in the task presented to participants. The purpose of the assignment was to judge whether two characters contained in an image are distinct or confusing. Before conducting crowd sourcing experiments, we measured the average time to finish an assignment by ourselves and found an assignment to require approximately 15 seconds to

Figure 9: Boxplot of the confusability scores for pairs of *SimChar* homographs for different values of the threshold, Δ . The notch in each box represents the median and the dashed lines represent the mean values. Whiskers represent 1.5 IQR.

complete, including the time to select an answer, submit it via the web interface, and wait for the page transition to the next assignment. On the basis of this observation, we set the reward per assignment as 0.05 USD, implying that the reward is equivalent to an average hourly compensation of 12 USD. As the minimum wage in the USA is in the range of 7–12 USD / hour [41] (as of March 2019), we believe our payment configuration was appropriate, i.e., it was neither too low nor too high.

To ensure the quality of experiments, we used the following two criteria when recruiting participants: (1) the number of approved tasks of a participant should exceed 50 and (2) the participant should have a task approval rate greater than 97%.

To check whether a participant was careful when completing the task, we inserted dummy images that contain two completely distinct random characters. A participant who judged a dummy image as being either “4: *confusing*” or “5: *very confusing*” had all their responses removed, assuming that the reliability of the participant was low. We likewise removed all the responses from participants who answered “1: *very distinct*” or “2: *distinct*” to a homograph contained in *SimChar* with the threshold of $\Delta = 0$, i.e.,

Figure 10: Boxplot of the confusability score for the pairs of three sets: Random (left), SimChar (middle), and UC (right). The threshold of SimChar was set to $\Delta \leq 4$. The boxplot configuration is the same as in Figure 10.

when the glyphs of the two characters were perfectly identical with the font we used (GNU Unifont). Although this strategy may have aggressively removed the useful responses by a participant who accidentally made a single mistake, we decided to overcome the drawback by simply increasing the number of responses/samples. **Experiment 1: Threshold of SimChar** We first studied the way in which the threshold, Δ , affects human perception. In this experiment, we used homoglyphs of the Basic Latin letters (lowercase), the numbers of which are listed in Table 3. For each letter, we extracted the glyphs with a distance of $\Delta \in \{0, \dots, 8\}$. For each Δ , we randomly sampled 20 pairs, where a pair consists of a letter and its potential homoglyph detected with the threshold Δ . In addition, we added 30 of dummy pairs that contain two distinct letters randomly generated. These $20 \times 9 = 180$ pairs of potential homoglyphs and 30 random pairs were judged by 10 participants (after the removal of unreliable participants). In total, we obtained 900 effective responses for the 180 pairs.

Figure 9 presents the result. As expected, the confusability score decreases as the threshold increases. When $\Delta = 4$, the mean and median of the confusability score were 3.57 and 4, respectively. This observation implies that the homoglyphs detected with the threshold were mostly perceived as “confusing.” When $\Delta = 5$, the mean and median of the confusability score were 2.57 and 2, respectively, implying that the detected homoglyphs were mostly perceived as “distinct.” On the basis of these observations, we adopted $\Delta = 4$ as the threshold for extracting homoglyphs; i.e., glyphs with $\Delta \leq 4$ were detected as homoglyphs. Although several pairs with the threshold of $\Delta = 5$ had a high confusability score, we adopted a conservative decision. Extracting further confusable homoglyphs from these potential homoglyphs remains as a future task.

Experiment 2: Confusability of UC and SimChar. Next, we studied the confusability of UC in comparison with *SimChar* for which we repeated the same procedure shown above. We sampled 30 of the homoglyphs of the Basic Latin letters (lowercase) listed

Figure 11: Three pairs of homoglyphs listed in UC. They were most frequently judged as “very distinct” by the participants. They are homoglyphs of ‘u’ (left), ‘y’ (middle), and ‘y’ (right), respectively.

in UC. These 30 pairs were judged by 28 participants (after the removal of unreliable participants). In total, we obtained 513 effective responses for the 30 pairs sampled from UC. For *SimChar*, we compiled 486 effective responses for the $20 \times 5 = 100$ pairs of homoglyphs detected with $\Delta \leq 4$.

Figure 10 shows the result. For comparison, 513 of the effective responses for the 30 dummy pairs (Random) are also plotted. Although the confusable scores of the random pairs were mostly concentrated near the lowest option (“very distinct”), for both *SimChar* and UC, the median of the confusable score was 4, i.e., the homoglyphs of both databases were perceived as “confusing” on average. Note that the average confusable score for *SimChar* was larger than 4, whereas that for UC was smaller than 4, implying that the homoglyphs contained in *SimChar* were more confusable than those contained in UC.

Figure 11 presents three examples of UC pairs that attracted the lowest confusability score. As these examples imply, several homoglyphs listed in UC have glyphs that could be perceived as distinct from the original letter, although some of the pairs could be semantically close. On the other hand, the homoglyphs listed in *SimChar* should have small differences by definition. These results led us to conclude that the homoglyphs listed in *SimChar* are actually perceived as confusable.

4.2 Computation Cost of the ShamFinder Framework

We first measured the time taken for constructing *SimChar*. Table 5 summarizes the results. As expected, the time for computing Δ for the pairwise combination of 52,457 characters, which is provided in Table 2, was the most time-consuming step of the computation. For this computation, we used a multi-processing approach with the number of concurrent processes set to 15. We used an off-the-shelf server with an Intel Xeon CPU E5-2620 v2 (2.10 GHz) and 62 GB memory. In practice, we would need to update *SimChar* when the Unicode standard adds a new set of glyphs or we incorporate a new set of fonts to be analyzed. That is, the frequency of updating *SimChar* should be reasonably low; e.g., Unicode version 12.0 was released one year after the release of version 11.0. The new version added 553 characters to those in the previous one.

Next, we measured the time to extract IDN homographs using the *ShamFinder* framework. To extract IDN homographs of the Alexa top-10k domains from the 141 M of .com TLD domain names (see Table 6 for reference) required 743.6 seconds, i.e., on average, each reference domain name was inspected in $0.07 (= 743.6 / 10,000)$ seconds, which is sufficiently fast to block a suspicious, newly found IDN homograph attack in real time.

Table 5: Time taken for constructing SimChar.

Process	Time
Generating images	79.2 seconds
Computing Δ for all the pairs	10.9 hours
Eliminating sparse characters	18.0 seconds

5 DATA SOURCES

In this section, we describe the data sources used for our analysis of IDN homographs.

5.1 Reference Domain Names

The aim of an IDN homograph attack is to attract a victim to a malicious website by using a homograph that is visually identical to the domain name of a legitimate website. As such, the natural assumption is that an attacker creates an IDN homograph of a domain name used for a popular website. In fact, other deception techniques such as “typosquatting” or “brandjacking” also target widely recognized domain names [1, 34]. As a reference of well-known popular domain names, we adopted Alexa Top Sites [2]; namely, we extracted the top-10K of .com domains from the Alexa ranking list.

5.2 Extracting IDNs

Although many domain name spaces are available in the Internet, in this study, we focused on domain names under the .com TLD for the following three reasons. First, the majority of popular websites are attributed to this TLD. As the word “dot-com bubble” symbolizes, .com has become the most popular TLD since the early 2000s. Although .com was originally intended for commercial usage, it eventually became available for general purposes. Second, as shown below, the majority of malicious IDNs are also attributed to this TLD. Finally, as .com TLD is globally popular, it permits a large number of Unicode blocks to be used for IDNs. According to IANA’s IDN tables [28], under the .com TLD, characters across 97 different Unicode blocks can be used for IDNs as of May 2019. This fact implies that for .com TLD, an attacker can register an IDN homograph that contains homoglyphs sampled from various Unicode blocks.

To search for IDN homographs, we first needed to extract registered IDNs. To this end, we used the DNS zone file maintained by the registries of the .com TLD — Verisign [42]. The DNS zone file lists all the registered domain names with their NS records. We complemented the zone file by using another list of domain names named domainlists.io [7]. The union set of the two lists contains 141.2 M of unique domain names. As mentioned above, we can extract IDNs by searching for domain names starting with the prefix “xn--”.

Table 6 summarizes the number of domain names/IDNs for each dataset. We first notice that a non-negligible number of IDNs are currently registered in the .com TLD, implying the widespread adoption of IDN in the wild. Next, we examined the languages used in those IDNs to understand which Unicode blocks are widely used in the IDNs. To detect the language used in a string, we leveraged a tool known as LangID [20], which is a Python module that can detect the most plausible language out of 97 distinct languages for

Table 6: Summary of domain name lists and the number of IDNs they contained.

Data	Number of domain names	Number of IDNs	Collection time
zone file [42]	140,900,279	952,352 (0.67%)	May 2019
domainlists.io [7]	139,667,014	953,209 (0.73%)	May 2019
Total (union)	141,212,035	955,512 (0.67%)	–

Table 7: Top languages used for IDNs.

Rank	Language	Number	Fraction (%)
1	Chinese	443,865	46.5
2	Korean	101,711	10.6
3	Japanese	88,970	9.3
4	Germany	53,776	5.6
5	Turkish	35,288	3.6

Table 8: Number of detected IDN homographs for ASCII domains.

Homoglyph DB	Number
UC	436
<i>SimChar</i>	3,110
$UC \cup SimChar$	3,280

a given string. Table 7 presents the results. We see that east Asian languages (Chinese, Japanese, and Korean) are dominantly used for composing IDNs whereas several European languages are also popular for this purpose. This observation implies that the demand for the use of native languages is ubiquitous.

6 DETECTING IDN HOMOGRAPHS WITH THE SHAMFINDER FRAMEWORK.

In this section, we apply the *ShamFinder* framework to the data we described in the previous sections. We first studied the IDN homographs that targeted popular domain names that reside in the .com TLD. We then studied the malicious IDN homographs detected by our approach. In comparison to the existing approach, we compared the number of detected malicious IDN homographs by changing the homoglyph database. As discussed in Section 8, the previous approach to detecting IDN homographs proposed by Quinkert et al. [26] leveraged UC as their homoglyph database. That is, we can directly compare the IDN homograph detection performance between their approach (UC only) and ours (UC and *SimChar*).

6.1 Statistics of the IDN Homographs

Table 8 presents the number of detected IDN homographs targeting ASCII-character domain names. When we used UC, the *ShamFinder* framework detected 436 IDN homographs out of the 955 K IDNs registered in the .com TLD. On the other hand, when we used *SimChar*, more than 3,110 of IDN homographs were detected. In total, by using both homoglyph databases, we detected 3,280 IDN

Table 9: Top-5 ASCII domain names that have the most IDN homographs.

Rank	Domain name	# homographs
1	myetherwallet.com	170
2	google.com	114
3	amazon.com	75
4	facebook.com	72
5	allstate.com	68

Table 10: Port scan results for the detected IDN homographs.

Ports	# domain names
TCP/80	1,642
TCP/443	700
TCP/80 & TCP/443	695
Total (unique)	1,647

homographs, which is approximately eight times more than those detected with UC. Thus, the adoption of *SimChar* as the homograph database enables us to detect more IDN homographs than existing approaches such as that of Quinkert [26].

Table 9 presents the the top-5 domain names that have the most IDN homographs. Three of these domains, google.com, amazon.com, and facebook.com are all popular domains; however, the two other domains, myetherwallet.com and allstate.com are not that popular compared to the other three domains. In fact, the first three domains are ranked among the top-10 domains in the Alexa ranking, whereas the other two domains are ranked 7,400th and 5,148th among the .com TLD domains in the Alexa ranking, respectively. This observation demonstrates that IDN homograph attacks not only target very popular websites, they also target other moderately popular websites, implying that starting with a small list of reference domains may not be effective for IDN homographs that target minor domains. We discuss this issue below (Section 6.4).

In the following, we analyze the IDNs that are currently active. First, we checked the NS records for the 3,280 homograph IDNs we detected. We found 2,294 domain names with NS records, while other domain names did not have NS records due to some reasons such as expiration, non-registration, etc. Of the 2,294 domain names, 385 domain names did not have A records. For the remaining 1,909 domain names, we performed port scans to the ports TCP/80 and TCP/443. Table 10 shows the results. We found that the 1,647 IDN homographs we detected were reachable through the HTTP or HTTPS; i.e., roughly half of the detected IDN homographs were active.

6.2 Deep Inspection of the Active IDN Homographs

In this section, we further inspect the characteristics of the active IDN homographs we found in the previous subsection. In the following, we show the analyses from two aspects: (1) analysis of the popular IDN homographs and (2) classification of IDN homographs.

(1) Analysis of the popular IDN homographs

To study how the active IDN homographs have been accessed by end users, we focus on the “popular” IDN homographs that likely attracted large number of end-users. To this end, we performed the analysis using the passive DNS [9], which is a DNS monitoring system that is composed of several working DNS cache servers. A passive DNS system provides useful statistics such as the number of cumulative name resolutions for each domain name. We note that the statistics provided by a passive DNS system reflects sampled data collected at a set of cache servers contributing to the system. Therefore, the actual numbers of DNS lookups over the entire Internet should be much larger than those obtained from a passive DNS system. We also note that the number of web accesses and number of DNS resolutions are different. However, we believe that the number of DNS resolutions is correlated with the popularity of a domain name, given that every first web query should be preceded by a DNS query.

Table 11 shows the top-10 domain names that had the largest numbers of DNS lookups. We studied the categories of the websites running on the IDNs by manual inspection. We found that of the top-10 IDNs, four of them targeted gmail.com. In particular, the top IDN, gmail[.]com was an active phishing site and had a large number of name resolutions, implying that there have been a large number of end-users who accessed the phishing website². We found that the website under the IDN employed a cloaking technique to redirect a visitor to the different websites according to the User-agent of the visitor’s browser. We also found that the majority of the IDNs were parked domains; these were used for monetizing through advertisements and/or were reserved for resale.

In Table 11, the columns “MX,” “Web link,” and “SNS” represent, whether there was a generic website linking to the IDN homograph, and whether there was a web link pointing to the IDN homograph on popular SNS websites such as Twitter. We used the search engines for the latter two analyses. We found that the IDN homographs that target domain names used for email services such as gmail.com and yahoo.com have MX records either in the past or in the present. We also saw that several IDN homographs have appeared in public web space, including SNS. These observations imply that the owners of these IDN homographs have attempted to make the IDN homographs publicly visible.

(2) Classification of IDN homographs

We now attempt to classify the 1,647 active IDN homographs that responded to either TCP/80 or TCP/443. To this end, we make use of a list of NS records for the domain parking companies, screenshots of the websites, and VirusTotal [43], which is an online virus scanner. To compile a list NS records for the domain parking companies, we leverage the list and methods proposed in [4, 44]. We added several NS records and ended up 17 of NS records used for domain parking.

Next, for the remaining IDN homographs that were not attributed to domain parking, we accessed to the corresponding websites via the two schemes, HTTP and HTTPS, and took screenshots using the puppeteer [25], which is a headless browser that provides APIs to control Chrome or Chromium. Based on the characteristics

²As of September 2019, this website was still in operation. We have reported about the website to the security vendors.

Table 11: Top-10 active IDNs that had the largest numbers of DNS resolutions. ● indicates that there is an active MX record registered. ○ indicates that there was a MX record in the past.

Domain name	Category	#resolutions	MX	Web link	SNS
gmail[.]com	Phishing	615,447	○		
döviz[.]com	Portal	127,417		✓	✓
gmail[.]com	Parked	74,699	○		
gmàil[.]com	Parked	63,233	●		
expansi3n[.]com	Parked	56,918	○	✓	✓
gmàil[.]com	Parked	49,248	●		
yàhoo[.]com	Parked	44,368	○		
shàdbase[.]com	Parked	38,556		✓	✓
youtubè[.]com	Sale	37,713	●		
perú[.]com	Parked	36,405	●		✓

Table 12: Classification of the active IDN homographs.

Category	Numbers
Domain parking	348
For sale	345
Redirect	338
Normal	281
Empty	222
Error	113
Total	1,647

of screenshots and HTTP responses, we classified the websites into the following five categories: “For sale,” “Redirect,” “Normal,” “Empty,” and “Error,” which represent a website that encourages you to buy the domain, a website that redirects to another website, a website that displays something legitimate successfully, a website that displays nothing, and a website that failed to get a screenshot due to a timeout or other reasons, respectively.

Table 12 shows the results. We found that 693 (42%) of the websites running on IDN homographs were used for business (“Domain parking” or “For sale”). We also found that 338 (21%) of the websites running on IDN homographs were redirected to other websites having different domain names. We further analyzed these 338 websites using VirusTotal and manual inspection of the screenshots. Table 13 shows the breakdown of the websites with redirect. Brand protection indicates that a website running on a homograph domain name is redirected to the website running on the corresponding original domain name. That is, the owner of the original domain name has registered the homograph to protect their brand. We found that while the majority of the redirected domain names were attributed to either brand protection or legitimate websites, we found 35 of them were detected as malicious websites.

6.3 Malicious IDN Homographs

To check whether the detected IDN homographs have been used for malicious purposes, we leveraged three different sources of blacklists, hpHosts [17], Google Safe Browsing (GSB) [14], and Symantec DeepSight [33]; of the three lists, hpHosts, which is a

Table 13: Classification of the active IDN homographs.

Category	Numbers
Brand protection	178
Legitimate website	125
Malicious website	35
Total	338

Table 14: Number of malicious IDN homographs.

Homoglyph DB	hpHosts	GSB	Symantec
UC	28	2	1
<i>SimChar</i>	222	12	7
$UC \cup SimChar$	242	13	8

community-based database, had the largest number of entries as we collected data spanning several years. As GSB and Symantec DeepSight are databases maintained by commercial companies, they provide lists of malicious domains that have been inspected by security experts with high confidence. We applied the blacklists to 3,280 of detected IDN homographs, which include non-active domains. Table 14 lists the results. We note that the numbers shown in the table do not include ones shown in the previous subsection; the previously found malicious websites had redirected URLs. By incorporating *SimChar* into the homoglyph DB, the number of detected malicious IDN homographs increased.

6.4 Reverting to Original Domains

Although we begin with a reference domain name list to search for IDN homographs, this approach may not detect IDN homographs if a non-popular website is targeted. Therefore, if we find a malicious domain name, which is composed as an IDN, it is useful to be able to identify the original domain name targeted by the IDN homograph attack. Otherwise, we cannot trace the possible damage caused by the attack. Thanks to the homoglyph database we developed, we can revert to the possible original domain name by replacing a homoglyph with the corresponding Basic Latin letter. We reverted the malicious IDNs to the original domain names and removed those were contained in the Alexa top-1k domains. We ended up 91 of malicious IDNs whose original domains were *not* contained in the Alexa list. This observation indicates that there were non-negligible number of malicious IDN that targeted non-popular websites. Our approach can automatically revert such domains.

7 DISCUSSION

In this section, we first discuss the limitations of our work, after which we consider effective countermeasures against the threat of an IDN homograph attack.

7.1 Limitations

The primary contribution of this study was to build an automated framework that can detect a Unicode homoglyph and an IDN homograph. Below we discuss several limitations of the approaches

we followed for evaluating our framework as well as their future extensions.

Confusability Test In this work, we evaluated the confusability of homoglyphs by a single character, i.e., participants judged whether a potential homoglyph is confusable or distinct by viewing a pair of characters. However, as homoglyphs are generally abused in a word or even in a sentence, we may also need to study the confusability of homoglyphs by using words or sentences because this context may affect the user's perception. The context-aware evaluation of the confusability of a homoglyph is left for future study.

Font Type In this work, we leveraged GNU Unifont, which is a bitmap-based font. GNU Unifont is one of the widely available Unicode glyphs with a wide range of coverage, but many other Unicode fonts are available in the wild, e.g., Noto font [13], which is a scalable font. As our framework is automated, it would be straightforward to extend our evaluation to other font families. This would be a future task.

Measurement Target Our measurement study focused on the world's most popular TLD, .com, yet many other TLDs are used in the wild. For instance, the blacklists we used in this work contain 1,054 of domain names attributed to the 'р ϕ' TLD, which is the Cyrillic country code TLD for the Russian Federation. Studying such class of malicious IDNs from the viewpoints visual deception is left for future study. In addition, although current IDN homoglyph attacks are mainly targeted at ASCII domains, IDNs that contain non-ASCII characters are emerging. Such IDNs may contain ideographs such as Hieroglyphs. Our approach can cover homoglyphs consisting of any characters including the ideographs. Studying these potential targets of homoglyph attacks and their threats would also be a future topic.

7.2 Countermeasures against IDN Homograph Attacks

As we have shown in Section 2, countermeasures against an IDN homoglyph attack implemented in modern browsers have the following drawbacks: if an IDN violates a rule of permitted characters, the countermeasure forcibly represents the IDN in the form of Punycode, which is not a user-friendly expression. This countermeasure may not provide a user with any indication of the *context* behind such transcoding. Moreover, a countermeasure is not effective against non-IDN homoglyphs where homoglyphs reside in the same Unicode block; i.e., the IDN conforms with the rule of permitted characters.

To explicitly inform the user of the possibility of an IDN homoglyph attack with a reasonable context would require the user to be presented with the Unicode representation, instead of forcibly converting the IDN to Punycode. To this end, we could adopt a user interface (UI) that emphasizes the difference between the original domain name and the potential IDN homoglyph. Figure 12 presents an image of such a UI, which could be implemented with the aid of homoglyph databases such as *SimChar* and UC. We note that sizes of *SimChar* and UC are small enough to be embedded into a client program such as Browser extension/plugin-in.. This UI would enable a user to understand which part of a domain name is replaced by which character. This information would be expected to play a vital role in informing the user about the possible threat of a phishing

Figure 12: Image of the UI presenting countermeasures against an IDN homoglyph attack based on the homoglyph database.

attack caused by an IDN homoglyph. More importantly, as an IDN is designed to provide a user-friendly expression of a domain name by using native languages, forcibly converting an IDN to Punycode would significantly impair the user experience. We expect the adoption of such an interface to improve users' awareness of the possible threats posed by an IDN homoglyph attack; i.e., they would become more knowledgeable regarding the context of the presented domain names and be more aware of possible threats. Implementation and evaluation of such a method could be the subject of further study.

7.3 Ethical Considerations

In Section 4, we performed human study to assess the human perception on the detected homoglyphs. Before conducting our human study experiments, we carefully followed the checklist provided by our institutional IRB and concluded that our experiments conformed with the principles of the research ethics. The fact that our user study does not collect any personally identifiable information nor privacy-sensitive information also justifies our conclusion. We also cared the amount of reward for the participants, considering the time to complete a task and minimum wage.

8 RELATED WORK

In this section, we discuss related work in terms of IDN homoglyph detection methods and their measurement studies.

8.1 IDN Homograph Detection

Several studies have led to the proposal of methods to detect IDN homoglyphs. The approaches they followed are broadly classified into two types: image-based and character-based.

Image-based IDN Homograph Detection As an IDN homoglyph exploits the visual similarity between characters, it is natural to apply image-based analysis for detecting these homoglyphs. Liu et al. [21] generated images corresponding to 1.4 million registered IDNs and reference domain names extracted from the top 1,000 domain names listed on Alexa Top Sites. They then detected 1,516 IDN homoglyphs based on the visual similarities between images. Furthermore, they found an additional 42,671 IDNs that were visually

similar to the reference domain names but were still unregistered. Unfortunately, details of their detection methods and settings are not provided in their paper. Sawabe et al. [30] developed a method to detect IDN homographs by leveraging optical character recognition (OCR). The method replaced non-ASCII characters in IDNs with similar ASCII characters using OCR-based image recognition and detected IDN homographs if the replaced IDNs corresponded with reference domain names on Alexa Top Sites.

Character-based IDN Homograph Detection A few researchers adopted the character-based approach. To the best of our knowledge, only two previous studies [26, 37] attempted to apply this approach. Quinkert et al. [26] searched IDN homographs based on a list of homograph pairs, which is equivalent to the homoglyph DB using UC in our study, and detected 2,984 IDN homographs targeting 810 reference domain names. Tian et al. [37] developed a detection method based on UC to identify IDN homographs. As shown in Section 4, our homoglyph DB, *SimChar* outperformed UC-based detection in the sense that the homoglyphs of *SimChar* were perceived to be more confusing than those of UC while maintaining high coverage of homoglyphs; thus, our method complements previous work to cover IDN homographs more comprehensively.

8.2 Measurement Study of IDN Homograph Attacks

Apart from the IDN homograph detection method described above, several researchers have performed measurement studies of IDN homograph attacks in the wild. In 2006, Holgers et al. [15] conducted a passive measurement study on a campus network to search for IDN homographs accessed by users. They also used active DNS probing to detect registered IDN homographs for a limited number of reference domains. Tian et al. [37] studied domains created by various types of domain squatting techniques including IDN homographs to detect phishing websites that exploit homographs in the wild. Le Pochat et al. [24] defined the concept of IDNs that owners of brands with diacritical marks would like to use and generated 15,276 such IDNs. They found that 43% of them were available for registration in 2019. Chiba et al. [5] performed a measurement study to demonstrate that there are many IDN homograph attacks targeting non-English brands or combining other domain squatting methods.

These previous studies mainly focused on the measurement of IDN homographs. We believe our character-based approach to comprehensively detect IDN homographs could be readily applied to these studies, and thus could complement them to provide a more comprehensive understanding of IDN homographs.

9 CONCLUSION

This work led to the development of a new framework named *ShamFinder*, which is useful for detecting IDN homographs efficiently. The key technical contribution of our work was the construction of a new homoglyph database named *SimChar*, which can be updated without requiring time-consuming manual efforts. As *SimChar* is portable, it can be implemented in various systems/platforms as a key component of countermeasures against the threat of IDN homograph attacks. Noteworthy is that *SimChar*

could be used for other promising security applications such as detecting obfuscated plagiarism, which exploits Unicode homoglyphs. We release the code and data of *ShamFinder* [32]. Our future work includes the extension of our study; i.e., extending the domain name space to be explored, extending the font sets, studying the confusability of non-ASCII homoglyphs, etc.

ACKNOWLEDGEMENT

We thank our shepherds, Taejoong Chung and Kimberly Claffy, for their thoughtful suggestions and feedback. We also thank the anonymous reviewers for their fruitful comments. A part of this work was supported by JSPS Grant-in-Aid for Scientific Research B, Grant Number 16H02832.

REFERENCES

- [1] Pieter Agten, Wouter Joosen, Frank Piessens, and Nick Nikiforakis. 2015. Seven Months' Worth of Mistakes: A Longitudinal Study of Typosquatting Abuse. In *Proc. Network and Distributed System Security Symposium (NDSS)*. <http://www.internetsociety.org/doc/seven-months%E2%80%99-worth-mistakes-longitudinal-study-typosquatting-abuse>
- [2] Alexa Top Sites [n. d.]. Alexa Top Sites. <https://aws.amazon.com/alexa-top-sites/>.
- [3] Binance. [n. d.]. Summary of the Phishing and Attempted Stealing Incident on Binance. <https://support.binance.com/hc/en-us/articles/360001547431>.
- [4] Daiki Chiba, Mitsuaki Akiyama, Takeshi Yagi, Kunio Hato, Tatsuya Mori, and Shigeki Goto. 2018. DomainChroma: Building actionable threat intelligence from malicious domain names. *Computers & Security* 77 (2018), 138–161. <https://doi.org/10.1016/j.cose.2018.03.013>
- [5] Daiki Chiba, Ayako Akiyama Hasegawa, Takashi Koide, Yuta Sawabe, Shigeki Goto, and Mitsuaki Akiyama. 2019. DomainScouter: Understanding the Risks of Deceptive IDNs. In *Proc. Research in Attacks, Intrusions and Defenses (RAID)*. 413–426.
- [6] Adam M. Costello. 2003. Punycode: A Bootstring encoding of Unicode for Internationalized Domain Names in Applications (IDNA). RFC 3492. <https://doi.org/10.17487/RFC3492>
- [7] DOMAINLISTS.IO. [n. d.]. Lists of all domains updated daily. <https://domainlists.io/>.
- [8] Patrik Fältström. 2019. IDNA2008 and Unicode 12.0.0. Internet-Draft draft-faltstrom-unicode12. Internet Engineering Task Force. <https://tools.ietf.org/html/draft-faltstrom-unicode12-00> Work in Progress.
- [9] Farsight DNSDB [n. d.]. Farsight DNSDB. <https://www.farsightsecurity.com/solutions/dnsdb/>.
- [10] Patrik Fältström. 2010. The Unicode Code Points and Internationalized Domain Names for Applications (IDNA). RFC 5892. <https://doi.org/10.17487/RFC5892>
- [11] Patrik Fältström and Paul E. Hoffman. 2003. Internationalizing Domain Names in Applications (IDNA). RFC 3490. <https://doi.org/10.17487/RFC3490>
- [12] Evgeniy Gabrilovich and Alex Gontmakher. 2002. The homograph attack. *Commun. ACM* 45, 2 (2002), 128.
- [13] Google Noto Fonts [n. d.]. Google Noto Fonts. <https://www.google.com/get/noto/>.
- [14] Google Safe Browsing [n. d.]. Google Safe Browsing. <https://developers.google.com/safe-browsing/>.
- [15] Tobias Holgers, David E. Watson, and Steven D. Gribble. 2006. Cutting through the Confusion: A Measurement Study of Homograph Attacks. In *Proc. USENIX Annual Technical Conference (ATC)*. 261–266. <http://www.usenix.org/events/usenix06/tech/holgers.html>
- [16] Alain Horé and Djemel Ziou. 2010. Image Quality Metrics: PSNR vs. SSIM. In *Proc. Int. Conf. Pattern Recognition (ICPR)*. 2366–2369.
- [17] hpHosts [n. d.]. hpHosts. <http://www.hosts-file.net/>.
- [18] IDN World Report. [n. d.]. IDN Totals by Year. <https://idnworldreport.eu/charts/idn-totals-by-year/>.
- [19] Internationalization of Domain Names. [n. d.]. <https://tools.ietf.org/html/draft-duerst-dns-18n-00>.
- [20] langid.py [n. d.]. langid.py. <https://github.com/saffsd/langid.py>.
- [21] Baojun Liu, Chaoyi Lu, Zhou Li, Ying Liu, Haixin Duan, Shuang Hao, and Zaifeng Zhang. 2018. A Reexamination of Internationalized Domain Names: The Good, the Bad and the Ugly. In *Proc. IEEE/IFIP Dependable Systems and Networks (DSN)*. 654–665.
- [22] Majestic Million [n. d.]. Majestic Million. <https://majestic.com/reports/majestic-million>.
- [23] Mozilla. [n. d.]. IDN Display Algorithm. https://wiki.mozilla.org/IDN_Display_Algorithm.

- [24] Victor Le Pochat, Tom van Goethem, and Wouter Joosen. 2019. Funny Accents: Exploring Genuine Interest in Internationalized Domain Names. In *Proc. Passive and Active Measurement Conference (PAM)*. 178–194. https://doi.org/10.1007/978-3-030-15986-3_12
- [25] Puppeteer [n. d.]. Puppeteer. <https://pptr.dev/>.
- [26] Florian Quinkert, Tobias Lauinger, William Robertson, Engin Kirda, and Thorsten Holz. 2019. It's Not What It Looks Like: Measuring Attacks and Defensive Registrations of Homograph Domains. In *Proc. IEEE Conf. Communications and Network Security (CNS)*. 259–267.
- [27] Repository of IDN Practices. [n. d.]. <https://www.icann.org/resources/pages/idn-guidelines-2003-06-20-en>.
- [28] Repository of IDN Practices. [n. d.]. <https://www.iana.org/domains/idn-tables>.
- [29] Walter Rweyemamu, Tobias Lauinger, Christo Wilson, William K. Robertson, and Engin Kirda. 2019. Clustering and the Weekend Effect: Recommendations for the Use of Top Domain Lists in Security Research. In *Proc. Passive and Active Measurement Conference (PAM)*. 161–177. https://doi.org/10.1007/978-3-030-15986-3_11
- [30] Yuta Sawabe, Daiki Chiba, Mitsuki Akiyama, and Shigeki Goto. 2019. Detection Method of Homograph Internationalized Domain Names with OCR. *Journal of Information Processing (JIP)* 27, 5 (2019).
- [31] Quirin Scheitle, Oliver Hohlfeld, Julien Gamba, Jonas Jelten, Torsten Zimmermann, Stephen D. Strowes, and Narseo Vallina-Rodriguez. 2018. A Long Way to the Top: Significance, Structure, and Stability of Internet Top Lists. In *Proc. ACM Internet Measurement Conference (IMC)*. 478–493. <https://dl.acm.org/citation.cfm?id=3278574>
- [32] shamfinder [n. d.]. shamfinder. <https://github.com/shamfinder/shamfinder>.
- [33] Symantec. [n. d.]. DeepSight Intelligence. <https://www.symantec.com/services/cyber-security-services/deepsight-intelligence>.
- [34] Janos Szurdi, Balazs Kocso, Gabor Cseh, Jonathan Spring, Márk Félegyházi, and Chris Kanich. 2014. The Long "Tail" of Typosquatting Domain Names. In *Proc. USENIX Security Symposium*. 191–206. <https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/szurdi>
- [35] The Chromium Projects. [n. d.]. IDN in Google Chrome. <https://www.chromium.org/developers/design-documents/idn-in-google-chrome>.
- [36] The Unicode Consortium. [n. d.]. Confusables Data Collection. <http://unicode.org/reports/tr39/>.
- [37] 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 *Proc. ACM Internet Measurement Conference (IMC)*. 429–442. <https://dl.acm.org/citation.cfm?id=3278569>
- [38] Unicode fonts [n. d.]. Unicode fonts. https://en.wikipedia.org/wiki/List_of_typefaces.
- [39] Unicode Inc. [n. d.]. Unicode 12.0.0. <http://unicode.org/versions/Unicode12.0.0/>.
- [40] Unifoundry.com. [n. d.]. <http://unifoundry.com/unifont/index.html>.
- [41] U.S. Department of Labor. [n. d.]. Minimum Wage Laws in the States. <https://www.dol.gov/whd/minwage/america.htm>.
- [42] Verisign. [n. d.]. Top-Level Domain Zone File Information. https://www.verisign.com/en_US/channel-resources/domain-registry-products/zone-file/index.xhtml.
- [43] VirusTotal [n. d.]. VirusTotal. <https://www.virustotal.com/>.
- [44] Thomas Vissers, Wouter Joosen, and Nick Nikiforakis. 2015. Parking Sensors: Analyzing and Detecting Parked Domains. In *Proc. Network and Distributed System Security Symposium (NDSS)*. <https://www.ndss-symposium.org/ndss2015/parking-sensors-analyzing-and-detecting-parked-domains>
- [45] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Processing* 13, 4 (2004), 600–612. <https://doi.org/10.1109/TIP.2003.819861>
- [46] Xudong Zheng. 2017. Phishing with Unicode Domains. <https://www.xudongz.com/blog/2017/idn-phishing/>.