

Better Passwords through Science (and Neural Networks)

WILLIAM MELICHER, BLASE UR, SEAN M. SEGRETI, LUJO BAUER,
NICOLAS CHRISTIN, AND LORRIE FAITH CRANOR



William Melicher is a PhD student in the College of Electrical and Computer Engineering at Carnegie Mellon University. He works on passwords, Web security, and online privacy. He received his undergraduate degree in computer engineering from the University of Virginia. billy@cmu.edu



Blase Ur is Neubauer Family Assistant Professor of Computer Science at the University of Chicago. His research focuses broadly on usable security and privacy, including authentication, privacy transparency, and tools for helping users make better security decisions. He received his PhD and MS from Carnegie Mellon University and his AB from Harvard University. blase@uchicago.edu



Sean Segreti is a Security Consultant, Developer, and Passwords Researcher at KoreLogic. Segreti maintains Carnegie Mellon University's Password Guessability Service, which is used by over 30 universities to estimate password strength. Segreti holds a master's degree in electrical and computer engineering (ECE) from Carnegie Mellon University, and a bachelor's degree in electrical engineering from the University of Maryland. ssegreti@cmu.edu

In this article, we discuss how we use neural networks to accurately measure password strength, and how we use this capability to build effective password meters. First, we show how neural networks can be used to guess passwords and how we leveraged this method to build a password guesser to better model guessing attacks. We report our measurements of the effectiveness of neural networks at guessing passwords, demonstrating that they outperform other popular methods of modeling adversarial password guessing. We then show how we developed a password guesser that can be compressed so that it is practical for client-side use inside a Web page [1]. Finally, we describe how we designed and built a password meter, based on neural networks, that gives more accurate and helpful guidance to users for creating passwords that are resistant to guessing attacks [2].

Passwords are the most common authentication mechanism in use today. We all use passwords every day and will likely continue to do so for the foreseeable future. Unfortunately, human-chosen passwords often follow predictable patterns. For example: exclamation points are at the end; capital letters are at the beginning of passwords; dictionary words, well-known phrases, keyboard patterns, and names of people and places are all common. Such predictable patterns allow attackers to break into accounts by guessing passwords.

Guessing attacks can take the form of online attacks in which attackers make guesses while trying to log in to a live system. Online attacks are sometimes defended against by limiting the rate at which attackers may make guesses against the system. In contrast, in offline guessing attacks, attackers can make large numbers of guesses without limits. This commonly happens when a database of hashed passwords is stolen, an event that occurs with disappointing regularity. Attackers guess candidate passwords and compare them against hashed passwords in the database, limited only by the amount of computer resources they have. The widespread incidence of password reuse makes such attacks more dangerous because attackers who crack a user's password that was leaked from a stolen database may use that cracked password—or common variations of the password—to guess the credentials for that user's other accounts. A common and effective defense against both online and offline guessing attacks is to urge users to create less predictable passwords that are more resistant to guessing.

To understand how to guide users to make less guessable passwords, our research group has studied methods for modeling how attackers guess passwords. Previous approaches for modeling password-guessing attacks include statistical approaches, and tools used in adversarial password cracking. Statistical methods, such as Markov models and probabilistic context-free grammars, work by deriving statistical properties from lists of training passwords. Adversarial password cracking tools, such as John the Ripper and Hashcat, are typically used in practice for their ability to crack hashed passwords quickly; often they are configured by experts to craft special password cracking rules for specific password sets. Prior work from our group has studied these approaches and shown how the combination

Better Passwords through Science (and Neural Networks)



Lujo Bauer is an Associate Professor in the Electrical and Computer Engineering Department and in the Institute for Software Research at

Carnegie Mellon University. His research interests span many areas of computer security and privacy, and include building usable access-control systems with sound theoretical underpinnings, developing languages and systems for run-time enforcement of security policies on programs, and generally narrowing the gap between a formal model and a practical, usable system. His recent work focuses on developing tools and guidance to help users stay safer online and in examining how advances in machine learning can lead to a more secure future. lbauer@cmu.edu



Nicolas Christin is an Associate Research Professor at Carnegie Mellon University, jointly appointed in the School of Computer Science and in

Engineering and Public Policy. He holds MS and PhD degrees in computer science from the University of Virginia. His research interests are in computer and information systems security; most of his work is at the boundary of systems and policy research. He has most recently focused on security analytics, online crime modeling, and economics and human aspects of computer security. nicolasc@cmu.edu



Lorrie Faith Cranor is a Professor of Computer Science and of Engineering and Public Policy at Carnegie Mellon University where she

is director of the CyLab Usable Privacy and Security Laboratory (CUPS). She is Associate Department Head of the Engineering and Public Policy Department and Co-Director of the MSIT-Privacy Engineering masters program. In 2016 she served as Chief Technologist at the US Federal Trade Commission. She is also a co-founder of Wombat Security Technologies, Inc., a security-awareness training company. She is a fellow of the ACM and IEEE and a member of the ACM CHI Academy. lorrie@cmu.edu

of multiple automated approaches approximates the ability of professional human experts to guess passwords [3]. However, modeling a guessing attack in which attackers can make large numbers of guesses often requires servers with tens of CPU cores and with gigabytes of disk space for storing models of password guessing. Such models are not practical for giving real-time feedback to users during password creation; users can't download gigabytes of data or wait days or weeks to get feedback for creating a password.

Due to the challenges of accurately modeling password attacks, most password meters are unable to provide data-driven, principled feedback to users during password creation. Meters will typically calculate some combination of a variety of heuristics—such as the number of special characters used or the length of the password—which often has little correlation to the resistance of passwords to guessing attacks [4]. When faced with such meters, users often make predictable modifications in order to satisfy the meter's strength estimate, such as adding an exclamation point to the end of their password. However, because attackers are also aware of the predictable patterns in password construction, such modifications do little to improve the password's resistance to guessing. In addition, meters are often incapable of providing positive advice or giving users suggestions about how to make passwords better, instead rating a password as simply “weak” or “fair.”

Designing a Neural Network Guesser

Neural networks are a machine-learning technique that is particularly adept at fuzzy classification problems and problems dealing with computer processing of natural language. The intuition for our approach was that, because the task of guessing passwords in an adversarial attack is conceptually related to generating natural language, neural networks would be well suited to our goal of modeling guessing attacks. Recently, the machine-learning community has showed how to use neural networks to generate text, which our approach leverages [5]. Generating a password with a neural network involves repeatedly predicting the next character of a password to build up the password one character at a time. This process can be extended to generate large numbers of probable passwords. During training, the neural network is taught to predict the next character when given a real password fragment. The neural network can then learn to recognize high-level patterns that often arise in password construction, such as keyboard patterns or exclamation points at the end of a password.

We tried many different variations and tunings for training our neural network guesser. When training neural networks, there is a large design space of different parameters and design decisions to explore for better performance. We experimented with a wide range of different parameters including: the number of parameters in the model; the method of representing password characters; different recurrent neural-network architectures; using different types of training data; and using a technique called transference learning, which specializes neural network predictions for different situations. At the end of these experiments, we had a neural-network training methodology that we found was most accurate for our application of guessing passwords. Additionally, we used a technique of modeling password guessing to arbitrarily high numbers of guesses by employing Monte Carlo methods [6], allowing us to accurately model password guessability against nation-states or other extremely powerful adversaries who have huge resources for cracking passwords.

When designing our neural-network guessing method, we tested it against the best tunings of other methods for guessing passwords. In addition, during development of our neural-network guesser, we comprehensively tested various different versions of the neural-network guesser against each other to find the best method. We measured the performance of our guessing approaches both on real passwords collected in recent password leaks and on passwords we have collected in our research studies, allowing us to compare the performance of guessing methods in a wide variety of password policies and situations. To train our guessing

Better Passwords through Science (and Neural Networks)




Figure 1: Comparison of the ability of different password methods to guess passwords. The x-axis of each graph shows the number of guesses made in log scale. The y-axis shows the percent of passwords guessed. Higher lines on the graph represent more accurate guessing. “Neural” shows the performance of our neural-network approach; “Markov” the Markov model approach; “PCFG” probabilistic context-free grammars; “JTR” John the Ripper; “Hashcat” shows the performance of Hashcat; and “MinGuess” shows a combination of all approaches, where a password receives the minimum guess number from all approaches. Each graph shows passwords created under a different policy—requiring a different minimum length and different mix of character classes (uppercase and lowercase characters, digits, and symbols).

methods in our experiments, we required large numbers of real passwords, which we obtained from leaked password lists. In total, our data set of passwords contained over 100 million passwords from more than 20 password leaks. This huge amount of data on real-world passwords allows machine-learning techniques to infer deep insights into password construction and to have the predictive power to model common password patterns.

We found that the neural networks guessed passwords more accurately than any other individual method. However, while our best-performing neural networks often performed close to an optimal guessing strategy, the combination of all methods including neural networks (MinGuess in Figure 1) performed better than just neural networks alone, showing that a combination of many models is still better than any individual method. Nonetheless, if one is limited to only one method for estimating password strength, neural networks are the most accurate. Figure 1 shows a selection of some of our results on guessing accuracy for different conditions; the neural network approach guesses a larger proportion of passwords over the same number of guesses than other methods. This finding holds to various degrees across all of our test sets, although we find that neural networks are particularly accurate when guessing passwords made under the more exotic, stronger password policies, which are becoming increasingly common as password guessing abilities increase.

Designing a Client-Side Strength Estimator

Besides increasing the accuracy of existing password strength models, we also strove to develop more practical models. Previous methods for modeling adversarial password cracking require large amounts of disk space or bandwidth—hundreds of megabytes or gigabytes—and take hours or days to calculate measures of password strength. In contrast, to give real-time feedback

to users during password creation, models must be smaller to download and give quick results. For this application, we wanted a model that was less than one megabyte to download, which is roughly half the size of an average Web page. Additionally, in the context of real-time feedback, a model must calculate a measurement of password strength within a fraction of a second—ideally below the threshold of human recognition, which is roughly 100 ms. In addition to these properties, the measurement should be accurate, and the model should run inside of a Web browser, which means that JavaScript is the most viable execution platform.

Given the challenges of implementing accurate password-strength measurement on resource-constrained clients, it might be tempting to use a system architecture where the password model is stored on a server and only measurement results are communicated to the client. However, in many situations, the user’s password should never be sent to the server for security reasons, for example, in the case of device encryption software, keys that protect cryptographic credentials, or the master password for a password manager. Even in cases where the user’s password is eventually sent to an external server, using a remote password-strength measurement mechanism may allow powerful side channels based on keyboard timing, message size, and caching [7]. For these reasons, we preferred architectures where password modeling and strength estimation are done entirely on the client side. This design decision has the added benefit of being easier for Web administrators to deploy.

To summarize our technical approach to meeting these goals: we started by training a neural network with fewer parameters—the features of the model that define how to predict the next character. Using this less complex model made the network smaller, but did not sacrifice much accuracy compared to our best-performing network. Then we reduced the precision of the already shrunken neural network’s parameters, again trading off

Better Passwords through Science (and Neural Networks)

space for some accuracy. Finally, we used standard lossless compression methods to further shrink the size of the model, eventually reaching a model size of 850 KB. To make our network produce low-latency results, we pre-computed an approximate mapping for estimating the strength of the password, which was sent to the client along with the network. In addition, we cached specific intermediate computations, so that the common case, in which a character is added to the end of the password, is quicker because the strength estimator only needs to update its previous computation. We were able to get the average response time to be 17 ms for this common case. Some of our optimizations sacrificed accuracy for the sake of quicker results or a smaller model; we empirically measured the impact that such optimizations introduced and found the error rate to be small enough to be acceptable for our purposes. In addition, we tuned the network so that it was much more likely that we would make safe errors—underestimating a password’s strength—than unsafe errors.

We compared the accuracy of our client-side strength estimation based on neural networks to existing password meters: “zxcvbn” and Yahoo’s password meter. zxcvbn, in particular, measures password strength using a number of highly tuned heuristics for password strength. We found our method of measuring password strength to be more accurate—correlating more highly with password strength measured by simulating a guessing attack—than either meter, having between 39% and 83% fewer unsafe errors, depending on the meter and the password policy. At the same time, our strength measurement also had fewer safe errors. In addition, our more principled method of simulating adversarial guessing entirely on the client-side has the benefit that it can be easily reconfigured—by re-training the neural network—for new password policies or new situations. We know that certain password sets often have special patterns that are unique to that set: for example, passwords for a sports Web site may contain more sports terminology than other password sets. Our method would be able to be easily retargeted to learn such patterns.

Design of a Password Meter

While the development of an accurate client-side strength-estimation tool is necessary for a password meter, it is not sufficient. There is a gap between a practical measurement of strength and providing effective real-time feedback about how to make a better password. We wanted to bridge this gap. Our main goal was to give human-understandable feedback about password creation; our neural-network strength estimation by itself can tell the user that a password is weak or strong, but it cannot say how to improve the password to be more resistant to guessing. To accomplish this, we aimed to give two types of suggestions: first, we wanted to be able to provide concrete suggestions for specific passwords that are stronger; second, we wanted to provide users with high-level guidance specific to their exact




Figure 2: Screenshot of our password meter’s interface. The bar shows the strength estimate of the user’s password. The popup dialog shows specific password feedback based on the user’s password.

situation—for example, notifying users that using capital letters at the beginning of the password is a common pattern and does not meaningfully improve the strength of their password.

We developed a password meter that achieves these goals. Our meter combines the accuracy of our neural-network strength measurement with a series of data-driven heuristics that provide human-understandable feedback about the user’s password. Figure 2 shows an example of our meter in action. Our meter uses the neural network to control the bar that shows how strong the user’s password is, while data-driven heuristics additionally give the user specific feedback about how to improve their password. The meter can also provide a concrete suggestion for how to change the password so that it will be stronger. It does so by creating several candidate suggestions that are similar to the user’s chosen password and then using the neural network to gauge their strength. Only those candidate passwords that are judged stronger by the network are shown to the user.

We tested whether the meter helps users to create stronger passwords. We recruited participants to create a password for a hypothetical high-value online account in a variety of different conditions—some participants used our meter during password creation, some used modified versions of our meter, and some did not have the benefit of any meter. Similar methodology has been used in prior work by our group for measuring the impact of a variety of different conditions on the security and usability of human-chosen passwords [8, 9].

We found that participants who used the meter created passwords that were 44% more resistant to guessing attacks than those who did not. Interestingly, we also found that participants who saw the human-readable suggestions produced even stronger passwords than those who only saw the measurement of strength. This implies that not only does providing real-time

Better Passwords through Science (and Neural Networks)

strength estimates help users make stronger passwords, but also that providing actionable suggestions about what users should do provides additional benefit.

Conclusion

We showed how neural networks can be used to guess passwords and that they can do so more accurately than other methods for adversarial password guessing. We also showed how leveraging neural networks can lead to more practical estimations of password strength on resource-constrained client machines in real time. Finally, we built and tested a password meter, based on neural networks, that gives human-understandable feedback and guides users to make better passwords. We have released our meter as open source software (at https://github.com/cupslab/neural_network_cracking and https://github.com/cupslab/password_meter) and invite people to use it.

Acknowledgments

We would like to thank Mahmood Sharif for participating in discussions about neural networks and Dan Wheeler for his feedback. This work was supported in part by gifts from the PNC Center for Financial Services Innovation, Microsoft Research, John & Claire Bertucci, and a gift from NATO through Carnegie Mellon CyLab.

References

- [1] W. Melicher, B. Ur, S. M. Segreti, S. Komanduri, L. Bauer, N. Christin, L. F. Cranor, “Fast, Lean, and Accurate: Modeling Password Guessability Using Neural Networks,” in *Proceedings of 25th USENIX Security Symposium*, 2016: <http://bit.ly/2fB18Jd>.
- [2] B. Ur, F. Alfieri, M. Aung, L. Bauer, N. Christin, J. Colnago, L. F. Cranor, H. Dixon, P. E. Naeini, H. Habib, N. Johnson, W. Melicher, “Design and Evaluation of a Data-Driven Password Meter,” in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, ACM, 2017: <https://doi.org/10.1145/3025453.3026050>.
- [3] B. Ur, S. M. Segreti, L. Bauer, N. Christin, L. F. Cranor, S. Komanduri, D. Kurilova, M. L. Mazurek, W. Melicher, R. Shay, “Measuring Real-World Accuracies and Biases in Modeling Password Guessability,” in *Proceedings of the 24th USENIX Security Symposium*, 2015: <https://www.usenix.org/system/files/conference/usenixsecurity15/sec15-paper-ur.pdf>.
- [4] X. de Carné de Carnavalet and M. Mannan. “From Very Weak to Very Strong: Analyzing Password-Strength Meters,” in *Proceedings of the 18th Network and Distributed System Security Symposium*, 2014: <https://www.ndss-symposium.org/ndss2014/programme/very-weak-very-strong-analyzing-password-strength-meters/>.
- [5] I. Sutskever, J. Martens, and G. E. Hinton. “Generating Text with Recurrent Neural Networks,” in *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, http://www.icml-2011.org/papers/524_icmlpaper.pdf.
- [6] M. Dell’Amico and M. Filippone, “Monte Carlo Strength Evaluation: Fast and Reliable Password Checking,” in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 2015: <https://doi.org/10.1145/2810103.2813631>.
- [7] D. X. Song, D. Wagner, and X. Tian, “Timing Analysis of Keystrokes and Timing Attacks on SSH,” in *Proceedings of the 10th USENIX Security Symposium*, 2001: https://www.usenix.org/legacy/events/sec01/full_papers/song/song.pdf.
- [8] S. Komanduri, R. Shay, P. G. Kelley, M. L. Mazurek, L. Bauer, N. Christin, L. F. Cranor, and S. Egelman, “Of Passwords and People: Measuring the Effect of Password-Composition Policies,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 2011: <https://doi.org/10.1145/1978942.1979321>.
- [9] R. Shay, S. Komanduri, A. L. Durity, P. Huh, M. L. Mazurek, S. M. Segreti, B. Ur, L. Bauer, N. Christin, and L. F. Cranor, “Can Long Passwords Be Secure and Usable?” in *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*, 2014: <https://doi.org/10.1145/2556288.2557377>.