

Comparative Analysis and Framework Evaluating Mimicry-Resistant and Invisible Web Authentication Schemes

Furkan Alaca, AbdelRahman Abdou, Paul C. van Oorschot



Abstract—Many password alternatives for web authentication proposed over the years, despite having different designs and objectives, all predominantly rely on the knowledge of some secret. This motivates us, herein, to provide the first detailed exploration of the integration of a fundamentally different element of defense into the design of web authentication schemes: a *mimicry-resistance* dimension. We analyze web authentication mechanisms with respect to new *usability* and *security* properties related to mimicry-resistance (augmenting the UDS framework), and in particular evaluate *invisible* techniques (those requiring neither user actions, nor awareness) that provide some mimicry-resistance (unlike those relying solely on static secrets), including device fingerprinting schemes, PUFs (physically unclonable functions), and a subset of Internet geolocation mechanisms.

1 Introduction

None of the many schemes proposed over the years to replace password-based web authentication have offered sufficient usability and deployability benefits to displace passwords at the Internet scale [1]. Since passwords do not seem to be disappearing [2], a prominent avenue of improvement is to reinforce their security by parallel mechanisms [3] without further burdening users.

A challenge in providing sufficient security guarantees for web authentication, *i.e.*, *user-to-web* and *device-to-web* (versus *user-to-device*), is that the security of many schemes, including those relying on *something-you-have* or *something-you-are*, requires the ability to protect a secret. For example, a physical biometric such as a fingerprint can be captured in transit and replayed by an attacker without possession of the physical fingerprint, resulting in security properties similar to other stored secrets such as passwords. The reliance of many schemes on an element of secrecy is reflected in

- *F. Alaca is with the Department of Mathematical and Computational Sciences at the University of Toronto Mississauga, Canada. This work was done while he was a PhD candidate at the School of Computer Science at Carleton University.*
- *A. Abdou is with the School of Computer Science, Carleton University, Canada.*
- *P.C. van Oorschot is with the School of Computer Science, Carleton University, Canada.*

Version created on: July 9, 2019. The final version of this paper will appear in the IEEE Transactions on Dependable and Secure Computing (TDSC). This is the authors' copy for personal use.©IEEE.

the Usability-Deployability-Security (UDS) evaluation framework [1], where eight of nine security properties assess a scheme's resilience against the exposure of a secret.

We revisit the process of compromising an account from an attacker's perspective, now viewing it as a two-stage process involving both exposure and mimicry. Exposure refers to the capture of information that enables account access, such as a password, session cookie, or a cryptographic key; mimicry refers to actions performed by the attacker to impersonate the legitimate user's behavior or associated characteristics, such as replaying a password or spoofing a user's geographic location (*i.e.*, in location-based authentication [4]). Accordingly, a scheme's authentication token may resist attacks by both: resisting exposure and resisting mimicry.

The term *mimicry* has been used previously in the context of intrusion detection systems, referring to the attacker's ability to mimic legitimate traffic [5]. In contrast, the mimicry-resistance dimension has been relatively little-studied in the context of web authentication. We therefore investigate mimicry-resistance in web authentication herein and systematize elements not previously addressed in the literature. We first define a suitable set of criteria and then rank schemes across a continuum of three classes of resistance to mimicry, as detailed in Section 3. Mimicry-resistance has been explored within the context of *user-to-device* authentication; for example, Khan et al. [6] analyze mimicry attacks on *user-to-device* authentication schemes that rely on behavioural biometrics, such as touchscreen finger movement patterns. In contrast, herein we analyze web authentication schemes; our analysis finds that most of these schemes offer little to no resistance to mimicry.

To construct a more comprehensive evaluation framework for authentication schemes, we augment the existing UDS security properties, which concentrate on exposure-resistance, with new properties measuring mimicry-resistance. We leverage the UDS framework security properties to systematically rate authentication schemes across an additional continuum ranging from lowest to highest resistance to exposure, and use these as orthogonal axes (exposure and mimicry) to plot a two-dimensional chart. Along both dimensions, our evaluation also reflects the scalability of attacks required to defeat a scheme.

The lack of mimicry-resistant schemes among those previously evaluated under UDS [1] motivates us to select and

evaluate a representative set of techniques for reinforcing web authentication that demonstrate the benefits of mimicry resistance. Our constructed framework dissects various degrees of mimicry resistance, and helps analyze currently known methodologies that may provide some degree of resistance to mimicry attacks. We then evaluate techniques that fall under four approaches, namely device fingerprinting (FP) [7], Internet geolocation [8], Physically Unclonable Functions (PUFs) [9], and One Time Passwords (OTPs) [10]; some variations of these offer resistance to mimicry attacks, and/or are also *invisible*, in that they do not require any user effort to configure or use (see Section 2). We supplement the UDS framework with two usability properties and four security properties (Table 3). Under this revised framework, we evaluate techniques which fall under the four aforementioned approaches when combined with passwords, and find that invisible and mimicry-resistant schemes combined with passwords provide significantly higher resistance to attack. This constitutes an initial step towards identifying mimicry-resistant web authentication schemes that can enhance security with minimal usability penalties.

In summary, the following contributions are made:

- Investigating the mimicry-resistance dimension in web authentication, including ranking schemes under three sub-classes of mimicry resistance.
- Exploring newer invisible techniques, and evaluating their degree of mimicry resistance when used for web authentication.
- Constructing a comprehensive evaluation framework, which includes (a) a two-dimensional chart combining the exposure and mimicry resistance dimensions, to visually reflect the ability of a scheme to resist scalable attacks; (b) an augmented UDS framework.
- Using the augmented UDS framework herein for the first detailed exploration of the benefits of combining mimicry-resistant web authentication techniques with ubiquitous password-based *user-to-web* authentication.

The remainder of this paper is organized as follows. Section 2 provides background. Section 3 introduces the mimicry-resistance dimension in web authentication, and plots a representative subset of authentication schemes onto an Exposure-Mimicry two-dimensional space. Section 4 evaluates relatively little-explored schemes found to have mimicry resistance, using a modified UDS framework with new properties addressing the mimicry resistance dimension. Section 5 analyzes benefits when the aforementioned schemes are combined with passwords. Section 6 concludes.

2 Background and Context

We briefly provide background and define terms for modes of authentication (Fig. 1) and review related work on the evaluation of authentication schemes.

The standard method of user authentication on the web is for the user to type a password into a web form, which is submitted over HTTP(S) for the web server to verify. We refer to this as *user-to-web* (see Fig. 1) authentication—even though the password is physically entered into the user’s local device, the verifier is a remote web server. The device used for authentication, which we call the *access device* herein,




Fig. 1. A comparison between *user-to-device*, *device-to-web*, and *user-to-web* (directly).

passively conveys the authentication token,¹ and does no checking on behalf of the web server. In contrast, when a user authenticates to their smartphone, it is *user-to-device*. Some authentication schemes that appear to be *user-to-web* are actually *two-stage* authentication schemes that combine a *user-to-device* scheme and a *device-to-web* scheme; for example, a mobile payment app may authenticate the user via a biometric (*e.g.*, fingerprint or iris scan), which unlocks a locally-stored cryptographic key used by the mobile device to authenticate to the remote server.

Implicit vs. invisible authentication. Implicit authentication [11] schemes reduce user burden by authenticating users without requiring any deliberate user effort at login time by, *e.g.*, measuring and using users’ biometric attributes or physiological behaviors for authentication. Many such schemes have been proposed [12] [13] [6], and are typically used for *user-to-device* authentication; the accuracy and resistance to mimicry of some of these schemes have been evaluated [6]. In contrast, herein we consider web authentication schemes, thus focusing on *user-to-web* and *device-to-web* authentication. As such, *user-to-device* (*e.g.*, [14] [15]) and similarly other paradigms such as *device-to-device* (*e.g.*, [16] [17]), are out of the scope of our work.

We also make the distinction between *implicit* and *invisible* schemes. The former refers to schemes that do not require extra user effort during login, but do require some initial user effort for setup; those are generally *user-to-device* schemes that authenticate the user based on their behaviour as measured through various sensors (*e.g.*, accelerometer, swipe patterns), where the initial user effort is often downloading an app, or calibrating input sensors. On the other hand, we define *invisible* schemes to be *device-to-web* authentication schemes requiring no user involvement at all, neither during set-up nor login. Note that not all *device-to-web* schemes are invisible, as some require a user to carry out an action; for example, in a *device-to-web* scheme that fingerprints the access device’s accelerometer while at rest [18], users may need to place their device on a flat surface.

3 The Mimicry Resistance Dimension

User authentication typically relies on *something-you-know* (*i.e.*, some secret the user knows), *something-you-have*, or *something-you-are*. Biometric-based *user-to-web* authentication mechanisms (*i.e.*, *something-you-are*) are similar in server-side implementation to *something-you-know*, since biometric data is (preferably) stored as a digital secret. Since the transmission path from the user’s biometric sensor (*e.g.*,

1. We use the terms *token* (by default indicating a digital token such as a password) and *credential* interchangeably. Hardware tokens will be explicitly identified.

fingerprint reader) to the authentication server is often untrusted, or at least less so than the path from the sensor to an authenticating application in *user-to-device* authentication, an attacker may defeat authentication by simply replaying an exposed secret (*e.g.*, a fingerprint). Exposure can occur through, *e.g.*, guessing or capture. Most *something-you-know* and *something-you-are* web authentication schemes offer limited resistance to mimicry, since exposure of a user secret typically allows attackers to trivially defeat the scheme.

User authentication via *something-you-have* requires verifying servers to both authenticate the hardware token *and* to verify user possession of the hardware token. Hardware tokens are generally electronic devices (*e.g.*, a USB OTP token or smartphone) that can be authenticated by the server using, *e.g.*, cryptographic techniques.² Because authenticating these hardware tokens will almost always rely on secrets, *something-you-have* often boils down to a *something-you-know*, *i.e.*, something *the hardware token* knows. Most known variations of hardware authenticator tokens are defeated once that secret is exposed [1], *e.g.*, by capture or theft, again providing little to no resistance to mimicry.

3.1 The Exposure-Mimicry Duality

Defeating web authentication typically requires two actions: exposing a token and mimicking certain behavior. That behavior is any action the legitimate user (or device) normally performs while authenticating to a service; for passwords, that behavior is trivially mimicked by simply submitting a static string. A scheme’s resistance to compromise thus depends on its ability to independently resist *exposure* and *mimicry*. To evaluate a scheme’s resistance across these orthogonal components, we construct a two-dimensional space in Fig. 2 and plot various web authentication schemes on it. A marker (*i.e.*, dot) represents a scheme’s authentication token. Height along the y-axis indicates resilience to exposure; distance from the origin along the x-axis indicates resilience to mimicry. Schemes placed on the chart are explained in Section 3.3.

3.1.1 Vertical axis

The y-axis of Fig. 2 is split into three segments: V1-Negligible-resistance (*i.e.*, guessable), V2-Guess-resistant, and V3-Leak-resistant. Guessing a credential is the easiest form of exposure. Digital theft (*leak* henceforth) is generally easier than physical theft in that (1) leaks can often be arranged at scale, *e.g.*, by phishing, and (2) leaks can be by remote access, *e.g.*, an Internet-facing authentication server is subject to attack from anywhere in the world (an adversary need not be in physical proximity).

V1. Within the lower-vertical segment, a credential requiring more guesses is placed higher. The guessability of a scheme’s credential may depend on several factors. A very weak password, for example, could be easier to guess than a randomly-generated PIN. Unless stated otherwise, schemes are plotted according to their strength in typical scenarios (*e.g.*, passwords are assumed to be user-chosen, which limits their resilience to guessing). Guessing attacks are assumed to follow common guessing strategies for *trawling attacks*

2. Physically Unclonable Functions (PUFs) do not hide a key, but still have a component that may be exposed albeit harder to reproduce/mimic. See Section 3.3.4 for further discussion.

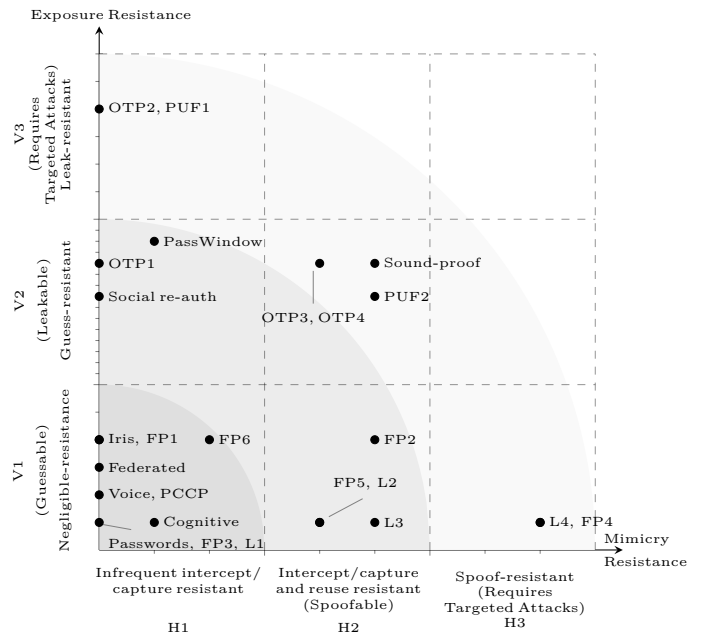


Fig. 2. Exposure resistance and mimicry resistance as two dimensions to measure authentication-scheme resistance to compromise. Along each axis, distance from the origin reflects the scalability of attacks required to defeat a scheme: those plotted closer to the origin (darker regions) can be defeated by attacks that can be scaled with ease (such as online guessing) to break a large number of accounts; defeating schemes plotted further from the origin requires more targeted attacks (such as device theft) that are highly unscalable. Section 3.3 gives detailed explanations of different classes of one time passwords (OTP_n), device fingerprinting (FP_n), geolocation (Ln) techniques, and physically unclonable functions (PUF_n).

(capturing as many accounts as possible, often by guessing the most common password across all accounts, then moving down a list of candidate passwords).

V2. The middle-vertical segment comprises schemes with secrets impractical to guess, such as cryptographic keys or randomly-generated passwords with sufficient length to withstand offline guessing attacks. Those may still be leaked (*e.g.*, client-side malware). Malware may either (1) steal login credentials for later user impersonation, or (2) remain dormant until a user logs into their account, after which attacker-created transactions can be authorized by the malware on the victim user’s device [19]. Herein we are interested in evaluating an authentication scheme’s resilience to impersonation, not to malicious authorization. Defending against the latter may require mechanisms such as out-of-band authorization for sensitive transactions (which may be independent from the method of user authentication being used).

We consider within this segment four sublevels, based on the number of sources from which leaks may occur: the human user, the browser, the user device (*e.g.*, laptop or smartphone), and a public server whose compromise defeats authentication; *e.g.*, a trusted third party of the main authentication server, or a party that also stores the same credential (*e.g.*, a same password used across multiple websites). A scheme subject to leaks from all four sources is placed at the lowest of V2’s four sublevels; one subject to leaks from any three places is second-lowest, and so on. A scheme not subject to leaks from any of these four sources is placed in V3.

V3. Schemes in the upper-vertical segment are those resilient to exposure by digital theft (*e.g.*, leak or capture)—thus mostly physical tokens of some sort. Compromising these requires targeted attacks (most commonly physical theft) involving physical proximity to specifically pre-identified users. A scheme’s vertical position within V3 varies with vulnerability to theft. A smartphone for example, relatively small in size and carried around more often, is easier to steal than a desktop PC—the latter may require physical break-in to an office and effort to conceal the escape.

Sorting rationale. The intuition behind arranging the three segments in the above manner follows logically from the trawling attack strategy of maximizing the number of accounts broken into per unit of attack effort. Assuming, for example, that a website does not throttle online password guessing, a good attack strategy is to try guessing passwords. If online guessing fails, an attacker often moves to digital theft (V2 segment on Fig. 2), *e.g.*, stealing credentials via XSS attacks, phishing, or client-side malware. Since random passwords may be leaked but not efficiently guessed, they are harder to expose. Weak passwords or PINs, for example, can be guessed and often also captured,³ and are rated lower in resistance to exposure. If all forms of leaks (digital theft) fail, the attacker is left with physically stealing, *e.g.*, a smartphone or hardware authentication token. These attacks also become gradually harder to scale (for an adversary) in the aforementioned order.

Relationship to UDS framework. As summarized in Table 1, the vertical axis in Fig. 2 addresses the S (Security) benefits in the UDS framework, excluding *Unlinkable* and *Requiring-Explicit-Consent*.⁴ The UDS framework did not intend to provide an overall summary rating for a scheme, as that would require subjective weighting of the usability, security, and deployability benefits (which may be context-dependent). In contrast, we use our sorting rationale (as discussed above) in conjunction with the UDS security benefits to rate a scheme’s overall resistance to exposure; for example, if a scheme only fails to provide *Resilient-to-Unthrottled-Guessing* but provides the remaining eight security properties, it is placed near the bottom of the vertical axis within V1 (despite having a virtually full row of security bullets). Passwords fail to provide *Resilient-to-Throttled-Guessing* and *Resilient-to-Unthrottled-Guessing* and are therefore placed in V1, despite providing *Resilient-to-Physical-Theft* (which corresponds to our highest exposure resistance category, segment V3).

Method for Plotting Schemes (Vertical Axis). Schemes are plotted along the vertical axis (see Section 3.3 for placement rationale) based on their security benefits offered. For schemes already evaluated in the UDS paper [1] we use

3. An example where a password can be guessed but not leaked would be challenge-response schemes where the password is never typed on the keyboard nor stored anywhere (neither on the server nor any client device), such that the user computes the response from the challenge in their head, or offline using a calculator. That would be resilient to phishing, theft, malware, leaks from verifiers, requires no trusted third parties, and possibly physical observation.

4. These two security benefits from the original UDS framework do not share the focus of Fig. 2 on exposure-resistance. *Unlinkable* is a privacy benefit and *Requiring-Explicit-Consent* relates to malicious authorization, *e.g.*, a malicious RFID-based card reader embedded in a sofa that authorizes a transaction without user knowledge [1].

TABLE 1

Original UDS security benefits [1] evaluating schemes by susceptibility to exposure, and relation to new framework. This table lists the properties a scheme must offer to move to the next-higher vertical segment in Fig. 2; *i.e.*, both S3 and S4 must be offered to be placed in V2; all of S1, S3-S7 and S9 must be offered to be placed in V3. The benefits are listed in descending order of the scalability of carrying out the corresponding attack; *e.g.*, guessing attacks (S3) are highly scalable, but physical theft (S8) is highly unscalable.

UDS SECURITY PROPERTY	SEGMENT
S3. <i>Resilient-to-Throttled-Guessing</i> S4. <i>Resilient-to-Unthrottled-Guessing</i>	Lower (V1)
S1. <i>Resilient-to-Physical-Observation</i> S5. <i>Resilient-to-Internal-Observation</i> S6. <i>Resilient-to-Leaks-from-Other-Verifiers</i> S7. <i>Resilient-to-Phishing</i> S9. <i>No-Trusted-Third-Party*</i>	Middle (V2)
S2. <i>Resilient-to-Targeted-Impersonation†</i> S8. <i>Resilient-to-Physical-Theft</i>	Upper (V3)

* Additional third parties increase chances of leaks from public servers. From the adversary’s perspective, this property is similar to *Leaks-from-Other-Verifiers*. See inline for details.

† S2 encompasses a variety of targeted attacks, *e.g.*, deceiving a human acting as a trusted party as in social re-authentication [20], locating a user’s password written on a post-it note, or lifting fingerprints from a doorknob [1]. The success of these attacks require attackers to be in physical or logical proximity (*e.g.*, social engineering attacks against social re-authentication) to victims. The difficulty in scaling such targeted attacks is similar to attacks involving physical theft (S8). These differ from more scalable attacks such as educated guessing based on information scraped *en-masse* from social media—these are covered by the more general guessing attacks (S3 and S4).

the rating given therein. The following criteria (based on the mappings in Table 1) are used to decide which of the three vertical segments (V1, V2, V3) the scheme will fall into:

- 1) Schemes subject to guessing attacks (*i.e.*, not offering both S3 and S4) are placed in V1.⁵
- 2) Schemes offering S3 and S4, but lacking any of S1, S5-7, or S9, are placed in V2.
- 3) Schemes offering all of the above benefits are placed in V3.

Moreover, each vertical segment has sublevels that are determined based on the number of security benefits captured within each segment, as specified in Table 1. Each segment has precisely $\sum_{i=2}^{N+1} i$ sublevels, where N is the number of benefits associated with that segment; thus, V1 and V3 each have five sublevels, and V2 has 14 sublevels. Each scheme is assigned to a sublevel such that a security benefit offered (*i.e.*, represented by a filled-circle “●” in Table 3) has a greater score than any number of partially-offered benefits (*i.e.*, represented by an empty-circle “○” in Table 3). For example, a scheme that offers all V1 benefits but no V2 benefits would be placed at the first sublevel of V2; a scheme that partially offers four V2 benefits would be placed at the fifth sublevel of V2; a scheme

5. When evaluating specific implementations of authentication schemes, resistance to guessing attacks can be more precisely calculated (*e.g.*, based on the specific throttling techniques used); V1 could then be divided into more granular sublevels based on the number of guesses required to break a scheme. Schemes breakable with zero guesses would be plotted at the origin, and schemes breakable with an offline guessing attack of just under 10^{20} guesses (the approximate threshold at which a scheme would be considered to be resistant to offline guessing attacks [21]) would be plotted at the upper boundary of V1.

TABLE 2

Three new UDS-type security properties (given in Section 4) directly related to mimicry. A scheme must achieve these to move further right across the horizontal segments in Fig. 2. Example: Both M2 and M3 must be achieved to be placed in H3.

NEW UDS SECURITY PROPERTIES	SEGMENT
M2. <i>Resilient-to-Infrequent-Capture-or-Intercept</i>	Left-most (H1)
M3. <i>Verifies-Non-Static-Information</i>	Middle (H2)
M4. <i>Resilient-to-Spoofing</i>	Right-most (H3)

that offers one V2 benefit and partially offers one V2 benefit would be placed at the seventh sublevel of V2; and so on.⁶

3.1.2 Horizontal axis

The x-axis (Mimicry Resistance) reflects burdens on the attacker to mimic the verifier-expected behavior, *after* exposing a scheme’s credential. Unfortunately, the security of most authentication schemes proposed to date rely on some form of resisting exposure, with compromise complete once an underlying secret token is revealed. This includes passwords, *user-to-web* physical biometrics (see Section 2), and poorly managed/generated private keys [22]. If leaking a credential does not allow an attacker access by direct replay of the credential, the scheme exhibits some degree of mimicry resistance, and gains horizontal distance from the origin in Fig. 2.

We split the horizontal axis into three segments: H1-Negligible-resistance (*i.e.*, easily replayable), H2-Replay-resistant, and H3-Spoof-resistant. In Section 4, we expand the UDS framework by six new benefits; three of these benefits (M2, M3, and M4) correspond to the aforementioned horizontal segments (H1, H2, and H3) as specified by Table 2.

H1. If an attacker can simply submit or replay a credential after exposing it, the scheme provides no resistance to replay and thus gains no horizontal distance along the Mimicry Resistance dimension (*i.e.*, lies on the vertical axis). For example, replaying a captured password is trivial; thus, passwords lie directly on the vertical axis in Fig. 2.

In some cases, there is no clear-cut answer as to whether exposing information directly helps the attacker gain access. For example, answers to personal knowledge questions based on recent account activity (*e.g.*, *who was the last person you emailed?*) likely remain unchanged for a few hours. Capturing an answer allows the attacker to login only within that window. Such schemes are given some horizontal distance from the origin, but remain within H1. Another example is password expiration policies [23], [24] forcing resets every, *e.g.*, one hour, which would provide some mimicry resistance (albeit highly unusable), and placed within H1 on Fig. 2.

Some challenge-response schemes are also candidates for H1; whether or not a captured response directly allows the attacker to compromise the account is conditional on whether the captured response remains valid. Schemes where the set of challenge-response pairs is finite and relatively small, or when capturing a handful of challenge-response pairs suffices for defeating authentication, would thus fall within H1. Cognitive schemes like Weinshall [25], and some challenge-response

schemes such as PassWindow [26], are prominent examples. Other challenge-response schemes are placed in H2.

H2. After a credential is leaked, schemes that either require additional attacker actions beyond simply replaying a string or conducting a straightforward operation (*e.g.*, cryptographic signing), or where attackers have a limited time window (*e.g.*, < 2 mins) in which the credential can be used, are considered replay-resistant and placed within H2. An example is DNS resolver fingerprinting (FP5 in Fig. 2), which determines the DNS resolver used by the client (see Section 3.3.2); knowing the resolver’s address is insufficient to defeat authentication, as the attacker must also be able to use it to resolve domain names.

One-time password (OTP) over SMS is another example scheme placed within H2, as it enables the attacker to compromise the account only if the user has not already used the OTP. An attacker that captures the OTP in clear text [27] must use it before the legitimate (victim) user. Additionally, the server may set a 2-minute time window where an OTP token is valid for usage, and expires afterwards. Such schemes would be placed within H2 on Fig. 2.

H3. We place within H3 any scheme that requires additional equipment and/or systems (*e.g.*, hardware chip manufacturers, or large scale distributed botnets) to mimic the behavior that the server measures and expects from the legitimate user. An example is robust location verification (see L4 in Section 3.3.1).

Because spoof-resistant schemes rely on measurements of various phenomena (*e.g.*, a user behavior or habit), the measurement process must typically allow some degree of tolerance to account for (1) imperfections in the measuring apparatus, or (2) the phenomena’s natural instability over time.

For (1), consider an example location-based authentication scheme [28], where a server grants access to a user only if the user is at an expected geographic location. Ideally, the user’s location would be identified to such a high granularity that no two human beings could exist at the same {latitude, longitude, altitude} coordinates.

For (2), again using location-based authentication as an example, a user’s daily commute [29] from one geographic location to another (*e.g.*, work to home) makes it impractical to grant account access only when the user is geographically present in, *e.g.*, a 1 m^2 area, even if the geolocation mechanism being used allows for such high accuracy. Geolocating users with coarser granularity (*e.g.*, city-level) would thus seem more practical (*e.g.*, to avoid false rejects) for generic location-based authentication purposes.

The degree of “fuzziness” (*i.e.*, lack of precision) introduced while measuring a user’s behavior in spoof-resistant schemes is thus likely to result in *false accepts*, *i.e.*, falsely accepting another user as the intended one. In practice, some websites use IP-address based location look-ups [30] to implement location-aware authentication, which returns locations at the city or state level. Thus, all attackers physically present in the user’s city (or even a legitimate user mistyping the username with no malicious intent, *i.e.*, if geolocation is used as a stand-alone authentication scheme) are falsely accepted.

Note the distinction between false accepts (a result of imprecision) and mimicry resistance (a result of resistance to attack): A scheme that uses high-precision GPS coordinates

6. To avoid visual ambiguity, no sublevels are plotted at the boundary between any two segments, namely the V1-V2 and V2-V3 boundaries on the vertical axis or the H1-H2 and H2-H3 boundaries on the horizontal axis. Therefore, no schemes are placed at these boundaries.

(*e.g.*, within a radius of $1m$) reported by the user browser (*e.g.*, L1, Section 3.3.1) may offer low false accept rate, but offers no mimicry-resistance since attackers can use browser extensions that report forged coordinates.

False accepts are not directly reflected in the Mimicry-Resistant dimension (though we capture them in the evaluation framework in Section 4). They are indirectly captured however by a scheme’s position vertically. This is because resilience to guessing attacks and resilience to false accepts are both related to the size and distribution of the credential space (*e.g.*, cryptographic keys with sufficient length and chosen uniformly at random are not subject to guessing attacks and false accepts).

Method for Plotting Schemes (Horizontal Axis).

Schemes are plotted along the horizontal axis based on their mimicry-resistance benefits offered. The following criteria (based on the mappings in Table 2) are used to decide which of the three horizontal segments (H1, H2, H3) the scheme will fall into:

- 1) Schemes that do not offer benefits M2, M3, and M4 offer no mimicry-resistance, and thus are placed on the vertical axis (*i.e.*, with zero horizontal displacement).
- 2) Schemes offering M2 are placed in segment H1.
- 3) Schemes offering M3 are placed in segment H2.
- 4) Schemes offering M4 are placed in segment H3.

Since H1, H2, and H3 are each associated with a single benefit (see Table 2), each horizontal segment has two sublevels, where the first sublevel represents a partially-offered benefit and the second sublevel represents a fully-offered benefit.

3.2 Interpreting Attack Scalability

Figure 2 illustrates the relative resilience of schemes against attackers that aim to maximize the number of accounts broken into (*i.e.*, trawling attacks). Along each individual axis, schemes further from the origin are less susceptible to scalable attacks. Scalable attacks are defined as attacks that can be scaled to break into a large number of accounts. Non-scalable attacks are typically highly targeted (*e.g.*, physical device theft); such attacks impose a higher cost on the attacker in terms of time, effort, and/or money, as a function of the total number of accounts targeted.

As our evaluation is qualitative, it is not intended that absolute distance from the origin be a quantitative measure suitable for comparing schemes. However, the resilience of two schemes (*i.e.*, scalability of the attacks that a scheme can withstand) can be compared if (1) they share one of the two coordinates (*i.e.*, identical x - or y -coordinates), or (2) one of the schemes is further from the origin along both its x - and y -coordinates. When comparing any two schemes that satisfy one of these conditions, a scheme that offers both high exposure-resistance (*i.e.*, high y -value in Fig. 2) and high mimicry-resistance (*i.e.*, high x -value in Fig. 2) is more costly (less scalable) to attack, since it requires the attacker to successfully mount both an exposure attack and mimicry attack, as elaborated on below.

Along the vertical (exposure-resistance) axis, schemes most vulnerable to scalable attacks are those where the attacker can guess user credentials (V1). Attacks against schemes requiring information theft (*e.g.*, via software vulnerabilities or malware) are less scalable (V2), and schemes

requiring physical device theft (V3) are the least scalable (*i.e.*, most costly) to attack.

Along the horizontal (mimicry-resistance) axis, attacks against schemes where an exposed secret can be used as-is by the attacker (see Section 3.1) are most scalable; attack scalability decreases for schemes that increase the attacker burden (in terms of hardware costs or control of network infrastructure) for mimicking account-holders beyond a simple replay attack or capture of a static secret (see Table 2).

3.3 Scheme Placement on Figure 2

We use the mapping method of Section 3.1.1 (“Relationship to UDS framework”) to position a representative subset of schemes along the exposure-resistance axis using their previously assessed (*cf.* [1, p.11]) security benefits. Upon evaluating the mimicry-resistance of these schemes, we find they offer little to no resistance to mimicry as we explain below.

PassWindow [26] is a visual crypto technique whereby the user overlays a transparency card on the screen to visually read and enter a 4-digit pass-code, which proves possession of the user’s card. It was shown [31] that observing 20–30 challenge/response pairs can leak the card’s secret; we thus place the scheme within H1. Likewise, cognitive schemes, including GrIDSure [32], Weinshall [25], Hopper Blum [33], and Word Association [34] are placed in H1 since the schemes can be compromised after few observations (*e.g.*, Weinshall can be compromised after about seven observations [35]).

In social re-authentication [20], a trusted friend can vouch for a user who, *e.g.*, has lost their credentials. An attacker may capture a vouch code (*e.g.*, via malware on a trusted friend’s device) and simply replay it to authenticate as the victim; we thus place it in H1.

Federated schemes, including OpenID [36], Microsoft Passport [37], Facebook Connect, and BrowserID [38], offer no mimicry-resistance if password authentication is used, and we thus place them in H1. Their position on the chart would change if the identity provider utilizes mimicry-resistant schemes. Persuasive Cued Click-Points (PCCP) [39] and biometric schemes on Fig. 2 (Iris and Voice) are also placed in H1, since captured biometric samples are typically replayable [1].

We explore and discuss example schemes which do provide certain mimicry-resistance, and position them along the two axes (the remaining schemes on Fig. 2). To avoid redundancy, we only explain the horizontal position of these schemes in this section; see Section 4 for their vertical position.

3.3.1 Geolocation (selected methods)

We review four broad classes of Internet geolocation techniques that can be used for location-based authentication, placing a marker for each class on Fig. 2 to rate their ability to resist compromise. Location-based authentication typically requires server-side storage of the user’s expected location (*e.g.*, city-level, nation/state level, or latitude/longitude coordinates, depending on the geolocation method employed) in plain text, hashed, or cloaked [40] to a certain degree to preserve users’ privacy. Any of the four classes of geolocation methods described below may be used by a verifying server to obtain and/or verify the user’s current location at the time of authentication, and grant access if the user is verified to be at the location expected by the server.

L1: GPS and WPS. GPS and WiFi Positioning System (WPS) geolocation are commonly used in practice. WPS uses multi-lateration based on the signal strengths between the device and nearby WiFi access points with known locations. These techniques are usually selected by the user’s browser in the W3C geolocation API [41], whereby the browser obtains the coordinates from the device’s GPS driver, or from a location provider after submitting a list of nearby WiFi access points and their signal strengths to the location provider. L1 gains no depth across the Mimicry Resistance dimension (see Fig. 2), since techniques rely on browser-reported information that can be substituted and replayed by an attacker, so knowledge of the user’s location is sufficient to break authentication.

L2: IP-address based tabulation. Tabulation-based geolocation service providers such as Maxmind⁷ and ipinfo⁸ maintain lookup tables, which map IP address blocks to cities and countries, possibly through publicly available information such as IP address registries (*e.g.*, whois⁹) and the geography of IP address allocation. Geolocation based on IP address table lookups is unreliable due to outdated entries [30], and is evadable through use of middleboxes and virtual private networks (VPNs) [42]. L2 is rated slightly more resilient to compromise than forgeable GPS/WPS coordinates, and is placed within H2; it is not in H3 since even an attacker unable to forge source IP addresses may use public HTTP proxies or bots in close proximity to the user.

L3: Measurement-based geolocation. In this class, network measurements, such as Round-Trip Times (RTTs), are conducted from a set of landmarks (*e.g.*, cloud-based or CDN servers) to the target user, and are then mapped to geographic distances using a pre-calibrated delay-to-distance function. The user’s location is estimated through multi-lateration, relative to the landmarks’ locations. Examples include Spotter [43] and CBG [44]. Other proposals have suggested mapping the network topology for higher accuracy [45]. To date, measurement-based geolocation can achieve an accuracy on the order of a few tens of meters. Manipulating L3 requires more advanced techniques than simply submitting forged coordinates, but can be achieved with enough knowledge of the network topology and the landmarks/verifiers being used [46], [47]; they are limited to H2.

L4: Robust location verification. Some techniques are designed to verify location information obtained by other Internet geolocation techniques (*e.g.*, L1-L3, above), and/or are designed to be resilient to common adversarial manipulation tactics. The result of a preliminary geolocation is treated as an asserted location (analogous to a username asserting identity), to be verified by a measurement-based proof (analogous to proof of knowledge of a secret). Examples of such schemes include Client Presence Verification (CPV) [48] (see Section 4) and Trusted Platform Module (TPM)-supported GPS drivers [49]; the former cryptographically protects network delay measurements used for verifying location assertions, and the latter communicates coordinates securely. We call such techniques robust location verification and rate them as spoof-resistant (H3) because manipulation requires attackers to expend more effort than simply reporting a false location;

successfully spoofing legitimate client locations requires using specialized proxy machines (for attacks that “co-locate” using a machine nearby to the victim) or GPS satellite signal-spoofing devices [50].

3.3.2 Device Fingerprinting

Device fingerprinting refers to techniques by which a server collects information on a device’s hardware/software configuration for the purpose of identification [7]. From 29 methods recently surveyed [51], we derive six representative categories for evaluation.

FP1: System parameters/preferences. This class includes software and hardware information about the device to be authenticated, provided by the web browser’s JavaScript API (*e.g.*, the navigator and windows JavaScript BOM objects), such as operating system version, screen resolution, time zone, system language, and supported WebGL capabilities. FP1 lies on the y-axis, since it can simply be mimicked by replaying the information.

FP2: Audio and canvas challenge/response. This class includes two techniques that fingerprint the client’s graphics and audio subsystems, respectively. HTML5 canvas fingerprinting renders a variety of text and graphics in an HTML5 canvas on the client’s browser, which results in subtly different images (*e.g.*, due to differences in anti-aliasing or font smoothing) depending on the graphics driver and hardware on the device being fingerprinted. Audio processing fingerprinting leverages the HTML AudioContext API that provides real-time frequency- and time-domain analysis of audio playback, mainly used for creating audio visualizations. Playing the same sound on different devices results in subtly different waveforms, depending on the sound driver and hardware. To improve resistance to replay attacks, these two techniques can be used in a challenge-response scheme, where the server can store the client’s responses to many different challenges [52], and are thus placed within H2.

FP3: Hardware sensors. This class, suitable to fingerprint smartphones, leverages the inherent variation in the manufacturing and factory calibration of typical smartphone sensors, such as accelerometer and speaker-microphone systems. Similar to FP1, FP3 lies on the y-axis since information can simply be replayed.

FP4: Clock skew. The server uses TCP timestamps to measure the clock skew of the device being fingerprinted, which differs across devices due to manufacturing variation. FP4 is placed within H3, since clock skew spoofing attacks are highly sophisticated; *e.g.*, Arackaparambil *et al.* show how timestamp manipulation can be detected to identify rogue wireless LAN access points [53]. However, clock-skew spoofing over a WAN connection is not well-studied, and therefore this rating may change subject to further study and experimentation.

FP5: DNS resolver. The server determines a client’s DNS resolver(s) by presenting to the client a page that contains a reference to a randomly-generated (non-existent) subdomain, triggering a client DNS lookup; as a result the server will receive a DNS query from the client’s resolver. The DNS resolver’s IP address serves as a fingerprint. FP5 provides partial replay-resistance but can be defeated via the use of proxies, similar to L2. In some cases, organizations run their own DNS servers, which resolve domain names only to

7. <https://www.maxmind.com/>

8. <http://ipinfo.io/>

9. <https://www.whois.net/>

machines within the department’s network. To use a victim’s DNS resolver, an attacker would need to be able to resolve domain names using the organization’s private DNS server. Since identifying the server is not enough (and is easier than using the organization’s DNS), FP5 is placed within H2.

FP6: Protocol-based fingerprinting. This class includes schemes that glean information from network-, transport-, and application-layer protocol fields. For example, the TLS library of the client device can be fingerprinted using the Client Hello packet received during the handshake sequence, which includes information such as the device’s supported TLS version, supported ciphersuites (and their order of presentation), compression options, and list of supported extensions (along with associated parameters such as elliptic curve parameters). FP6 lies within H1, since it is susceptible to mimicry, but attacks are less scalable than simply replaying a static string; OS- or library-level modifications may be required.

3.3.3 OTP Schemes

One-time password (OTP) schemes generate short-lived credentials, often used as a second factor alongside conventional passwords. Depending on implementation (four follow), an attacker may aim to capture either the seed or a challenge-response pair.

OTP1: OTP mobile apps. This class (*e.g.*, Google Authenticator) generates OTPs to be manually typed into a user’s access device, using a combination of a locally-stored shared secret and either the current time (TOTP [10]) or a counter (HOTP [54]). With malware on the device, the attacker can capture the locally-stored shared secret; because this directly enables the attacker to compromise the account, OTP1 provides no resistance to mimicry, and is placed on the vertical axis of Fig. 2. Note: A variant of this class is a mobile app that stores a public-private key pair, *e.g.*, as used by Duo Security [55] and Google Prompt (a more recent iteration of Google two-step verification [56]). When the user types in their password, the server sends a cryptographic challenge to the mobile app. The mobile app then requests user consent, via simply tapping a button, to send the cryptographic response to the server to complete the user authentication process. The advantage of this approach compared to OTP1 is that it is more efficient to use. The disadvantages are that the mobile app requires an Internet connection, and there is a chance of the user accidentally consenting to a malicious authentication attempt.

OTP2: OTP USB tokens. This class (*e.g.*, FIDO U2F keys [57]) is similar to OTP1, but requires less effort since the user can press a button on the hardware token to automatically enter the OTP into a browser window. Assuming hardware tokens can resist malware, it becomes relatively challenging for an attacker to capture challenge-response pairs. Similar to OTP1, the attacker can thus target the seed, *i.e.*, hardware token theft. This gives the attacker direct account compromise, and therefore is placed on the vertical axis (*i.e.*, no horizontal distance from the origin) in Fig. 2.

OTP3: SMS OTP. The server sends a randomly-generated OTP (*i.e.*, no reliance on a shared secret/seed) to the user via SMS. Contrary to the previous two OTP classes, the seed here is only stored on the server, which makes it harder to capture. The attacker can however capture an OTP

in transit, *e.g.*, by exploring weaknesses in the cellular network [27], but will be required to use it before it expires (and before the user uses it). Since the time window for a successful attack is limited, OTP3 is placed further (horizontally) from the origin, in H2 on Fig. 2.

OTP4: E-mail OTP. The server sends a randomly-generated OTP (again no reliance on a shared secret/seed) to the user via e-mail. Because an attacker can capture an OTP, *e.g.*, via malware on the user’s machine, its mimicry-resistance is similar to OTP3.

3.3.4 PUFs

Physically Unclonable Functions (PUFs) are hardware modules (typically manufactured into silicon-based chips) that leverage an underlying unique physical structure to generate challenge-response pairs; ideal PUFs are impossible to clone, since the unique structure of each individual PUF results from manufacturing variation [9]. For our evaluation herein, we assume the use a challenge-response protocol (as described by Yu *et al.* [9]) in conjunction with PUFs that are built into the user’s device (as opposed to, *e.g.*, a USB-based token).

PUF1: Strong PUFs. These theoretically generate an endless supply of challenge-response pairs, allowing verifying servers to store any number of challenge-response pairs (*e.g.*, at user account creation) to be used for user authentication. Several implementations of strong PUFs are available [58]; *e.g.*, some PUFs generate challenge-response pairs by shining a laser on a scattering object at selected angles and points of incidence [59]. A major challenge in developing strong PUFs is to avoid susceptibility to model-building attacks that collect challenge-response pairs and apply machine-learning algorithms to build a mathematical model of the PUF [60]. PUF1 lies in V3, since the only security property it lacks from Table 1 is *Resilient-to-Physical-Theft*; device theft allows an attacker to directly impersonate the user (*i.e.*, defeat authentication), and therefore no mimicry-resistance is offered.

PUF2: Weak PUFs. These can only generate a limited number of challenge-response pairs. They are suitable for authentication when augmented with appropriate mechanisms, *e.g.*, when restricted to only responding to challenges sent from a single trusted verifier over an authenticated channel [9]; this limits an attacker’s ability to exhaustively capture all possible challenge/response pairs to mount a mimicry attack. PUF devices are commercially available, *e.g.*, for device authentication and cryptographic key storage [61]. PUF2 is placed within H2—it lacks *Resilient-to-Internal-Observation* from Table 1, but a successful mimicry attack can be mounted by building a mathematical model of the PUF only after capturing a sufficient number of challenge/response pairs (*e.g.*, via man-in-the-middle); however, since every PUF is unique, the mathematical model would be device-specific. Counterintuitively, PUF2 exhibits higher mimicry-resistance than PUF1—however, recall that the strength of a scheme is determined by both exposure- and mimicry-resistance. While PUF2 can be mimicked by a determined and targeted attacker even without device possession, a successful attack against PUF1 requires device possession (thereby obviating the need for mimicry), justifying PUF1’s higher placement within V3.

3.3.5 Sound-Proof

Sound-Proof [62] is an authentication scheme that determines whether the user’s smartphone and access device (*i.e.*, another device from which the user is authenticating to access their online account) are in close proximity. The access device records ambient sound from the microphone and transmits it to the smartphone (via the Internet), which compares with the sound recorded by its own microphone. A match indicates that both devices are in an identical noise environment, and therefore likely in close proximity, resulting in the smartphone sending a cryptographically-signed assertion to the web server to approve the user authentication.

We include Sound-Proof in our evaluation as it provides a degree of mimicry-resistance if a trusted channel is established between the smartphone and web server, *e.g.*, by means of a TPM (via a trusted execution environment and secure key storage). A TPM-augmented Sound-Proof would address attacks where, *e.g.*, the attacker can steal the key via root-privileged malware and use it to sign assertions. Note that while Sound-Proof was originally proposed and evaluated as a second factor alongside passwords, we place it on Fig. 2 as a single-factor scheme (without passwords)—Section 3.4 explains our rationale.

Although a TPM-based implementation may render it impractical for an attacker to steal the Sound-Proof app’s cryptographic key from the user’s smartphone and thus gain permanent access to the account, the attacker may still use a malicious app on the user’s device to relay recorded audio. For example, the attacker can reproduce the ambient sound from the environment in which the user’s smartphone is currently located, *e.g.*, by leveraging malware on the user’s smartphone or a nearby device to record and relay the ambient sound to the attacker. Sound-Proof is thus placed in H2 because this process is more sophisticated than simply replaying a static token. It is not placed in H1, since repeated (physical or internal) observations do not seem to help the attacker gain permanent access—the authors show [62] that the scheme is resilient to guessing attacks even when the attacker knows the user’s environment (*e.g.*, by using typical background noise from a Starbucks coffee shop), and that attacks are also made more difficult by requiring the sound files recorded by both devices to be timestamped (with NTP-synchronized clocks).

3.4 Further Insights

We briefly discuss further insights from our analysis of mimicry resistance, and the relative strengths of schemes plotted in Fig. 2.

Multi-factor authentication. Fig. 2 should be used to evaluate individual authentication schemes. To evaluate a multi-factor scheme, the scheme must be broken down to its constituent factors, each represented individually as an independent marker. Two independent markers can then be combined as a two-factor scheme as follows: the (x, y) position of the resultant (combined) scheme is the greater of both x -values and both y -values. That position should however be carefully interpreted. For example, although combining L4 with OTP2 would result in a marker in {H3, V3} (see Fig. 2), an attacker capable of physically stealing the device to defeat OTP2 will already be geographically co-located with the user, and thus need not expend any additional effort to defeat L4.

Lack of schemes in top-right corner of Fig. 2. None of the schemes analyzed have strong resistance to both mimicry and exposure—see the three empty squares in the top-right ({H2, V3}, {H3, V2}, {H3, V3}). Schemes in this region would strongly resist scalable attacks. This motivates combining complementary schemes, as discussed immediately above.

Bands versus markers. Instead of marker representation, schemes on Fig. 2 could be represented using bands (*i.e.*, lines/curves or shaded areas). For example, passwords could be represented by a vertical band from V1 (user-chosen passwords) to somewhere in V2 (system-assigned passwords that are resilient to guessing attacks, but still subject to leaks). Such representation can help identify how different implementations of a scheme may alter security. We chose the simpler representation as it is easier to interpret, and to avoid cluttering the chart with intersecting bands.

4 Comparative Evaluation

In Table 3, we evaluate selected classes of invisible and mimicry-resistant authentication schemes (selected baseline schemes are also included for comparison). We augment the original 25 usability, deployability, and security properties (benefits) of the UDS framework [1] with six new properties relevant to schemes that are invisible and have mimicry-resistance. These new properties were not present in the original UDS paper. The original UDS properties (see Appendix A) and our new properties defined below are italicized when referenced herein. The new properties are now described:

U9. No-False-Rejects is a usability benefit concerning authentication failures resulting from *system* error (*e.g.*, due to measurement error). False rejects arise due to *fuzzy* or non-binary matching functions employed by some mimicry-resistant and/or invisible authentication schemes, and are thus relevant to penalize (by withholding this benefit) schemes that suffer from false rejects (*e.g.*, measurement-based Internet geolocation methods). This differs from *Infrequent-Errors* (U7, Appendix A), wherein authentication may fail due to *user* action (*e.g.*, incorrectly typing a password) or attempts to authenticate under unusual circumstances (*e.g.*, from unexpected locations).

U10. Easy-to-Change-Credentials is a usability benefit for schemes where a user may easily change credentials (*e.g.*, in event of a server database leak). Intuitively, since invisible schemes require no user action upon login or account set-up (see Section 2), these schemes rely on remotely (*cf. device-to-web*) observing habitual user/device attributes and/or behaviours, and transparently verifying these upon login. Changing these credentials (regardless of the reason for changing) is likely to impose user burden, *i.e.*, changing physical habits and/or habitual behaviors. This new property (U10) allows appropriate penalization for such schemes. Note that in contrast, *Easy-Recovery-From-Loss* (U8) reflects how easily a user can recover from a credential loss (*e.g.*, forgotten password). Credential loss requires a fallback mechanism to verify the user’s identity; changing credentials does not, since the user remains in possession of valid credentials. *Easy-to-Change-Credentials* is inherent to the authentication mechanism itself, whereas *Easy-Recovery-From-Loss* may also depend on the fallback mechanism.

M1. No-False-Accepts is a mimicry-related (hence, the ‘M’ index) security benefit of schemes that have a sufficiently large credential space and/or measurement precision such that two different sets of credentials (*e.g.*, iris scans from two different individuals) are always distinguishable. False accepts may include both non-malicious users and attackers mistaken for legitimate users, *e.g.*, due to close proximity in a geolocation scheme, or a device fingerprint similar (within a margin of error) to that of a legitimate user. Similarly to false rejects, the fuzzy nature of many mimicry-resistant schemes explored in Section 3 also introduces the possibility of false accepts, justifying the importance of this new property (M1). Note that false accepts exclude attacks on the integrity of the authentication system (these are covered by other security properties), such as manipulating delay measurements to spoof a location or tampering with client-side code to spoof a device fingerprint. For example, a location-based scheme lacks this benefit if it is susceptible to colocation attacks, *i.e.*, where the attacker travels and colocates himself with the user in a highly targeted attack.

M2. Resilient-to-Infrequent-Capture-or-Intercept is a security benefit of schemes in which credentials are not static, but change relatively slowly, *e.g.*, personal knowledge questions based on the user’s account activity, such as recent transactions. Such non-static credentials, when compromised, limit the duration for which the attacker retains account access. Schemes with horizontal distance from the origin in Fig. 2 (*i.e.*, placed in H1, H2, or H3) offer this benefit; see Section 3.1.2 for further discussion of schemes that satisfy this benefit, under the heading **H1**.

M3. Verifies-Non-Static-Information is a security benefit of challenge-response based schemes where the server issues a new challenge per authentication; thus an attack that simply captures and uses a static string should not succeed. This benefit is offered by schemes where the server verifies a client’s ability to generate a valid response to a challenge, where the response is based on non-static information as opposed to, *e.g.*, a static secret. For example, Sound-Proof [62] (see Section 3.3.5) verifies that the ambient sound (which is non-static information) recorded by the user’s smartphone matches the ambient sound recorded from the user’s access device. OTP schemes may offer this benefit, if (1) the OTPs are not generated based on a static client-stored seed (which would be at risk of being captured by an attacker), and (2) the OTPs expire either upon first use or within a short time frame (*e.g.*, 2 minutes), thereby substantially limiting the window for successful replay. Schemes within H2 or H3 in Fig. 2 offer this benefit; see Section 3.1.2 for further discussion of schemes that satisfy this benefit, under the heading **H2**.

M4. Resilient-to-Spoofing is a security benefit of schemes that leverage measurement techniques (*e.g.*, hardware or network-based) that are impractical for an attacker to defeat at scale. For example, CPV [48] is *Resilient-to-Spoofing*, since the measurements cannot be manipulated to make an attacker appear to be in a different location (*i.e.*, that of the victim user). Only schemes in H3 of Fig. 2 offer this benefit; see Section 3.1.2 for further discussion of schemes that satisfy this benefit, under the heading **H3**.

Table 3 cells corresponding to benefits M2, M3, and M4 are populated based on the analysis in Section 3.3 (and visualized in Fig. 2). The following sections evaluate the schemes with

respect to the remaining benefits.

4.1 Impact of Account Recovery on Security

In practice, authentication schemes are paired with backup mechanisms to help users recover account access if they lose their credentials. The recovery mechanism should not be easier to defeat than the primary scheme, but may sacrifice usability since it is used less frequently.

For conventional password-based authentication, e-mail based password reset is the standard recovery mechanism. Its security relies on the implicit assumption that the user’s e-mail account is at least as well-secured as any systems that rely on it for password reset. It is ideally expected that users choose a stronger password for primary e-mail accounts; security-conscious users may also use two-factor authentication. E-mail based recovery can also be used for the geolocation and device fingerprinting schemes evaluated herein; alternatives such as SMS-based OTP are also suitable. For example, in a geolocation-based scheme, a user wishing to login from a new location could be e-mailed a link through which they could confirm their new location.

4.2 Evaluation of Stand-alone Schemes

Table 3 summarizes our evaluation of four main categories of schemes: geolocation, device fingerprinting, OTPs and PUFs as both stand-alone authentication schemes and in combination with passwords, based on the augmented UDS criteria from Section 4. Password authentication is included as reference; OTP schemes are included as they are widely used in combination with passwords. While additional schemes could have been included, those selected were chosen as examples to demonstrate the comparative framework. Each row in the table corresponds to an authentication scheme, and each column to a benefit; a cell with a bullet represents a benefit offered by the scheme, an empty circle represents a benefit partially provided, and an empty cell indicates that the benefit is not provided.

4.2.1 Web Passwords

For the new properties, web passwords provide *No-False-Rejects* since a correctly-typed password will never be rejected, *Easy-to-Change-Credentials* since passwords can be easily changed, and *No-False-Accepts* since an exact match is required. However, passwords lack *Resilient-to-Infrequent-Capture-or-Intercept*, *Verifies-Non-Static-Information* and *Resilient-to-Spoofing*, since it is trivial to replay a captured password.

4.2.2 Location-based Schemes

Since L1-L4 in Table 3 are invisible to the user, they provide most of the usability benefits. L3 (measurement-based) and L4 (location verification) may sometimes take longer than conventionally considered convenient, thus providing only a partial *Efficient-to-Use* benefit. Additionally, all but L1 (GPS/WPS) may miscalculate the location and falsely reject users in some cases, and therefore do not fulfill *No-False-Rejects*. Although L1 (GPS/WPS) may sometimes result in a small error in location calculation, the calculated location will generally remain in the same city, and thus gets a bullet. None of L1-L4 are *Easy-to-Change-Credentials*, since changing credentials would require the legitimate user to change their location.

they require server-side changes, but are *Browser-Compatible* (no client-side changes needed). L3 is partially *Mature* since there are indications that it is being used in practice [63]; L4 is not *Mature*. Finally, *Non-Proprietary* variations for all of L1-L4 are available.

L1-L4 are largely similar to each other in terms of security properties offered, with the exception of *Replay-Resistance* and *Spoofing-Resistance*. Due to the small guessing space, L1-L4 are not *Resilient-to-Throttled-Guessing-Attempts* and *Resilient-to-Unthrottled-Guessing-Attempts*. Since location information must be stored server-side, they are susceptible to the same exposure threats as passwords, namely *Resilient-to-Physical-Observation*, *Resilient-To-Internal-Observation*, *Resilient-to-Leaks-from-Other-Verifiers*, and *Resilient-to-Phishing*. *No-Trusted-Third-Party* is provided by L3-L4 if the website runs their own geolocation infrastructure (e.g., deploys their own verifier servers), but not by L2 (IP address tabulation) which typically rely on a third-party service provider, and partially by L1 since GPS does not require a third-party service whereas WPS may. None of L1-L4 are fully *No-False-Accepts*, since other users (both legitimate users and attackers) could be residing near the legitimate user, and thus could be indistinguishable to the server. They are not *Resilient-to-Targeted-Impersonation*, since they are susceptible to targeted colocation attacks. They are not *Requiring-Explicit-Consent*, since they are invisible to the user—except for L1, since a GPS/WPS location request may trigger a browser permission prompt to the user, but if the user allows the browser to remember the preference, no future prompts are displayed. They are all partially *Unlinkable*, since leaking server-stored user location information may narrow the search space for linking together multiple accounts across websites (the attacker may still need to collect additional information).

4.2.3 Device fingerprinting

FP1-FP6 deliver most usability benefits, since they are invisible to the user. However, as significant changes in the device configuration (e.g., software or hardware upgrade) may substantially change a device fingerprint, they do not provide *No-False-Rejects*. They are *Easy-Recovery-from-Loss* since e-mail based recovery can be used, as with password authentication. They are not *Easy-to-Change-Credentials*, since changing a device fingerprint may require the user to obtain a different device.

For deployability, device fingerprinting schemes are *Accessible*, since they do not require any explicit user action. They are *Negligible-Cost-Per-User*, since the cost of implementation is essentially independent of the number of users. They are not *Server-Compatible*, as server-side implementation is required, but are *Browser-Compatible*. As device fingerprinting has been used for anti-fraud applications [64], [65], but not widely for user authentication, we consider FP1, FP2, FP5, and FP6 partially *Mature*; FP3 and FP4 are not, as they have been demonstrated academically but are not used in practice to our knowledge. FP1-FP6 are generally available via *Non-Proprietary* implementations.

Many of the security properties are shared across FP1-FP6. They lack *No-False-Accepts*, since users that own identical devices (or share the same device) may be indistinguishable from each other (and for techniques such as

clockskew, the overall credential space is not large enough to rule out collisions). They lack *Resilient-to-Unthrottled-Guessing-Attempts*, since none of the data collected thus far to our knowledge indicates that the overall distribution of device fingerprints would offer distinguishability of more than about 30 bits [7], [18], [51], [66], [67], [68]. They lack *Resilient-to-Internal-Observation* and *Resilient-to-Phishing*, since an attacker may collect device information by running their own device fingerprinting scripts via XSS attacks or phishing websites.¹⁰ They lack *Resilient-to-Leaks-from-Other-Verifiers*, since a user’s device fingerprint will be similar across websites (varying only based on the particular fingerprinting techniques that each website uses, and the method of storage used), thereby leaking information that can be used to attack users’ accounts on different websites. They are each individually partially *Unlinkable* since there may well be multiple users with colliding device fingerprints—the likelihood of this may diminish substantially when combining multiple techniques, however. *Requiring-Explicit-Consent* is not provided, since device fingerprinting is invisible to the user.

The remaining security properties differ across FP1-FP6, as follows. FP3 and FP4 are *Resilient-to-Physical-Observation* and *Resilient-to-Targeted-Impersonation*, since they rely on device-specific manufacturing variations (even across devices of the same model) and therefore can only be determined via measurement; the remaining schemes only partially fulfill these properties since an attacker that visually observes a user’s device may obtain or mimic the same device model. FP1 and FP2 have been shown to provide enough distinguishing information to provide *Resilient-to-Throttled-Guessing-Attempts*—to our knowledge the remainder have not, but may collectively provide it if combined together [51].

4.2.4 OTP Schemes

For usability, OTP1-OTP4 are *Memorywise-Effortless* since they do not require the user to memorize anything; they are *Scalable-for-Users* since they allow the user to set up multiple accounts without impacting usability. OTP4 (e-mail) is *Nothing-to-Carry*, since the user just needs to be able to access their e-mail account, whereas OTP1 (mobile app) and OTP3 (SMS) are partially *Nothing-to-Carry*, assuming users are typically in possession of their mobile phones at all times; OTP2 (USB token) does not fulfill this property. All schemes are *Easy-to-Learn*, but *Inefficient-to-Use* since they require the user to type an extra code. They all provide *Infrequent-Errors* since typos are much less likely with a short 6-digit numerical code, and *No-False-Rejects*. OTP1 and OTP2 lack *Easy-Recovery-From-Loss* since the secret is stored on the device/token; OTP3 partially provides it since there is no secret stored on the phone; OTP4 provides it, since a back-up e-mail address can be used. They all offer *No-False-Rejects*, since a valid OTP will not be rejected. *Easy-to-Change-Credentials* is provided in full by all; a user that needs to switch to a new OTP authenticator could, e.g., login with their old authenticator, and then register the new one.

For deployability, all of OTP1-OTP4 are partially *Accessible*, as blind users would require screen reading software to read the OTP code. Only OTP1 and OTP4 are *Negligible-Cost-per-User*, since SMS incurs a per-message cost to the

¹⁰ Collecting device information helps the attacker, but alone is not enough to defeat mimicry-resistant schemes such as L4.

server,¹¹ and USB tokens have a per-user cost. All schemes except for OTP2 (USB tokens) are *Browser-Compatible*; in addition to requiring browser support (at this time, only Google Chrome supports FIDO U2F), OTP2 tokens require a hardware interface (*e.g.*, USB-A, USB-C, NFC) that is compatible with the user’s devices—while a variety of hardware tokens are available, they typically only possess one or two interfaces. All of OTP1-OTP4 are *Mature* and with *Non-Proprietary* implementations available.

For security, all of OTP1-OTP4 are *No-False-Accepts*, *Resilient-to-Physical-Observation*, *Resilient-to-Throttled-Guessing*, and *Resilient-to-Unthrottled-Guessing*. OTP1-OTP2 and OTP4 are *Resilient-to-Targeted-Impersonation*, but OTP3 only partially provides this benefit—attackers may transfer the victim’s phone number to a new device via a social engineering attack on the user’s mobile operator [69]; the ease of conducting such an attack is subject to the security measures put in place by the mobile operator. They are *Resilient-to-Leaks-from-Other-Verifiers* and *Resilient-to-Phishing*. Only OTP2 (USB tokens) is *Resilient-to-Internal-Observation*—OTP1 is susceptible to theft of the shared seed via malware, and OTP3/OTP4 are susceptible to malware- and network-based [27] capture attacks.¹² Only the e-mail mechanism (OTP4) is *Resilient-to-Physical-Theft*. They all provide *No-Trusted-Third-Party* and *Requiring-Explicit-User-Consent*. OTP1 and OTP2 are *Unlinkable* since a unique key is used for each website; phone numbers (OTP3) are linkable across sites; e-mail addresses (OTP4) are partially *Unlinkable*, since users may freely create different e-mail aliases.

4.2.5 PUFs

For usability, PUF1 and PUF2 lack *Nothing-to-Carry* and *Easy-Recovery-From-Loss*, since they are tied to the device. PUF1 is *Easy-to-Change-Credentials* since virtually endless challenge/response pairs can be generated; PUF2 is not, since the hardware needs replacement if subjected to a model-building attack. All other usability properties are fulfilled, since no user effort is required.

For deployability, PUF1 and PUF2 are not *Negligible-Cost-per-User*, since hardware would need to be deployed for each user, and they are not *Server-Compatible*, *Browser-Compatible*, or *Mature*. *Non-Proprietary* designs for PUF2 are available [9], but not for PUF1 (to our knowledge).

For security, assuming a large enough space of challenge-response pairs, PUF1 and PUF2 fulfill *No-False-Accepts*, *Resilient-to-Throttled-Guessing*, and *Resilient-to-Unthrottled-Guessing*. PUF1 provides *Resilient-to-Internal-Observation* and *Resilient-to-Phishing* since challenge/response pairs can be exposed without consequence; PUF2 does not provide the former, since it is susceptible to model-building attacks, but it provides the latter since it only responds to challenges

11. Bulk SMS rates can be on the order of a penny each (varying by country). The total cost will vary based on whether the user base is small or in the millions, and on average login frequency.

12. Since the UDS framework combines both network- and device-based eavesdropping/interception into a single property (*Resilient-to-Internal-Observation*), OTP1 and OTP3-OTP4 share the same *y*-coordinate in Fig. 2, though OTP3 should be lower since it appears to be susceptible to a larger subset of network-based interception attacks [70].

originating from a verifier server. Both are *Resilient-to-Leaks-from-Other-Verifiers*. PUF1 and PUF2 are neither *Resilient-to-Physical-Theft* nor *Requiring-Explicit-Consent*. PUF1 has *No-Trusted-Third-Party* and *Unlinkable* since every website can store a separate set of challenge-response pairs, but PUF2 lacks these properties since it typically requires a single server to verify challenge-response pairs (see Section 3.3.4).

4.2.6 Sound-Proof

For usability, Sound-Proof provides *Memorywise-Effortless* since there is no secret for the user to memorize. It is partially *Scalable-for-Users*, since only a single device is required, but it appears that under the current architecture each service requires a separate app. It partially provides *Nothing-to-Carry*, since users with a smartphone will typically carry it with them at all times. It is *Easy-to-Learn*, *Efficient-to-Use*, and *Infrequent-Errors* as (aside from installing a smartphone app at setup time) it does not require any user effort to use. It partially provides *No-False-Rejects* and *No-False-Accepts*—the false accept and false reject rates are tunable via a threshold value, and the authors show that the intersection between the two (the equal error rate) is about 0.02%. It lacks *Easy-Recovery-From-Loss*, as the consequence of losing the phone is similar to that of OTP1. For deployability, it lacks *Server-Compatible* and *Mature*.

For security, Sound-Proof partially provides *Resilient-to-Physical-Observation*, since being able to record sound in the vicinity of the user can break the scheme (it also lacks *Resilient-to-Targeted-Impersonation*, for this reason). It is *Resilient-to-Throttled-Guessing-Attempts* and *Resilient-to-Unthrottled-Guessing-Attempts*. It is not *Resilient-to-Internal-Observation* or *Resilient-to-Theft*, since the scheme can be broken by malware on the phone or by device theft. It is *Resilient-to-Leaks-From-Other-Verifiers* (since there is no static secret to be stored), *Resilient-to-Phishing*, and *No-Trusted-Third-Party*. It is not *Requiring-Explicit-Consent* since it requires no user action. It is *Unlinkable*, since different services can use their own apps on the user’s smartphone.

5 Evaluation of Combined Schemes

The evaluation above naturally suggests that schemes offering some degree of mimicry resistance may complement passwords by adding a new security dimension; and invisible schemes such as device fingerprinting and geolocation do so without further burdening users. None of the latter (invisible) schemes, however, seem suitable as a sole mechanism for *user-to-web* authentication because either the invisibility aspect makes them limited to *device-to-web* authentication (see Section 2), or their security space is so small that they lack the ability to uniquely identify users (*e.g.*, suffer from false accepts). This motivates exploring the resultant benefits when these schemes are combined with passwords. This is also useful since some password usability drawbacks might be ameliorated in the medium to long-term by use of browser-based and/or stand-alone password managers.

When combining two authentication schemes, both sets of credentials should be correct for the user to gain access. This then strengthens security. Ideally, the implementation should limit any partial feedback, which might be beneficial to an attacker using a divide-and-conquer strategy to independently

defeat each scheme individually. When password authentication is combined with an invisible scheme, if the correct password is provided but the supporting invisible scheme fails, the server can immediately fall back to a backup scheme such as e-mail OTP. Otherwise, allowing retry authentication attempts benefits attackers, but not legitimate users (except in cases of system error where a retry might succeed, since invisible credentials are not entered by the user). In the event of an attack, this serves to notify the user that their password has been compromised; otherwise, it gives the user an opportunity to reset their supporting authentication mechanism (e.g., by registering a new location or new device).

In Table 3, a combined authentication scheme as described above only inherits the intersection of the usability and deployability benefits, and the union of the security benefits—except for *No-Trusted-Third-Party* and *Unlinkable*.

The bottom half of Table 3 evaluates the schemes from the top half when used in combination with passwords. The evaluation suggests that either geolocation or device fingerprinting can be combined with passwords to improve security properties, namely *Resilient-to-Physical-Observation*, *Resilient-to-Throttled-Guessing*, *Resilient-to-Infrequent-Capture-or-Intercept*, *Verifies-Non-Static-Information*, and/or *Resilient-to-Spoofing*. To further improve security, geolocation can be combined with one or more device fingerprinting schemes. While combining schemes, select those that maximize the resulting overall benefits, e.g., combining L4 with FP2 and/or FP4. As a downside, this would increase the probability of false rejects, assuming the failure of any scheme causes authentication failure. PUFs provide many security benefits, suggesting that they could be sufficient as a sole authentication mechanism. However, their major drawback is that they lack most deployability benefits (in the web authentication context).

The usability drawbacks observed for the combined schemes in Table 3 are in *Efficient-to-Use* for schemes that introduce a perceptible delay in authenticating the user, *No-False-Rejects* primarily due to variations in measurement or fuzzy matching functions, and *Easy-to-Change*. The deployability drawbacks are in *Server-Compatible* and *Mature*, for schemes that have not yet been widely used in practice.

6 Concluding Remarks

While it is not our main objective to provide recommendations—rather the primary focus is in providing the first in-depth treatment of the mimicry-resistance dimension in web authentication schemes and taking this dimension into account in a comparative evaluation framework—we nonetheless give some guidelines which we hope can help advance web authentication research. Our evaluation shows that leveraging schemes with some degree of mimicry-resistance is advantageous to security because they could reduce attack scalability, and can achieve this with minimal impact on usability if the scheme is also invisible. In multi-factor authentication, combining highly usable schemes often enhances the resultant (combined) benefits, even when these offer relatively few security benefits individually since security benefits are often additive (complementary). While multi-factor authentication typically requires all factors to pass in order to grant access, account recovery schemes

provide an alternative authentication path that attackers can exploit to bypass the primary authentication mechanism. Therefore, while combining multiple factors strengthens security, the security level of any alternative attack path (e.g., account recovery mechanisms) must be correspondingly strengthened.

To the best of our knowledge, this work is the first to systematically investigate the mimicry-resistance aspect of *user-to-web* and *device-to-web* authentication schemes. Segmenting the mimicry-resistance dimension into levels of replayability and spoofability helps differentiate the ability of various schemes to resist mimicry attacks. Evaluating schemes using the new dimension, we found that most web authentication schemes have limited to no mimicry resistance as they rely fundamentally on the knowledge of some secret that (once exposed) can be directly used by the attacker for account break-in.

To incorporate mimicry-resistance into an evaluation framework, we add mimicry-resistance properties and explore how mimicry-resistance offers a second dimension orthogonal to exposure-resistance. Plotting these dimensions (Fig.2) gives a novel representation of attack scalability, where a scheme’s ability to resist scalable attacks is conveyed by its distance from the origin. The new framework thus allows for a more comprehensive evaluation of web authentication schemes by their ability to resist both exposure and mimicry, and helps visualize relative resistance to attack scalability, represented as a function of both components. Positioning many representative web authentication schemes on the 2D chart highlights that none explored to date occupy the top-right corner (where schemes highly resistant to scalable attacks would be placed).

Using the new framework, we explore mechanisms that offer mimicry-resistant web authentication, and find several such schemes such as some variations of device fingerprinting, several OTP-based schemes, robust Internet geolocation methods, and certain variations of Physically Unclonable Functions (PUFs). Our analysis highlights the degree to which different schemes within the same family of schemes (e.g., different implementations of geolocation) can vary in mimicry-resistance, and our evaluation framework facilitates scheme selection (including from within a family) for maximizing security benefits. For example, in Table 3, some geolocation schemes (e.g., L4) are less susceptible to spoofing than others (e.g., L1).

A challenging obstacle in reinforcing password authentication with additional schemes is doing so without deployability or usability penalties [2]. To evaluate the benefits of combining these four approaches with passwords, we augment the UDS framework with new usability and security properties to reflect mimicry resistance. Using the augmented UDS framework (Table 3) to evaluate the resultant combined schemes revealed that *invisible* techniques—those neither requiring user involvement at set-up nor at login times such as geolocation and device fingerprinting—offer substantial usability advantages and have better deployability than PUFs and OTP-based schemes.

Our work highlights the advantages of utilizing mimicry-resistant techniques in web authentication, motivating further exploration into invisible and mimicry-resistant techniques, and providing directions on how to explore and evaluate the

ability of schemes to resist scalable attacks.

7 Acknowledgements

We thank the anonymous referees for their helpful feedback. The third author acknowledges funding from the Natural Sciences and Engineering Research Council of Canada (NSERC) for both his Canada Research Chair in Authentication and Computer Security, and a Discovery Grant.

References

- [1] J. Bonneau, C. Herley, P. van Oorschot, and F. Stajano, “The quest to replace passwords: A framework for comparative evaluation of web authentication schemes,” in *IEEE Symposium on Security and Privacy (S&P)*, 2012, pp. 553–567.
- [2] C. Herley and P. van Oorschot, “A research agenda acknowledging the persistence of passwords,” *IEEE Security & Privacy*, vol. 10, no. 1, pp. 28–36, 2012.
- [3] J. Bonneau, C. Herley, P. van Oorschot, and F. Stajano, “Passwords and the evolution of imperfect authentication,” *Communications of the ACM*, vol. 58, no. 7, pp. 78–87, 2015.
- [4] D. E. Denning and P. F. MacDoran, “Location-based authentication: Grounding cyberspace for better security,” *Computer Fraud & Security*, vol. 1996, no. 2, pp. 12–16, 1996.
- [5] D. Wagner and P. Soto, “Mimicry attacks on host-based intrusion detection systems,” in *ACM CCS*, 2002, pp. 255–264.
- [6] H. Khan, A. Atwater, and U. Hengartner, “A comparative evaluation of implicit authentication schemes,” in *RAID*. Springer, 2014, pp. 255–275.
- [7] P. Eckersley, “How unique is your web browser?” in *Privacy Enhancing Technologies Symposium (PETS)*. Springer, 2010, pp. 1–18.
- [8] Z. Dong, R. D. Perera, R. Chandramouli, and K. Subbalakshmi, “Network measurement based modeling and optimization for IP geolocation,” *Computer Networks*, vol. 56, no. 1, pp. 85–98, 2012.
- [9] M.-D. M. Yu and S. Devadas, “Pervasive, Dynamic Authentication of Physical Items,” *Communications of the ACM*, vol. 14, no. 6, 2017.
- [10] D. M’Raihi, S. Machani, M. Pei, and J. Rydell, “TOTP: Time-Based One-Time Password Algorithm,” RFC 6238 (Informational), Internet Engineering Task Force, May 2011. [Online]. Available: <http://www.ietf.org/rfc/rfc6238.txt>
- [11] M. Jakobsson, E. Shi, P. Golle, and R. Chow, “Implicit authentication for mobile devices,” in *USENIX Workshop on Hot Topics in Security (HotSec)*, 2009.
- [12] M. Frank, R. Biedert, E. Ma, I. Martinovic, and D. Song, “Touchalytics: On the applicability of touchscreen input as a behavioral biometric for continuous authentication,” *IEEE TIFS*, vol. 8, no. 1, pp. 136–148, 2013.
- [13] T. Feng, J. Yang, Z. Yan, E. M. Tapia, and W. Shi, “TIPS: Context-aware implicit user identification using touch screen in uncontrolled environments,” in *ACM HotMobile*, 2014.
- [14] S. Mare, A. M. Markham, C. Cornelius, R. Peterson, and D. Kotz, “ZEBRA: Zero-effort bilateral recurring authentication,” in *IEEE Symposium on Security and Privacy (S&P)*, 2014.
- [15] A. Acar, H. Aksu, A. S. Uluagac, and K. Akkaya, “WACA: Wearable-Assisted Continuous Authentication,” in *IEEE Security and Privacy Workshops (SPW)*, 2018.
- [16] H. Aksu, A. S. Uluagac, and E. Bentley, “Identification of wearable devices with bluetooth,” *IEEE Transactions on Sustainable Computing (T-SUSC)*, 2018, (Early Access).
- [17] S. V. Radhakrishnan, A. S. Uluagac, and R. Beyah, “GTID: A technique for physical device and device type fingerprinting,” *IEEE Transactions on Dependable and Secure Computing (TDSC)*, vol. 12, no. 5, pp. 519–532, 2015.
- [18] H. Bojinov, Y. Michalevsky, G. Nakibly, and D. Boneh, “Mobile Device Identification via Sensor Fingerprinting,” in *arXiv preprint arXiv:1408.1416 [cs.CR]*, 2014.
- [19] M. Mannan and P. van Oorschot, “Leveraging personal devices for stronger password authentication from untrusted computers,” *Journal of Computer Security*, vol. 19, no. 4, pp. 703–750, 2011.
- [20] J. Brainard, A. Juels, R. L. Rivest, M. Szydlo, and M. Yung, “Fourth-factor authentication: somebody you know,” in *ACM CCS*, 2006, pp. 168–178.
- [21] D. Florêncio, C. Herley, and P. van Oorschot, “An Administrator’s Guide to Internet Password Research,” in *USENIX LISA*, 2014, pp. 35–52.
- [22] N. Heninger, Z. Durumeric, E. Wustrow, and J. A. Halderman, “Mining Your Ps and Qs: Detection of Widespread Weak Keys in Network Devices,” in *USENIX Security*, 2012.
- [23] Y. Zhang, F. Monrose, and M. K. Reiter, “The security of modern password expiration: An algorithmic framework and empirical analysis,” in *ACM CCS*, 2010, pp. 176–186.
- [24] S. Chiasson and P. van Oorschot, “Quantifying the security advantage of password expiration policies,” *Designs, Codes and Cryptography*, vol. 77, no. 2-3, pp. 401–408, 2015.
- [25] D. Weinshall, “Cognitive authentication schemes safe against spyware,” in *IEEE Symposium on Security and Privacy (S&P)*, 2006.
- [26] PassWindow, “Two factor authentication cards for secure online login,” <https://www.passwindow.com>, Aug 2017.
- [27] D. Goodin, “Thieves drain 2FA-protected bank accounts by abusing SS7 routing protocol,” <https://arstechnica.com/security/2017/05/thieves-drain-2fa-protected-bank-accounts-by-abusing-ss7-routing-protocol/>, May 2017.
- [28] P. van Oorschot and S. G. Stubblebine, “On identity theft and a countermeasure based on digital uniqueness and location cross-checking,” in *Financial Cryptography and Data Security*, 2005, see also extended version: Technical report TR-05-12, School of Computer Science, Carleton University, Canada.
- [29] M. Jakobsson, E. Shi, P. Golle, and R. Chow, “Implicit authentication for mobile devices,” in *USENIX HotSec*, 2009.
- [30] I. Poese, S. Uhlig, M. A. Kaafar, B. Donnet, and B. Gueye, “IP geolocation databases: Unreliable?” *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 2, pp. 53–56, 2011.
- [31] S. Nettle, S. O’Neil, and P. Lock, “PassWindow: A new solution to providing second factor authentication,” *VEST Corporation*, 2009.
- [32] R. Jhawar, P. Inglesant, N. Courtois, and M. A. Sasse, “Make mine a quadruple: Strengthening the security of graphical one-time PIN authentication,” in *IEEE Network and System Security (NSS)*, 2011, pp. 81–88.
- [33] N. J. Hopper and M. Blum, “Secure human identification protocols,” in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2001, pp. 52–66.
- [34] S. L. Smith, “Authenticating users by word association,” *Computers & Security*, vol. 6, no. 6, pp. 464–470, 1987.
- [35] P. Golle and D. Wagner, “Cryptanalysis of a cognitive authentication scheme,” in *IEEE Symp. Security & Privacy*, 2007.
- [36] D. Recordon and D. Reed, “OpenID 2.0: a platform for user-centric identity management,” in *ACM workshop on Digital identity management*, 2006, pp. 11–16.
- [37] D. P. Kormann and A. D. Rubin, “Risks of the Passport single signon protocol,” *Computer Networks*, vol. 33, no. 1, pp. 51–58, 2000.
- [38] M. Hanson, D. Mills, and B. Adida, “Federated browser-based identity using email addresses,” in *W3C Workshop on Identity in the Browser*, 2011.
- [39] S. Chiasson, E. Stobert, A. Forget, R. Biddle, and P. van Oorschot, “Persuasive cued click-points: Design, implementation, and evaluation of a knowledge-based authentication mechanism,” *IEEE TDSC*, vol. 9, no. 2, pp. 222–235, 2012.
- [40] M. L. Damiani, C. Silvestri, and E. Bertino, “Fine-grained cloaking of sensitive positions in location-sharing applications,” *IEEE Pervasive Computing*, vol. 10, pp. 64–72, 2011.
- [41] A. Popescu, “Geolocation API specification,” *World Wide Web Consortium, Candidate Recommendation CR-geolocation-API-20100907*, vol. 201, 2010.
- [42] J. A. Muir and P. van Oorschot, “Internet geolocation: Evasion and counterevasion,” *ACM Computing Surveys*, vol. 42, no. 1, 2009.
- [43] S. Laki, P. Mátray, P. Hága, T. Sebök, I. Csabai, and G. Vattay, “Spotter: A model based active geolocation service,” in *IEEE INFOCOM*, 2011, pp. 3173–3181.
- [44] B. Gueye, A. Ziviani, M. Crovella, and S. Fdida, “Constraint-based geolocation of Internet hosts,” *IEEE/ACM Trans. Networking*, vol. 14, no. 6, pp. 1219–1232, 2006.
- [45] D. Li, J. Chen, C. Guo, Y. Liu, J. Zhang, Z. Zhang, and Y. Zhang, “IP-geolocation mapping for moderately connected

- Internet regions,” *IEEE Trans. Parallel and Distributed Systems, TPDS*, vol. 24, no. 2, pp. 381–391, 2013.
- [46] A. Abdou, A. Matrawy, and P. van Oorschot, “Accurate Manipulation of Delay-based Internet Geolocation,” in *ACM AsiaCCS*, 2017, pp. 887–898.
- [47] P. Gill, Y. Ganjali, B. Wong, and D. Lie, “Dude, where’s that IP?: circumventing measurement-based IP geolocation,” in *USENIX Security*, 2010.
- [48] A. Abdou, A. Matrawy, and P. van Oorschot, “CPV: Delay-based Location Verification for the Internet,” *IEEE TDSC*, vol. 14, no. 2, pp. 130–144, 2017.
- [49] S. Park, J. N. Yoon, C. Kang, K. H. Kim, and T. Han, “TGVisor: a tiny hypervisor-based trusted geolocation framework for mobile cloud clients,” in *IEEE International Conference on Mobile Cloud Computing, Services, and Engineering (MobileCloud)*, 2015, pp. 99–108.
- [50] N. O. Tippenhauer, C. Pöpper, K. B. Rasmussen, and S. Capkun, “On the requirements for successful GPS spoofing attacks,” in *ACM CCS*, 2011, pp. 75–86.
- [51] F. Alaca and P. van Oorschot, “Device fingerprinting for augmenting web authentication: classification and analysis of methods,” in *ACM ACSAC*, 2016, pp. 289–301.
- [52] E. Bursztein, A. Malyshev, T. Pietraszek, and K. Thomas, “Picasso: Lightweight device class fingerprinting for web clients,” in *ACM SPSM*, 2016, pp. 93–102.
- [53] C. Arackaparambil, S. Bratus, A. Shubina, and D. Kotz, “On the Reliability of Wireless Fingerprinting Using Clock Skews,” in *ACM WiSec*, 2010.
- [54] D. M’Raihi, M. Bellare, F. Hoornaert, D. Naccache, and O. Ranen, “HOTP: An HMAC-Based One-Time Password Algorithm,” RFC 4226 (Informational), Internet Engineering Task Force, Dec. 2005. [Online]. Available: <http://www.ietf.org/rfc/rfc4226.txt>
- [55] Duo Security, “Secure Two-Factor Authentication App,” <https://duo.com/product/trusted-users/two-factor-authentication/duo-mobile>.
- [56] R. Amadeo, “Google makes it much easier to use 2FA on your account,” <https://arstechnica.com/gadgets/2016/06/google-new-two-factor-authentication-system-tap-yes-to-log-in/>, June 2016.
- [57] Nitrokey, “USB Dongle Auth List,” Available at <http://www.dongleauth.info/dongles/>, 2017, Accessed: 2017-05-19.
- [58] U. Ruhrmair and D. E. Holcomb, “PUFs at a glance,” in *Design, Automation and Test in Europe (DATE)*. IEEE, 2014, pp. 1–6.
- [59] R. Pappu, B. Recht, J. Taylor, and N. Gershenfeld, “Physical one-way functions,” *Science*, vol. 297, no. 5589, pp. 2026–2030, 2002.
- [60] J. Delvaux, D. Gu, D. Schellekens, and I. Verbauwhede, *Secure Lightweight Entity Authentication with Strong PUFs: Mission Impossible?* Springer, 2014, pp. 451–475.
- [61] I. ID, Available at <https://www.intrinsic-id.com/>, 2017, Accessed: 2017-05-14.
- [62] N. Karapanos, C. Marforio, C. Soriente, and S. Capkun, “Sound-Proof: Usable Two-Factor Authentication Based on Ambient Sound,” in *USENIX Security*, 2015, pp. 483–498.
- [63] Nanjee, <http://www.nanjee.net/>, Nanjee, Inc., 2017.
- [64] The 41st Parameter, “The 41st parameter announces new real-time product as world’s first standard for online PC identification,” <http://www.the41.com/buzz/announcements/41st-parameter-announces-new-real-time-product-worlds-first-standard-online-pc-0>, Jul. 2006, Accessed: 2017-05-14.
- [65] Maxmind Developer Site, “Device tracking add-on for min-fraud and proxy detection services,” <http://dev.maxmind.com/minfraud/device/>, 2017, Accessed: 2017-05-14.
- [66] P. Laperdrix, W. Rudametkin, and B. Baudry, “Beauty and the Beast: Diverting modern web browsers to build unique browser fingerprints,” in *IEEE Symposium on Security and Privacy (S&P)*, 2016.
- [67] S. Englehardt and A. Narayanan, “Online tracking: A 1-million-site measurement and analysis,” in *ACM CCS*, 2016, pp. 1388–1401.
- [68] T. Kohno, A. Broido, and K. Claffy, “Remote physical device fingerprinting,” *IEEE TDSC*, vol. 2, no. 2, pp. 93–108, 2005.
- [69] N. Popper, “Identity Thieves Hijack Cellphone Accounts to Go After Virtual Currency,” <https://www.nytimes.com/2017/08/21/business/dealbook/phone-hack-bitcoin-virtual-currency.html>, Aug 2017.
- [70] B. Schneier, “NIST is No Longer Recommending Two-Factor Authentication Using SMS,” https://www.schneier.com/blog/archives/2016/08/nist_is_no_long.html, Aug 2016.



Furkan Alaca is an Assistant Professor, Teaching Stream at the Department of Mathematical and Computational Sciences at the University of Toronto Mississauga. He completed his PhD (2018) at the School of Computer Science at Carleton University. His research interests include web authentication and device fingerprinting.



AbdelRahman Abdou is an Assistant Professor of Computer Science at Carleton University, Canada. In 2018, he was a postdoctoral researcher at the Institute of Information Security, ETH Zürich, Switzerland. He received his PhD (2015) in Systems and Computer Engineering from Carleton University. His research interests span SDN security, Internet security (including TLS, DNSSEC, authentication, secure BGP, and secure Internet geolocation), and using Internet measurements to understand and solve problems

related to Internet systems’ security.



Paul C. van Oorschot is a Professor of Computer Science at Carleton University, and the Canada Research Chair in Authentication and Computer Security. He was the program chair of USENIX Security 2008 and NDSS 2001-2002; a co-author of the *Handbook of Applied Cryptography*; and a past Associate Editor of *IEEE TDSC*, *IEEE TIFS*, and *ACM TISSEC*. He is a Fellow of the ACM, IEEE, and Royal Society of Canada. His research interests include authentication and Internet security.

Appendix A

UDS Properties

The definitions of the UDS benefits (properties) can be found in Bonneau et al. [1]. Here we list their labels and names for reader convenience. They are grouped in three categories: usability (U), deployability (D), and security (S).

U1 *Memorywise-Effortless*
U2 *Scalable-for-Users*
U3 *Nothing-to-Carry*
U4 *Physically-Effortless*
U5 *Easy-to-Learn*
U6 *Efficient-to-Use*
U7 *Infrequent-Errors*
U8 *Easy-Recovery-from-Loss*
D1 *Accessible*
D2 *Negligible-Cost-per-User*
D3 *Server-Compatible*
D4 *Browser-Compatible*
D5 *Mature*
D6 *Non-Proprietary*
S1 *Resilient-to-Physical-Observation*

S2 *Resilient-to-Targeted-Impersonation*
S3 *Resilient-to-Throttled-Guessing*
S4 *Resilient-to-Unthrottled-Guessing*
S5 *Resilient-to-Internal-Observation*
S6 *Resilient-to-Leaks-from-Other-Verifiers*
S7 *Resilient-to-Phishing*
S8 *Resilient-to-Physical-Theft*
S9 *No-Trusted-Third-Party*
S10 *Requiring-Explicit-Consent*
S11 *Unlinkable*

The new benefits introduced herein are listed below. The mimicry-resistance (M) benefits are a new sub-category of security benefits. For definitions of these new properties, see Section 4.

U9 *No-False-Rejects*
U10 *Easy-to-Change-Credentials*
M1 *No-False-Accepts*
M2 *Resilient-to-Infrequent-Capture-or-Intercept*
M3 *Verifies-Non-Static-Information*
M4 *Resilient-to-Spoofing*