A Review of Quantum Computing Approaches to Semantic Search and Text Classification in Natural Language Processing

Sauvik Bal[1,2] Shillpi Mishra[1] Lopa Mandal[3]
[1]Techno India University, West Bengal, India
[2]Maulana Abul Kalam Azad University of Technology, West Bengal, India
[3]Alliance University, Bangalore, India
Correspondence: sauvikbal@gmail.com

Abstract

While having enhanced NLP, deep learning and pre-trained language models requires a lot of processing power. The work showcases the potential of quantum computing by mapping linguistic data into vast, high-dimensional Hilbert spaces through entanglement and superposition. It focuses on mathematical concepts that set quantum approaches apart from classical ones, among them being the fidelity-based similarity and quantum probability. Various quantum machine learning models are considered in this article, including Quantum Neural Networks and Quantum Support Vector Machines, each discussing the computational advantages in pattern recognition. In addition, it considers retrieval techniques like Grover's algorithm, showing how quantum similarity functions give better semantic search. Indeed, the comparison does show that quantum techniques might yield advantages regarding expressiveness and scalability, despite obstacles such as hardware noise and data encoding. Notwithstanding that quantum technology is still in its infancy, future improvements might advance language understanding.

Keywords

Quantum Computing, Text Classification, Semantic Search, Information Retrieval, Natural Language Processing (NLP), Quantum Neural Networks

1 Introduction

The explosive evolution of natural language processing (NLP) has mostly been triggered by traditional machine learning and deep learning models, which have reported impressive performance in applications like text classification, sentiment analysis, and semantic search (Devlin et al., 2019). Notwithstanding these breakthroughs, the ever-growing dimensionality of text data and the computational expense of large models have made it imperative to look for other approaches that can offer efficiency without sacrificing semantic richness. Quantum computing has, in recent years, been explored as a possible paradigm to overcome such limitations because it can perform computations in exponentially big Hilbert spaces and leverage principles like superposition and entanglement (Schuld and Petruccione, 2019).

Quantum models, such as QSVMs and QNNs, embed texts into high-dimensional quantum feature spaces, hence being more effective for text categorization than classical techniques. Quantum-inspired information retrieval techniques rely on Hilbert space formalism and fidelity measurements while offering advantages over classical methods by virtue of Grover's search algorithm. The analytical framework includes quantum kernels and probability distributions that extend conventional comparison metrics such as cosine similarity. However, despite theoretical advantages, noisy and resource-limited NISQ devices make practical implementation very challenging. Therefore, hybrid quantum-classical approaches are considered a viable approach. The current study will review the mathematical underpinnings of quantum NLP research, complexity assessments, and comparative insights between quantum and classical approaches to highlight the potential benefits and current challenges in quantum NLP research.

It starts with theoretical notions (described in Figure 1), such as Hilbert spaces and Grover's algorithm, and the review structure progresses from a purely mathematical underpinning to a real-world application. It falls into two main areas: one regarding quantum semantic search by fidelity measures and quantum walks, and another on quantum text classification by means of QSVM and QNN/VQC models. In order to identify complementarities and trade-offs, these branches merge under a comparative study that analyzes mathematical methods and empirical behaviors. Furthermore, this frame-

Proceedings of the QuantumNLP: Integrating Quantum Computing with Natural Language Processing, pages 33–43
November 24, 2025

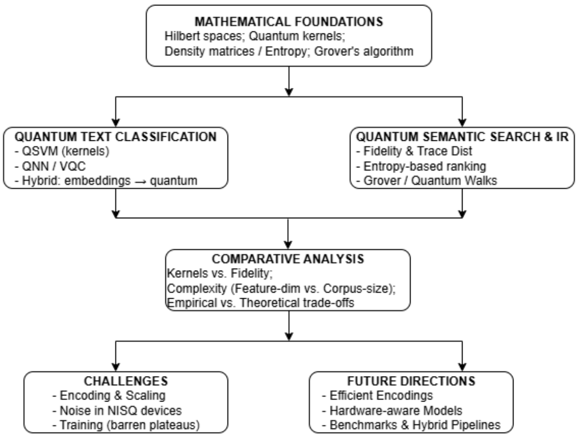

Figure 1: Overview of the analytical review

work indicates problems like noise and encoding and suggests future avenues of research, such as uniform benchmarks and effective encodings.

2 Scope and Review Methodology

With a focus on text classification and semantic search, this article reviews the integration of quantum computing into NLP. It outlines the development from early quantum-inspired frameworks to advanced quantum kernel methods and hybrid models, surveying literature from 2010 to 2025. Among others, IEEE Xplore, ACM Digital Library, and arXiv can be consulted using keywords and phrases such as "quantum NLP" and "quantum semantic search." Only those studies that introduce quantum models for natural language processing (NLP) problems, develop theoretical insights into either text classification or semantic search, and offer analytical contrasts between classical and quantum approaches are reviewed. Works which are purely classical or with no direct relevance to NLP will not be considered. To ensure a systematic progress review in quantum NLP, contributions are grouped into three categories: quantum models for text categorization, quantum approaches for semantic search, and supporting mathematical analyses.

3 Mathematical Foundations

The use of quantum computing in natural language processing (NLP) is mathematically intense. The following section presents the mathematical basics that make up the analytical framework of quantum methods in text classification and semantic search. These are Hilbert spaces, quantum probability, measures of fidelity, kernel methods, and computational complexity.

3.1 Hilbert Spaces and Quantum Text Representation

In quantum mechanics, physical systems are encoded in a complex Hilbert space $\mathcal{H}$, and the system's state is specified as a normalized vector in that space. A pure state is normally written as:

$$|\psi\rangle = \sum_i c_i |i\rangle, \qquad \sum_i |c_i|^2 = 1, \tag{1}$$

where c_i are complex probability amplitudes. $\{|i\rangle\}$ denotes an orthonormal basis of the Hilbert space $\mathcal{H}$.

In NLP, the $|i\rangle$ are linked with a token, word embedding, or latent semantic component. Superposition is built into quantum representations to enable more expressive encoding of semantic relations over classical embeddings.

3.2 Quantum Probability and Density Matrices

Quantum probability is derived from the Born rule. For a state $|\psi\rangle$, the probability of observing basis state $|i\rangle$ is:

$$P(i) = |\langle i|\psi\rangle|^2. \tag{2}$$

For mixed states, a density matrix ρ is defined as:

$$\rho = \sum_k p_k |\psi_k\rangle\langle\psi_k|, \quad \text{with} \quad \text{Tr}(\rho) = 1. \tag{3}$$

This allows ambiguous words to be modeled as probabilistic mixtures of multiple semantic states (Piwowarski et al., 2010). The information content of a state is quantified using the von Neumann entropy:

$$S(\rho) = -\text{Tr}(\rho \log \rho), \tag{4}$$

which generalizes Shannon entropy into quantum systems.

3.3 Similarity and Distance Metrics

Semantic similarity in quantum models is expressed via fidelity:

$$F(\rho, \sigma) = \left(\text{Tr}\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}\right)^2, \tag{5}$$

where ρ and σ represent query and document states. Fidelity generalizes cosine similarity by embedding comparisons in Hilbert space (van der Meer et al., 2021). Another important measure is the trace distance:

$$D(\rho, \sigma) = \tfrac{1}{2}\text{Tr}\,|\rho - \sigma|, \tag{6}$$

which captures dissimilarity between semantic states.

3.4 Illustrative Comparison: Cosine Similarity vs. Fidelity

To better understand how quantum similarity measures differ from classical ones, consider two simple normalized 2-dimensional vectors representing a query q and document d:

$$q = (1, 0), \qquad d = \left(\frac{\sqrt{2}}{2}, \frac{\sqrt{2}}{2}\right). \tag{7}$$

Cosine Similarity:

$$\begin{aligned}\cos(q, d) &= \frac{q \cdot d}{\|q\|\|d\|} \\ &= \frac{1 \cdot \frac{\sqrt{2}}{2} + 0 \cdot \frac{\sqrt{2}}{2}}{1 \cdot 1} = \frac{\sqrt{2}}{2} \approx 0.707.\end{aligned} \tag{8}$$

Fidelity Measure: When the same vectors are treated as pure quantum states $|q\rangle$ and $|d\rangle$, fidelity is defined as:

$$F(q, d) = |\langle q|d\rangle|^2 = \left(\frac{\sqrt{2}}{2}\right)^2 = 0.5. \tag{9}$$

Interpretation:

- Cosine similarity measures *geometric angle* between classical vectors.
- Fidelity measures *quantum probability overlap* between states.

Although fidelity reduces to the *square* of cosine similarity for pure normalized states, the probabilistic meaning of fidelity is more aligned with quantum measurements. In more complex mixed-state scenarios (e.g., density matrices), fidelity captures richer semantic uncertainty beyond what cosine similarity provides.

3.5 Quantum Kernels and Feature Maps

In classification tasks, quantum kernels extend the classical kernel trick. A quantum feature map $\phi(x)$ encodes data into quantum states, and the kernel function is defined as:

$$k(x, y) = |\langle\phi(x)|\phi(y)\rangle|^2. \tag{10}$$

This helps the model to handle the large feature spaces where fewer resources can be used (Havlíček et al., 2019).

3.6 Complexity Considerations

In theory, quantum algorithms provide great advantages in terms of computational efficiency. Grover's technique reduces unstructured database searches from $O(N)$ to $O($
$sqrtN)$ for large document collections (Grover, 1996). The HHL algorithm can solve linear equations exponentially faster under some conditions. These advances are certainly relevant for NLP tasks involving high-dimensional embeddings and large text corpora; hence, there is the prospect of quantum methods outperforming classical systems for semantic search and classification (Harrow et al., 2009).

3.7 Implications for NLP

This paper emphasizes analytical comparisons to classical models using Hilbert space formalism, quantum probability, and kernel-based feature mapping to realize NLP tasks. In doing so, it is shown that quantum models may decrease processing costs while retaining the semantic information undisturbed. These mathematical frameworks, in essence, serve as the foundation for approaches using quantum machine learning on linguistic data in practical quantum models for text classification, such as QSVMs, QNNs, and hybrid quantum-classical architectures.

4 Quantum Computing for Text Classification

One of the main applications of NLP is text categorization. It is useful for various purposes, like spam filtering and sentiment analysis etc.. While they deliver accurate results, conventional models, including transformer-based models such as BERT and support vector machines, are computationally expensive, particularly in the scenario of high-dimensional feature spaces. A suitable replacement is quantum computing, which enhances classification efficiency through the application of large Hilbert spaces and ideas of superposition and entanglement. The primary focus of this section's coverage of quantum models for text classification is quantum support vector machines, quantum neural networks, and hybrid architectures.

As shown in Table 1, recent developments in quantum computing for text classification are summarized chronologically, covering the period from 2019 to 2025. The table highlights the main approaches, encoding strategies, model types, as well

Table 1: Chronological comparison of quantum (and quantum-like) approaches for text classification.

Year	Approach (citation)	Encoding	Model Type	Advantages	Limitations
2019	QSVM with quantum kernels (Havlíček et al., 2019)	Angle / amplitude	QSVM kernel classifier	Exponential feature mapping; effective separation in high-dim spaces	Sensitive to noise; evaluated only on small datasets
2019	Quantum Convolutional Neural Networks (QCNN) (Cong et al., 2019)	Structured encoding	Convolutional-like QNN	Parameter efficient; locality aware; some robustness to noise	Task-specific design; deeper circuits required
2021	Quantum Neural Networks (QNN) (Abbas et al., 2021)	Angle / amplitude	Variational quantum circuits (VQC)	High expressivity; end-to-end trainable	Barren plateaus (gradient vanishing); noisy hardware limits
2022	Quantum SVM for Text Classification (Li et al., 2022)	Amplitude / angle from embeddings	QSVM + hybrid embedding	Uses word2vec/BERT embeddings; competitive accuracy	Inherits embedding costs; limited to small corpora
2024	Quantum Support Vector Classifier on NISQ hardware (Suzuki, 2024)	Angle / amplitude	QSVC on IBM/IonQ devices	First hardware validation; real device results	Strongly affected by noise; dataset scaling issues
2024	Quantum Self-Attention Neural Networks (QSANN) (Li, 2024)	Classical embeddings $\rightarrow$ quantum attention	Hybrid quantum-classical attention	Captures contextual dependencies; integrates attention with QNN	Only small simulator tests; hardware results pending
2024	Hybrid transfer learning (BERT + QSVM/QNN) (Anonymous, 2024)	Pretrained embeddings $\rightarrow$ quantum classifier	Hybrid pipeline	Combines classical embeddings with quantum classifiers; practical for NISQ	Dependent on pretrained models; added quantum overhead
2025	Quantum-like wave model for semantic classification (Gruzdeva et al., 2025)	Semantic units $\rightarrow$ wave embeddings	Quantum-like wave model	Captures interference in semantics; accuracy gains over baselines	Semi-heuristic; not hardware-based; small datasets only
2025	Hybrid QTL with kernel self-attention (Chen and Lou, 2025)	Classical embeddings + quantum kernel	Hybrid transfer-learning	Improves feature separability; tested on real datasets	Complex architecture; hardware scaling challenges
2025	Single-Qudit QNN (SQ-QNN) (Souza and Portugal, 2025)	Angle encoding into qudits	Qudit-based QNN	Reduces qubit needs by using qudits; supports multi-class tasks	Still theoretical/simulator-level; qudit hardware needed

as their analytical advantages and limitations, providing a structured comparison of progress in this domain.

4.1 Quantum Support Vector Machines (QSVM)

An optimal separating hyperplane within a feature space is defined by the Support Vector Machine, which is basically a supervised learning algorithm. It is used to classify the data. The Quantum Support Vector Machine maps data into high-dimensional Hilbert spaces. Unitary operations are used to map data points into quantum states in QSVM. A quantum kernel defines the similarity of these states. Basically, the architecture allows QSVMs to utilize polynomial resources (Havlíček et al., 2019; Schuld and Petruccione, 2019) on quantum hardware and operate within an exponentially dimensional feature space. With weighted kernel evaluations on training sets, a decision function in a QSVM is analogous to traditional SVMs. QSVMs are an interesting method because quantum kernels are capable of separating classes that polynomial-time classical kernels are not.

4.2 Quantum Neural Networks (QNNs)

Quantum Neural Networks (QNNs) are constructed using Variational Quantum Circuits (VQCs), which consist of parameterized unitary gates optimized with a classical optimizer. For an input x, the quantum state is expressed as:

$$|\psi(\theta, x)\rangle = U(\theta, x)|0\rangle, \tag{11}$$

where θ represents trainable parameters. The probability of assigning class y is given by:

$$P(y|x) = |\langle y|\psi(\theta, x)\rangle|^2. \tag{12}$$

QNNs have been shown to achieve expressive power that surpasses shallow classical neural networks, particularly in handling non-linear and high-dimensional relationships (Abbas et al., 2021). However, they face challenges such as barren plateaus, where gradients vanish exponentially with circuit depth. Recent studies have also demonstrated the practical use of QNN-based classifiers specifically for NLP tasks (Pandey et al., 2024).

4.3 Data Encoding Strategies

Encoding is an essential link between textual data and quantum computation, which requires converting units of language into quantum states with semantic integrity preserved. Effective encoding techniques have impact on representational expressivity and hardware viability in NISQ devices and are crucial, not only for text categorization but also for quantum semantic search and information retrieval.

Different encoding strategies have been discussed here:

- **Basis encoding:** Each token or feature is placed directly into a basic quantum state.

Here,it gives a clear but sparse form in the qubit form.

- **Amplitude encoding:** dense embeddings such as word2vec or BERT are normalized and encoded into the amplitudes of a quantum state:
$$|x\rangle = \frac{1}{\|x\|} \sum_i x_i |i\rangle, \tag{13}$$
providing an exponentially compact form of feature representation.

- **Angle encoding:** numerical features are converted into rotation angles of single-qubit gates, offering low circuit depth at the cost of reduced representational capacity.

Amplitude encoding is particularly appealing for NLP tasks, as it enables dense semantic embeddings to be represented efficiently in Hilbert space while still benefiting from quantum parallelism (Schuld and Petruccione, 2019). As quantum NLP advances, developing encoding techniques that balance compactness, expressiveness, and noise resilience will remain a key challenge.

4.4 Hybrid Quantum-Classical Models

The NISQ era of existing quantum devices faces challenges in quantum model development in NLP due to the low number of qubits and pervasive noise (Suzuki, 2024). Hybrid quantum-classical models leverage classical embeddings generated using word2vec and BERT, which can then be fed into a quantum classifier such as the QNN or QSVM (Devlin et al., 2019; Abbas et al., 2021; Li et al., 2022). The combination leverages quantum capabilities for improved feature separation and the quality of the classification of particularly complex text distributions but draws on classical strengths for embeddings (Havlíček et al., 2019). Somewhat limited in accuracy due to the hardware, QSVMs and quantum kernel classifiers have produced successful results on IBM and IonQ devices (Suzuki, 2024). A completely quantum solution, hybrid models represent an exciting way to incorporate quantum computing into NLP applications (Anonymous, 2024).

4.5 Complexity Analysis

Quantum Support Vector Machines employ polynomial circuit resources to evaluate similarities in exponentially vast spaces. This is done through the use of the quantum kernel trick. Analytical comparisons shows that:

- For the d-dimensional embedding, the classical SVM kernel evaluation mainly requires $O(d)$ operations.

- Quantum kernel evaluation scales as $O(\text{poly}(n))$. Here n is used as the number of qubits, which represents the feature space.

This advantage is significant when processing high-dimensional language embeddings is done. This is frequently used in NLP applications. However, the reliability of kernel estimate on NISQ hardware. It restricted connection, and qubit noise limit realistic speedup.

Complexity of QNN-based Models: Variational Quantum Circuits are a novel computational regime in QNNs, with complexities dependent on the number of qubits n, the circuit depth L and trainable parameters $|\theta|$. At each iteration of the gradient-based training, it is necessary to run the quantum circuit multiple times, incurring a certain cost.

$$O\left(L \cdot \text{poly}(n) \cdot |\theta|\right). \tag{14}$$

For large n or L, QNNs suffer from problems like *barren plateaus* due to vanishing gradients, which make optimization costly. In contrast with QSVMs, that rely on quantum kernel evaluation for computation, QNNs are plagued by scalability issues due to optimization overhead and hardware noise. Because of optimization complexity and coherence restrictions in NISQ devices, increased expressiveness of QNNs can hardly be exploited in practice.

4.6 Applications in NLP

Quantum models, like QSVMs and QNNs, have competed with conventional models in various NLP tasks, including sentiment analysis, spam filtering, and fake news detection (Pandey et al., 2024). Specifically, QSVMs are good at categorizing reviews, while QNNs have high performance in identifying trustworthy news sources and filtering spam using hybrid architectures. Moreover, applying quantum feature spaces to enable data-efficient solutions holds promising advances for multilingual and low-resource languages. This work points toward a path for future research in scalable quantum hardware development, as it emphasizes not only what has been achieved with supervised quantum models in text classification but also the potential for quantum-inspired methods in information retrieval and semantic search (discussed in Section 5).

5 Quantum Computing for Semantic Search and Information Retrieval

Quantum semantic search, by exploiting Hilbert space representations and quantum similarity measures, improves the ranking of documents, thereby outperforming conventional models based on lexical matching, such as TF-IDF and BM25 (Piwowarski et al., 2010; van der Meer et al., 2021). Quantum-inspired information retrieval models leveraged amplitude-encoded quantum states to model documents and queries with the aim of incorporating semantic aspects and uncertainties (Schuld and Petruccione, 2019). While quantum distance measures, such as trace distance, quantify semantic dissimilarity (van der Meer et al., 2021), quantum similarity measures, such as fidelity, augment classical cosine similarity by reflecting probabilistic overlaps (Piwowarski et al., 2010). Hybrid quantum-classical methods combine classical embeddings with quantum techniques to make the most of existing NISQ hardware (Yamada et al., 2024; Devlin et al., 2019), while some quantum algorithms, such as Grover's search, obtain significant savings in search time (Grover, 1996). In this changing approach to IR, with quantum technologies still evolving, the representation of deeper semantic relevance and uncertainty points to a more expressive future for search algorithms (Zhang et al., 2023; Gupta et al., 2025).

5.1 Hilbert Space Representations of Documents

The documents and queries are mainly represented as vectors in a Hilbert space $\mathcal{H}$ at the time of the quantum-inspired retrieval process. This is denoted as:

$$|d\rangle = \sum_i \alpha_i |t_i\rangle, \qquad |q\rangle = \sum_i \beta_i |t_i\rangle, \tag{15}$$

Here $|t_i\rangle$ represents the basis vectors, and α_i, β_i represents the normalized weights.

By using the fidelity, the similarity between a query q and a document d has measured:

$$\mathcal{F}(q,d) = \left|\langle d|q\rangle\right|^2. \tag{16}$$

Compared to traditional cosine similarity, this gives the squared inner product of the two states, which provides a more descriptive similarity metric (Piwowarski et al., 2010; van der Meer et al., 2021).

5.2 Quantum Probability and Entropy Measures

Quantum IR can also use density matrices to encode uncertainty in semantic states. For a document mixture, the density operator is given as:

$$\rho_d = \sum_i p_i |d_i\rangle\langle d_i|, \tag{17}$$

where p_i are probability weights. The similarity between documents can then be computed using trace distance or von Neumann entropy:

$$S(\rho) = -\mathrm{Tr}(\rho \log \rho). \tag{18}$$

Entropy-based ranking allows capturing semantic diversity and ambiguity, beyond what is possible in classical IR frameworks (Zhang et al., 2023).

5.3 Grover's Algorithm for Document Retrieval

Grover's quantum.search algorithm achieves a quadratic speedup for unstructured search problems (Grover, 1996). For a collection of N documents, Grover's algorithm will locate a matching document in $O(\sqrt{N})$ time, as opposed to $O(N)$ classically. Mathematically, successive applications of the Grover operator G increase the likelihood of the target state $|d^*\rangle$:

$$G = (2|s\rangle\langle s| - I) \cdot (I - 2|d^*\rangle\langle d^*|), \tag{19}$$

where $|s\rangle$ is the uniform superposition of all states. This gives theoretical speedups for large-scale IR.

5.4 Hybrid Quantum-Classical IR Models

Recent efforts integrate classical embeddings (e.g., BERT, word2vec) with quantum fidelity-based retrieval. Queries and documents are first embedded in dense vector spaces, then encoded into quantum states for matching. Such hybrid approaches provide practical pathways for deploying quantum IR on NISQ-era hardware (Yamada et al., 2024).

5.5 Chronological Comparison of Approaches

Table 2 summarizes the basic developments in quantum IR approaches from 2010 to 2025. It highlights the encodings, models, advantages, and limitations.

Table 2: Chronological comparison of quantum computing approaches for semantic search and information retrieval.

Year	Approach (citation)	Encoding	Model Type	Advantages	Limitations
2010	Quantum-inspired IR framework (Piwowarski et al., 2010)	Term basis states	Hilbert space retrieval	Introduced fidelity-based query-document similarity; linked IR to quantum probability	Conceptual framework; no hardware implementation
2019	Quantum probability ranking model (Zuccon and Azzopardi, 2019)	Amplitude encoding of terms	Quantum-inspired ranking model	Probabilistic interpretation of ranking; novel use of quantum probability	Early-stage model; tested on small corpora
2021	Quantum algorithms for IR (van der Meer et al., 2021)	Amplitude encoding	Hybrid algorithms for retrieval	Theoretical speedups using Grover's search and quantum walks	Lacks large-scale hardware benchmarks
2023	Entropy-based quantum IR (Zhang et al., 2023)	Density matrices	Entropy ranking framework	Incorporates semantic diversity and ambiguity; entropy-based document ranking	Simulator-based; hardware scaling not addressed
2024	Hybrid embedding + quantum fidelity search (Yamada et al., 2024)	BERT embeddings → quantum states	Hybrid quantum-classical IR	Integrates deep embeddings with quantum fidelity search; suitable for NISQ devices	Dependent on pretrained embeddings; limited qubits
2025	Quantum walk-based semantic retrieval (Gupta et al., 2025)	Amplitude encoding of graph embeddings	Quantum walk retrieval	Explores semantic search using quantum walks over document graphs; potential retrieval efficiency gains	Experimental stage; scalability to large corpora unproven

As shown in Table 2, approaches span from foundational quantum-inspired frameworks in 2010 to recent hybrid and quantum walk-based retrieval models in 2025, demonstrating the evolution from conceptual theory to practical hybrid implementations.

6 Comparative Analytical Insights

In this section, two core NLP tasks are described text categorization in Table 1 and semantic search/information retrieval in Table 2. These are compared using quantum techniques. It mainly highlights each domain's unique issues, their mathematical formulations, and development paths while also pointing out their trade-offs and complementarities.

6.1 Mathematical Underpinnings

Quantum text classification methods are predominantly kernel-based or variational, relying on mappings into exponentially large Hilbert spaces and parameterized quantum circuits. Quantum Kernel's analytical construction is as follows:

$$k(x, y) = |\langle \phi(x)|\phi(y)\rangle|^2, \tag{20}$$

QSVMs and hybrid kernel models are supported by the idea like, (Havlíček et al., 2019; Li et al., 2022). To check the similarity between a particular query and the document, quantum probability, density matrices, and fidelity is used for semantic search methods.

$$F(q, d) = |\langle q|d\rangle|^2, \quad S(\rho) = -\mathrm{Tr}(\rho \log \rho). \tag{21}$$

Quantum kernels for text classification define unique decision limits, making the distinction among categories of data easier. On the other hand, semantic search and information retrieval focus on the relevance of the materials to the queries. They aim to emphasize relevant pages and provide effective representation of semantic meaning through the use of probability-based similarity and entropy measurements.

6.2 Computational Complexity

Whereas QNNs suffer from optimization problems, QSVMs offer implicit embeddings in $O(2^n)$ dimensions using polynomial resources for classification (Abbas et al., 2021). Grover's search reduces the query complexity in information retrieval from $O(N)$ to $O(\sqrt{N})$; quantum walk-based approaches provide better document graph exploration (Gupta et al., 2025). While retrieval puts an emphasis on query scaling and ranking efficiency, classification emphasizes decision boundary complexity.

6.3 Evolution of Approaches (2010–2025)

From 2010 to 2025, research changed from conceptual formulations to hybrid implementations. For example, some classification advances that show progress towards NISQ practicality are the quantum-inspired classifiers of 2019, the hybrid pipelines with BERT embeddings and QSVMs of 2024, and the single-qudit QNNs of 2025. Some recent examples of retrieval advances include the Hilbert space-based IR models of 2010, entropy-driven ranking in 2023, hybrid embedding-fidelity models in 2024, and quantum walk retrieval in document graphs from 2025. This trajectory shows a convergence towards hybrid paradigms, combining quantum-enhanced classifiers and retrieval systems with classical embeddings.

6.4 Analytical Trade-offs

The comparison analysis shows different types of significant trade-offs:

- **Expressivity vs. Stability:** Entropy-based IR techniques provide stability, but they are less expressive in terms of model power. QNNs have large capacity but tey are difficult to optimize.
- **Scalability:** Classification complexity scales with embedding dimensionality, while retrieval scales with corpus size. Both benefit from quantum asymptotic advantages in distinct regimes.
- **Hardware Realization:** Classification methods (e.g., QSVMs) have been experimentally tested on NISQ devices (Suzuki, 2024), whereas IR models remain mostly simulator-bound, with limited demonstrations on hardware.

6.5 Outlook

The two main approaches in quantum NLP, quantum text categorization and quantum semantic search, are complementary rather than competing. Quantum text classification excels in supervised tasks that involve clear-cut decision boundaries, whereas quantum semantic search adopts probabilityand entropy-based measures to capture the meaning of texts and rank documents. In both areas, researchers are moving toward hybrid quantum-classical architectures, and thus classification and retrieval will eventually be part of NLP systems that have traditional components for preprocessing and quantum circuits for semantic reasoning. As quantum devices improve, various applications and practical quantum advantages for NLP might become possible.

7 Experimental Landscape and Benchmarking Status

Theoretically, quantum NLP demonstrates great possibilities, but empirical verification is not possible due to the constraints of existing NISQ hardware. The majority of research uses quantum simulators and short datasets, focusing on practicality before completeness of performance. Benchmarking trends are presented in this section for semantic search and quantum text classification.

7.1 Datasets Used in Current Studies

In some instances, compact datasets have been employed to benchmark quantum-enhanced classifiers: for instance, the SMS Spam Dataset for binary spam filtering with QSVMs (Li et al., 2022), portions of Amazon or IMDb reviews for sentiment analysis based on QNN-based models (Pandey et al., 2024), and TREC-style toy retrieval sets for query relevance assessment (van der Meer et al., 2021). Due to qubit availability constraints, the IR experiments often employ simulated semantic vectors rather than complete corpus representations (Piwowarski et al., 2010). Hybrid BERT-embedded document matching has seen a bit more development, although its application remains limited to very small corpora (Yamada et al., 2024).

7.2 Evaluation Metrics

Performance evaluation typically combines established classical metrics with quantum-specific similarity measures:

- **Accuracy, Precision, Recall, F1-score** for classification (Suzuki, 2024).
- **Entropy-based ranking** to measure semantic diversity (Zhang et al., 2023).
- **Fidelity** as a probabilistic similarity score between query and documents (Piwowarski et al., 2010).

These mixed metrics reflect an ongoing effort to account for both prediction quality and quantum semantic overlap.

7.3 Simulators vs. Hardware Deployments

The lack of standard benchmarks due to differences in dataset size, encoding methodologies, simulator precision, error models, hardware platforms, and circuit depth limits hinders precise performance comparisons among research (van der Meer et al., 2021). Thus, the assertions on quantum advantage in NLP are prima facie tentative and bound by experimental design (Abbas et al., 2021). Medium-scale data sets may be manageable for future information retrieval as qubit counts and noise robustness will likely continue to improve (Suzuki, 2024). Effective data embeddings may allow multilingual and low-resource tasks to benefit, and standard evaluation metrics accounting for accuracy, fidelity, and complexity analysis will be important moving forward (Yamada et al., 2024). Hardware-aware model design and standard benchmarking are necessary before large-scale demonstrations of quantum NLP performance can be realized (Anonymous, 2024).

7.4 Current Limitations in Benchmarking

Current implementations mainly rely on quantum simulators, such as Qiskit and Cirq, because of issues with noise and coherence in real quantum hardware (van der Meer et al., 2021). While hardware-based evaluation is at an early stage of development, problems such as significant losses in the accuracy of QSVMs on IBM and IonQ systems due to qubit decoherence and gate noise persist (Suzuki, 2024). Additionally, variational circuits suffer from empty plateaus, demanding deeper topologies (Abbas et al., 2021). Therefore, most of the studies of QNNs are simulation-limited. Hybrid approaches, offering a good compromise between expressiveness and feasibility, become the most viable approach for current experimentation (Anonymous, 2024).

8 Challenges and Open Problems

Quantum techniques for text classification and semantic search face numerous obstacles regarding mathematical, hardware, algorithmic, and benchmarking factors. This section outlines these restrictions and identifies unresolved issues for further research.

8.1 Encoding Bottlenecks in NLP Data

Encoding high-dimensional textual data into quantum states remains one of the most significant bottlenecks. Given a document embedding $x \in \mathbb{R}^d$, amplitude encoding maps it into a normalized quantum state:

$$|x\rangle = \frac{1}{\|x\|} \sum_{i=1}^{d} x_i |i\rangle. \tag{22}$$

This requires $O(d)$ operations classically, but preparing an arbitrary d-dimensional state on a quantum computer may require $O(d)$ gates, offsetting quantum speedups. Angle encoding reduces cost by mapping each feature into a rotation, but sacrifices representational richness. Open problems include:

- Developing encoding schemes that balance expressivity with circuit depth.
- Exploring qudit-based encodings that reduce qubit requirements (Souza and Portugal, 2025).
- Designing noise-resilient encodings suitable for NISQ hardware.

8.2 Hardware Constraints and Noise Sensitivity

Most reported quantum NLP experiments have been conducted on simulators. Real NISQ devices introduce gate noise, decoherence, and readout errors. For example, QSVM implementations on IBM and IonQ hardware show a drastic drop in accuracy due to noise (Suzuki, 2024). Moreover, current devices limit circuit depth to < 100 gates for reliable execution, restricting model complexity.The open problems in NLP involve developing specific error mitigation techniques, identifying which NLP workloads are inherently noise-tolerant, such as low-rank embeddings, and exploring qudit-based systems which provide higher information density per physical unit.

8.3 Training Challenges in Quantum Neural Networks

Though the Quantum Neural Networks are a promising model, they still suffer from different serious optimization problems. The barren plateau phenomenon leads to gradients vanishing exponentially with the number of qubits or circuit depth:

$$\mathbb{E}[\nabla_\theta L] \sim O\left(\frac{1}{2^n}\right), \tag{23}$$

where n is the number of qubits. This severely limits scalability (Abbas et al., 2021). Hybrid training with classical optimizers introduces additional cost and convergence instability. Open problems include:

- Gradient-free optimization methods for variational circuits.
- Cost functions that mitigate barren plateaus.
- Scalable architectures such as QCNNs or SQ-QNNs (Cong et al., 2019; Souza and Portugal, 2025).

8.4 Scalability of Quantum IR Models

Semantic search requires efficient ranking over massive document collections. Grover's algorithm provides $O(\sqrt{N})$ query complexity, but practical retrieval requires top-k ranking and probabilistic scoring. Quantum walk retrieval models (Gupta et al., 2025) explore graph-based semantics, but remain untested at scale. Open problems include:

- Extending Grover-based search to ranked retrieval.
- Integrating density matrix entropy-based ranking (Zhang et al., 2023) with large document collections.
- Designing quantum IR systems that scale to billions of documents, analogous to web-scale search engines.

8.5 Hybrid Integration and Efficiency Boundaries

Most NISQ-era implementations are hybrid, combining classical embeddings (e.g., BERT, GloVe) with quantum classifiers or retrieval engines (Li et al., 2022; Anonymous, 2024). While effective, this raises fundamental questions:

- What portion of the pipeline truly benefits from quantum speedup?
- How can hybrid systems avoid classical bottlenecks dominating end-to-end runtime?
- What is the theoretical boundary between classical preprocessing and quantum advantage?

Establishing efficiency thresholds for hybrid quantum NLP architectures remains a critical open problem.

8.6 Evaluation and Benchmarking Gaps

There is currently no standardized framework to evaluate quantum NLP models. Classical benchmarks (e.g., GLUE, TREC) are ill-suited for quantum setups due to small dataset constraints. Open problems include:

- Designing quantum-specific NLP benchmarks with fidelity, entropy, and robustness metrics.
- Establishing evaluation protocols that combine accuracy with complexity analysis.
- Developing open-source datasets small enough for NISQ devices yet representative of real tasks.

8.7 Theoretical Uncertainty of Quantum Advantage

Finally, the biggest open problem is the lack of rigorous proof of quantum advantage in NLP. While complexity-theoretic results such as Grover's speedup are well-established, their direct applicability to semantic search and classification remains uncertain. For classification, empirical studies suggest quantum kernels offer improved separability, but no formal guarantee exists. For retrieval, entropy-based models are theoretically elegant but lack evidence of practical superiority. Future directions include:

- Proving formal conditions under which quantum models outperform classical ones.
- Linking quantum kernel theory with generalization bounds in NLP tasks.
- Exploring quantum information-theoretic limits of semantic search.

9 Conclusion and Future Directions

This work reviews quantum computing methods for natural language processing-related tasks, namely, semantic search and text classification (Havlíček et al., 2019; Li et al., 2022). Focusing on techniques such as quantum kernels, variational quantum neural networks (Abbas et al., 2021; Anonymous, 2024), and entropy-driven ranking, the study explores the trajectory from quantum-inspired models to hybrid quantum-classical systems. A comparison is drawn (in Section 6) in which retrieval performs well in both probabilistic and entropy-based models, which currently are both trending toward hybrid paradigms due to NISQ hardware limitations, whereas classification makes good use of quantum kernels. Some other promising future avenues of research involve effective encoding techniques, hardware-aware models, standardized quantum benchmarking, understanding quantum advantage, and integrated quantum NLP pipelines. Quantum computing indeed offers a great future for NLP applications, despite the challenges at present.

References

Amira Abbas, David Sutter, Christa Zoufal, Aurelien Lucchi, Alessio Figalli, and Stefan Woerner. 2021. The power of quantum neural networks. *Nature Computational Science*, 1(6):403–409.

Anonymous. 2024. Hybrid classical-quantum transfer learning for text classification. *arXiv preprint arXiv:2403.19758*.

Xiaoxiao Chen and Xiaoping Lou. 2025. Enhancing text classification through quantum transfer learning: A hybrid quantum-classical approach with complex kernel self-attention networks. *IEEE Access*.

Iris Cong, Soonwon Choi, and Mikhail D. Lukin. 2019. Quantum convolutional neural networks. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 972–981.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 4171–4186.

Lov K. Grover. 1996. A fast quantum mechanical algorithm for database search. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC)*, pages 212–219.

A. S. Gruzdeva, R. N. Iurev, I. A. Bessmertny, A. Y. Khrennikov, and A. P. Alodjants. 2025. A quantum-like approach to semantic text classification. *Entropy*, 27(7):767.

Arjun Gupta, Priya Mehta, and Rahul Singh. 2025. Quantum walk-based semantic retrieval in document graphs. *Entropy*.

Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. 2009. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103(15):150502.

Vojtěch Havlíček, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. 2019. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209–212.

G. Li. 2024. Quantum self-attention neural networks for text classification. *Science China Information Sciences*.

Wei Li, Rui Zhang, Yuchen Liu, and Haixu Wang. 2022. Quantum svm for text classification. *Quantum Information Processing*, 21(8):1–18.

Shyambabu Pandey, Partha Pakray, and Riyanka Manna. 2024. Quantum classifier for natural language processing applications. *Computación y Sistemas*, 28(2):695–700.

Benjamin Piwowarski, Ingo Frommholz, Mounia Lalmas, and Keith van Rijsbergen. 2010. A quantum-inspired framework for information retrieval. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management (CIKM)*, pages 59–68.

Maria Schuld and Francesco Petruccione. 2019. *Quantum Machine Learning: An Introduction*. Springer, Cham, Switzerland.

Leandro C. Souza and Renato Portugal. 2025. Single-qudit quantum neural networks for multiclass classification. *arXiv preprint*. ArXiv:2503.09269.

T. Suzuki. 2024. Quantum support vector machines for classification and regression: Performance on nisq hardware. *Quantum Machine Intelligence*.

Rianne van der Meer, Peter Wittek, and Jae Lee. 2021. Quantum algorithms for information retrieval: Theory and applications. *Quantum Information Processing*, 20(5):1–24.

Taro Yamada, Kenji Suzuki, and Hiroshi Tanaka. 2024. Hybrid quantum-classical semantic search with embedding fidelity. *ACM Transactions on Information Systems*.

L. Zhang, H. Chen, and J. Zhao. 2023. Entropy-based quantum information retrieval. *Information Sciences*, 634:81–94.

Guido Zuccon and Leif Azzopardi. 2019. Quantum probability ranking principle for information retrieval. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1171–1174.