

PILOT: Password and PIN Information Leakage from Obfuscated Typing Videos *

Kiran Balagani ^{a,1}, Matteo Cardaioli ^{b,d}, Mauro Conti ^b,
Paolo Gasti ^a, Martin Georgiev ^{c,2}, Tristan Gurtler ^{a,3}, Daniele Lain ^{b,4}, Charissa Miller ^{a,5},
Kendall Molas ^a, Nikita Samarin ^{a,6}, Eugen Saraci ^b, Gene Tsudik ^c, and Lynn Wu ^{a,7}

^a *New York Institute of Technology*

^b *University of Padua*

^c *University of California, Irvine*

^d *GFT Italy*

Abstract. This paper studies leakage of user passwords and PINs based on observations of typing feedback on screens or from projectors in the form of masked characters (* or ●) that indicate keystrokes. To this end, we developed an attack called *Password and Pin Information Leakage from Obfuscated Typing Videos (PILOT)*. Our attack extracts inter-keystroke timing information from videos of password masking characters displayed when users type their password on a computer, or their PIN at an ATM. We conducted several experiments in various attack scenarios. Results indicate that, while in some cases leakage is minor, it is quite substantial in others. By leveraging inter-keystroke timings, *PILOT* recovers 8-character alphanumeric passwords in as little as 19 attempts. When guessing PINs, *PILOT* significantly improved on both random guessing and the attack strategy adopted in our prior work [4]. In particular, we were able to guess about 3% of the PINs within 10 attempts. This corresponds to a 26-fold improvement compared to random guessing. Our results strongly indicate that secure password masking GUIs must consider the information leakage identified in this paper.

Keywords: Authentication, Information leakage, Shoulder-surfing attacks, PIN inference, Password inference

1. Introduction

Passwords and PINs are susceptible to shoulder surfing attacks [26] of which there are two main types: (1) input-based and (2) output-based. The former is more common; in it, the adversary observes an input device (keyboard or keypad) as the user enters a secret (password or PIN) and learns the key-presses. The latter involves the adversary observing an output device (screen or projector) while the user enters a secret which is displayed in cleartext. The principal distinction between the two types is the adversary's

* Submitted to the ESORICS 2018 special issue.

¹ Authors are listed in alphabetical order.

² Current affiliation: University of Oxford.

³ Current affiliation: University of Illinois at Urbana-Champaign.

⁴ Current affiliation: ETH Zurich.

⁵ Current affiliation: Rochester Institute of Technology.

⁶ Current affiliation: University of California, Berkeley.

⁷ Current affiliation: Bryn Mawr College.

proximity: observing input devices requires the adversary to be closer to the victim than observing output devices, which tend to have larger form factors, i.e., physical dimensions.

Completely disabling on-screen feedback during password/PIN entry (as in, e.g., the Unix `sudo` command) mitigates output-based shoulder-surfing attacks. Unfortunately, it also impacts usability: when deprived of visual feedback, users cannot determine whether a given key-press was registered and are thus more likely to make mistakes. In order to balance security and usability, user interfaces typically implement password masking by displaying a generic symbol (e.g., “•” or “*”) after each keystroke. This technique is commonly used on desktops, laptops and smartphones as well as on public devices, such as Automated Teller Machines (ATMs) or Point-of-Sale (PoS) terminals at shops or gas stations.

Despite the popularity of password masking, little has been done to quantify how visual keystroke feedback impacts security. In particular, masking assumes that showing generic symbols does not reveal any information about the corresponding secret. This assumption seems reasonable, since visual representation of a generic symbol is independent of the key-press. However, in this paper we show that this assumption is incorrect. By leveraging precise inter-keystroke timing information leaked by the appearance of each masking symbol, we show that the adversary can significantly narrow down the password/PIN’s search space. Put another way, the number of attempts required to brute-force decreases appreciably when the adversary has access to inter-keystroke timing information.

There are many realistic settings where visual inter-keystroke timing information (leaked via appearance of masking symbols) is readily available while the input information is not, i.e., the input device is not easily observable. For example, in a typical lecture or classroom scenario, the presenter’s keyboard is usually out of sight, while the external projector display is wide-open for recording. Similarly, in a multi-person office scenario, an adversarial co-worker can surreptitiously record the victim’s screen. The same holds in public scenarios, such as PoS terminals and ATMs, where displays (though smaller) tend to be easier to observe and record than entry keypads.

In this paper we consider two representative scenarios: (1) a presenter enters a password into a computer connected to an external projector; (2) a user enters a PIN at an ATM in a public location. The adversary is assumed to record keystroke feedback from the projector display or an ATM screen using a dedicated video camera or a smartphone. We note that a human adversary does not need to be present during the attack: recording might be done via an existing camera either pre-installed or pre-compromised by the adversary, possibly remotely, e.g., as in the infamous Mirai botnet [17].

Contributions. The main goal of this paper is to quantify the amount of information leaked through video recordings of on-screen keystroke feedback. To this end, we conducted extensive data collection experiments that involved 84 subjects.¹ Each subject was asked to type passwords or PINs while the screen or projector was video-recorded using either a commodity video camera and a smartphone camera. Based on this, we determined the key statistical properties of the resulting data, and set up an attack called *PILOT: Password and Pin Information Leakage from Obfuscated Typing Videos*. It allows us to quantify reduction in brute-force search space due to timing information. *PILOT* leverages multiple publicly available typing datasets to extract population timings, and applies this information to inter-keystroke timings extracted from videos.

Our results show that video recordings can be effective in extracting precise inter-keystroke timing information. Experiments show that *PILOT* substantially reduces the search space for each password, even when the adversary has no access to user-specific keystroke templates. When run on passwords,

¹Where required, IRB approvals were duly obtained prior to the experiments.

PILOT performed better than random guessing between 87% and 100% of the time, depending on the password and the machine learning technique used to instantiate the attack. The resulting average speedup is between 25% and 385% (depending on the password), compared to random dictionary-based guessing; some passwords were correctly guessed in as few as 68 attempts. A single password timing disclosure is enough for *PILOT* to successfully achieve these results. However, when the adversary observes the user entering the password three times, *PILOT* can crack the password in as few as 19 attempts. Clearly, *PILOT*'s capabilities depend in part on the strength of a specific password. With very common passwords, benefits of *PILOT* are limited. Meanwhile, we show that *PILOT* substantially outperforms random guessing with less common passwords. With PINs, disclosure of timing poses an effective risk. The PIN guessing algorithm can reduce the number of attempts by up to 26 times compared to random guessing.

Paper Organization. Section 2 reviews the state-of-the-art in password guessing based on timing attacks. Section 3 presents *PILOT* and the adversary model. Section 4 discusses our data collection and experiments. We then present the results on password guessing using *PILOT* in Section 5, and on PIN guessing in Section 6. The paper concludes with the summary and future work directions in Section 7.

2. Related Work

There is a large body of prior work on timing attacks in the context of keyboard-based password entry. Song et al. [24] demonstrated a weakness that allows the adversary to extract information about passwords typed during SSH sessions. The attack relies on the fact that, to minimize latency, SSH transmits each keystroke immediately after entry, in a separate IP packet. By eavesdropping on such packets, the adversary can collect accurate inter-keystroke timing information. Song et al. [24] showed that this information can be used to restrict the search space of passwords. The impact of this work is significant, because it shows the power of timing attacks on cracking passwords.

There are several studies of keystroke inference from analysis of video recordings. Balzarotti et al. [5] addressed the typical shoulder-surfing scenario, where a camera tracks hand and finger movements on the keyboard. Text was automatically reconstructed from resulting videos. Similarly, Xu et al. [33] recorded user's finger movements on mobile devices to infer keystroke information. Unfortunately, neither attack applies to our sample scenarios, where the keyboard is invisible to the adversary.

Shukla et al. [23] showed that text can be inferred even from videos where the keyboard/keypad is not visible. This attack involved analyzing video recordings of the back of the user's hand holding a smartphone in order to infer which location on the screen is tapped. By observing the motion of the user's hand, the path of the finger across the screen can be reconstructed, which yields the typed text. In a similar attack, Sun et al. [25] successfully reconstructed text typed on tablets by recording and analyzing the tablet's movements, rather than movements of the user's hands.

The closest work to the paper is our prior work [4], in which we show that passwords can be inferred at a higher probability than random guesses using the timing information from onscreen keystroke feedback. However, in [4] we concluded that the timing information is not helpful in inferring PINs. In this paper, we revisit our earlier conclusion on inferring PINs and show that it is incorrect. In fact, the attack strategy employed in this paper yielded a 26-fold improvement in inferring PINs over random guesses and significantly outperforms [4] in terms of number of PINs recovered within a small number of attempts.

Another line of work aimed to quantify keystroke information inadvertently leaked by motion sensors. Owusu et al. [19] studied this in the context of a smartphone's inertial sensors while the user types using

the on-screen keyboard. The application used to implement this attack does not require special privileges, since modern smartphone operating systems do not require explicit authorization to access inertial sensors data. Similarly, Wang et al. [30] explored keystroke information leakage from inertial sensors on wearable devices, e.g., smartwatches and fitness trackers. By estimating the motion of a wearable device placed on the wrist of the user, movements of the user’s hand over a keyboard can be inferred. This allows learning which keys were pressed during the hand’s path. Compared to our work, both [19] and [30] require a substantially higher level of access to the user’s device. To collect data from inertial sensors the adversary must have previously succeeded in deceiving the user into installing a malicious application, or otherwise compromised the user’s device. In contrast, *PILOT* is a fully passive attack.

Acoustic emanations represent another effective side-channel for keystroke inference. This class of attacks is based on the observation that different keyboard keys emit subtly different sounds when pressed. This information can be captured (1) locally, using microphones placed near the keyboard [3, 35], or (2) remotely, via Voice-over-IP [10]. Also, acoustic emanations captured using multiple microphones can be used to extract the locations of keys on a keyboard. As shown by Zhou et al. [34], recordings from multiple microphones can be used to accurately quantify the *time difference of arrival (TDoA)*, and thus triangulate the positions of pressed keys.

3. System and Adversary Model

We now present the system and adversary model used in the rest of the paper.

We model a user logging in (authenticating) to a computer system or an ATM using a PIN or a password (*secret*) entered via a keyboard or keypad (*input device*). The user receives immediate feedback about each key-press from a screen, a projector, or both (*output device*) in the form of dots or asterisks (*masking symbols*). The shape and/or location of each masking symbol does not depend on which key is pressed. The adversary can observe and record the output device(s), though not the input device or the user’s hands. An example of this scenario is shown in Figure 1. The adversary’s goal is to learn the user’s secret.

The envisaged attack setting is representative of many real-world scenarios that involve low-privilege adversaries, including:

- (1) A presenter in a lecture or conference who types a password while the screen is displayed on a projector. The entire audience can see the timing of appearance of masking symbols, and the adversary can be anyone in the audience.
- (2) An ATM customer typing a PIN. The adversary who stands in line behind the user might have an unobstructed view of the screen, and the timing of appearance of masking symbols (see Figure 2).
- (3) A customer enters her debit card PIN at a self-service gas-station pump. In this case, the adversary can be anyone in the surroundings with a clear view of the pump’s screen.

Although these scenarios seem to imply that adversary is located near the user, proximity is not a requirement for our attack. For instance, the adversary could watch a prior recording of the lecture in scenario (1); or, could be monitoring the ATM machine using a CCTV camera in (2); or, remotely view the screen in (3) through a compromised IoT camera.

Also, we assume that, in many cases, the attack involves multiple observations. For example, in scenario (1), the adversary could observe the presenter during multiple talks, without the presenter changing passwords between talks. Similarly, in scenario (2), customers often return to the same ATM.

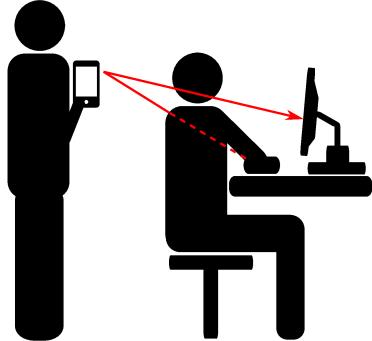


Fig. 1. Example attack scenario.



Fig. 2. Attack example – ATM setting. (a) Adversary's perspective. (b) Outsider's perspective.

4. Overview and Data Collection

We assume that the adversary can capture only inter-keystroke timings leaked by the output device while the user types a secret. The goal is to analyze differences between the distribution of inter-keystroke timings and infer corresponding keypairs. This data is used to identify the passwords that are most likely to be correct, thus restricting the brute-force search space of the secret. To accurately extract inter-keystroke timing information, we analyze video feeds of masking symbols, and identify the frame where each masking symbol first appears. In this setting, the accuracy and resolution of inter-keystroke timings depends on two key factors: the refresh frequency of the output device, and the frame rate of the video camera. Inter-keystroke timings are then fed to a classifier, where classes of interest are keypairs. Since we assume that the adversary has no access to user-specific keystroke information, the classifier is trained on population data, rather than on user-specific timings.

In the rest of this section, we detail the data collection process. We collected password data from two types of output devices: a VGA-based external projector, and LCD screens of several laptop computers. See Section 4.1 for details of these devices and corresponding procedures. For PIN data, we video-recorded the screen of a simulated ATM. Details can be found in Section 4.2. The dataset is published at <https://lamp.soecs.nyit.edu/datasets/silk-tv/>

4.1. Passwords

We collected data using an EPSON EMP-765 projector, and using the LCD screens of the subjects' laptops computers. In the projector setting, we asked the subjects to connect their own laptop so they would be using a familiar keyboard. The refresh rate of both laptop and projector screens were set to 60 Hz – the default setting for most systems. This setting introduces quantization errors of up to about $1/60$ s ≈ 16.7 ms. Thus, events happening within the same refresh window of 16.7ms are indistinguishable. We recorded videos of the screen and the projector using the rear-facing camera of two smartphones: a Samsung Galaxy S5 and an iPhone 7 Plus. With both phones, we recorded videos at 120 frames per second, i.e., 1 frame every 8.3 ms. To ease data collection, we placed the smartphones on a tripod. When recording the projector, the tripod was placed on a table, filming from a height of about 165 cm, to be horizontally aligned with respect to the projected image. When recording laptop screens, we placed the smartphone above and to the side of the subject, in order to mimic the adversary sitting behind the subject.

All experiments took place indoors, in labs and lecture halls at the authors' institutions. We recruited a total of 62 subjects, primarily from the student population of two large universities. Most participants

were males in their 20s, with a technical background and good typing skills. We briefed each subject on the nature of the experiment, and asked them to type four alphanumerical passwords: “jillie02”, “william1”, “123brian”, and “lamondre”. We selected these passwords uniformly at random from the RockYou dataset [1] in order to simulate realistic passwords. The subjects typed each password three times, while our data collection software recorded ground-truth keystroke timings of correctly typed passwords with millisecond accuracy. Timings from passwords that were typed incorrectly were discarded, and subjects were prompted to re-type the password whenever a mistake was made. The typing procedure lasted between 1 and 2 minutes, depending on the subject’s typing skills. All subjects typed with the “touch typing” technique, i.e., using fingers from both hands.

4.2. PINs

We recorded subjects entering 4-digit PINs on a simulated ATM, shown in Figure 3. Our dataset was based on experiments with 22 participants; 19 subjects completed three data collection sessions, while 4 subjects completed only one session, resulting in a total of 61 sessions. At the beginning of each session, the subject was given 45 seconds to get accustomed with the keypad of the ATM simulator. During this time, they were free to type as they pleased. Next, a subject was shown a PIN on the screen for ten seconds (Figure 4a), and, once it disappeared from the screen, asked to enter it four times (Figure 4b). Subjects were advised not to read the PINs out loud. This process was repeated for 15 consecutive PINs. During each session, subjects were presented with the same 15-PIN sequence 3 times. Subjects were given a 30-second break at the end of each sequence.

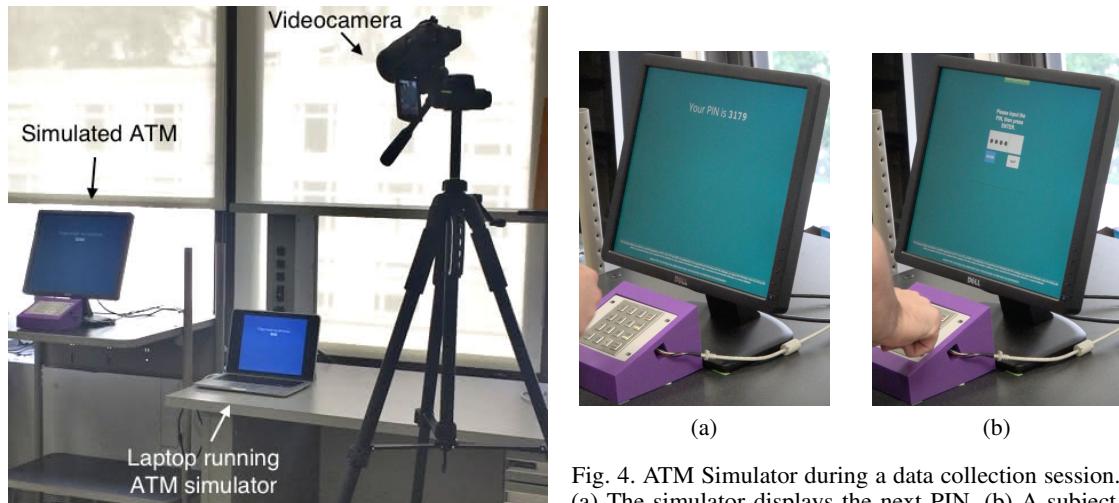


Fig. 3. Setup used in PIN inference experiments.

Specific 4-digit PINs were selected to test whether: (1) inter-keypress time is proportional to Euclidean Distance between keys on the keypad; and (2) the *direction of movement* (up, down, left, or right) between consecutive keys in a keypair impacts the corresponding inter-key time. We show an example of these two situations on the ATM keypad in Figure 5. We chose a set of PINs that allowed collection of a significant number of key combinations appropriate for testing both hypotheses. For instance, PIN 3179 tested horizontal and vertical distance two, while 1112 tested distance 0 and horizontal distance 1.

Sessions were recorded using a Sony FDR-AX53 camera, at a pixel resolution of $1,920 \times 1,080$ pixels, and 120 frames per second. At the same time, ATM simulation software collected millisecond-accurate inter-key distance ground truth by logging each keypress. PIN feedback was shown on a DELL 17" LCD screen with a refresh rate of 60 Hz, which resulted in each frame being shown for 16.7 ms.

4.3. Timing Extraction from Video

We developed software that analyzes video recordings to automatically detect the appearance of masking symbols and log corresponding timestamps. This software uses OpenCV [20] to infer the number of symbols present in each image. All frames are first converted to grayscale, and then processed through a bilateral filter [28] to reduce noise arising from the camera's sensor. The resulting images are analyzed using Canny Edge detection [11] to capture the edges of the masking symbol. External contours are compared with the expected shape of the masking symbol. When a masking symbol is detected, software logs the corresponding frame number.

Our experiments show that this technique leads to fairly accurate inter-keystroke timing information. We observed average discrepancy of 8.7 ms ($stdev$ of 26.6 ms) between the inter-keystroke timings extracted from the video and ground truth recorded by the ATM simulator. Furthermore, 75% of inter-keystroke timings extracted by the software had errors under 10 ms, and 97% had errors under 20 ms. Similar statistics hold for data recorded on keyboards for the passwords setting. Figure 6 shows the distribution of error discrepancies.

5. Password Guessing

PILOT treats identifying digraphs from keystroke timings as a multi-class classification problem, where each class represents one digraph, and input to the classifier is a set of inter-keystroke times. Without loss of generality, in this section, we assume that the user's password is a sequence of lowercase alphanumeric characters typed on a keyboard with a standard layout.

To reconstruct passwords, we compared two classifiers: Random Forest (RF) [16] and Neural Networks (NN) [22]. RF is a well-known classification technique that performs well for authentication based on

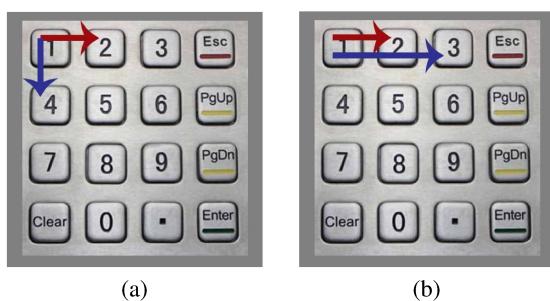


Fig. 5. ATM keypad in our experiments. (a) To type keypairs 1-2 and 1-4, the typing finger travels the same distance in different directions. (b) Keypairs 1-2 and 1-3 require the typing finger to travel different distances in the same direction.

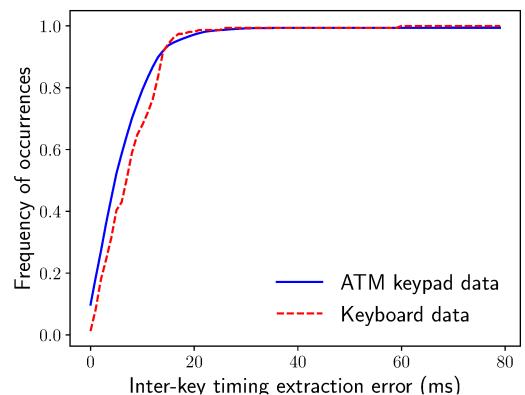


Fig. 6. CDF showing error distribution of inter-keystroke timings extracted from videos.

keystroke timings [7]. Input to RF is one inter-keystroke timing, and its output is a list of N digraphs ranked based on the probability of corresponding to input timing. NN is a more complex architecture designed to automatically determine and extract complex features from the input distribution. In our experiments, the input to NN is a list of inter-keystroke timings corresponding to a password. This enables NN to extract features, such as arbitrary n -grams, or timings corresponding to non-consecutive characters. NN’s output is a guess for the entire password.

We instantiated NN using the following parameters:

- number of units in the hidden layer – 128 (with ReLU activation functions);
- inclusion probability of the dropout layer – 0.2;
- number of input neurons – 25;
- number of output layers – 25 which represents one character in one-hot encoding. Output layers use softmax activation function;
- training was performed using batch sizes of 40 and 100 epochs. We used the Adam optimizer with a learning rate of 0.001.

Classifier Training. We trained *PILOT* on three public datasets [6, 21, 29] that contain keystroke timing information collected from English free-text. Using these datasets for training, we modeled an attack that relies exclusively on *population data*. Without loss of generality, we filtered the datasets to remove all timings that do not correspond to digraphs composed of alphanumeric lowercase characters. This is motivated by the datasets’ limited availability of digraph samples that contain special characters. In practice, the adversary could collect these timings using, for instance, crowdsourcing tools such as Amazon Mechanical Turk. To take care of uneven frequencies of digraphs, we under-represented the most frequent digraphs in the dataset. Data in public datasets was often gathered from free-text typing of volunteers. Therefore, more frequent digraphs in English were represented more than rarer ones. For example, considering lamondre, digraph `re` appears 43,606 times in the population dataset, while `am` has only 6,841. Similarly, in 123brian, digraph `ri` occurs 19,782 times, while `3b` has only 138 occurrences. We therefore under-sampled each digraph appearing more than 1,000 times to 1,000 randomly selected occurrences. Similarly, we excluded infrequent digraphs that appeared under 100 times in the whole dataset.

5.1. Attack Process.

To infer the user’s secret from inter-keystroke timings, *PILOT* leverages a dictionary of passwords (e.g., a list of passwords leaked by online services [1, 2, 12, 27]), possibly expanded using techniques such as probabilistic context-free grammars [32] and generative adversarial networks [15]. When evaluating *PILOT*, we assume that the user’s secret is in the dictionary. In practice, this is often the case, as many users use the same weak passwords (e.g., only 36% of the password of RockYou is unique [18]), and reuse them across many different services [14, 31]. Given that the size of a reasonable password dictionary is on the order of billions of entries,² the goal of *PILOT* is to narrow down the possible passwords to a small(er) list, e.g., to perform online attacks. This list is then ranked by the probability associated with each entry, computed from inter-keystroke timing data.

²See for example the lists maintained by <https://haveibeenpwned.com/>.

Specifically:

(1) Using RF, for each inter-key time extracted from video (corresponding to a digraph), *PILOT* returns a list of N possible guesses, sorted by the classifier's confidence. Next, *PILOT* ranks the passwords in the dictionary by resulting probabilities as follows: for each password, *PILOT* identifies the position in the ranked list of predictions for the first digraph of the password being guessed, and assigns that position as a “penalty” to the password. By performing these steps for each digraph, *PILOT* obtains a total penalty score for each password, i.e., a score that indicates the probability of the password given the output of the RF.

For example, to rank the password `jillie02`, *PILOT* first considers the digraph `ji`, and the list of predictions of RF for the first digraph. It notes that `ji` appears in such list as the X -th most probable; therefore, it assigns X as the penalty for `jillie02`. Then, it considers `il`, which appears in Y -th position in the list of predictions for the second digraph. The penalty for `jillie02` is thus updated to $X + Y$. This operation is repeated for all the 7 digraphs, thus obtaining the final penalty score.

(2) Using NN, *PILOT* computes a list of N possible guesses, sorted by the classifier's confidence of each guess. In this case, the *PILOT* processes the entire list of flight times at once, rather than refining its guess with each digraph.

We considered the following attack settings: *single-shot*, and *multiple recordings*. With the former, the adversary trains *PILOT* with inter-keystroke timings from population data, i.e., from users other than the target, e.g., from publicly available datasets, or by recruiting users and asking them to type passwords. In this scenario, the adversary has access to the video recording of a single password entry session. With multiple recordings, the adversary trains *PILOT* as before, and additionally, has access to videos of multiple login instances by the same user.

Training *PILOT* exclusively with population data leads to more realistic attack scenarios than training it with user-specific data, because usually the adversary has limited access to keystrokes samples from the target user. Further, access to user-specific data will likely improve the success rate of *PILOT*.

5.2. Results

In this section, we report on *PILOT* efficacy in reducing search time on the RockYou [1] password dataset compared to random choice, weighted by probability. We restricted experiments to the subset of 8-character passwords from RockYou, since the adversary can always determine password length by counting the number of masking symbols shown on the screen. This resulted in 6,514,177 passwords, of which 2,967,116 were unique.

Attack Baseline. To establish the attack baseline, we consider an adversary that outputs password guesses from a leaked dataset in descending order of frequency. (Ties are broken using random selection from the candidate passwords.) Because password probabilities are far from uniform (e.g., in RockYou, the top 200 8-character passwords account for over 10% of the entire dataset), this is the best adversarial strategy given no additional information on the target user.

Passwords selected for our evaluation represent a mix of common and rare passwords. Thus, they have widely varying frequencies of occurrence in RockYou and the expected number of attempts needed to guess each password using the baseline attack varies significantly. For example, the expected number of attempts for:

- `123brian` (appears 6 times) is 93,874;

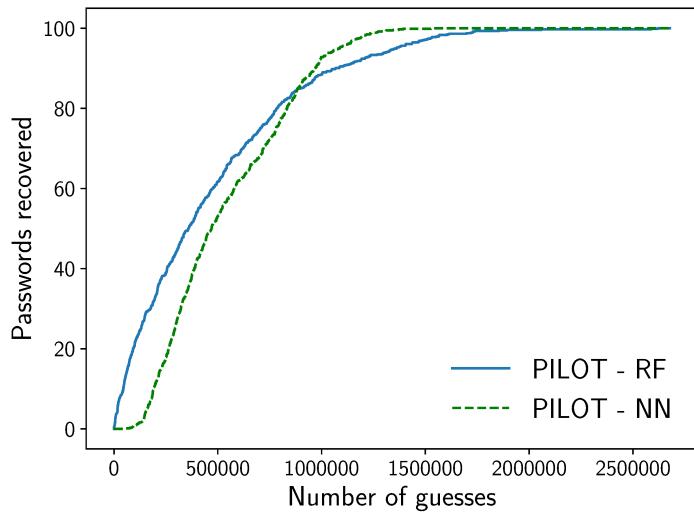


Fig. 7. CDF of the amount of passwords recovered by *PILOT*—Population Data attack scenario.

- `jillie02` (appears only once) is 1,753,571;
- `lamondre` (appears twice) is 397,213;
- `william1` (appears 1,164 times) is only 187.

Single-shot. Results in the single-shot setting are summarized in Table 1. The cumulative Distribution Function (CDF) of successfully recovered passwords is reflected in Figure 7, and the breakdown of results (by target password) is shown in Figure 8.

Results show that, for uncommon passwords (`jillie02` and `lamondre`), *PILOT* consistently outperforms random guessing. In particular, for `jillie02` both RF and NN greatly exceed random guessing, since both their curves in Figure 8 are above the random guess baseline. For `lamondre`, RF shows an advantage over random guess in 76% of the instances, while NN never beats the baseline.

For common passwords, sorted random guess wins over *PILOT*. In particular, `123brian` is both popular (i.e., the 93,874-th most popular password of the set, corresponding to the top 3% of the RockYou dataset) and very hard to recover with *PILOT*. This can be observed from Figure 8, where the curves corresponding to `123brian` are least steep. Finally, `william1`, being the 187-th most popular password, is always recovered early in our baseline attack, with the notable exception of one instance by RF.

In general, *PILOT* wins over the sorted random guess on infrequent passwords, such as `jillie02` and `lamondre`, that appear only once or twice, respectively. Such infrequent passwords exhibit the same random guess baseline curve and average, reported in Table 1 and shown in Figure 8. Given the similar steepness of CDF curves in Figure 8, which hint that *PILOT*’s performance might be similar for many other passwords, *PILOT* can probably outperform the baseline for uncommon passwords. We also note that uncommon passwords represent the vast majority of user-chosen passwords: 90% of RockYou passwords appear at most twice, and 80% exactly once. We expect that a realistic adversary would first generate password guesses based on their frequency alone (as in our baseline attack), and then switch to *PILOT* once these frequencies drop below some threshold.

Finally, we highlight that random guess baseline is computed on the distribution of passwords in RockYou. Other datasets might have different distributions: for example, in the *10 million password list*

Table 1

PILOT—Single-shot setting. *Avg*: average number of attempts to guess a password; *Stdev*: standard deviation; *Rnd*: number of guesses for the baseline adversary; *<Rnd*: how often PILOT outperforms random guessing; *Best*: number of attempts of the best guess; *< n*: how many passwords are successfully guessed within first $n = 20,000/100,000$ attempts.

	Avg	Stdev	Med	Rnd	<Rnd	Best	<20k	<100k
Random Forest								
123brian	581,743	414,761	508,332	93,874	8.7%	5,535	1.1%	9.3%
jillie02	749,718	448,319	656,754	1,753,571	97.8%	28,962	0.0%	2.7%
lamondre	301,906	334,681	199,344	397,213	75.0%	145	13.0%	33.7%
william1	246,437	264,090	145,966	187	0.5%	68	10.9%	39.9%
Neural Network								
123brian	923,534	165,454	886,802	93,874	0.0%	577,739	0.0%	0.0%
jillie02	456,811	210,512	383,230	1,753,571	100.0%	164,754	0.0%	0.0%
lamondre	517,472	189,355	493,713	397,213	28.8%	148,403	0.0%	0.0%
william1	265,813	140,753	215,840	187	0.0%	45,176	0.0%	3.8%

dataset [9], jillie02, lamondre, and 123brian appear only once, while william1 appears 176 times.

We believe that the discrepancy between the performance of PILOT on various passwords is due to how frequently the digraphs in each password appear in training data. Specifically, even with our under-representation, all digraphs in william1, with the exception of m1, are far more frequent in the training data than 12, 23, 3b, or 02.

Regarding specific classifiers, RF overtakes NN in most instances. For example, when guessing 123brian (Figure 8a), NN performs worse than random guessing for the first 800,000 attempts. Afterwards, NN outperforms both random guessing and RF. Furthermore, while RF can guess a substantial percentage of passwords within 20,000, 50,000 and 100,000 attempts, NN cannot achieve the same result.

In terms of minimum number of guesses per password, RF recovered william1 in 68, lamondre in 145, 123brian in 5,535, and jillie02 in 28,962 attempts. NN required a consistently higher minimum number of attempts for each password.

Multiple Recordings. Information from three login instances was used as follows. We averaged classifiers’ predictions over three login instances for a given user, and ranked passwords accordingly.

The results are summarized in Table 2 and Figure 9. Although PILOT still consistently outperforms random guessing, using data from multiple authentication recordings leads to mostly identical results overall with both RF and NN. PILOT’s guessing success rate for 123brian and jillie02 is slightly improved compared to the previous setting and minimum number of attempts to recover each password diminished slightly. We recovered william1 in 19, lamondre in 404, 123brian in 13,931, and jillie02 in 67,875 attempts. Overall, the results show that there are no substantial benefits in using timing data from three recordings from the same user.

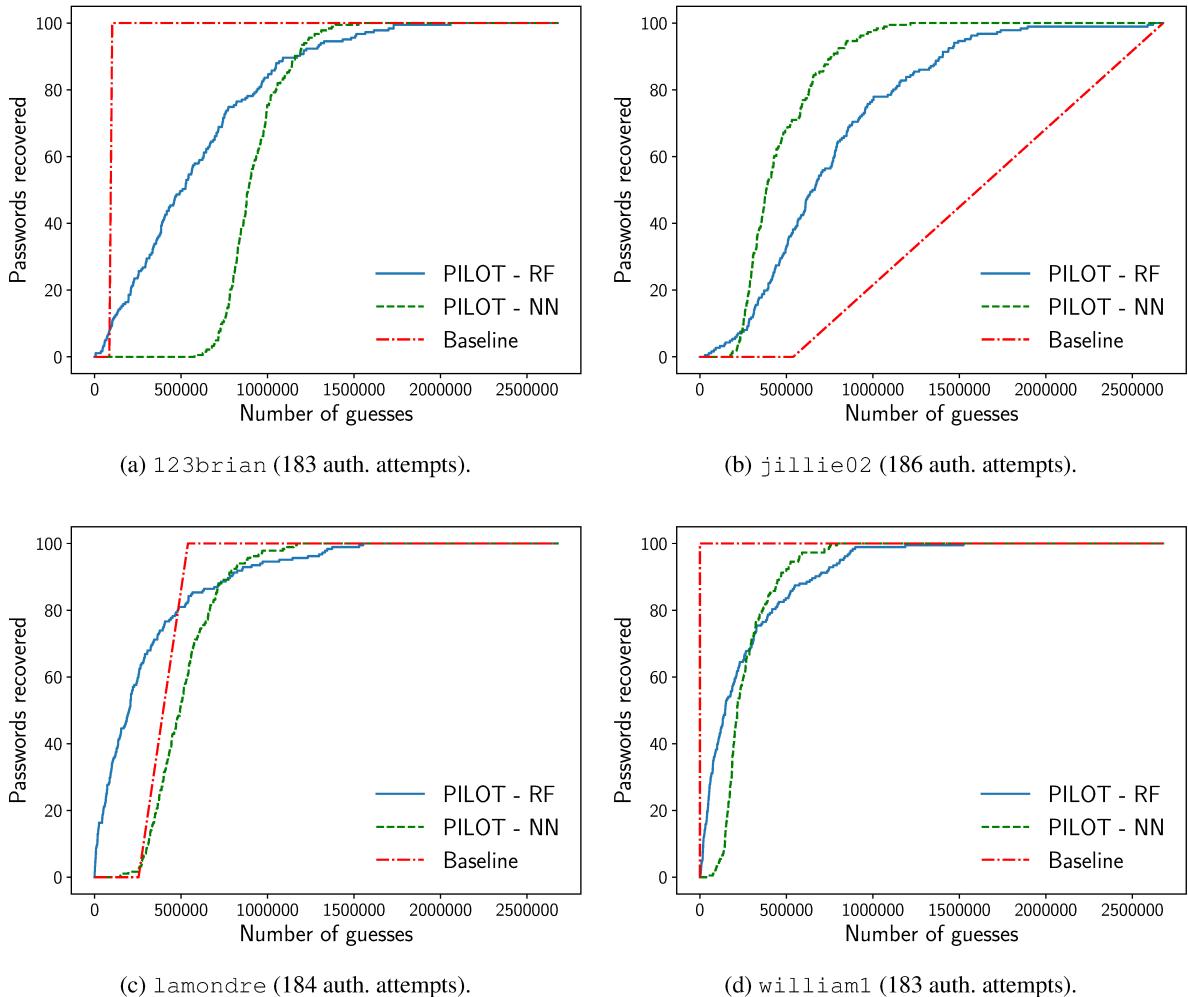


Fig. 8. CDF for the number of passwords recovered by *PILOT*, for each target password. Plots also show the baseline attack for the corresponding password.

6. PIN Guessing

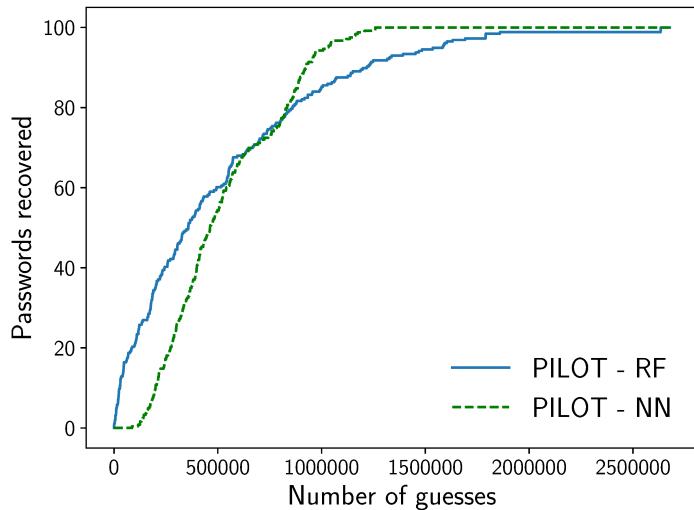
In this section, we present our results on PIN-guessing. First, we analyze the relationship between inter-keystroke timings and Euclidean Distance between consecutive keys, and between inter-keystroke timings and direction of movement on the keypad. We then show how the adversary can use timing information to infer key distances, and therefore to predict PINs.

We are not aware of any publicly-available PIN timing datasets that can be used to train *PILOT*. To address this issue, we collected data from 22 users. To compute the attack baseline, we considered all PINs to be equally likely, i.e., we are modeling PINs as random four-digit strings. This is consistent with how many European banks assign PINs to bank cards [13], and with the work of Bonneau et al. [8], which showed that users are reluctant to change the random PIN provided by their bank.

Table 2

PILOT—*Multiple recordings* setting. *Avg*: average number of attempts to guess a password; *Stdev*: standard deviation; *Rnd*: number of guesses for the baseline adversary; *<Rnd*: how often *PILOT* outperforms random guessing; *Best*: number of attempts of the best guess; *< n*: how many passwords are successfully guessed within first $n = 20,000/100,000$ attempts.

	Avg	Stdev	Med	Rnd	<Rnd	Best	<20k	<100k
Random Forest								
123brian	552,574	468,539	402,166	93,874	14.1%	13,931	4.7%	14.1%
jillie02	713,895	410,225	606,403	1,753,571	100.0%	67,875	0.0%	1.6%
lamondre	398,186	425,811	236,905	397,213	65.6%	404	6.2%	25.0%
william1	370,933	602,654	148,405	187	1.6%	19	17.2%	42.2%
Neural Network								
123brian	922,655	129,927	889,406	93,874	0.0%	676,418	0.0%	0.0%
jillie02	439,414	155,385	402,332	1,753,571	100.0%	205,645	0.0%	0.0%
lamondre	503,248	137,276	504,493	397,213	21.3%	182,123	0.0%	0.0%
william1	248,769	103,240	216,630	187	0.0%	86,213	0.0%	1.6%

Fig. 9. CDF showing number of passwords recovered by *PILOT* in the *Multiple recordings* scenario.

Distance. Across all subjects, we observed that distributions of inter-keystroke latencies were distinct in all cases (for p -value $< 5 \cdot 10^{-6}$), with the following exceptions: (1) latencies for distance 2 (e.g., keypair 1-3) were close to latencies for distance 3 (keypair 2-0); (2) latencies for distance 2 were close to latencies for diagonal 1×1 (e.g., keypair 4-8); latencies for distance 3 were close to latencies for 2×1 diagonal (i.e. “2” to “9”, “1” to “6”, etc.), and diagonal 2×2 (e.g., keypair 7-3), and diagonal 3×2 (e.g., keypair 3-0). Figure 10a shows the various probability distributions, while Figure 10b models these different probability

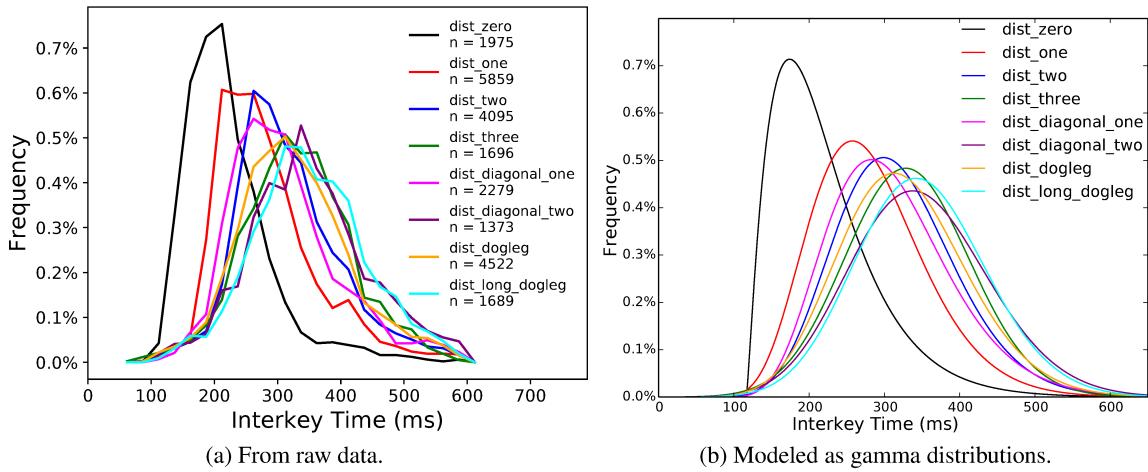


Fig. 10. Inter-keystroke timings of all possible distances for ATM keypad typing.

distribution functions as gamma distributions. In Figure 10a, *dist_zero* indicate keypairs composed of the same two digits. *dist_one*, *dist_two*, and *dist_three* shows timings distributions for keypairs with horizontal or vertical distance one (e.g., keypair 2-5), two (e.g., 2-8), and three (2-0), respectively. *dist_diagonal_one* and *dist_diagonal_two* indicates keypairs with diagonal distance one (e.g., 2-4) and distance two (e.g., 1-9), respectively. *dist_dogleg* and *dist_long_dogleg* show timing distributions of keypairs such as 1-8 and 0-3. In Figure 10b, *dist_one_horizontal* and *dist_one_vertical* indicate Euclidean Distance right in the left/right directions, and up/down directions, respectively, while *dist_one_up*, *dist_one_down*, *dist_one_left*, and *dist_one_right* indicate distances one in the up, down, left, and right directions.

Direction. The relative orientation of keypairs characterized by the same Euclidean distance (e.g., 2-3 vs. 2-5) has a negligible impact on the corresponding inter-key latency. We observed that the distributions of keypress latencies observed from each possible direction between keys were not significantly different (for p -value $< 10^{-4}$). Figure 11 shows different probability distributions relative to various directions for Euclidean distance 1.

PIN Inference Using the data we collected, we mapped the distribution of inter-keypress latencies, and used the resulting probabilities to test the effectiveness of PINs prediction from inter-key latencies. Our PIN guessing algorithm is composed of two parts: (1) an algorithm that estimates distances from keystroke timings; and (2) an algorithm that ranks PINs based on the estimated distances. The source code of the algorithm can be found at <https://spritz.math.unipd.it/datasets/PILOT/> The core idea is to consider the PIN pad as a weighed multigraph. The graph nodes represent the keys, and are labeled 0-9. Keys are connected by weighted edges. The weight of an edge corresponds to the Euclidean distance between the corresponding keys, using the distance between two adjacent keys (e.g., 1 and 2) as unit. We identified 8 possible distances: zero distance (e.g., key 3 followed by key 3, $weight = 0$); horizontal or vertical distance one (e.g. keys 1-2, $weight = 1$); horizontal or vertical distance two (e.g., keys 1-3, $weight = 2$); vertical distance three (e.g., keys 2-0, $weight = 3$); diagonal distance one (e.g., keys 1-5, $weight = \sqrt{1^2 + 1^2}$); diagonal distance two (e.g., keys 1-9, $weight = \sqrt{2^2 + 2^2}$); short diagonal distance (e.g., keys 1-8, $weight = \sqrt{1^2 + 2^2}$) and long diagonal distance (e.g. keys 1-0, $weight = \sqrt{1^2 + 3^2}$).

For each PIN, we created three sets a subgraphs, indicated as S_1 , S_2 , and S_3 , composed only of the nodes connected by edges with the same weight as the estimated distance. Specifically, S_i contains all

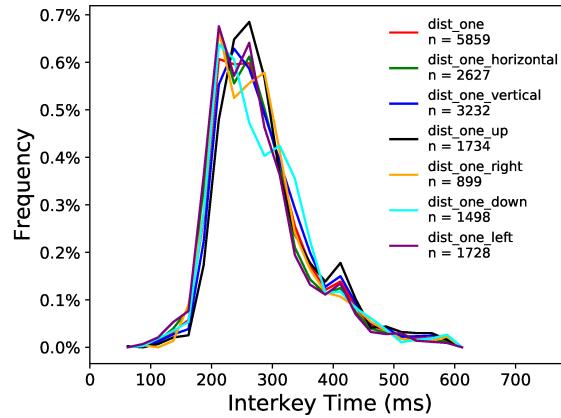


Fig. 11. Frequency of inter-keystroke timings for Euclidean Distance of one. dist_one indicates latency distribution for distance one in any direction.

the two-nodes subgraphs such that their edges have weight equal to the estimated distance between the keys in the i -th PIN digraph.

We combined the subgraphs in these sets by ensuring that, for $i = 1$ and $i = 2$, the second node of a graph from S_i is the same as the first node of a graph from S_{i+1} . For instance, given estimated distances 3, 0, and $\sqrt{2}$, our algorithm extracts the subgraphs shown in Figure 12. It then refines these choices by removing all subgraphs from S_2 which do not have nodes 2 and 0 as their first node. The same rule is applied to S_3 . The two resulting graphs are shown in Figure 13.

Not all estimated distances correspond to possible PINs. For instance, estimated distances 3, 0, and $\sqrt{8}$ do not match any PIN that can be typed on the pad used for the experiments: distance 3 indicates that the second PIN digit must be either 0 or 2; as a consequence, distance 0 associated to the second PIN digraph restricts the third PIN digit to 0 or 2; however, the set of keypairs with a relative distance of $\sqrt{8}$ (i.e., $\{(1, 9), (7, 3), (9, 1), (3, 7)\}$) does not include keys 0 or 2. Therefore, estimated distances 3, 0, and $\sqrt{8}$ do not lead to any valid PIN. Figure 14 shows a visual representation of this example.

For each triplet of estimated distances, the number of associated PINs may differ. For example, 58 triplets have no associated PIN (e.g., distances 3, 0, and $\sqrt{8}$). The remaining 454 combinations vary from a minimum of 2 associated PINs (57 combinations; e.g., distances 3, 3, and 3 correspond only to PINs 2020 and 0202) to a maximum of 216 PINs (distances 1, 1, and 1).

If the adversary is able to reconstruct the distances between digraphs without errors, this process drastically reduces the number of attempts needed to guess the PIN compared to a random guessing. Figure 15 shows the benefit of this approach in terms of percentage of PINs guessed within a fixed number of attempts. However, due to the overlapping between timing distributions shown in figures 10 and 11, the adversary cannot always estimate distances correctly. To evaluate the impact of distance evaluation errors on PIN guessing, we split our keystroke dataset in two sets. The first (training set) consists of 5195 PIN, typed by 11 participants. The second (testing set) consists of 5135 PIN, typed by a distinct set of 11 participants.

For each PIN in the testing set, we associated a list of triplets of distances sorted by their probability as determined using the gamma distributions. Figure 16 shows the effectiveness of this algorithm compared to random guessing, and compared to the technique used in [4]. Our results show that our technique significantly improves on both random guessing and [4] (see Table 3). In particular, we were able to

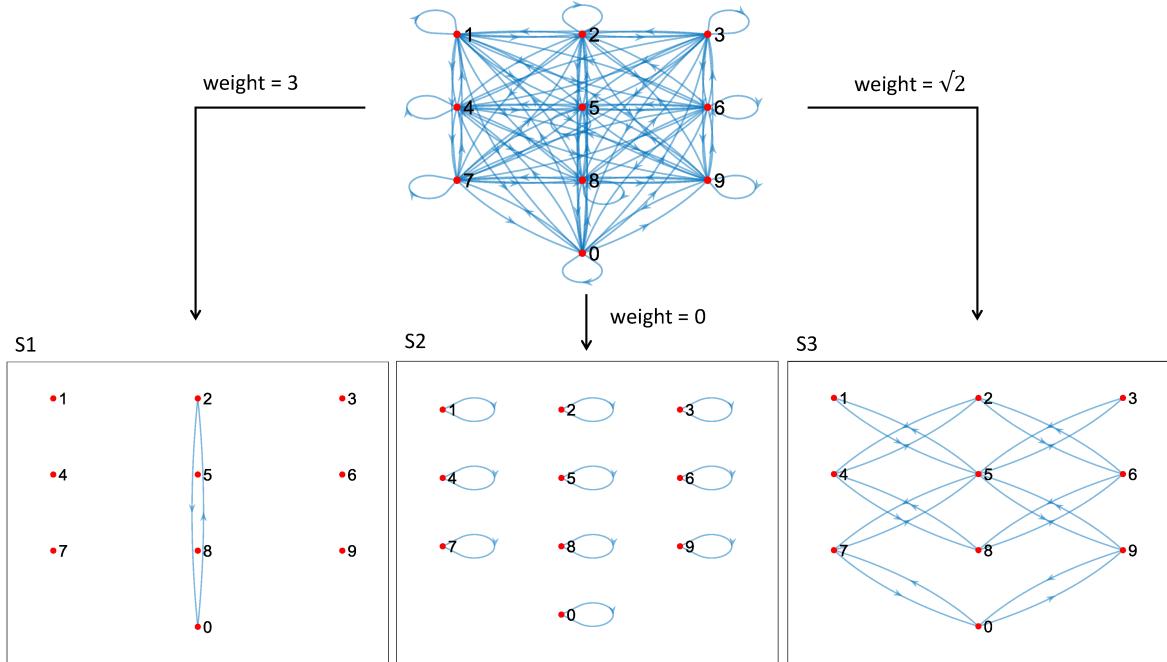


Fig. 12. Full graph, and subgraphs $h_1 \in S_1$ (weight = 3), $h_2 \in S_2$ (weight = 0), and $h_3 \in S_3$ (weight = $\sqrt{2}$).

guess about 3% of the PINs within 10 attempts. This corresponds to a 26-fold improvement compared to random guessing. Moreover, we were able to significantly outperform [4] in terms of number of PINs recovered within a small number of attempts. Specifically, we were able to recover more than 1% of the PINs within 5 attempts, while [4] could guess only 0.32% of the PINs within the same number of attempts. In summary, our technique demonstrates that inter-keystroke timings contain considerable information about the physical distance between consecutive keys in a PIN, and thus they can substantially reduce the number of attempts required to guess a PIN.

7. Conclusion

In this paper, we have shown that the inter-key timing information disclosed by showing password masking symbols can be effectively used to reduce the cost of password guessing attacks. To determine the impact of this side channel, we recorded videos from 84 subjects, typing several passwords and PINs under different conditions: in a lecture hall, while their laptop was collected to a projector; in a classroom setting; and using a simulated ATM machine. Our results show that: (1) it is possible to infer very accurate timing information from videos of LCD screens and projectors (the average error was 8.7ms, which corresponds to the duration of a frame when the refresh rate of a display is set to 60 Hz); (2) inter-keystroke timings reduce the number of attempts to recover a password by 25% and 385%, with some passwords guessed within 19 attempts. We consider this a substantial reduction in the cost of password guessing attacks, to the point that we believe that masking symbols should not be publicly displayed when typing passwords; and (3) disclosing inter-keystroke timings have a significant impact on PIN guessing attacks. In particular, we were able to predict about 1 in 14 PINs within 40 attempts,

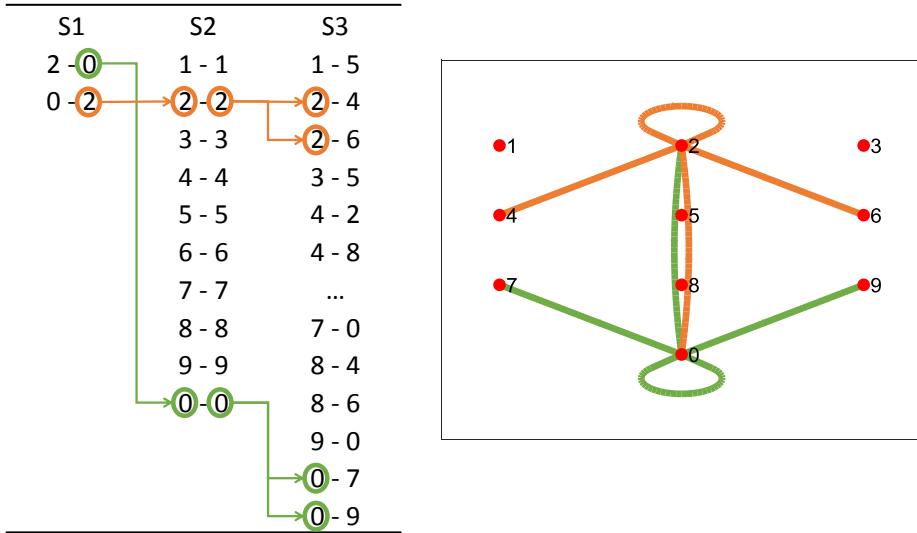
Fig. 13. Nodes sequences for input distances 3, 0, $\sqrt{2}$.

Table 3
Percentage of PIN guessed using our technique, Random Guessing (RG), PIN guessing with exact distance (Exact), and [4]

Attempts	Our work	RG	Exact	SILK-TV [4]	Improvement of our
					work compared to RG
5	1.09%	0.05%	13.05%	0.32%	21.81×
10	2.65%	0.10%	26.37%	0.47%	26.48×
20	3.93%	0.20%	45.57%	0.79%	19.67×
40	7.05%	0.40%	67.37%	1.31%	17.62×
80	9.41%	0.80%	88.85%	2.28%	11.76×
160	12.89%	1.60%	98.34%	3.82%	8.06×
320	19.69%	3.20%	100.00%	6.34%	6.15×
640	28.33%	6.40%	100.00%	10.05%	4.43×
1280	42.30%	12.80%	100.00%	16.34%	3.30×
2560	59.39%	25.60%	100.00%	27.81%	2.32×
5120	78.53%	52.20%	100.00%	51.84%	1.50×

compared to one in 250 with no timing information. While this result is not as dramatic as the one with passwords, it suggests that keystroke timing information should be carefully concealed by ATMs.

Clearly, the benefits of *PILOT* compared to our baseline attack vary depending on how common the user's password is. For very common (and therefore very easy to guess) passwords, our results show

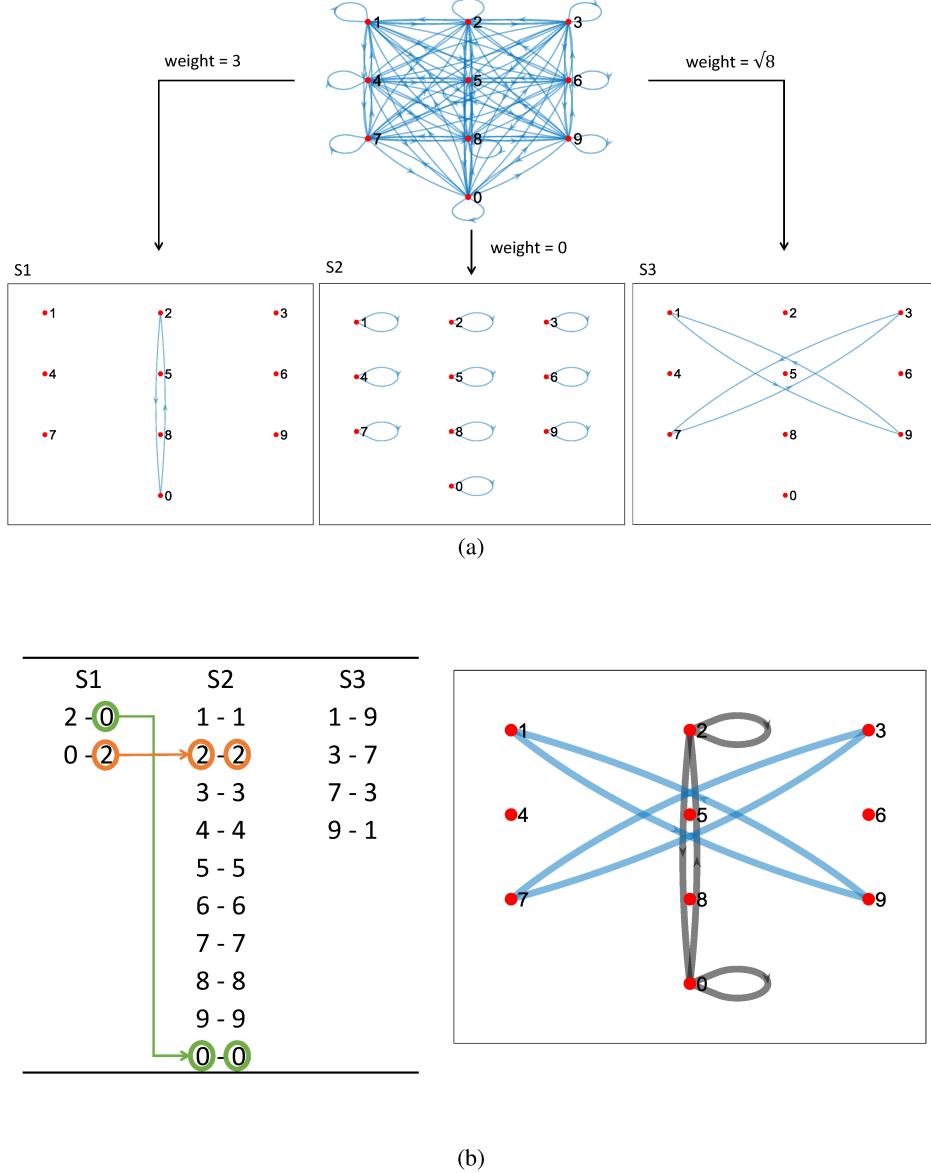


Fig. 14. (a) Subgraphs corresponding to distances 3, 0, and $\sqrt{2^2 + 2^2}$. (b) Nodes connected by vertices weights 3, 0, and $\sqrt{2^2 + 2^2}$. No common nodes are in S_2 and S_3 , and therefore this combination of distances does not correspond to any PIN.

that PILOT might not be needed. On the other hand, the speedup offered by PILOT when guessing rare passwords is substantial. Given the effectiveness of this attack on password guessing, we think that future work should consider countermeasures that strike the right balance between usability and security when displaying masking symbols. For instance, GUIs may not display masking symbols on a secondary screen (e.g., projectors), or may display new masking symbols at fixed intervals (say, every 250ms). Clearly, both countermeasures have usability implications, and we leave the quantification of this impact to future work.

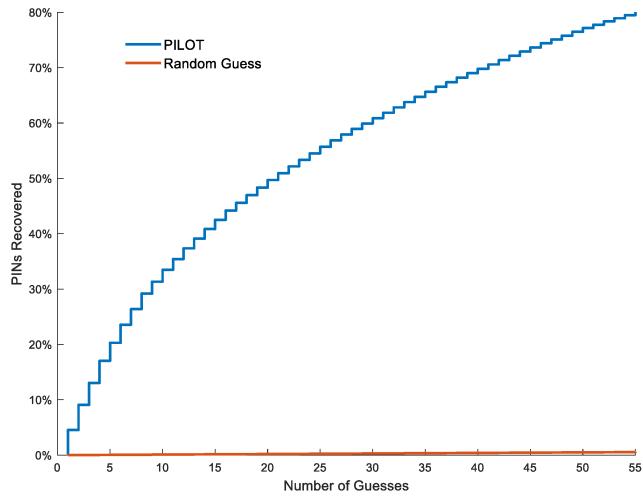


Fig. 15. CDF showing the number of PINs recovered under the assumption that distances are recovered without error.

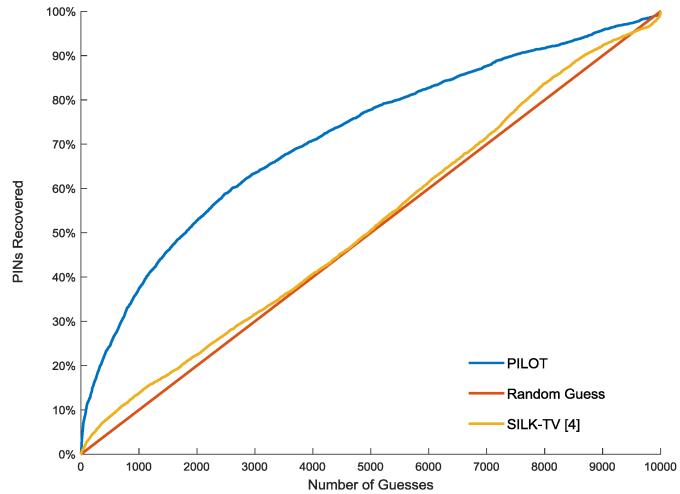


Fig. 16. CDF showing the number of PINs recovered on our dataset, compared to the baseline and the algorithm of [4].

Acknowledgements

Kiran Balagani and Paolo Gasti were supported by the National Science Foundation under Grant No. CNS-1619023. Tristan Gurtler, Charissa Miller, Kendall Molas, and Lynn Wu were supported by the National Science Foundation under Grant No. CNS-1559652. This work is partially supported by the EU TagItSmart! Project (agreement H2020-ICT30-2015-688061), and the EU-India REACH Project (agreement ICI+/2014/342-896).

References

- [1] Rockyou password leak (2010), <http://downloads.skullsecurity.org/passwords/rockyou.txt.bz2>
- [2] Linkedin password leak (2016), <https://hashes.org/leaks.php>
- [3] Asonov, D., Agrawal, R.: Keyboard acoustic emanations. In: IEEE Symposium on Security and Privacy, 2004. Proceedings. 2004. pp. 3–11. IEEE (2004)
- [4] Balagani, K.S., Conti, M., Gasti, P., Georgiev, M., Gurtler, T., Lain, D., Miller, C., Molas, K., Samarin, N., Saraci, E., et al.: Silk-tv: Secret information leakage from keystroke timing videos. In: European Symposium on Research in Computer Security. pp. 263–280. Springer (2018)
- [5] Balzarotti, D., Cova, M., Vigna, G.: Clearshot: Eavesdropping on keyboard input from video. In: 2008 IEEE Symposium on Security and Privacy (sp 2008). pp. 170–183. IEEE (2008)
- [6] Banerjee, R., Feng, S., Kang, J.S., Choi, Y.: Keystroke patterns as prosody in digital writings: A case study with deceptive reviews and essays. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP). pp. 1469–1473 (2014)
- [7] Bartlow, N., Cukic, B.: Evaluating the reliability of credential hardening through keystroke dynamics. In: 2006 17th International Symposium on Software Reliability Engineering. pp. 117–126. IEEE (2006)
- [8] Bonneau, J., Preibusch, S., De Andeson, R.: A birthday present every eleven wallets? the security of customer-chosen banking pins. In: International Conference on Financial Cryptography and Data Security. pp. 25–40. Springer (2012)
- [9] Burnett, M.: Today i am releasing 10 million passwords (2015), <https://xato.net/today-i-am-releasing-ten-million-passwords-b6278bbe7495>
- [10] Compagno, A., Conti, M., Lain, D., Tsudik, G.: Don't skype & type!: Acoustic eavesdropping in voice-over-ip. In: Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security. pp. 703–715. ACM (2017)
- [11] Ding, L., Goshtasby, A.: On the canny edge detector. Pattern Recognition **34**(3), 721–725 (2001)
- [12] Fiegerman, S.: Yahoo says 500 million accounts stolen (2017), <http://money.cnn.com/2016/09/22/technology/yahoo-data-breach/index.html>
- [13] FINECO: Le carte fineco (2019), <https://help.finecobank.com/it/conto-e-carte/le-carte-fineco.html#nuove-commissioni>
- [14] Florencio, D., Herley, C.: A large-scale study of web password habits. In: Proceedings of the 16th international conference on World Wide Web. pp. 657–666. ACM (2007)
- [15] Hitaj, B., Gasti, P., Ateniese, G., Perez-Cruz, F.: Passgan: A deep learning approach for password guessing. arXiv preprint arXiv:1709.00440 (2017)
- [16] Ho, T.K.: Random decision forests. In: Proceedings of 3rd international conference on document analysis and recognition. vol. 1, pp. 278–282. IEEE (1995)
- [17] Koliاس, C., Kambourakis, G., Stavrou, A., Voas, J.: Ddos in the iot: Mirai and other botnets. Computer **50**(7), 80–84 (2017)
- [18] Ma, J., Yang, W., Luo, M., Li, N.: A study of probabilistic password models. In: 2014 IEEE Symposium on Security and Privacy. pp. 689–704. IEEE (2014)
- [19] Owusu, E., Han, J., Das, S., Perrig, A., Zhang, J.: Accessory: password inference using accelerometers on smartphones. In: Proceedings of the Twelfth Workshop on Mobile Computing Systems & Applications. p. 9. ACM (2012)
- [20] Pulli, K., Baksheev, A., Korniyakov, K., Eruhimov, V.: Real-time computer vision with opencv. Communications of the ACM **55**(6), 61–69 (2012)
- [21] Roth, J., Liu, X., Metaxas, D.: On continuous user authentication via typing behavior. IEEE Transactions on Image Processing **23**(10), 4611–4624 (2014)
- [22] Schalkoff, R.J.: Artificial neural networks, vol. 1. McGraw-Hill New York (1997)
- [23] Shukla, D., Kumar, R., Serwadda, A., Phoha, V.V.: Beware, your hands reveal your secrets! In: Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security. pp. 904–917. ACM (2014)
- [24] Song, D.X., Wagner, D.A., Tian, X.: Timing analysis of keystrokes and timing attacks on ssh. In: USENIX Security Symposium. vol. 2001 (2001)
- [25] Sun, J., Jin, X., Chen, Y., Zhang, J., Zhang, Y., Zhang, R.: Visible: Video-assisted keystroke inference from tablet backside motion. In: NDSS (2016)
- [26] Tari, F., Ozok, A., Holden, S.H.: A comparison of perceived and real shoulder-surfing risks between alphanumeric and graphical passwords. In: Proceedings of the second symposium on Usable privacy and security. pp. 56–66. ACM (2006)
- [27] The Password Project: (2017), http://thepasswordproject.com/leaked_password_lists_and_dictionaries
- [28] Tomasi, C., Manduchi, R.: Bilateral filtering for gray and color images. In: Iccv. vol. 98, p. 2 (1998)
- [29] Vural, E., Huang, J., Hou, D., Schuckers, S.: Shared research dataset to support development of keystroke authentication. In: IEEE International Joint Conference on Biometrics. pp. 1–8. IEEE (2014)
- [30] Wang, C., Guo, X., Wang, Y., Chen, Y., Liu, B.: Friend or foe?: Your wearable devices reveal your personal pin. In: Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security. pp. 189–200. ACM (2016)

- [31] Wang, C., Jan, S.T., Hu, H., Bossart, D., Wang, G.: The next domino to fall: Empirical analysis of user passwords across online services. In: Proceedings of the Eighth ACM Conference on Data and Application Security and Privacy. pp. 196–203. ACM (2018)
- [32] Weir, M., Aggarwal, S., De Medeiros, B., Glodek, B.: Password cracking using probabilistic context-free grammars. In: 2009 30th IEEE Symposium on Security and Privacy. pp. 391–405. IEEE (2009)
- [33] Xu, Y., Heinly, J., White, A.M., Monroe, F., Frahm, J.M.: Seeing double: Reconstructing obscured typed input from repeated compromising reflections. In: Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security. pp. 1063–1074. ACM (2013)
- [34] Zhu, T., Ma, Q., Zhang, S., Liu, Y.: Context-free attacks using keyboard acoustic emanations. In: Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security. pp. 453–464. ACM (2014)
- [35] Zhuang, L., Zhou, F., Tygar, J.D.: Keyboard acoustic emanations revisited. ACM Transactions on Information and System Security (TISSEC) **13**(1), 3 (2009)