

Revisiting Monte Carlo Strength Evaluation

Martin Stanek¹

¹Department of Computer Science, Faculty of Mathematics, Physics and Informatics, Comenius University

Abstract

The Monte Carlo method, proposed by Dell’Amico and Filippone, estimates a password’s rank within a probabilistic model for password generation, i.e., it determines the password’s strength according to this model. We propose several ideas to improve the precision or speed of the estimation. Through experimental tests, we demonstrate that improved sampling can yield slightly better precision. Moreover, additional precomputation results in faster estimations with a modest increase in memory usage.

Keywords

Password, Monte Carlo, Strength Evaluation

1. Introduction

Passwords remain a frequently used authentication method, despite numerous initiatives, technologies, and implementations aiming for passwordless authentication. Although the popularity of methods such as Windows Hello, Passkey, and WebAuthn has increased, the security of passwords continues to be a significant topic in many application areas.

Evaluating the strength of a password is useful for providing users with feedback on their chosen passwords. This feedback can assist users in selecting stronger passwords. Often, the strength is calculated as a password’s rank, i.e., how many passwords will be generated by some chosen algorithm until our password is produced. There are various tools that calculate the strength of the password, for example `zxcvbn` [1], or password scorer tool in PCFG cracker [2].

Dell’Amico and Filippone proposed a Monte Carlo algorithm that estimates a password’s rank within a probabilistic model [3]. The algorithm works for any probabilistic password generation model, and the authors proved that estimated results converge to the actual ranks.

The Monte Carlo estimator is also used to evaluate and compare different probabilistic models for password generation. The original paper compares n -grams models [4], the PCFG model using probabilistic context-free grammar [5], and the Backoff model [6]. Recent examples of using the Monte Carlo estimator are the evaluation of a password guessing method that employs a random forest [7].

Our contribution. We propose three ideas for improving the precision or speed of the Monte Carlo estimator.

The first idea is to interpolate password’s rank within the sampled interval it belongs, according to its probability. The second idea aims to reduce probability overlap in sampled passwords. Both these ideas, presented in Section 3.2, seek to improve the estimator’s precision. The estimation speed for a password, originally based on binary search, can be enhanced with some additional data computed in advance (the third idea, see Section 3.3). All ideas have been tested experimentally to assess their merit. The results are presented in Section 4. Our experiments demonstrate that improved sampling can yield slightly better precision. However, the effect of interpolation on precision is inconclusive, and we cannot rely on this technique to improve precision.

We utilize the reference implementation of the Monte Carlo estimator, which was published by one of the authors of the original paper on GitHub [8], and we employ the RockYou dataset for our experiments. Given that our focus lies on the estimator itself, the choice of dataset is relatively unimportant.

2. How the Monte Carlo Estimator works


We mostly follow [3] in this section. Let Γ be a set of all allowed passwords. A probabilistic password model aims to capture how humans select passwords, assigning higher probabilities to more frequently chosen passwords and lower probabilities to less common ones. Let $p(\alpha)$ denote a probability assigned to password α by the model, such that $\sum_{\alpha \in \Gamma} p(\alpha) = 1$. Different models yield different probability distributions.

When the model is used for an attack, it enumerates passwords in descending order of probability. Therefore, the strength of a password α is the number of passwords with a higher probability:

$$S_p(\alpha) = |\{\beta \in \Gamma; p(\beta) > p(\alpha)\}|.$$

ITAT’24: Workshop on Applied Security, September 20–24, 2024, Drienica, SK

 martin.stanek@fmph.uniba.sk (M. Stanek)

 © 2024 Copyright for this paper by its authors. Use permitted under Creative Commons Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

Remark. In this context, the authors do not address the possibility that the model may assign identical probabilities to multiple passwords, resulting in a non-monotonic p . The definition of $S_p(\alpha)$ assigns all passwords that share the same probability the lowest rank in their group. This approach can be considered prudent from a security standpoint.

Computing the exact value of $S_p(\alpha)$, for a random α , has prohibitively large time complexity. The Monte Carlo estimator uses sampling and approximation to provide efficient and sufficiently accurate estimation. It relies on two properties of the underlying model:

- The model allows for efficiently computing $p(\alpha)$ for any password α .
- There is an efficient sampling method that generates a password according to the model’s distribution.

Precomputation. The estimator generates a sample Θ of n passwords (sampling with replacement). The sample $\Theta = \{\beta_1, \dots, \beta_n\}$ is sorted by descending probability, i.e., $p(\beta_1) \geq \dots \geq p(\beta_n)$. The cumulative ranks of sampled passwords are calculated as follows:

$$c_i = \frac{1}{n} \sum_{j=1}^i \frac{1}{p(\beta_j)} \quad \text{for } i = 1, \dots, n.$$

The estimator needs to store the probabilities. The cumulative ranks can be easily recomputed. However, both these arrays are usually significantly smaller than representation of the model, see Section 3.

Remark. The implementation [8] uses negative \log_2 probabilities, i.e., scaling $p(\beta_j)$ to $-\log_2 p(\beta_j)$.

Estimation. In order to estimate $S_p(\alpha)$ for some password α , the probability $p(\alpha)$ is computed first. Then the binary search is used to compute the largest index j such that $p(\beta_j) > p(\alpha)$. The result, estimated rank of α is $S_p(\alpha) \approx c_j$. Hence, the time complexity of the estimator is $O(\log n)$.

3. Areas for improvement

3.1. Memory requirement

The RockYou dataset contains more than 14 million unique passwords. The more passwords are used to train a model, the better and more precise results we can expect, such as in our case for password strength estimation. However, there is a point beyond which additional training data provide only negligible improvement, while further increasing the model’s size. Notably, even the set of 10,000 most frequent passwords generates models of

substantial size: 3.17 MB for 4-gram, 7.45 MB for 5-gram, 43.5 MB for Backoff, and 0.99 MB for PCFG. An attempt to use up to 10% of the RockYou dataset for training leads to unacceptable model sizes, where Backoff model being the largest, as shown in Figure 1.

The model defines how passwords are represented, generated, and how their probabilities are calculated. Since these methods are specific for each model, we do not aim to improve the model size. However, the Monte Carlo estimator utilizes an additional arrays, where probabilities and ranks of sampled passwords are precomputed. The original paper [3] experiments with various sample sizes up to 100,000 (having “relative error 1%”), but mostly uses the default sample size of 10,000. The default sample size requires 160 kB of memory¹ and its dominated by the memory required for any model trained on a dataset of reasonable length.

3.2. Precision

The estimator assigns the same rank c_j to any password α for which the probability falls within the range $p(\beta_j) > p(\alpha) \geq p(\beta_{j+1})$. Intuitively, passwords with distinct probabilities should not get the same numeric estimate. Certainly, this is not an issue when the password strength is presented on a reduced scale using descriptive characteristics like *weak – medium – strong – very strong*, or using a traffic lights metaphor *red – amber – green*.

Idea 1. Interpolate rank values within intervals using an appropriate function. The most basic approach, without additional parameters, is linear interpolation. This has no impact on memory complexity and a negligible impact on time complexity. Figure 2 shows a graph of password ranks, on a logarithmic scale, for various models and the sample size of 10,000. It appears that linear interpolation on the logarithmic scale should perform well for these models.

The precision of the estimator depends on the sample size. More specifically, it depends on the number of unique probabilities in the set $P = \{p(\beta_1), \dots, p(\beta_n)\}$. We define the overlap of Θ as the fraction of probability values that are already in the set, and therefore do not contribute to the estimator’s precision: $1 - |P|/n$. Table 1 shows the average overlap for different models and sample sizes. Surprising differences in overlap are observed among different models. An expected increase in overlap is observed with an increasing sample size, since the overlap depends substantially on password probability distribution, given by the model from which the passwords are generated. On the other hand, a larger training dataset results in greater diversity of passwords, leading to slightly lower overlap.

¹Real numbers are represented as the `numpy.float64` datatype.

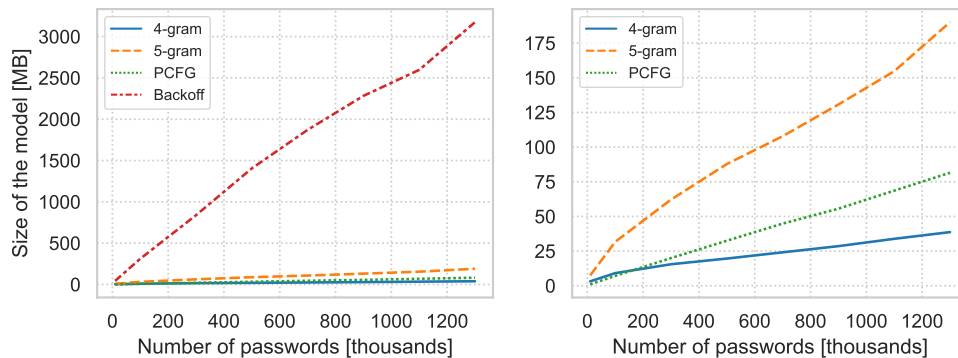


Figure 1: Size of the model reflecting the number of passwords in a training dataset. The graph on the right excludes the Backoff model to show other three models more clearly.




Figure 2: Password ranks corresponding to the position in the sample

Idea 2. The estimator will sample random passwords for Θ until it gets n unique probabilities. It *compresses* sample by discarding duplicate probabilities in such a way that preserves the cumulative sum of the entry with the largest index. Hence, the rank calculation remains intact, and the overlap of the resulting Θ is be 0. Since the sampling is done in precomputation phase, it does not impact the estimation time or memory complexity in any way.

Table 2 shows how many passwords must be sampled using a trained model to achieve the target size of the sample with distinct probabilities.

3.3. Estimation speed

The binary search employed in the original estimator is fast enough for assessing individual passwords. However, when the estimator is used to evaluate or compare different models and their variants, the ranks of a large number

of passwords need to be estimated. An optimization can be relevant in these scenarios.

Idea 3. Divide the interval of possible probability values $p(\beta_i)$, in our case expressed as negative \log_2 values, into t intervals (bins): $[0, \tau_1), [\tau_2, \tau_3), \dots, [\tau_{t-1}, \infty)$, where $0 < \tau_1 < \dots < \tau_{t-1}$. For each interval, for $1 \leq i \leq t$, we calculate minimal and maximal look-up indices that narrow interval for binary search (we use $\tau_0 = 0$ in the following equations):

$$\begin{aligned} \text{LU}_{\min}(i) &= \max\{1 \leq j \leq n \mid -\log_2 p(\beta_j) \geq \tau_{i-1}\}, \\ \text{LU}_{\max}(i) &= \min\{1 \leq j \leq n \mid -\log_2 p(\beta_j) < \tau_i\}. \end{aligned}$$

The estimator is adapted accordingly. Given a password α , we calculate an appropriate interval such that $-\log_2 p(\alpha) \in [\tau_{i-1}, \tau_i)$. Then, the binary search is performed within the set of indices $\{\text{LU}_{\min}(i), \dots, \text{LU}_{\max}(i)\}$, instead of full set $\{1, \dots, n\}$. We expect to narrow the interval for the binary search substantially, so the benefit of fewer comparisons will be measurable. Trivially, the precision of the estimator remains unchanged.

The price paid is the cost of computing LU_{\min} and LU_{\max} arrays, which is simple one-time precomputation, and small memory needed to store these arrays in the estimator².

4. Experiments

We implement the ideas presented in the previous section and present the results of our experiments.

²For example, 100 intervals “cost” approximately 7.8 kB, even with a wasteful representation using Python’s `int` objects for stored indices and lists for the arrays

training set	sample size	4-gram	5-gram	Backoff	PCFG
500,000	10,000	13.6%	16.5%	20.4%	44.6%
	30,000	20.6%	25.8%	31.6%	60.8%
	50,000	24.8%	30.5%	37.3%	67.1%
1,000,000	10,000	12.4%	14.7%	17.2%	43.1%
	30,000	19.1%	22.4%	27.6%	58.5%
	50,000	22.4%	26.7%	33.2%	64.8%

Table 1

Overlap percentage for different models and sample sizes. Models are trained on 500,000 and 1,000,000 passwords using the most frequent passwords from the RockYou dataset. Every number is an average of 3 experiments.

target sample size	Sampled passwords			
	4-gram	5-gram	Backoff	PCFG
10,000	11,689	12,239	12,894	23,483
30,000	38,795	42,032	47,178	122,865
50,000	68,358	75,953	90,123	258,489

Table 2

Average number of sampled passwords required to achieve the desired sample size with distinct probabilities. Models are trained on 500,000 passwords using the most frequent passwords from the RockYou dataset. Every number is an average of 10 experiments, rounded to the nearest integer.

4.1. Precision

The ideas aimed at improving precision apply to the Monte Carlo Estimator, regardless of the underlying model. We do not attempt to modify the models. For example, if password `-1-1-1-1` is assigned inf^3 as negative \log_2 probability in the PCFG model, because the pattern is outside of the trained grammar, we do not try to “fix this”. Moreover, we do not compare the performance of the models to each other.

We assess the impact of our ideas on the real ranks of password generated by the models. Similarly to the original paper [3], we generate all passwords up to some probability threshold. The rank of a password is its position in the list sorted by the probabilities assigned by the model to the passwords.

The first experiment uses the PCFG model trained on 10 million passwords from the RockYou dataset. The threshold for password generation was set at 20, i.e., all passwords with probability at least 2^{-20} were generated – there were 91,693 passwords in this dataset (let’s denote it T). We consider various combinations of proposed ideas:

- original – a reference implementation of the estimator [8];

³Python’s `float(‘inf’)` value

- interpolation – interpolate rank calculation within the interval between two adjacent probabilities (Idea 1);
- sampling – improved sampling with n unique probabilities (Idea 2);
- all – a combination of interpolation and sampling.

Let $\text{rr}(\alpha)$ denote the real rank of password $\alpha \in T$, and let $\text{er}(\alpha)$ denote the rank estimated by a particular variant of the estimator. The *weighted error* of the estimator on the password set T is calculated as follows:

$$\sum_{\alpha \in T} p(\alpha) |\text{er}(\alpha) - \text{rr}(\alpha)|.$$

The weighted error assumes that the estimators are used to assess passwords chosen by humans, following the original distribution. We also consider a *simple error* for completeness:

$$\frac{1}{|T|} \sum_{\alpha \in T} |\text{er}(\alpha) - \text{rr}(\alpha)|.$$

variant	weighted error	simple error
original	16.54	101.11
interpolation	15.33	90.63
sampling	11.79	70.63
all	10.86	63.10

Table 3

Weighted and simple errors of various estimator variants. Every number is an average of 100 experiments.

Table 3 shows the results of our experiment. We performed 100 experiments. We have to warn the reader – the reported errors are sensitive to the particular password distribution sampled into Θ . Unsurprisingly, the sampling (Idea 2) helps to reduce estimation errors in general. The situation with interpolation (Idea 1) is mixed, with a substantial fraction of experiments showing worse




Figure 3: The simple difference error (on the left) and the relative error (on the right) of the original and the "all" strategy. Both graphs display results for 50,000 of the most probable passwords from the PCFG model. The numbers are the average values from 100 experiments.

statistics. The reason is that the interpolation makes the error worse when passwords in Θ already "overshoot" their true ranks. Taking the same rank without interpolation compensates for this. Therefore, interpolation cannot be recommended for improving the precision of the estimator. On the other hand, it helps with the "same rank" problem, when different passwords are assigned the same rank by the estimator.

Figure 3 compares visually the original variant with the "all" variant. It illustrates the simple difference of calculated rank and estimated rank. It also shows the relative error of the estimators. As expected, based on the convergence proof in [3], the relative error is rather small in both cases.

4.2. Estimation speed

We tested two configurations: the first one with 100 intervals (bins), and the second with 1,000 intervals. Negative log probabilities are divided into fixed intervals $[0, 1)$, $[1, 2)$, \dots , $[99, \infty)$ for the first case, and into $[0, 0.1)$, $[0.1, 0.2)$, \dots , $[99.9, \infty)$ for the second case, respectively. Both configurations were tested with four different sizes of Θ . Table 4 shows the relative speed of different variants with respect to the baseline, which is the original algorithm with $|\Theta| = 10000$. The results confirm a moderate speed-up for 100 intervals and a substantial speed-up for 1,000 intervals.

5. Additional observation and conclusion

Since passwords in Θ are generated according to their probability, with sufficiently large sample size, we expect that for some k , the top- k most probable passwords will be in the correct order at the beginning of Θ . Therefore,

sample size	Estimation performance		
	original	100 bins	1000 bins
10,000	1.00	0.92	0.37
30,000	1.08	1.00	0.39
50,000	1.13	1.04	0.40
100,000	1.20	1.10	0.40

Table 4

Average relative estimation performance, where the baseline 1.00 is the estimation performance of the original binary search for the sample size 10,000. Experiment uses 10^6 randomly generated passwords by the PCFG model. Every number is an average of 10 experiments, and rounded to the two decimal places.

simply reporting the order of these top- k passwords by the estimator can be beneficial to the precision. Figure 4 illustrates this phenomenon for the PCFG model and the sample size of 10,000, where approximately the top 180 passwords have the exact rank as their position in Θ . However, further down the precision quickly deteriorates.

An interesting question is if we can improve the estimator's precision by compensating for unusually large or small jumps (differences) between adjacent probabilities in the sampled passwords.

An area outside this paper that deserves further focus is the precision of the estimator for low-probability passwords. The estimator's precision worsens for passwords with high ranks. A potential approach might use a different or additional sampling methods that focus on less probable passwords, so that we can cover this part of the probability space better.




Figure 4: Comparison of original password’s rank estimate and estimate according to the position in sampled passwords (denoted as *fixed*).

[//www.usenix.org/conference/usenixsecurity23/presentation/wang-ding-password-guessing](https://www.usenix.org/conference/usenixsecurity23/presentation/wang-ding-password-guessing).

- [8] M. Dell’Amico, *montecarlopwd*: Monte carlo password checking, 2016. URL: <https://github.com/matteodellamico/montecarlopwd>.

References

- [1] D. L. Wheeler, Zxcvbn: low-budget password strength estimation, in: Proceedings of the 25th USENIX Conference on Security Symposium, SEC’16, USENIX Association, 2016, pp. 157–173. URL: <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/wheeler>.
- [2] PCFG cracker, 2024. URL: https://github.com/lakiw/pcfg_cracker.
- [3] M. Dell’Amico, M. Filippone, Monte carlo strength evaluation: Fast and reliable password checking, in: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, CCS ’15, Association for Computing Machinery, 2015, p. 158–169. URL: <https://doi.org/10.1145/2810103.2813631>.
- [4] A. Narayanan, V. Shmatikov, Fast dictionary attacks on passwords using time-space tradeoff, in: Proceedings of the 12th ACM Conference on Computer and Communications Security, CCS ’05, Association for Computing Machinery, 2005, pp. 364–372. URL: <https://doi.org/10.1145/1102120.1102168>.
- [5] M. Weir, S. Aggarwal, B. Medeiros, B. Glodek, Password cracking using probabilistic context-free grammars, in: 30th IEEE Symposium on Security and Privacy, 2009, pp. 391–405. URL: <https://doi.org/10.1109/SP.2009.8>.
- [6] J. Ma, W. Yang, M. Luo, N. Li, A study of probabilistic password models, in: IEEE Symposium on Security and Privacy, 2014, pp. 689–704. URL: <https://doi.org/10.1145/1102120.1102168>.
- [7] D. Wang, Y. Zou, Z. Zhang, K. Xiu, Password guessing using random forest, in: 32nd USENIX Security Symposium (USENIX Security 23), USENIX Association, 2023, pp. 965–982. URL: <https://www.usenix.org/conference/usenixsecurity23/presentation/wang-ding-password-guessing>.