

VaultIME: Regaining User Control for Password Manager through Auto-correction

Le Guan¹, Sadegh Farhang¹, Yu Pu¹, Pinyao Guo¹,
Jens Grossklags², and Peng Liu¹

¹ Pennsylvania State University, PA, USA
{lug14, farhang, yxp134, pug132, pliu}@ist.psu.edu

² Technical University of Munich, Germany
jens.grossklags@in.tum.de

Abstract. Users are often educated to follow different forms of advice from security experts. For example, using a password manager is considered an effective way to maintain a unique and strong password for every important website. However, user surveys reveal that most users are not willing to adopt this tool. They feel uncomfortable or even threatened, when they grant password managers the privilege to automate access to their digital accounts. Likewise, they are worried that individuals close to them may be able to access important websites holding personal data by using the stored passwords in the password manager.

We propose VaultIME to nudge more users towards the adoption of password managers by offering them a tangible benefit, while only slightly interfering with their current usage practices. Instead of “auto-filling” password fields, we propose and develop a new mechanism to “auto-correct” passwords in the presence of minor typo errors. VaultIME innovates by integrating the functionality of a password manager into the input method editor. Specifically, running as an app on mobile phones, VaultIME remembers user passwords on a per-app basis, and corrects mistyped passwords within a typo-tolerant set. We show that VaultIME achieves high levels of usability *and* security. With respect to usability, VaultIME is able to correct as many as 47.8% of password typos in a real-world password typing dataset. In terms of security, simulated attacks reveal that the security loss brought by VaultIME against a brute-force attacker is at most 0.43% in the worst case.

Key words: Password manager, Auto-correction, IME, Usable security

1 Introduction

To keep their digital accounts safe, Internet users are advised to adopt strong passwords that are hard to crack and guess [14]. However, long and random passwords are also difficult for users to remember [11]. Further, the sizable number of online accounts users need to manage has introduced an additional burden [5]. Using a *password manager* (e.g., 1password, lastpass and keepassdroid), which saves user credentials into a database, is a highly recommended approach by

security experts. Contents in the credential database are encrypted for data protection, where the encryption/decryption key is generated from a master password only known to the user [5].

Unfortunately, adoption of password managers is behind expectations despite the benefits apparent to security experts: 1) enhancing convenience by “auto-filling” password fields on behalf of the user [14], and 2) improving security by allowing for long and complex passwords. In addition, a password manager would reduce the perceived need for insecure practices such as storing passwords in clear-text as a memory help etc. Nevertheless, surveys indicate adoption figures as low as 6% [2] and at most as high as 21% [4], which leave a lot to be desired. Further, since password manager adopters are generally more security-savvy [4], this leaves behind those users who would most benefit from the technology.

Prior survey research has shown a split between the perceptions of adopters of password managers and those that hesitate [4]. While adopters echo the security benefits lauded by experts, 78% of non-adopters perceive “some” or “a lot” of individual risk from using a password manager [4]. Some factors for hesitation are quite reasonable, and hard to address. For example, some people simply do not trust providers of password managers [19], and software vulnerabilities may lead to exposure of all user passwords to hackers [6].

Other impeding factors are more amendable to solution approaches. Specifically, the threat of a lost phone or merely unmonitored access to the phone may be perceived quite disconcerting if high value data and important services such as social networking and online banking are left more vulnerable due to the stored credentials in a password manager. In fact, otherwise trusted individuals such as family members are often the cause of such invasions [18]. According to a Javelin Research study, in 2014, there were 550,000 reports of identity theft caused by someone the victim knew [7]. Taking advantage of the bond of trust, individuals are able to more easily access family members’ digital accounts and use the stolen identities to gain financial benefits [18, 12, 7]. Further, trust is especially impeded when the provider stores the password file on the cloud [19], rather than on the user’s machine. In addition, empirical work shows that people prefer a high *degree of control* when completing form-fields with personal information over having the same done by auto-fill [10]; we anticipate that a similar finding could be made in the highly related context of passwords.

With our work, we want to provide a stepping stone to nudge people towards adopting a password manager by providing an easy-to-understand benefit, while limiting interference with their habituated usage practices. Further, we target adoption hesitation due to the aforementioned reasons by allowing for a higher degree of control by the user.

Concretely, we propose a mechanism to *auto-correct* passwords in the presence of minor typo errors by utilizing a client-side password vault. While the user is still required to input a “near correct” password to activate the auto-correction feature, the approach allows users to apply longer and less trivial passwords. At the same time, user frustration can be substantially reduced by a tangible reduction of failed attempts. In this sense, our solution provides a

potentially sensible middle-ground for the adoption of password managers by leaving full control over authentication in the hands of the user, and reducing the threat of stolen data when a mobile device is lost or individuals with access to the device betray the trust of the user.

While the first systematic work of password auto-correction appears in [3], it is implemented on the server-side with the purpose of increasing the password acceptance rate. The authors found that almost 10% of failed login attempts are caused by simple, easily correctable typos that should otherwise be accepted. Following this observation, the authors proposed an auto-correction framework that can be integrated into existing password-based authentication systems on the server-side. In particular, a set of correctors¹ are first defined, and a received password is adjusted by each of the correctors to generate a set of candidate passwords. The login attempt is granted provided that at least one of the candidate passwords results in a password hash value that matches the one stored on the authentication server. When it comes to the security of the typo-tolerant authentication scheme, the authors show that it does not downgrade the security of user passwords by offering a formal proof of a free correction theorem.

Different from previous server-side auto-correction, we aim to provide added convenience of password typing on the client-side to further enhance user control. We propose VaultIME, a mobile-centric password manager granting users control of password input. VaultIME integrates the functionality of a password manager into an *Input Method Editor (IME)*, which is an app that displays a software keyboard and enables users to enter text. In particular, VaultIME remembers a user password on a per-app basis. If a password input interface is detected, the auto-correction feature is activated, which replaces a mistyped password (within an acceptable set) with the correct one.

The design goals of the new password manager are as follows. First, to meaningfully reduce user frustration, the auto-correction mechanism should cover a wide range of mistypes. Second, our mechanism should not downgrade password security even if an attacker has access to the phone and could perform a brute-force attack to stored passwords. To achieve the first goal, we conducted a mobile-centric password typing analysis. Based on it, we developed a new set of password correctors, which differ from the previous work [3] and cover 26.3% more typos. To achieve the second goal, we designed VaultIME to be compatible with the free auto-correction theory of [3], which states that with a certain filter policy, auto-correction introduces *zero* security loss. To measure the security loss, we ran simulation attacks to our auto-correction scheme. In the worst case, we show that the security loss is 0.43%, assuming that a brute-force attacker has 10 tries. When configured with the filter complying with the free auto-correction theory, VaultIME introduces *zero* security loss as expected. We have developed a proof-of-concept prototype of VaultIME. With reasonable optimization, the prototype results in no user-perceivable delay when auto-correcting passwords. However, interface features could be added to increase awareness of the benefits of auto-correction.

¹ For example, switching caps status, removing the last character, etc.

Contributions. Our work provides the following key contributions:

1. We propose a design for password managers addressing user concerns substantiated in related work. Without losing control to the login process, our design ameliorates users' concerns for using password manager in a "too open" way and maintains users' habituated login process.
2. To cover a maximum range of typos, while maintaining tight control over security, we analyze the nature of typos on a mobile platform in a systematic way. Based on the analysis results, we develop a new set of correctors, and run simulation attacks to measure the security loss introduced by VaultIME.
3. We implement a prototype of VaultIME as a normal Android IME app. Therefore, VaultIME can be instantly deployed on existing mobile platforms.

2 Background

This section explains the concept and design of the input method framework in the Android mobile OS as well as password managers.

Input Method Editor. Since API level 3.0, Android, the most popular mobile operating system, provides an extensible input-method framework. By extending the `InputMethodService` class, developers are able to implement a customized soft keyboard for better experience and capabilities. Besides, extending the `KeyboardView` class allows for the rendering of a personalized keyboard layout. These classes are packaged together to compose an *Input Method Editor (IME)* which provides user control to enable users to enter text.

When a user inputs text for an app, the default IME pops up. The framework allows an IME to completely control user input, including reading current input, and making arbitrary modifications. These functions are supported by operating on an `InputConnection` class. In particular, method `getTextBeforeCursor` and `getTextAfterCursor` can be invoked to read input before and after the current cursor, while an app ultimately receives an input string determined by the `commitText` methods.

Password Managers. Memorizing passwords has become a significant challenge for users. Although difficult to crack by attackers, strong passwords that are sufficiently long and random are also hard for users to remember [11].

Using a password manager is one of the most recommended approaches that can free users from the duty of memorizing lots of complex passwords. Mainly developed as a plug-in for web browsers, or as stand-alone web/smartphone applications, password managers save user credentials into a database, and later automatically auto-complete requests for the credentials on behalf of users [14]. In order to ensure security of the credential database, a user controls access to the password manager database via a master password. Specifically, contents in a credential database are typically encrypted for data protection, where the encryption/decryption key is generated from a master password [5].

3 Server Side Typo-tolerant Checking Scheme

To allow for a direct comparison, our work follows the formalization of a password authentication system proposed in [3], and also applies the same model for evaluating security loss in the presence of a brute-force attacker. To begin with, we review some of the important concepts and notations.

3.1 System Model

Checking Passwords. Two phases are involved in a password authentication process. In the registration phase, a user registers his password, e.g., w , with the server, and the server stores another string, s , derived from a hash function mixing a random salt value and w . In the checking phase, a user submits a password, \tilde{w} , to the authentication server, and the server verifies the request by calculating on \tilde{w} and the stored value s . The request is granted only if it returns true. In an *exact checker* (**ExChk**), the checker returns true only if the typed password \tilde{w} is exactly equal to w , i.e., $\tilde{w} = w$.

Typo-tolerant Scheme. Contrary to an exact checker **ExChk**, a typo-tolerant scheme runs a *relaxed checker*, which may return a true value for multiple strings other than w . When a user submits \tilde{w} , the authentication algorithm, instead of only examining \tilde{w} , examines a set of strings neighboring \tilde{w} . This set is represented by a ball of \tilde{w} denoted by $B(\tilde{w})$. If any element in the ball passes the exact checker **ExChk**, \tilde{w} is accepted. Formally, the ball is derived by applying a set of correctors (or transformation functions) $\mathbf{C} = \{f_0, f_1, \dots, f_c\}$ to \tilde{w} .

Brute-force Attacker and Security Loss. Before formalizing a brute-force attacker, we first model the password distribution and typo distribution. The theoretical analysis of security loss introduced by a brute-force attacker against a relaxed checker assumes an attacker with exact knowledge of these distributions.

We associate a distribution p to a set of all possible passwords. Therefore, $p(w)$ is the probability that a user selects a string w as a password. A user with password w may type a password \tilde{w} upon authentication. The probability of this event is represented by $\tau_w(\tilde{w})$. If $w \neq \tilde{w}$, a typo occurred. Furthermore, we say \tilde{w} is a neighbor of w if $\tau_w(\tilde{w}) > 0$.

Let $\{w_1, w_2, w_3, \dots\}$ be a non-increasing sequence of passwords ordered by their probabilities. $\lambda_q = \sum_{i=1}^q p(w_i)$ is called the q -success rate. The success rate of an attacker \mathbf{A} trying to guess a user's password is denoted by $\text{Att}(\text{checker}, \mathbf{A}, q)$, in which checker is the checking algorithm, and q represents the maximum number of tries attacker \mathbf{A} can make. For an exact checker, it is obvious that $\text{Att}(\text{ExChk}, \mathbf{A}, q) \leq \lambda_q$. To achieve λ_q , a brute-force attacker must choose the password with the highest probability in each round.

Regarding a relaxed checker, we define an *optimal attacker* to be able to achieve the maximum password guessing probability. Formally, the probability that an optimal attacker successfully guesses a password in q time is denoted by $\lambda_q^{\text{fuzzy}} = \max_{\mathbf{A}} \text{Att}(\text{Chk}, \mathbf{A}, q)$. Similar to the case of an exact checker, where

the attacker chooses the passwords with the highest probabilities, an optimal attacker against a relaxed checker tries to guess a password \tilde{w} , so that the corresponding ball $B(\tilde{w})$ has the highest aggregate probability in each round. The construction of such an optimal attacker is NP-hard. However, in [3], the authors proposed a greedy algorithm to realize this attacker in practice. As a result, the security loss caused by such a greedy attacker against a relaxed checker can be calculated by $\Delta_q^{greedy} = \lambda_q^{greedy} - \lambda_q$.

3.2 Secure Typo-tolerant Checker

The naïve relaxed checker downgrades the security of an authentication system in the presence of an optimal attacker, i.e., $\Delta_q > 0$. However, there exists an optimal relaxed checker, **OpChk**, that avoids causing security degradation (*free corrections*), i.e., $\Delta_q = 0$ [3]. When a user submits a string \tilde{w} as password, the relaxed checker creates a set of candidate passwords based on a set of correctors \mathbf{C} , and thereby a candidate set $\hat{B}(\tilde{w}) = \{w' | w' = f_i(\tilde{w}), p(w')\tau_{w'}(\tilde{w}) > 0, f_i \in \mathbf{C}\}$. To guarantee security, the optimal checker **OpChk** further rules out some of the candidate passwords by solving an optimization problem with a brute-force algorithm. **OpChk** maximizes the password acceptance rate without losing security. For the detailed explanation of the algorithm see [3]; Section V.D.

3.3 Limitations of Server-side Password Auto-correction

Previous work is invaluable as it provides a theoretical basis for a secure typo-tolerant authentication scheme, in contradiction to the common belief that accepting more than the one correct password would significantly degrade security. However, as shown in the paper, the proposed scheme cannot handle proximity typos, which, however, are the most prevalent form of all typos (21.8%). Their occurrence is even more pronounced for mobile clients (29.6%). Proximity typos occur when a user accidentally hits a key adjacent to the intended one (e.g., hitting an ‘a’ instead of an ‘s’). The reason for this limitation is that correcting a proximity typo necessitates the coverage of a larger space of possible passwords, and running the hash-based authentication algorithm for each possible password requires considerable computational resources. For enterprises, this requires more infrastructure investments to enhance computing capability. For customers, the introduced latency can be unacceptable.

Drawing on the specific situational context of the mobile environment and ecosystem, we design VaultIME to overcome innate limitations of the previous work, and enable VaultIME to cover more typos. Specifically, implemented as a password manager on smartphones, VaultIME is aware of the correct password. Therefore, checking a candidate password is as simple as performing a string matching, as opposed to the complex hash calculations needed by previous work. Since computationally intensive hash computation is avoided, covering proximity typos becomes possible.

4 Empirical Study of Typos on Mobile Devices

Prior studies have shown that strong passwords are difficult to type [8, 9, 16]. For example, users could easily mistype a character by slipping to an adjacent position on the keyboard, or they may forget to switch off the caps lock status. These human problems are further exacerbated on mobile devices. In particular, the cramped, and less tactile virtual keypad, which is widely used on today’s mobile phones, has a negative influence on error-free typing [13, 15]. As a result, it has been reported that the error rate is 8% higher for text typed on virtual keypads than for physical keyboards [13].

To understand the most frequent types of typos on mobile devices, we need to analyze a sizable number of real-world password-typing observations. For this purpose, we work on publicly available password-typing datasets from the previous work [3], and particularly focus on the data collected on touchscreen mobile devices.² In this section, we first briefly introduce these datasets. Then, we present our analysis results. Our results uncover several new findings, which guide us in designing new mobile-centric auto-correction schemes.

4.1 Password-typing Dataset on Touchscreens

In [3], the authors carried out two experiments on the Amazon Mechanical Turk (MTurk) platform to collect typo records during the entering of passwords. One experiment collected data from either PC or mobile platforms, while the other only collected data from mobile devices with touchscreens. In collecting the latter dataset, human-intelligence tasks (HITs) were assigned to participants over the web, where each participant was required to type 10~14 passwords in an HTML password input box within 300 seconds. The participants could only use touchscreen mobile devices. The results were later verified by the `user-agent` field in the HTTP header of the workers’ browsers. The passwords were sourced from the RockYou password leak [17], one of the largest leaked password databases. In total, 24,000 password-typing records were collected by 1,987 HITs.

4.2 Understanding Typos on Mobile Devices

In this section, we explain our findings by analyzing the dataset mentioned above. We first list top typos and their corresponding correctors in Table 1. A corrector is the reverse operation of the corresponding typo. It returns a set of passwords that could potentially contain the intended one. For example, corrector `rm-last` removes the last character in the received password, which effectively corrects typo `ins-last`. While the definitions of many correctors can be found in work [3], the newly introduced ones are quite self-explanatory. For example, `rep-prox-rs` means for each character, replace it with each of the adjacent ones in the correct keyboard status.

² The dataset collected on touchscreen devices can be downloaded from <https://www.cs.cornell.edu/~rahul/data/mturk15-touchonly.json.bz2>.

Table 1. Top typos and their corresponding correctors.

Typo explanation	Typo	Corrector
Proximity errors, i.e., hitting an adjacent key regardless of the intended keyboard status [†] , e.g., typing an ‘a’ as an ‘S’.	prox [‡]	n/a
Proximity errors with correct status, i.e., hitting an adjacent key in the same keyboard status with the intended one, e.g., typing an ‘a’ as an ‘s’.	prox-rs	rep-prox-rs
All letters are flipped.	swc-all [‡]	swc-all
First letter is flipped.	swc-first [‡]	swc-first
An extra character is added to end.	ins-last [‡]	rm-last
An extra character is added to front.	ins-first [‡]	rm-first
Forget pressing shift for symbol at the end.	n2s-last [‡]	n2s-last
Miss a character at an arbitrary location.	rm-any	ins-any
Insert an extra character at an arbitrary location.	ins-any	rm-any
An arbitrary letter is flipped.	swc-any	swc-any

[†]: The keyboard statuses are “normal”, “capitalized”, and “symbolized” in the AOSP keyboard.

[‡]: The definition of the typo is also used in [3].

Table 2. Top typos that occur in the mobile dataset and general dataset.

Environment	Typo Percentages						
	prox	swc-all	ins-last	swc-first	ins-first	n2s-last	others
Any	21.8	10.9	4.6	4.5	1.3	0.2	56.6
Mobile	prox-rs	rm-any	ins-any	swc-all	swc-any	ins-last	others
	21.4	20.4	10.8	8.0	7.6	1.2	32.6

1. The “Any” row covers the results drawn directly from [3]. The dataset is collected from participants with PC or mobile devices.
2. The “Mobile” row covers the results obtained from mobile devices only.
3. The sum of all items in the mobile environment is greater than 1. This is because our definitions of typos are not exclusive. For example, `ins-last` is a special case of `ins-any`.

In Table 2, we show top typos that occur in both the mobile and general datasets. Let us first have a look at the “any” row drawn directly from previous work [3]. Their solution can handle all typos except for `prox` and `others`, resulting in a coverage rate of 21.5%. However, `prox` alone contributes 21.8% of all typos, which the previous solution does not address. We have discussed the reason why previous work cannot handle proximity errors in Section 3.3.

We independently conducted a typo distribution analysis on the mobile dataset, the results of which are shown in the “Mobile” row in Table 2. Our study differs from the previous work as we are more concerned with specifics in the mobile environment. We differentiate between a virtual keyboard and a physical one, and pay more attention to the respective influences on typing.

We explain our new findings in the following. First, we find that PC users frequently make proximity typos with incorrect keyboard status, such as typing

‘a’ as ‘S’. This can be explained by the combined effect of finger slipping and unnoticed caps status. However, mobile users seldom make such mistakes. The reason is that a virtual keyboard typically reflects the keyboard status directly on the display of each key, which a user is likely to notice. Therefore, we define a new mobile-centric proximity error, i.e., **prox-rs**. The difference to the general **prox** is that **prox-rs** only considers proximity errors with correct caps and symbol status.³ Therefore, typing ‘a’ as ‘S’ or ‘@’ is not considered as a proximity error in our analysis.⁴

Apart from proximity errors, we found that mobile users frequently miss (20.4%) or insert (10.8%) a character at arbitrary locations. In addition, they may also ignore capitalization, either completely (8.0%) or only for a single letter (7.6%). Compared with the “any” environment, where the users frequently *add* an additional character, mobile users are more likely to *miss* a character. Indeed, unintentional extra key-strokes can happen due to inertia in high-speed input on physical keyboards. Among these typos, we found that correcting a missing character is challenging, i.e., a huge number of password candidates would need to be examined. This number is roughly estimated as the number of all possible characters (over 100) multiplied by the length of a password. Therefore, we do not consider this kind of typo in this work. It is also interesting to mention that both of **swc-all** and **swc-any** contribute substantially to mobile typos. While the previous work only handles **swc-all**, we argue that people are equally likely to flip only one letter, which has already been validated by our experiments. In the next section, we show how we auto-correct these typos. In total, our correctors can handle as many as 47.8% of the typos, which is the union of typos of type **prox-rs**, **ins-any**, **swc-all**, and **swc-any**.

5 Password Auto-correction for Mobile

VaultIME implements a password auto-correction scheme on the mobile client side. Instead of letting the authentication algorithm on the server judge whether a password should be accepted or not, VaultIME directly auto-corrects the passwords on the mobile client’s side if only minor typos occur. To achieve this, VaultIME, as a special IME, stores the correct password for users on a per-app basis, and runs a password checker as defined in Section 3. Before a typed password is fed to the corresponding app, the checker checks the received input. If the checker returns true, the stored correct password is forwarded to the app, otherwise, the received input is forwarded as is.

More specifically, after the user is done with password input, the checker in IME first checks the received password \tilde{w} . If it matches with the correct password, w , recorded in the password vault, the IME leaves the password as is

³ In the default AOSP keyboard layout, there are three statuses (“normal”, “capitalized”, and “symbolized”), which map the letter ‘a’ to ‘a’, ‘A’, and ‘@’ respectively.

⁴ As a result, the results of the previous work exhibit a higher proportion of proximity error (29.6%) than measured with **prox-rs** (21.4%) on the same raw data.

and returns. Otherwise, a ball $B(\tilde{w})$ of candidate passwords is derived from a predefined transformation function set $\mathbf{C} = \{f_1, \dots, f_c\}$, where f_i is a corrector defined in Section 4. Then, w is compared with each element in the ball. If a match is found, \tilde{w} is replaced by w ; otherwise, \tilde{w} is left as is.

This section first defines the used transformation function sets. Then, we present how these functions influence the ball size under different checking policies. A checking policy is a filter applied to the candidate ball obtained by the naïve relaxed checker. A stricter filter leads to a reduced ball size, but retains more security of the password. Our results show that the optimal checker, **OpChk**, does not reduce the ball size significantly. Since **OpChk** has been proven to lose zero security of a password, our system can achieve both high security and high usability. Finally, we also run simulation experiments to demonstrate that our scheme is secure against a greedy attacker.

5.1 Transformation Function Sets

A transformation function is also called a corrector, which is the reverse operation of a typo, and can be used to recover the correct password. We have listed top-rated mobile correctors in Table 2. Based on their capabilities (i.e., coverage of typos) to correct typos, we define four transformation function sets. They are $\mathcal{C}_{top1} = \{rep-prox-rs\}$, $\mathcal{C}_{top2} = \mathcal{C}_{top1} \cup \{rm-any\}$, $\mathcal{C}_{top3} = \mathcal{C}_{top2} \cup \{suc-all\}$, and $\mathcal{C}_{top4} = \mathcal{C}_{top3} \cup \{suc-any\}$, respectively.

5.2 Ball Size Estimation

In [3], three checking policies are discussed. In **Chk-All**, the algorithm tries all the derived passwords in the ball $B(\tilde{w})$. In **Chk-wBL**, the ball is filtered by a predefined blacklist that is comprised of a set of frequently used passwords. In **Chk-AOp**, based on empirical distributions of passwords and typos (p, τ), a brute-force algorithm is executed to filter the ball. The algorithm maximizes the password acceptance rate without losing security against a greedy attacker who knows both the distribution (p, τ) and the algorithm of the checker.

Table 3. Average ball size for all RockYou passwords over different checker policies and transformation function sets.

	\mathcal{C}_{top1}	\mathcal{C}_{top2}	\mathcal{C}_{top3}	\mathcal{C}_{top4}
Chk-All	59.25	69.61	70.54	79.16
Chk-wBL	59.24	69.60	70.53	79.14
Chk-AOp	53.80	58.77	57.87	64.06

To understand the effect of policies applied to the ball, we run a simulation to calculate the averaged ball size after filtering. As shown in Table 3, the ball size decreases when policies are applied (**Chk-All** can be viewed as an all-pass policy), and increases as more transformation functions are added to the set \mathbf{C} . Each increase is a reflection of the added corrector. From \mathcal{C}_{top1} to \mathcal{C}_{top2} , we observe an increment of around 10, indicating that **rm-any** produces 10 password candidates, which conforms to the length of a password. From \mathcal{C}_{top2} to \mathcal{C}_{top3} , only

one new password is produced. This is expected because `swc-all` is a one-to-one mapping. Lastly, `swc-any` produces less than 9 new passwords as there are around 9 letters in a password on average.

Statistically, all the checkers in Table 3 significantly increase the number of candidate passwords to be checked. On the one hand, this indicates that our checkers could achieve a high auto-correction rate, because more passwords are examined in each query. On the other hand, security could be degraded because an attacker gains more information about the real password in each query. Interestingly, from `Chk-A11` to `Chk-A0p`, we do not observe an abrupt shrinkage of the ball size. Since `Chk-A0p` leaks no more information about the real password than an unmodified exact checker leaks to an optimal brute-force attacker, this proves that our checker can achieve both a high auto-correction rate and a low security loss. In the next section, we show results from our simulation experiments. We emulate a greedy attacker who has complete knowledge about the implementation details of the used typo-tolerant checker.

5.3 Security Evaluation

We begin by clarifying the threats we consider in this work. Then, we show the measured security losses under a set of simulated attacks.

Threats in Scope. We consider an attacker who has physical access to an unlocked victim phone. This is particularly likely to happen considering an in-house betrayer. However, we do not consider a fully compromised mobile OS. In a compromised mobile OS, the attacker may retrieve user’s credential data (including all keystrokes) remotely.

We consider a brute-force attacker who is given q chances to query the authentication system. Such an attacker has been formalized in Section 3.1. Specifically, the attacker follows the greedy algorithm mentioned in Section 3.1, and the security loss can be represented by $\Delta_q^{greedy} = \lambda_q^{greedy} - \lambda_q$.

Results. In Figure 1, we show the security loss of each checker for different query numbers. We set the upper bound of q to 10, because it is a reasonable upper bound for queries given observations in practice before a device is locked. Mobile devices often enforce a long waiting time if consecutive failed login attempts are detected.

It is obvious that the security loss increases with q . However, `Chk-A0p` remains zero throughout our experiments, because it is an optimal checker that suffers no security loss in theory. For `Chk-A11` and `Chk-wBL` shown in the figure, there is a clear gap between the transformation function set \mathcal{C}_{top1} and others. This indicates that the security loss caused by applying `rep-prox-rs` alone can be quite limited – as low as 0.085% ($\lambda_q^{greedy} = 0.02937$ and $\lambda_q 0.02852$) in the worst case when $q = 10$ using checker `Chk-A11`. This can be explained by the fact that a proximity typo often leads to low probability passwords, which do not increase the overall aggregate probability of the attacker’s ball. For example, when checking the password ‘password’, `rep-prox-rs` will derive a huge ball containing candidate

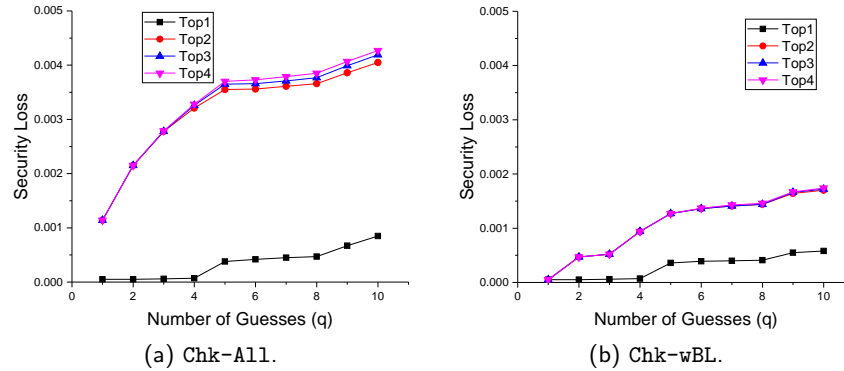


Fig. 1. Security loss measured for different checkers and query numbers. Note that the security loss for **Chk-A0p** is zero, so we omit it for the sake of fine typography.

passwords such as ‘oassword’ and ‘pssword’, which are rarely used by humans. On the other hand, applying **swc-all** will obtain ‘PASSWORD’, which is also a frequently used password. In the worst case, the security loss is 0.427% ($\lambda_q^{greedy} = 0.03279$ and $\lambda_q = 0.02852$) when $q = 10$ and using checker **Chk-A11** under the transformation function set \mathcal{C}_{top4} .

6 Implementation

We have implemented a proof-of-concept prototype of VaultIME for the Android OS. A user is able to customize the transformation function set ranging from \mathcal{C}_{top1} to \mathcal{C}_{top4} , and the checking algorithms among **Chk-A11**, **Chk-wBL**, and **Chk-A0p**.

The prototype uses the standard QWERTY US keyboard layout. It automatically detects the attribute of the current `TextView`, and inserts an “AuCo” key in the bottom right of the keyboard for the `TYPE_TEXT_VARIATION_PASSWORD` and `TYPE_TEXT_VARIATION_VISIBLE_PASSWORD` input types. VaultIME records a new password entry when the “AuCo” key is pressed. We use the package name of a login app and the account information as the key to index the password. Once a correct password has been recorded, subsequent login attempts will go through the typo-tolerant checker to auto-correct possible typos. As with traditional password vaults, the file storing passwords is encrypted by a secure master key [1]. The master key is randomly generated, and managed by the Android `KeyStore` provider.

7 Future Work

In the future, we plan to conduct user studies to investigate the usability of the VaultIME app as well as adoption intentions in detail. Specifically, by empirically

evaluating how users interact with our system, we aim to deliver a more usable and secure user experience for mobile phone users. Moreover, we are interested to learn to which degree users prefer our method to the traditional auto-fill password manager, whether users feel less threatened, have less frustration, and whether the correction process fits users' habituated login process.

In evaluating the security loss imposed by VaultIME, we mainly focus on a brute-force attacker who attempts to maximize the possibility coverage in each guessing. However, given that some personal data is publicly available (e.g., user name, birthday, etc.), particularly to family members or close friends, a targeted guessing attack could be more efficient [20]. Building an attack model which incorporates personal information into the on-line guessing and designing a new free auto-correction schema specific to this model constitutes an interesting research topic.

8 Conclusion

In this paper, we present VaultIME, a new password auto-correction scheme for mobile platforms. Our work ameliorates concerns of password manager users that they lack control over the use of their credentials. We achieve this by requiring the user to type a "near correct" password, which is automatically replaced with the correct one.

In designing the auto-correction policies, we conduct a mobile-centric password typo analysis, and are able to categorize the observed typos occurring while using virtual keyboards. Based on these empirical observations, we are able to develop a customized set of password correctors, which can cover as much as 47.8% of the detected password typos on mobile systems. This substantial coverage is made possible through a client-side implementation of our password-correction scheme as an app which allows for the treatment of the most common typographical errors, i.e., proximity typos. Moreover, the proposed auto-correction scheme is secure against a brute-force attacker under the formal model proposed in [3]. Our experimental results reveal that in the worst case, our scheme causes a security loss of 0.43%, indicating our auto-correction scheme has a high level of security robustness.

Acknowledgments: We would like to thank the anonymous reviewers for their insightful comments that helped to improve our paper. This work was supported by NSF CNS-1422594, ARO W911NF-13-1-0421 (MURI), and the German Institute for Trust and Safety on the Internet (DIVSI).

References

1. AgileBits, Inc. 1password security. <https://support.1password.com/1password-security/>.

2. R. Butler and M. Butler. The password practices applied by South African online consumers: Perception versus reality. *South African Journal of Information Management*, 17(1):1–11, 2015.
3. R. Chatterjee, A. Athalye, D. Akhawe, A. Juels, and T. Ristenpart. pASSWORD tYPOS and how to correct them securely. In *IEEE Security and Privacy (S&P)*, 2016.
4. M. Fagan and M. Khan. Why do they do what they do?: A study of what motivates users to (not) follow computer security advice. In *SOUPS '16*, 2016.
5. P. Gasti and K. Rasmussen. On the security of password manager database formats. In *ESORICS '12*, 2012.
6. A. Gott. Important security updates for our users, 2017. <https://blog.lastpass.com/2017/03/important-security-updates-for-our-users.html/>.
7. K. Grant. Identity theft victims: You might know the culprit, 2015. <http://www.cnbc.com/2015/07/21/identity-theft-victims-may-know-the-culprit.html>.
8. M. Keith, B. Shao, and P. Steinbart. The usability of passphrases for authentication: An empirical field study. *International Journal of Human-Computer Studies*, 65(1):17–28, 2007.
9. M. Keith, B. Shao, and P. Steinbart. A behavioral analysis of passphrase design and effectiveness. *Journal of the Association for Information Systems*, 10(2):63–89, 2009.
10. B. Knijnenburg, A. Kobsa, and H. Jin. Counteracting the negative effect of form auto-completion on the privacy calculus. In *ICIS '13*, 2013.
11. S. Komanduri, R. Shay, P. Kelley, M. Mazurek, L. Bauer, N. Christin, L. Cranor, and S. Egelman. Of passwords and people: Measuring the effect of password-composition policies. In *ACM CHI '11*, 2011.
12. S. Kossman. Familiar fraud: When family and friends steal your identity, 2014. http://www.creditcards.com/credit-card-news/familiar_fraud-damage-1282.php.
13. S. Lee and S. Zhai. The performance of touch screen soft buttons. In *CHI '09*, 2009.
14. Z. Li, W. He, D. Akhawe, and D. Song. The emperor’s new password manager: Security analysis of web-based password managers. In *USENIX Security '14*, 2014.
15. Y. Park, S. Han, J. Park, and Y. Cho. Touch key design for target selection on a mobile phone. In *MobileHCI '08*, 2008.
16. R. Shay, S. Komanduri, A. Durity, P. Huh, M. Mazurek, S. Segreti, B. Ur, L. Bauer, N. Christin, and L. Cranor. Can long passwords be secure and usable? In *ACM CHI '14*, 2014.
17. M. Siegler. One of the 32 million with a RockYou account? You may want to change all your passwords. Like now. *TechCrunch*, <http://techcrunch.com/2009/12/14/rockyou-hacked>, 2009.
18. J. Stroup. Who Commits Identity Theft?, 2016. <https://www.thebalance.com/who-commits-identity-theft-1947637>.
19. M. Tabini. Review: Lastpass takes your passwords to the cloud, 2013. <http://www.macworld.com/article/2032046/review-lastpass-takes-your-passwords-to-the-cloud.html>.
20. D. Wang, Z. Zhang, P. Wang, J. Yan, and X. Huang. Targeted online password guessing: An underestimated threat. In *ACM CCS '16*, 2016.