

A New High-Performance Approach to Approximate Pattern-Matching for Plagiarism Detection in Blockchain-Based Non-Fungible Tokens (NFTs)

Ciprian Pungilă¹, Darius Galiş¹, Viorel Negru^{1,2}

¹ Faculty of Mathematics and Informatics, Department of Informatics

² ICAM - Environmental Advanced Research Institute

{ciprian.pungila,darius.galis,viorel.negru}@e-uvvt.ro
West University of Timișoara, Romania

Abstract. We are presenting a fast and innovative approach to performing approximate pattern-matching for plagiarism detection, using an NFA-based approach that significantly enhances performance compared to other existing similarity measures. We outline the advantages of our approach in the context of blockchain-based non-fungible tokens (NFTs). We present, formalize, discuss and test our proposed approach in several real-world scenarios and with different similarity measures commonly used in plagiarism detection, and observe significant throughput enhancements throughout the entire spectrum of tests, with little to no compromises on the accuracy of the detection process overall. We conclude that our approach is suitable and adequate to perform approximate pattern-matching for plagiarism detection, and outline research directions for future improvements.

Keywords: plagiarism detection · pattern-matching · approximate · blockchain · NFT · automaton · Aho-Corasick · sliding window · similarity measurement

1 Introduction

Plagiarism is an important issue with respect to protecting intellectual property, a crucial centerpiece in the academic community, as well as in various business aspects. Plagiarism occurs whenever material is copied without the author's permission or approval. With the spread of the Internet, plagiarism becomes a growing concern to authors developing original work. The blockchain industry aims to resolve this problem through the use of non-fungible tokens (NFTs), a market that has grown exponentially in the past few years [1].

Detecting plagiarism could leverage costs significantly for businesses with respect to protecting their intellectual property, and could identify theft of ideas and methodologies commonly used in the academia. When it comes to protecting its rightful author, intellectual property is surely a form of art on its own.

Blockchain-driven architectures with inherent NFT support, aim to protect digital art, in all its forms, through NFT-ready technology. This provides NFT-ready blockchain technology with the means to protect against various forms of intellectual property infringement.

This paper is organized as follows: Section 2 provides a deeper insight into the related work that lies as the ground foundation of this paper; Section 3 outlines the intrinsic details of our own implementation, and discusses various design choices made throughout its inception; Section 4 describes the testing methodology and real-world experiments that we have employed in order to test our hypotheses, and discusses the outcome of the approach as compared to other existing approaches; finally, Section 5 summarizes the work we have employed, and describes our vision for future directions on it.

2 Related work

2.1 Plagiarism detection

Related research. Plagiarism detection is a serious issue of interest to the academia, as well as to businesses in general. Numerous papers and approaches have been proposed to leverage this issue, and improve existing techniques to identify potential scammers. For example, in [2], the authors propose an approach based on tokenization and the computation of a similarity score, making use of natural language processing (NLP) techniques. In [3], authors focus on a hybrid approach: by combining the Jaccard and cosine distances, they recurse to machine learning (ML) and NLP, corroborated with text mining and similarity analysis. Similar to the previous approach, the authors recurse to tokenization. Saini et al [4] had build a plagiarism detection using text mining methods and the cosine distance for computing similarity. Putri et al [5] propose a text-mining-based approach to plagiarism detection using a Karp-Rabin variant of pattern-matching, involving a slightly modified Jaccard similarity computation method and tokenization, to identify the number of common words in the two texts being compared. While the approach is able to identify matches in the two texts being compared, the approach itself is restricted to identifying common words. Approximate pattern-matching techniques in text mining help leverage the gap between partially modified copycats and original work, through a more accurate identification of close-enough wording to cause suspicion of plagiarism.

Plagiarism detection through text mining. There are several classes and types of plagiarism. From a lexical perspective, text mining is a common approach to performing plagiarism detection. From a semantic perspective, as outlined before, there are approaches based on semantics, particularly NLP and ML, that bring the concept of "understanding" of the text contents into the picture, just as well. Our main focus in this paper is the approach based on text mining utilizing the tokenization of input data, in order to perform the mining process and assign a plagiarism degree between two different texts $A = a_1a_2...a_n$ and $B = b_1b_2...b_n$ over an alphabet σ .

Similarity measures for approximate pattern-matching. A few similarity measures are commonly used to perform text mining, in order to find approximate pattern-matches. Of these, the following are the approaches we will focus on in this paper:

a) *The Euclidean distance* [6] is a granular approach to identifying similarities between two sequences of characters, by taking into account the numeric values of their corresponding ASCII codes, and computing a numerical similarity. It is defined, in normalized form, as:

$$d(A, B) = 1 - \frac{\sqrt{\sum_{i=1}^n (a_i - b_i)^2}}{\sqrt{n \times |\sigma|}}$$

b) *The Hamming distance* [7] measures the number of differences, position-wise, between the two texts. It is defined, in normalized form, as:

$$d(A, B) = 1 - \frac{\sum_{i=1}^n (d_i = \begin{cases} 1, & a_i \neq b_i \\ 0, & a_i = b_i \end{cases})}{n}$$

c) *The Levenshtein distance* [8] is a string metric for identifying the number of single-character edits required to transform one sequence of characters into another. Assuming costs c are the various costs for the different operations possible (insertion, delete, substitution - with an implicit cost value of 1 for each of these), it is defined as:

$$d_{i,j} = \min \begin{cases} d_{i-1,j} + c_{delete}(b_i) \\ d_{i,j-1} + c_{insert}(a_j) \\ d_{i-1,j-1} + c_{substitute}(a_j, b_i) \times [a_j \neq b_i] \end{cases}$$

Performance analysis. Not many research papers outline the performance of their approach in real-world experiments, especially when using very large datasets. With the growing demand for fast solutions in almost all fields of technology nowadays, there is no doubt that speed is becoming a major constraining factor with respect to existing implementations. Therefore, one particular emphasis in our approach lies on the acceleration of the approximate matching process overall.

2.2 Blockchain technology

Blockchain technology has come into focus in recent years with the growing public interest in Bitcoin [9]. A blockchain is a ledger used to store data in a sequential manner time-wise, enforcing data immutability using security measures, of which the most common are cryptographic hashes. Data is organized into blocks, with each block having its own cryptographic hash that is dependent on the hash of the previous block. A block is usually comprised of various

operations. Due to its inherent design, alterations of any part of data, in any given block, would determine inconsistency at the hash level for all consecutive blocks, ensuring therefore data non-repudiation.

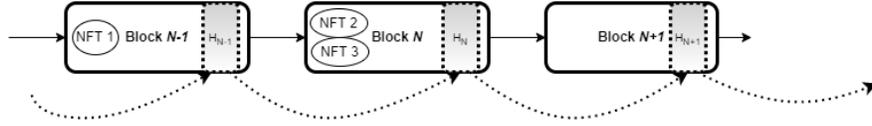


Fig. 1: A blockchain-driven ledger, where block N 's hash (H_N) is cryptographically and computationally dependent on the hash of the previous block, H_{N-1} . Also depicted are various NFTs that are minted throughout blocks.

In recent years, with the growing interest in digital currencies making use of blockchain technology (such as Bitcoin), commonly known today as cryptocurrencies, the NFT market has also grown significantly in just a few months from its inception, outlining interest in use-cases of the technology for protection against copycats. Our focus in this paper is how we can apply our approach for approximate pattern-matching, to NFT technology, in particular to protecting intellectual property stored through the form of an NFT that was already previously minted in the blockchain.

2.3 Challenges of NFT-driven plagiarism detection

As blockchain technology provides nowadays the adequate means to uniquely identify and, therefore, protect intellectual property (in particular, identify plagiarized material), we have to also consider the potential drawbacks of such an approach. For example, given how hashes are able to uniquely identify a given piece of data, they are not able to differentiate between close wording, which is what approximate pattern-matching aims to achieve. With NFTs becoming more popular, a growing concern becomes the need to parse large sets of NFTs (in order to identify potential clashes or inconsistencies with other existing NFTs) whenever new NFTs are being minted. While hash techniques can offer protection against immediate copycats of data, they do not offer protection against slightly altered versions of the same data. For example, the SHA256 [10] hash of the text "The sky is beautiful" differs significantly from that of the slightly altered version of the same text: "The sky's beautiful" (please refer to Table 1). Not only that, but there is no discernible link or connection between the two hashes that are computed for the two texts: by all means and purposes, different hashes mean different texts, with no indication of any semantic relationship between the texts that were hashed.

A similar principle applies to digital data in all its form (e.g. not just text), including digital art stored and valued through NFT technology. For example, a PNG picture without alpha transparency, as a form of digital art, could be

digitally altered with the change of a single bit (e.g. for example, corresponding to the unused alpha transparency channel of any pixel), without any visual changes of any kind to the actual picture itself - but with completely different resulting hashes, NFT-wise, determined by the two different pieces of data. This means that the problem of uniqueness is not handled entirely in the NFT world, but could potentially be solved through an approximate pattern-matching technique, where a similarity comparison would be able to properly identify the high degree of similarity, from a lexical perspective at least, between the two files, and classify the altered file as a plagiarized version of the first one, implicitly denying the minting of the new NFT therefore.

Table 1: Comparison of hashes of an original text, and a slight alteration of the same original text.

Type of text	Text contents	SHA256 hash computed for the text
Original text	The sky is beautiful	d5217a507abdf43516559face8b9f51cecd463fe2f4542a53ea6de3363642a62
Altered text	The sky's beautiful	f452385958a7281aff97bcd2d52f9d43c29be3bcd24481f87d23e6a9b626a2a8

2.4 The Aho-Corasick automaton

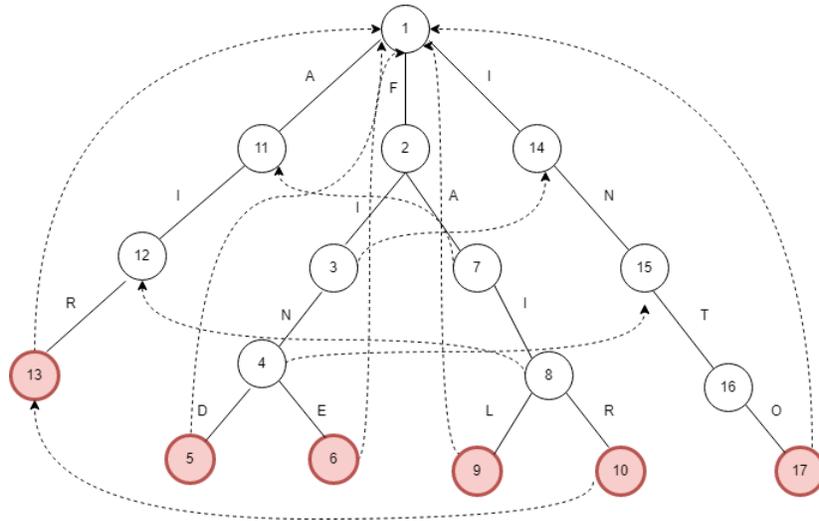


Fig. 2: The Aho-Corasick multiple pattern matching automaton for the set of keywords $\{AND, FIND, FINE, FAIL, FAIR, INTO\}$. Mismatch transitions are depicted through dashed lines. Leaves are depicted as bolded and colour marked.

The deterministic finite automaton (DFA) proposed by Aho-Corasick [11] is a very popular approach to solving the multiple pattern-matching problem. The Aho-Corasick automaton is built from a trie tree, constructed from the set of keywords that need to be looked up in an input data set. The trie tree is then transformed into a deterministic finite automaton (DFA) by computing a "mismatch" transition (failure function) that points to the longest suffix of the word at the current node that is also present in the trie. The mismatch transition is used whenever there is no match of the input set at the current node. Every node will save information about the word associated with the respective position inside the trie (prefix), and the transition towards the next state. For example, for the automaton in Figure 2, node 4 is associated with the word *FIN*, has the suffixes $[IN, N, <empty>]$, and while *IN* is also a word in the trie, pointing to node 15, the mismatch transition of node 4 will therefore point to node 15.

One drawback of the Aho-Corasick automaton is that, while it performs very well in tracking accurate matches of patterns in an input string, it is not suitable to track alterations of those patterns in the input data. Nonetheless, its almost linear runtime speed makes it one of the top choices for exact pattern-matching nowadays, and that was the starting premises for the non-deterministic finite automaton (NFA) approach that we propose in this paper.

3 Implementation

3.1 Our NFA-based approach

Our approach is inspired by the Aho-Corasick [11] finite state machine, however it differs significantly from it as it creates a new NFA state machine, using a sliding window concept at node-level, as well as a locally-applied similarity measurement for the sliding window concept that we introduce. We modify the failure transitions of the NFA, and we are changing the entire heuristics of those transitions so that we can apply inexact pattern-matching. In order to achieve this, we allow the transitions from a certain state to have multiple outcomes when parsing the same input characters. In order to start determining the output of any given transition, we employed the usage of the sliding window concept, accompanied by the computation of every possible suffix for every node in our automaton.

In order to achieve this, we improve upon the classic Aho-Corasick DFA, which formally can be described as a finite ordered list of 5 symbols $M = \{Q, \sigma, \delta, q_i, F\}$, and transform it into a NFA 5-tuple $M' = \{Q, \sigma, \delta, q_i, F\}$. In both approaches, the Q symbol represents the finite set of states, σ is representative of the used alphabet, q_i is the start state, and F is the set of final states. The main difference between the formally described aforementioned methods is implementation of the transition set δ . In the classical approach, the transitions can occur from a single state to another single state, based on the input characters, and can be formally described as $\delta: Q \times \sigma \rightarrow Q$. In our approach, the transitions from a certain state can have multiple outcomes by parsing the same input characters, thus the transition definition will translate to $\delta: Q \times \sigma \rightarrow P(Q)$, where

$P(Q)$ is a power-set of Q , or a set of possible combinations of states in which the transition will end up. In order to start determining the output of any given transition, we employed the usage of the sliding window concept, accompanied by the computation of every possible suffix for every node in our automaton.

From an implementation point-of-view, our proposed NDFEA will add a new pre-processing step, computed after the failure functions calculation inside of the automaton. In this newly added pre-processing stage, for every node in our automaton, we will start parsing every possible transition, or when that is not possible, the failure function linked to that node, so that in the end, we computed all the viable suffixes of a length equal to that of a predetermined sliding window length, able to start from that node. This approach effectively produces, for every node, a list (as long as the sliding window length) of all possible matches from that particular node forward, significantly reducing matching comparisons when the local optimum is employed.

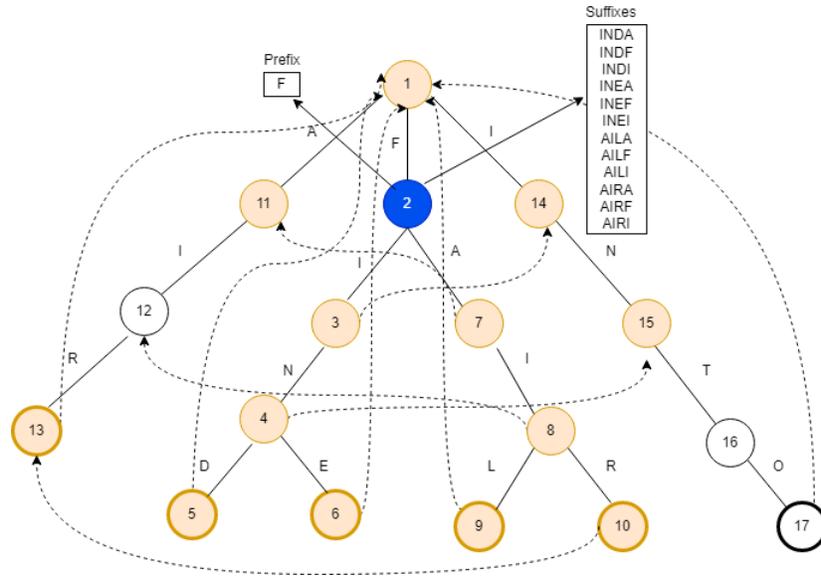


Fig. 3: The non-deterministic multiple pattern matching automaton for the same set of keywords $\{AND, FIND, FINE, FAIL, FAIR, INTO\}$, using a sliding window of length 4. Parsed nodes used for computing suffixes are colour marked.

For example, in figure 3, for the node 2, the preprocessing stage of computing the sliding window length suffixes will generate the suffixes list $[INDA, INDF, INDI, INEA, INEF, INEI, AILA, AILF, AILI, IRN]$. In memory-constrained environments, the introduction of this additional storage will prove a challenge of its own, one that opens new opportunities for future work and development.

GREAT	IDEA +
	INDA = 0.5
	INDF = 0.25
	INDI = 0.25
	INEA = 0.75

	AIRI = 0

Fig. 4: Computation of Hamming distance at node 2 for all the possible suffixes

3.2 Local optimum threshold computation

The nondeterministic behavior is properly highlighted in the actual processing stage. At this stage, the novelty brought by this NDFA approach is the possibility of applying local similarity measurements for approximate pattern-matching at any given step inside the parsing of the automaton for the length of the sliding window, instead of computing them at every position inside the input, for the length of the entire keyword at the processing stage. We chose to focus on the Euclid, Hamming and Levenshtein distances, as the similarity measures used for the local optimum threshold computation process. Figure 4 shows the comparisons done at node 2 when the input *GREATIDEA* is parsed, and also the select of the best result from all similarity scores computed. By obtaining a locally determined result of the similarity metric applied at any given node (between a concatenation of the node’s prefix and the suffixes, one at a time), and the string of the same length resulting by parsing the input to get to the respective node, we can decide if we move throughout the automaton with an existing transition (similar to the exact match of the Aho-Corasick DFA), or we jump to the mismatch transition, thus allowing multiple outcome states for the same input. The distinction between the aforementioned cases is done by an *a priori* definition of a local optimum threshold value for all the similarity metrics results. We hypothesize that choosing a high value for this local optimum threshold may result in certain keywords not being detected (false negatives), while setting the value too low, might produce other matches than expected (false positives). In order to try to mitigate the occurrence of false positives, we also added an additional verification so no match is detected if more than a given empirical percentage of the length is disjoint. The percentage is set based on the maximum length of the keywords and the sliding windows. For example, for a maximum keyword length of 10 characters, the existing similarity measurements for approximate pattern-matching are using a percentage of 20%, while for the same keyword length size and a sliding window of 5, in order to have related similarities, the percentage of our approach need to be set as 10%.

3.3 Blockchain NFT minting

While NFTs may have become popular due to digital art (in particular, photographs), they may also hold other types of information, including books, various texts, novels, short stories, etc., all of which are meant to ensure that specific

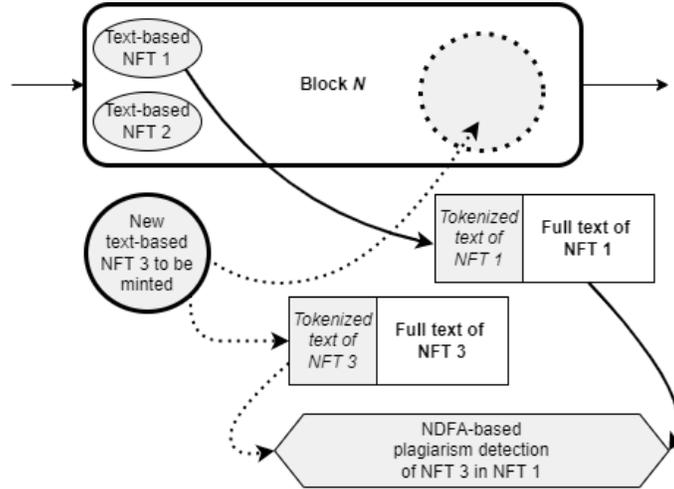


Fig. 5: Proposed implementation and architectural layout of a text-based NFT in a blockchain ledger.

form of digital art can be later traded or otherwise preserved in the blockchain. In our proposed approach, in order for an NFT to be minted, its associated digital data needs to be verified against plagiarism. After it passes verification, it can be stored in the ledger as minted. For text-based NFTs, we propose an approach where the original text is stored alongside the tokenized text. This way, tokenized text can be verified using our proposed NFA-based approach against plagiarism, with other already-minted NFTs. As speed is a concerning factor in all blockchain-drive platforms, our focus on this aspect becomes stronger and more clear. A proposed layout of the text-based NFTs in a blockchain ledger is outlined in Figure 5.

3.4 Challenges and results interpretation

The ideal situation for this scenario, as opposed to employing the traditional similarity measures in text mining, would be to obtain no false negatives and no false positives. The next best fit here would be the situation where we have false positives, but no false negatives - as the false positives can be trialed in a separate, post-processing step. The worst-case scenario is the situation where false negatives occur in large amounts, offsetting by a large margin the number of results of the classic similarity measures.

4 Experimental results

We have performed our experiments on an AMD Ryzen 5 5600H CPU, with 16 GB of DDR4 RAM and a 512 GB SSD drive, running on Windows 10. For

the experimental dataset, we have used slightly modified versions of datasets created by Paul Clough (Information Studies) and Mark Stevenson (Computer Science), at the University of Sheffield [12,13]. For scenario 1, we benchmarked the performance of our approach using a pair of datasets of 75 patterns (of variable length ranging between 4 to 15 characters), and an input size of 2,450 KB. We also ran consecutive tests in order to empirically determine the best local optimum threshold values, so that the results produced by our approach match the ones produced by the classic similarity measurements used: Euclid, Hamming and Levenshtein. Figure 7 showcase the improvement in terms of throughput, as well as the speed-up when benchmarking the performances of our approach and the basic similarity measurements for the Euclid, Hamming and Levenshtein distances computation.

Figures 6 and 8 showcase how the outcome is influenced by changing the local optimum threshold values. We correctly presumed that setting the value too high will cause false negatives. Also, due to the increased number of false positives, when the local optimum threshold is set too low, we would be dealing with an additional set of false negatives. This is caused by partial matches both in the sliding window and the input data, which produce better local similarity scores as opposed to the global score, associated with the entire pattern. Therefore, we recommend that the sliding window length is at least as high as the length of the longest pattern in our set of keywords.

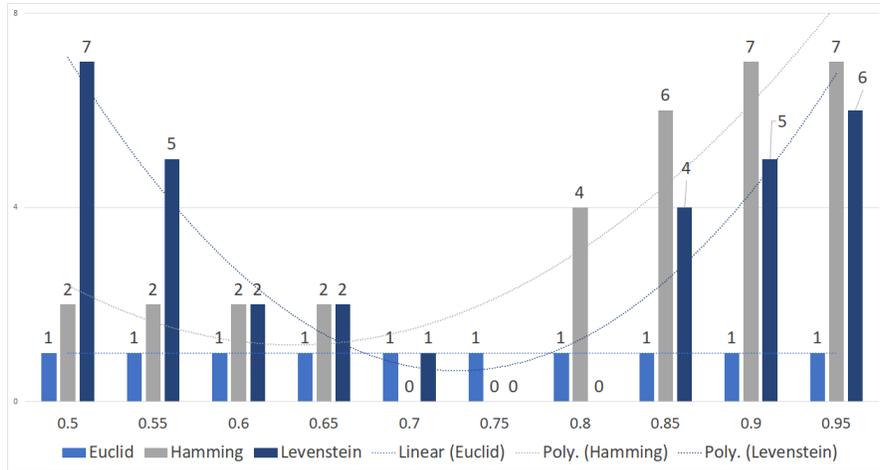


Fig. 6: Number of false negatives, for various local threshold values (0.5 to 0.95).

In terms of accuracy, the number of false negatives in Figure 6 is a constant 1 achieved for the Euclidean similarity, throughout the entire experiment. Both Hamming and Levenshtein similarities return no false negatives when the threshold value is set to 0.75, which seems to be the best fit for the scenarios

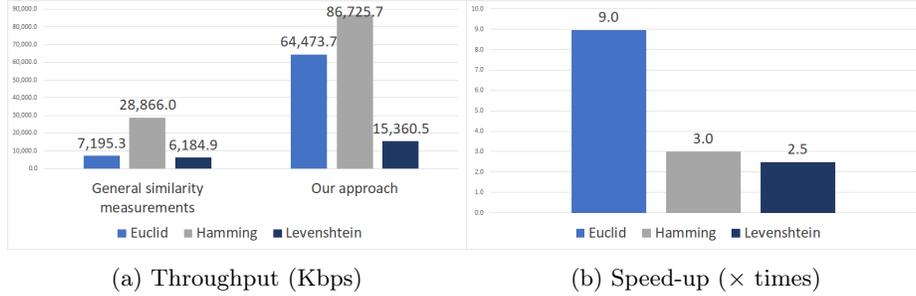


Fig. 7: Performance comparison between classic similarity measures and our proposed approach in Scenario 1

tested. The number of false positives remains constant for the Euclidean similarity, but follows a decreasing trend as the threshold values get closer to 1 for the others (which, in theory, would produce the same exact matching results as in the original Aho-Corasick DFA) - particularly, starting with a threshold value of 0.85, both Hamming and Levenshtein similarities produce no false positives.

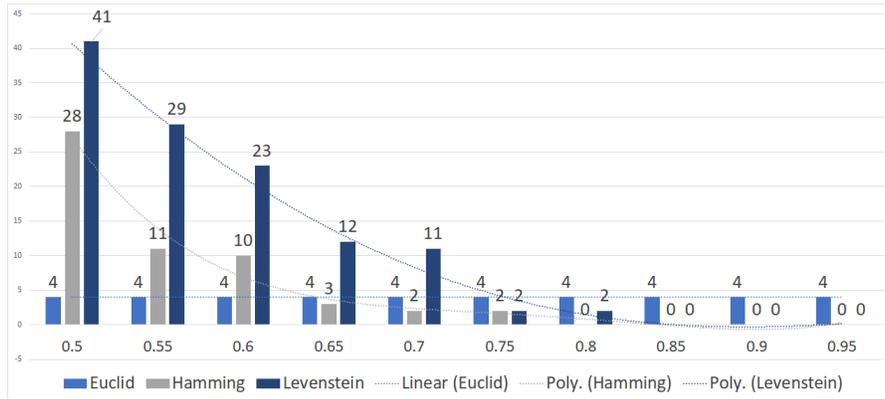


Fig. 8: Number of false positives, for various local threshold values (0.5 to 0.95).

For thorough testing, we have empirically observed that the longer the patterns in the keyword set tend to be, the faster the speed-up achieved for the Levenshtein similarity. The results of this particular setup (which we call here as scenario 2), with patterns of double lengths as in the original scenario 1, have been outlined in Figure 9. We have observed significant run-time speed-ups for our proposed approach, compared to the classic approaches. In particular, for the dataset tested in scenario 2, the Levenshtein similarity measurement throughput has improved to about $9.5\times$ higher throughput, up from $2.5\times$ higher throughput as in scenario 1. The Euclid-driven similarity measurement has achieved a speed-

up of about $9\times$ in both scenarios tested, and similarly the Hamming distance provided around $3\times$ higher throughput in our approach.

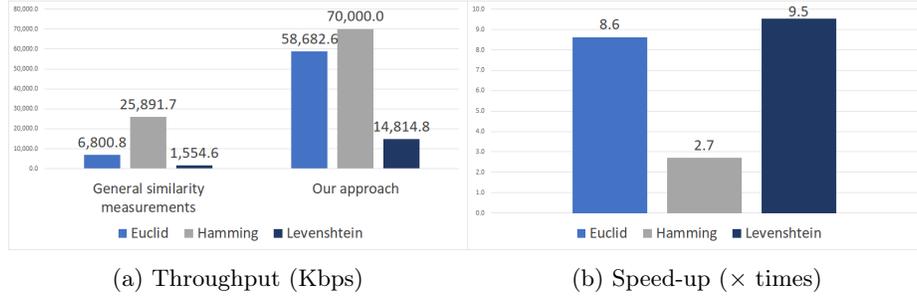


Fig. 9: Performance comparison between classic similarity measures and our proposed approach in Scenario 2

One additional challenge that we are currently looking to overcome, is the reduction of the storage size required for our proposed approach. While the throughput performance is significantly better than that the classic approaches, our approach uses significantly more memory due to the size of the NDFA automaton, including the storage of all suffixes at each node level. We are already working on an improved storage model for our approach therefore, which will be the subject of a future paper.

5 Conclusions

In this paper, we have presented an innovative fast approach to performing approximate-pattern matching for plagiarism detection, with particular applicability to blockchain-driven, NFT-ready platforms and ecosystems. For our proposed implementation, we have used our own NDFA-based automaton along with a sliding window concept and local thresholds at node-level, for tracing partial matches faster. We have tested our approach and concluded that it behaves suitably similar to existing various other similarity measures used in text mining for plagiarism detection, while obtaining a significant speed-up improvement of up to $9.5\times$ faster throughput. This makes our approach significantly faster than classic similarity measurements, and provides a potential approach to integrating the technology with blockchain-driven NFTs, in order to further secure the uniqueness of the non-fungible tokens and reduce the risk of plagiarism.

5.1 Future work

Future work includes various aspects of intrinsic functionality being tested, including: a) more similarity measurements being used, tested and potentially

developed; b) the corroboration of a semantic parser to the approximate pattern-matching performed in the lexical parsing process; c) the reduction of storage requirements for large sliding window sizes, and applicability of our approach to heterogeneous architectures; d) applicability of our methodology to other fields of research and data science.

5.2 Acknowledgements

This research was supported by the virtualLedger project [14] and the MOISE project number 240/2020, ID POC/398/1/1, financed by EU, Romanian government and West University of Timisoara. The views expressed in this paper do not necessarily reflect those of the corresponding project's partners.

References

1. Howcroft, E., Reuters. NFT sales hit \$25 billion in 2021, but growth shows signs of slowing. <https://www.reuters.com/markets/europe/nft-sales-hit-25-billion-2021-growth-shows-signs-slowing-2022-01-10/>, last accessed: 2nd of April, 2022
2. Anzelmi, D., Carlone, D., Rizzello, F., Thomsen, R., Hussain, D. M. A. (2011). Plagiarism Detection Based on SCAM Algorithm. In Proceedings of the International MultiConference on Engineers and Computer Scientists 2011 (Vol. Volume I, pp. 272-277). Newswood Limited, International Association of Engineers, IAENG.
3. Farah K. AL-Jibory, Mohammed S. H. Al- Tamimi (2021). Hybrid System for Plagiarism Detection on A Scientific Paper. In Turkish Journal of Computer and Mathematics Education 2021 (Vol.12 No 13 (2021), 5707-5719).
4. Saini A., Kumari S., Bahl A., Singh M.. Plagiarism Checker: Text Mining. In International Journal of Computer Applications 2016 (Volume 134 – No.3).
5. Putri, R. E., Putera, A. and Siahaan, U. Examination of Document Similarity Using Rabin-Karp Algorithm. In International Journal of Recent Trends in Engineering and Research 2017 (Vol. 3(8), pp. 196–201). DOI: 10.23883/IJRTER.2017.3404.4SNDK.
6. The Euclidean distance. Wikipedia. https://en.wikipedia.org/wiki/Euclidean_distance, last accessed: 2nd of April, 2022
7. The Hamming distance. Wikipedia. https://en.wikipedia.org/wiki/Hamming_distance, last accessed: 2nd of April, 2022
8. Levenshtein, Vladimir I. (February 1966). Binary codes capable of correcting deletions, insertions, and reversals. Soviet Physics Doklady. 10 (8): 707–710
9. Nakamoto, S. Bitcoin: A Peer-to-Peer Electronic Cash System (2008). <https://bitcoin.org/bitcoin.pdf>, last accessed: 2nd of April, 2022
10. SHA-2: Secure Hash Algorithm 2 (2001). <https://en.wikipedia.org/wiki/SHA-2>, last accessed: 2nd of April, 2022
11. Aho, A. V., & Corasick, M. J. (1975). Efficient String Matching: An Aid to Bibliographic Search. Commun. ACM, 18(6), 333–340. <https://doi.org/10.1145/360825.360855>
12. Paul Clough and Mark Stevenson. 2011. Developing a corpus of plagiarised short answers. Lang. Resour. Eval. 45, 1 (March 2011), 5–24. DOI:<https://doi.org/10.1007/s10579-009-9112-1>

13. Paul Clough and Mark Stevenson. 2011. Developing a corpus of plagiarised short answers. https://s3.amazonaws.com/video.udacity-data.com/topher/2019/January/5c4147f9_data/data.zip, last accessed: 2nd of April, 2022
14. The virtuaLedger project: <https://virtualedger.com>, last accessed: 2nd of April, 2022