MagHacker: Eavesdropping on Stylus Pen Writing via Magnetic Sensing from Commodity Mobile Devices

Yihao Liu, Kai Huang, Xingzhe Song, Boyuan Yang and Wei Gao
University of Pittsburgh

ABSTRACT

Stylus pens have been widely used with today's mobile devices to provide a convenient handwriting input method, but also bring a unique security vulnerability that may unveil the user's handwriting contents to a nearby eavesdropper. In this paper, we present *MagHacker*, a new sensing system that realizes such eavesdropping attack over commodity mobile devices, which monitor and analyze the magnetic field being produced by the stylus pen's internal magnet. MagHacker divides the continuous magnetometer readings into small segments that represent individual letters, and then translates these readings into writing trajectories for letter recognition. Experiment results over realistic handwritings from multiple human beings demonstrate that MagHacker can accurately eavesdrop more than 80% of handwriting with stylus pens, from a distance of 10cm. Only slight degradation in such accuracy is produced when the eavesdropping distance or the handwriting speed increases. MagHacker is highly energy efficient, and can well adapt to different stylus pen models and environmental contexts.

CCS CONCEPTS

• **Security and privacy → Mobile and wireless security**.

KEYWORDS

Magnetic sensing, Eavesdropping, Stylus pen, Smartphones, Coordinate transformation

ACM Reference Format:
Yihao Liu, Kai Huang, Xingzhe Song, Boyuan Yang and Wei Gao. 2020. MagHacker: Eavesdropping on Stylus Pen Writing via Magnetic Sensing from Commodity Mobile Devices. In *The 18th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '20), June 15–19, 2020, Toronto, ON, Canada.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3386901.3389030

1 INTRODUCTION

In recent years, stylus pens have emerged as an important accessory to many mobile devices, including laptop computers, tablets and smartphones. As shown in Figure 1, a stylus pen can be used for handwriting on the LCD touchscreen as an efficient input method in many scenarios, such as quick notes, art design, business demonstration and emerging edge computing applications [29, 36]. Commercial stylus pens, such as Microsoft Surface pen [3], are equipped with internal magnets and can be attached to the metal shell of mobile devices for good mobility.

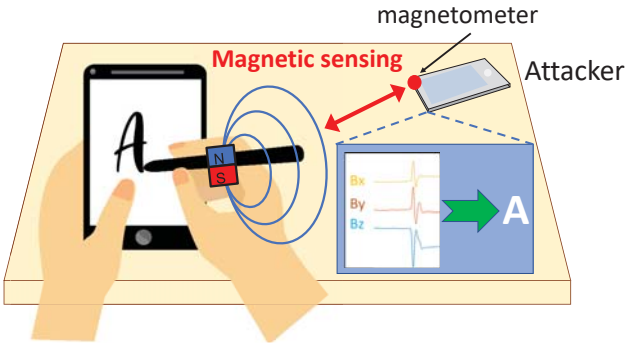

Figure 1: Evaesdropping on stylus pen writing

However, using stylus pens may also bring potential security vulnerability, which allows an attacker to eavesdrop on handwriting inputs. The basic rationale of such vulnerability, as shown in Figure 1, is that the movement of stylus pen's internal magnet changes the nearby magnetic field. This change could be captured by a magnetometer and analyzed to infer the handwriting. Compared to existing methods of eavesdropping attacks that retrieve and analyze the IMU data from the victim's on-body devices [22, 37], this new attack can be reliably launched from remote without compromising the victim's personal mobile devices.

The unique characteristics of digital magnetometers, on the other hand, significantly reduce the difficulty of launching such attack in practice. First, the magnetic field is known to be consistent against most types of environmental dynamics. Hence, being different from current eavesdropping attacks that either capture the audible sound of handwriting [15, 16, 42] or the handwriting's disturbance to wireless signals [28, 41, 44], the magnetic-based attack can effectively resist against the impact of ambient noise, multipath effects and nearby moving objects, ensuring high accuracy in eavesdropping. Second, the low cost, small form factor and low power consumption of digital magnetometer lead to its wide availability on today's smartphones and wearables (e.g., wristbands and smartwatches). An attacker can then easily conceal the eavesdropping device even being very close to the victim, without requiring line of sight on the attacker as current camera attacks do [13, 23, 24, 34].

In this paper, we present *MagHacker*, a new sensing system that realizes such eavesdropping attack, which assumes that the victim writes in English and aims to recognize each individual English letter of his/her handwriting from distance. Being different

The first two authors contributed equally to the paper.

MobiSys '20, June 15–19, 2020, Toronto, ON, Canada

ACM ISBN 978-1-4503-7954-0/20/06…$15.00
https://doi.org/10.1145/3386901.3389030

	Accuracy	Environment adaptability	Hardware cost	Concealment	Attack distance	Response delay
IMU sensing [22, 37]	Low	Low	Low	None	None	Low
Audio capture [15, 16, 42]	Low	Very Low	Medium	Medium	Medium	Medium
Camera capture [13, 23, 24, 34]	High	Medium	High	Very Low	Low	Very High
Wireless sensing [28, 41, 44]	High	Low	Very High	High	Medium	High
Magnetic sensing [11, 12, 32, 40]	Medium	High	High	Medium	Low	Medium
MagHacker	**High**	**Very High**	**None**	**Very High**	**Medium**	**Low**

Table 1: Comparison of the approaches to eavesdropping on handwriting

from existing work on magnetic sensing that uses multiple custom magnetometers to track the magnet's 3D position [11, 12, 32, 40], MagHacker only uses the readings from a single magnetometer on the commodity mobile device operated by the attacker, and requires no extra hardware.

To ensure accuracy, MagHacker divides the continuous magnetometer readings into small segments, each of which corresponds to a letter being written and is individually recognized. This segmentation, however, is challenging because of the heterogeneous efforts to write different letters[1], which make it hard to estimate the duration of writing each letter. Instead, MagHacker exploits the fluctuation of humans' speed of hand movement in writing, which changes less frequently in transition between letters [27, 31]. MagHacker calculates the victim's speed of hand movement from magnetometer readings and applies Continuous Wavelet Transform (CWT) to the time series data of movement speed. The produced spectrogram, then, represents the frequency of speed changes and can be used for segmentation.

The readings from digital magnetometers, on the other hand, are given in form of 3D magnetic field strengths and may not directly reflect the handwriting trajectory (i.e., the movement of magnet in stylus pen). In particular, the correspondence between magnetometer readings and magnet movement depends on the relative positioning between the magnetometer and the magnet, and is non-linear in most cases. MagHacker addresses this challenge with coordinate projection and transformation: the 3D magnetometer readings of each handwritten word is first projected to a 2D plane with the minimum distortion, and the 2D magnetometer readings of each letter in the word are then separately transformed to the magnet's coordinate system. Afterward such projection and transformation, the magnetometer readings represent the handwriting trajectories and are applied to a Convolutional Neural Network (CNN) classifier for recognition.

To the best of our knowledge, MagHacker is the first work that uses a single magnetometer on commodity mobile devices to eavesdrop on humans' handwriting. It widely applies to different types of mobile devices ranging from smartphones to digital wearables, and unveils an important security vulnerability that may potentially leak mobile users' personal sensitive information at very low attacking costs. For example, people's personal communication with others may be exposed to a nearby attacker, and the leakage of important credentials could be vital to the victim's personal and financial safety. Our detailed technical contributions are as follows.

- We quantitatively investigated different frequency components on the CWT spectrogram to ensure appropriate segmentation of magnetometer readings, and effectively removed the impact of humans' heterogeneous hand motions and writing speeds.
- We minimized the distortion of handwriting trajectories being translated from magnetometer readings, by adapting the coordinate projection and transformation to the unique condition of each handwritten letter.
- We incorporated the humans' heterogeneous handwriting patterns and habits into account when designing the CNN classifier for letter recognition, and effectively prevented the classifier from overfitting by augmenting the training data.

We have implemented and tested MagHacker over a wide collection of commodity mobile devices (e.g., Apple iPhone Xs and Google Pixel 2 XL smartphones) and stylus pen products (e.g., Microsoft Surface Pen, Adonit Snap 2 Pen and Maglus Pen), and evaluated the eavesdropping performance of MagHacker with five student volunteers who handwrote all the possible combinations and transitions of English letters. From our experiment results, we have the following conclusions:

- MagHacker is highly *accurate*. MagHacker can correctly recognize more than 80% of lowercase and uppercase letters being written by all experiment participants, with different writing patterns and habits.
- MagHacker is highly *adaptive*. MagHacker can well adapt to the heterogeneous writing speeds of humans, and can achieve good accuracy of eavesdropping over different stylus pen models. It can also retain the accuracy of eavesdropping in different environmental conditions, even with nearby metal objects.
- MagHacker is widely *applicable*. MagHacker allows the attacker to effectively eavesdrop within the physical proximity of the victim, and achieve good accuracy with an eavesdropping distance up to 20cm. It also incurs the minimum power consumption on the attacking device and hence allows several hours of continuous eavesdropping.

2 BACKGROUND AND DESIGN CHALLENGES

To better understand the design of MagHacker, we first describe the threat model and technical background of magnetic sensing. We

[1]Writing different English letters involves different numbers of strokes. For example, some simple letters like 'c' only involve one stroke but others like 'm' and 'w' involve three or more strokes.

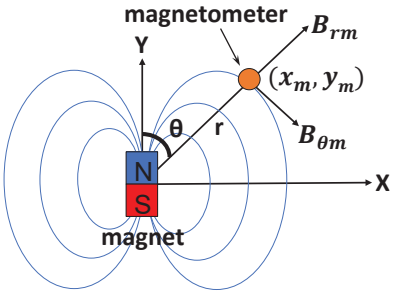

Figure 2: Magnetic sensing

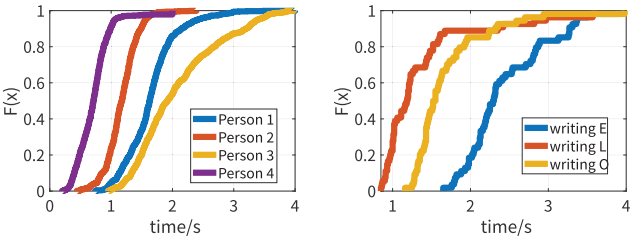

(a) Different persons on the same letter (b) One person writing different letters

Figure 3: Cumulative distribution functions (CDFs) of durations of writing different letters

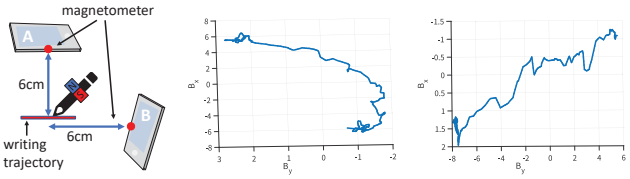

(a) Experiment setup (b) Reading from device A (c) Reading from device B

Figure 4: Magnetometer readings

then motivate our design by demonstrating the technical challenges raised by the characteristics of magnetic sensing.

2.1 Threat Model

MagHacker considers that an attacker uses a mobile device with digital magnetometer to eavesdrop in the physical proximity of victim. The attacking device is well concealed, so that the victim is unaware of eavesdropping even if the attacking device is in his/her line of sight. In practice, such concealment can be achieved by running a background process at the attacking device that retrieves and analyzes magnetometer readings without any screen display.

MagHacker does not require the attacker to compromise the victim's device, nor to physically contact the victim's body, stylus pen or the writing surface. We assume that the attacker has line of sight to the victim and can monitor the victim's body movement. Such knowledge about the victim's body movement is far from sufficient to derive the victim's handwriting, but allows the attacker to tell when the victim starts and stops handwriting with the stylus pen. The attacker can then remotely start and stop the eavesdropping accordingly, without physically touching the eavesdropping device and impairing the attacker's concealment.

2.2 Magnetic Sensing

Digital magnetometers on commodity mobile devices (e.g., smartphones and smartwatches) can produce readings with a high sampling rate of 100Hz, which is sufficient to capture the magnet's subtle movements. Figure 2 illustrates how the magnetic field emanated from a magnet is being sensed by a magnetometer, in the magnet's coordinate system where the magnet stands at $(0, 0)$. Since the magnet in the stylus pen is usually a bar magnet, its emanated magnetic field is symmetric about its Y axis and has zero strength at the Z axis. Hence, without loss of generality, we project the magnetometer's position to the magnet's X-Y plane as (x_m, y_m). Then, the magnetic field at (x_m, y_m) can be decomposed into tangential and radial components as

$$\begin{cases} B_{rm} & = \frac{\mu_0 M cos\theta}{2\pi r^3} = \frac{\mu_0 M r cos\theta}{2\pi r^4} = \frac{\mu_0 M}{2\pi r^4} y_m \\ B_{\theta m} & = \frac{\mu_0 M sin\theta}{4\pi r^3} = \frac{\mu_0 M r sin\theta}{4\pi r^4} = \frac{\mu_0 M}{4\pi r^4} x_m, \end{cases} \quad (1)$$

where $r = \sqrt{x_m^2 + y_m^2}$ is the distance between the magnet and the magnetometer, and θ is the magnetometer's orientation with respect to the magnet's north pole. M is the magnetic moment that is usually constant for permanent magnets [25]. μ_0 is the permeability constant.

The magnetometer's readings about the 3D magnetic field strengths, as $[B_x, B_y, B_z]$, are produced in the magnetometer's local coordinate system. They can be converted to the quantities in Eq. (1) via coordinate transformation as

$$[B_{rm}, B_{\theta m}, 0]' = \mathbb{T} \cdot [B_x, B_y, B_z]', \quad (2)$$

where $\mathbb{T}$ is the transformation matrix that is determined by the Euler angle between the magnetometer's and the magnet's coordinate systems.

2.3 Design Challenges

To divide the continuous magnetometer readings into segments that correspond to individual letters, one intuitive solution is to use the duration of writing individual letters. To investigate such writing durations in practice, we monitored four graduate student volunteers to hand write 27 times of the 26 uppercase English letters. The distributions of writing durations, as shown in Figure 3, are highly heterogeneous with >70% variation: one may spend different amounts of time to write the same letter, and different people exhibit different paces when writing the same letter. Such heterogeneity, then, motivates us to seek for other reliable characteristics in the magnetometer readings for segmentation.

After segmentation, the major challenge of eavesdropping is how to convert the magnetometer readings to the trajectory of magnet's movement, because the magnetometer readings are jointly determined by the position and orientation of both the magnetometer and the magnet of stylus pen. The same magnet's movement, hence, may result in different magnetometer readings. To demonstrate such difference, we conducted a preliminary experiment, as shown in Figure 4(a), to draw a straight line using a Microsoft Surface pen and sense the magnet's movement with two magnetometers in different orientations. The magnetometer readings, as shown in

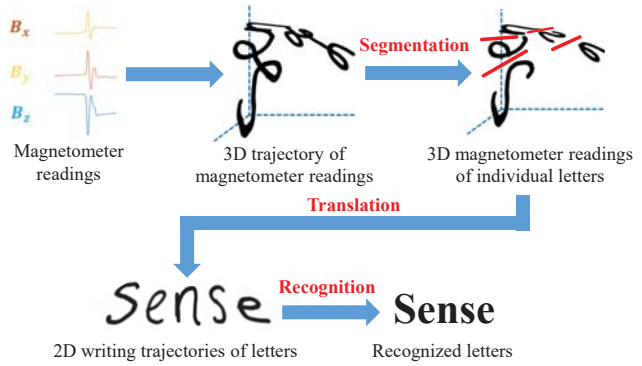

Figure 5: Overall design of MagHacker

Figure 4(b) and 4(c), are highly different from each other. In particular, the magnetometer readings from device A is also distorted from a straight line, because the angle between A's magnetometer and the magnet's north pole (θ in Figure 2) keeps changing and results non-linear fluctuations of magnetic field. MagHacker addresses this challenge by investigating appropriate coordinate transformations between the magnetometer's and the magnet's coordinate systems.

3 MAGHACKER DESIGN

As illustrated in Figure 5, MagHacker first divides the continuous 3D magnetometer readings into small segments, each of which corresponds to a letter being written. Afterwards, it individually translates the magnetometer readings for each letter to the corresponding handwriting trajectory via projection and coordinate transformation, to remove the distortions in the raw magnetometer readings. Such trajectory, then, is sent to a CNN classifier to be recognized.

3.1 Segmenting Magnetometer Readings

MagHacker segments the continuous magnetometer readings in two steps: first, it counts the number of letters and roughly locates each letter in the time series of magnetometer readings; then, it uses such rough knowledge about each letter to identify the boundaries of the corresponding segment of magnetometer readings. As shown in Figure 6(a), such segmentation builds on the fact that the speed of human's hand movement in writing increases in continuous strikes and decreases in turning points [27]. Since each English letter contains multiple strokes, the speed of hand movement changes more frequently when writing letters, but exhibits fewer changes in transition between letters and words.

To correctly measure heterogeneous frequencies of speed changes, for time series of magnetometer readings $(B_x(t), B_y(t), B_z(t))$, we first compute its changing speed as

$$v_t^2 = (B_x(t+1) - B_x(t))^2 + (B_y(t+1) - B_y(t))^2 + (B_z(t+1) - B_z(t))^2,$$

and then apply Continuous Wavelet Transform (CWT)[2] on the time series of $\{v_t\}$. From the spectrogram produced by CWT, let x_{ft} indicate the amplitude of its frequency component f at time t, we can compute the mean amplitude of all its N different frequency components as

$$\mu_t = \sum_{f=f_{\min}}^{f_{\max}} x_{ft}/N,$$

and each peak in the time series of $\{\mu_t\}$ corresponds to a segment with highly frequent speed changes that indicates a different letter. The valleys of $\{\mu_t\}$, on the other hand, separate between letters and words. For example, the spectrogram corresponding to the handwriting of two words 'NO' and 'WAY', as shown in Figure 6(b), exhibits obvious separation between letters and the two words. The mean amplitudes, furthermore, exhibit 5 peaks as shown in Figure 6(c).

Since each letter is located by a peak of $\{\mu_t\}$, an intuitive idea is to use the valleys of $\{\mu_t\}$ as boundaries of segments, as shown in Figure 6(c). This approach, however, may be ineffective in practice, because of the ignorance of transiting period between letters. Instead, more appropriate boundaries of each letter segment should be in the middle between each peak and valley of $\{\mu_t\}$. Details of deciding such boundaries will be described in Section 4.

3.2 Translation into Writing Trajectories

After segmentation, MagHacker projects the 3D magnetometer readings of each handwritten word onto a 2D plane, and then translates the segment of each letter in this word to the handwriting trajectory. An intuitive approach is to project to the $B_x - B_y$ plane in the magnetometer's coordinate system, but may largely distort the produced handwriting trajectory due to the inconsistency between magnetometer readings and magnet's movement as shown in Section 2.3. For example, when the 3D magnetometer readings in Figure 7(a) are projected to the $B_x - B_y$ plane, the projected trajectories can be hardly recognizable as shown in Figure 7(b).

Instead, our approach, as shown in Figure 7(a), is to choose the topmost (with the maximum B_y), bottommost (with the minimum B_y), leftmost (with the minimum B_x) and rightmost (with the maximum B_x) points (P_1, P_2, P_3, P_4) on 3D magnetometer readings, and use the 2D plane defined by vectors $\overrightarrow{(P_1, P_2)}$ and $\overrightarrow{(P_3, P_4)}$ for projection. Since this plane intersects with most of the 3D magnetometer readings, using it for projection significantly reduces the distortion as shown in Figure 7(c).

The major challenge of translating the 2D magnetometer readings of an individual letter, as shown in Figure 7(c), is that the centroids of different letters (indicated by the red dots) in the word form a curvy line, even though humans usually write each word along a straight line. Such distortion is caused by the non-linear correspondence between the magnetometer readings and magnet's movement as shown in Section 2.3, and requests for different coordinate transformations to be applied to different letters. Details of such coordinate transformation will be described in Section 5.

4 DECIDING LETTER BOUNDARIES

An intuitive approach to deciding the boundaries of magnetometer readings for different letters is to find the highest gradient among different frequency components in the CWT spectrogram.

[2] Although Short Time Fourier Transform(STFT) can analyze such speed changes over time, its performance depends on the choice of time window. However, using a fixed time window cannot analyze letters with heterogeneous writing durations. Instead, CWT provides better flexibility on the time scale of frequency-domain analysis through its scaling factor [18].

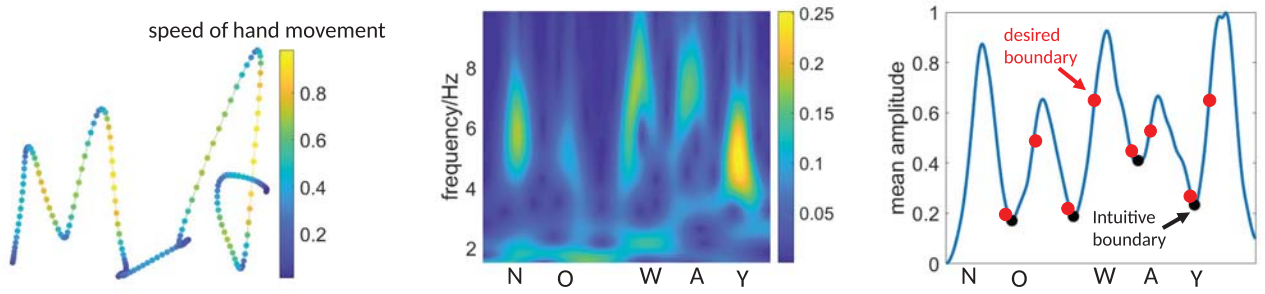

(a) Speed of hand movement while writing two letters 'MA'

(b) CWT spectrogram of magnetoemter readings for two words 'NO' and 'WAY'

(c) Mean amplitude of different frequency components in the CWT spectrogram

Figure 6: Segmentation of magnetometer readings

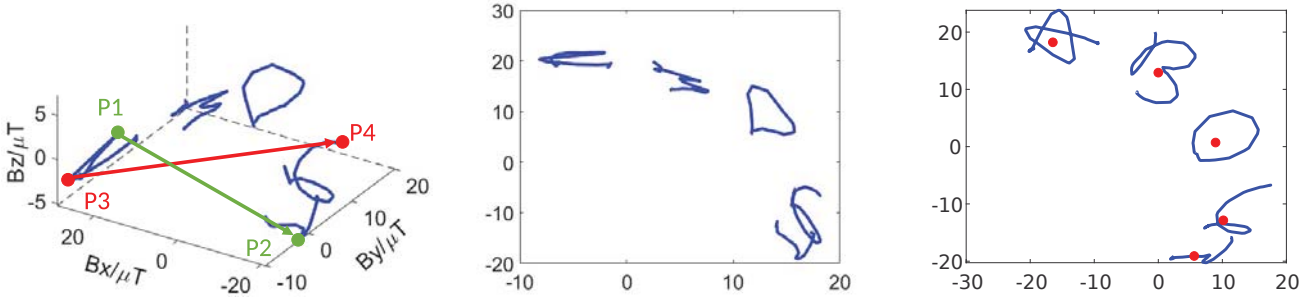

(a) 3D magnetometer readings of 'ABOUT'

(b) 2D projection onto the $B_x - B_y$ plane

(c) 2D projection in MagHacker

Figure 7: Translating the magnetometer readings of a word into handwriting trajectories

As shown in Figure 8(a), such highest gradients are mostly found when starting and finishing writing letters, due to the sudden change in the speed of hand movement. However, such gradient-based approach may be inaccurate in some cases, because the highest gradient only corresponds to a single frequency component in the CWT spectrogram. The produced boundaries, as shown in Figure 8(a), are hard to be converted to specific time points.

Instead, as shown in Figure 8(b), MagHacker addresses this problem by investigating the mean amplitudes of different frequency components in the CWT spectrogram, and decides letter boundaries as the point with the highest derivative over time. More specifically, the average writing frequency of humans is 68 letters per minute [9], and each letter has 2-5 strokes. Thus in practice, we calculate such mean amplitudes over frequency components between 1.5Hz and 10Hz in the CWT spectrogram. The highest derivative after each valley of mean amplitudes can be considered as the starting point of writing the next letter, and such before each valley of mean amplitudes is the ending point of the previous letter.

4.1 Removing Humans' Hand Motion

In some cases, humans' hand will have extra motions when transiting between writing different letters: after completion of the previous letter, the writer's hand will lift up the pen and move horizontally before pinning down to the writing surface for the next letter. As a result, as shown in Figure 8(c), the highest derivative is usually found in the period of hand motion due to the intense change in the speed of hand movement, and involves irrelevant magnetometer readings of hand motion into the segmented writing trajectories.

To solve this problem, we will detect such time periods of hand motions with a threshold. If the speed of hand movement at the highest derivative is higher than this threshold, the hand motion is detected. For example in Figure 8(c), the highest derivative at H_2 is found to be during hand motion. The valley points before and after H_2, namely A and B, are then detected as the starting and ending points of such hand motion, and B is used as the start of the next letter instead of H_2. The magnetometer readings during such hand motions (indicated in red as shown in Figure 8(d)), on the other hand, are removed from segmentation and letter recognition.

4.2 Heterogeneity of Writing Speeds

As shown in Section 2.3, durations of writing different letters are usually heterogeneous. On the mean amplitudes of CWT frequency components over time, a letter written with a longer duration corresponds to a deeper valley and wider peak, as shown in Figure 9. Therefore, when people are writing fast and cursively, the valleys will be too shallow to be recognized. We will further investigate the impact of such heterogeneous writing speeds on MagHacker's letter segmentation in Section 7.5.

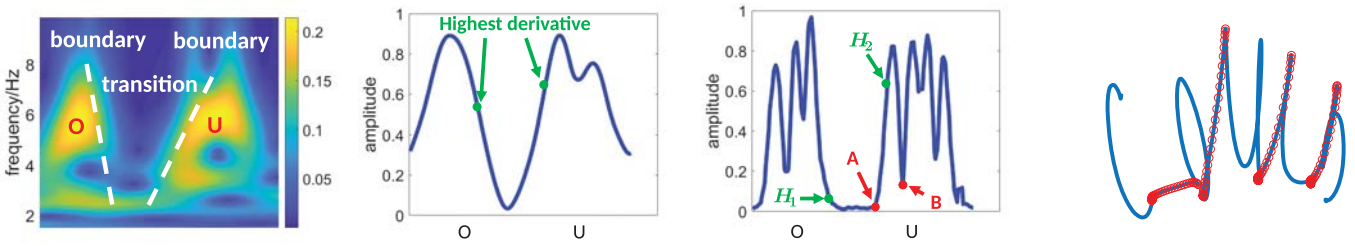

(a) CWT spectrogram of letters 'O' and 'U' (b) Mean amplitudes of frequency components (c) Humans' hand motion between letters (d) Removal of hand motions

Figure 8: Deciding letter boundaries

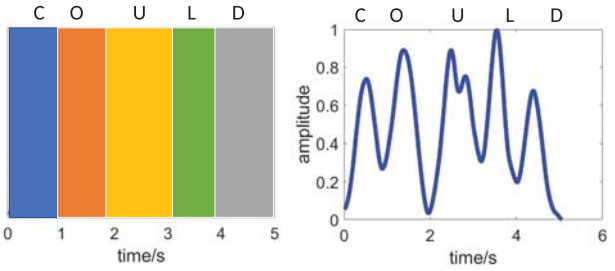

(a) Duration of writing letters (b) Mean amplitudes of frequency components

Figure 9: Writing speeds in the word "COULD"

5 COORDINATE TRANSFORMATION

In practice, humans usually write letters in each word along a straight line from left to right. Without loss of generality, we consider this straight line of writing direction as the x-axis in the magnet's coordinate system. Our objective of coordinate transformation, then, is to restore the distorted coordinates of letter centroids in the same word, as shown in Figure 7(c) after projection, back to this straight line. Our basic approach, as shown in Figure 10(a), is to fit the distorted letter centroids in a word into an arc [39], and then transform each letter along the tangential direction of its centroid (indicated as red arrows) on the arc.

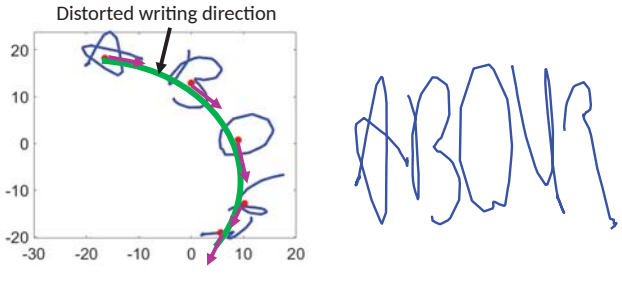

(a) Transformation approach (b) Trajectories after transformation

Figure 10: Coordinate transformations

In this way, for each 2D magnetometer reading (x, y) of a letter, its transformation will be

$$\begin{cases} x_{new} = xcos\theta - ysin\theta \\ y_{new} = xsin\theta + ycos\theta, \end{cases} \tag{3}$$

where (x_{new}, y_{new}) are the coordinates after transformation, and θ is the angle between the letter centroid's tangential direction and the x-axis in the magnet's coordinate system. For the magnetometer readings shown in Figure 10(a), the writing trajectories after transformation are shown in Figure 10(b) and visibly recognizable.

	Lowercase	Uppercase	Percentage
Same	'o','r','v','w'	'M','O','V','W'	15.4%
Lower	'i','j'	'A','N'	7.7%
Higher	others	others	79.8%

Table 2: The height of the letter's start point compared to the letter's end point

5.1 Flipping Writing Trajectories

After rotation, the letters might be upside down. MagHacker detects these cases and vertically flips the letters' writing trajectories as necessary, according to the height difference between the start point and end point of the letter's writing trajectory in the 2D plane. As shown in Table 2, most letters are written from top to bottom, and their start points will be at a higher position. For the other 20% of letters, most letters such as "O" and "i" can still be recognized even if they are upside down.

5.2 Impact of Magnet's Orientation

In handwriting, different humans tend to hold their pens with different orientations. The magnetometer readings, in this case, could be different because of the changes on the magnet's coordinate system. For example, Figure 11(a) shows the 3D magnetometer readings of writing the same word 'ABOUT' with a stylus pen orientation that has 120 degrees of difference from that in Figure 10. Nevertheless, the coordinate transformation approach in MagHacker can effectively remove such impact of different orientations, because different orientations of the pen only impact the rotation of magnetometer reading from the magnetometer's coordinate system to the magnet's coordinate system. Since different letters in a word are always being transformed back to a straight line along the x-axis

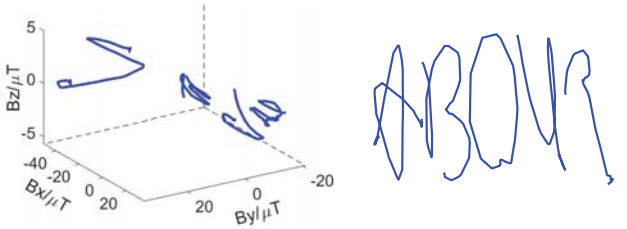

(a) Reading with a different pen orientation (b) Transformed writing trajectories

Figure 11: Impact of stylus pen's orientation

in the magnetometer's coordinate system, MagHacker can produce similar writing trajectories for letters being written with different pen orientations. For example, the transformed writing trajectories from the magnetometer readings in Figure 11(a), as shown in Figure 11(b), are very similar to the trajectories in Figure 10(b).

5.3 Varying Distances from Magnetometer

When writing multiple letters, the distance between the magnet and the magnetometer constantly varies. The magnetic field's strength at the magnetometer, hence, keeps changing and results in different contour sizes of letters' writing trajectories. For example as shown in Figure 12, if we write towards the magnetometer, the contour size of letters grows bigger. Otherwise, such size becomes smaller. To prevent such heterogeneity of contour sizes from affecting the accuracy of letter recognition, we scale them to the same size before letter recognition. For each point (x_i, y_i) on the letter's writing trajectory, the scaling equation is

$$\begin{cases} x_i' = (x_i - x_{\min})/(x_{\max} - x_{\min}) \\ y_i' = (y_i - y_{\min})/(y_{\max} - y_{\min}), \end{cases}$$

where $x_{\min}$ and $x_{\max}$ are the minimum and maximum X-axis coordinates on the trajectory, and $y_{\min}$ and $y_{\max}$ are similarly defined.

6 LETTER RECOGNITION

MagHacker applies the writing trajectories of individual letters to a CNN classifier to recognize these letters. One intuitive approach to training the CNN classifier is to use the handwriting samples produced by the attackers[3], but may easily overfit the trained CNN model due to the limited amount of training data. Instead, MagHacker augments the training data by incorporating 1) the printed letters with different fonts and 2) online datasets of handwritten letter images [4].

6.1 CNN Classifier Design

Being different from existing work that uses the time series data to train the classifier [37, 42], MagHacker converts the writing trajectory of each letter into a thumbnail image and then uses these images to train the CNN classifier. In this way, MagHacker can efficiently resist again the impact of humans' heterogeneous writing habits in practice. For example, some people use to write the letter 'O' clockwise and some others prefer to write 'O' counter-clockwise. Although the two cases produce the same thumbnail

[3]We assume that multiple attackers could collude to contribute such training data.

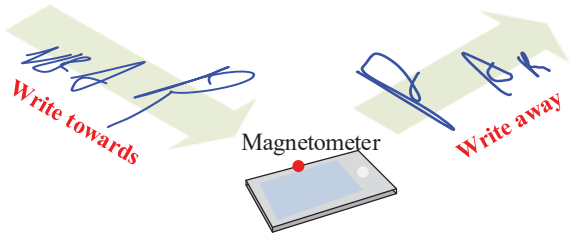

Figure 12: Different contour sizes of letters

images of writing trajectory, they correspond to totally different time series of magnetometer readings. Similarly, different people tend to write some letters with different sequences of strokes, such as 'E' and 'F'.

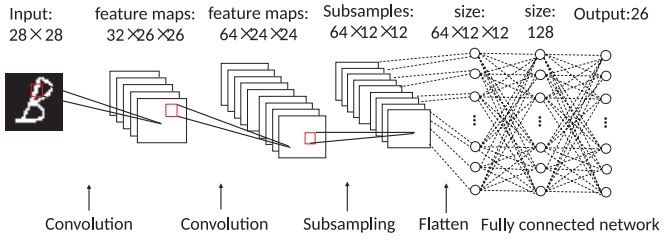

Figure 13: CNN classifier in MagHacker

Our CNN classifier design is shown in Figure 13. Following the specification in the MNIST dataset [26], we resize each image of writing trajectory to a 28×28 thumbnail for both training and inference phases. To train the CNN classifier, we calculate the convolutions twice to get 64 feature maps, and compress that data with a subsampling process. The subsampled feature maps are then flattened to be a fully connected neural network and trained with the input images until it converges. A similar process undergoes for inference, and the output is the one-hot code among the letter alphabet.

6.2 Augmenting the Training Data

We augment the training dataset by pre-training the CNN with printed letters of different fonts. We use all the default 138 font files provided by Windows 10 and 62 online files of handwriting fonts to generate 200 training samples for each letter. To better mimic the possible distortions of handwriting trajectories being translated from magnetometer readings, we randomly rotate, shift, shear and zoom each sample, before applying them for training. Note that such training is an one-time effort and could be done by the attacker offline, and there is hence no extra cost for the attacker at run-time. This pre-trained CNN model is then used to take inputs of magnetometer readings for further training, as described in Section 6.1.

7 PERFORMANCE EVALUATION

We evaluate the performance of MagHacker in recognizing the handwriting of five graduate student volunteers with different writing patterns and habits. All experiments are being conducted in a 10m×10m office. Every student volunteer, as shown in Figure 14, uses a stylus pen to write English letters over the LCD screen of a tablet, and a smartphone (iPhone Xs or Google Pixel 2 XL) is then placed nearby for eavesdropping. The size of letters being written is 2cm×2cm[4], and the distance between the eavesdropping device and the victim device varies between 5cm and 20cm[5]. Since the eavesdropping device does not produce any audible sound, screen display or disturbance to the victim device, in practice it could be casually placed by the attacker without producing any vigilance of the victim.

Figure 14: Experiment setup

In each experiment, the volunteer starts holding the stylus pen with a random orientation and then consecutively writes three English letters in the form of $L_1L_2L_1$, where both L_1 and L_2 iterate through the alphabet and each three-letter series is repetitively written for five times. In this way, every student volunteer covers all the possible transitions between letters, and contributes $26 \times 26 \times 3 \times 5 = 10,140$ handwriting trajectories of different letters. Writing trajectories of uppercase and lowercase letters are separately collected.

All the magnetometer readings are transmitted to a desktop PC for data processing and recognition. For each student volunteer, we use the other four volunteers' handwriting trajectories as the training data to train the CNN classifier using the Keras package of Tensorflow [35] and run it on Ubuntu 18.04. Then, we use the classifier to recognize this student volunteer's handwriting. The experiment result is averaged over the five volunteers.

7.1 Performance of Letter Segmentation

According to Section 4, the performance of letter segmentation in MagHacker is determined by whether the transition periods between letters can be completely removed from the writing trajectories. For example, compared with to ideal segmentation over the three-letter series 'BBB' in Figure 15(a), the practical segmentation results as shown in Figure 15(b) could be imperfect in some cases, and the remaining magnetometer readings during transitions between letters may lead to wrong recognitions of letters.

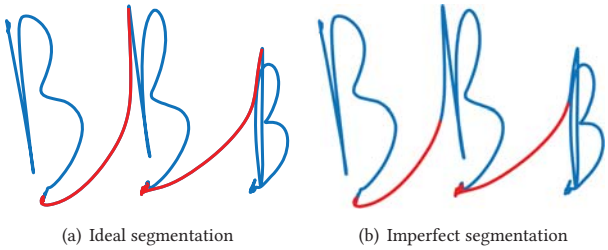

(a) Ideal segmentation (b) Imperfect segmentation

Figure 15: Illustration of practical letter segmentation. Lines in red indicate the transition between letters that should be removed from writing trajectories.

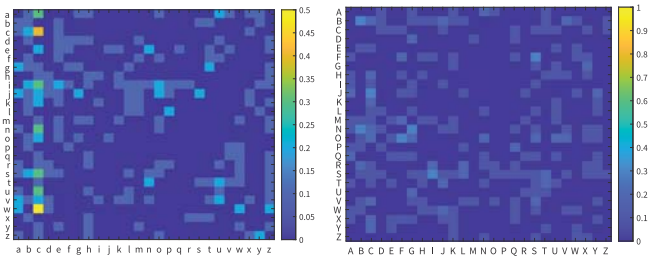

(a) Lowercase letters (b) Uppercase letters

Figure 16: Errors of Letter Segmentation

To evaluate the performance of letter segmentation, we compare the results of letter segmentation in MagHacker with those of ideal segmentation, and compute the segmentation errors as the percentage of the difference between the two. To obtain magnetometer readings with ideal segmentation, we manually monitor and record the starting and ending timestamps for every letter that the student volunteers are writing.

The average errors when segmenting different three-letter series are shown in Figure 16. For example, the number at 'c' in x-axis and 'w' in y-axis indicates the error when segmenting the three-letter series 'cwc'. From Figure 16 we can see that, our proposed approach to letter segmentation can correctly identify letter boundaries and remove transition periods in most cases, and the average segmentation errors for lowercase and uppercase letters are 3.51% and 3.77%, respectively.

On the other hand, segmentation errors are more likely to occur in transition between certain letters, due to the specific shapes of these letters. For example, when transiting from 'c' to 'e', people tend to move the pen directly from the end of 'c' to the start of 'e' and hence make it difficult to find letter boundaries. Similarly, when transiting from 'S' to 'I', the end of hand motion in transition and the first vertical stroke in writing 'I' are usually connected and hard to be distinguished from each other. Nevertheless, these special cases are very few and have negligible impact on the overall performance of MagHacker's letter segmentation.

[4]This size matches most cases of handwriting inputs using stylus pens. For example, iPhone Xs uses the bottom half of screen area for handwriting input, with a size of 6cm×7cm. Most mobile tablets (e.g., Microsoft Surface Pro), on the other hand, require the minimum size of handwritten letters to be large enough for accurate handwriting recognition.

[5]Such distance is measured between the eavesdropping device's magnetometer and the center of letter being written on the victim device.

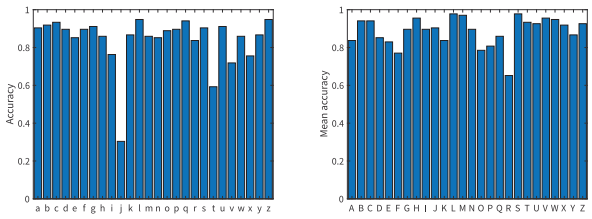

(a) Recognizing lowercase letters (b) Recognizing uppercase letters

Figure 17: Recognition accuracy with ideal segmentation

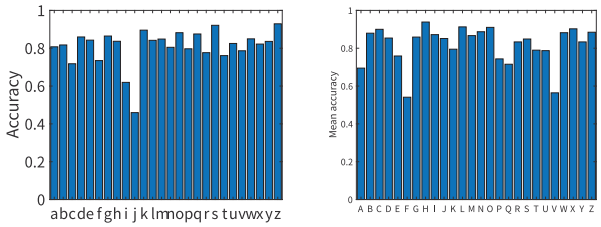

(a) Recognizing lowercase letters (b) Recognizing uppercase letters

Figure 18: MagHacker's eavesdropping accuracy

7.2 Performance of Coordinate Transformation

The eavesdropping accuracy of MagHacker is mainly determined by the performance of our proposed coordinate transformation approach in Section 5. We evaluate its performance by examining the accuracy of recognizing the transformed magnetometer readings with ideal segmentation. The recognition accuracy for lowercase and uppercase letters, as shown in Figure 17, is 84.16% and 88.69%, respectively. Such higher accuracy demonstrates that the distorted magnetometer readings, in most cases, can be correctly restored to recognizable writing trajectories after coordinate transformation.

On the other hand, the accuracy of letter recognition is affected by the ineffective coordinate transformation in certain cases. For example, when we did not completely remove the hand motion trajectory between the vertical stroke and the top dot in letter 'j', the shape of 'j' looks very differently from the printed letter and could be easily misclassified. Similarly, the hand motion trajectory from the end of the first stroke in 'K' to its second stroke makes 'K' look like "R". Thus, we have low accuracy on both cases.

7.3 Eavesdropping Accuracy

MagHacker's eavesdropping accuracy may be impaired by the imperfect segmentation as shown in Section 7.1 and could hence be lower from the results being shown in Section 7.2 with ideal segmentation. However, the results in Figure 18 show that such performance degradation is very limited: MagHacker's eavesdropping accuracy can achieve 77.74% for lowercase letters and 80.10% for uppercases letters, respectively, and only experiences 6% and 8% degradation due to letter segmentation errors. In particular, note that most segmentation errors concentrate on recognizing few letters such as 'j', 'F' and 'V'.

Furthermore, Table 3 shows that such eavesdropping accuracy could be consistently retained over each individual student volunteer, with a variance smaller than 5%. Such variance is mainly caused by their different writing habits, such as the sequence of strokes when writing letters, different orientations holding the stylus pen, and different speeds when writing the same letter. Nevertheless, in the worst case, the eavesdropping accuracy is always higher than 70% for lowercase letters and 75% for uppercase letters.

Volunteer index	1	2	3	4	5
Lowercase letters	85.60%	73.13%	72.58%	80.08%	77.30%
Uppercase letters	81.76%	75.01%	82.35%	81.90%	79.49%

Table 3: Recognition accuracy on individual student volunteers

In practice, humans' handwriting is mostly limited among legitimate English words, over which only a small portion of letter combinations could possibly appear. Individual errors of letter recognitions, in such cases, could be corrected by spell checking that is widely available in editing software (e.g., Microsoft Word). To investigate such recognition accuracy over legitimate words, we ask student volunteers to write the top-100 frequently used English words [5] with different lengths that range between 3 and 7, and only consider a word as correctly recognized if all its letters are correctly recognized after spell checking. Experiment results in Table 4 show that the recognition accuracy over these words ranges between 75% and 80%. In particular, even when the word length increases to 7 and greatly increases the difficulty in spell checking and correction, MagHacker can still have 75% accuracy in recognizing these long words. These results are important and demonstrate MagHacker's potential in eavesdropping and recognizing humans' handwriting inputs in practical conditions.

Word length	3	4	5	6	7
Accuracy	81.7%	81.7%	80.8%	78.7%	75.1%

Table 4: Recognition accuracy over legitimate words

7.4 Impact of the Eavesdropping Distance

The strength of magnetic field being produced by the stylus pen's magnet decays over distance. When the magnet is further away from the magnetometer, the magnetic field becomes weaker and hence brings more noise in the magnetometer readings. In all the experiments above, we fix the distance between the attacking device and victim device to be 10cm.

On the other hand, when the eavesdropping distance enlarges, its impact on the eavesdropping accuracy is shown in Figure 19. MagHacker can effectively retain the eavesdropping accuracy at 70% when the eavesdropping distance enlarges to 15cm. When the distance further enlarges to 20cm, the decay of magnetic field strength produces extra errors in letter segmentation, and the eavesdropping accuracy drops to 60%. In these cases, MagHacker can still further improve its recognition accuracy when being applied among legitimate English words.

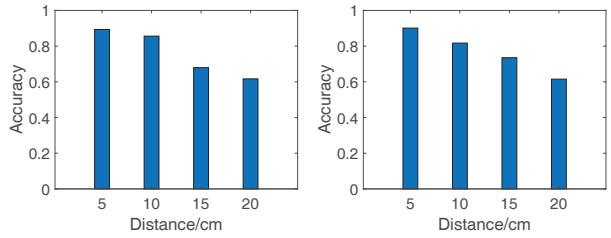

(a) Recognizing lowercase letters (b) Recognizing uppercase letters

Figure 19: The impact of eavesdropping distance

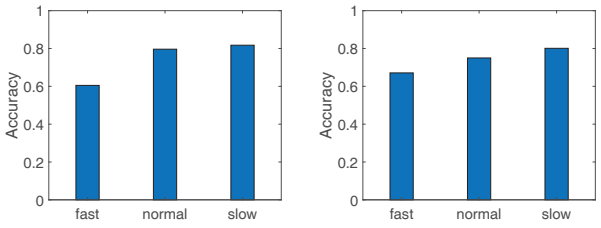

(a) Recognizing lowercase letters (b) Recognizing uppercase letters

Figure 20: The impact of different writing speeds

7.5 Impact of Writing Speed

The duration of writing each letter will significantly impact the detection of letter boundaries, and the segmentation algorithm may misdetect the number of letters in a word when people are writing fast. To evaluate the impact of such different writing speeds on the eavesdropping accuracy, we let the student volunteers to write the top-20 frequently used words [5], whose length ranges from 3 to 7, in fast, normal and slow speeds: the fast speed is to write three letters per two second, the slow speed is to write four letters per three seconds, and the normal speed is to write one letter per second.

Experiment results in Figure 20 show that the recognition accuracy of MagHacker can be reliably retained to be >75% when the writing speed does not exceed 1 letter/sec (normal speed), and only experience significant performance degradation if the handwriting speed is very fast. In practice, this fast speed already exceeds the normal range of humans' handwriting (especially using stylus pens over LCD screens), and only appears over casual cursive writing or fast short typing.

7.6 Eavesdropping with Different Positions and Orientations

As stated in Section 2.3, the magnetometer readings are jointly determined by the position and orientation of both the magnetometer and the magnet of stylus pen. To evaluate the impact of such different positions and orientations on the eavesdropping accuracy, we put the eavesdropping device at 4 different locations and also change its orientation at each location, as shown in Figure 21(a). The eavesdropping distance (d) is 10cm and normal writing speed is used in all cases.

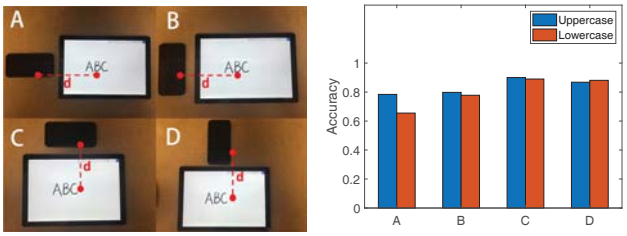

(a) Different positions and orientations (b) Eavesdropping accuracy

Figure 21: Eavesdropping with different positions and orientations

The average eavesdropping accuracy over all the three-letter series, as shown in Figure 21(b), exhibits little variance for both lowercase and uppercase letters. Noticeable degradation in eavesdropping accuracy (around 10%) is only noticed for position A, where the tablet's speaker is on the line of sight between the stylus pen and magnetometer and hence disturbs the magnetic field being produced. The eavesdropping accuracy for all other three cases, on the other hand, is always higher than 80%. These results demonstrate that our proposed approach to coordinate projection and transformation can effectively eliminate the impact of different magnet orientations.

7.7 Power Consumption of Eavesdropping

MagHacker requires continuous magnetometer readings to acquire the handwriting trajectories, and hence consumes extra power on the eavesdropping device. To evaluate such amount of extra power being consumed by eavesdropping, we keep the eavesdropping smartphone's magnetometer running at different sampling rates, and then compare the speed of power consumption with that of the smartphone's idle status. In both cases, the smartphone's screen is always on with 50% light intensity, and all the experiments are done without wireless connections and other background services.

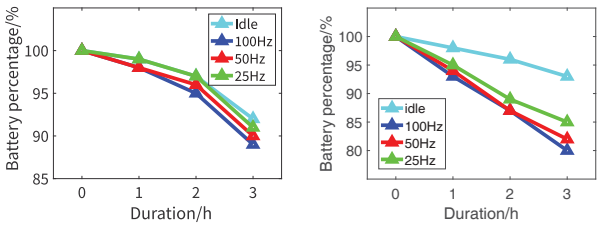

(a) iPhone Xs (b) Google Pixel 2 XL

Figure 22: Power consumption of MagHacker

The experiment results in Figure 22 show that MagHacker is very energy-efficient. With a fully charged smartphone, continuously using its magnetometer for 3 hours consumes less than 20% of the smartphone's battery power. In particular, such power consumption varies over different smartphone models, and Figure 22(b) shows that the Google Pixel 2 XL smartphone running Android 10 has faster power consumption. The basic reason for such difference

Figure 23: Other stylus pen models being evaluated

lies in their different mobile OSes being used: accessing magnetometer in Android OS has to go through its Sensor service that is implemented as a mobile middleware, and hence incurs higher system overhead. Such sensor access over iOS, comparatively, is more power efficient.

Besides, due to the low sampling rate of magnetometers, the data size of magnetometer readings is very small and hence incurs a negligible amount of energy consumption when being transmitted to the desktop PC for data processing and letter recognition.

7.8 Eavesdropping on Different Stylus Pen Models

We expect that MagHacker can be generally used to eavesdrop different stylus pen models. To examine such generality, we also tested the MagHacker's eavesdropping accuracy over Adonit Snap 2 Pen [1] and Maglus Pen [2], which are shown in Figure 23 and are both widely used products on market. As shown in Figure 24, when the eavesdropping distance is 10cm, the eavesdropping accuracy over these two stylus pen models is higher than 80% in most cases, and the average accuracy of these two pen models is 87.11% (Adonit Snap 2: 84.80% for lowercase and 89.41% for uppercase) and 89.39% (Maglus: 87.57% for lowercase and 91.20% for uppercase), respectively. These results demonstrate that MagHacker can be generally applied to different eavesdropping scenarios in practice.

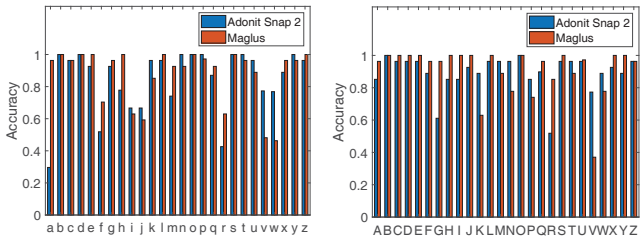

(a) Recognizing lowercase letters (b) Recognizing uppercase letters

Figure 24: Accuracy over other stylus pen models

7.9 Eavesdropping in Different Environments

Magnetic field sensing obviously will not be affected by most types of surrounding objects or obstacles, such as books, clothes, etc. In our previous experiments, each volunteer collects data at difference

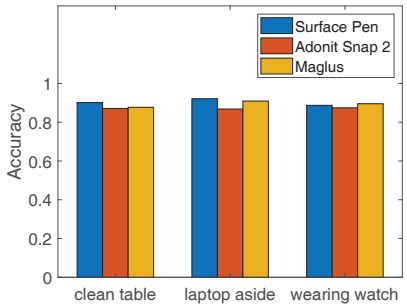

Figure 25: Eavesdropping in Different Environments

places, e.g. they collect the data on their lab table with different surrounding objects. These environments are usually surrounded by different types of electrical devices or other objects with metal components. We specifically select one volunteer and require him to wear a metal watch on his left wrist when collecting the first half of his data in the experiment at the beginning of Section 7.

The eavesdropping accuracy under these environments, as shown in Figure 25, experiences very little variation. These results demonstrate that MagHacker has very good adaptability to the environmental dynamics, even with nearby metal objects being present. In practice, the attacker can place the eavesdropping device under different covers (e.g., boxes, clothes, paper, etc), and achieve good concealment without impairing the eavesdropping accuracy.

8 RELATED WORK

Eavesdropping on handwriting. Early methods suggest to eavesdrop on humans' handwriting with hidden cameras [13, 23, 24, 34], but need to be in the victim's physical proximity with line of sight. The real-time processing of the captured images, on the other hand, also incurs high computing overhead. Research efforts have been made to eavesdrop on humans' handwriting by monitoring the audible sound or body motions produced by handwriting. For example, the rubbing sound between the pen nib and the writing surface could be captured and analyzed for eavesdropping [16, 42], which however, only works in highly quiet indoor environments with minimum ambient noise. IMU data readings from a smartwatch can also be applied to deep learning models for eavesdropping [22, 37], but requires the smartwatch to be compromised and worn on the victim's writing hand. In contrast, MagHacker can effectively launch the eavesdropping attack from distance, and can ensure the accuracy of eavesdropping in diverse environmental settings.

Advanced wireless sensing techniques validated that humans' hand gestures could lead to fluctuations of the wireless signals on the air, and can hence be derived from signal characteristics, such as Doppler shift [8, 20, 43] and Channel State Information (CSI) [28, 41, 44]. These characteristics could be potentially used for eavesdropping, but can only be captured with specialized RF hardware that is expensive. They are also susceptible to the environmental dynamics such as channel fading and multipath. In contrast, MagHacker can be easily executed from commodity mobile devices, and the magnetic field produced by the magnet has been proved to be less affected by the surrounding objects.

Magnetic sensing. Tracking the 2D or 3D movement of a magnet usually requires multiple magnetometers [11, 12, 19], so as to precisely estimate all the six degrees of freedom (DOF) in the magnet's movement. Existing work on tracking a magnet's movement with a single magnetometer, on the other hand, needs help from other sensors on the victim, such as IMU sensors [40] or a secondary magnet [6] to provide information about the primary magnet's orientation. These existing techniques, hence, cannot be applied to commodity mobile devices without adding extra hardware or physically contacting with the victim.

Letter segmentation in handwriting recognition. An intuitive approach to letter segmentation is to detect vertical movements of the pen, which only happen in the start and end points of writing letters [21, 30]. Such detection, however, requires IMU readings that are usually unavailable from stylus pens. Others use a pre-defined time window to detect transition between letters [14, 42], but are very sensitive to the heterogeneous writing habits among humans. MagHacker, instead, exploits the frequency-domain features of the magnet's movement, and can hence be generally applied to humans with different writing habits.

9 DISCUSSIONS AND FUTURE WORK

In this section, we discuss the potential pitfalls and limitations of MagHacker, as well as how such eavesdropping attack could be effectively defended.

The Eavesdropping Distance: The attack design of MagHacker focuses on scenarios in crowded public places, such as public libraries, markets or transportation systems, where people are close to each other even without any social interaction. In these cases, MagHacker enables eavesdropping without attracting victim's attention, because of its adaptability shown in Section 7.9 that brings great concealment. For example, the eavesdropping device could be carried by the attacker under cover, or be placed by the attacker in the public place that he/she shares with the victim (e.g., table).

On the other hand, results in Section 7.4 show that MagHacker's eavesdropping accuracy will drop to <70% when the eavesdropping distance extends to 15-20cm, and hence impairs MagHacker's applicability in certain scenarios. The major reason of such drop is the decay of magnetic field strength over distance. One possibility of combating such decay is to use multiple magnetometers with known positions for extra signal calibration, and this can be realized by using multiple mobile devices owned by the attacker. We will further explore this possibility in our future work.

Recognizing Non-Letter Characters: Although our evaluations and discussions in this paper focus on English letters, MagHacker can be easily extended to recognizing handwritten non-letter characters (e.g., numbers and special characters), by incorporating them into the training data of CNN classifier. Similarly, MagHacker can also recognize different English fonts and letters written with different variants (e.g., 'z' and '7' with an extra middle stroke).

Recognition of Cursive Writing: The design of MagHacker assumes that humans individually write each letter without cursive writing. This design can be further extended to recognizing cursive writing by being applied to existing approaches in handwriting recognition. For example, a common method to recognize cursive writing is to train a Hidden Markov Model that segments each letter into strokes and try different combinations of strokes for recognition [7, 14, 33], and can be well integrated with the letter segmentation approach in MagHacker.

Mixing Uppercases and Lowercases: In this paper, we train two CNN classifiers in MagHacker to recognize uppercase and lowercase letters, separately. However in practical handwriting, both uppercase and lowercase letters may appear in the same word, e.g., the first word in a sentence. Some traditional methods judge whether a letter is uppercase or lowercase according to its size [17], but may not be applicable to magnetometer readings that naturally result in different letter sizes as described in Section 5.3. Instead, MagHacker can efficiently address such mixed cases via its coordinate transformation and scaling, which remove such impact of heterogeneous letter sizes.

Eavesdropping from Victim's Device: If the victim's mobile device being used for handwriting with stylus pen (e.g., tablets) has a built-in magnetometer, MagHacker can also be implemented as a malware and injected to the victim's device for eavesdropping via the local magnetometer. Such eavesdropping can achieve higher accuracy due to the close distance between the pen magnet and the magnetometer. Furthermore, accessing the magnetometer, being similar to IMU sensor access, is unrestricted in most of today's mobile OSes (e.g., iOS and Android). Such malware injection, hence, is much easier and does not need to compromise the victim device's OS. We will investigate this attack as our future work.

Possible Defenses: This eavesdropping attack is purely passive and hence hard to be detected by the victim. The most effective defense against such eavesdropping attack, instead, is to apply magnetic shielding on stylus pens, but may be practically expensive or greatly increase the pen's weight. Another alternative is to intentionally obfuscate the magnetic field produced by the stylus pen's magnet, so that the trajectories of magnetometer readings are not recognizable. Being different from existing RF or location obfuscation techniques [10, 38], alternating the magnetic field is much harder and requires extra hardware. For example, another magnet could be attached to the victim's device for obfuscation but may affect certain device functionality (e.g., gyroscope sensing). An electromagnet is a better option, but may be power hungry.

10 CONCLUSION

In this paper, we present MagHacker, a new sensing system that allows an attacker to precisely eavesdrop humans' handwriting with their stylus pens from distance. MagHacker has been proved to achieve high eavesdropping accuracy over an eavesdropping distance up to 20cm, and can well adapt to different types of environmental dynamics and humans' writing patterns. MagHacker's applicability in practice could be further examined over a larger variety of English words and with more participants in different backgrounds, and these examinations could be our future work.

ACKNOWLEDGMENTS

We thank our shepherd Wenyao Xu and anonymous reviewers for their comments and feedback. This work was supported in part by the National Science Foundation (NSF) under grant number CNS-1617198, CNS-1812399 and CNS-1812407.

REFERENCES

[1] Adonit snap 2 pen. *https://www.adonit.net/snap2/*.
[2] Maglus pen. *http://maglusstylus.com/*.
[3] Microsoft surface pen. *https://www.microsoft.com/en-us/p/surface-pen/92fp8q09qhxc*, 2019.
[4] Urban fonts. *https://www.urbanfonts.com/fonts/handwritten-fonts.htm*, 2019.
[5] Word frequency. *https://www.wordfrequency.info/*, 2019.
[6] T. Abe, B. Shizuki, and J. Tanaka. Input techniques to the surface around a smartphone using a magnet attached on a stylus. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pages 2395–2402. ACM, 2016.
[7] C. Agarwal, D. P. Dogra, R. Saini, and P. P. Roy. Segmentation and recognition of text written in 3d using leap motion interface. In *2015 3rd IAPR Asian Conference on Pattern Recognition (ACPR)*, pages 539–543. IEEE, 2015.
[8] M. T. I. Aumi, S. Gupta, M. Goel, E. Larson, and S. Patel. Doplink: using the doppler effect for multi-device interaction. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 583–586, 2013.
[9] D. Bledsoe. Handwriting speed in an adult population. *Advance for Occupational Therapy Practitioners*, 27(22):10, 2011.
[10] A. Brush, J. Krumm, and J. Scott. Exploring end user preferences for location obfuscation, location-based services, and the value of location. In *Proceedings of the 12th ACM international conference on Ubiquitous computing*, pages 95–104. ACM, 2010.
[11] K.-Y. Chen, K. Lyons, S. White, and S. Patel. utrack: 3d input using two magnetic sensors. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*, pages 237–244. ACM, 2013.
[12] K.-Y. Chen, S. N. Patel, and S. Keller. Finexus: Tracking precise motions of multiple fingertips using magnetic sensing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 1504–1514. ACM, 2016.
[13] L. Chen, S. Wang, W. Fan, J. Sun, and S. Naoi. Beyond human recognition: A cnn-based framework for handwritten character recognition. In *2015 3rd IAPR Asian Conference on Pattern Recognition (ACPR)*, pages 695–699. IEEE, 2015.
[14] M. Chen, G. AlRegib, and B.-H. Juang. Air-writing recognitionâĂŤpart ii: Detection and recognition of writing activity in continuous stream of motion data. *IEEE Transactions on Human-Machine Systems*, 46(3):436–444, 2015.
[15] M. Chen, P. Yang, S. Cao, M. Zhang, and P. Li. Writepad: Consecutive number writing on your hand with smart acoustic sensing. *IEEE Access*, 6:77240–77249, 2018.
[16] H. Du, P. Li, H. Zhou, W. Gong, G. Luo, and P. Yang. Wordrecorder: Accurate acoustic-based handwriting recognition using deep learning. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*, pages 1448–1456. IEEE, 2018.
[17] D. Fiset, C. Blais, C. Ethier-Majcher, M. Arguin, D. Bub, and F. Gosselin. Features for identification of uppercase and lowercase letters. *Psychological science*, 19(11):1161–1168, 2008.
[18] A. Grossmann, R. Kronland-Martinet, and J. Morlet. Reading and understanding continuous wavelet transforms. In *Wavelets*, pages 2–20. Springer, 1990.
[19] X. Han, H. Seki, Y. Kamiya, and M. Hikizu. Wearable handwriting input device using magnetic field. In *SICE Annual Conference 2007*, pages 365–368. IEEE, 2007.
[20] W. Huang, Y. Xiong, X.-Y. Li, H. Lin, X. Mao, P. Yang, and Y. Liu. Shake and walk: Acoustic direction finding and fine-grained indoor localization using smartphones. In *IEEE INFOCOM 2014-IEEE Conference on Computer Communications*, pages 370–378. IEEE, 2014.
[21] W. Jeen-Shing, H. Yu-Liang, and C. Cheng-Ling. Online handwriting recognition using an accelerometer-based pen device. In *2nd International Conference on Advances in Computer Science and Engineering (CSE 2013)*. Atlantis Press, 2013.
[22] H. Jiang. Motion eavesdropper: Smartwatch-based handwriting recognition using deep learning. In *2019 International Conference on Multimodal Interaction*, pages 145–153. ACM, 2019.
[23] L. Jin, D. Yang, L.-X. Zhen, and J.-C. Huang. A novel vision-based finger-writing character recognition system. *Journal of Circuits, Systems, and Computers*, 16(03):421–436, 2007.
[24] D. Keysers, T. Deselaers, H. A. Rowley, L.-L. Wang, and V. Carbune. Multi-language online handwriting recognition. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1180–1194, 2016.
[25] M. B. Kraichman. *Handbook of electromagnetic propagation in conducting media*.
[26] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
[27] C.-C. Lee, C.-Y. Shih, and B.-S. Jeng. Fingertip-writing alphanumeric character recognition for vision-based human computer interaction. In *2010 International Conference on Broadband, Wireless Computing, Communication and Applications*, pages 533–537. IEEE, 2010.
[28] M. Li, Y. Meng, J. Liu, H. Zhu, X. Liang, Y. Liu, and N. Ruan. When csi meets public wifi: Inferring your mobile phone password via wifi signals. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 1068–1079. ACM, 2016.
[29] Y. Li and W. Gao. Interconnecting heterogeneous devices in the personal mobile cloud. In *Proceedings of IEEE INFOCOM*, 2017.
[30] M. Nakai, T. Sudo, H. Shimodaira, and S. Sagayama. Pen pressure features for writer-independent on-line handwriting recognition based on substroke hmm. In *Object recognition supported by user interaction for service robots*, volume 3, pages 220–223. IEEE, 2002.
[31] S. Patil, D. Kim, S. Park, and Y. Chai. Handwriting recognition in free space using wimu-based hand motion analysis. *Journal of Sensors*, 2016, 2016.
[32] G. Reyes, J. Wu, N. Juneja, M. Goldshtein, W. K. Edwards, G. D. Abowd, and T. Starner. Synchrowatch: One-handed synchronous smartwatch gestures using correlation and magnetic sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(4):158, 2018.
[33] S. Sagar and S. Dixit. Hmm segmentation approach for offline cursive handwritten words. *Int. J. Eng. Sci. Comput.(IJESC)*, 8, 2018.
[34] C.-L. Shih, W.-Y. Lee, and Y.-T. Ku. A vision-based fingertip-writing character recognition system. *Journal of Computer and Communications*, 4(04):160, 2016.
[35] Tensorflow. https://www.tensorflow.org/guide/keras, 2019.
[36] L. Tong, Y. Li, and W. Gao. A hierarchical edge cloud architecture for mobile computing. In *Proceedings of IEEE INFOCOM*, 2016.
[37] Q. Xia, F. Hong, Y. Feng, and Z. Guo. Motionhacker: Motion sensor based eavesdropping on handwriting via smartwatch. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 468–473. IEEE, 2018.
[38] T. Xiong, W. Lou, J. Zhang, and H. Tan. MIO: Enhancing wireless communications security through physical layer multiple inter-symbol obfuscation. *IEEE transactions on information forensics and security*, 10(8):1678–1691, 2015.
[39] X. Yang. Efficient circular arc interpolation based on active tolerance control. *Computer-Aided Design*, 34(13):1037–1046, 2002.
[40] S. H. Yoon, K. Huo, and K. Ramani. Tmotion: Embedded 3d mobile input using magnetic sensing technique. In *Proceedings of the TEI'16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*, pages 21–29. ACM, 2016.
[41] J. Yu, L. Lu, Y. Chen, Y. Zhu, and L. Kong. An indirect eavesdropping attack of keystrokes on touch screen through acoustic sensing. *IEEE Transactions on Mobile Computing*, 2019.
[42] T. Yu, H. Jin, and K. Nahrstedt. Writinghacker: audio based eavesdropping of handwriting via mobile devices. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 463–473. ACM, 2016.
[43] S. Yun, Y.-C. Chen, and L. Qiu. Turning a mobile device into a mouse in the air. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services*, pages 15–29, 2015.
[44] M. Zhou, Q. Wang, J. Yang, Q. Li, F. Xiao, Z. Wang, and X. Chen. Patternlistener: Cracking android pattern lock using acoustic signals. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 1775–1787. ACM, 2018.