

# Password Guessers Under a Microscope: An In-Depth Analysis to Inform Deployments

Zach Parish · Connor Cushing · Shourya Aggarwal · Amirali Salehi-Abari · Julie Thorpe

Received: date / Accepted: date

**Abstract** Password guessers are instrumental for assessing the strength of passwords. Despite their diversity and abundance, comparisons between password guessers are limited to simple success rates. Thus, little is known on how password guessers can best be combined with or complement each other. To extend analyses beyond success rates, we devise an analytical framework to compare the types of passwords that guessers generate. Using our framework, we show that different guessers often produce dissimilar passwords, even when trained on the same data. We leverage this result to show that combinations of computationally-cheap guessers are as effective in guessing passwords as computationally-intensive guessers, but more efficient. Our framework can be used to identify combinations of guessers that will best complement each other. To improve the success rate of any guesser, we also show how an effective training dataset can be identified for a given target password dataset, even when the target dataset is hashed. Our insights allow us to provide a concrete set of practical recommendations for password checking to effectively and efficiently measure password strength.

**Keywords** Password Checking, Password Guessers, Passwords, User Authentication

## 1 Introduction

Passwords are presently the most common form of user authentication, providing the first layer of defense in most sys-

tems. User authentication aims to confirm a user's claimed identity, typically by something the user knows (e.g., a password), something they have (e.g., a mobile device), or something they are (e.g., a biometric). Despite decades of research into more secure authentication methods, passwords remain dominant mainly due to their ease to implement and familiarity to most users [5]. Password systems, in spite of their popularity, suffer from many security issues as passwords are mistakenly given to attackers [46], reused across accounts [9], and cracked by guessing attacks [4, 16, 33, 37, 39, 51, 53, 57].

Password guessing attacks are a threat to both accounts (particularly after a data breach of hashed passwords) and hard-disk encryption (where passwords are used as a key). To protect against password guessing attacks, administrators are advised to perform password checking, either *proactively* at the time of password creation or *reactively* through attempting to crack their own password databases [3]. While there are many password guessers available, the administrator's choice of them is critical for effective password checking. However, there is uncertainty as to which guessers to use, and how to train them, for best results.

To make an informed decision, an administrator must understand how password guessers compare to and complement each other under different conditions. Unfortunately, the literature lacks methods to support such decisions and analyses of guesser combinations and training. Our work aims to fill this gap, by creating and applying a framework to put a set of password guessers "under a microscope", in order to support such decisions.

Our contributions are as follows: (1) We create an analytical framework to reveal insights into password guessers' behavior and their ability to complement and substitute each other. Our framework is an asset in identifying sets of complimentary guessers (as shown in our experiments). (2) We apply our framework to perform a comprehensive compari-

Zach Parish\* · Connor Cushing · Amirali Salehi-Abari · Julie Thorpe  
Ontario Tech University, Oshawa, Ontario, Canada

\*Corresponding author: E-mail: zachary.parish@ontariotechu.net

Shourya Aggarwal  
Indian Institute of Technology Delhi, New Delhi, Delhi, India

son between a set of six popular password guessers, across a variety of training conditions. This comparison is arguably the most comprehensive to date, as it compares many aspects of the password guessers, including how they complement each other, how well they generalize, how sensitive they are to training data size, and their success rate over six different training and testing datasets. (3) We show how practitioners can get more bang for their buck by using combinations of computationally-cheap guessers that, when used together, have comparable success rates to computationally-intensive guessers, but are more efficient (i.e., run faster). (4) We perform a comprehensive analysis of six publicly leaked password datasets, to support our investigation on how guesser performance is impacted by different aspects of training data. (5) We describe how a useful similarity metric can be applied to identify a similar (which our results support is best) training dataset for password guessers, even when the target dataset only contains hashed passwords.

Our work has two primary outcomes: (i) Our results allow us to provide a set of recommendations for practitioners performing password checking. (ii) Our analytical framework supports more comprehensive comparisons between password guessers. We discuss use cases regarding how researchers and practitioners can use our framework to understand how additional password guessers can compliment or substitute others.

## 2 Related Work

Unfortunately, it has been repeatedly shown that user passwords are often similar or identical, and are consequently guessable by an adversary [4, 7, 34, 35]. In this section, we review some security concerns with passwords, their counter-measures, and finally how our work fits into the literature.

*Patterns in Passwords.* Many users adopt common strategies for creating their passwords to help them remember their passwords. However, these strategies leave behind specific patterns, which often make passwords more guessable. These patterns include keyboard patterns [43], distribution of character classes (or password structures) [56], replacement of letters with resembling characters (e.g., e to 3) [29], popular topics (e.g., love) [51] and dates [52].

*Reuse of Passwords.* Password reuse weakens password strength. When a password is reused across multiple accounts, the breach of a password in one account could lead to a breach of other accounts. The average password is used for approximately 6 different websites [17], and 77% of users either reuse or modify an existing password [9]. These reused passwords have been exploited in targeted attacks (i.e., against a single target user), with success ranging from

16% in 1000 guesses [39] to 32-73% in 100 guesses when personal information is also incorporated [53].

*Password Composition Policies.* To prevent users from selecting weak passwords, many systems implement password composition policies—sets of rules that a new acceptable password must follow. Common examples of composition policies include a minimum password length and/or the inclusion of characters from multiple character classes (e.g., lowercase, uppercase, numbers, special characters). Despite their practical benefit in strengthening selected passwords [33, 45], overly strict password policies push users to insecure behaviors [6, 28, 33] including writing down passwords [28], reusing passwords [6, 33], or extending a weak password with a special character [33]. Partly due to this usability shortfall, many social-media websites, which are often targets of attacks, choose to adopt less restrictive policies [18].

*Password Meters.* Password meters, by estimating the strength of passwords during creation, encourage users to create stronger passwords [48]. However, most of the heuristic-based meters used in practice don't accurately reflect actual password strength [12]. Recent developments focus on various approaches, such as advanced heuristics-based methods [58], probabilistic methods (e.g., Markov model) [7], and neural networks [37, 39, 47]. Proposals based on neural networks, Markov models, and PCFGs have been found to outperform others [23]. Also, password meters can be personalized either by taking into account a user's personal information (e.g., user profile [30] or previously-leaked passwords [39]) in measuring the password strength, or by providing personalized feedback for password strength improvement [47].

*Password Guessing Tools.* There are many widely-studied guessing tools and techniques for guessing passwords. *Markov models* have been promising in password guessing [16, 38]. *Probabilistic context-free grammars (PCFGs)* [57] (and its extensions [27, 55]) create grammar structure-based password guesses, and has been widely-used (see, for example [4, 7, 9, 32, 35, 48]). The *semantic guesser* [51] expanded PCFGs to exploit semantic patterns in passwords. Recently, *neural network guessers* have drawn considerable attention [26, 37]. The use of multiple guessers has been proposed to measure password strength [49]. While some guessers employ a combination (e.g., PCFGv4 uses OMEN), it is not clear how to confirm they are using the most complementary guessers, nor are there any studies or methods to support their identification. Many password guessers need to be carefully tuned on training datasets to effectively guess passwords of a target dataset. Some password guessers are sensitive to language differences in training data [31], and the similarity between training and target datasets improves guessing success [30], a finding that we

corroborated in just one of our many experiments but using a different method, more data, and more guessers (see Section 4.2.2). However, it was not clear how to identify similar data sets when the target is hashed; we describe a method to do so using our methods in Section 5.1.

We note that practitioners need to make many decisions to implement effective password checking. These decisions include which subset of guessing tools to choose among many available options, and which training dataset to choose. To support these decisions, the literature falls short in systematically understanding guesser behaviors and their ability to complement or substitute one another. This work attempts to address this gap.

### 3 Analytical Framework

The analytical framework presented in this section can be applied to evaluate any set of guessers. It can also be used to evaluate a set of training datasets to identify the best training datasets for password checking. These two use cases are discussed further in Section 5.

We consider a set of  $m$  password guessers  $\mathcal{G} = \{g_1, \dots, g_m\}$  where each  $g_i$  represents a specific guesser (e.g., John the Ripper, OMEN, etc.). We aim to understand how each guesser  $g_i \in \mathcal{G}$  behaves when trained on or tested against particular password datasets, what types of passwords they guess, and how similar one guesser’s behavior is to others. To this end, each guesser  $g \in \mathcal{G}$  will be trained on and tested against a set of  $n$  password datasets  $\mathcal{D} = \{D_1, \dots, D_n\}$ , where each  $D_j$  is a publicly-available password dataset (e.g., RockYou, Twitter, etc.).<sup>1</sup> When a guesser  $g_i \in \mathcal{G}$  is trained on a dataset  $D_j \in \mathcal{D}$ , it can create a password guess list  $L_{ij}$ . To compare various guessers trained on various datasets, we develop some statistics (see Section 3.3) for comparing guessers’ guess lists. Our statistics deploy some pairwise-comparison metrics (see Section 3.2), which use either structural features (see Section 3.1) or the passwords shared between two lists.

#### 3.1 Password Features

For each password  $w$ , we extract two structural features: *password length*  $n_w$  (i.e., the number of its characters) and the *number of character classes*  $c_w$  that it contains. We focus on four distinct character classes: lowercase letters, uppercase letters, numbers, and symbols. For instance,  $w = \text{password!}$  has  $n_w = 9$  and  $c_w = 3$  with three character classes: lowercase letter, number, and symbol.

<sup>1</sup> We use the terminology of “testing against a dataset” when a guesser is guessing the passwords of a target password dataset.

To extract features from password list  $L'$  (e.g., leaked password database or guess list of a guesser), we first aggregate the extracted features of all  $w \in L'$  into a matrix  $\mathbf{V} = [v_{xy}]$  where  $v_{xy}$  is the fraction of passwords in password list  $L'$  which contains  $y$  characters covering  $x$  character classes:

$$v_{xy} = \frac{1}{|L'|} \sum_{w \in L'} \mathbb{1}[c_w = x \ \& \ n_w = y], \quad (1)$$

where  $\mathbb{1}[\cdot]$  is the indicator function, and  $|L'|$  represents the number of passwords in the list.<sup>2</sup> The matrix  $\mathbf{V}$  has a natural probability interpretation: when one selects a password  $w$  from the password list  $L'$  uniformly at random, the password  $w$  contains  $y$  characters from  $x$  character classes with a probability of  $v_{xy}$ . In other words, our matrix  $\mathbf{V}$  captures the joint probability distribution of passwords over character classes and the number of characters. To ease our notations and analyses, we collapse (i.e., flatten) the matrix  $\mathbf{V}$  into a feature vector  $\mathbf{v}$ . We refer to this feature vector as the *structural features* of a password list. This simple representation allows us to preserve the impact of password policies of each password list.

#### 3.2 Pairwise Comparison Metrics

Our deployed pairwise comparison metrics are symmetric, so are computed once for each pair of password lists. While these metrics can use any features, we use either the structural features described in Section 3.1 or the passwords shared between two lists. The use of cosine similarity combined with our proposed structural password features ease the interpretation of our analyses with regard to their connections to the password policies of password datasets. The use of Jaccard Index on the passwords shared between two lists allows us to analyze the degree to which password guessers have complementary behavior on the finest level of granularity. Our metrics have been widely used in information retrieval [1, 21, 44], data mining [2, 15], and other password research [30]. Of particular interest is the *generalized Jaccard index*, which as we show in Section 5 and the Appendix, can be used to not only measure similarity between each guesser’s output, but also measure similarity between a plaintext dataset and a hashed & salted password dataset.

*Cosine similarity* measures the angle between two non-zero vectors. For comparison of two password lists, one can extract structural features from each list, and then compute the cosine similarity on the corresponding feature vectors. The cosine similarity between two password lists  $A$  and  $B$  is given by

$$C(A, B) = \frac{\mathbf{v}_A \cdot \mathbf{v}_B}{\|\mathbf{v}_A\| \|\mathbf{v}_B\|}, \quad (2)$$

<sup>2</sup> The function  $\mathbb{1}[s]$  returns 1 if the statement  $s$  is true; otherwise 0.

where  $\mathbf{v}_A$  and  $\mathbf{v}_B$  are structural feature vectors of  $A$  and  $B$ , respectively.  $\|\cdot\|$  is the Euclidean norm, and  $\mathbf{v}_A \cdot \mathbf{v}_B$  is the dot product of those two vectors. The closer the cosine similarity value is to 1, the smaller the angle between the two vectors is, and the more similar they are. In other words, two lists of passwords with similar feature distributions have a high cosine similarity. We use cosine similarity combined with our proposed structural features for two purposes: (i) comparing the structure of leaked password databases with each other; (ii) comparing the structure of two guess lists.

*Jaccard index* measures the extent two sets overlap with each other, where the intersection of two sets is compared to their union. The Jaccard index between two password lists  $A$  and  $B$  can be computed by

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}. \quad (3)$$

The closer the Jaccard index is to 1, the closer in size the intersection of the sets is to their union, and consequently the more similar two sets are. In other words, two sets of passwords with high overlap will have a high Jaccard index. The Jaccard index also has a natural probabilistic interpretation: if one chooses a password uniformly at random from either password list, the Jaccard index captures the likelihood of selecting a password belonging to both sets.

When password lists have duplicates (e.g., leaked password datasets), we view the password list as a multiset, a modification of sets that allows for duplicated elements. In these cases, we apply a generalized version of the Jaccard index [30] to preserve the frequency information of password duplicates in password lists. Letting  $o(w, A)$  be the number of occurrences of password  $w$  in password list  $A$ , the *generalized Jaccard index* between two password lists  $A$  and  $B$  is given by

$$J(A, B) = \frac{\sum_{w \in U} \min(o(w, A), o(w, B))}{\sum_{w \in U} \max(o(w, A), o(w, B))}, \quad (4)$$

where  $U = (\Omega(A) \cup \Omega(B))$ , and  $\Omega(A)$  represents the set of unique passwords in the password list  $A$ .

### 3.3 Statistics

Our comparison metrics can be readily used for the comparison of a pair of password lists. However, to compare two guessers thoroughly, it is useful to summarize the comparison metrics of two guessers under different training and testing datasets. This section explains our proposed statistics for summarizing comparison metrics. Our statistics fall into two categories: relating to guessing behaviors and relating to guessing success.

#### 3.3.1 Statistics for Guessing Behaviors

This class of statistics is devised to either compare the guessing behaviors of password guessers with each other, or measure how different training datasets affect the guessing behavior of a given guesser.

Our *guessing similarity* statistic summarizes the similarity of two guessers' guess lists when trained on the same dataset by averaging the comparison metric (e.g, Jaccard or Cosine) of their guess lists over various training datasets. We calculate the *guessing similarity* of two guessers  $g_i$  and  $g_j$  by

$$G(g_i, g_j, M) = \frac{1}{n} \sum_{k=1}^n M(L_{ik}, L_{jk}), \quad (5)$$

where  $M \in \{C, J\}$  is either Cosine similarity (see Eq. 2) or Jaccard index (see Eq. 3), and  $L_{ik}$  is the list of password guesses (without any duplicates) generated by  $g_i$  trained on datasets  $D_k$ . Here,  $n$  is the number of datasets in  $\mathcal{D}$ . We also introduce *successful guessing similarity* to measure how two guessers' successful guesses are similar:

$$SG(g_i, g_j, M) = \frac{1}{n(n-1)} \sum_{k=1}^n \sum_{\ell \neq k}^n M(L_{ik} \cap D_\ell, L_{jk} \cap D_\ell). \quad (6)$$

One might be interested in measuring how similarly two different password datasets can train guessers. To this end, we introduce our *training similarity* statistic which calculates the extent two different training password datasets result in generating similar guess lists of passwords when used for training. We define *training similarity* between two datasets  $D_j$  and  $D_k$  by

$$T(D_j, D_k, M) = \frac{1}{m} \sum_{i=1}^m M(L_{ij}, L_{ik}), \quad (7)$$

where  $m$  is the number of different guessers in  $\mathcal{G}$ . This formula computes how similarly  $D_j$  and  $D_k$  can train guessers on average. By capturing the extent two various datasets are effectively similar in training guessers, one can identify training datasets which are as effective as another dataset in training guessers. This could be used to identify effective, yet small datasets, which could drastically speed up the training process.

#### 3.3.2 Statistics for Guessing Success

The *guessing success* statistics quantify the guessing accuracy of guessers under various settings (e.g., training and testing datasets), and also determine how training data affects guessing success for various guessers.

When each guesser  $g_i$  is trained on password dataset  $D_j$  and tested against password dataset  $D_k$ , one can compute

its *success rate*, as the portion of successfully guessed passwords, by

$$s_{ijk} = \frac{|L_{ij} \cap D_k|}{|D_k|}. \quad (8)$$

Note that  $s_{ijk} \in [0, 1]$ , where  $s_{ijk} = 1$  implies that all passwords in  $D_k$  are guessed successfully by  $g_i$  trained on  $D_j$ . To summarize the success rate for a specific guesser  $g_i$ , one can compute its mean success rate over all distinct training and testing datasets by

$$\langle s_{i::} \rangle = \frac{1}{n(n-1)} \sum_{j=1}^n \sum_{k \neq j}^n s_{ijk}. \quad (9)$$

We similarly compute the success rate of training dataset  $D_j$  by

$$\langle s_{:j} \rangle = \frac{1}{m(n-1)} \sum_{i=1}^m \sum_{k \neq j}^n s_{ijk}, \quad (10)$$

and the average success rate of a fixed dataset  $D_j$  and guesser  $g_i$  by

$$\langle s_{ij} \rangle = \frac{1}{n-1} \sum_{k \neq j}^n s_{ijk}. \quad (11)$$

## 4 Experiments

Our experiments aim to understand the impact of training dataset choice on guessers, the performance of guessers, and how guessers can complement or substitute one another.

### 4.1 Experimental Setup

We choose a variety of different password datasets and guessers.

#### 4.1.1 Password Datasets

Our experiments use a variety of publicly available leaked password datasets, which have been the subject of other password research studies (for example, [9, 22, 53, 54, 56, 59]). We have curated and cleaned these datasets by converting their passwords to Unicode. Table 1 shows the number of total and unique passwords in each dataset as well as the ratio between those values.<sup>3</sup>

<sup>3</sup> We exclusively use publicly available datasets and don't report any specific password information. Thus, there is no risk of exposing private user information. We keep only the passwords with no links to their original owner.

Table 1: The password datasets, their sizes, and the ratio between unique and total number of passwords. \*Merged contains all other plaintext datasets in this table.

Datasets	Number of Passwords			Type
	Total	Unique	Ratio	
ClixSense [24]	2,222,359	1,628,205	0.7326	Plaintext
Webhost [20]	15,292,021	10,589,775	0.6925	Plaintext
Mate1 [42]	27,403,932	11,988,154	0.4375	Plaintext
RockYou [8]	32,596,319	14,337,716	0.4399	Plaintext
Fling [10]	40,769,652	16,810,091	0.4123	Plaintext
Twitter [11]	40,872,901	22,579,065	0.5524	Plaintext
Merged*	159,157,184	67,628,637	0.4249	Plaintext
LinkedIn [25]	174,243,105	61,829,207	0.3548	Hashed

#### 4.1.2 Password Guessers

To include a wide variety of guesser behaviors, we focus on six guessers from three different classes of password guessers: Markov models, Probabilistic Context Free Grammars (PCFGs), and Neural Networks. All examined guessers are used with their recommended optimal/default settings, or tuned to perform their best on our datasets as described below. Our set of guessers was selected on the basis of two criteria: (i) each guesser is commonly deployed in previous studies and (ii) each guesser has publicly available code, allowing our work to be reproducible.

*John the Ripper (JtR-Markov)*. We use its community build (1.9.0-bleeding-jumbo) [40] in Markov mode. We restrict the maximum length of passwords to 12 characters, which provided the best results and is consistent with other studies [51]. JtR runs single-threaded during both training and guessing.

*Ordered Markov Enumerator (OMEN)*. We use OMEN [16, 41] with the default settings. OMEN produces only ASCII passwords and runs single-threaded during training and guessing.

*Probabilistic Context-Free Grammar (PCFGv4)*. We used PCFG version 4.0 [55], an extension of the original PCFG [57]. This version uses OMEN to generate a certain percentage of passwords and generate the remainder with PCFGs. We have disabled this feature to generate passwords exclusively from PCFGv4 as the use of OMEN decreased the success rate in most of our tests. PCFGv4 runs single-threaded.

*Semantic Guesser (Sem)*. We use the lite version of Sem [50, 51]. The grammars are trained as recommended using maximum likelihood estimation, the backoff algorithm is used for producing tags, and mangling rules are enabled for generating guesses. Sem uses multiprocessing during training, but runs single threaded for guessing.

*Neural Network (NN)*. We generate guesses using the NN's "human" mode [36, 37], and sort them in descending prob-

Table 2: Mean success rates (and std. deviations) for training password datasets. Datasets are ordered smallest to largest.

Training	Success Rate 1m	Success Rate 300m
ClixSense	15.929 (12.634)	33.737 (16.355)
000webhost	8.72 (5.387)	29.602 (11.968)
Mate1	18.337 (13.234)	38.167 (14.799)
RockYou	13.845 (14.037)	30.264 (17.592)
Fling	11.835 (9.158)	35.155 (16.393)
Twitter	20.59 (15.303)	42.815 (17.371)

ability order. We limit the length of passwords to 6–40 characters to maximize NN’s success rate for our datasets. We use a model consisting of three LSTM layers (with 1024 neurons each) and two dense layers (with 512 neurons each). The neural network is our only guesser that uses GPU resources along with CPU. The neural network runs multi-threaded during training and guessing.

*Identity Guesser (ID)*. This guesser takes a training dataset as input, removes its duplicates, and outputs its unique passwords in descending order of their frequency in the training dataset. In other words, this guesser computes the empirical probability distribution of the passwords in the training dataset (i.e., training phase), then outputs the passwords from the highest to the lowest probability (i.e., generation phase). This simple guesser is a valuable benchmark for understanding how well other guessers learn and generalize.

## 4.2 Impact of Training Data Choice

We investigate how guessing success rates are impacted by different aspects of training data. We train all six password guessers on each of the six individual plaintext datasets and test them against every other plaintext dataset, yielding 180 password cracking scenarios. For all guessers, we set the cutoff to 300 million guesses (see Appendix A for some preliminary analyses that motivate our analyses in this section).

### 4.2.1 Size of training dataset

We ask whether the success rate of a guesser, on average, increases with the size of training dataset. Table 2 shows the average success rates of each training dataset over all guessers and target datasets (computed by Eq. 10), with datasets ordered from smallest to largest size. While our largest dataset performs the best, our smallest dataset ClixSense outperforms both Webhost and RockYou, which are over six and fifteen times larger than it respectively. For a formal analysis, we calculated the statistical correlation between the number of passwords in the training dataset and the averaged success rate. The resulting Pearson coefficient of 0.189 ( $p = 0.315$ ) suggests insignificant correlation

between training dataset size and success rate. This result suggests that a larger dataset size isn’t necessarily a requirement for an effective training data set.

### 4.2.2 Similarity between training and target datasets

We next focus on how the similarity between training and target datasets impacts the success rate of guessers. We first compute the cosine similarity and generalized Jaccard index (see Eq. 2 and Eq. 4) between password datasets, and then explore the relationship of these similarities with success rates.

Figure 1a shows that that Mate1, Twitter, RockYou and ClixSense have high structural password similarity (i.e., cosine similarity). Fling and Webhost are dissimilar to other datasets, but similar to each other. Figure 1b suggests that the exact overlap between datasets (i.e., generalized Jaccard similarity) is often low with exceptions for larger datasets (i.e., Fling, Twitter, RockYou and Mate1), likely due to their sizes.

The cross-examination of Figures 5, 1a, and 1b suggest the datasets with higher similarity tend to have mutually higher success rates (e.g., Mate1 and RockYou share high similarity and mutual success rates). Thus, we hypothesize that the similarity between training and testing datasets has a positive effect on success rate. To test this hypothesis, we ran Pearson statistical tests between the similarity metric of any ordered pair of datasets and their success rates. Our cosine similarity and Jaccard metric have correlation coefficients of 0.597 ( $p = 0.00049$ ) and 0.596 ( $p = 0.00049$ ) respectively. Both are significant and large by Cohen’s convention. This further confirms that dataset similarity, structural (cosine) or overlap (Jaccard), is a key factor in success rate. These results complement previous findings [30] on the relationship between the similarity of training and testing datasets and guesser success rates. We note that this is our only experiment with partial overlap with other work [30] by computing cosine similarity and Jaccard index between datasets; however, we use a different set of datasets and guessers, different features for cosine similarity, and a different application of Jaccard index (between datasets rather than between their features). We also go on to show in Section 5.1 how Jaccard Index can be used to measure similarity even when the target dataset is hashed.

### 4.2.3 Training similarity between datasets

We next explore how similarly two datasets can train a guesser using our notion of training similarity (see Eq. 7). This exploration might not have a direct application in password checking, but offers interesting observations for password guessing. Our investigation is motivated by the surprising performance of ClixSense in Table 2. Despite

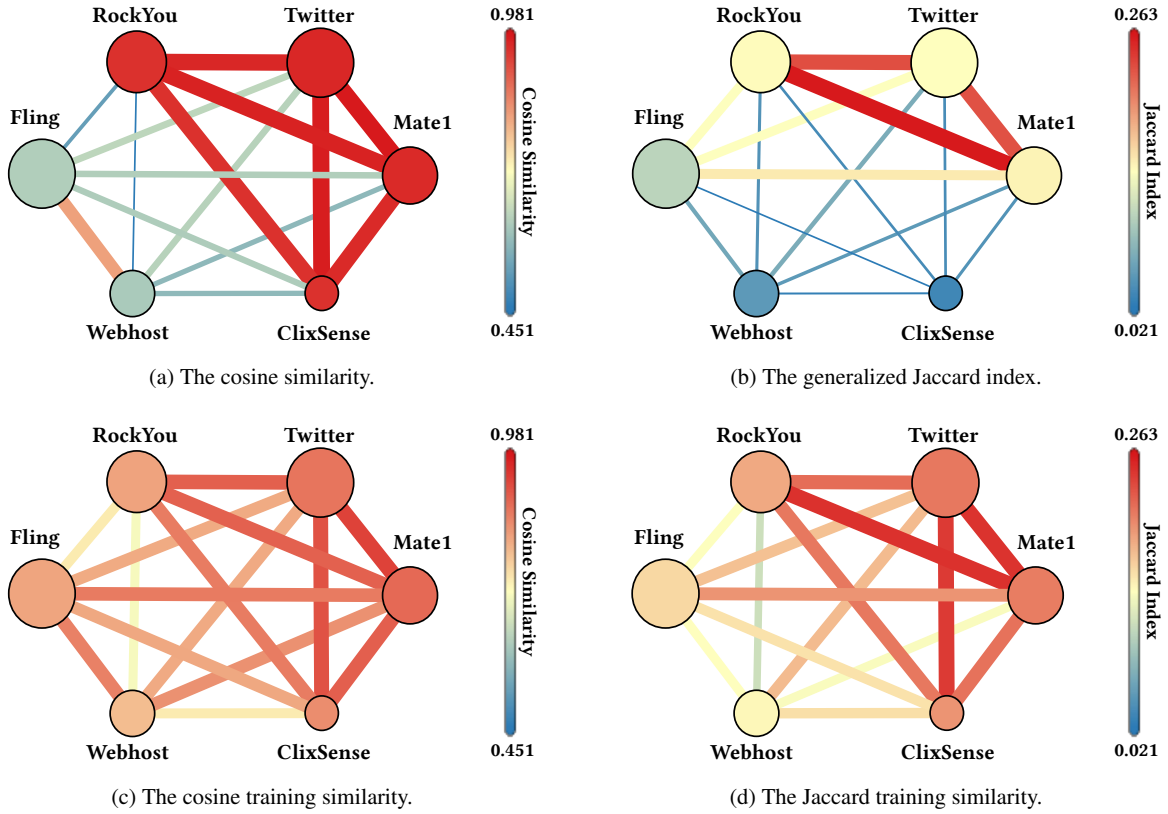


Fig. 1: Plaintext datasets with their pairwise (a) cosine similarity, (b) generalized Jaccard similarity, (c) cosine training similarity, and (d) Jaccard training similarity. The training similarity between datasets is computed by Eq. 7. The edge weights and colors are based on the corresponding metric value between two datasets. The node color captures the metric average for the corresponding dataset. The node size is proportional to the dataset size.

ClixSense’s small size, its performance raises the question of how similarly ClixSense and a bigger dataset can train a guesser, as smaller training datasets may be desirable in some cases to reduce training time. We exclude the Identity guesser in this analysis due to its simplicity in learning; also, its results mirror dataset similarity (see 4.2.2).

Figures 1c and 1d demonstrate the cosine and Jaccard training similarity between our datasets. The cosine training similarity is relatively high between most pairs of datasets. The cluster of RockYou, Twitter, Mate1, and ClixSense share relatively high overlap of generated passwords (see their pairwise Jaccard training similarity). This means passwords generated from training with ClixSense, despite its small size, have high overlap with passwords generated from training with other datasets.

### 4.3 Individual Guesser Performance

To evaluate the performance of each guesser, we compute its average success rate and runtime across varied training data,

target data, and password guessing scenarios (i.e., online and offline attacks).

#### 4.3.1 Guessing Success Rate

To gauge the average performance of each guesser, we train and test every guesser on each possible pair of non-merged plaintext datasets. Then, each guesser generates guess lists at cutoffs of 1 million and 300 million guesses; these cutoffs were selected to simulate online [19] and limited offline attacks, respectively. Table 3 shows the mean success rate of each guesser, computed by Eq. 9. At one million guesses, PCFGv4 and the Identity guesser outperform others, while JtR-Markov and OMEN perform the worst. Notably, only PCFGv4 is able to outperform the Identity guesser at this cutoff with a negligible margin.

For three-hundred million guesses, PCFGv4 performs the best, with a 6% lead over the second best guesser Sem. The Identity guesser performs surprisingly well, with an average of 30.5% (but a high standard deviation of 14.07%) in at most 21,653,268 guesses (compared to 300 million

Table 3: Guessers’ mean success rates at 1 Million and 300 Million guesses (standard deviations in parenthesis). The two best and worst are highlighted with green and red, resp.

Guesser	Success Rate@1M	Success Rate@300M
Identity	23.238 (11.859)	30.519 (14.079)
JtR-Markov	0.665 (0.993)	27.591 (11.563)
OMEN	5.921 (3.225)	22.121 (10.749)
Sem	18.219 (10.344)	41.343 (13.274)
PCFGv4	23.551 (11.545)	47.397 (12.364)
NN	17.662 (11.585)	40.768 (19.734)

Table 4: Guesser training and generation time. Training datasets are randomly sampled from the Merged Dataset. Guessers (except Identity) generated 300M guesses.

Guesser	Training		Generation
	1 Million	50 Million	
JtR-Markov	00h 00m 00.1s	00d 00h 00m 02.2s	00h 00m 33s
Identity	00h 00m 00.3s	00d 00h 00m 24.9s	00h 00m 18s
OMEN	00h 00m 03.0s	00d 00h 00m 23.0s	00h 07m 10s
Sem	00h 01m 38.3s	00d 00h 20m 14.6s	00h 55m 30s
PCFGv4	00h 03m 49.5s	00d 01h 03m 38.4s	00h 30m 58s
NN	01h 18m 08.0s	02d 17h 01m 49.0s	19h 44m 20s

guesses for other guessers).<sup>4</sup> In its best case, the Identity guesser trained on Twitter guesses 56.7% of RockYou, only 10.14% lower than the best guesser PCFGv4 on that same pair. The Identity guesser’s high success rate arises from a relatively large overlap between datasets, observed in Figure 1b. OMEN under-performs JtR-Markov, performing worst overall at this cutoff.

#### 4.3.2 Average Runtime

To help a system administrator understand the resource requirements of guessers, we next analyze their runtimes during training and guess list creation. Each guesser is trained and generates guesses on the same GPU-accelerated server which ran no other jobs. The server has 2 Intel(R) Xeon(R) Gold 6148 CPUs with 80 total cores @ 2.40GHz and 4 Nvidia GeForce 1080 Ti GPUs. We note that only the neural network benefits from multiple GPUs to parallelize computations.

Table 4 reports guesser training and generation time. For training, we created two datasets by sampling 1 million and 50 million passwords from the Merged dataset.<sup>5</sup> For each

<sup>4</sup> The upperbound for number of guesses in the Identity guesser is derived from the maximum number of unique passwords in our datasets.

<sup>5</sup> Our code for training the identity guesser (i.e., computing empirical distribution of unique passwords) and its guess generation (i.e., sorting passwords based on their probabilities) is written in Python without any optimization.

Table 5: Guessers’ generalizability, with 300M guess cutoff. A higher success rate indicates a better ability to generalize.

Identity	OMEN	JtR-Markov	Sem	NN	PCFGv4
15.378	15.664	30.265	33.099	39.585	43.618

guesser, the training time increases with the training dataset size. The Markov-based and Identity guessers perform the fastest (< 25 sec. for 50 million), with PCFGs taking longer (about one hour for 50 million) and the neural network taking the longest (more than 2.5 days for 50 million). For password generation, we observe that the Identity guesser and Markov models are again by far the fastest. Note that the Identity guesser only produced approximately 67M guesses, almost 4.5 times fewer guesses than produced by others. The NN is considerably slower than others: 2100 times slower than JTR-Markov, and even 39 times slower than PCFGv4.

#### 4.4 Guesser Behaviors

We investigate the behavior of each guesser (i.e., their generated guess lists) under various training and target datasets. We also explore how each guesser complements and substitutes others.

##### 4.4.1 Generalizability

One important characteristic of guessers is how well they can generalize, i.e., predict and generate previously unseen passwords. To measure this, we train each guesser on the Webhost dataset as it is the least similar to the other datasets, both in terms of structure (see Figure 1a) and actual password overlap (see Figure 1b). We then test the Webhost trained guessers against every other dataset and calculate each guesser’s mean success rate. Table 5 shows the mean success rate of each guesser: PCFGv4 and NN outperform others, demonstrating a relatively high degree of generalizability compared to others. The Identity guesser and OMEN perform notably worse. This is expected for the Identity guesser with its inability to generalize, but surprising for OMEN. There is a notable amount of variance in the success rates of guessers with similar approaches: 15% difference between Markov models JtR and OMEN, and 10% difference between PCFG-based guessers PCFGv4 and Sem. This highlights how even guessers with similar underlying approaches can display differing generalization behavior.

##### 4.4.2 Sensitivity to Training Size

We intend to learn how each guesser’s success rate is impacted by the size of training data, drawn from the same



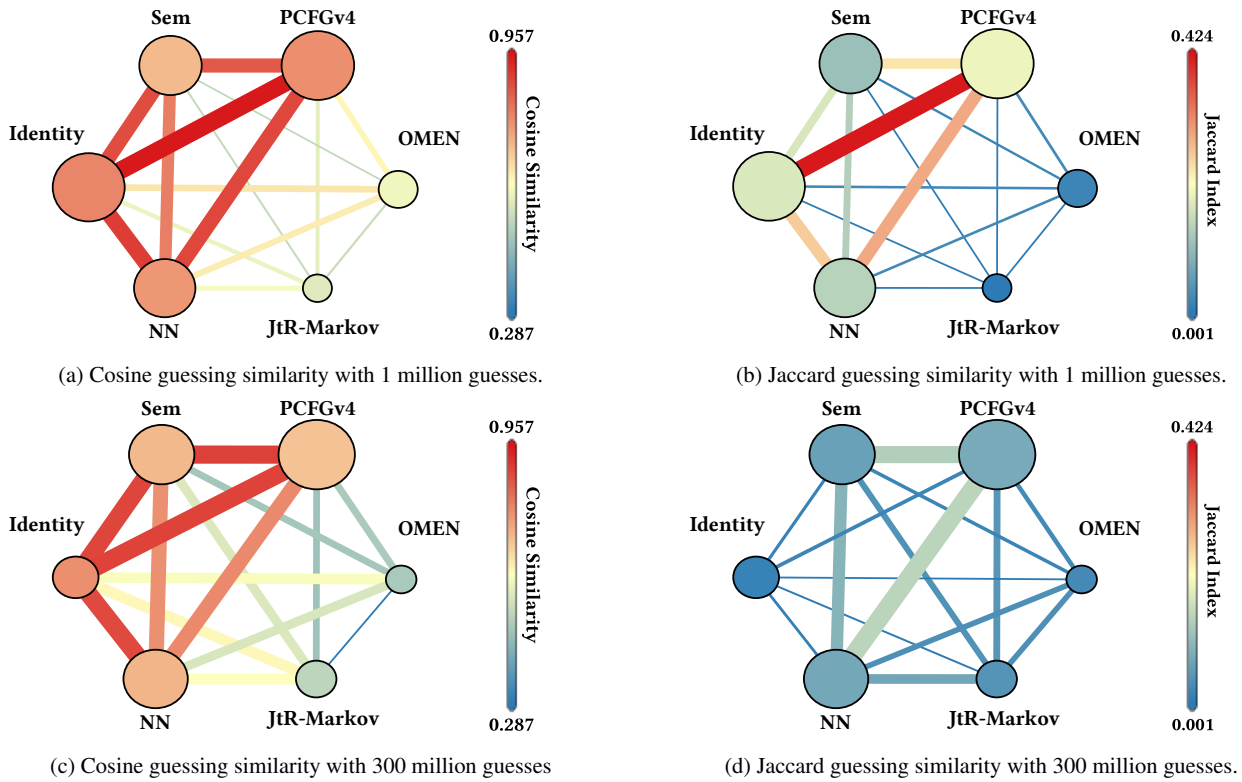


Fig. 2: The cosine and Jaccard guessing similarity (see Eq. 5) between guessers at the cutoffs of 1 million or 300 million guesses. The edge colors represent the similarity value between two guessers. The edge width further highlights the relative similarities within a figure (thicker means more similar). The node size represents the guesser’s average success rate. The node colors represent their average similarity.

Table 6: The mean guessing success rate (and standard deviation in parentheses) for each guesser when trained on different-sized subset of Twitter with a cutoff of 300M.

Guessers	Training Size		
	1 Million	10 Million	30 Million
Identity	21.194(10.172)	33.441(13.458)	39.853(14.186)
JtR	27.570(12.853)	27.541(12.846)	27.527(12.828)
OMEN	29.077(11.383)	29.216(10.916)	29.461(10.942)
Sem	41.493(13.669)	46.910(14.432)	48.021(14.832)
PCFGv4	41.517(11.469)	48.719(13.51)	51.178(14.242)
NN	43.688(13.420)	56.500(14.674)	58.259(14.970)

distribution. Sampling from the Twitter dataset<sup>6</sup>, we create three different datasets of sizes 1 million, 10 million and 30 million. After training guessers on each dataset, we generate guess lists at a cutoff of 300M and test them against all other datasets. Table 6 reports the mean success rates by Eq. 11. All guessers (except JtR-Markov) improve when trained on the larger dataset, but to various extents.

<sup>6</sup> We train on Twitter for this purpose, as opposed to the Merged dataset, since the Merged dataset would contain the testing (target) data.

The Identity guesser has the most drastic improvement with training size growth, from 21.2% to 39.85%. OMEN and JtR-Markov show the least improvement. Sem, PCFGv4, and NN have more modest, but notable improvements, increasing their success rates by 6.5%, 9.7%, and 14.6%, respectively.

#### 4.4.3 Guessing Similarity

Using our notion of guessing similarity (see Eq. 5), we analyze how similar the guess lists of two guessers are when they are trained on the same training data. Figure 2 shows the cosine and Jaccard guessing similarity between guessers at cutoffs of 1 million and 300 million guesses. For both cutoffs, PCFGv4, Sem, ID and NN share high structural (cosine) similarity when compared to OMEN and JtR (see Figure 2a and Figure 2c). Interestingly, despite both deploying a Markov approach, JtR and OMEN are dissimilar. This is likely because OMEN outputs guesses in probability order, whereas JtR-Markov does not.

Figures 2b and 2d show Jaccard guessing similarity between guessers, capturing the overlap of guessers’ guesses, at both cutoffs. Guessers with higher success rates (see Ta-




Fig. 3: The generalized Jaccard successful guessing similarity between guessers. The edge weights and colors represent the similarity of two guessers. The node size represents the guesser’s average success rate. The node color represents the guesser’s average similarity with other guessers.

ble 3) seem to have higher Jaccard guessing similarity (or overlap): At 1 million, the two best guessers PCFGv4 and ID share the highest overlap whereas PCFGv4, Sem and NN with the highest success rates at 300 million have highest overlaps. One can also readily observe that the Jaccard guessing similarities decrease as the cutoff increases. This change suggests that by generating more passwords, each guesser has begun demonstrating their own unique guessing behavior (i.e., the percentage overlap between guessers’ guess lists decreases).

#### 4.4.4 Successful Guessing Similarity

Our guessing similarity analyses showed that guessers trained on the same data, generate mostly unique guesses (see Figures 2b and 2d). However, it is possible that many of these unique guesses are unsuccessful. In this light, one might be interested in measuring the uniqueness of *successful* guesses between guessers. To achieve this, we use our successful guessing similarity metric in Eq. 6 with generalized Jaccard index.<sup>7</sup>

As shown in Figure 3, there is still a considerable degree of uniqueness in successful guesses. Even Sem and PCFGv4—with the highest similarity—have a generalized Jaccard index of 0.86, implying that 14% of their successful guesses are unique to one guesser. Similarly, NN and Sem, by sharing 72% of their successful guesses, owe 28% of their success to unique passwords. Interestingly, the Identity guesser seems to have moderate Jaccard similarity with any other guesser (i.e., its similarity values range from 0.529 to 0.725) despite its smaller guess lists sizes (i.e., ranging from 2.2 million to 40 million compared to 300 million for

<sup>7</sup> The generalized Jaccard allows us to weight the successful guesses of each guesser based on their frequencies in the target dataset.

Table 7: Percentage of LinkedIn passwords successfully guessed. Guessers are trained on the Merged dataset and cut-off at 2 billion guesses.

OMEN	JtR-Markov	Identity	Sem	PCFGv4	NN
35.641	37.028	47.561	55.159	58.798	63.145

all other guessers). These findings offer two important recommendations: (i) the use of one guesser does not make another guesser entirely redundant, even when the underlying approach or achieved success rates are similar; (ii) The cost-effective Identity guesser can complement any other guessers as it has a relatively high number of successful guesses. We explore the gains achieved by combining multiple guessers in our combination attack discussed below in Section 4.5.2.

## 4.5 Combining Guessers

We evaluate the ability of password guessers to complement one another on a previously unseen dataset (i.e., LinkedIn) in an offline attack scenario. We begin in Section 4.5.1 by evaluating each individual guesser against the LinkedIn dataset. Next we analyze different combinations of guessers in Section 4.5.2.

### 4.5.1 Individual Guessers

To compare guessers’ performance, we train each guesser on the Merged dataset, and allow them to each make 2 billion guesses against the LinkedIn dataset. As reported in Table 7, NN outperforms all others, with a 4.3% lead over PCFGv4. PCFG-based (PCFGv4 and Sem) and Identity guessers outperform Markov-based guessers (OMEN and JtR-Markov). Figure 4 depicts the percentage of guessed passwords over the number of guesses. JtR-Markov surpasses OMEN close to the end of the attack.<sup>8</sup> Notably, PCFGv4, Identity, and NN traded places for the best guesser before Identity ran out of guesses. We next apply our findings from our successful guess similarity experiments to further improve the results using combination attacks.

### 4.5.2 Combination Attacks

Our analyses shed light on how guessers complement each other by generating unique successful guesses. We also learn that the Identity guesser not only complements every other guesser, but also often outperforms some advanced guessers.

<sup>8</sup> One might think that JtR-Markov might outperform others with additional guesses. However, our further tests show that even after 20 billion guesses, JtR-Markov only reaches a success rate of 50.568%.




Fig. 4: Performance of guessers trained on the Merged dataset and tested against LinkedIn. The dotted line marks the Identity guesser’s last guess at 67 million guesses, each other guesser made 2 billion guesses.

These findings motivate us to design a *combination attack* where the Identity guesser is used to attack a password dataset prior to the application of a set of other guessers. This hybrid approach is recommended in John the Ripper where a traditional attack follows wordlist mode. We run many independent combination attacks on LinkedIn. Each guesser is trained on the Merged Dataset and produces two billion guesses.

Table 8 reports the result of our combination attacks. When ID is combined with any individual guesser (e.g., ID+O, ID+J, etc.), the combination attacks experience a notable degree of improvement compared to an individual guesser’s performance (compare the columns of sole guesser vs. ID + guesser). JtR-Markov experiences the largest improvement of 18.67%. Even guessers with high success rates (e.g., NN and PCFGv4) realize improvements of 1-4%. By dramatically increasing the success rate of weaker guessers (e.g., OMEN and JTR-Markov), this combined approach makes less resource intensive guessers more competitive.

As shown in Table 8, when more guessers are combined with the Identity guesser, the success rate increases, but with diminishing returns. For example, compare J to J+S (+7.562%), J+S to J+S+P (+3.359%), and J+S+P to J+S+P+N (+1.991%). There seems to be two factors in determining which additional guesser can improve an existing combination attack the most: the success rate of the candidate guesser, and its successful guessing similarities with each of the combined guessers. A candidate guesser with higher success rate has more potential to improve the combined guesser (e.g., compare O+J to O+S). However, a candidate guesser with low successful guessing similari-

ties can be a more effective addition. This interplay of success rate and successful guessing similarities might make a less successful guesser with lower successful guessing similarities more attractive. For example, the weaker JtR and stronger Sem have successful guessing similarities of 0.675 and 0.902 with PCFGv4. The addition of JtR to the combination attack of ID+P offers more improvement than the addition of Sem (3.88% vs. 2.71%).

Each additional guesser also incurs higher runtime and resource requirements. The attacks color-coded green in Table 8 could be completed within one workday (or 8 hours), whereas the yellow and red color-coded attacks must be run overnight (within 8-16 hours) and over two weeks, respectively. The neural network is the largest contributor to runtime in our combinations and also adds GPU requirements. Interestingly, unlike the sole guesser attacks, the slower combination attacks don’t always outperform the faster attacks. For example, the O+J+P attack (65.466%) runs in under 8 hours while S+P (63.866%) and O+S+P (65.250%) take between 8-16 hours, and N (64.875%) and O+N (65.241%) take over 2 weeks. This result implies that competitive success rates can be achieved by the combination of computationally-cheap guessers with less resources. These combination attacks serve as a competitive alternative for practitioners without access to GPU resources, or with time constraints to perform reactive checking (e.g., J+P attack outperforms N while running within a workday and without GPU resources).

## 5 Discussion and Recommendations

We provide further context by presenting use cases of our framework and a set of recommendations based on our empirical results described in Section 4.

### 5.1 Framework Use Cases

Our framework is useful for both practitioners and researchers in:

(1) *Evaluating new guessers and/or settings.* As new password guessers emerge, our framework can be applied to update our knowledge of how to best combine password guessers, by adding the new guesser to  $\mathcal{G}$  and recomputing the formulas in Sec. 3. Our framework equips practitioners and researchers to assess whether or not emerging password guessers offer some complementary power to existing and deployed guessers. This supports a more informed selection of sets of password guessers for password checking. Using our framework supports evaluation beyond typical practices of simply benchmarking individual guesser’s success rates.

Table 8: The percentage of LinkedIn passwords cracked by an offline attack using the Identity guesser followed by a combination of guessers, each making two billion guesses. The names of guessers are shortened to their first letters: (P)CFG, (O)MEN, (N)N, (S)em, and (J)tR-Markov. Each combination attack is color-coded by its runtime for training and guess generation: Green is less than 8 hours (i.e., a workday), yellow is less than 16 hours, and red is over two weeks.

Sole Guesser		ID + 1 Guesser		ID + 2 Guessers				ID + 3 Guessers				ID + 4 Guessers	
guesser	guessed	guesser	guessed	guessers	guessed	guessers	guessed	guessers	guessed	guessers	guessed	guessers	guessed
OMEN	35.641	O	52.272	O+J	57.628	J+P	65.038	O+J+S	63.825	O+P+N	67.336	O+J+S+P	66.931
JtR-M	37.028	J	55.693	O+S	61.536	J+N	65.909	O+J+P	65.466	J+S+P	66.614	O+J+S+N	67.484
Sem	55.159	S	59.773	O+P	62.907	S+P	63.866	O+J+N	66.199	J+S+N	67.247	O+J+P+N	68.060
PCFGv4	58.798	P	61.158	O+N	65.241	S+N	66.411	O+S+P	65.260	J+P+N	67.855	O+S+P+N	68.169
NN	63.145	N	64.876	J+S	63.255	P+N	67.082	O+S+N	66.705	S+P+N	67.943	J+S+P+N	68.605

(2) *Identifying effective training data for password checking.* For guessing scenarios where the target passwords are hashed and salted, our framework can still be applied with the generalized Jaccard index comparison metric, assuming that the salt of each hashed password is available to the administrator, as is typically the case. We describe how this can be accomplished in Proposition 1, the proof for which can be found in the Appendix.

**Proposition 1** *Assuming a candidate training password list  $A$  and salted & hashed password list  $B_h$ , the generalized Jaccard index between  $A$  and  $B_h$  can be computed by:*

$$J(A, B_h) = \frac{F_{\min}(A, B_h)}{|A| + |B_h| - F_{\min}(A, B_h)}, \quad (12)$$

where

$$F_{\min}(A, B_h) = \sum_{w \in \text{supp}A} \min(o(w, A), g(w, B_h)).$$

Here,  $g(w, B_h) = \sum_{y \in B_h} \mathbb{1}[y = H(w + s_y)]$ ,  $\mathbb{1}[\cdot]$  is the indicator function,  $s_y$  and  $H(\cdot)$  are, respectively, the salt and hash function originally used for computation of the salted & hashed password  $y$ . Here,  $|A|$  and  $|B_h|$  are the number of passwords in  $A$  and the number of salted & hashed passwords in  $B_h$ . Also,  $\Omega(A)$  is the set of unique passwords in the training dataset  $A$ , and  $o(w, A)$  is the number of occurrences of  $w$  in  $A$ .

While the offline Identity attack success rate can be used as a proxy for measuring the similarity of a candidate training dataset  $A$  with a salted & hashed password list  $B_h$ , the generalized Jaccard index is more informative. For example, consider two candidate training datasets:  $Y$  with hundreds of millions of entries, and dataset  $Z$  with one million entries. If they each achieve a 50% success rate, dataset  $Z$  should be considered more similar and selected as the best training set. However, the pure Identity attack success rate falls short in distinguishing  $Z$  from  $Y$  as opposed to the generalized Jaccard index which would give a higher similarity score to  $Z$ .

(3) *Identifying complementary guessers for longer offline attacks.* Guessers might offer different complementary patterns for short and long guessing sessions (e.g., online vs. offline attacks). Our experiments show the complimentary patterns of guessers for shorter sessions (up to 300 million guesses). To identify complementary guessers for longer attacks (e.g., approx.  $10^{14}$  guesses), our successful guessing similarity (recall Eq. 6) can be employed, in combination with Monte Carlo methods [13]), to approximate complementary effects of guessers. Rather than directly computing  $L_{ik} \cup D_\ell$  in Eq. 6 by allowing guesser  $g_i$  (trained on  $D_k$ ) to generate the guess list  $L_{ik}$ , one can approximately determine the elements of  $L_{ik} \cup D_\ell$  (i.e., the passwords that would be successfully guessed in a leaked plaintext testing dataset  $D_\ell$ ) by: (1) setting the threshold  $\tau$  for maximum number of guesses (2) for each password  $w \in D_\ell$ , estimate its *guess number* (i.e., the minimum number of required guesses) using Dell’Amico et al.’s approach [13] (3) if the guess number is less than the threshold  $\tau$ , it belongs to  $L_{ik} \cup D_\ell$ . We note that to apply this procedure, each guesser  $g_i$  should be able to assign a probability to a password.

## 5.2 Recommendations

Our work provides a number of practical recommendations (R1-R4) for practitioners auditing their passwords. Of course, this set of recommendations may update as more guessers and training datasets are analyzed using our framework. While our work can be directly applied to reactive checking, it has a natural extension to proactive checking, as guessers that generate probability scores for a given password can be applied as password meters.

*R1: Try publicly-available leaked passwords first.* Our results show that an attacker can be relatively successful by applying the Identity guesser (i.e., the training data of leaked passwords as a guess list) before considering any advanced guessers. This might seem a familiar concept, occasionally applied in practice (e.g., John the Ripper [14]). However, to the best of our knowledge, the impact and benefits of us-

ing an Identity guesser vs. other guessers has not been extensively quantified. For the first 1 million guesses (a number considered feasible for online attacks [19]), the Identity guesser along with PCFGv4 outperform more advanced guessers. For offline attacks, the Identity guesser performed surprisingly well; with only 22 million guesses, on average it achieved 64% of the success rate of the top offline guesser PCFGv4 with 300 million guesses (see Table 3). Additionally, in our LinkedIn experiments, the Identity guesser, with 67 million guesses, had 75% the success rate of the top guesser NN, with 2 billion guesses (see Table 7). These experiments strongly suggest that the Identity guesser can achieve high guessing success rates, comparable to the top guessers, while using at least an order of magnitude fewer guesses. Thus, we strongly recommend that leaked password datasets should be the first priority in password checking.

*R2: Apply combinations of guessers.* Our results for guessing similarity show that the majority of guesses produced by each guesser are unique, even when the underlying approach or success rate is similar. Even for *successful* guesses, each tested guesser is able to crack passwords that others overlook (e.g., the Identity guesser found millions of LinkedIn passwords overlooked by other guessers). Our analysis indicates that no single guesser is able to completely substitute another, and they can complement each other when used together. However, some combinations are more effective than others. Our framework can be used to assist identification of complementary guesser combinations. We also show how some combinations of guessers can have comparably high success rates with lower computing requirements. For example, in less than 8 hours, Identity + PCFGv4 + JtR-Markov can achieve a success rate that compares to Identity + NN (which takes about 2 weeks). Considering both success rate and computing requirements, our results from targeting LinkedIn passwords suggest that a reasonable strategy is to apply this ordering of guessers: Identity, PCFGv4, JtR-Markov, Sem, OMEN, NN. As discussed in Section 5.1, our framework can be used to identify complementary combinations involving additional guessers, and also for long guessing sessions.

*R3: Train with datasets similar to target.* Our results show that when choosing training data, the similarity to the target data is an important factor.<sup>9</sup> Thus, our dataset similarity metric can be used to decide on the most effective training dataset. The most effective dataset can be identified, even when the target dataset is hashed, as outlined in Section 5.1.

<sup>9</sup> These results confirm and complement previous findings [30] by employing different features, more and larger datasets, and more password guessers. We also show how similarity can be measured between a hashed & salted target dataset and a plaintext candidate training set.

*R4: Consider using less training data.* Using more training data takes more computing resources and longer training times. Our results indicate that training dataset size does not correlate with guessing success rates. Although when sampling from the same dataset (Twitter), we observed that data size can increase training effectiveness, the gains between 1 million and 30 million training passwords are not as large as one might expect. Therefore, if time or space constraints exist, a reasonable compromise would be to use a sample of training data from a dataset with high similarity (such as Twitter in our experiments).

## 6 Conclusion and Future Work

We provide an in-depth analysis of password guessers, revealing insights regarding when and how to use them (both alone and in combination). This work demonstrates that combinations of computationally-cheap guessers can be comparably effective to more resource-intensive guessers. Our work also points towards a set of recommendations for practitioners who use password checking tools.

Our framework (i.e., various metrics and statistics) for comparing password guessers and training datasets can be utilized or extended by practitioners and researchers for future password studies. While our present work supports decisions about how to best combine password guessers, there remains some human interpretation of the results—i.e., our framework can help identify the guessers that are most dissimilar and have the highest success rate; however the final decision of which to combine should be made by the human involved (and consider computational efficiency as well). As such, an interesting direction is to develop artificial intelligence algorithms to automate finding combinations of guessers with a maximum success rate under budgeted time and resource requirements. Our present work lays the foundations for such future directions. Another interesting direction for future work is to explore how to summarize a large training dataset into a smaller dataset that trains guessers just as well. Such a smaller training dataset would decrease training time and aim to maximize success rate.

## 7 Compliance with Ethical Standards

### 7.1 Funding

This research was supported by Natural Sciences and Engineering Research Council of Canada (NSERC).

### 7.2 Conflict of Interest

The authors have no conflicts of interest to declare that are relevant to the content of this article.

### 7.3 Ethical Approval

This article does not contain any studies with human participants or animals performed by any of the authors.

### 7.4 Informed Consent

Not Applicable.

## References

- Baeza-Yates, R.A., Ribeiro-Neto, B.: Modern Information Retrieval. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA (1999)
- Berkhin, P.: Survey of clustering data mining techniques. In: Grouping multidimensional data, pp. 25–71 (2006)
- Bishop, M., Klein, D.V.: Improving system security via proactive password checking. *Computers & Security* **14**(3), 233–249 (1995)
- Bonneau, J.: The science of guessing: Analyzing an anonymized corpus of 70 million passwords. In: Proceedings of the 2012 IEEE Symposium on Security and Privacy (S&P), pp. 538–552 (2012)
- Bonneau, J., Herley, C., van Oorschot, P.C., Stajano, F.: The quest to replace passwords: A framework for comparative evaluation of web authentication schemes. In: Proceedings of the 2012 IEEE Symposium on Security and Privacy (S&P), pp. 553–567 (2012)
- Campbell, J., Ma, W., Kleeman, D.: Impact of restrictive composition policy on user password choices. *Behaviour and Information Technology* **30**(3), 379–388 (2011)
- Castelluccia, C., Dürmuth, M., Perito, D.: Adaptive password-strength meters from markov models. In: Proceedings of the 2012 Network and Distributed System Security Symposium (NDSS) (2012)
- Cubrilovic, N.: Rockyou hack: From bad to worse – techcrunch (2009). <https://techcrunch.com/2009/12/14/rockyou-hack-security-myspace-facebook-passwords/>
- Das, A., Bonneau, J., Caesar, M., Borisov, N., Wang, X.: The tangled web of password reuse. In: Proceedings of the 2014 Network and Distributed System Security Symposium (NDSS), pp. 23–26 (2014)
- Das, S.: 40 million fling.com users' passwords, sexual preferences stolen | hacked: Hacking finance (2016). <https://hacked.com/40-million-fling-com-users-passwords-sexual-preferences-stolen/>
- Databases Today: twitter.7z (2019). <https://databases.today/search-nojs.php>
- de Carné de Carnavalet, X., Mannan, M.: From very weak to very strong: Analyzing password-strength meters. In: Proceedings of the 2014 Network and Distributed System Security Symposium (NDSS), pp. 23–26 (2014)
- Dell'Amico, M., Filippone, M.: Monte carlo strength evaluation: Fast and reliable password checking. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, pp. 158–169 (2015)
- Designer, S.: John the ripper password cracker (2002). <https://www.openwall.com/john/>
- Dunham, M.H.: Data Mining: Introductory and Advanced Topics. Prentice Hall PTR, Upper Saddle River, NJ, USA (2002)
- Dürmuth, M., Angelstorf, F., Castelluccia, C., Perito, D., Chaabane, A.: OMEN: Faster password guessing using an ordered markov enumerator. In: Proceedings of the International Symposium on Engineering Secure Software and Systems, pp. 119–132 (2015)
- Florencio, D., Herley, C.: A large-scale study of web password habits. In: Proceedings of the 16th International Conference on World Wide Web (WWW), pp. 657–666 (2007)
- Florêncio, D., Herley, C.: Where do security policies come from? In: Proceedings of the Sixth Symposium on Usable Privacy and Security (SOUPS), pp. 10:1–10:14 (2010)
- Florêncio, D., Herley, C., Van Oorschot, P.C.: Pushing on string: The 'don't care' region of password strength. *Commun. ACM* **59**(11), 66–74 (2016)
- Fox-Brewster, T.: 13 million passwords appear to have leaked from this free web host (2017). URL <https://www.forbes.com/sites/thomasbrewster/2015/10/28/000webhost-database-leak/>
- Frakes, W.B., Baeza-Yates, R. (eds.): Information Retrieval: Data Structures and Algorithms. Prentice-Hall, Inc., Upper Saddle River, NJ, USA (1992)
- Furnell, S.: Assessing password guidance and enforcement on leading websites. *Computer Fraud & Security* **2011**(12), 10–18 (2011)
- Golla, M., Dürmuth, M.: On the accuracy of password strength meters. In: Proceedings of ACM CCS, pp. 1567–1582 (2018)
- Goodin, D.: 6.6 million plaintext passwords exposed as site gets hacked to the bone (2016). <https://arstechnica.com/information-technology/2016/09/plaintext-passwords-and-wealth-of-other-data-for-6-6-million-people-go-public/>
- Hackett, R.: LinkedIn lost 167 million account credentials in data breach (2016). <http://fortune.com/2016/05/18/linkedin->

- [data-breach-email-password/](#)
26. Hitaj, B., Gasti, P., Ateniese, G., Perez-Cruz, F.: Passgan: A deep learning approach for password guessing. In: *Applied Cryptography and Network Security*, pp. 217–237. Springer International Publishing (2019)
  27. Houshmand, S., Aggarwal, S., Flood, R.: Next gen pcfg password cracking. *IEEE Transactions on Information Forensics and Security* **10**(8), 1776–1791 (2015)
  28. Inglesant, P.G., Sasse, M.A.: The true cost of unusable password policies. In: *Proceedings of the 2010 Conference on Human Factors in Computing Systems (CHI)*, pp. 383–392 (2010)
  29. Jakobsson, M., Dhiman, M.: The benefits of understanding passwords. In: *Mobile Authentication*, pp. 5–24. Springer (2013)
  30. Ji, S., Yang, S., Das, A., Hu, X., Beyah, R.: Password correlation: Quantification, evaluation and application. In: *Proceedings of the IEEE Conference on Computer Communications*, pp. 1–9 (2017)
  31. Ji, S., Yang, S., Hu, X., Han, W., Li, Z., Beyah, R.: Zerosum password cracking game: A large-scale empirical study on the crackability, correlation, and security of passwords. *IEEE Transactions on Dependable and Secure Computing* **14**(5), 550–564 (2017)
  32. Kelley, P.G., Komanduri, S., Mazurek, M.L., Shay, R., Vidas, T., Bauer, L., Christin, N., Cranor, L.F., Julio, L.: Guess again (and again and again): Measuring password strength by simulating password-cracking algorithms. In: *Proceedings of the 2012 IEEE Symposium on Security and Privacy (S&P)*, pp. 523–537 (2012)
  33. Komanduri, S., Shay, R., Kelley, P.G., Mazurek, M.L., Bauer, L., Christin, N., Cranor, L.F., Egelman, S.: Of passwords and people: Measuring the effect of password-composition policies. In: *Proceedings of the 2011 Conference on Human Factors in Computing Systems (CHI)*, pp. 2595–2604 (2011)
  34. Malone, D., Maher, K.: Investigating the distribution of password choices. In: *Proceedings of the 21st International Conference on World Wide Web (WWW)*, pp. 301–310 (2012)
  35. Mazurek, M.L., Komanduri, S., Vidas, T., Bauer, L., Christin, N., Cranor, L.F., Kelley, P.G., Shay, R., Ur, B.: Measuring password guessability for an entire university. In: *Proceedings of the 2013 ACM SIGSAC conference on Computer & Communications Security (CCS)*, pp. 173–186 (2013)
  36. Melicher, W.: The neural network password meter (2019). [https://github.com/cupslab/neural\\_network\\_cracking](https://github.com/cupslab/neural_network_cracking)
  37. Melicher, W., Ur, B., Segreti, S.M., Komanduri, S., Bauer, L., Christin, N., Cranor, L.F.: Fast, lean, and accurate: Modeling password guessability using neural networks. In: *Proceedings of the 25th USENIX Security Symposium*, pp. 175–191 (2016)
  38. Narayanan, A., Shmatikov, V.: Fast dictionary attacks on passwords using time-space tradeoff. In: *Proceedings of the 2005 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pp. 364–372 (2005)
  39. Pal, B., Daniel, T., Chatterjee, R., Ristenpart, T.: Beyond credential stuffing: Password similarity models using neural networks. In: *IEEE Symposium on Security and Privacy*, pp. 417–434 (2019)
  40. Peslyak, A.: John the ripper community build (1.9.0-bleeding-jumbo) (2019). <https://github.com/magnumripper/JohnTheRipper>
  41. Ruhr University Bochum, RUB-SysSec: OMEN: Ordered markov enumerator (2019). <https://github.com/RUB-SysSec/OMEN>
  42. Russon, M.A.: Mate1.com hack: 27 million account passwords and emails have been leaked and sold on dark web (2016). <https://www.ibtimes.co.uk/mate1-com-hack-27-million-account-passwords-emails-have-been-leaked-sold-dark-web-1547166>
  43. Schweitzer, D., Boleng, J., Hughes, C., Murphy, L.: Visualizing keyboard pattern passwords. *Information Visualization* **10**(2), 127–133 (2011)
  44. Singhal, A.: Modern information retrieval: A brief overview. *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering* **24**(4), 35–43 (2001)
  45. Summers, W.C., Bosworth, E.: Password policy: The good, the bad, and the ugly. In: *Proceedings of the Winter International Symposium on Information and Communication Technologies*, pp. 1–6 (2004)
  46. Thomas, K., Moscicki, A., Margolis, D., Paxson, V., Bursztein, E., Li, F., Zand, A., Barrett, J., Ranieri, J., Invernizzi, L., Markov, Y., Comanescu, O., Eranti, V.: Data breaches, phishing, or malware?: Understanding the risks of stolen credentials. In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pp. 1421–1434 (2017)
  47. Ur, B., Habib, H., Johnson, N., Melicher, W., Alfieri, F., Aung, M., Bauer, L., Christin, N., Colnago, J., Cranor, L.F., Dixon, H., Emami Naeini, P.: Design and evaluation of a data-driven password meter. In: *Proceedings of the 2017 Conference on Human Factors in Computing Systems (CHI)*, pp. 3775–3786 (2017)
  48. Ur, B., Kelley, P.G., Komanduri, S., Lee, J., Maass, M., Mazurek, M.L., Passaro, T., Shay, R., Vidas, T., Bauer, L., Christin, N., Cranor, L.F.: How does your password measure up? the effect of strength meters on password creation. In: *Proceedings of the 21st USENIX Security Symposium*, pp. 65–80 (2012)

49. Ur, B., Segreti, S.M., Bauer, L., Christin, N., Cranor, L.F., Komanduri, S., Kurilova, D., Mazurek, M.L., Melicher, W., Shay, R.: Measuring real-world accuracies and biases in modeling password guessability. In: Proceedings of the 24th USENIX Security Symposium, pp. 463–481 (2015)
50. Veras, R.: Semantic password guesser (lite) (2019). <https://github.com/vialab/semantic-guesser/tree/lite>
51. Veras, R., Collins, C., Thorpe, J.: On the semantic patterns of passwords and their security impact. In: Proceedings 2014 Network and Distributed System Security Symposium (NDSS), pp. 23–26 (2014)
52. Veras, R., Thorpe, J., Collins, C.: Visualizing semantics in passwords: the role of dates. In: Proceedings of the Ninth International Symposium on Visualization for Cyber Security, pp. 88–95 (2012)
53. Wang, D., Zhang, Z., Wang, P., Yan, J., Huang, X.: Targeted online password guessing: An underestimated threat. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS), pp. 1242–1254 (2016)
54. Wei, M., Golla, M.: The password doesn't fall far: How service influences password choice. In: Proceedings of the 2018 Who Are You?! Adventures in Authentication Workshop (2018)
55. Weir, C.M.: Pretty cool fuzzy guesser (4.0) (2019). [https://github.com/lakiw/pcfg\\_cracker](https://github.com/lakiw/pcfg_cracker)
56. Weir, M., Aggarwal, S., Collins, M., Stern, H.: Testing metrics for password creation policies by attacking large sets of revealed passwords. In: Proceedings of the 2010 ACM SIGSAC Conference on Computer and Communications Security (CCS), pp. 162–175 (2010)
57. Weir, M., Aggarwal, S., De Medeiros, B., Glodek, B.: Password cracking using probabilistic context-free grammars. In: Proceedings of the 2009 IEEE Symposium on Security and Privacy (S&P), pp. 391–405 (2009)
58. Wheeler, D.L.: zxcvbn: Low-budget password strength estimation. In: Proceedings of the 25th USENIX Security Symposium, pp. 157–173 (2016)
59. Zhou, H., Liu, Q., Zhang, F.: Poster: An analysis of targeted password guessing using neural networks. In: Proceedings of the 2017 IEEE Symposium on Security and Privacy (S&P) (2017)



## Appendix A: Preliminary Analyses and Visualization

Figure 5 captures the average success rates for various pairs of training and testing datasets. One can make two important observations: (i) some datasets (e.g., Twitter, Mate1) are more effective as training data than others (e.g., Webhost); (ii) some pairs of datasets are effective for training and testing against each other, i.e., when one dataset can train guessers well against another dataset (e.g., RockYou-Mate1, ClixSense-Mate1, etc.). These two observations motivate us towards a deeper analysis of the characteristics of effective training datasets, discussed in Section 4.2.




Fig. 5: The average success rates over all guessers for training and testing dataset pairs. The edges are directed clockwise from training to testing dataset, with colors matching the training dataset color. The edge width is proportional to the average success rate of guessers for a fixed training and testing dataset pair. The node size shows the dataset size.

## Appendix B: Proofs

*Proof (Proof of Proposition 1)* By Lemma 1 and Lemma 2, the generalized Jaccard index between the password list  $A$  and unhashed password list  $B$  (which is not accessible) can be computed by:

$$J(A, B) = \frac{\sum_{w \in \text{supp} A} \min(o(w, A), o(w, B))}{|A| + |B| - \sum_{w \in \text{supp} A} \min(o(w, A), o(w, B))}, \quad (13)$$

Defining  $g(w, B_h) = \sum_{y \in B_h} \mathbb{1}[y = H(w + s_y)]$  for counting the number of occurrences of password  $w$  in the salted & hashed password list  $B_h$ , we note that  $o(w, B) = g(w, B_h)$  and  $|B| = |B_h|$ . So Eq. 13 is equivalent to:

$$J(A, B_h) = \frac{\sum_{w \in \text{supp} A} \min(o(w, A), g(w, B_h))}{|A| + |B_h| - \sum_{w \in \Omega(A)} \min(o(w, A), g(w, B_h))}.$$

Letting  $F_{\min}(A, B_h) = \sum_{w \in \text{supp} A} \min(o(w, A), g(w, B_h))$ , we derive Eq. 12.  $\square$

**Lemma 1** Let  $o(w, A)$  and  $o(w, B)$  be the number of occurrences of password  $w$  in password lists  $A$  and  $B$  respectively. We have

$$\sum_{w \in (\Omega(A) \cup \Omega(B))} \min(o(w, A), o(w, B)) = |A| + |B| - \sum_{w \in (\Omega(A) \cup \Omega(B))} \max(o(w, A), o(w, B)).$$

Here,  $|A| = \sum_{w \in \Omega(A)} o(w, A)$  and  $|B| = \sum_{w \in \Omega(B)} o(w, B)$  are the number of passwords in  $A$  and  $B$  respectively. Also,  $\Omega(A)$  is the set of unique passwords in  $A$ .

*Proof* One can observe that for any two numbers  $a$  and  $b$ :  $\min(a, b) + \max(a, b) = a + b$ . Using this equality, we can derive

$$\begin{aligned} & \sum_{w \in (\Omega(A) \cup \Omega(B))} [\min(o(w, A), o(w, B)) + \max(o(w, A), o(w, B))] \\ &= \sum_{w \in (\Omega(A) \cup \Omega(B))} o(w, A) + o(w, B) = \sum_{w \in (\Omega(A) \cup \Omega(B))} o(w, A) + \sum_{w \in (\Omega(A) \cup \Omega(B))} o(w, B) \\ &= \sum_{w \in \Omega(A)} o(w, A) + \sum_{w \in \Omega(B)} o(w, B). \end{aligned}$$

The last equality holds as  $o(w, A) = 0$  when  $w \notin A$  and  $o(w, B) = 0$  when  $w \notin B$ . By decomposing the first summation, we have shown

$$\sum_{w \in (\Omega(A) \cup \Omega(B))} \min(o(w, A), o(w, B)) + \sum_{w \in (\Omega(A) \cup \Omega(B))} \max(o(w, A), o(w, B)) = |A| + |B|,$$

where  $|A| = \sum_{w \in \Omega(A)} o(w, A)$  and  $|B| = \sum_{w \in \Omega(B)} o(w, B)$ . By rearranging the terms of this equality, we derive

$$\sum_{w \in (\Omega(A) \cup \Omega(B))} \min(o(w, A), o(w, B)) = |A| + |B| - \sum_{w \in (\Omega(A) \cup \Omega(B))} \max(o(w, A), o(w, B)).$$

$\square$

**Lemma 2** Letting  $o(w, A)$  and  $o(w, B)$  be the number of occurrences of password  $w$  in password lists  $A$  and  $B$  respectively,

$$\sum_{w \in \Omega(A) \cup \Omega(B)} \min(o(w, A), o(w, B)) = \sum_{w \in \text{supp} A} \min(o(w, A), o(w, B)), \quad (14)$$

where  $\Omega(A)$  is the set of unique passwords in  $A$ .

*Proof* Partitioning  $\Omega(A) \cup \Omega(B)$  to two disjoint sets of  $\Omega(A)$  and  $\Omega(B) - \Omega(A)$ , we have

$$\begin{aligned} & \sum_{w \in \Omega(A) \cup \Omega(B)} \min(o(w, A), o(w, B)) = \\ & \sum_{w \in \Omega(A)} \min(o(w, A), o(w, B)) + \sum_{w \in \Omega(B) - \Omega(A)} \min(o(w, A), o(w, B)). \end{aligned}$$

As  $o(w, A) = 0$  for  $w \in \Omega(B) - \Omega(A)$ , we have  $\min(o(w, A), o(w, B)) = 0$  for all  $w \in \Omega(B) - \Omega(A)$ . So we have derived

$$\sum_{w \in \Omega(A) \cup \Omega(B)} \min(o(w, A), o(w, B)) = \sum_{w \in \Omega(A)} \min(o(w, A), o(w, B)).$$

$\square$