An AI Method to Reveal ‘Shielded’ PIN Entries at ATMs

Originally published at https://www.unite.ai/an-ai-method-to-reveal-shielded-pin-entries-at-atms/

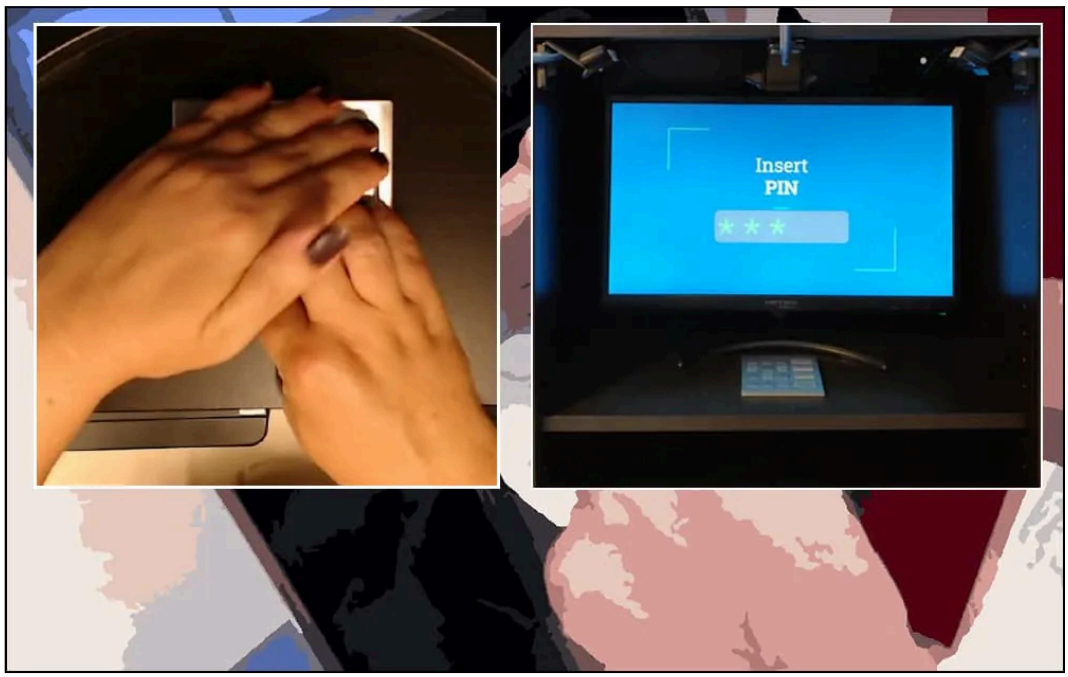

Researchers in Italy and the Netherlands have developed a machine learning method capable of inferring the PIN number that a bank customer puts into an ATM, based on captured video – even in cases where the customer shields their hand to protect against shoulder-surfing.

The method involves training a Convolutional Neural Network (CNN) and a Long Short-Term Memory (LSTM) module on videos of ‘covered hand’ PIN entries at a ‘shadow’ ATM that has been fitted with the same keypad as the target ATM – equipment that can be purchased, as the researchers were able to do for the project, recreating a ‘mirror’ ATM in order to gather the data.

The fake ATM can be trained in private, as the researchers have done, obviating the risk of public installations of fake ATMs, a common *modus operandi* in this particular type of crime.

Left, two pin pad models used for the Italian research. Pictured right, the ‘shadow’ ATM that the researchers constructed in laboratory conditions. Source: https://ar

The system, which keys on hand movements and positioning during PIN input, can currently predict 41% of 4-digit pins and 30% of 5-digit PIN numbers inside three attempts (generally the maximum number of attempts a bank will allow before locking the customer’s account). Tests involved 58 volunteers using random PIN numbers.

The research, whose data is publicly available, finds that the proposed system offers a four-fold improvement on a human’s ability to guess a PIN by shoulder-surfing a victim.

The paper is titled *Hand Me Your PIN! Inferring ATM PINs of Users Typing with a Covered Hand*, and comes from five researchers at the University of Padua, and one from the Delft University of Technology.

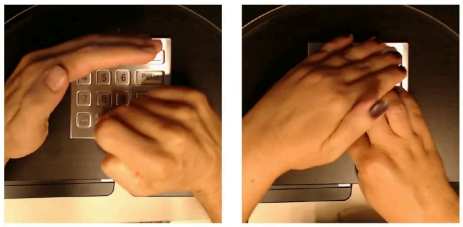

The researchers excluded captures where the subjects did not adequately cover the PIN pad (left).

The researchers contend that their system achieves superior results to prior work that keys on timing, sound and thermal signatures, without a video analysis component.

They also note that increased awareness of 'skimming' devices center around the card input slot, since this is a traditional method of attack, and that customers have no reason to believe that any similarly-hidden micro-cameras could 'see through' their covered hands, or that the generic clatter of keys and the identical feedback sound for each key-press could disclose any information.

The 'additional' equipment of the ATM would therefore appear in a place where no-one is currently expecting it, under the upper interior surface of the ATM recess, as a molded enclosure concealing the camera equipment – or even entirely outside the ATM surface, attached to a nearby building or post.

PIN Money

Despite the severe consequences of a breach, PIN numbers are among the shortest and most easily guessable passwords we use; it's been estimated that an attacker already has a 1-in-10 chance of guessing a PIN correctly. Social engineering isn't even always a necessary adjunct to more sophisticated AI-based attacks, since *1234* has been estimated to represent 11% of all pins, while *19* (as the first part of a birth year) represents the first two numbers in over 80% of PIN numbers.

Nonetheless, the authors of the new paper have not given themselves this advantage, but rather have set out to investigate whether the hand-wriggling of 'shielded' PIN entries has a decipherable pattern that can indicate which numbers are being pressed.

To establish a baseline, the researchers constructed a fake ATM for the purposes of data-gathering (see first image above). This represents the proposed hypothetical attack method, where a malefactor will passively analyze typical PIN input characteristics over a long period of time to prepare for a later 'swoop' on accounts.

Though this very 'studied' approach is common in sophisticated ATM fraud crime, with many instances of bogus ATMs exfiltrating customer data over a long period, in this case the attacker can set up the fake ATM in their own space, and train it without public input.

Since the ATM's screen is not likely to be concealed during PIN input, the timing of a key press can be established by synchronizing hand movements to the appearance of the 'masked' digits (usually asterisks) that appear in the ATM screen in response to user input, and also to generic feedback noises (such as beeps) that coincide with the strokes. This synchronization reveals the exact hand disposition in a 'shielded' scenario at the moment of input.

Targeting Specific Keypads

First, a model must be developed by observation and recording of shielded PIN inputs. Ideally, the keypad should be a specific industry standard model, though some variation of millimeters will not stop the method working. Key-press timings can be obtained by audio and visual cues (i.e. feedback beeps, key clatter, and asterisk feedback).

With these breakpoints, the attacker can automate the extraction of a training set and go on to train a model capable of identifying representative hand configurations for the pressing of a specific key. This will produce a ranked list of probabilities for the card's PIN, out of which the top three will be selected for the attack when authentic customer data is identified by the system in a real-world scenario.

Methodology

Data gathering was conducted over two sessions, using right-handed volunteers for the study. Each participant typed 100 randomly-generated 5-digit PIN numbers, to ensure even coverage of all ten possible keypad presses. In this way, the researchers gathered 5,800 individual PIN entries.

PIN pads used in the tests were the DAVO LIN Model D-8201F and the DAVO LIN Model D-8203 B models. They're commercial models used in ATMs, and are available, respectively, here and here (among numerous other vendors).

The gathered video segments were converted to grayscale and normalized and cropped, before being resized to 250×250 pixels for inclusion in the machine learning training sessions. Clips were segmented to obtain sub-sequences of frames pertaining to key presses. Audio cues (as mentioned above) were used as timestamp markers for press events.

Training

The datasets were split into training, validation and test sets, with the training taking place on a Xeon(R) Intel CPU running at E5-2670 2.60GHz, and equipped with 128GB of RAM. The data was implemented on Keras2.3.0-tf (TensorFlow 2.2.0) and Python 3.8.6 on three Tesla K20m GPUs with 5gb of VRAM each.

To account for variability in capture environments (lighting, slight differences in camera angles, etc.), synthetic examples and perturbations (such as rotation and view shift) were generated, and the authors report that this kind of data augmentation is a great aide in improving the effectiveness of the model.

Results

The model was tested against three scenarios: 'single PIN pad', where the attacker knows the model of pin pad, and trains specifically for it; 'PIN pad independent', where the model is trained on a pad that's similar (but not identical) to the target PIN pad; and a 'mixed scenario', where the attacker has a copy of both PIN pads.

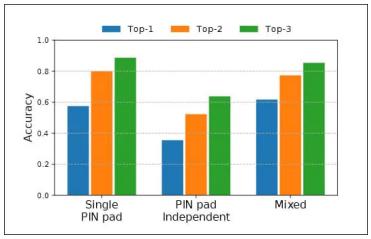

General results across the three scenarios, where Top-N signifies a guess of the digit within N *attempts.*

There's a noted difference in accuracy for inference of 5-digit versus 4-digit PINs:

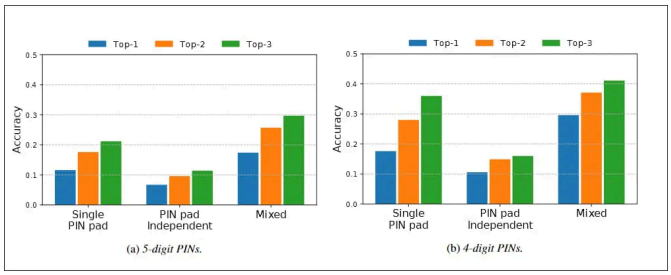

Countermeasures

In considering countermeasures to existing systems (i.e. without a radical rethink of the entire PIN/ATM security infrastructure), the researchers consider that there are no really workable defenses against this kind of attack.

Lengthening the minimum required numbers in a PIN will make the numbers more difficult to remember; randomizing the order of the numbers keypad with a touch-screen software keyboard, though increasingly happening in ATM deployments, also produces usability issues; and screen protectors would not only be expensive to deploy on existing ATMs, but would arguably make the paper's attack method even easier to implement, depending on how much coverage it might give. The researchers assert that their attack is workable even where 75% of the PIN pad is covered (and covering more would make it difficult for the customer to type).

In devising a human-based equivalent to the automated PIN extraction, real people were, by contrast, only able to achieve a fraction of the AI system's accuracy in guessing PINs, based on the same information.

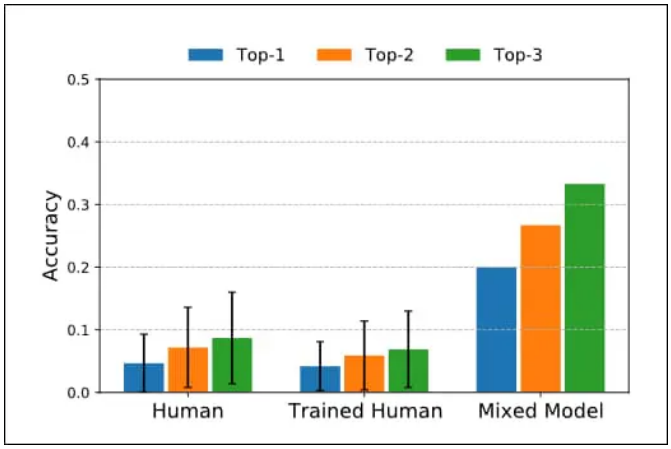

In future development of the work, the researchers intend to examine results from non-right handed people, and to investigate hand-covering strategies that might mitigate the attack. They also intend to repeat the experiments with a greater diversity of ages and races, since they observe that older people make more significant and telling hand movements when entering a PIN, and that the attack 'will have difficulties working for people from other races' (than Caucasian).

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.
Portfolio site: martinanderson.ai
Contact: martin@martinanderson.ai

Discover More