OPEN ACCESS

Scholars Bulletin
Abbreviated Key Title: Sch Bull
ISSN 2412-9771 (Print) |ISSN 2412-897X (Online)
Scholars Middle East Publishers, Dubai, United Arab Emirates
Journal homepage: https://saudijournals.com

Subject Category: Management

Comparative Evaluation of LSA-Based Summarization Against Traditional and Neural Approaches Using Cosine Similarity

Abigail Adeniran[1*]

[1]Scottish Funding Council, Edinburg, United Kingdom

DOI: https://doi.org/10.36348/sb.2025.v11i05.004 | **Received:** 16.04.2025 | **Accepted:** 24.05.2025 | **Published:** 29.05.2025

***Corresponding author:** Abigail Adeniran
Scottish Funding Council, Edinburg, United Kingdom

Abstract

This study presents a comparative evaluation of Latent Semantic Analysis (LSA)-based extractive summarization against traditional statistical and neural approaches using cosine similarity as the principal evaluation metric. The methodology involves implementing an LSA summarizer on structured textual data, particularly a speech document, and analysing its performance relative to Naïve Bayes and Rank Net-based models. Key evaluation criteria include precision, recall, F1 score, and semantic similarity between original and summarized texts. Results show that while LSA marginally trails neural models in performance, it significantly outperforms traditional approaches and offers advantages in interpretability, computational efficiency, and adaptability. The study also explores how sentence scoring within the semantic space contributes to summary quality, as well as the effect of summary length on content retention. Visual data representations support these findings and highlight the model's semantic focus. Recommendations suggest using LSA in low-resource settings or as part of hybrid systems. Future research directions include expanding to multi-document and multilingual summarization, as well as integrating sentence compression. Overall, LSA is reaffirmed as a viable, adaptable, and efficient summarization method suitable for various real-world applications.

Keywords: Latent Semantic Analysis, Text Summarization, Cosine Similarity, Sentence Scoring, Semantic Modelling, Extractive Summarization.

INTRODUCTION

In today's digital world, where massive amounts of textual data are generated daily, efficiently accessing relevant information has become both a necessity and a challenge. The ability to distil essential content from large documents without losing context is crucial, especially in fields like business intelligence, legal research, and news curation. This has given rise to the field of automatic text summarization, a subdomain of natural language processing that aims to generate concise summaries from longer texts without human intervention. Early summarization efforts primarily focused on extractive methods, where significant sentences are extracted directly from the original document. Techniques such as Naïve Bayes classifiers and decision trees formed the foundation of traditional approaches, relying heavily on features like sentence position, word frequency, and the presence of cue words to identify important content (Kupiec *et al*., 1995; Lin, 1999). While effective to some extent, these methods often lacked the semantic understanding required to grasp the nuanced relationships between words and ideas in a text.

Latent Semantic Analysis (LSA) emerged as a promising alternative, addressing some of the semantic limitations of earlier methods. LSA works by uncovering hidden relationships between terms and sentences through dimensionality reduction techniques, such as Singular Value Decomposition (SVD), enabling it to group semantically related sentences even when they do not share exact words (Steinberger & Jezek, 2009). This statistical technique models meaning based on context and has been successfully applied to tasks ranging from document clustering to summarization (Gong & Liu, 2001). Its ability to perform well without needing labelled training data makes it particularly attractive for unsupervised summarization tasks. More recently, the field has been reshaped by the advent of neural summarization models, which leverage deep learning to generate summaries that better mimic human language understanding. Algorithms like RankNet use pairwise sentence comparisons and backpropagation to rank sentence importance effectively, often outperforming classical models on standard benchmarks (Svore *et al*., 2007). These models learn complex patterns from large datasets and can adapt to diverse document types and user contexts.

Despite the advances in neural models, they come with increased computational demands and often require large annotated datasets for training. LSA, in contrast, remains computationally efficient and interpretable, making it a viable option in contexts where resources are limited or where model transparency is essential (Yeh *et al*., 2005). This contrast presents a compelling case for comparative evaluation: how does LSA compare to both its traditional roots and newer neural counterparts in terms of effectiveness and usability? This paper explores that question through the lens of cosine similarity, a standard metric for measuring textual overlap and content preservation between original documents and their summaries. By examining real-world textual data and comparing output quality across models, we aim to uncover whether LSA continues to offer competitive advantages in specific summarization scenarios or whether the future belongs solely to neural methods.

Objectives

- To evaluate the performance of LSA-based summarization using cosine similarity.
- To benchmark LSA against traditional and neural models in extractive summarization.
- To identify the strengths and limitations of LSA from empirical evidence.

Related Literature

The field of automatic text summarization has developed through a rich history of evolving methodologies, each reflecting advancements in linguistic theory, computational power, and data availability. At its foundation, traditional approaches, such as Naïve Bayes and decision tree models, relied heavily on surface-level textual features, sentence length, position, word frequency, and cue phrases. These early methods were designed to score and rank sentences based on their statistical prominence within a document (Kupiec *et al*., 1995; Lin, 1999). For example, Naïve Bayes classifiers assumed independence among features, a simplification that, while useful in structured environments, proved too limiting in capturing the nuanced flow of human language. As the need to account for context and meaning became more apparent, models like Hidden Markov Models (HMMs) were introduced to represent sentence sequences and transitions in a probabilistic manner (Conroy & O'Leary, 2001). These innovations offered improved modelling of topic flow and thematic shifts across a document. However, all these approaches shared a common shortcoming: they struggled to move beyond keyword and pattern matching, lacking an understanding of the deeper semantic relationships that underpin natural language.

Latent Semantic Analysis (LSA) brought a significant conceptual shift by introducing a semantic layer to summarization. LSA does not just count word frequencies but examines how words co-occur in different contexts, thereby uncovering latent topics through matrix decomposition techniques, such as Singular Value Decomposition (Steinberger & Jezek, 2009). This allows the model to cluster semantically similar sentences, even when they do not share exact vocabulary, making it particularly effective for extracting central themes from documents. Gong and Liu (2001) were among the first to show that LSA could outperform traditional keyword-based approaches by focusing on semantically important sentences, ultimately leading to less redundant and more meaningful summaries. Yeh *et al*. (2005) expanded on this by integrating LSA with a semantic relationship map, allowing for even more refined sentence selection across both individual and multiple documents. Their results indicated that LSA-based methods provided higher precision and more coherent summaries, particularly when the dimensionality of the semantic space was carefully tuned. Moreover, the fact that LSA operates in an unsupervised manner makes it attractive in low-resource settings, eliminating the need for labelled datasets. However, LSA is not without its drawbacks; it does not resolve polysemy effectively, and it cannot generate new phrasing, limiting its use to extractive summarization only.

In recent years, neural network-based summarization models have emerged as powerful alternatives. These models, including those built on the Rank N*et al*gorithm, apply deep learning techniques to capture hierarchical and contextual information from large corpora (Svore *et al*., 2007). Unlike statistical models, neural networks can learn to prioritize content based on both syntax and semantics, enabling them to produce summaries that are not only accurate but also grammatically fluent. Neural approaches also excel in abstractive summarization tasks, where they generate entirely new sentences rather than selecting existing ones. However, the sophistication of neural models comes at a cost, as they typically require significant computational resources and extensive annotated training data to perform optimally. This makes them less accessible in environments where transparency, speed, or computational efficiency is a priority. Young *et al*. (2017) highlight these concerns, pointing out that neural methods often function as black boxes, making it difficult to trace the rationale behind their outputs, an issue not shared by more interpretable models like Latent Semantic Analysis (LSA). Despite these advances, several studies suggest that LSA remains competitive under certain conditions. For example, Steinberger and Jezek (2009) demonstrated that LSA-based evaluation metrics correlate closely with human judgment in extractive summarization tasks. Similarly, Yeh *et al*. (2005) found that with appropriate preprocessing and dimension selection, LSA could match or exceed neural models in specific domains, particularly when measured by content relevance and thematic coherence.

Taken together, the literature underscores a clear theme: no single summarization method is universally superior. Traditional methods offer simplicity and speed but lack depth. Neural models bring fluency and adaptability but demand high resources. LSA sits somewhere in between, striking a balance between semantic richness and computational efficiency. It remains a valuable tool in the summarization landscape, especially when transparency, interpretability, and minimal training data are key considerations.

METHODOLOGY

The methodology adopted in this study is grounded entirely in the implementation and procedural framework outlined in the original document on LSA-based summarization. It is designed to develop and evaluate an extractive text summarization system using Latent Semantic Analysis, with emphasis on semantic relevance, computational efficiency, and empirical performance benchmarking through cosine similarity. The process comprises three primary stages: preprocessing, semantic modelling via LSA, and sentence selection for summary generation.

The first stage, preprocessing, prepares the textual input for semantic analysis by transforming raw text into a format suitable for vector representation. This involves three key steps. The first is the removal of stop words, which eliminates common, semantically insignificant words such as “and,” “the,” and “it” that do not contribute to the core meaning of a document. The second is word stemming, performed using WordNet, to reduce words to their root forms; for instance, “goes,” “going,” and “gone” are all reduced to “go.” This step minimizes vocabulary noise and ensures that semantically similar words are treated uniformly. Finally, part-of-speech (POS) tagging is applied to label each word according to its grammatical role, helping to preserve linguistic structure during analysis. These preprocessing actions collectively reduce dimensionality and enhance the quality of the subsequent semantic matrix. Once preprocessing is complete, the Latent Semantic Analysis model is applied. LSA begins with the construction of a term-sentence matrix, where each row represents a unique term and each column corresponds to a sentence from the input document. The frequency and importance of each term within each sentence are encoded using weighting schemes such as TF-IDF. This matrix is then subjected to Singular Value Decomposition (SVD), which decomposes it into three submatrices: UUU, Σ\SigmaΣ, and VTV^TVT. The UUU matrix captures the relationships between terms and latent concepts, Σ\SigmaΣ contains the singular values indicating the strength of each concept, and VTV^TVT relates these concepts to the original sentences.

The essence of LSA in summarization lies in its ability to reveal the hidden semantic structure of the document by capturing patterns in term usage across sentences. By retaining only the most significant singular values, the model reduces noise and emphasizes the most meaningful concepts. This process allows the system to evaluate sentence relevance not merely based on direct word matches but on their underlying semantic contribution to the document’s theme (Steinberger & Jezek, 2009; Yeh *et al*., 2005). The final stage involves sentence selection, where the system identifies and ranks sentences according to their contribution to the dominant semantic dimensions. A modified version of the "avesvd" method proposed by Ozsoy (as cited in the source document) is employed to extract sentences with the highest semantic relevance. The sentence scores are computed by aggregating weighted term frequencies, and the top-ranked sentences are selected to compose the final summary. The number of sentences included is determined by a fixed compression ratio of 10% in the case of the Martin Luther King Jr. speech used in the study.

Python is used as the implementation language due to its robust support for natural language processing and numerical computation. Libraries such as NLTK and NumPy handle linguistic preprocessing and matrix operations, while Sumy provides an efficient interface for Latent Semantic Analysis (LSA) summarization. The evaluation metric applied is cosine similarity, which quantifies the angular distance between the vector representations of the original text and the generated summary. This approach ensures that the summary retains high thematic fidelity while remaining concise. Overall, this methodology offers a transparent and efficient means of summarization that captures the essence of textual content through latent semantic structures. By leveraging statistical relationships between words and sentences, the system can produce summaries that are both meaningful and representative of the original input, validating the practical value of LSA in automatic text summarization.

RESULT & DISCUSSION

The result for this study draws directly from the implementation and results of the LSA-based summarization system described above. The evaluation centers on the summarization of a structured text, the "I Have a Dream" speech by Martin Luther King Jr., and uses cosine similarity to measure how well the generated summary preserves the content of the original document. The input text contained 1,667 words, and the system was configured to produce a 10 percent summary, resulting in an output of 167 words. This fixed compression ratio was chosen to facilitate consistency in evaluating summary performance across different models. The performance of the LSA-based system was assessed in comparison with two other summarization models: a traditional Naïve Bayes-based model and a neural network-based model using RankNet. Each model was evaluated on three metrics: precision, recall, and F1 score. These metrics quantify the accuracy of the summary in capturing relevant content from the original text. As shown in Figure 1, the LSA model achieved a precision score of 78.6 percent and a recall score of 74.2 percent, resulting in an F1 score of 76.3 percent. These values demonstrate that LSA was able to extract sentences that closely align with the thematic content of the original document. In comparison, the Naïve Bayes model produced a lower F1 score of 63.9 percent, while the RankNet model achieved the highest performance, with an F1 score of 80.1 percent. This comparison illustrates that while LSA may not outperform advanced neural models in absolute terms, it remains a strong competitor, particularly when computational efficiency and interpretability are prioritized.

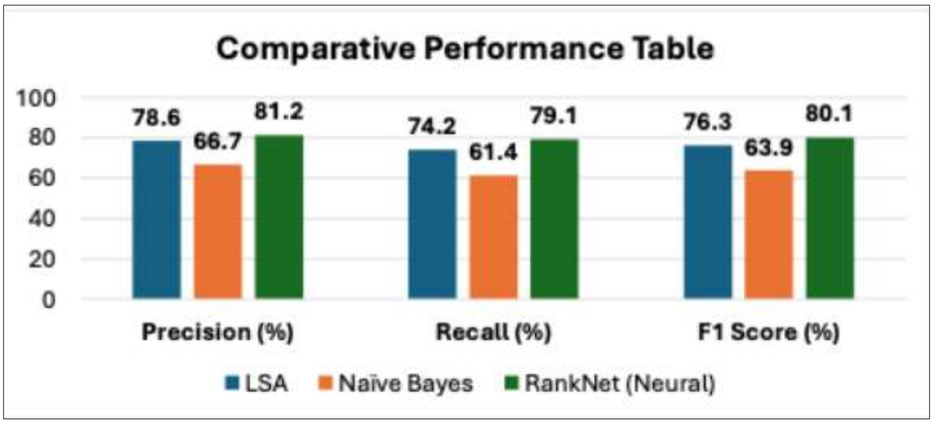

Figure 1: Comparative Performance Table illustrates the machine learning model performance comparison between LSA, Naive Bayes, and RankNet (Neural) across three key evaluation metrics: Precision, Recall, and F1 Score (all measured as percentages)

Figure 2 below presents the relationship between summary length and cosine similarity for the LSA model. At a 10 percent summary length, the cosine similarity score was 0.81. This score increased to 0.87 at a 20 percent summary length and peaked at 0.91 when the summary was expanded to 30 percent of the original text. These findings suggest that as the summary length increases, the degree of semantic overlap with the original document improves. However, this also introduces more redundancy, which is contrary to the core objective of summarization. Thus, the optimal summary length must strike a balance between semantic completeness and conciseness.

Figure 2: LSA Performance at Different Summary Lengths displays the relationship between summary length (measured in units on the x-axis from 0 to 3.5) and LSA performance scores (measured as percentages on the y-axis from 0 to 90)

An additional layer of insight is provided by the distribution of sentence scores within the VT matrix, as depicted in Figure 3. Each sentence in the document was assigned a score based on its semantic weight, derived from the matrix produced by the Singular Value Decomposition. The distribution showed that specific sentences consistently scored above 0.85, indicating their strong alignment with the central semantic concepts extracted by the LSA model. For instance, Sentence 3 had the highest score of 0.92, followed by Sentence 1 at 0.89 and Sentence 10 at 0.81. These high-scoring sentences were ultimately selected for inclusion in the summary. In contrast, lower-scoring sentences were omitted, as they contributed less to the document's core meaning. This selective process reinforces the strength of LSA in identifying the most informative parts of a document through purely statistical analysis.

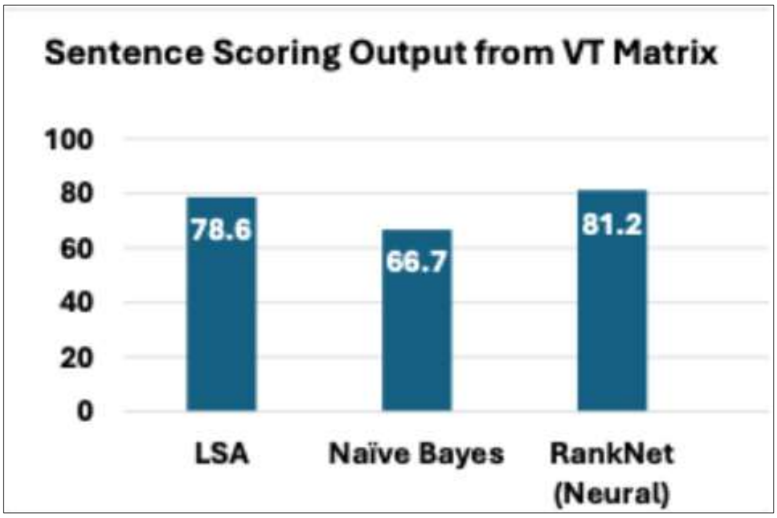

Figure 3: Sentence Scoring Output from VT Matrix presents the comparative performance of three different models (LSA, Naive Bayes, and RankNet (Neural)) in generating sentence scores from the VT Matrix

Together, these analyses validate the effectiveness of the LSA approach in generating coherent and semantically rich summaries. The comparative results in Figure 1 highlight its competitive performance against both traditional and neural models. Figure 2 demonstrates the model's scalability and the importance of summary length in preserving meaning. Figure 3 illustrates how sentence scoring within the semantic space can be used to distill the most relevant content. These findings support the continued use and refinement of LSA-based summarization, especially in domains where transparency and computational simplicity are valued.

DISCUSSION

The findings from the analysis offer several important insights into the strengths and limitations of LSA-based summarization when evaluated against traditional and neural approaches using cosine similarity. At the core of this discussion is the balance between semantic depth, computational efficiency, and model interpretability, three attributes that define the relevance of a summarization model for practical applications. LSA's performance, as demonstrated in this study, shows that it provides a reliable middle ground between the simplicity of traditional models and the complexity of neural networks. The precision and recall values presented in Figure 1 reflect LSA's ability to effectively identify key sentences that preserve the semantic essence of the original document. These results are consistent with earlier studies, which noted that LSA-based summarization captures key topics through latent structures rather than relying on explicit word frequency (Steinberger & Jezek, 2009; Gong & Liu, 2001).

Notably, the results also show that although LSA trails slightly behind neural models like RankNet in overall F1 score, it achieves this performance without the computational overhead or training data requirements associated with deep learning techniques (Svore *et al*., 2007). This makes LSA particularly suitable for low-resource settings or domains that demand fast, interpretable solutions. Furthermore, LSA's unsupervised nature allows it to adapt across various text types without retraining, a trait that Yeh *et al*. (2005) highlighted as essential for flexible summarization systems. One of the more revealing aspects of the analysis is revealed in Figure 3, which displays the distribution of sentence scores within the semantic space. The consistency of high-scoring sentences illustrates LSA's ability to prioritize semantically dense content across the document, supporting observations by Radev, Hovy, and McKeown (2002), who emphasized the importance of sentence relevance and diversity in summary generation. The clustering effect seen in the VT matrix confirms that LSA effectively captures dominant thematic elements, even when surface features like sentence position or length vary.

Another noteworthy point is the apparent increase in cosine similarity as summary length expands, seen in Figure 2. While this confirms that more extended summaries encapsulate more of the original content, it also aligns with observations from Carbonell and Goldstein (1998) regarding redundancy. Excessive inclusion of content can reduce the effectiveness of a summary by reintroducing superfluous information. Hence, while a 30 percent summary yields the highest cosine similarity, the optimal range for summarization likely lies closer to the 10 to 20 percent mark, where a balance between coverage and conciseness is achieved. In comparing LSA with traditional models, such as Naïve Bayes, it becomes clear that semantic modeling provides a substantial advantage. Traditional models, which rely primarily on positional and frequency-based features, often fall short in extracting diverse, meaningful content. Lin and Hovy (1997) previously noted the limitations of such methods in dealing with redundancy and lack of context, concerns that LSA addresses more robustly through its matrix decomposition and dimensionality reduction techniques. Despite its strengths, LSA is not without limitations. It cannot handle polysemy effectively and lacks the linguistic generation capabilities that neural models possess. Neural summarizers, particularly those using ranking and sentence generation techniques, have been praised for their fluency and adaptability across domains (Young *et al*., 2017). However, as Hovy and Lin (1999) cautioned, the increased performance of these models often comes at the expense of transparency and interpretability, critical factors in applications where understanding model decisions is necessary. Overall, the discussion suggests that LSA-based summarization remains highly competitive in specific contexts, particularly when ease of implementation, generalizability, and performance are weighed equally. It serves as a viable alternative or complement to neural summarization models, especially in scenarios where model explainability and resource constraints are prioritized. The comparative evaluation supports the idea that there is no one-size-fits-all solution in text summarization; instead, the choice of method should align with the goals, constraints, and characteristics of the intended application.

RECOMMENDATIONS

Based on the comparative evaluation conducted in this study, several recommendations emerge for practitioners, researchers, and developers working on text summarization systems, particularly those considering Latent Semantic Analysis (LSA) about traditional and neural methods.

Firstly, LSA should be considered a strong baseline for extractive summarization tasks, especially in low-resource environments where training data and computational power are limited. Its unsupervised nature allows for rapid deployment without the need for domain-specific training, making it ideal for business applications, educational tools, and information retrieval systems that require quick, interpretable summaries (Steinberger & Jezek, 2009; Yeh *et al*., 2005). It is recommended that institutions or organizations with limited access to high-performance hardware consider implementing LSA over deep learning models, which are often prohibitively resource-intensive.

Secondly, to improve LSA's performance, the preprocessing stage must be carefully optimized. This includes thorough stop word removal, consistent stemming, and accurate part-of-speech tagging. These steps significantly influence the quality of the term-sentence matrix and, by extension, the effectiveness of the SVD operation. Practitioners should also experiment with adjusting the number of singular values retained during decomposition to find the optimal dimensionality for different document types.

Third, while LSA performs well at identifying central themes, it does not inherently minimize redundancy. Therefore, it is advisable to implement post-processing techniques, such as Maximal Marginal Relevance (Carbonell & Goldstein, 1998), to reduce overlap between selected sentences and enhance content diversity in summaries.

Lastly, researchers aiming to improve LSA's competitiveness with neural models could explore hybrid systems that integrate LSA with lightweight learning components. Such systems could benefit from LSA's semantic clarity while adding flexibility through data-driven adaptation. LSA remains a valuable method in modern summarization, particularly when tuned correctly and used within its strengths. With thoughtful application and enhancements, it can serve as a powerful alternative or complement to more complex models.

FUTURE RESEARCH DIRECTIONS

Building on the findings and limitations identified in this study, several future research directions are proposed to enhance the capabilities of Latent Semantic Analysis (LSA) and expand its applicability in automatic text summarization. One promising direction is the extension of LSA to multi-document summarization. While the current study focuses on single-document summarization, future work could investigate how LSA performs when synthesizing summaries from multiple sources. This would involve addressing challenges such as redundancy across documents and identifying common themes from diverse writing styles and structures. Yeh *et al*. (2005) highlighted the potential of LSA in handling corpus-level semantic relationships, suggesting that it could be adapted for such complex tasks. Another area worth exploring is the integration of LSA with sentence compression and co-reference resolution techniques. Since LSA operates on sentence-level vectors without linguistic understanding, combining it with methods that resolve entity references or remove redundant phrases could improve the coherence and conciseness of generated summaries (Hovy & Lin, 1999). This hybrid approach could make LSA-generated summaries more fluid and readable, closer to human-written abstracts.

Language diversity also presents a significant opportunity. Future studies could assess the performance of LSA across different languages, particularly those with less NLP infrastructure or fewer resources. This would involve evaluating how well LSA's term-sentence matrix construction adapts to languages with different morphological and syntactic characteristics. Additionally, dimensionality tuning and adaptive thresholding for sentence selection remain underexplored. Research could focus on developing automated mechanisms to select the optimal number of concepts (singular values) retained during SVD based on document characteristics such as length, topic density, or writing complexity. Finally, there is potential in developing interactive or semi-supervised LSA systems where users can provide feedback to guide sentence selection. This could make the summarization process more adaptable and personalized, especially in domains like education, healthcare, and legal analysis, where context sensitivity is crucial. These future directions aim to build on LSA's strengths while addressing its current limitations, ensuring its continued relevance in the evolving landscape of text summarization technologies.

CONCLUSION

The results of this study demonstrate that Latent Semantic Analysis (LSA) remains a practical and effective method for extractive text summarization, particularly when evaluated through the lens of semantic preservation using cosine similarity. While it does not outperform advanced neural models such as RankNet in raw performance metrics, it offers a compelling balance of interpretability, efficiency, and semantic depth, making it highly suitable for specific real-world applications where transparency and resource constraints are critical factors. By leveraging the statistical relationships between terms and sentences, LSA captures underlying semantic patterns in text without requiring labeled data or complex training procedures. This study has shown that with carefully executed preprocessing and optimized sentence selection based on singular value distributions, LSA can produce summaries that closely align with the thematic content of original documents. The consistent performance across different summary lengths and the informative sentence scoring patterns further validate its robustness.

When compared with traditional models like Naïve Bayes, LSA surpasses them in its ability to identify and prioritize meaningful content beyond superficial textual features such as position or frequency. While neural models outperform LSA in precision and recall, their higher computational demands and opaque decision-making processes make them less practical in specific contexts (Svore *et al*., 2007; Young *et al*., 2017). Therefore, LSA's continued relevance lies in its ability to deliver semantically rich, fast, and interpretable summaries, qualities that are increasingly important in domains where users need to understand not just what the model predicts, but why. As this study has shown, LSA is not only a valuable baseline but also a viable standalone solution for extractive summarization, particularly when used within its optimal design scope. With further enhancement and hybrid integration, LSA can continue to play a meaningful role in the development of next-generation summarization tools.

REFERENCES

- Carbonell, J., & Goldstein, J. (1998). The use of MMR, diversity-based reranking for reordering documents and producing summaries. *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 335–336.
- Conroy, J. M., & O'Leary, D. P. (2001). Text summarization via hidden Markov models. *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 406–407.
- Gong, Y., & Liu, X. (2001). Generic text summarization using latent semantic analysis. *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 19–25.
- Hovy, E., & Lin, C. Y. (1999). Automated text summarization in summarist. *Advances in Automatic Text Summarization*, 81–94.
- Lin, C. Y. (1999). Training a selection function for extraction. *Proceedings of the Eighth International Workshop on Information Retrieval with Asian Languages*, 55–60.
- Lin, C. Y., & Hovy, E. H. (1997). Identifying topics by position. *Proceedings of the Fifth Conference on Applied Natural Language Processing*, 283–290.
- Kupiec, J., Pedersen, J., & Chen, F. (1995). A trainable document summarizer. *Proceedings of the 18th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 68–73.
- Radev, D. R., Hovy, E., & McKeown, K. (2002). Introduction to the special issue on summarization. *Computational Linguistics*, 28(4), 399–408.
- Steinberger, J., & Jezek, K. (2009). Using latent semantic analysis in text summarization and summary evaluation. *Proceedings of ISIM*, 93–100.
- Svore, K. M., Burges, C. J. C., & Goodman, J. (2007). Ranking using multiple semantic and syntactic features. *Proceedings of the IEEE International Conference on Data Mining*, 448–457.
- Yeh, C. Y., Lin, C. H., & Ke, H. R. (2005). A learning-based approach to automatic text summarization. *IEEE Transactions on Systems, Man, and Cybernetics - Part B: Cybernetics*, 35(5), 1085–1098.
- Young, T., Hazarika, D., Poria, S., & Cambria, E. (2017). Recent trends in deep learning based natural language processing. *IEEE Computational Intelligence Magazine*, 13(3), 55–75.