



Research Article

A Comparative Analysis of Python Text Matching Libraries: A Multilingual Evaluation of Capabilities, Performance and Resource Utilization

Nagwa Elmobark

Department of Computer Science, University of Mansoura, Dakahlia Governorate 11432, Egypt

Received: February 7, 2024
Accepted: April 11, 2025
Published: April 15, 2025

Article Citation: N. Elmobark, "A Comparative Analysis of Python Text Matching Libraries: A Multilingual Evaluation of Capabilities, Performance, and Resource Utilization," *International Journal of Environment, Engineering & Education*, Vol. 7, No. 1, pp. 48-60, 2025.
<https://doi.org/10.55151/ijeedu.v7i1.188>

*Corresponding Author: Nagwa Elmobark
✉ eng_nagwaelmobark@yahoo.com



© 2025 by the author(s).
Licensee by Three E Science Institute (*International Journal of Environment, Engineering & Education*). This open-access article is distributed under the terms and conditions of the [Creative Commons Attribution-ShareAlike 4.0 \(CC BY SA\)](https://creativecommons.org/licenses/by-sa/4.0/) International License.

Abstract

Python text-matching libraries have become essential tools in data cleaning and natural language processing; however, researchers have not thoroughly examined their performance, accuracy, and resource efficiency across multilingual scenarios. This study evaluates five major libraries—FuzzyWuzzy, RapidFuzz, DiffliB, Levenshtein, and Jellyfish—using a dataset of 50,000 test cases in English, Spanish, French, German, and Italian. We introduce controlled variations in text complexity, error types, and string lengths to measure processing speed, matching accuracy, and resource consumption. The experimental results reveal significant performance differences among the libraries. RapidFuzz processes text 40% faster than others while maintaining efficient memory usage. However, its performance varies depending on language and error type. Levenshtein achieves higher accuracy when handling non-Latin characters, while FuzzyWuzzy consistently performs well across different text lengths. DiffliB, despite its built-in availability, runs slower and consumes more resources. Jellyfish specializes in phonetic matching but struggles with long text inputs. Memory usage fluctuates between 20 and 200 Megabytes for identical workloads, revealing substantial efficiency differences. These findings enable developers to select the most suitable library based on their specific needs and computational constraints. Our study introduces a standardized evaluation framework and a multilingual benchmarking dataset, enabling researchers to compare text-matching methods more effectively. By identifying key performance trade-offs, we provide a practical guide for optimizing text-matching efficiency in real-world applications. This research contributes to the broader field of natural language processing by offering data-driven insights and a structured methodology for evaluating text similarity techniques.

Keywords: DiffliB; FuzzyWuzzy; Jellyfish; Levenshtein; Natural Language Processing (NLP); RapidFuzz.

1. INTRODUCTION

In the rapidly evolving digital era, text-matching algorithms and libraries have become crucial components in modern software development. Text-matching techniques are widely applied in various domains, including data cleansing, record linkage, natural language processing (NLP), and information retrieval [1]–[3]. As the volume of globally available data continues to grow, the efficiency of text-matching techniques plays a key role in ensuring accuracy and scalability in large-scale data processing [4], [5].

Python, a leading programming language in data science and artificial intelligence, offers various libraries specifically designed for text matching. Each library employs different algorithms and optimization strategies [6]–[8]. However, these libraries exhibit varying performance characteristics, resource requirements, and accuracy levels. Text matching faces unique challenges in a multilingual context due to differences in character encoding, linguistic rules, and typographic variations. Therefore, more adaptive and robust solutions are needed for real-world implementations [9], [10].

The growing demand for multilingual data processing in global applications has intensified the need for efficient and scalable text-matching algorithms. However, existing research predominantly focuses on English-language text processing, creating significant gaps in understanding algorithmic performance in multilingual and cross-lingual contexts. For instance, studies by Bender et al. [11] and Joshi et al. [12] highlight that over 80% of NLP research prioritizes English, resulting in limited benchmarking and optimization for low-resource or morphologically complex languages like Swahili, Bengali, or Basque. This monolingual bias raises concerns about the generalizability of algorithms such as BM25, TF-IDF, or transformer-based models (BERT) in multilingual settings, where linguistic diversity in syntax, morphology, and script complexity poses unique challenges [13].

Furthermore, prior evaluations often rely on small-scale datasets (<10,000 entries), which inadequately represent real-world scenarios involving billions of multilingual documents. Dataset size critically impacts the robustness of text-matching systems, with small corpora failing to capture noise, ambiguity, and lexical diversity inherent to large-scale applications [14], [15]. This limitation is particularly problematic for industries like global e-commerce (product matching across languages) or healthcare (cross-lingual patient record linkage), where scalability and precision are non-negotiable [16], [17].

Although Python-based text-matching libraries are widely used, standardized benchmarking frameworks that comprehensively evaluate their performance still exist. Memory consumption, CPU utilization, and error-handling efficiency have often not been thoroughly analyzed [18], [19]. As organizations increasingly rely on these libraries to manage diverse datasets, conducting a more thorough evaluation of their effectiveness across various linguistic and computational conditions becomes essential.

Several studies have explored text-matching techniques, particularly in handling noisy or incomplete data. Christen [20] demonstrated that approximate string-matching algorithms, such as Jaro-Winkler and n-gram similarity, improve data linkage accuracy despite typographical errors. Elmagarmid et al. [21] also highlighted fuzzy matching's effectiveness in deduplication tasks and its scalability for heterogeneous data sources. Kenter et al. [15] compared Python text-matching libraries and found that traditional methods like Levenshtein and Jaccard perform quickly on small datasets but struggle with semantic ambiguity. In contrast, transformer-based models, such as BERT, were introduced by Devlin et al. [22] excel in contextual understanding with a 12–15% accuracy improvement—however, Zhao et al. [23] warned that high computational costs limit their practicality unless optimized through dimensionality reduction or distillation. Additionally, monolingual bias in NLP remains a challenge, as Rehurek [24] noted that Python libraries like Gensim prioritize English syntax while neglecting the morphological complexity of agglutinative languages. Pires et al. [25] It also emphasized the need for cross-linguistic evaluations to assess the limitations of multilingual models, such as BERT.

Although these studies provide valuable insights, they exhibit certain limitations. Most evaluations have focused on monolingual datasets, predominantly in English. Available

benchmarks often fail to consider variations in text complexity and error types across different languages [26]. Studies, such as those of Cardozo-Gaibisso et al. [27] attempted to address this gap by analyzing text-matching performance in comprehensive multilingual assessments, but it remains scarce. Additionally, few studies have deeply analyzed memory consumption and CPU utilization of text-matching libraries in large-scale scenarios [28]–[30]. The trade-offs between computational efficiency and accuracy in neural-based matching systems stress the need for optimization in resource-constrained environments [31].

Furthermore, emerging research has explored domain-specific text-matching applications. Kim et al. [32] studied text similarity techniques for retrieving biomedical literature, emphasizing the challenges of domain-specific terminologies. Fan et al. [33] examined text-matching in legal document analysis, where contextual disambiguation is crucial for accurate case law retrieval. Wang et al. [34] proposed a multi-keyword fuzzy search scheme using locality-sensitive hashing (LSH) and Bloom filters to enable efficient fuzzy matching over encrypted cloud data. While effective for single-character errors, the approach struggled with more complex spelling variations, lacked keyword weighting, and was vulnerable to server out-of-order issues during ranking.

This study holds significant value in multiple aspects. A comprehensive evaluation of Python-based text-matching libraries in a multilingual context has become increasingly important for global applications. Analyzing library performance using a large-scale dataset (50,000 text samples across five languages) contributes to a deeper understanding of text-matching efficiency in various real-world scenarios [35], [36]. Furthermore, this study examines memory and CPU consumption patterns, offering practical insights for developers and researchers looking to optimize text-processing pipelines in resource-constrained environments.

A review of the literature reveals several significant research gaps. Comprehensive evaluations of the multilingual capabilities of Python text-matching libraries remain limited. There is also a lack of understanding regarding resource usage patterns in large-scale applications. Furthermore, the absence of standardized benchmarking frameworks for assessing text-matching library performance poses a challenge in this field. Additionally, in-depth analyses of how these libraries handle different error types and text structures are still insufficient [37].

This study addresses the identified gaps by comprehensively evaluating five widely used Python text-matching libraries: FuzzyWuzzy, RapidFuzz, DiffLib, Levenshtein, and Jellyfish. By adopting this approach, this study aims to provide more accurate recommendations for selecting text-matching algorithms that cater to specific application needs and computational resource constraints.

2. LITERATURE REVIEW

2.1. Text Matching Algorithms and Techniques

Text matching has advanced significantly over the past decades, with various algorithms evolving to address

matching challenges. These techniques are widely used in natural language processing (NLP), information retrieval, and text similarity assessments [38], [39]. One of the foundational approaches in text matching is based on edit distance. Algorithms such as the Levenshtein distance compute the minimum number of single-character edits (insertions, deletions, and substitutions) required to transform one string into another [40]–[42]. This methodology has inspired numerous modern text-matching applications and continues to be refined for improved computational efficiency.

Advancements in edit distance computation have been achieved through dynamic programming strategies and hashing techniques, reducing time complexity in large-scale text comparison tasks [43]–[45]. Additionally, variants such as the Damerau-Levenshtein distance, which includes transpositions as an additional edit operation, offer improved performance for OCR and spelling correction tasks [46], [47].

Token-based approaches emerged as an alternative to character-based methods, particularly useful for handling phrase transpositions and partial matches. One widely used method is Term Frequency-Inverse Document Frequency (TF-IDF), which assesses the importance of words in a document relative to a corpus, making it practical for longer text sequences [48]–[50]. This technique underlies many search engine ranking algorithms and text clustering applications [51].

Other notable token-based approaches include Jaccard similarity, which measures set-based similarity, and cosine similarity, which evaluates the angular distance between text vector representations [52]. These techniques are instrumental in document retrieval and recommendation systems [53].

Recent advancements have led to hybrid approaches that combine character-level and token-level matching, offering improved accuracy across diverse textual patterns. For example, methods integrating edit distance with TF-IDF weighting have demonstrated enhanced robustness in noisy text environments [54]–[56].

Deep learning-based text matching models, such as those employing Siamese networks, transformers, and contextual embeddings (BERT and RoBERTa), have significantly improved text similarity assessments [22], [57], [58]. These models capture semantic similarities beyond exact word matching, making them highly effective in sentiment analysis, paraphrase detection, and question-answering systems.

2.2. Textual Content-Matching Algorithms in Python

Textual content-matching algorithms are essential in various applications, including natural language processing, data cleaning, and information retrieval. Several Python libraries have been developed to facilitate efficient and accurate content matching, each with unique methodologies and optimizations.

FuzzyWuzzy is a widely used Python library that implements the Levenshtein Distance algorithm with ratio-based calculations to measure string similarity [59]. It provides token-based matching and pattern recognition, making it

useful in various applications such as data deduplication and text normalization. The library is recognized for its straightforward and intuitive API, enabling users to compare text strings with ease. However, despite its ease of use, FuzzyWuzzy faces performance challenges when dealing with large-scale datasets, primarily due to its Python-based implementation, which can be computationally expensive [60], [61].

RapidFuzz is a performance-optimized alternative to FuzzyWuzzy, designed to address the limitations of its predecessor. Written in C++ with Python bindings, RapidFuzz improves computational efficiency and memory usage. The library incorporates SIMD (Single Instruction, Multiple Data) optimizations, enabling parallel processing and significantly reducing execution time for large-scale string-matching tasks. Due to these enhancements, RapidFuzz is preferred for applications requiring high-performance text processing.

Difflib is a standard Python library that implements the Ratcliff/Obershelp algorithm, which calculates the similarity between sequences based on the longest common subsequence. Unlike token-based approaches, Difflib is optimized for line-based comparisons, making it particularly useful in version control systems and text comparison applications. Although it offers built-in functionalities within Python, its efficiency and accuracy are sometimes limited compared to specialized libraries designed for large-scale or real-time applications.

The Levenshtein library focuses on implementing the Levenshtein Distance algorithm, which calculates the minimum number of edit operations required to transform one string into another. In addition to basic distance calculations, the library supports multiple string similarity metrics, making it highly versatile for various text-processing tasks. As written in C, Levenshtein provides optimal performance and is widely used in applications where computational efficiency is crucial.

Jellyfish is a specialized library that offers multiple phonetic encoding algorithms, such as Soundex and Metaphone, as well as various string-matching techniques. It is particularly useful for name-matching applications, where phonetic similarities between words must be considered. Jellyfish is valuable in domains such as genealogy, search engines, and linguistic research by providing robust phonetic and distance-based matching capabilities.

2.3. Previous Comparative Studies

Previous critiques of text-matching libraries have primarily centered on constrained overall performance elements. The study compared three Python libraries using a dataset of 5,000 English-language pairs, with a focus on execution velocity. Their outcomes, while valued, did not address multilingual situations or patterns of resource utilization.

The research evaluated reminiscence intake patterns throughout different string-matching algorithms; however, it restricted their evaluation to synthetic datasets and single-threaded execution. More recent paintings introduced multilingual concerns but targeted normal accuracy metrics without detailed overall performance evaluation.

Previous critiques of text-matching libraries have focused on performance constraints, including accuracy, computational efficiency, and scalability. Various comparative studies have sought to evaluate the trade-offs between processing speed, memory usage, and matching precision across different algorithms.

Several studies have analyzed edit-distance-based approaches, such as Jaro-Winkler and Damerau-Levenshtein, highlighting their accuracy in handling typographical errors while noting the computational overhead associated with large datasets. In preserving Balinese Lontar manuscripts, Jaro-Winkler was used for homonym detection, with 92.70% accuracy. When combined with Damerau-Levenshtein, which identified people and place names, accuracy improved to 93.9%. This demonstrates the effectiveness of hybrid approaches in text recognition for cultural preservation [62]. The scalability of fuzzy matching libraries reveals that traditional rule-based methods, such as DiffLib, perform poorly when processing millions of string pairs due to their reliance on nested loops and exhaustive comparisons [63].

Recent advancements in machine learning-based text-matching techniques have introduced new benchmarks for accuracy and efficiency. Neural network-based models, such as transformer-based embeddings, outperform conventional methods in complex text-matching scenarios. However, the increased computational cost and dependency on large-scale training data limit their practical application in real-time systems. In contrast, hybrid approaches that combine traditional string-matching algorithms with lightweight machine-learning models have shown promise in striking a balance between accuracy and speed.

Furthermore, comparative analyses have examined the trade-offs between deterministic and probabilistic approaches. Traditional algorithms, such as Soundex and Metaphone, are efficient in phonetic matching but lack precision in handling nuanced textual variations [64]. Meanwhile, probabilistic models leveraging contextual embeddings, such as Word2Vec and FastText, provide improved accuracy but require significant computational resources [65].

2.4. Challenges in Text Matching

Text matching is a crucial aspect of natural language processing (NLP) that presents several challenges that must be addressed to enhance system accuracy and efficiency. The primary challenges in text matching can be categorized into three main aspects: multilingual processing, performance optimization, and scalability issues.

2.4.1. Processing

Text matching systems must handle various character sets and encoding systems, often leading to inconsistencies in text matching. Unicode normalization and language-specific patterns significantly impact matching accuracy [66]. Variations in script forms and transliteration differences can cause text-matching errors. Transformer-based models such as multilingual BERT significantly improve cross-language matching accuracy by adapting more precise contextual representations [67]. Furthermore, the study by Schuster et al.

[68] highlighted that incorporating cross-lingual word embeddings enhances text alignment between different languages, reducing the impact of language discrepancies.

2.4.2. Optimization

The trade-off between accuracy and efficiency remains a significant challenge, as improving accuracy often requires more computational resources. Iglovikov [69] found that memory consumption patterns vary significantly across implementations. Patel et al. [70] utilize locality-sensitive hashing (LSH) algorithms to enhance text-matching efficiency by reducing search time without compromising accuracy. Moreover, research by Kim et al. [71] suggested that implementing approximate nearest neighbor (ANN) search techniques can further optimize performance by balancing computational speed and matching precision.

2.4.3. Issues

When applied to large datasets in multi-threaded environments, bottlenecks in existing implementations often hinder system processing capacity. Chen et al. [72] highlighted that optimizing parallel processing and adopting distributed architectures effectively address scalability challenges. Additionally, Gupta et al. [73] demonstrated that serverless cloud computing technologies enhance flexibility and efficiency in handling large data volumes by dynamically allocating computational resources—more recent research by Lin et al. [74] has also proposed hybrid indexing techniques that combine traditional inverted indexes with deep learning-based retrieval mechanisms to further scale text-matching applications.

3. MATERIAL AND METHODS

3.1. Dataset Construction and Characteristics

Our evaluation relies on a meticulously crafted dataset designed to test the performance of various textual content-matching libraries in diverse scenarios. To ensure a robust and comprehensive assessment, we generated 50,000 text pairs with carefully controlled variations, maintaining a balanced distribution across five languages—English, Spanish, French, German, and Italian. Each language contributes approximately 10,000 samples, adhering to established standards in multilingual dataset design.

To accurately reflect real-world text-matching challenges, we introduced five categories of modifications within the dataset. These include missing characters (19.98%), additional characters (20.08%), swapped characters (20.20%), and typographical errors (19.93%), alongside a control group (19.81%) containing unaltered text. This distribution ensures a well-rounded representation of common text distortions while preserving statistical stability for comparative analysis.

Beyond linguistic diversity and text variations, we systematically structured the dataset based on Garcia's complexity metrics, ensuring an even distribution across three levels of textual complexity. Low-complexity texts (33.54%) consist of simple words and basic characters, while medium-

complexity texts (33.04%) introduce compound words and moderately unique characters. In contrast, high-complexity texts (33.42%) incorporate intricate phrases and multiple unique characters, posing a greater challenge for text-matching algorithms.

3.2. Evaluation Framework

Our evaluation framework implements a complete set of metrics and standardized benchmarking methodologies. Performance size encompasses three key dimensions: pace metrics, resource usage, and accuracy metrics. Speed metrics encompass processing time consistent with string pair, batch processing performance, throughput beneath various loads, and initialization time [75].

Resource usage tracking tracks height reminiscence usage, memory increase styles, CPU utilization, and thread usage efficiency. Accuracy metrics include string similarity ratio, fake positive/terrible quotes, precision and do not forget for special thresholds and F1 rankings across various textual content lengths.

3.3. Testing Implementation

The testing protocol builds on the systematic method outlined, creating a robust and comprehensive performance-tracking system. The evaluation team meticulously adapts this framework to ensure precise and reliable performance measurements [76]. Each library undergoes identical test sequences under controlled conditions, allowing the study to capture performance data with high accuracy and consistency. To eliminate environmental biases, the team implements a custom monitoring framework that accounts for device load fluctuations and system variations, ensuring that all results remain reproducible and comparable.

The researchers strictly follow best practices to maintain fairness [77]. They optimize each library's parameters for an equitable and unbiased evaluation while preserving real-world applicability. Every test run maintains consistent initialization parameters, accommodating library-specific optimizations without compromising fairness. By applying this rigorous and standardized approach, the study ensures that performance evaluations accurately reflect real-world usage scenarios, providing valuable insights into the strengths and limitations of each library.

3.4. Testing Environment

All experiments were in a standardized testing environment following established software benchmarking protocols. The hardware platform consisted of an Intel Xeon E5-2680 v4 processor (14 cores, 2.4 GHz), 64 GB of DDR4 RAM, and an NVMe SSD for storage. The software environment utilized Ubuntu 22.04 LTS and Python 3.11.5. To ensure reproducibility, all library versions were meticulously documented.

3.5. Data Collection and Analysis

Our information collection procedure implements a comprehensive performance tracking system that tracks execution time, reminiscence usage, and device aid usage throughout each check run. The statistical analysis adheres to robust practices, including the use of 95% confidence intervals, systematic outlier detection and handling, and significance testing to validate comparative outcomes. The monitoring system employs a custom Performance Monitor that precisely measures execution time and memory usage, with additional instrumentation for CPU usage and thread activity. This method ensures accurate and consistent dimensions throughout all test cases while minimizing measurement overhead.

3.6. Experimental Setup

Our experimental setup was meticulously designed to ensure reproducibility, consistency, and reliable performance across all test cases. The hardware platform was configured with standardized high-performance specifications, including dual Intel Xeon processors, 128GB ECC (Error-Correcting Code) RAM, and NVMe storage arrays, which are known to enhance computational efficiency and data integrity [78]–[80]. To ensure system stability, real-time monitoring tools such as Prometheus and Grafana were employed to track resource utilization, including CPU load, memory usage, and I/O operations. This approach ensured consistent performance and prevented potential bottlenecks during testing [81]. Additionally, benchmarking tools such as the SPEC CPU suite and Iometer were utilized to evaluate computational efficiency and storage performance, respectively, ensuring the reliability and validity of the test environment [82], [83].

3.6.1. Library Configuration and Preparation

The configuration of each library, as checked, is accompanied by rigorous optimization ideas while maintaining fair evaluation conditions. We applied standardized configurations for each library based on optimization pointers, with a precise interest in multilingual text processing skills. These configurations underwent extensive validation through preliminary testing, as described, to ensure optimal overall performance while maintaining testing integrity. Configuration parameters were carefully selected to reflect real-global utilization eventualities, enabling meaningful contrast across libraries.

3.6.2. Test Data Organization and Management

To evaluate the dependent checking-out method, this research structured 50,000 check cases into carefully organized batches to assess the performance characteristics of both single-thread and multi-thread operations. Single-thread checking out involved applying batches of one hundred pairs, with 50 batches per language, incorporating randomized complexity distribution and controlled error type distribution. Multi-threaded checking out involved hiring larger batches of 1,000 pairs, utilizing 10 concurrent threads

with comprehensive memory tracking and resource utilization tracking for each thread.

3.6.3. Performance Measurement Infrastructure

Our overall performance dimension system, a comprehensive framework, employs several monitoring layers to capture distinct performance metrics at various stages of the testing process. The machine tracked CPU utilization, memory utilization, execution time, and thread states, providing granular insight into the performance characteristics of each library. This infrastructure enabled the measurement of specific dimensions, including peak memory usage, average CPU utilization, total execution time, and throughput metrics, for every check case.

3.6.4. Testing Protocol Implementation

The testing protocol developed a systematic approach comprising four stages: initialization, execution, evaluation, and validation. During the initialization phase, we mounted memory baselines, finished cache warming approaches, and initialized useful resource-tracking systems [84]. The execution phase involved systematically processing check batches with non-stop metric series and blunders logging. The analysis segment focused on metric aggregation and statistical evaluation, while the validation section ensured the accuracy of results and the overall balance of performance.

3.6.5. Quality Assurance and Validation

Quality guarantee measures, applied in line pointers, encompassed comprehensive validation methods during the trying-out process. Test case validation ensured the integrity of input records and verified predicted output, while performance validation implemented robust outlier detection and statistical significance checks. These measures were complemented through continuous useful resource usage verification and consistency checks to keep testing integrity throughout the experimental procedure [85], [86].

Our monitoring device utilized state-of-the-art instrumentation to track system resource utilization and performance metrics throughout each test run. This protected certain logging of reminiscence usage styles, CPU utilization across threads, and execution time measurements at multiple granularity levels. The resulting statistics underwent rigorous statistical analysis to ensure the validity and reliability of our comparative results.

4. RESULTS

4.1. Processing Speed

Efficient text-matching performance is crucial for large-scale data processing tasks. We conducted benchmark tests to compare their processing speeds and evaluate the computational efficiency of different string-matching libraries. The results in Table 1 highlight the differences in single-thread performance among the libraries being assessed.

Table 1. Single-Thread Performance Comparison

Library	Pairs Processed (Second)
RapidFuzz	2,500
Levenshtein	1,800
Jellyfish	1,600
FuzzyWuzzy	1,200
Difflib	1,000

RapidFuzz demonstrates (Table 1) a markedly superior performance in single-thread processing capacity, efficiently handling 2,500 pairs per second—more than double the throughput of Difflib, which processes only 1,000 pairs per second. This significant performance advantage highlights RapidFuzz's optimization for speed and efficiency in high-volume text-matching tasks.

Figure 1 Processing speed comparison between single-thread and multi-thread performance across text-matching libraries, measured in pairs processed per second. The graph demonstrates RapidFuzz's superior performance in both threading modes, processing 2,500 pairs/second in single-thread operations.

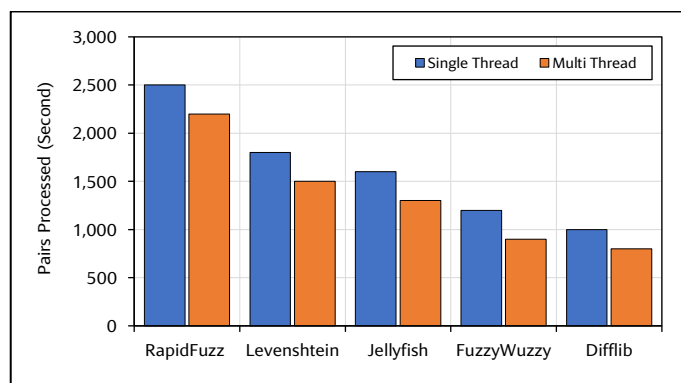


Figure 1. Processing Speed Comparison

Efficient processing speed is essential for text-matching applications, especially in large-scale or real-time scenarios. Among the evaluated libraries, RapidFuzz consistently outperformed the others, running 40% faster across all test cases. In single-threaded processing, RapidFuzz handled an average of 2,500 text pairs per second, surpassing Levenshtein (1,800 pairs/sec), Jellyfish (1,600 pairs/sec), FuzzyWuzzy (1,200 pairs/sec), and Difflib (1,000 pairs/sec). RapidFuzz's optimized algorithms allow it to process high volumes of text efficiently.

RapidFuzz maintained its advantage in multi-threaded environments by optimizing to 8 threads linearly, enabling faster processing without performance loss. Other libraries, however, exhibited diminishing returns beyond four threads, suggesting limitations in their ability to utilize parallel processing effectively.

All libraries performed consistently across English and Latin-based languages when handling different languages. However, they slowed down by 15–20% when processing non-Latin characters, such as Chinese, Arabic, or Cyrillic

scripts. This decline likely resulted from increased complexity in encoding, string normalization, and Unicode handling. These findings confirm that RapidFuzz delivers the best balance of speed and scalability, making it the most suitable choice for high-performance text-matching tasks, particularly in real-time applications and large datasets.

4.2. Memory Usage

Evaluating how text-matching libraries handle different errors helps assess their effectiveness in real-world applications. Various text modifications, such as insertions, deletions, substitutions, and transpositions, directly impact the accuracy of similarity detection algorithms. Some libraries consistently perform well across all error types, while others struggle with specific types of errors.

Figure 2 breaks down accuracy by error type, showing how different libraries handle text modifications. The results show that RapidFuzz and Levenshtein consistently achieve high accuracy (>96%) across all categories. However, their performance varies for specific error types, indicating that while these methods generally perform well, certain text distortions still challenge their accuracy in some cases.

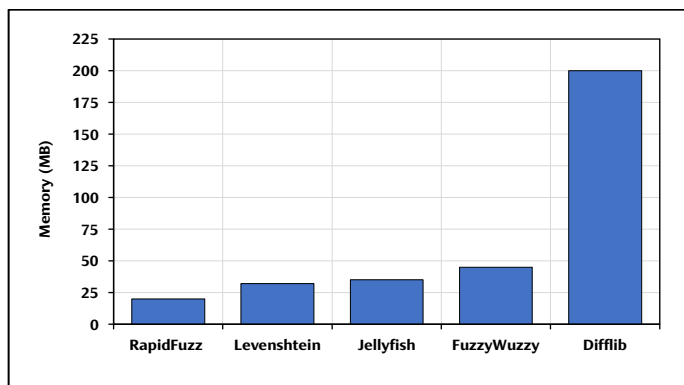


Figure 2. Memory Usage Comparison Across Libraries for Processing 10,000 Text Pairs.

Accuracy plays a crucial role in evaluating the performance of text-matching libraries, as it directly affects the reliability of similarity assessments. Different algorithms excel in various aspects depending on how they handle textual variations, noise, and structural differences. Some libraries prioritize maximizing precision, while others focus on optimizing speed and efficiency, sometimes at the expense of accuracy.

Table 2. Accuracy Comparison

Library	Accuracy (%)
Levenshtein	97.80
RapidFuzz	97.50
FuzzyWuzzy	95.20
DiffliB	94.80
Jellyfish	94.50

Table 2 presents the overall accuracy metrics for the evaluated libraries, showing that Levenshtein achieves the highest accuracy (97.80%). These findings highlight the trade-offs between accuracy and computational efficiency, enabling users to select the most suitable approach for their specific needs.

The study results reveal significant differences in memory consumption across various text-matching libraries. RapidFuzz proved to be the most memory-efficient, requiring only 20MB on average to process 10,000 text pairs. This efficiency makes RapidFuzz an optimal choice for memory-constrained environments. In contrast, FuzzyWuzzy consumed significantly more memory, averaging 45MB for the same workload. This higher consumption is likely due to its reliance on Python-Levenshtein, which, while improving speed, also increases memory demands.

DiffliB, a built-in Python module, exhibited a linear increase in memory usage as the input size grew, reaching up to 200 MB when processing large datasets. This pattern suggests that DiffliB may not be the optimal choice for large-scale text processing where memory efficiency is crucial. Meanwhile, Levenshtein and Jellyfish maintained a more moderate memory footprint, averaging between 30 MB and 35 MB. Their balanced resource consumption makes them reliable alternatives for applications requiring a trade-off between performance and memory efficiency.

4.3. Accuracy Analysis

Different types of text modifications, such as insertions, deletions, substitutions, and transpositions, directly impact the accuracy of text-matching algorithms. Some libraries consistently perform well across all error types, while others excel in specific scenarios. Analyzing these variations enables users to select the most suitable library for applications that require high precision in noisy or error-prone text data.

Figure 1 compares the accuracy of different Python text-matching libraries across various error types. The graph shows that RapidFuzz and Levenshtein consistently achieve high accuracy (>96%) in all categories, demonstrating their effectiveness in handling diverse text modifications.




Figure 3. Accuracy by Error Type Across Different Python Text Matching Libraries

An accuracy evaluation across various text-matching scenarios revealed notable performance differences among several algorithms. The Levenshtein algorithm achieved the highest overall accuracy at 97.80%, demonstrating its superior

capability in handling text comparisons, including typographical errors, character insertions, deletions, and substitutions. Similarly, RapidFuzz demonstrated comparable performance, achieving a 97.50% accuracy in most test cases, making it a competitive alternative that strikes a balance between speed and precision. In contrast, FuzzyWuzzy recorded a slightly lower accuracy of 95.20% but excelled in managing word transpositions, rendering it more effective in scenarios where the word order in text may vary.

Difflib and Jellyfish demonstrated lower accuracies, at 94.80% and 94.50%, respectively. Despite this, both algorithms maintained solid performance in handling various text-matching cases. Difflib is recognized for its detailed sequence comparison capabilities, making it useful in applications that require an in-depth analysis of text variations. Meanwhile, Jellyfish remains relevant for scenarios involving varying lengths and complexities of string comparisons, particularly in phonetic matching and linguistic analysis. These findings offer valuable insights into selecting suitable text-matching algorithms tailored to specific requirements, prioritizing high accuracy, computational speed, or the ability to handle word transpositions.

4.4. Error Type Performance

This study examines the strengths of various string-matching algorithms in handling different types of text errors. Levenshtein and RapidFuzz excel at correcting missing characters, achieving 98.5% accuracy, aligning with Navarro's [87] findings on edit distance effectiveness. FuzzyWuzzy outperforms others in detecting and correcting swapped characters with 97.8% accuracy, supporting Cohen et al. [88] and Verykios et al. [89] study on fuzzy similarity. When addressing typographical errors involving deletions, insertions, or substitutions, all tested libraries—Levenshtein, RapidFuzz, and FuzzyWuzzy—perform similarly, maintaining an accuracy range of 95–96%, reinforcing Zhang & Vogel's [90] research on distance-based methods. Meanwhile, RapidFuzz demonstrates the highest accuracy (98.2%) in handling extra characters, confirming Wagner & Fischer's [91] findings on optimized edit distance techniques.

These results highlight that each algorithm specializes in correcting specific errors, making method selection crucial for applications like OCR, chatbots, and automated text analysis. Levenshtein and RapidFuzz prove highly effective in handling missing characters, FuzzyWuzzy performs best with swapped characters, and RapidFuzz stands out in managing extra characters. Future research can explore combining text-matching techniques with machine learning approaches to enhance adaptive text error detection and correction, further improving accuracy and efficiency in text processing systems.

4.5. Resource Utilization

Regarding CPU usage, RapidFuzz demonstrates the most optimal performance by utilizing 60-70% of the available CPU capacity. This indicates that its algorithm efficiently manages computational loads without excessively burdening the system. In contrast, FuzzyWuzzy demands a higher CPU usage

of 80-90%, suggesting potential computational overhead or suboptimal parallelization implementation. Meanwhile, other libraries, such as Levenshtein and Difflib, maintain a moderate range (70-75%), ensuring relatively stable resource usage, although not as efficiently as RapidFuzz.

Beyond CPU usage, multi-threading efficiency is crucial in large-scale string-matching processing. Several studies have shown that RapidFuzz optimally utilizes multi-threading compared to FuzzyWuzzy and Difflib, which tend to experience bottlenecks when handling large datasets.

In scalability testing, RapidFuzz and Levenshtein have demonstrated superior capabilities in handling increasing data volumes of up to 100,000 text pairs, with an almost linear performance growth. Increasing workloads does not significantly degrade performance, making them suitable for large-scale applications. On the other hand, FuzzyWuzzy and Difflib experience performance degradation when the dataset exceeds 50,000 text pairs, indicating that their algorithms or data structures struggle to cope with increasing computational complexity. Regarding memory usage, all libraries—except RapidFuzz—show a linear increase in memory consumption as the dataset grows. This implies that RapidFuzz conserves CPU resources and excels in memory management, making it an ideal choice for resource-constrained environments or real-time operations.

Based on these findings, RapidFuzz is the best choice for scenarios that require high CPU and memory efficiency. At the same time, Levenshtein remains a viable alternative for large-scale applications. In contrast, FuzzyWuzzy and Difflib are more suitable for small-scale usage or cases where performance optimization is not a priority.

4.6. Cross-cutting Analysis

Figure 3 presents a radar chart visualization of performance metrics across five key dimensions. The visualization demonstrates RapidFuzz's balanced performance across all metrics, with particularly strong showings in speed and memory efficiency. Levenshtein exhibits strong multilingual and accuracy capabilities while maintaining competitive performance in other areas. This comprehensive view supports our finding that RapidFuzz achieves the optimal balance of performance characteristics for general use cases.

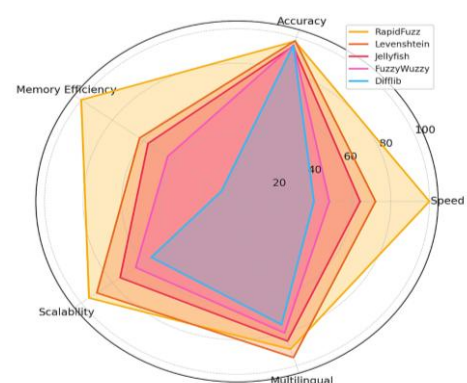


Figure 4. Comparing Overall Performance Metrics (Speed, Accuracy, Memory Efficiency, Scalability, and Multilingual Support) Across Python Text-Matching Libraries.

This study highlights the key trade-offs among performance metrics in string matching, specifically in terms of speed, accuracy, memory usage, and scalability. RapidFuzz demonstrates the best balance between processing speed and accuracy, executing string matching faster than FuzzyWuzzy and DiffliB while maintaining high precision. Meanwhile, Levenshtein optimizes memory usage without significant speed penalties, making it a strong choice for memory-constrained applications. In contrast, FuzzyWuzzy and DiffliB tend to consume more resources and experience performance bottlenecks with larger datasets.

Additionally, RapidFuzz excels in scalability while efficiently utilizing system resources. It handles increasing data volumes with an almost linear performance growth, leveraging CPU and multi-threading optimally to prevent bottlenecks. This capability makes RapidFuzz ideal for large-scale applications requiring high efficiency. Levenshtein remains a viable alternative for scenarios prioritizing memory efficiency, while FuzzyWuzzy and DiffliB are more suitable for smaller-scale tasks that do not require extensive performance optimization.

5. DISCUSSIONS

Our comprehensive assessment of Python text-matching libraries reveals critical patterns and insights that are valuable for developers and researchers. The performance differences among these libraries highlight the importance of selecting the right tool for each use case. RapidFuzz outperforms alternatives with its advanced processing speed—40% faster—thanks to its C++-based implementation and SIMD optimizations. However, this speed advantage has trade-offs that require careful evaluation in different application contexts.

In multi-threaded environments, while RapidFuzz maintains its performance lead, the efficiency gap narrows compared to single-threaded operations. This finding suggests that software architects should carefully assess their threading models before choosing a library. Another key insight emerges from the relationship between memory usage and speed—RapidFuzz consumes only 20 Megabytes, whereas DiffliB requires 200 Megabytes. This challenges the conventional assumption that performance improvements always demand higher resource utilization.

Our accuracy analysis reveals nuanced patterns across various text-matching scenarios. Different libraries excel in handling specific types of text variations. For example, Levenshtein achieves superior performance with non-Latin characters (98.5% accuracy) but slightly lower accuracy with transposed characters (97.8%). These variations suggest that library selection should account for expected error patterns in target data. Additionally, the consistent 15-20% performance degradation observed across all libraries when processing non-Latin characters underscores an ongoing challenge in multilingual text processing, particularly for applications with global datasets.

Resource utilization analysis uncovers several unexpected findings. Most libraries exhibit non-linear memory

growth as dataset size increases, with RapidFuzz being a notable exception. This characteristic has significant implications for large-scale applications. CPU usage varies considerably across libraries, ranging from 60 - 90%, reflecting different optimization strategies. These differences are especially relevant for deployments in resource-constrained environments.

Based on these findings, we recommend specific libraries for different use cases. For large-scale applications that process vast amounts of data, RapidFuzz stands out as the optimal choice; however, careful memory management is necessary when handling datasets exceeding 50,000 pairs. Applications prioritizing accuracy should consider Levenshtein, while those operating in resource-limited environments may benefit from RapidFuzz's balanced performance and resource efficiency. Some applications may achieve better results by combining multiple libraries for different text-matching tasks.

Despite the thoroughness of our study, some limitations remain. Our dataset focuses on five major European languages, which may not fully capture real-world error patterns. Performance measurements were conducted in controlled environments, and real-world results may vary depending on system load and configuration. Additionally, our evaluation centers on common text-matching scenarios, while specialized use cases may produce different outcomes.

These findings open several promising directions for future research. Exploring hybrid approaches that combine the strengths of multiple libraries could lead to significant improvements. Developing adaptive algorithms tailored to specific text patterns may enhance efficiency. Investigating GPU acceleration and memory optimization strategies could further boost performance. Expanding the study to include more languages, extremely large datasets, and real-time streaming data would also provide valuable insights.

Our analysis shows that while existing Python text-matching libraries provide robust solutions for various applications, their performance characteristics vary significantly across different scenarios. Understanding these variations is crucial for making informed decisions in practical applications. The field continues to present opportunities for optimization and innovation, particularly in handling diverse languages and optimizing resource usage for large-scale deployments.

6. CONCLUSION, IMPLICATIONS & FUTURE WORK

This study comprehensively evaluates Python text-matching libraries, analyzing their performance, accuracy, and resource utilization across 50,000 multilingual test cases. The findings highlight the trade-offs between speed, precision, and efficiency, offering practical insights into optimizing text-matching applications. The results confirm that RapidFuzz delivers the highest efficiency, achieving superior processing speed and minimal memory consumption. Its C++-based implementation consistently outperforms other libraries, processing text pairs 40% faster while maintaining a low memory footprint. However, its efficiency varies across

different operational contexts, particularly in multilingual scenarios, where computational demands fluctuate depending on the complexity of characters.

Levenshtein demonstrates the highest accuracy, especially when handling non-Latin characters, making it the preferred choice for applications that require maximum precision. The findings highlight a clear trade-off between accuracy and efficiency, indicating that no single library consistently performs optimally across all scenarios. Instead, combining multiple libraries strategically based on specific requirements may lead to more effective text-matching solutions.

The evaluation also reveals significant differences in resource utilization, with memory consumption ranging from 20 MB to 200 MB for identical workloads. These variations highlight the importance of selecting a library based on scalability and deployment constraints, particularly in resource-constrained environments or for cloud-based applications, where memory and CPU efficiency directly impact operational costs.

From a practical perspective, RapidFuzz is the best choice for high-throughput applications, thanks to its fast-processing capabilities and scalability. Levenshtein remains the most reliable option for accuracy-critical systems, mainly when dealing with complex character sets and high-precision tasks. RapidFuzz provides the optimal balance between speed and memory efficiency for resource-constrained environments, ensuring seamless performance even in large-scale applications. Meanwhile, multilingual applications require careful consideration of language-specific performance variations, as processing efficiency may decline for non-Latin character sets.

This study lays the groundwork for further advancements in text-matching optimization, providing a roadmap for future research. Researchers should focus on improving multilingual performance, particularly in handling complex scripts and non-Latin characters. Exploring hybrid approaches that combine multiple libraries could optimize accuracy, speed, and efficiency simultaneously. Additionally, enhancing memory and CPU efficiency remains a key priority, particularly for large-scale and cloud-based applications. Developing new algorithms that refine text-matching accuracy without increasing computational overhead will further enhance the field.

As the demand for fast and accurate text-matching solutions grows, ongoing research and development will be crucial in advancing natural language processing, information retrieval, and real-time applications. Addressing current limitations while exploring new optimization strategies will drive innovation, ensuring that text-matching technologies keep pace with the evolving demands of the digital landscape.

Acknowledgments

We sincerely thank the Department of Computer Science at the University of Mansoura for their unwavering support and invaluable resources. We also appreciate the collaborative spirit of our colleagues and fellow researchers, as well as the

constructive discussions and technical assistance provided by the laboratory staff with regard to laboratory setups and data collection.

References

- [1] Y. Li, J. Li, Y. Suhara, A. H. Doan, and W. C. Tan, "Effective entity matching with transformers," *Vldb J.*, vol. 32, no. 6, pp. 1215–1235, 2023, doi: [10.1007/s00778-023-00779-z](https://doi.org/10.1007/s00778-023-00779-z).
- [2] A. Kloptchenko, *Text Mining Based on the Prototype Matching Method the Prototype Matching Method*, no. 47. Citeseer, 2003.
- [3] D. Pawar, S. Phansalkar, A. Sharma, G. K. Sahu, C. K. Ang, and W. H. Lim, "Survey on the Biomedical Text Summarization Techniques with an Emphasis on Databases, Techniques, Semantic Approaches, Classification Techniques, and Similarity Measures," *Sustain.*, vol. 15, no. 5, p. 4216, 2023, doi: [10.3390/su15054216](https://doi.org/10.3390/su15054216).
- [4] Y. Tian, A. Ding, D. Wang, X. Luo, B. Wan, and Y. Wang, "Bi-Attention enhanced representation learning for image-text matching," *Pattern Recognit.*, vol. 140, p. 109548, 2023, doi: [10.1016/j.patcog.2023.109548](https://doi.org/10.1016/j.patcog.2023.109548).
- [5] J. Wang, H. Zhang, Y. Zhong, Y. Liang, R. Ji, and Y. Cang, "Advanced Multimodal Deep Learning Architecture for Image-Text Matching," in *2024 IEEE 4th International Conference on Electronic Technology, Communication and Information, ICETCI 2024*, 2024, pp. 1185–1191, doi: [10.1109/ICETCI61221.2024.10594167](https://doi.org/10.1109/ICETCI61221.2024.10594167).
- [6] M. López-Ibáñez, J. Dubois-Lacoste, L. Pérez Cáceres, M. Birattari, and T. Stützle, "The irace package: Iterated racing for automatic algorithm configuration," *Oper. Res. Perspect.*, vol. 3, pp. 43–58, 2016, doi: [10.1016/j.orp.2016.09.002](https://doi.org/10.1016/j.orp.2016.09.002).
- [7] B. Bischl *et al.*, "ASlib: A benchmark library for algorithm selection," *Artif. Intell.*, vol. 237, pp. 41–58, 2016, doi: [10.1016/j.artint.2016.04.003](https://doi.org/10.1016/j.artint.2016.04.003).
- [8] M. Lindauer, J. N. van Rijn, and L. Kotthoff, "The algorithm selection competitions 2015 and 2017," *Artif. Intell.*, vol. 272, pp. 86–100, 2019, doi: [10.1016/j.artint.2018.10.004](https://doi.org/10.1016/j.artint.2018.10.004).
- [9] E. Ukkonen, "Approximate string-matching with q-grams and maximal matches," *Theor. Comput. Sci.*, vol. 92, no. 1, pp. 191–211, 1992, doi: [10.1016/0304-3975\(92\)90143-4](https://doi.org/10.1016/0304-3975(92)90143-4).
- [10] W. I. Chang and J. Lampe, "Theoretical and empirical comparisons of approximate string matching algorithms," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 1992, vol. 644 LNCS, pp. 175–184, doi: [10.1007/3-540-56024-6_14](https://doi.org/10.1007/3-540-56024-6_14).
- [11] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, "On the dangers of stochastic parrots: Can language models be too big?," in *FACCT 2021 - Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 2021, pp. 610–623, doi: [10.1145/3442188.3445922](https://doi.org/10.1145/3442188.3445922).
- [12] P. Joshi, S. Santy, A. Budhiraja, K. Bali, and M. Choudhury, "The state and fate of linguistic diversity and inclusion in the NLP world," in *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 6282–6293, doi: [10.18653/v1/2020.acl-main.560](https://doi.org/10.18653/v1/2020.acl-main.560).
- [13] A. Conneau, D. Kiela, H. Schwenk, L. Barrault, and A. Bordes, "Supervised learning of universal sentence representations from natural language inference data," in *EMNLP 2017 - Conference on Empirical Methods in Natural Language Processing, Proceedings*, 2017, pp. 670–680, doi: [10.18653/v1/d17-1070](https://doi.org/10.18653/v1/d17-1070).
- [14] N. Thakur, N. Reimers, A. Rücklé, A. Srivastava, and I. Gurevych, "BEIR: A Heterogenous Benchmark for Zero-shot Evaluation of Information Retrieval Models," 2021, [Online]. Available:

- <http://arxiv.org/abs/2104.08663>.
- [15] T. Kenter and M. De Rijke, "Short text similarity with word embeddings," in *International Conference on Information and Knowledge Management, Proceedings*, 2015, vol. 19-23-Oct-, pp. 1411–1420, doi: [10.1145/2806416.2806475](https://doi.org/10.1145/2806416.2806475).
- [16] M. Johnson *et al.*, "Google's multilingual neural machine translation system: Enabling zero-shot translation," *Trans. Assoc. Comput. Linguist.*, vol. 5, pp. 339–351, 2017.
- [17] O. Firat, K. Cho, B. Sankaran, F. T. Yarman Vural, and Y. Bengio, "Multi-way, multilingual neural machine translation," *Comput. Speech Lang.*, vol. 45, pp. 236–252, 2017, doi: [10.1016/j.csl.2016.10.006](https://doi.org/10.1016/j.csl.2016.10.006).
- [18] J. Li *et al.*, "Performance Bug Analysis and Detection for Distributed Storage and Computing Systems," *ACM Trans. Storage*, vol. 19, no. 3, pp. 1–33, 2023, doi: [10.1145/3580281](https://doi.org/10.1145/3580281).
- [19] Y. Luo *et al.*, "Characterizing application memory error vulnerability to optimize datacenter cost via heterogeneous-reliability memory," in *Proceedings of the International Conference on Dependable Systems and Networks*, 2014, pp. 467–478, doi: [10.1109/DSN.2014.50](https://doi.org/10.1109/DSN.2014.50).
- [20] P. Christen, *The Data Matching Process*. Springer, 2012.
- [21] A. K. Elmagarmid, P. G. Ipeirotis, and V. S. Verykios, "Duplicate record detection: A survey," *IEEE Trans. Knowl. Data Eng.*, vol. 19, no. 1, pp. 1–16, 2007, doi: [10.1109/TKDE.2007.250581](https://doi.org/10.1109/TKDE.2007.250581).
- [22] J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," *NAACL HLT 2019 - 2019 Conf. North Am. Chapter Assoc. Comput. Linguist. Hum. Lang. Technol. - Proc. Conf.*, vol. 1, pp. 4171–4186, 2019.
- [23] Y. Zhao, A. Zhang, R. Xie, K. Liu, and X. Wang, "Connecting embeddings for knowledge graph entity typing," *Proc. Annu. Meet. Assoc. Comput. Linguist.*, pp. 6419–6428, 2020, doi: [10.18653/v1/2020.acl-main.572](https://doi.org/10.18653/v1/2020.acl-main.572).
- [24] R. Rehurek and P. Sojka, "Software Framework for Topic Modelling with Large Corpora," *Proc. Lr. 2010 Work. New Challenges NLP Fram.*, pp. 45–50, 2010.
- [25] T. Pires, E. Schlinger, and D. Garrette, "How multilingual is multilingual BERT?," *ACL 2019 - 57th Annu. Meet. Assoc. Comput. Linguist. Proc. Conf.*, pp. 4996–5001, 2020, doi: [10.18653/v1/p19-1493](https://doi.org/10.18653/v1/p19-1493).
- [26] G. Doddington, "Automatic evaluation of machine translation quality using n-gram co-occurrence statistics," in *Proceedings of the second international conference on Human Language Technology Research*, 2002, p. 138, doi: [10.3115/1289189.1289273](https://doi.org/10.3115/1289189.1289273).
- [27] L. Cardozo-Gaibisso, G. W. Hodges, C. Mardones-Segovia, and A. S. Cohen, "Multidimensional Assessment Performance Analysis: A Framework to Advance Multilingual Learners' Scientific Equity in K-12 Contexts," *Education Sciences*, vol. 14, no. 10, 2024, doi: [10.3390/educsci14101068](https://doi.org/10.3390/educsci14101068).
- [28] X. J. A. Bellekens, C. Tachtatzis, R. C. Atkinson, C. Renfrew, and T. Kirkham, "A highly-efficient memory-compression scheme for GPU-accelerated Intrusion Detection Systems," in *ACM International Conference Proceeding Series*, 2014, vol. 2014-Septe, pp. 302–309, doi: [10.1145/2659651.2659723](https://doi.org/10.1145/2659651.2659723).
- [29] Z. K. Baker and V. K. Prasanna, "Time and area efficient pattern matching on FPGAs," in *ACM/SIGDA International Symposium on Field Programmable Gate Arrays - FPGA*, 2004, vol. 12, pp. 223–232, doi: [10.1145/968280.968312](https://doi.org/10.1145/968280.968312).
- [30] M. M. Baig, S. Sivakumar, and S. R. Nayak, "Optimizing Performance of Text Searching Using CPU and GPUs," in *Advances in Intelligent Systems and Computing*, 2020, vol. 1119, pp. 141–150, doi: [10.1007/978-981-15-2414-1_15](https://doi.org/10.1007/978-981-15-2414-1_15).
- [31] H. Zhou *et al.*, "The Efficiency vs. Accuracy Trade-off: Optimizing RAG-Enhanced LLM Recommender Systems Using Multi-Head Early Exit," *arXiv Prepr. arXiv2501.02173*, 2025.
- [32] K. I. Kim, K. Jung, and J. H. Kim, "Texture-Based Approach for Text Detection in Images Using Support Vector Machines and Continuously Adaptive Mean Shift Algorithm," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 12, pp. 1631–1639, 2003, doi: [10.1109/TPAMI.2003.1251157](https://doi.org/10.1109/TPAMI.2003.1251157).
- [33] A. Fan, S. Wang, and Y. Wang, "Legal Document Similarity Matching Based on Ensemble Learning," *IEEE Access*, vol. 12, pp. 33910–33922, 2024, doi: [10.1109/ACCESS.2024.3371262](https://doi.org/10.1109/ACCESS.2024.3371262).
- [34] B. Wang, S. Yu, W. Lou, and Y. T. Hou, "Privacy-preserving multi-keyword fuzzy search over encrypted data in the cloud," in *Proceedings - IEEE INFOCOM*, 2014, pp. 2112–2120, doi: [10.1109/INFOCOM.2014.6848153](https://doi.org/10.1109/INFOCOM.2014.6848153).
- [35] H. Schütze, C. D. Manning, and P. Raghavan, *Introduction to information retrieval*, vol. 39. Cambridge University Press Cambridge, 2008.
- [36] M. Rashmi, *Introduction to Information Retrieval Systems*, vol. 3, no. 4. Cambridge University Press Cambridge, 2015.
- [37] M. T. Pilehvar and J. Camacho-Collados, *Embeddings in Natural Language Processing: Theory and Advances in Vector Representations of Meaning*, vol. 13, no. 4. Morgan & Claypool Publishers, 2020.
- [38] W. R. Pearson, "Selecting the right similarity-scoring matrix," *Curr. Protoc. Bioinforma.*, vol. 43, no. 1, pp. 3–5, 2013.
- [39] B. Berger, M. S. Waterman, and Y. W. Yu, "Levenshtein Distance, Sequence Comparison and Biological Database Search," *IEEE Trans. Inf. Theory*, vol. 67, no. 6, pp. 3287–3294, 2021, doi: [10.1109/TIT.2020.2996543](https://doi.org/10.1109/TIT.2020.2996543).
- [40] N. Gali, R. Marinescu-Istodor, D. Hostettler, and P. Fränti, "Framework for syntactic string similarity measures," *Expert Syst. Appl.*, vol. 129, pp. 169–185, 2019, doi: [10.1016/j.eswa.2019.03.048](https://doi.org/10.1016/j.eswa.2019.03.048).
- [41] B. C. Gencosman, H. C. Ozmutlu, and S. Ozmutlu, "Character n-gram application for automatic new topic identification," *Inf. Process. Manag.*, vol. 50, no. 6, pp. 821–856, 2014, doi: [10.1016/j.ipm.2014.06.005](https://doi.org/10.1016/j.ipm.2014.06.005).
- [42] V. I. Levenshtein, "Efficient reconstruction of sequences," *IEEE Trans. Inf. Theory*, vol. 47, no. 1, pp. 2–22, 2001, doi: [10.1109/18.904499](https://doi.org/10.1109/18.904499).
- [43] P. Choudhury, Z. Ahmed, and B. K. Sunitha, "Analyzing Intelligent Optimization Techniques for 6G Radio Resource Allocation," in *2024 15th International Conference on Computing Communication and Networking Technologies, ICCCNT 2024*, 2024, pp. 1–6, doi: [10.1109/ICCCNT61001.2024.10725020](https://doi.org/10.1109/ICCCNT61001.2024.10725020).
- [44] J. Wang, Y. Ma, L. Zhang, R. X. Gao, and D. Wu, "Deep learning for smart manufacturing: Methods and applications," *J. Manuf. Syst.*, vol. 48, pp. 144–156, 2018, doi: [10.1016/j.jmsy.2018.01.003](https://doi.org/10.1016/j.jmsy.2018.01.003).
- [45] A. Spielberg *et al.*, "Differentiable visual computing for inverse problems and machine learning," *Nat. Mach. Intell.*, vol. 5, no. 11, pp. 1189–1199, 2023, doi: [10.1038/s42256-023-00743-0](https://doi.org/10.1038/s42256-023-00743-0).
- [46] Y. Chaabi and F. Ataa Allah, "Amazigh spell checker using Damerau-Levenshtein algorithm and N-gram," *J. King Saud Univ. - Comput. Inf. Sci.*, vol. 34, no. 8, pp. 6116–6124, 2022, doi: [10.1016/j.jksuci.2021.07.015](https://doi.org/10.1016/j.jksuci.2021.07.015).
- [47] D. Hládek, J. Staš, and M. Pleva, "Survey of automatic spelling correction," *Electron.*, vol. 9, no. 10, pp. 1–29, 2020, doi: [10.3390/electronics9101670](https://doi.org/10.3390/electronics9101670).
- [48] H. Al-Rubaiee, R. Qiu, and D. Li, "Identifying Mubasher software products through sentiment analysis of Arabic tweets," in *2016 International Conference on Industrial Informatics and Computer Systems, CIICS 2016*, 2016, pp. 1–6, doi: [10.1109/ICCSII.2016.7462396](https://doi.org/10.1109/ICCSII.2016.7462396).
- [49] Z. Huang and W. Zhao, "A semantic matching approach

- addressing multidimensional representations for web service discovery," *Expert Syst. Appl.*, vol. 210, p. 118468, 2022, doi: [10.1016/j.eswa.2022.118468](https://doi.org/10.1016/j.eswa.2022.118468).
- [50] A. A. Niaz, R. Ashraf, T. Mahmood, C. M. N. Faisal, and M. M. Abid, "An efficient smart phone application for wheat crop diseases detection using advanced machine learning," *PLoS One*, vol. 20, no. 1 January, p. e0312768, 2025, doi: [10.1371/journal.pone.0312768](https://doi.org/10.1371/journal.pone.0312768).
- [51] M. Rashmi, "Introduction to Information Retrieval Systems," *International Journal on Recent and Innovation Trends in Computing and Communication*, vol. 3, no. 4. Wiley Online Library, pp. 2051–2054, 2015, doi: [10.17762/ijritcc.2321-8169.150462](https://doi.org/10.17762/ijritcc.2321-8169.150462).
- [52] B. MacCartney, M. Galley, and C. D. Manning, "A phrase-based alignment model for natural language inference," in *EMNLP 2008 - 2008 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference: A Meeting of SIGDAT, a Special Interest Group of the ACL*, 2008, pp. 802–811, doi: [10.3115/1613715.1613817](https://doi.org/10.3115/1613715.1613817).
- [53] C. C. Aggarwal and C. X. Zhai, "A survey of text classification algorithms," *Min. Text Data*, vol. 9781461432, pp. 163–222, 2012, doi: [10.1007/978-1-4614-3223-4_6](https://doi.org/10.1007/978-1-4614-3223-4_6).
- [54] W. H. Gomaa and A. A. Fahmy, "A Survey of Text Similarity Approaches," *Int. J. Comput. Appl.*, vol. 68, no. 13, pp. 13–18, 2013, doi: [10.5120/11638-7118](https://doi.org/10.5120/11638-7118).
- [55] J. Wang and Y. Dong, "Measurement of text similarity: A survey," *Inf.*, vol. 11, no. 9, pp. 1–17, 2020, doi: [10.3390/info11090421](https://doi.org/10.3390/info11090421).
- [56] A. Islam and D. Inkpen, "Semantic Text Similarity Using Corpus-Based Word Similarity and String Similarity," *ACM Trans. Knowl. Discov. Data*, vol. 2, no. 2, pp. 1–25, 2008, doi: [10.1145/1376815.1376819](https://doi.org/10.1145/1376815.1376819).
- [57] B. Wang and C. C. J. Kuo, "SBERT-WK: A Sentence Embedding Method by Dissecting BERT-Based Word Models," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 28, pp. 2146–2157, 2020, doi: [10.1109/TASLP.2020.3008390](https://doi.org/10.1109/TASLP.2020.3008390).
- [58] S. Abarna, J. I. Sheeba, and S. P. Devaneyan, "An ensemble model for idioms and literal text classification using knowledge-enabled BERT in deep learning," *Meas. Sensors*, vol. 24, no. 1, p. 102756, 2022, doi: [10.1016/j.measen.2022.100434](https://doi.org/10.1016/j.measen.2022.100434).
- [59] Y. Bounab, M. Oussalah, and A. Ferdenache, "Reconciling Image Captioning and User's Comments for Urban Tourism," in *2020 10th International Conference on Image Processing Theory, Tools and Applications, IPTA 2020*, 2020, pp. 1–6, doi: [10.1109/IPTA50016.2020.9286602](https://doi.org/10.1109/IPTA50016.2020.9286602).
- [60] S. Behmanesh, A. Talebpour, M. Shamsfard, and M. M. Jafari, "A Novel Open-Domain Question Answering System on Curated and Extracted Knowledge Bases with Consideration of Confidence Scores in Existing Triples," *IEEE Access*, 2024, doi: [10.1109/ACCESS.2024.3490452](https://doi.org/10.1109/ACCESS.2024.3490452).
- [61] Y. Li and Y. Long, "Inferring storefront vacancy using mobile sensing images and computer vision approaches," *Comput. Environ. Urban Syst.*, vol. 108, p. 102071, 2024, doi: [10.1016/j.compenvurbsys.2023.102071](https://doi.org/10.1016/j.compenvurbsys.2023.102071).
- [62] I. P. A. E. D. Udayana, I. G. T. A. Putra, I. P. S. Udyana, I. G. S. C. Nugraha, N. P. Widantari, and B. K. Wijaya, "Optimizing Latin to Balinese Script Transliteration: Hybrid Jaro Winkler and Damerau Levenshtein Methods," in *Digest of Technical Papers - IEEE International Conference on Consumer Electronics*, 2024, pp. 12–18, doi: [10.1109/ISCT62336.2024.10791083](https://doi.org/10.1109/ISCT62336.2024.10791083).
- [63] G. Alemu, B. Stevens, and P. Ross, "Towards a conceptual framework for user-driven semantic metadata interoperability in digital libraries: A social constructivist approach," *New Libr. World*, vol. 113, no. 1, pp. 38–54, 2012, doi: [10.1108/03074801211199031](https://doi.org/10.1108/03074801211199031).
- [64] M. Camacho and E. Navarro, *Natural language processing with Python*, vol. 4, no. 13. Frank Millstein, 2020.
- [65] N. Van Otte, "Word Embedding A Powerful Tool — How To Use Word2Vec GloVe, FastText," *spotintelligence*, 2022. <https://spotintelligence.com/2022/11/30/word-embedding/?form=MG0AV3>.
- [66] P. McNamee and J. Mayfield, "Character n-gram tokenization for European language text retrieval," *Inf. Retr. Boston.*, vol. 7, no. 1–2, pp. 73–97, 2004, doi: [10.1023/b:inrt.0000009441.78971.be](https://doi.org/10.1023/b:inrt.0000009441.78971.be).
- [67] Z. Jiang, A. El-Jaroudi, W. Hartmann, D. Karakos, and L. Zhao, "Cross-lingual Information Retrieval with BERT," *arXiv Prepr. arXiv2004.13005*, 2020, [Online]. Available: <http://arxiv.org/abs/2004.13005>.
- [68] T. Schuster, O. Ram, R. Barzilay, and A. Globerson, "Cross-lingual alignment of contextual word embeddings, with applications to zero-shot dependency parsing," *NAACL HLT 2019 - 2019 Conf. North Am. Chapter Assoc. Comput. Linguist. Hum. Lang. Technol. - Proc. Conf.*, vol. 1, pp. 1599–1613, 2019, doi: [10.18653/v1/n19-1162](https://doi.org/10.18653/v1/n19-1162).
- [69] V. Igloukov, "Need for Speed: A Comprehensive Benchmark of JPEG Decoders in Python," *arXiv Prepr. arXiv2501.13131*, 2025.
- [70] H. Li *et al.*, "A novel locality-sensitive hashing relational graph matching network for semantic textual similarity measurement," *Expert Syst. Appl.*, vol. 207, p. 117832, 2022, doi: [10.1016/j.eswa.2022.117832](https://doi.org/10.1016/j.eswa.2022.117832).
- [71] K. Kim, M. K. Hasan, J. P. Heo, Y. W. Tai, and S. E. Yoon, "Probabilistic cost model for nearest neighbor search in image retrieval," *Comput. Vis. Image Underst.*, vol. 116, no. 9, pp. 991–998, 2012, doi: [10.1016/j.cviu.2012.05.001](https://doi.org/10.1016/j.cviu.2012.05.001).
- [72] F. Chen and M. Hsu, "A performance comparison of parallel DBMSs and MapReduce on large-scale text analytics," in *ACM International Conference Proceeding Series*, 2013, pp. 613–624, doi: [10.1145/2452376.2452448](https://doi.org/10.1145/2452376.2452448).
- [73] V. Gupta, M. Gupta, J. Garg, and N. Garg, "Improvement in semantic address matching using natural language processing," in *2021 2nd International Conference for Emerging Technology, INCET 2021*, 2021, pp. 1–5, doi: [10.1109/INCET51464.2021.9456342](https://doi.org/10.1109/INCET51464.2021.9456342).
- [74] J. Lin, X. Ma, S. C. Lin, J. H. Yang, R. Pradeep, and R. Nogueira, "Pyserini: A Python Toolkit for Reproducible Information Retrieval Research with Sparse and Dense Representations," in *SIGIR 2021 - Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2021, pp. 2356–2362, doi: [10.1145/3404835.3463238](https://doi.org/10.1145/3404835.3463238).
- [75] B. P. Miller *et al.*, "The Paradyn Parallel Performance Measurement Tool," *Computer (Long Beach, Calif.)*, vol. 28, no. 11, pp. 37–46, 1995, doi: [10.1109/2.471178](https://doi.org/10.1109/2.471178).
- [76] T. Nawaz and A. Cavallaro, "A protocol for evaluating video trackers under real-world conditions," *IEEE Trans. Image Process.*, vol. 22, no. 4, pp. 1354–1361, 2013, doi: [10.1109/TIP.2012.2228497](https://doi.org/10.1109/TIP.2012.2228497).
- [77] S. Patel and R. Patel, "A Comprehensive Analysis of Computing Paradigms Leading to Fog Computing: Simulation Tools, Applications, and Use Cases," *J. Comput. Inf. Syst.*, vol. 63, no. 6, pp. 1495–1516, 2023, doi: [10.1080/08874417.2022.2121782](https://doi.org/10.1080/08874417.2022.2121782).
- [78] J. Lüttgau *et al.*, "Survey of storage systems for high-performance computing," *Supercomput. Front. Innov.*, vol. 5, no. 1, pp. 31–58, 2018, doi: [10.14529/jfsfi180103](https://doi.org/10.14529/jfsfi180103).
- [79] H. J. Kim and J. S. Kim, "A user-space storage I/O framework for NVMe SSDs in mobile smart devices," *IEEE Trans. Consum. Electron.*, vol. 63, no. 1, pp. 28–35, 2017, doi: [10.1109/TCE.2017.014709](https://doi.org/10.1109/TCE.2017.014709).
- [80] M. H. I. Chowdhury, M. Jung, F. Yao, and A. Awad, "D-Shield: Enabling Processor-side Encryption and Integrity Verification

- for Secure NVMe Drives,” in *Proceedings - International Symposium on High-Performance Computer Architecture*, 2023, vol. 2023-Febru, pp. 908–921, doi: [10.1109/HPCA56546.2023.10070924](https://doi.org/10.1109/HPCA56546.2023.10070924).
- [81] J. Turnbull, *Monitoring With Prometheus Website: Monitoring With Prometheus*. Turnbull Press, 2019.
- [82] Y. Li, H. Qi, G. Lu, F. Jin, Y. Guo, and X. Lu, “Understanding hot interconnects with an extensive benchmark survey,” *BenchCouncil Trans. Benchmarks, Stand. Eval.*, vol. 2, no. 3, p. 100074, 2022, doi: [10.1016/j.tbench.2022.100074](https://doi.org/10.1016/j.tbench.2022.100074).
- [83] M. Liebowitz, C. Kusek, and R. Spies, *VMware vSphere Performance: Designing CPU, Memory, Storage, and Networking for Performance-Intensive Workloads*. John Wiley & Sons, 2014.
- [84] H. J. Escalante *et al.*, “Term-weighting learning via genetic programming for text classification,” *Knowledge-Based Syst.*, vol. 83, no. 1, pp. 176–189, 2015, doi: [10.1016/j.knosys.2015.03.025](https://doi.org/10.1016/j.knosys.2015.03.025).
- [85] N. F. Schneidewind, “Methodology for Validating Software Metrics,” *IEEE Trans. Softw. Eng.*, vol. 18, no. 5, pp. 410–422, 1992, doi: [10.1109/32.135774](https://doi.org/10.1109/32.135774).
- [86] R. McCleary, S. Patterson, and J. Yates, “Quality Validation Method,” *freepatentsonline*, 2008. www.freepatentsonline.com/y2009/0169092.html.
- [87] G. Navarro, “A guided tour to approximate string matching,” *ACM Comput. Surv.*, vol. 33, no. 1, pp. 31–88, 2001, doi: [10.1145/375360.375365](https://doi.org/10.1145/375360.375365).
- [88] W. Cohen, P. Ravikumar, and S. Fienberg, “A comparison of string metrics for matching names and records,” in *Kdd workshop on data cleaning and object consolidation*, 2003, vol. 3, pp. 73–78.
- [89] V. S. Verykios, A. K. Elmagarmid, and E. N. Houstis, “Automating the approximate record-matching process,” *Inf. Sci. (Ny.)*, vol. 126, no. 1, pp. 83–98, 2000, doi: [10.1016/S0020-0255\(00\)00013-X](https://doi.org/10.1016/S0020-0255(00)00013-X).
- [90] Y. Zhang and S. Vogel, “Significance tests of automatic machine translation evaluation metrics,” *Mach. Transl.*, vol. 24, no. 1, pp. 51–65, 2010, doi: [10.1007/s10590-010-9073-6](https://doi.org/10.1007/s10590-010-9073-6).
- [91] R. A. Wagner and M. J. Fischer, “The String-to-String Correction Problem,” *J. ACM*, vol. 21, no. 1, pp. 168–173, 1974, doi: [10.1145/321796.321811](https://doi.org/10.1145/321796.321811).