Identifying document similarity using a fast estimation of the Levenshtein Distance based on compression and signatures

Peter Coates^{*a*}, Frank Breitinger^{*b*,*}

^aKingswood Road, Weehawken, NJ, 07086, United States ^bSchool of Criminal Justice, University of Lausanne, 1015 Lausanne, Switzerland

ARTICLE INFO

Keywords: Levenshtein Distance Edit Distance Estimation Document Similarity Approximate String Matching Fingerprint Digest

ABSTRACT

Identifying document similarity has many applications, e.g., source code analysis or plagiarism detection. However, identifying similarities is not trivial and can be time complex. For instance, the Levenshtein Distance is a common metric to define the similarity between two documents but has quadratic runtime which makes it impractical for large documents where large starts with a few hundred kilobytes. In this paper, we present a novel concept that allows estimating the Levenshtein Distance: the algorithm first compresses documents to signatures (similar to hash values) using a user-defined compression ratio. Signatures can then be compared against each other (some constrains apply) where the outcome is the estimated Levenshtein Distance. Our evaluation shows promising results in terms of runtime efficiency and accuracy. In addition, we introduce a significance score allowing examiners to set a threshold and identify related documents.

1. Introduction

A crucial task for digital forensic investigations is the identification of relevant artifacts (evidence) where examiners often rely on automation to cope with the sheer amounts of data. One example is file classification, which is often done using hash functions: compute the hash value of a file and compare the hash against a whitelist (in case of a match the file can be ignored) or blacklist (in case of a match the file is regarded as potential evidence that requires further investigation). As hash functions can only identify exact duplicates, they have been complemented by approximate matching (AM) a.k.a. similarity hashing or fuzzy hashing (Breitinger, Guttman, McCarrin, Roussev, White et al., 2014a). Instead of providing a binary decision, most AM algorithms provide a certainty score indicating their confidence that two artifacts are similar or related. For instance, on a scale of 0 to 100, where 0 indicates no similarity, a score of 80 indicates that the tool is confident that the two artifacts have similarity where the definition of similarity is defined by the tool and can be on various levels, e.g., byte-level similarity vs. visual similarities in case of images. While these algorithms provide great value to practitioners, there are two problems: First is the interpretation of the output as a score of 80 does not mean there is a similarity of 80 percent. In fact, two artifacts can have a score of 100 but their cryptographic hashes are different. Second, AM algorithms work best for larger artifacts but are less reliable for medium sized texts, e.g., comparing spam Emails with each other to identify spam campaigns. Which, depending on the length of the text,

Copyright remains with the authors.

URL: https://hadoopoopadoop.com (P. Coates); https://FBreitinger.de (F. Breitinger) is best solved using approximate string matching algorithms such as Levenshtein distance.

Levenshtein distance (LD) LD is precise and has an intuitive meaning regarding the dissimilarity of two strings. Given a pair of strings, the LD is the number of singlecharacter edits (i.e., insertions, deletions, or substitutions) that are required to turn one string into the other. For example, the LD of pat and mat is one (replace m with p); the LD of pats and mat is two (replace m with p and delete the s). The LD of pairs of randomly chosen longer texts of equal length tends to be narrowly distributed. Therefore, an observed divergence of the LD of a pair of texts from the expected mean can be interpreted mathematically and also has an intuitive meaning to a human. Unfortunately, in practice, LD is primarily relevant for short strings (e.g., comparing two strings with 50'000 characters each requires already 3.5sec on a modern Laptop) because the time complexity of the algorithm is quadratic¹ being proportional to the product of the lengths of the two strings, i.e., $(O(|a| \times |b|))$. Quadratic performance does not present a hard size limit, but comparing larger strings, such as Web pages, long articles, or books, becomes impractically slow, and faster computers are of little help when confronted with a quadratic performance curve.

Contribution The heuristic described in this article offers a partial solution to the problem by offering a way to *estimate* the LD of texts pairs over a size range many hundreds of times larger than is practical for calculating an exact value. The algorithm works by first compressing the input (the result is called signature) and then by applying the Levenshtein Distance to the signatures rather than to the original

^{*}Corresponding author.

Email addresses: coatespt@gmail.com (P. Coates); frank.breitinger@unil.ch (F. Breitinger)

ORCID(s): 0000-0001-5261-4600 (F. Breitinger)

¹It has been shown that LD cannot be calculated in better than quadratic time for the worst case, but it is possible to do better in special cases, for instance, for strings that are known to be similar. For simplicity, we ignore the special cases in this article.

texts. Note, as the heuristic described herein itself uses LD, it suffers from the same quadratic limitation. However, signatures can be typically many hundreds of times shorter than the original texts and therefore the algorithm produces estimates faster. For example: A size reduction factor of C results in a speedup of approximately C^2 . Viewed the other way around, if a given computation time is deemed acceptably fast for files of up to size X, the heuristic extends the tolerable size to CX (where C is a parameter and can be adjusted by a user).

Abbreviations The following abbreviations are used

- LD Levenshtein Distance
- eLD estimated Levenshtein Distance
- LCA Lossy Compression Algorithm

For simplicity, we will use LD and eLD in the text as well as during the algorithmic description. Thus, LD(A, B)means generating the Levenshtein distance between two documents A and B. Note that we differentiate between 'signature' and 'digest': a digest is the return value of our compression algorithm (similar to a hash value) where the signature includes the 'digest' as well as other information (e.g., filename and file size).

Structure The remainder of the paper is organized as follows: first, we describe our algorithm in Sec. 2, which is the core contribution of this article, followed by a brief summary of the reference implementations in Sec. 3. Sec. 4.5 provides an assessment of our algorithm as well as a discussion. In Sec. 5 we discuss some forensic applications from a high-level perspective and introduce the significance score allowing to filter for related documents. The last two sections are the Related work and the Conclusions.

2. The Algorithm

We estimate (approximate) the LD of two documents by (1) compressing each document into a signature using a Lossy Compression Algorithm (LCA), (2) applying the LD algorithm to the compressed signatures, and (3) scaling the result back by the compression rate. As LD is well-defined, this chapter discusses compression algorithm in Sec. 2.1, the layout of the *eLD* signature in Sec. 2.2 and the scaling / estimation of the LD in Sec. 2.3.

2.1. Compression Algorithms

As outlined in the subsequent paragraph, conventional compression algorithms such as LZ77 or BZIP2 are not suited to this purpose (Sec. 2.1.1). Given these limitations, we present properties necessary for our lossy compression algorithm in Sec. 2.1.2 followed by the algorithmic description in Sec. 2.1.3.

2.1.1. Conventional Compression

One reason conventional compression algorithms are not suitable is that they perform 'lossless' compression, which greatly limits their compression rate. Lossless compression is inherently limited by the amount of information present in the data to be compressed. In information theory terms, English text contains about 1.3 bits of information per 8bit character, which makes typical text about 0.1625 information and 0.8375 'air' that can in principle be squeezed out without rendering the original irretrievable. This limits text compression to a factor of approximately 6:1. A second reason is that we want the compressed version to be something like the thumbnail version of a photograph. This means that despite being many times smaller (in size) than the original, a signature must preserve some kind of recognizable similarity to the original if direct comparison of signatures is to provide a meaningful result. Ordinary compression algorithms systematically eliminate this kind of local similarity because maximum compression means that the result looks uniformly random.

2.1.2. Properties for a lossy compression algorithm

When estimating the LD, we rarely if ever are concerned with rates as low as the highest rates conventional compression can achieve. A compression rate of 25:1 would be quite small for an estimation application; compression factors from 100 to the low thousands are more usual for text and even larger compression factors may be used for binary data such as videos.

For our approach, the compression function requires the following properties (note: the term (lossy compression) *digest* is used to describe the output of our compression function):

- Compression: The digest must be much smaller than the original input where 'much' is a factor of 100 to the low thousands.
- Determinism: The digest must be identical for identical inputs.
- Runtime efficiency: Ideally, the compression algorithm runs in linear time with respect to the length of the input. At least, it has to be fast.
- Concatenation: The result of concatenating the digests of two strings shall be identical to concatenating the two strings first and then calculating the digest, i.e., dig(A) + dig(B) = dig(A + B).

Our proposed algorithm does not fully satisfy the fourth property and therefore we settle for "good enough." That is, if we concatenate strings first and then generate a digests (dig(A + B)), the difference from the concatenated digests (dig(A) + dig(B)) must be limited to differences arising from a bounded number of characters on either side of the point of concatenation of the originals. As we will elaborate on below, the number of characters around the point where strings abut that can affect the output is determined by one of the two key compression parameters.

2.1.3. Lossy Compression Algorithm (LCA)

Our LCA uses a rolling window of size-N (neighborhood) that slides through the text (character by character)

and generates a pseudo-random value (hash) at each position (if the document has L characters, L - N + 1 neighborhoods are processed). Depending on a neighborhood's hash, our LCA either (a) does nothing, which is by far the most common decision or (b) uses the hashed value to select a single character from a fixed set 'replacing' the neighborhood in the digest. This keyhole view of the data ensures that any substring of the input can have only a localized effect on the output. The exact procedure is as follows:

LCA description Let *S* be the input of length *L*, and let S_p denote the current position *p* in *S* with $0 \le p \le L - N + 1$. $H_N(S_p)$ denotes the hash function on a substring of length *N* starting at *p* in *S* (details about *H* are discussed in Sec. 2.1.4). Let *ALPHABET* be an array of unique characters (e.g., [a...z,A...Z,0...9]). Then our LCA works as follows (a visual description of the implementation is provided in Fig. 1):

- 1. dig (lossy compression digest) denotes a string² containing the digest, which is initially empty. The compression rate *C* is an integer and discussed later.
- 2. If $p + N \ge L$, print *dig* and quit (*p* is initially zero).
- 3. Compute $H_N(S_p)$ and store the result in T.
- 4. If $(T \mod C) == 0^3$,
 - (a) generate tmp = T mod len(ALPHABET), and
 (b) append ALPHABET[tmp] to dig.
- 5. Set p = p + 1 and return to step 2.

The three required parameters (C, N, ALPHABET) impact the algorithm as follows:

- The nominal compression rate *C*: The larger *C*, the less likely the if-statement (step 4) is triggered and thus the smaller the digest will be. The choice of *C* is subject to the following limitations: if *C* is too large, too much information is lost and the digest provides little value (or may even be an empty string); if *C* is too small, digests will be bulky requiring more storage and comparing digests using LD will be slow. The optimal choice depends on the use case.
- The neighborhood size N: Compression operates on size-N substrings of the input, one neighborhood starting with each successive character. The larger N, the more sensitive the heuristic will be to small differences because each character is part of the N distinct neighborhoods (except within N characters of the beginning or end of the input). For instance, let us assume an input with L = 1000 and N = 500. Performing two changes, e.g., at positions $p = \{300, 700\}$, will likely change the resulting digest completely. On the other hand, if N is too small, the lack of diversity in the neighborhood substrings results in an uneven distribution of output characters. Common values for N range from 11 to 21.

²String was chosen for simplicity and readability; any other format such as binary, hex or base64 would be possible and only requires minor changes. ³While we chose 0, any other arbitrarily chosen value between 0 and C - 1 will work.

• The ALPHABET describes the characters used to build the digests, e.g., printable ASCII characters except newline, tab or space, and are chosen to be conveniently readable for the user. The larger the ALPHABET, the less likely two different T's are mapped to the same character. Note, the length of ALPHABET has to be mutually prime to C, (for simplicity one may use an ALPHABET of length prime). If they are not mutually prime, different strings will be mapped to the same output character more likely. ALPHABET cannot include ',' (comma) as it is used as a separator in the final signature.

Given the ALPHABET = [a...z, A...Z, 0...9], a digest generated by the LCA may look like: AeVVgCAe2aZUpa6dnEkK...

Figure 1: A simple visualization of our compression algorithm; at p = 3 the sequence LLO_WOR is compressed to the letter 'a' (hash values are randomly chosen and do not match the actual implementation).

 $p = 0; sig = "; N = 7; C = 101; ALPHABET = \{A, C, a, b, 1, 2, +\}$



2.1.4. Selection of Neighborhood Hash Function

As every *N*-size neighborhood is hashed resulting in L - N + 1 calls of the hash function *H*, its runtime efficiency is crucial. Naturally, a rolling hash function (e.g., Rabin Karb) or a non-cryptographic hash function (e.g., djb2) seem to be a good choice as they are performant. To analyze the impact of *H*, several different hash functions were tested and the results are summarized in Sec. 4.1. Our final implementation uses Rabin Karb fingerprint but the source code allows to easily replace this function. Note, Rabin Karb is a non-cryptographic hash function and does not have crypto properties such as preimage resistance or collision resistance. However, the lossy compression ensures that the preimage cannot be recovered but it is possible to generate an input that produces a given digest.

2.2. eLD signature

The final *eLD* signature comprises of a *header* as well as the LCA *digest* (in the following we will use the term digest instead of LCA digest). The header is necessary, as only a comparison of signatures with identical parameters is possible. Furthermore, additional information is required to do the scaling as well as assessing the quality of the potential

match. The final *eLD* signature includes the following values separated with a comma (* indicates mandatory values):

- The path/filename of the original file.
- The length of the original file*.
- The C value used to generate the signature.
- The *N* value used to generate the signature.
- The length of the LCA digest
- The LCA digest*.

Values that are not essential can be suppressed in the output, e.g., if C and N are immutable they are not required in the signature. A sample output is given in Fig. 2

Figure 2: Sample eLD signature output including all six values; the signature has been shortened for better readability; comment line (#) was added manually for readability and is usually not printed.

```
$ go run *.go -in test-data/doc.txt
#filename, fileLength, C, N, digestLength, digest
test-data/doc.txt,13680,101,20,152,7n(n[(9&dU7;aRZaPgSGWzoFCC_r
[FUB;A]v?dGrL.bQ2!3GCT)V0r>XWpPNmQ6>#8YhTN]h-X598u+qw1Q1K:+&CVI
gThCK//i2 ...
```

2.3. Estimating the LD

To estimate the LD of two documents given their *eLD* signatures, we first compute the LD between the digests and then perform a scaling where ideally $LD(A, B) \approx eLD(A, B)$. Before explaining the estimation, it is essential to realize that there are two scenarios:

- 1. The documents have roughly the same size and therefore the generated digests shall also have approximately the same size.
- 2. The documents have a large difference in size (e.g., ratio 1:5) which also results in a large difference of the digests length.

In addition, we have to consider that two unrelated documents still have a LD shorter than their length. Meaning that, if A and B are two texts with 1000 characters each, then the maximum LD(A, B) would be 1000 but usually it is less as there is an overlap by chance. This phenomenon will be called *expected overlap* R. As a consequence, it requires more than only multiplying the result by the compression ratio.

Expected Overlap Ratio R The expected overlap ratio R is mostly influenced by the underlying alphabet and the letter frequency. For instance, the alphabet in an English book is almost all 100 printable ASCII characters⁴ but the distribution and order is not random, e.g., there will be spaces

```
<sup>4</sup>Python $ len(string.printable)
```

between words, letter 'e' is most frequently used, 'the' is a frequent word, etc. To obtain R, we programmed a python script⁵ that, given an alphabet, creates two test-strings (TS) of length 30'000 using randomly chosen alphabet elements and calculates the LD. The three alphabets are:

- random_chars: This alphabet consists of 83 elements
 string.ascii_letters, string.digits and
 ()[]+#_-!?%<>@.:;&/{}.
- **eng_chars:** For this alphabet we parsed a book, removed newlines and generated one large string with 686'250 elements. As we randomly select one element at a time, this reflects the frequency of letters in English texts, i.e., the *TS* has many spaces and 'e' but only a few '!' or 'X'.
- **eng_words:** Similar to before, we parsed a book but converted it into a list containing words instead of one large string. Thus, the alphabet is a list of 97'591 words. When creating *TS*, we added a space after each word.

Example TS for each alphabet are provided in Fig. 3. The estimated overlap ratio R is then calculated as follows:

$$R_{alphabet} = 1 - \frac{LD(TS_1, TS_2)}{len(TS_1)} \tag{1}$$

where $len(TS_1) = len(TS_2) = 30'000$. For each $R_{alphabet}$, we performed 10 runs and averaged the ratio resulting in: $R_{random} = 0.0417$, $R_{chars} = 0.1593$ and $R_{words} = 0.1902$.

eLD scaling Let *A* and *B* be two documents with $len(A) \ge len(B)$ (we expect the same for the digests, i.e., $len(dig_A) \ge len(dig_B)$). Then the estimated LD is calculated as follows (for readability | | replaces the *len* function):

$$\begin{aligned} \text{digDiff} &= |dig_A| - |dig_B| \\ \text{effectiveC} &= (|A| + |B|)/(|dig_A| + |dig_B|) \\ \text{digLD} &= LD(dig_A, dig_B) \\ \text{scaledDigLD} &= (\text{ digLD} - \text{ digDiff}) * \text{ effectiveC}/(1 + R_{words}) \end{aligned}$$

$$\begin{aligned} \text{fileLengthDiff} &= |A| - |B| \\ eLD &= \text{round}(\text{ scaledDigLD} + \text{ fileLengthDiff}) \end{aligned}$$

That means: eLD consists of the difference in document length fileLengthDiff as for every missing character, one LD operation is neede (the document lengths (|A|, |B|) are stored as part of the signatures). The second part of the estimation, scaledDigLD, which has to be added to fileLengthDiff, focuses on the portion of the digests that do *not* match, i.e., $dig_A = 12ABCD$ and $dig_B = 12FE$, this focuses on the *FE* (digLD - digDiff); the 12-part can be ignored as the LD=0. The effective compression rate effectiveC is used for scaling. Lastly, we know that there is an expected overlap which we factor in by multiplying $(1 + R_{words})$.

⁵getR.py in the helper-directory.

Figure 3: Shortened examples of the test-strings used to calculate R based on the different alphabets.

random_chars: 9SRCv8A{CK08h{Ly0u7Nu)gw0(0E(0e/!AMm#ss-IGu&CIAoB@lCbwF#Rn:/(Br5GC.6U??iMe.]]t{/VQcd#jW0sTpS?+Xi8GK00uNVf(6qX50f[2TA&z##X1. eng_chars: stwbeulahf yafyes tagokku seuca niuaasdnnicollneptdCannags t gahoni hiyneet ntdtngn eio;awgptmpao RrD eu1lc;ccdsLnPtrs7r il eng_words: are Penn still Logan? the John of compensation Nicola sins Wilcox,along booklet on 9, Louis for spur trees, Caleb of in volume

Sample calculation Let us assume two documents (700 and 500 bytes) that have the first 260 bytes in common. They have the following signatures (C and N are identical thus signatures can be compared):

#filename,fileLength,C,N,digLength,digest
docA,700,51,20,15,AABBCFF00192192
docB,500,51,20,10,AABBCCDDEE

We will now follow Eq. 2 and calculate the eLD:

- 1. digDiff = 15 10 = 5 (to get the length of the digests we can either count or use the digLength-column
- 2. effective C = (700 + 500) / (15 + 10) = 48 (we had a compression rate of C = 51 but the actual compression is likely different so we compute the actual compression rate).
- 3. digLD = (AABBCFF00192192, AABBCCDDEE) = 10 (replace CDDEE with FF001 and add 92192).
- 4. scaledDigLD = (10 5) * 48 = 240 (the five characters CDDEE corresponds to half the signature and should roughly cause an LD of 250); however docA and docB have likely some overlap in this region wherefore we reduce 250 / (1 + 0.19) = 201.68
- 5. fileLengthDiff = 700 500 = 200
- 6. eLD = round(200 + 201.68) = 402

If we compare eLD to our starting situation (260 bytes in common), we know that the LD must be between $200 \le LD(docA, docB) \le 440$ where 200 is impossible as this would mean they have the first 500 bytes in common which is a contradiction.

3. Reference Implementation(s)

There are three reference implementations of the algorithm available - Java 11, Python 3.9 and GoLang 1.17 (thus all are platform independent). As they were developed and modified in parallel by different developers, they all work slightly different but were used to validate results. In the following we will focus on the GoLang implementation; details about the Java implementation can be found in the A; the Python implementation turned out to be significantly slower and will not be discussed further. All implementations are open-source and available online⁶.

The zip file contains our go implementation in the eLD folder, some test data, as well as an overview of the test

results. Options can be found in eLD/config/eLDconfig.go To generate signatures for the files-folder run

\$ go run *.go -in /path/to/dir > signatures.csv

and to compare all signatures within a signature file run

\$ go run *.go -source signatures.csv

The help menu laying out the other options can be access using the -h option.

Remark: As it is an early-stage prototype, error handling and logging is currently limited. It may also be possible to optimize code and make the implementations more efficient.

4. Evaluation

To assess our implementation, we ran experiments looking into Signature generation runtime efficiency (Sec. 4.1), Signature comparison runtime efficiency (Sec. 4.2), Estimated LD vs. LD (unrelated/loosely related files) (Sec. 4.3), and Estimated LD vs. LD (related files) (Sec. 4.4). The last section is a Discussion of the evaluation (Sec. 4.5). All experiments were conducted using the GoLang reference implementation summarized in the previous section.

Device Test results were obtained on a MacBook Pro running BigSur version 11.06. The system used for testing has an 8-Core Intel Core i9 @ 2.3 GHz with 64 GB Memory and an Apple SSD AP1024N 1TB drive. During all tests, the system was in use meaning that other processes like Browser, Email client, etc. were running.

Dataset As our implementation is geared towards text files, we utilize the Gutenberg project and downloaded 18,216 files ranging from 10 KB to 12 MB totaling in approximately 7.2 GB. Subsets were created as needed and are explained in the corresponding section.

4.1. Signature generation runtime efficiency

The efficiency of generating signatures is dominated by the speed of the hashing algorithm as it has to process all neighborhoods. To test the performance of different hashing algorithms, we created four versions of our tool utilizing Rabin Karb, djb2, FNV and MD5. For each setting (N ={7, 14, 21}), we did three runs and averaged times. The time command was used for measurement and we provide the user+sys time which denotes the actual CPU time of the process⁷:

\$ time go run *.go -in ../allfiles/ > res.db

⁶The GoLang version can be downloaded here: https: //www.fbreitinger.de/?page_id=218; the Java implementation is available here: https://coatespt.github.io/Fast-Levenshtein-Distance/; for Python please contact the authors as it is not maintained.

⁷\$ go run will compile and run the program; compiling first using \$ go build and then running the application would be slightly faster.

Table 1

Average runtime for different hashing algorithms using three different N and a constant C = 301.

Ν	7	14	21
Rabin Karb	42s	43s	42s
djb2	1m 55s	2m 35s	3m 18s
FNV*	2m 13s	2m 39s	3m 21s
MD5*	19m 18s	19m 41s	19m 39s

*Tests for FNV and MD5 utilized GoLang libraries and required type-casts which may have impacted runtime.

Table 2

Efficiency for comparing the sample signatures for various C (N = 11); csv size is the file size of the signature file used for comparison.

С	duration	csv size
51	2m 12s	342 KB
101	34s	175 KB
301	4.2s	59 KB
501	2.0s	37 KB
1001	0.9s	19 KB

The results are shown in Table 1; as a Benchmark we also processed the dataset with shasum which required 14s.

As expected, N does have no impact on Rabin Karb fingerprint but on other algorithms that process sequences and not char-by-char. It also has no impact on MD5 as long as N is smaller than the MD5 block size of 512 bit. A constant C was utilized since it has little/no impact on the runtime: Cinfluences the digest length (we confirmed our assumption by running some tests using different C). Lastly, we also ran some tests with the Java as well as Python implementation but both were significantly slower and will not be discussed here, e.g., Java using the string.hashcode() required usually around 5 minutes.

4.2. Signature comparison runtime efficiency

While *C* has no impact on the signature generation, it influences the signature comparisons. Remember, the larger *C*, the smaller the final digest and the digest comparison uses LD with $O(L^2)$. For this test, we randomly selected 40 files (between 1.2 MB and 22 KB) from our dataset with a total of 17.4 MB and created a signature file. We then measured the time for an all-against-all comparison (i.e., 780 comparisons total) with various *C*'s. Times and signature file sizes are summarized in Table 2.

4.3. Estimated LD vs. LD (unrelated/loosely related files)

This section compares our estimated LD vs. the original LD with respect to efficiency and accuracy for unrelated/loosely related files (as all documents originate from the Gutenberg project, they share common header and footer information which we did not remove for this test). We created another subset using 20 random files from the original dataset where each document was required to be between 20 KB and 40 KB ensuring that LD completes in a reasonable time. In total, the 20 files had 627 KB. An overview of the experiment results is given Table 3 where the LDcolumn shows the original LD followed by various C's. Our test ignored self-comparisons, i.e., docA is only compared against other documents and not itself, resulting in a total of 190 comparisons.

To determine the avg error (row 2), we summed all absolute values of our estimate LD minus the LD and divided it by 190. In addition, we provide the min error, max error as well as the standard deviation. This raw error does not consider the ratio of the LD to the file length which is necessary to interpret the meaning of the results. For instance, let us assume pairs documents (size 32'000) having LD 1, 1000 and 15'000 and their eLD are 2, 2000 and 30'000. In all cases the relative error $(\frac{eLD-LD}{eLD})$ is 50% but the impact is different: LD 1 / eLD 2 both indicate a strong correlation between the files while LD 15'000 indicates correlation but eLD 30'000 indicates unrelated documents. Therefore, we also put the error rates in perspective (we use the error rate as defined by MacKenzie & Soukoreff (2002, Eq. 1) but we then deduct the rates to obtain a single value):

$$ER = abs\left(\frac{LD(A, B)}{max(|A|, |B|)} - \frac{eLD(A, B)}{max(|A|, |B|)}\right).$$
 (3)

Consequently, if LD and eLD have similar error rates (which is desired), *ER* is low. For our example, this would result in the following error rates: $(abs(\frac{1}{32'000} - \frac{2}{32'000}) =) 0.00, 0.03$ and 0.47, respectively.

As shown by our results, our algorithm is significantly faster and results have a low 'ER'. The larger C the less accurate is the algorithm which may also be due to the rather small sample files (e.g., the average digest length is 158). Detailed results are also included in the downloadable zip file.

4.4. Estimated LD vs. LD (related files)

For this second test we reused the 20-random-file dataset from the previous section, created a modified copy of the file and then compared the original version against the modified version. A summary of the outcome is provided in Table 4.

The proposed estimation works well for deletions as long as they are not spread out across the complete document (compare L1-12 vs. L13). Similarity, random insertions (L14+15) have a similar impact on the performance although if there are only a few insertions and a higher compression rate, this may be equalized. Swapping (L16+17) is also handled well where it is eye-catching that the estimated distances are lower than the original distance. Clearly, our algorithm has difficulties if there are many minor changes throughout the document (L18-20). However, these may be reduced by some additional pre-processing such as trimming / stripping or making it case insensitive.

Table 3

LD-column shows the time for a traditional LD implementation from GoLangPrograms.com; the remaining columns outline eLD results for various compression rates C where N = 11 and EXPECTED_OVERLAP_TEXT=0.19 remained unchanged. Columns provide the absolute value and the percentage value (rounded to one digit).

	LD	<i>C</i> = 11	<i>C</i> = 21	<i>C</i> = 51	<i>C</i> = 101	<i>C</i> = 201
Duration	9min 09s	3.5s	1.2s	0.6s	0.4s	0.3s
Avg Error (abs/%)	_	1223 / 6.5%	1166 / 6.4%	1519 / 9.0%	1578 / 9.0%	1650 / 9.4%
Min Error (abs/%)	_	16 / 0.1%	10 / 0.1%	10 / 0.1%	3 / 0.0%	16 / 0.1%
Max Error (abs/%)	_	3498 / 23.2%	3484 / 23.1%	5177 / 34.3%	4187 / 40.7%	6942 / 35.6%
Std Dev. (abs/%)	_	822 / 3.6%	767 / 3.8%	1114 / 6.4%	1075 / 7.0%	1424 / 7.3%
Error rate (avg/std dev)		0.03 / 0.02	0.03 / 0.02	0.04 / 0.03	0.04 / 0.02	0.05 / 0.04

Table 4

Results of comparing one document against a modified version of itself manipulated according to the operationcolumn(modifications were done manually and randomly). FS and LD represent the original file size and the Levenshtein distance. The remaining columns show the *eLD* for various *C* (again N = 11 and EXPECTED_OVERLAP_TEXT=0.19 remained unchanged).

#			Operation	org. FS	LD	C=11	C=21	C=51	C=101
1	10348.t×t	10348mod.txt	Deleted 10 lines	24226	663	718	798	709	663
2	ltplt10.txt	ltplt10mod.txt	Deleted 10 lines	22869	634	672	652	717	719
3	7563.t×t	7563mod.txt	Deleted 51 lines in the beginning	23414	1024	1024	1024	1024	1024
4	7565.t×t	7565mod.txt	Deleted 101 lines in the middle	25983	2234	2234	2234	2234	2330
5	7906.t×t	7906mod.txt	Deleted 90 lines (some beginning; some end)	21586	3329	3329	3329	3329	3329
6	10630.txt	10630mod.txt	Deleted 7 chunks	32185	7085	7095	7103	7086	7086
7	17282.txt	17282mod.txt	Deleted 3 large chunks	25534	7436	7436	7436	7436	7436
8	18232.txt	18232mod.txt	Deleted 15 times 3 lines	31635	2964	3117	3022	3034	3165
9	8528.t×t	8528mod.txt	Deleted approx first half	33625	15811	15811	15811	15811	15811
10	11592.txt	11592mod.txt	Deleted 10 paragraphs	36253	4224	4262	4259	4224	4224
11	16637.t×t	16637mod.txt	Deleted 5 paragraphs at the beginning	39864	1731	1740	1731	1731	1731
12	wlett10.t×t	wlett10mod.txt	Deleted big middle chunk	32346	8968	8968	8968	8968	8968
13	lf17w10.txt	lf17w10mod.txt	Deleted all (450) 'the'	38902	1215	5409	5749	7128	7156
14	haw4610.t×t	haw4610mod.t×t	Inserted 10 random 'A'	31340	10	93	113	96	10
15	haw7810.txt	haw7810mod.t×t	Inserted 40 random 'A'	36281	40	531	562	819	390
16	lf20w10.txt	lf20w10mod.txt	Swapped 6 paragraphs around	37765	3218	2831	2969	2897	1559
17	14814.t×t	14814mod.txt	Swapped first and second half (approx)	28050	17517	14128	14454	15134	17721
18	12337.txt	12337mod.txt	All 'b' (338) replaced with 'B'	31241	338	4204	3172	3004	3248
19	17195.t×t	17195mod.txt	All 'e' (3080) replaced with 'E'	35158	3080	21629	21605	22510	20106
20	13322.t×t	13322mod.txt	All spaces (6572) replaced with double-space	38838	6572	34083	35883	31128	32218

4.5. Discussion of the evaluation

The initial assessment primarily focused on the *C*-parameter and runtime; more testing is required to assess the influence of other parameters such as N or R (EXPECTED_OVERLAP_TEXT). However, this preliminary evaluation shows promising results and probably can be optimized by adjusting other parameters. Currently, our approach performs well as long as the files are of sufficient size to generate adequate signatures – for small files one can fall back to the original LD.

Estimating signatures has both advantages and limitations: Most obviously, if the application at hand requires a precise LD value, the proposed heuristic is not useful. For instance, the heuristic will often estimate an LD of zero for texts that in fact vary slightly. It will never, however, do the opposite, i.e., estimate that two documents have an LD greater than zero when in fact they are identical.

Algorithmic performance depends upon specific signaturegeneration parameters as well as the text itself but estimates are typically within a few percent of the LD and tend to be more accurate with relatively similar documents, which is often the most useful category. Meaning that there are more efficient ways to identify very/completely different documents such as approximate matching (see Sec. 6). The level of accuracy is sufficient for tasks like detecting nearduplication and partial duplication, estimating how much two documents vary, filtering through large numbers of documents to look for a near-match to some substantial block of text, and detecting documents that are related in more complex ways.

The effectiveness also depends upon variability in the input; there can be inputs that produce pathological results. For instance, the same short (relative to N) character sequence endlessly repeated may result in a 'zero' length output or in an excessive output because one or more of the N-length substrings that are repetitively encountered causes an output character to be generated. While repetitive input of any great length would be unusual in natural language text, it can show up in various kinds of machine-generated text such as CSV files, log files, experimental data, etc. From a forensic perspective, signatures for such data are unlikely to be useful (i.e., one would probably not compare them against other documents). If excessively short or empty, they would fail to match even nearly identical files, and if they were excessively long, they would result in excessive processing time. For realistic *N* and *C*, neither of these results will be close to the expected output size of $\frac{1}{C}$ of the input size and thus can be easily identified.

The signature of a fragment of a document is identical with the corresponding portion of the signature of a larger document in which it is embedded (except possibly near the points where they abut.) Therefore, signatures of related documents differ in locations that are in approximately the same relative positions as the corresponding differences in the originals. This means, not only can signatures be used to estimate how much two documents differ, they can also be used to infer useful information about where and how they differ. For example, they can indicate that two documents differ mostly at the beginning and end but differ little elsewhere, or that many differences are sprinkled throughout, or that one document appears to be appended to the other.

One of the most useful aspects of the signature-based approach is that neither the target documents nor the search documents need be in hand at the time an estimate is made. This has the obvious benefit of allowing the cost of signature generation to be amortized over many uses, but it also means that it is possible to search a collection of signatures for similar documents and assess their degree of textual similarity of candidates without having direct access to any of the documents involved.

5. Digital forensic application

While many files on computers are binary, there is still plenty of text based information that must be assessed during an investigation which is the application of our prototype implementation. Having an efficient implementation of the Levenshtein distance allows the algorithm not only to be applied on short sequences but also larger files such as complete documents or source code files to cluster files, detect plagiarism or utilize it for black-/whitelisting. In the following we outline four scenarios where our algorithm can contribute:

- **Source code analysis:** During an investigation it may be necessary to analyze and compare large software projects in order to identify if there are similarities in the source code. Identified similarities could mean that source code was copied (stolen) and reused. On the other hand, this also allows clustering similar files together. An example here would be the analysis of malicious code where an investigator received access to the source code of thousands of malware samples and wants to cluster similar samples together reducing the time of the manual analysis (e.g., we assume that similar malware applications have a similar behavior and only one sample per family, i.e., cluster, needs to be analyzed).
- **Fuzzy string matching:** Instead of looking for exact matches, it may be necessary to identify strings that roughly

match a given string or are a substring of a given string. For instance, let us assume an examiner is given a Spam-Email and has to proof that the suspect is responsible for an Email spam campaign wherefore s/he decides to search in all documents for the given Email. It is likely that the Email has been modified throughout the campaign to avoid detection by spam filters.

- Similarities between cybercrimes: Besides processing evidence, our approach may also be useful to identify similarities between cybercrimes as discussed by Bollé & Casey (2018). In their work, the authors propose "new approach to finding linkages and repetitions across cases in a cyber-investigation context using near similarity calculation of distinctive digital traces". In their study, the authors use the Levenshtein distance on email addresses. Using a more performant algorithm allows to not only analyze email addresses but to cross compare additional (more comprehensive) traces.
- **Anonymous searches:** As little useful information about the contents of documents can be divined from the signatures, one can have a service that allows an outside party to check whether partial matches to a document or other files are present without the service having access to the library to be searched or to the text to be searched for. The client does not have to provide the query and the service only has access to the signatures, not the originals.

Note, while the current prototype has been developed for text files only, we plan to expand it and see if it can be used for binary files as well.

5.1. Significance score

The estimated LD is computed for each pair of signatures, but in most cases, an estimate is not informative without context. To assess the importance of a comparison and filter out pairs that are not in a relationship (which is hard to do from the estimated LD score), we introduce the *significance* δ^8 . This value puts the *eLD* in relation to the digest lengths and allows an examiner to set a threshold *T* ($0.0 \le T \le 1.0$) so that only matches with $\delta \ge T$ are shown. Note, T = 0.0 will return all matches and T = 1.0 will only return 'exact' matches with eLD = 0.

As discussed in Sec. 2.3, there are two different scenarios where digests have a similar length or may vary widely. Ideally, δ considers both cases. That is, *delta* shall be high if a smaller digest is (almost) included in a larger sequence as well as it shall be large if there is a low LD between two similar sized digests. Currently, the significance score δ for two digests is computed as follows (we assume $len(dig_A) \ge$ $len(dig_B)$):

⁸This is work in progress and as will be seen the current score has some weaknesses requiring a workaround.

Table 5Sample δ values for various digest length and LDs.

	$len(dig_A)$	$len(dig_B)$	$LD(dig_A, dig_B)$	delta
1	700	700	0	1.000
2	700	700	10	0.986
3	700	350	400	0.857
4	700	100	600	1.000
5	700	700	600	0.143
6	700	350	650	0.143
7	700	100	696	0.040
8	700	200	700	0.000
9	70'000	700	70'000	0.000
10	70'000	700	69'650	0.500
11	70'000	700	69'300	1.000

 $ld = LD(dig_A, dig_B)$ Levenshtein distance of digests $\delta = \frac{len(dig_A) - ld}{len(dig_B)}$

However, this calculation has the weakness that it does not work for very large differences in digest size (e.g., ratios of 1 to 20). Table 5 provides some sample *delta* values for given digests length and overlap where the top third contains 'good' matches, the middle third contains 'poor matches' and the last third demonstrates where δ fails. The problem is, if the difference is too large, dig_A can be converted to dig_B just by deleting characters resulting in $LD(dig_A, dig_B) = len(dig_A) - len(dig_B)$.

There are several possible solutions: (i) define a multiplier so that both digest lengths have to be in this range from each other (e.g., factor 10); (ii) adjust the equation to consider this behavior (e.g., using R_{Random}); or (iii) require that sequences share a common substring of a certain length (e.g., if dig_{R} is a substring of dig_{A} , they are likely related).

5.2. Assessment significance score (δ)

To assess δ 's raison d'être, we performed two test runs: (1) do an all-against-all comparison of 72 unrelated files and (2) finding 10 originals in the full dataset.

Test1 To create the 72-unrelated-files dataset, we used the large set, randomly selected 72 files and shrank them to 30 KB by removing text from the beginning and the end (i.e., deleting potential header/footer information). We then calculated the LD between all documents and sorted the 2556-pairs according to their LD ranging from to 22'255 to 25'762. Lastly, we selected some pairs with a low distance and analyzed them manually to ensure they are different. While further testing is required, we believe these still significant overlaps originate from the many leading spaces in the documents (see example in Fig. 4)⁹.

Next, we ran our algorithm (C = 51, N = 11) and computed the significance score δ for each pair. The **Figure 4:** Extract from document 13649.txt showing the leading spaces which cause relatively low LD's.

Of the Hills of the Chankly Bore,
Then, through the vast and gloomy dark
There moves what seems a fiery spark,
A lonely spark with silvery rays
Piercing the coal-black night,
A Meteor strange and bright:
Hither and thither the vision strays,
A single lurid light.

Table 6

Summary of the δ values for unrelated files. Right-half of the table shows accumulated results where all matches where almost all matches where lower than 0.12.

Min	Max	Avg	≤ 0.03	≤ 0.06	≤ 0.09	≤ 0.12	≤ 0.15
0.025	0.122	0.058	19	1470	2479	2555	2556

Table 7

Distribution of the δ score from comparing 10 documents against the dataset (18'196).

$\begin{array}{l} \delta \leq 0.1 \\ 93'359 \end{array}$	$\begin{array}{c} 0.1 < \delta \leq 0.2 \\ 9357 \end{array}$	$0.2 < \delta \le 0.3$ 18'193	$\begin{array}{c} 0.3 < \delta \leq 0.4 \\ 25'743 \end{array}$	$\begin{array}{c} 0.4 < \delta \leq 0.5 \\ 24'298 \end{array}$
$0.5 < \delta \le 0.6$ 10'998	$\begin{array}{c} 0.6 < \delta \leq 0.7 \\ 2 \end{array}$	$\begin{array}{c} 0.7 < \delta \leq 0.8 \\ 0 \end{array}$	$\begin{array}{c} 0.8 < \delta \leq 0.9 \\ 0 \end{array}$	0.9 < δ 10

summarized results are shown in Table 6. As expected, we obtain low scores for all matches where only one match is above 0.12. To conclude, δ works for dissimilar documents that are of same/similar size.

Test 2 Our second test focused on identifying 10 of these unrelated documents in the whole dataset. Therefore, we created two signature files and used the source-destination mode of our application. As the dataset is rather large, we set C = 101 (*N* remained 11). Additionally, we manipulated the application to only consider files that are up to 10 times larger than the source file (if larger, δ is set to zero). The distribution of the δ score is shown in Table 7 where the top rows indicate the δ -score and the bottom rows the absolute number of pairs.

As indicated by the table, the 10 matches were clearly identified with a gap to any other matches. We manually investigated the two matches for $0.6 < \delta \le 0.7$ and did not see any similarity. However, both files were close to 300 KB which is maximum difference in size to be considered and again highlights the weakness of the current δ -formula.

6. Related work

Given the sheer amount of work that has been done in terms of string similarity and approximation, we only included a summary where we decided to split the section into two domains: string similarity / approximate string matching and approximate matching (as defined from the digital forensics community).

 $^{^{9}\}mbox{For one pair, we removed the double-space which increased the LD by about 2200.$

String similarity String similarity metrics have a long history and may different techniques exist. Wu (2020) proposes to divide them in four categories: (1) character-based / edit distance (examples are Hamming Distance, Levenshtein Distance, Damerau-Levenshtein Distance or Jaro Distance); (2) Sequence-based (examples are Longest Common Subsequunce, Longest Common Substring, or Geshtalt Pattern Matching); (3) Token-based (examples are Q-Grams, Overlap Coefficient or Jakkard Similarity); and (4) Phoneticsensitive (examples are Smith-Waterman or Editex). Our work falls into the first category, which is also known as the string-to-string correction problem (Wagner & Fischer, 1974). Many of these techniques have been analyzed in detail, optimized algorithms have been published and improvements have been presented. A thorough analysis of work related to the Levenshtein Distance has been done by Navarro (2001) who compared the performance of different algorithms doing a number of experiments. However, most works focus on improving and not on estimating. For instance, Ukkonen (1985) proposes an improvement that runs in $O(s \cdot min(len(A), len(B)))$ where s is the edit distance of the two strings compared. Thus, the algorithm runs in approximately linear runtime for highly similar strings.

Approximate matching (AM) From a digital forensics perspective, these algorithms are interesting but despite their improvements, too slow for practical usage on large amounts of data. Therefore, the community invented its own AM (a.k.a. fuzzy hashing for similarity hashing) which later has been summarized by the National Institute of Standards and Technologies in SP800-168 with the goal to "to provide a definition and terminology to describe approximate matching in order to promote discussion, research, tool development and tool acquisition" (Breitinger et al., 2014a). The research in this domain evolves around developing novel algorithms (e.g., Kornblum (2006); Oliver, Cheng & Chen (2013); Chang, Ghosh, Sanadhya, Singh & White (2019)), comparing algorithms (e.g., Roussev (2011)) or improving/assessing existing implementations (e.g., Baier & Breitinger (2011)). Discussing all existing works is beyond the scope of this article but a somewhat up to date summary is provided by Harichandran, Breitinger & Baggili (2016) as well as Martín-Pérez, Rodríguez & Breitinger (2021) who published a categorization of AM algorithms including a discussion of various algorithms. A point that is raised in many publications is the similarity score (i.e., the result when comparing two fingerprints) which is only an indicator of similarity but the similarity is not clearly defined, i.e., a similarity score of 50 does not meant that two inputs are 50 percent identical or that 50 paragraphs are identical. While this indicator (two inputs are similar with a high certainty) is sufficient for some use cases, other scenarios may require a more precise result. One example would be checking for plagiarisms as often done by university during online submissions.

Breitinger, Ziroff, Lange & Baier (2014b) published an application named saHash (statistical analysis hashing) which is based on four sub-hash functions, each creating its own sub-hash value. The final, fixed-length hash is then the concatenation of all sub-hashes. While the authors claim that their implementation is almost as fast as SHA-1, it only "enables the detection of 'small' changes between files – up to several hundred Levenshtein operations".

7. Conclusions

The Levenshtein Distance (LD) is a well-known distance metric to describe the dissimilarity between two strings / documents. While it works well for short sequences, it requires a quadratic runtime and thus is impractical for long sequences (e.g., hundreds of kilobytes). In this article we presented a heuristic that allows estimating the LD also for larger sequences. The algorithm works by first compressing a given document into a signature. and then comparing signatures against each other to estimate the LD. Consequently, it is not necessary to possess the original documents to calculate their LD. To assess the quality of our algorithm, we performed several experiments and show that our estimation is significantly faster and, depending on the modifications done to a source element, obtains accurate estimations (e.g., the experiments showed that estimations are very accurate for dissimilar documents of the same size as well as that 'deletions' are handled well.

Future Research While the effectiveness of the heuristic seems clear, there are several areas that need further investigation, especially the following four: (1) As outlined in the article, the calculation of δ is not perfect and we had to introduce a multiplier that documents have to be of similar size (e.g., within factor 10). It requires more experiments to see what mechanism provides the best results from a digital forensic perspective. (2) Documents can have a large LD despite sharing substantial blocks of identical or nearly identical text. One way to detect pairs that are related in this way is to divide large documents into uniform sized blocks and estimate the LD's of all pairwise combinations of blocks from the two originals. This would allow relatively small segments, that are similar, to stand out against a background of unrelated text. (3) File signatures are like thumbnail images, with the differences between a pair of signatures providing a coarse map of the differences between the corresponding originals. A better understanding how to characterize and quantify different kinds of similarity, and understand the limitations of this technique would make this property more useful. (4) All experiments have been on books in English that have been coarsely mutilated, usually with many blocks of lines either being deleted or duplicated. A more systematic survey of the quality of estimates based on files with quantified degrees and kinds of differences would be useful. The behavior on particular types of data would also be useful. This would be of particular interest in applying the algorithm to binary data such as video or compiled computer programs.

References

- Baier, H., & Breitinger, F. (2011). Security aspects of piecewise hashing in computer forensics. In 2011 Sixth International Conference on IT Security Incident Management and IT Forensics (pp. 21–36). IEEE.
- Bollé, T., & Casey, E. (2018). Using computed similarity of distinctive digital traces to evaluate non-obvious links and repetitions in cyberinvestigations. *Digital Investigation*, 24, S2–S9.
- Breitinger, F., Guttman, B., McCarrin, M., Roussev, V., White, D. et al. (2014a). Approximate matching: definition and terminology. *NIST Special Publication*, 800, 10–6028.
- Breitinger, F., Ziroff, G., Lange, S., & Baier, H. (2014b). Similarity hashing based on levenshtein distances. In *IFIP International Conference on Digital Forensics* (pp. 133–147). Springer.
- Chang, D., Ghosh, M., Sanadhya, S. K., Singh, M., & White, D. R. (2019). Fbhash: A new similarity hashing scheme for digital forensics. *Digital Investigation*, 29, S113–S123.
- GoLangPrograms.com (). Golang program for implementation of Levenshtein distance. https://www.golangprograms.com/ golang-program-for-implementation-of-levenshtein-distance.html.
- Harichandran, V. S., Breitinger, F., & Baggili, I. (2016). Bytewise approximate matching: the good, the bad, and the unknown. *Journal of Digital Forensics, Security and Law*, 11, 4.
- Kornblum, J. (2006). Identifying almost identical files using context triggered piecewise hashing. *Digital investigation*, 3, 91–97.
- MacKenzie, I. S., & Soukoreff, R. W. (2002). A character-level error analysis technique for evaluating text entry methods. In *Proceedings* of the second Nordic conference on Human-computer interaction (pp. 243–246).
- Martín-Pérez, M., Rodríguez, R. J., & Breitinger, F. (2021). Bringing order to approximate matching: Classification and attacks on similarity digest algorithms. *Forensic Science International: Digital Investigation*, 36, 301120.
- Navarro, G. (2001). A guided tour to approximate string matching. ACM computing surveys (CSUR), 33, 31–88.
- Oliver, J., Cheng, C., & Chen, Y. (2013). Tlsh–a locality sensitive hash. In 2013 Fourth Cybercrime and Trustworthy Computing Workshop (pp. 7–13). IEEE.
- Roussev, V. (2011). An evaluation of forensic similarity hashes. *digital investigation*, 8, S34–S41.
- Ukkonen, E. (1985). Algorithms for approximate string matching. *Information and control*, 64, 100–118.
- Wagner, R. A., & Fischer, M. J. (1974). The string-to-string correction problem. *Journal of the ACM (JACM)*, 21, 168–173.
- Wu, G. (2020). String Similarity Metrics Edit Distance. https://www. baeldung.com/cs/string-similarity-edit-distance.

A. eLD Java Implementation

The implementation of our algorithm is written entirely in Java 11 and compiled using Maven 3.6.3 and OpenJDK 64-Bit Java version: 11.0.11. The source code is generic Java using no language features not found in any standard imperative computer language. To keep descriptions of performance simple, the current version is single-threaded only. It is open-source and can be downloaded here: https:// coatespt.github.io/Fast-Levenshtein-Distance/. The application supports command line arguments or the usage of a configuration file. To generate signatures run

\$ java -jar Fast-LD.jar -p config_gen.properties

and to compare the files-folder against an existing set of signatures (sigs.csv) run

\$ java -jar Fast-LD.jar -p config_cmp.properties

In these examples, our Fast-LD implementation utilizes configuration files (CLI arguments overwrite the config) which is listed in Fig 5. Parameters *C*, *N*, and *CH* (Alphabet) have been a explained in Sec. 2.1.3; f points to a CSV file whose content is a list of files to be processes (one per line); Id=true indicates comparison mode (Levenshtein Distance Estimation); and ft is the signature database in the CSV format. In order to generate signatures (no comparison), one only has to set Id=false (for better readability, ft can be removed as it is not needed / will be ignored).

Figure 5: Content of the configuration file config_cmp.properties; the alphabet has been shortened for better readability.

c = 251

- n = 11
- f = input.csv
 ch = abcdefghijklnmnoprstuv[...]234567890

ld = true

ft = sigs.csv