

Quantum LLMs Using Quantum Computing to Analyze and Process Semantic Information

Timo Aukusti Laine*

Financial Physics Lab, Finland

*Corresponding Author

Timo Aukusti Laine, Financial Physics Lab, Finland.

Submitted: 2025, Nov 20; **Accepted:** 2025, Dec 29; **Published:** 2026, Jan 16

Citation: Laine, T. A. (2026). Quantum LLMs Using Quantum Computing to Analyze and Process Semantic Information. *OA J Applied Sci Technol*, 4(1), 01-19.

Abstract

We present a quantum computing approach to analyzing Large Language Model (LLM) embeddings, leveraging complex-valued representations and modeling semantic relationships using quantum mechanical principles. By establishing a direct mapping between LLM semantic spaces and quantum circuits, we demonstrate the feasibility of estimating semantic similarity using quantum hardware. One of the key results is the experimental calculation of cosine similarity between Google Sentence Transformer embeddings using a real quantum computer, providing a tangible demonstration of a quantum approach to semantic analysis. This work reveals a connection between LLMs and quantum mechanics, suggesting that these principles can offer new perspectives on semantic representation and processing, and paving the way for future development of quantum algorithms for natural language processing.

1. Introduction

Large Language Models (LLMs) have achieved remarkable success in natural language processing, yet their underlying mechanisms remain only partially understood. This paper explores a fundamental question: Can quantum mechanics, as a theoretical framework, provide valid and useful insights into the structure and behavior of LLMs? We hypothesize that quantum mechanics, with its inherent ability to represent complex relationships and probabilistic behavior, can offer a valuable lens for understanding and potentially manipulating LLM semantic spaces. This is particularly relevant given the increasing computational demands of LLMs, which motivates the exploration of alternative computational paradigms. Furthermore, we envision that quantum algorithms may offer significant advantages, such as the potential for exponential speedups in calculating semantic similarity within high-dimensional embedding spaces if more efficient quantum algorithms can be developed, or the ability to uncover subtle semantic relationships that are difficult to detect using classical methods. However, we acknowledge that current quantum hardware and algorithms face significant limitations, particularly in terms of qubit scaling and coherence.

Current LLMs primarily rely on real-valued embeddings, which may limit their capacity to capture subtle semantic nuances such as antonymy, sentiment, or relational information. To address this, we introduce an approach that incorporates complex-valued embeddings, leveraging principles from quantum mechanics to represent semantic information with both magnitude and phase. This framework establishes a direct mapping between LLM semantic spaces and quantum circuits, enabling the potential application of quantum algorithms to language-based tasks. This work aims to address the following research questions:

1. Can quantum mechanics, as a theoretical framework, provide valid and useful insights into the structure and behavior of LLMs, leading to a deeper understanding of their emergent properties?
2. How accurately can quantum circuits compute semantic similarity between LLM embeddings compared to classical methods, providing concrete evidence for the applicability of quantum principles to semantic analysis?

-
3. What are the key challenges and opportunities in utilizing current real quantum computers for natural language processing tasks involving LLM embeddings, considering the limitations of qubit availability and coherence?

One of the key contribution of this paper is the demonstration, using a real-world LLM (Google's Sentence Transformer), of a direct, experimentally verifiable connection between LLM embedding spaces and quantum circuits. We leverage this connection to explore real embedding spaces and calculate semantic similarity on a real quantum computer. We illustrate the principles of our approach by simulating the double-slit experiment within a quantum circuit, providing a tangible demonstration of quantum mechanics applied to semantic modeling. This provides evidence supporting the validity of a quantum-inspired approach to LLMs.

The paper is structured as follows: We begin by reviewing relevant background on LLMs, complex embeddings, and quantum mechanics. We then describe our approach for mapping LLM embeddings to quantum circuits and calculating semantic similarity. Next, we present the results of our experiments on a real quantum computer, comparing them with classical calculations. Finally, we discuss the implications of our findings, limitations, and future research directions, focusing on the broader potential of quantum mechanics as a framework for understanding and enhancing LLMs.

2. Background and Related Work

Large Language Models (LLMs) have revolutionized Natural Language Processing (NLP), achieving state-of-the-art results in various tasks such as text generation, translation, and question answering. At the heart of these models lies a sophisticated mechanism for representing text: high-dimensional vector embeddings. These embeddings map words, phrases, and even entire documents into a continuous semantic space, where geometric relationships reflect semantic similarities. The Transformer architecture, with its self-attention mechanisms, has been key to capturing long-range dependencies in text and enabling LLMs to scale to unprecedented sizes, as seen in models like BERT and GPT-3 [1-3].

While these embedding spaces are often treated as continuous for practical purposes, a fundamental aspect of LLMs hints at an underlying discreteness: their reliance on a finite vocabulary of tokens. This discrete foundation suggests that the seemingly continuous semantic space might possess a quantized structure, analogous to the discrete energy levels observed in quantum systems. This inherent quantization prompts a compelling question: can we leverage the theoretical frameworks of mathematical physics and tools of quantum mechanics to gain a deeper understanding of the organization and dynamics of these semantic spaces, and potentially uncover previously hidden relationships?

Several approaches have been proposed for analyzing LLM embedding spaces. Dimensionality reduction techniques, such as Principal Component Analysis (PCA) and t-distributed Stochastic Neighbor Embedding (t-SNE), are often used to visualize the embedding space in lower dimensions [4,5]. Geometric analysis techniques, such as calculating cosine similarity between vectors, are used to quantify the semantic similarity between words and phrases [6]. These approaches have provided valuable insights into the organization of semantic information in LLMs, but they often fail to capture more nuanced relationships such as antonymy, analogy, or hierarchical structures. The geometry of these embedding spaces has been explored by researchers like Mimno and Thompson, while Tenney et al. have investigated how context influences these embeddings [7,8].

The field of quantum computing offers a different paradigm for computation, leveraging the principles of quantum mechanics to solve problems intractable for classical computers [9]. Quantum circuits, composed of qubits and quantum gates, are used to implement quantum algorithms [10]. Key quantum algorithms, such as Shor's algorithm for factoring and Grover's algorithm for searching, demonstrate the potential for speedups in certain computational tasks [11,12]. The development of quantum hardware is rapidly advancing, with increasing qubit counts and improved coherence times [13]. However, the noisy intermediate-scale quantum (NISQ) era presents significant challenges for realizing the full potential of quantum algorithms [14].

Quantum-inspired models have also found applications in cognitive science, offering a framework for representing concepts and relationships in a more flexible manner than traditional models. These models have been used to address phenomena such as contextuality and ambiguity in human cognition [15–18]. Approach utilizes quantum formalism to represent uncertainty and superposition in cognitive processes. Furthermore, the use of complex embeddings in machine learning, particularly in knowledge graph embedding models like ComplEx, RotatE, and QuatE, provides a richer representation of data by encoding both magnitude and phase information [19-21].

More recently, the field of quantum machine learning (QML) has emerged, exploring the potential of quantum algorithms to enhance machine learning tasks. Variational Quantum Algorithms (VQAs) are a prominent approach, using parameterized quantum circuits to approximate solutions to optimization problems [14]. Quantum feature maps and kernels offer the potential to create feature spaces that are intractable for classical computers, potentially leading to improved machine learning performance [22]. Quantum neural networks (QNNs) are also being investigated as a way to leverage quantum computation for neural network training and inference [23]. However,

the development of practical QML algorithms faces significant challenges, including the limitations of current quantum hardware and the difficulty of demonstrating a quantum advantage [24].

In the realm of quantum natural language processing (QNLP), researchers are exploring ways to use quantum algorithms for NLP tasks. Quantum language models are being developed to represent and process language data in a quantum framework. Compositional QNLP focuses on combining the meanings of words and phrases using quantum-inspired mathematical structures [25]. While QNLP is still in its early stages, it holds the potential to unlock new capabilities in natural language understanding and processing.

Our work builds upon these existing approaches by exploring the analogy between LLM embedding spaces and quantum mechanics, with a focus on experimentally demonstrating the applicability of quantum principles. We extend the standard, real-valued embedding space to the complex domain to model semantic interference effects, drawing parallels to phenomena such as the double-slit experiment in quantum mechanics. We build upon the theoretical foundation laid by previous works by Laine, which introduced the concept of "semantic wave functions" to describe the representation of meaning in LLMs, drawing an analogy between words and sentences and quantum wave functions [26,27]. This paper provides a concrete implementation of these ideas using quantum circuits, demonstrates the feasibility of calculating LLM embeddings on a real quantum computer, and explores qubit reduction using complex embeddings.

3. LLM Embeddings

This section provides a review of LLM embeddings and cosine similarity, concepts that are central to this paper. LLM embeddings are vector representations of text designed to capture semantic meaning. They enable tasks such as semantic search and enhance downstream machine learning models by providing a richer and more compact representation of text compared to raw words. These embeddings are crucial for understanding and processing language beyond simple keyword matching.

In LLMs, embedding vectors are typically real-valued

$$\mathbf{E}_N(\text{text}) = [e_1, e_2, e_3, \dots, e_N], \quad (1)$$

where e_i are real numbers and N is the dimension of the embedding space. The relative similarity between two embedding vectors is often quantified using cosine similarity, a technique, for example, employed in sentence transformers and Retrieval-Augmented Generation (RAG) implementations for retrieving semantically similar context or the best-matching chunks from a vector database.

Given two real-valued vectors $\mathbf{E}_1$ and $\mathbf{E}_2$, their inner product is defined as

$$\mathbf{E}_1 \cdot \mathbf{E}_2 = \|\mathbf{E}_1\| \|\mathbf{E}_2\| \cos \theta, \quad (2)$$

which leads to the definition of cosine similarity

$$S_C(\mathbf{E}_1, \mathbf{E}_2) = \cos \theta = \frac{\mathbf{E}_1 \cdot \mathbf{E}_2}{\|\mathbf{E}_1\| \|\mathbf{E}_2\|}. \quad (3)$$

Cosine similarity is a normalized measure, with values ranging from -1 to 1. A value of 1 indicates perfectly aligned (highly similar) vectors, while -1 signifies the most dissimilar vectors. In real-valued vector embedding spaces, cosine similarity provides a valuable measure of the semantic similarity between words, sentences, or documents, assessing similarity based on subject matter, independent of the length of the documents.

3.1. Complex Embeddings

This paper explores complex representations, specifically complex number embeddings. Complex embeddings are vector representations that utilize complex numbers to encode both magnitude and phase, offering a richer data representation compared to real-valued embeddings. They are particularly effective in modeling relationships and capturing nuanced information, especially in knowledge graphs and relational reasoning tasks. While providing increased representational capacity, they also introduce computational complexity relative to real-valued embeddings. Complex embeddings are defined as

$$\mathbf{E}_{CN} = [e_{c1}e^{i\varphi_1}, e_{c2}e^{i\varphi_2}, e_{c3}e^{i\varphi_3}, \dots, e_{cN}e^{i\varphi_N}], \quad (4)$$

where e_{ci} is a real number and φ_i is an angle of rotation. Real-valued embedding vectors can be derived from complex vectors by setting $\varphi_i = 0$

$$e_i = e_{ci} e^{i\varphi_i=0} \quad (5)$$

For complex vectors, denoted as $\mathbf{a}$ and $\mathbf{b}$, the complex cosine similarity is defined as

$$\mathbf{S}_C(\mathbf{a}, \mathbf{b}) = \frac{\mathbf{a}^H \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|}, \quad (6)$$

where $\mathbf{a}^H$ denotes the conjugate transpose (Hermitian transpose) of the vector $\mathbf{a}$. The complex cosine similarity can be a real number, but it is generally a complex number.

The relative importance of the phase and magnitude components of the complex cosine similarity depends on the specific application and the definition of semantic similarity. In some cases, the magnitude is dominant, indicating that phase is less significant. This occurs when alignment strength is paramount, where the magnitude, $|\mathbf{S}_C(\mathbf{a}, \mathbf{b})|$, reflects the strength of alignment, irrespective of relative phase. Phase-invariant similarity also falls into this category, where vectors are considered similar despite phase shifts, emphasizing the magnitude. Similarly, in shape or pattern recognition, the overall shape or pattern is more important than absolute phase, making the magnitude more relevant.

Conversely, the phase component can be of primary importance. This is evident in relative timing or synchronization problems, such as with signals or oscillations, where relative timing is crucial, and the phase reflects the phase difference. Interference phenomena in physics also highlight the importance of phase, as phase differences determine interference patterns. Furthermore, when information is encoded in signal phase, phase comparison is essential, as is the case with phase-encoded information. If complex numbers represent directions, the phase encodes the angle, making it a crucial factor.

4. Double-Slit Experiment

To illustrate the concepts of embedding spaces and complex phases, this paper utilizes the double-slit experiment as an analogy. As discussed in reference [26], the double-slit experiment provides a framework for understanding potential semantic interference effects. The key aspects are illustrated in Figures 1 and 2.




Figure 1: Real-Valued Amplitude Distribution in the Double-Slit Experiment. This Serves as an Analogy for the Representation of Semantic Information in Traditional, Real-Valued LLM Embeddings

The classic double-slit experiment involves directing a stream of particles (e.g., photons or electrons) towards a barrier with two openings. The resulting pattern observed on a detection screen exhibits an interference pattern, a characteristic signature of wave behavior. Instead of a simple sum of probabilities from each slit acting independently, the probability distribution reveals alternating regions of high and

low intensity, known as interference fringes. This phenomenon arises from the superposition, or combination, of the waves emanating from each slit.

Mathematically, the wave functions describing particles passing through each slit can be expressed as

$$\psi_1(x, y) = A(x, y)e^{i\varphi}, \quad (7)$$

$$\psi_2(x, y) = B(x, y)e^{i\phi}, \quad (8)$$

where $A(x, y)$ and $B(x, y)$ represent the amplitudes of the waves at a specific location (x, y) on the detection screen, and φ and ϕ represent their respective phases. The probability of detecting a particle at that location is then given by

$$P(x, y) = |\psi_1(x, y) + \psi_2(x, y)|^2 = |A(x, y)|^2 + |B(x, y)|^2 + 2|A(x, y)||B(x, y)|\cos(\varphi - \phi). \quad (9)$$




Figure 2: Intensity Distribution on the Detection Screen in the Double-Slit Experiment, Showing the Alternating Regions of High and Low Intensity Resulting from Wave Interference

The term $2|A(x, y)||B(x, y)|\cos(\varphi - \phi)$, known as the interference term, is crucial. It arises from the superposition of complex probability amplitudes. A description relying solely on real numbers would fail to capture the phase relationships essential for explaining the interference pattern. The double-slit experiment demonstrates the waveparticle duality inherent in light and matter. The phase difference $\varphi - \phi$ is key to the interference pattern. In the context of Large Language Models, each "slit" can be interpreted as representing a distinct semantic context, and the wave function represents the word, sentence, or paragraph being analyzed. The interference pattern then reflects the potential semantic relationships between these contexts.

4.1. Double-Slit Experiment in a Quantum Circuit

This section explores a quantum circuit model of the double-slit experiment. Quantum circuits are composed of qubits and quantum gates. A qubit, the fundamental unit of quantum information, exists in a superposition of two distinct states, denoted as $|0\rangle$ and $|1\rangle$. The Hadamard gate is a single-qubit quantum gate that transforms a definite state into a superposition of equally probable $|0\rangle$ and $|1\rangle$ states. The Hadamard gate is used in quantum algorithms and is important for quantum computational advantage. In the context of the double-slit experiment, qubits can represent the possible paths of a particle, and the Hadamard gate can be used to create a superposition of these paths, analogous to the particle existing in both slits simultaneously.

To illustrate how interference manifests within a quantum circuit, consider a qubit with basis states $|0\rangle$ and $|1\rangle$. We begin with an arbitrary initial complex state

$$|\psi\rangle_{initial} = \psi_1|0\rangle + \psi_2|1\rangle = \alpha_1 e^{i\varphi_1}|0\rangle + \beta_1 e^{i\phi_1}|1\rangle, \quad (10)$$

where α_1 and β_1 are real numbers and satisfy $\alpha_1^2 + \beta_1^2 = 1$. This can be interpreted as encoding two values, with amplitudes α_1 and β_1 and phases φ_1 and ϕ_1 .

We apply a Hadamard gate, H , to this state, resulting in a superposition

$$H|\psi\rangle = H(\alpha_1 e^{i\varphi_1}|0\rangle + \beta_1 e^{i\phi_1}|1\rangle) = \frac{1}{\sqrt{2}}[\alpha_1 e^{i\varphi_1}(|0\rangle + |1\rangle) + \beta_1 e^{i\phi_1}(|0\rangle - |1\rangle)] \quad (11)$$

$$= \frac{1}{\sqrt{2}}[(\alpha_1 e^{i\varphi_1} + \beta_1 e^{i\phi_1})|0\rangle + (\alpha_1 e^{i\varphi_1} - \beta_1 e^{i\phi_1})|1\rangle]. \quad (12)$$

The probability of measuring $|0\rangle$ and $|1\rangle$ after the Hadamard gate application are

$$P(|0\rangle) = |\langle 0|H|\psi\rangle|^2 \quad (13)$$

$$= \frac{1}{2}|\alpha_1 e^{i\varphi_1} + \beta_1 e^{i\phi_1}|^2 = \frac{1}{2}(1 + 2\alpha_1\beta_1 \cos(\varphi_1 - \phi_1)), \quad (14)$$

$$P(|1\rangle) = |\langle 1|H|\psi\rangle|^2 \quad (15)$$

$$= \frac{1}{2}|\alpha_1 e^{i\varphi_1} - \beta_1 e^{i\phi_1}|^2 = \frac{1}{2}(1 - 2\alpha_1\beta_1 \cos(\varphi_1 - \phi_1)). \quad (16)$$

The probability of measuring the state $|0\rangle$ directly corresponds to the interference term at the origin in the doubleslit experiment, specifically the term $2|A||B|\cos(\varphi - \phi)$ in Eq. (9). The probability $P(|1\rangle)$ corresponds to destructive interference or dissimilarity between the two initial quantum states. The probabilities can also be written as

$$P(|0\rangle) = \frac{1}{2}|\psi_1 + \psi_2|^2, \quad (17)$$

$$P(|1\rangle) = \frac{1}{2}|\psi_1 - \psi_2|^2. \quad (18)$$

In the classical double-slit experiment, only $P(|0\rangle)$ is directly visible. However, because the total probability of the waves must be 1 after passing through the wall, $P(|1\rangle)$ must also be present after the wall. The quantum circuit diagram that models this double-slit phenomenon is shown in Figure 3.



Figure 3: A single-qubit quantum circuit demonstrating superposition. The Hadamard gate creates an equal superposition of the $|0\rangle$ and $|1\rangle$ states, analogous to the superposition of waves in the double-slit experiment

The measurement of the qubit $q[0]$ collapses the quantum state and yields values 0 or 1, depending on the state of the system. By running the quantum circuit with multiple shots and summing over the results, we obtain probabilities that correspond to the values on the states. These probabilities reflect the "similarity" or "dissimilarity" between the initial states, providing a quantum-inspired way to assess similarity.

It is important to note that this quantum circuit does not directly compute the exact complex cosine similarity. Instead, it transforms the input state (encoding amplitudes and phases) and uses measurement probabilities as indicators of similarity. Higher probabilities for specific outcomes suggest greater alignment between the components of the input state. The results are statistical estimates and are subject to fluctuations. The values obtained are not numerically equal to the exact cosine similarity but capture the concept of similarity in a different way, leveraging quantum superposition and interference. The encoding method, circuit design, and interpretation of probabilities are crucial for relating the quantum computation to semantic similarity. The potential quantum advantage lies in achieving

more efficient or accurate similarity estimation compared to classical methods, particularly for high-dimensional data, although this remains a topic for further investigation.

5. Quantum Circuit for Complex Cosine Similarity

This section details the construction of a quantum circuit designed to estimate complex cosine similarity. Building upon the connection established in the previous section between the interference term and a simple quantum circuit, we now aim to connect complex cosine similarity to quantum systems. This requires: (1) identifying the necessary quantum circuit components for the measurement, and (2) deriving equations that relate the measurement results to the complex cosine similarity.

5.1. Quantum Circuit Components for Similarity Measurement

We begin by identifying the quantum circuit components required for the measurement. We start from the definition of complex cosine similarity in Eq. (6)

$$S_C(\mathbf{a}, \mathbf{b}) = \frac{1}{\|\mathbf{a}\| \|\mathbf{b}\|} \sum_{i=1}^N a_{ci}^* b_{ci} \quad (19)$$

We assume that vectors $\mathbf{a}$ and $\mathbf{b}$ are expressed as

$$\mathbf{a} = [a_{c1}, a_{c2}, \dots, a_{cN}] = [a_1 e^{i\varphi_1}, a_2 e^{i\varphi_2}, \dots, a_N e^{i\varphi_N}]. \quad (20)$$

$$\mathbf{b} = [b_{c1}, b_{c2}, \dots, b_{cN}] = [b_1 e^{i\phi_1}, b_2 e^{i\phi_2}, \dots, b_N e^{i\phi_N}], \quad (21)$$

where a_i and b_i are real numbers. We also assume L2 normalization of the vectors

$$\|\mathbf{a}\|^2 = \sum_{i=1}^N a_i^2 = 1, \quad (22)$$

$$\|\mathbf{b}\|^2 = \sum_{i=1}^N b_i^2 = 1. \quad (23)$$

Expanding the equation, we obtain

$$S_C(\mathbf{a}, \mathbf{b}) = \sum_{i=1}^N (a_i b_i \cos(\phi_i - \varphi_i)) + i a_i b_i \sin(\phi_i - \varphi_i) \quad (24)$$

We now consider quantum circuits capable of generating these terms. The circuit in Figure 3 can be used to produce results for the cosine term. To generate the sine term, we observe that applying a $\pi/2$ rotation to one of the initial states yields the desired result. Therefore, we consider the circuit shown in Figure 4. This circuit first applies the S -gate, which rotates the phase of state $|1\rangle$, followed by the Hadamard gate, producing a superposition between the states $|0\rangle$ and $|1\rangle$.




Figure 4: Quantum circuit demonstrating phase manipulation. The S -gate introduces a $\pi/2$ phase shift to the $|1\rangle$ state, enabling the estimation of both cosine and sine components of complex similarity

Using an arbitrary initial complex state as defined in Eq. (10) and applying the steps in the circuit, we obtain the following state probabilities

$$P(|0\rangle) = \frac{1}{2}(1 + 2\alpha_1\beta_1 \sin(\varphi_1 - \phi_1)), \quad (25)$$

$$P(|1\rangle) = \frac{1}{2}(1 - 2\alpha_1\beta_1 \sin(\varphi_1 - \phi_1)). \quad (26)$$

The probability $P(|0\rangle)$ provides a quantum-inspired estimate related to the imaginary part of the component-wise complex cosine similarity, specifically the term involving $\sin(\varphi_1 - \phi_1)$. Note that

$$\sin(\phi_i - \varphi_i) = -\sin(\varphi_i - \phi_i). \quad (27)$$

Accurate calculations require careful attention to the signs of the phases. We now have the quantum circuit components, shown in Figures 3 and 4, from which we can construct a circuit to estimate the complex cosine similarity.

5.2. Derivation of the Complex Cosine Similarity Equation

Our next aim is to derive an equation for the complex cosine similarity, $\mathbf{S}_C$, that can be effectively implemented using quantum circuit measurements. We start from Eq. (24)

$$\mathbf{S}_C(\mathbf{a}, \mathbf{b}) = \sum_{i=1}^N a_i b_i e^{i(\phi_i - \varphi_i)} = \sum_{i=1}^N a_i b_i [\cos(\phi_i - \varphi_i) + i \sin(\phi_i - \varphi_i)]. \quad (28)$$

The key challenge lies in mapping these complex vector components to qubit states. Qubit states are represented by amplitudes, and these amplitudes must satisfy the normalization condition (sum of squares equals 1). To achieve this, we introduce a scaling factor, c_i , for each dimension i to ensure that the scaled values can represent valid qubit amplitudes.

Let's define scaled values α_i and β_i as

$$a_i = c_i \alpha_i \quad (29)$$

$$b_i = c_i \beta_i \quad (30)$$

The crucial requirement is that for each dimension i , these scaled values satisfy the qubit normalization condition

$$\alpha_i^2 + \beta_i^2 = 1. \quad (31)$$

This ensures that α_i and β_i can be interpreted as the amplitudes of a qubit state. Substituting the definitions of α_i and β_i , we get

$$c_i^2 = a_i^2 + b_i^2 \quad (32)$$

Using the results from Eqs. (14) and (25), we can express the cosine and sine terms in terms of the probabilities obtained from quantum circuit measurements. Specifically, we can rewrite Eq. (14) as

$$P(|0\rangle)_i = \frac{1}{2}(1 + 2\alpha_i\beta_i \cos(\varphi_i - \phi_i)), \quad (33)$$

from which

$$a_i b_i \cos(\varphi_i - \phi_i) = \frac{c_i^2}{2}(2P(|0\rangle)_i - 1). \quad (34)$$

Similarly, we can derive an equation for the sine term, and $\mathbf{S}_C$ becomes

$$\mathbf{S}_C(\mathbf{a}, \mathbf{b}) = \frac{1}{2} \sum_{i=1}^N c_i^2 (2P(|0\rangle)_{2i-2} - 1) - \frac{i}{2} \sum_{i=1}^N c_i^2 (2P(|0\rangle)_{2i-1} - 1), \quad (35)$$

where $P(|0\rangle)_{2i-2}$ and $P(|0\rangle)_{2i-1}$ are the probabilities of measuring the $|0\rangle$ state from the quantum circuits in Figures 3 and 4 for dimension i ,

respectively. Using quantum circuit controlled rotations, or multiple quantum circuits, one could potentially eliminate c_i from the right-hand side of Eq. (35), but we leave it as is.

The magnitude of the complex cosine similarity can be calculated as

$$|S_C(\mathbf{a}, \mathbf{b})| = \sqrt{\left[\frac{1}{2} \sum_{i=1}^N c_i^2 (2P(|0\rangle)_{2i-2} - 1) \right]^2 + \left[\frac{1}{2} \sum_{i=1}^N c_i^2 (2P(|0\rangle)_{2i-1} - 1) \right]^2} \quad (36)$$

This equation provides a quantum-inspired method for estimating the complex cosine similarity between two vectors, leveraging the principles of quantum superposition and interference.

5.3. Density Matrix Interpretation and Semantic Feature Analysis

Equation (35) demonstrates that the coefficients c_i^2 are not arbitrary. Let's use the matrix formulation to write the same

$$\begin{aligned} S_C(\mathbf{a}, \mathbf{b}) &= \frac{1}{2} \begin{bmatrix} c_1^2 & 0 & 0 & 0 & 0 \\ 0 & c_2^2 & 0 & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 \\ 0 & 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & 0 & c_N^2 \end{bmatrix} \begin{bmatrix} 2P(|0\rangle)_0 - 1 \\ 2P(|0\rangle)_2 - 1 \\ \vdots \\ \vdots \\ 2P(|0\rangle)_{2N-2} - 1 \end{bmatrix} - i \frac{1}{2} \begin{bmatrix} c_1^2 & 0 & 0 & 0 & 0 \\ 0 & c_2^2 & 0 & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 \\ 0 & 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & 0 & c_N^2 \end{bmatrix} \begin{bmatrix} 2P(|0\rangle)_1 - 1 \\ 2P(|0\rangle)_3 - 1 \\ \vdots \\ \vdots \\ 2P(|0\rangle)_{2N-1} - 1 \end{bmatrix} \\ &= \boldsymbol{\rho}_c \mathbf{P}_c - i \boldsymbol{\rho}_c \mathbf{P}_i \end{aligned} \quad (37)$$

$$(38)$$

where

$$\boldsymbol{\rho}_c = \frac{1}{2} \begin{bmatrix} c_1^2 & 0 & 0 & 0 & 0 \\ 0 & c_2^2 & 0 & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 \\ 0 & 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & 0 & c_N^2 \end{bmatrix} \quad (39)$$

The matrix $\boldsymbol{\rho}_c$ can be interpreted as a quantum density matrix, describing the statistical state of a quantum system. A key property of density matrices is that their trace is equal to 1

$$\text{Tr}(\boldsymbol{\rho}_c) = 1. \quad (40)$$

In this context, $\boldsymbol{\rho}_c$ represents the uncertainty in our knowledge of the embedding vectors. Each diagonal element, $c_i^2/2$, weights the contribution of the i -th dimension to the similarity calculation, reflecting the combined magnitude ($a_i^2 + b_i^2$) of the corresponding components in vectors $\mathbf{a}$ and $\mathbf{b}$. The diagonal form of the density matrix indicates that the dimensions are treated as statistically independent in the component-wise similarity calculation.

If the original embedding vectors $\mathbf{a}$ and $\mathbf{b}$ are L2 normalized, then

$$\text{Tr}(\boldsymbol{\rho}_c) = \frac{1}{2} \sum_{i=1}^N c_i^2 = \frac{1}{2} \sum_{i=1}^N (a_i^2 + b_i^2) = \frac{1}{2} \sum_{i=1}^N a_i^2 + \frac{1}{2} \sum_{i=1}^N b_i^2 = 1. \quad (41)$$

This constraint, imposed by the properties of density matrices and L2 normalization, suggests that the individual components a_i^2 and b_i^2 might also have a specific meaning beyond simply being squared components of $\mathbf{a}$ vector. Specifically, if we interpret the a_i^2 values as the diagonal elements (and therefore eigenvalues) of a matrix representing the embedding vector $\mathbf{a}$ in semantic space, then it suggests that each dimension of the embedding space contributes a specific, quantifiable amount to the overall semantic representation of the word.

Based on quantum density matrix definitions, many established results from quantum mechanics become applicable [28]. In quantum mechanics, an observable represents a physical quantity that can be measured, while its expectation value is the average value one would

obtain from many measurements on identically prepared systems. If S_C is an observable of the system, then its expectation value is given by

$$\langle S_C \rangle = \text{Tr}(\rho_c S_C) \quad (42)$$

with the spectral resolution

$$S_C = \sum_{i=1}^N c_i^2 P_i \quad (43)$$

where P_i is the projection operator onto the eigenspace corresponding to eigenvalue c_i^2 . Analogously, the vectors

$$\begin{bmatrix} 2P(|0\rangle)_0 - 1 \\ 2P(|0\rangle)_2 - 1 \\ \vdots \\ 2P(|0\rangle)_{2N-2} - 1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 2P(|0\rangle)_1 - 1 \\ 2P(|0\rangle)_3 - 1 \\ \vdots \\ 2P(|0\rangle)_{2N-1} - 1 \end{bmatrix}$$

can be interpreted as projection operators in the eigenspace of the qubits.

The observable S_C has eigenvalues c_i^2 , representing the contribution of each dimension to the complex cosine similarity. The density matrix ρ_c , on the other hand, also possesses eigenvalues. These eigenvalues of ρ_c represent the probabilities of finding the system in a particular eigenstate, which can be interpreted as the prominence of specific semantic features within the embedded words. The magnitude of each eigenvalue reflects the importance of the corresponding semantic feature in characterizing the words. The eigenvectors of ρ_c provide a basis for the semantic space, with each eigenvector corresponding to a direction or axis, and its components indicating the contribution of each dimension to that axis. A detailed exploration of these consequences will be addressed in future work.

6. Illustrative Example: Semantic Similarity between "DOG" and "CAT"

This section provides an illustrative example of calculating the semantic similarity between the words "dog" and "cat" using the quantum circuit approach described above. To simplify the demonstration, we represent each word using a two-dimensional complex embedding. The embedding vectors are defined as


$$\mathbf{E}_1(\text{dog}) = [a_1 e^{i\varphi_1}, a_2 e^{i\varphi_2}], \quad (44)$$

$$\mathbf{E}_2(\text{cat}) = [b_1 e^{i\phi_1}, b_2 e^{i\phi_2}]. \quad (45)$$

The embedding vectors are normalized such that $|\mathbf{E}_i|^2 = 1$. The complex cosine similarity is then given by

$$S_C(\mathbf{a}, \mathbf{b}) = a_1 b_1 \cos(\phi_1 - \varphi_1) + i a_1 b_1 \sin(\phi_1 - \varphi_1) + a_2 b_2 \cos(\phi_2 - \varphi_2) + i a_2 b_2 \sin(\phi_2 - \varphi_2). \quad (46)$$

To calculate the complex cosine similarity, we require a quantum circuit to estimate each term in the equation. Several circuit architectures are possible. We could use a single qubit and execute the quantum job four times, use two qubits and execute separate jobs for the real and imaginary parts, or use four qubits to calculate all values in a single job. For this example, we select the four-qubit approach. The quantum circuit is shown in Figure 5. The qubits are organized to correspond to the terms in Eq. (46).






Figure 5: Quantum circuit for calculating the complex cosine similarity between simplified "dog" and "cat" embeddings. Qubits 0 and 2 estimate the real components, while qubits 1 and 3, incorporating S-gates, estimate the imaginary components.

In this circuit, the first qubit estimates the real part of the first term in Eq. (46), and the second qubit estimates the imaginary part, and so on. The qubits are independent and do not interact with each other. For larger embeddings with $2N$ dimensions, the circuit can be constructed using $2N$ qubits, following the same logic.

The initial states for the quantum circuit are

$$|q[0]\rangle_{\text{initial}} = \alpha_1 e^{i\varphi_1} |0\rangle + \beta_1 e^{i\phi_1} |1\rangle \quad (47)$$

$$|q[1]\rangle_{\text{initial}} = \alpha_1 e^{i\varphi_1} |0\rangle + \beta_1 e^{i\phi_1} |1\rangle \quad (48)$$

$$|q[2]\rangle_{\text{initial}} = \alpha_2 e^{i\varphi_2} |0\rangle + \beta_2 e^{i\phi_2} |1\rangle \quad (49)$$

$$|q[3]\rangle_{\text{initial}} = \alpha_2 e^{i\varphi_2} |0\rangle + \beta_2 e^{i\phi_2} |1\rangle \quad (50)$$

where α_i and β_i are real numbers that satisfy $\alpha_i^2 + \beta_i^2 = 1$. We encode the word embeddings into the quantum circuit. After applying the phase and Hadamard gates, we obtain

$$|q[0]\rangle = \frac{1}{\sqrt{2}} [(\alpha_1 e^{i\varphi_1} + \beta_1 e^{i\phi_1})|0\rangle + (\alpha_1 e^{i\varphi_1} - \beta_1 e^{i\phi_1})|1\rangle] \quad (51)$$

$$|q[1]\rangle = \frac{1}{\sqrt{2}} [(\alpha_1 e^{i\varphi_1} + i\beta_1 e^{i\phi_1})|0\rangle + (\alpha_1 e^{i\varphi_1} - i\beta_1 e^{i\phi_1})|1\rangle] \quad (52)$$

$$|q[2]\rangle = \frac{1}{\sqrt{2}} [(\alpha_2 e^{i\varphi_2} + \beta_2 e^{i\phi_2})|0\rangle + (\alpha_2 e^{i\varphi_2} - \beta_2 e^{i\phi_2})|1\rangle] \quad (53)$$

$$|q[3]\rangle = \frac{1}{\sqrt{2}} [(\alpha_2 e^{i\varphi_2} + i\beta_2 e^{i\phi_2})|0\rangle + (\alpha_2 e^{i\varphi_2} - i\beta_2 e^{i\phi_2})|1\rangle] \quad (54)$$

The final state is a tensor product of these qubit states

$$|\psi\rangle_{\text{final}} = |q[0]\rangle \otimes |q[1]\rangle \otimes |q[2]\rangle \otimes |q[3]\rangle. \quad (55)$$

The probabilities of measuring the qubit states after a single shot are

$$P(q[0] \rightarrow 0) = \frac{1}{2}(1 + 2\alpha_1\beta_1 \cos(\varphi_1 - \phi_1)) \quad (56)$$

$$P(q[0] \rightarrow 1) = \frac{1}{2}(1 - 2\alpha_1\beta_1 \cos(\varphi_1 - \phi_1)) \quad (57)$$

$$P(q[1] \rightarrow 0) = \frac{1}{2}(1 + 2\alpha_1\beta_1 \sin(\varphi_1 - \phi_1)) \quad (58)$$

$$P(q[1] \rightarrow 1) = \frac{1}{2}(1 - 2\alpha_1\beta_1 \sin(\varphi_1 - \phi_1)) \quad (59)$$

$$P(q[2] \rightarrow 0) = \frac{1}{2}(1 + 2\alpha_2\beta_2 \cos(\varphi_2 - \phi_2)) \quad (60)$$

$$P(q[2] \rightarrow 1) = \frac{1}{2}(1 - 2\alpha_2\beta_2 \cos(\varphi_2 - \phi_2)) \quad (61)$$

$$P(q[2] \rightarrow 1) = \frac{1}{2}(1 - 2\alpha_2\beta_2 \cos(\varphi_2 - \phi_2)) \quad (61)$$

$$P(q[3] \rightarrow 0) = \frac{1}{2}(1 + 2\alpha_2\beta_2 \sin(\varphi_2 - \phi_2)) \quad (62)$$

We use Qiskit and the quantum Aer-simulator to calculate the results. Qiskit is an open-source Python SDK for quantum computing that enables users to design, simulate, and execute quantum algorithms on both simulators and real quantum hardware. It offers a modular approach with tools for quantum algorithm development, simulation, error correction, and applications in fields like chemistry, machine learning, finance, and optimization. To demonstrate the approach, we assume an embedding model produces the following vectors

$$\mathbf{E}_1(\text{dog}) = [0.4e^{i\pi/6}, \sqrt{1 - 0.4^2}e^{i\pi/2}], \quad (64)$$

$$\mathbf{E}_2(\text{cat}) = [0.5e^{i\pi/4}, \sqrt{1 - 0.5^2}e^{i\pi/3}]. \quad (65)$$

These vectors have been selected to be arbitrary, solely to illustrate the approach and assess its accuracy. In the next section, we present a more realistic case. We shoot the words to the double-slit experiment. Applying Qiskit, the simulator provides estimates for the probabilities in Eqs. (56)-(63), given the initial conditions in Eqs. (47)-(50).

	Exact results	Simulation
q[0], state 0	0.971	0.973
q[0], state 1	0.029	0.027
q[1], state 0	0.374	0.371
q[1], state 1	0.626	0.629
q[2], state 0	0.932	0.936
q[2], state 1	0.068	0.064
q[3], state 0	0.750	0.751
q[3], state 1	0.250	0.249

Table 1: Measurement probabilities obtained from theoretical calculations and a quantum simulator for the "dog" and "cat" example, demonstrating the ability of quantum simulation to approximate the expected quantum behavior

We have calculated the exact final probabilities of the qubit states using Eqs. (56)-(63). The results are shown in Table 1, along with the experimental results obtained using the Qiskit quantum simulator. We used 10000 shots for the simulation.

The complex cosine similarity is calculated as

$$\begin{aligned} \mathbf{S}_C(\mathbf{a}, \mathbf{b}) = & \frac{1}{2} [c_1^2(2P(q[0] \rightarrow 0) - 1) + c_2^2(2P(q[2] \rightarrow 0) - 1)] \\ & - \frac{i}{2} [c_1^2(2P(q[1] \rightarrow 0) - 1) + c_2^2(2P(q[3] \rightarrow 0) - 1)]. \end{aligned} \quad (66)$$

Using the exact probabilities, the complex similarity is

$$\mathbf{S}_C(\text{exact}) = 0.881 - i0.345 \quad (67)$$

Using the probabilities obtained from the simulation, the complex similarity is

$$\langle \mathbf{S}_C(\text{simulated}) \rangle = 0.886 - i0.346 \quad (68)$$

The results from the simulation are in close agreement with the exact calculation. The exact results are based on direct calculation using the equations, while the simulated results are obtained from the quantum computer simulator, which provides probabilities and statistical averages.

Depending on the specific use case, the simulated result may be sufficiently accurate. For example, cosine similarity is often used to retrieve the top k best matches in RAG systems, where high precision of S_c is not necessary required. The accuracy also depends on the corpus used to train the embedding model. If the training corpus is large, the embedding space is dense, meaning there are many states that are close to each other. In such cases, small differences in the decimal places of the cosine similarity can be meaningful.

7. Quantum Computation of Cosine Similarity with a Real LLM

This section presents the results of computing cosine similarity between embedding vectors generated by a realworld LLM on a real quantum computer. We used Google’s Sentence Transformer, ‘google/embedding-gemma-300m’ (available on Hugging Face), to generate the embedding vectors. EmbeddingGemma, a 300M parameter model, leverages Gemma 3 (with T5Gemma initialization) and the technology used in Gemini models to produce vector representations of text suitable for tasks like search, retrieval, classification, clustering, and semantic similarity search. Trained on data in over 100 languages, EmbeddingGemma outputs vectors with a default dimension of 768. This model was selected as a representative example of a modern, accurate, and widely used sentence embedding model. The specific choice of model is not critical to the demonstration of the quantum computation; any comparable model producing suitable embedding vectors could be used. To reduce the qubit requirements for this demonstration, we employed Matryoshka Representation Learning (MRL) to truncate the output to 128 dimensions. MRL was chosen to reduce the dimensionality of the embeddings to match the qubit constraints of the available quantum hardware. The quantum circuit used is shown in Figure 6.




Figure 6: Quantum circuit used to calculate cosine similarity between embeddings generated by Google’s Sentence Transformer. The 128 qubits represent the 128 dimensions of the MRL-truncated embedding vectors.

The quantum computations were performed using the IBM Quantum Cloud, which provides access to IBM’s quantum computing systems and resources. This platform allows researchers, developers, and educators to explore and advance quantum computing without the need for specialized hardware. It offers access to quantum processors, the Qiskit software development kit, and a collaborative environment.

As a concrete example, we calculated the cosine similarity between the following sentences:

“the quick brown fox jumps over the lazy dog”
and
“the brown dog jumps over the lazy cat”

These sentences were chosen to illustrate the model’s ability to capture semantic similarity despite variations in wording. The Google Sentence Transformer generated 128-dimensional embedding vectors for these sentences (after MRL projection). These vectors were then normalized to unit length, a standard procedure before calculating cosine similarity. Calculating the cosine similarity required 128 qubits. The initial state was prepared using amplitude encoding, and a quantum circuit with Hadamard gates was employed. The Python program, developed using Qiskit, was executed on the IBM Quantum Cloud, utilizing 4096 shots with the ibm fez backend. The choice of the backend was based on its availability and sufficient qubit count at the time of execution. The resulting quantum computation yielded a cosine similarity of

$$\langle S_C(\text{quantum}) \rangle = 0.8635 \quad (69)$$

The corresponding theoretical result, calculated classically, is

$$S_C(\text{exact}) = 0.8682 \quad (70)$$

The cosine similarity is a real number in this case due to the use of real-valued embeddings. The results demonstrate reasonable agreement between the quantum computation and the classical calculation, suggesting that the quantum circuit is capturing the essential features of the semantic similarity between the sentences. However, a small discrepancy exists between the quantum and classical results. It's important to acknowledge that current quantum hardware is susceptible to noise and errors, which can impact the accuracy of computations.

7.1. Addressing Noise and Error in Quantum Computations

The experimental validation of our quantum-inspired approach to LLM embedding analysis, while promising, is inherently limited by the presence of noise and errors in current quantum hardware. These imperfections can significantly impact the accuracy and reliability of quantum computations, leading to deviations from the theoretically expected results. Several factors contribute to these errors, including decoherence, where qubit coherence times are finite, meaning that the quantum state of a qubit can degrade over time due to interactions with the environment; gate errors, where quantum gates are not perfect and can introduce errors during their execution; measurement errors, where the measurement process itself can introduce errors; and crosstalk, where unintended interactions between qubits can lead to errors. The overall fidelity of the quantum computation, a measure of how closely the actual quantum state matches the ideal state, is thus reduced by these imperfections.

To mitigate the effect of noise and errors, several techniques can be employed. Error mitigation techniques aim to reduce the impact of errors without requiring full quantum error correction. Examples include zero-noise extrapolation (ZNE), probabilistic error cancellation (PEC), and readout error mitigation. Quantum error correction (QEC) involves encoding quantum information in a redundant manner, allowing for the detection and correction of errors. Advancements in quantum hardware, such as longer coherence times, more accurate gates, and reduced crosstalk, will also contribute to reducing noise and errors in quantum computations.

Future research should focus on implementing and evaluating these error mitigation techniques in the context of quantum-enhanced LLM embedding analysis, and rigorously assessing the statistical significance of the differences between quantum and classical results, accounting for noise and error. Furthermore, as quantum hardware continues to improve, the impact of noise and errors will be reduced, making quantum computations more accurate and reliable.

7.2. Code Listing

This subsection shows the Python code that was used in quantum computer calculation.

Listing 1: Qiskit code for quantum cosine similarity calculation.

```

1 import numpy as np
2 import pickle
3 from qiskit import QuantumCircuit, transpile
4 from qiskit_ibm_runtime import QiskitRuntimeService, Session, Sampler
5
6 def run_and_store_results(vectora_q, vectorb_q, filename='quantum_results.pkl'):
7
8     if len(vectora_q) != len(vectorb_q):
9         raise ValueError("Vectors must have the same length.")
10
11     num_qubits = len(vectora_q)
12     qc = QuantumCircuit(num_qubits, num_qubits)
13
14     for i in range(num_qubits):
15         a = vectora_q[i]
16         b = vectorb_q[i]
17
18         initial_state = np.array([a, b], dtype=np.complex128)
19         norm = np.linalg.norm(initial_state)
20         normalized_state = initial_state / norm
21
22         qc.initialize(normalized_state, i)

```

```

23
24 qc.h(range(num_qubits))
25 qc.measure(range(num_qubits), range(num_qubits))
26
27 try:
28     service = QiskitRuntimeService()
29
30     backend = service.least_busy(simulator=False, operational=True, min_num_qubits=num_qubits)
31     print(f"Using backend: {backend.name}")
32
33     transpiled_qc = transpile(qc, backend)
34
35     with Session(backend=backend) as session:
36         sampler = Sampler()
37
38         job = sampler.run([transpiled_qc], shots=4096)
39         print(f"Job ID: {job.job_id()}")
40         result = job.result()
41
42         with open(filename, 'wb') as f:
43             pickle.dump(result, f)
44
45         print(f"Results stored in {filename}")
46
47 except Exception as e:
48     print(f"Error running the circuit: {e}")
49     print(f"Error type: {type(e)}")
50     print(f"Error message: {e}")

```

7.3. Explanation

The `run_and_store_results` function takes two vectors, `vectora_q` and `vectorb_q`, as input, representing the amplitudes used to initialize the qubits. The function begins by checking if the input vectors have the same length. It then creates a `QuantumCircuit` with the number of qubits equal to the length of the input vectors. For each qubit, the state is initialized using the corresponding amplitudes from `vectora_q` and `vectorb_q` through amplitude encoding, where the values from the vectors are directly used to define the initial quantum state, normalizing the initial state to ensure it represents a valid quantum state. Subsequently, a Hadamard gate (`qc.h`) is applied to each qubit, creating a superposition essential for leveraging quantum interference to estimate the cosine similarity, and all qubits are measured. The function then connects to the IBM Quantum Cloud using `QiskitRuntimeService()`, selects a real quantum backend using `service.least_busy()`, and transpiles the circuit for the selected backend using `transpile()`. A session is created with the backend, and a `Sampler` instance is created. The circuit is run using the `Sampler` within the session, which efficiently estimates probabilities without returning the full state vector. Finally, the results are stored in a pickle file. A try-except block is included to catch any errors during the quantum circuit execution, printing the error type and message if an error occurs. The key aspects of this implementation include the use of amplitude encoding for qubit initialization, the application of the Hadamard gate to leverage quantum interference, and the use of the `Sampler` primitive for efficient probability estimation.

8. Complex Embeddings and the Full Quantum Mechanical Framework

Given that LLM embeddings can be real-valued, one might question the necessity of quantum mechanics, which relies on complex values. This section explores the potential for leveraging complex embeddings to more fully utilize the quantum mechanical framework when calculating semantic similarity. A $2N$ -dimensional real vector can be represented as an N -dimensional complex vector using the following transformations:

$$\alpha_i = \sqrt{a_{2i-1}^2 + a_{2i}^2} \quad (71)$$

$$\varphi_i = \begin{cases} \arccos(a_{2i-1}/\alpha_i) & \text{if } a_{2i} \geq 0 \\ -\arccos(a_{2i-1}/\alpha_i) & \text{if } a_{2i} < 0 \end{cases} \quad (72)$$

$$a_{2-i1} = \alpha_i \cos(\varphi_i) \quad (73)$$

$$a_{2i} = \alpha_i \sin(\varphi_i) \quad (74)$$

While this representation allows a more compact encoding, the computational cost on a quantum computer is contingent on the embedding type.

Consider complex embeddings with real coefficients, $\mathbf{a} = [a_1 e^{i\varphi_1}, \dots, a_N e^{i\varphi_N}]$ and $\mathbf{b} = [b_1 e^{i\phi_1}, \dots, b_N e^{i\phi_N}]$. The complex cosine

similarity can be expressed as

$$\mathbf{S}_C(\mathbf{a}, \mathbf{b}) = \sum_{i=1}^N a_i b_i e^{i(\phi_i - \varphi_i)} \quad (75)$$

We can factor out the first term

$$\mathbf{S}_C(\mathbf{a}, \mathbf{b}) = a_1 b_1 e^{i(\phi_1 - \varphi_1)} + \sum_{i=2}^N a_i b_i e^{i(\phi_i - \varphi_i)} \quad (76)$$

$$= e^{i(\phi_1 - \varphi_1)} \left[a_1 b_1 + \sum_{i=2}^N a_i b_i e^{i(\phi_i - \varphi_i) - i(\phi_1 - \varphi_1)} \right]. \quad (77)$$

Let's define $\tilde{\mathbf{S}}_C(\tilde{\mathbf{a}}, \tilde{\mathbf{b}})$ as

$$\tilde{\mathbf{S}}_C(\tilde{\mathbf{a}}, \tilde{\mathbf{b}}) = a_1 b_1 + \sum_{i=2}^N a_i b_i e^{i(\phi_i - \varphi_i) - i(\phi_1 - \varphi_1)} \quad (78)$$

Then

$$\mathbf{S}_C(\mathbf{a}, \mathbf{b}) = e^{i(\phi_1 - \varphi_1)} \tilde{\mathbf{S}}_C(\tilde{\mathbf{a}}, \tilde{\mathbf{b}}). \quad (79)$$

Taking the magnitude gives

$$|\mathbf{S}_C(\mathbf{a}, \mathbf{b})| = |\tilde{\mathbf{S}}_C(\tilde{\mathbf{a}}, \tilde{\mathbf{b}})|. \quad (80)$$

This allows us to focus on calculating $|\tilde{\mathbf{S}}_C(\tilde{\mathbf{a}}, \tilde{\mathbf{b}})|$. By calculating $a_1 b_1$ using one qubit, we can, in principle, reduce the total qubit count requirement by one. However, the primary significance of this transformation lies not in the qubit reduction itself, but in the fact that it allows us to operate within the full quantum mechanical framework. Even starting with real embeddings, by transforming to a complex representation, we gain access to the full power of quantum mechanics, including the ability to manipulate phases and exploit interference effects.

When using a complex representation, if the input vectors are real, the resulting cosine similarity will also be a real number. The imaginary component can be viewed as a helper dimension. It provides additional degrees of freedom that can be exploited within a quantum circuit. It allows for operations like phase flips and rotations that can manipulate the quantum state in ways not possible with purely real-valued representations. This added flexibility could potentially lead to the design of more efficient and tailored quantum circuits for semantic similarity calculations, even if the initial qubit savings are minimal.

While this initial approach focuses on utilizing the full quantum mechanical framework, significant qubit reduction ultimately requires more sophisticated techniques that exploit the underlying structure of the problem. Several techniques can be employed to achieve qubit reduction. Exploiting sparsity, if the embedding vectors exhibit a high degree of sparsity, techniques like compressed sensing or sparse quantum algorithms could substantially reduce the qubit requirement. The potential reduction depends on the degree of sparsity in the embedding vectors. Quantum Phase Estimation (QPE) offers another approach, providing a way to estimate eigenvalues related to cosine similarity. The number of qubits required for QPE scales logarithmically with the desired precision, potentially offering significant savings compared to direct amplitude encoding. Variational Quantum Algorithms (VQAs) present a third possibility, as they can be designed to directly learn the cosine similarity between embedding vectors. The qubit requirement for VQAs depends on the complexity of the chosen ansatz (parameterized quantum circuit). Quantum Singular Value Decomposition (QSVD) can decompose the matrix formed by the embedding vectors, allowing for cosine similarity estimation based on singular values. The qubit requirement for QSVD depends on the size of the matrix being decomposed.

9. Potential Applications and Future Research Directions

The presented cosine similarity calculation using a quantum computer alone is not currently practical, as classical methods are faster for real-valued embeddings. However, the quantum-inspired approach presented in this paper could become relevant as part of a larger

quantum-classical process, such as training sentence transformers.

Sentence transformers generate dense vector representations for sentences, capturing semantic meaning. They are fine-tuned on large datasets to outperform generic transformer models on tasks like semantic search and text classification. Training involves optimizing sentence embeddings using sentence pairs or triplets and evaluating performance on downstream tasks. Quantum algorithms could accelerate the training process by efficiently optimizing the embedding space or providing better initial parameter estimates. Quantum optimization algorithms could be used to find the optimal parameters for the sentence transformer model. This is a relatively unexplored area.

Other potential applications and future research directions include: Implicit feature maps and quantum kernels for enhanced representation. If the qubit amplitudes are the result of a complex feature map, a quantum computer could implement this map implicitly using a quantum kernel. This could lead to a more efficient feature representation, as quantum kernels can perform computations in high-dimensional spaces intractable for classical computers. Variational Quantum Eigensolver (VQE) or Quantum Approximate Optimization Algorithm (QAOA) could be used to find the optimal feature map. Designing effective quantum kernels remains a challenge.

Entanglement for high-dimensional semantic analysis is another possibility. Quantum computers can create and manipulate entangled states in high-dimensional Hilbert spaces, potentially allowing cosine similarity calculations between complex quantum states that are impossible to process classically. The Hilbert space of N qubits has dimension 2^N . Building and controlling large numbers of qubits is a major technological challenge.

Quantum amplitude estimation for high-precision similarity measurement is also promising. If high precision is required, Quantum Amplitude Estimation (QAE) can estimate probabilities with a speedup compared to classical Monte Carlo methods. QAE uses Grover's search algorithm to estimate the amplitude of a quantum state. QAE requires a relatively large number of qubits and coherent gate operations.

Hybrid quantum-classical approaches for semantic analysis offer a promising direction. Combining the strengths of both quantum and classical computation, hybrid algorithms could leverage quantum computers for specific tasks, such as similarity calculations or feature extraction. Designing effective hybrid algorithms requires careful consideration of the strengths and weaknesses of both quantum and classical computation.

10. Discussion

This paper has explored the application of quantum formalism to the analysis of LLM embeddings, providing theoretical insights and experimental validation of a connection between these fields. We demonstrated the estimation of cosine similarity between LLM embeddings using quantum circuits and a real quantum computer, suggesting a potential pathway for quantum-enhanced semantic analysis. Quantum algorithms, specifically designed for this task, could unlock advantages in discovering subtle semantic relationships or accelerating similarity calculations in high-dimensional spaces.

Our exploration suggests that LLM embedding spaces exhibit characteristics that resonate with quantum mechanics. The finite vocabulary of tokens hints at an underlying discreteness, aligning with the concept of quantization. By extending the embedding space to the complex domain and drawing parallels to the double-slit experiment, we modeled potential semantic interference effects.

The quantum-inspired formalism offers a new perspective for analyzing LLM embeddings. The analogy to quantum mechanics suggests that concepts like superposition and interference may offer insights into semantic representation and processing. The use of quantum circuits provides experimental tools to explore these concepts.

Several limitations must be acknowledged. The current implementation relies on amplitude encoding, leading to a qubit requirement that scales linearly with embedding dimensionality. This presents a challenge for scaling to larger LLMs. Experiments on real quantum computers are susceptible to noise and limited qubit coherence times, affecting accuracy. The relatively small LLM and truncated embedding dimension also limit the generalizability of our findings. Despite these limitations, this work suggests promising directions for future research. These include developing more efficient quantum algorithms for semantic similarity, such as those based on quantum phase estimation or variational quantum eigensolvers, and investigating quantum error correction techniques.

Beyond semantic similarity, this approach could potentially be extended to other NLP tasks, such as text generation and machine translation, by modeling probabilistic relationships between words and phrases using quantum circuits. The density matrix representation of semantic information could also be used as a feature representation for downstream machine learning models.

This work highlights the potential for quantum computing to contribute to a deeper understanding of natural language and intelligence. By demonstrating a connection between semantic spaces and quantum mechanics, we hope to inspire further research into quantum-inspired models for cognitive science and artificial intelligence. Assuming that quantum mechanics provides a valid framework for LLM semantic spaces, and if quantum computing can be effectively applied to LLM training and predictions, it could significantly reduce the reliance on large GPU farms, potentially democratizing LLM development and usage by lowering the barrier to entry.

While high accuracy is crucial for many LLM applications, current expectations often exceed the capabilities of these models, leading to disappointment. If LLMs indeed operate according to quantum mechanical principles, as suggested by our findings, then a degree of inherent uncertainty is unavoidable. However, a good understanding of quantum mechanics provides the tools to manage and control this uncertainty, potentially enabling developers to achieve more robust and reliable results from LLMs, even in the face of inherent probabilistic behavior. In conclusion, this paper provides evidence for the potential of applying quantum formalism to the analysis of LLM embeddings. While challenges remain, the combination of theoretical insights and experimental validation suggests that quantum circuits may offer new tools for exploring these systems. Further research is needed to realize the potential of quantum natural language processing.

Reference

1. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
2. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019, June). 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, volume 1 (long and short papers)* (pp. 4171-4186).
3. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33, 1877-1901.
4. Jolliffe, I., T. (2002). Principal Component Analysis. *Springer*.
5. L. Van der Maaten and G. Hinton. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9(86):2579-2605, 2008.
6. G. Salton and M. McGill. Introduction to Modern Information Retrieval, *McGraw-Hill*, 1986.
7. D. Mimno and L. Thompson. The strange geometry of skip-gram with negative sampling. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2298-2307, 2017.
8. Tenney, I., Xia, P., Chen, B., Wang, A., Poliak, A., McCoy, R. T., ... & Pavlick, E. (2019). What do you learn from context? probing for sentence structure in contextualized word representations. *arXiv preprint arXiv:1905.06316*.
9. Nielsen, M. A., & Chuang, I. L. (2010). Quantum computation and quantum information. *Cambridge university press*.
10. DiVincenzo, D. P. (2000). The physical implementation of quantum computation. *Fortschritte der Physik: Progress of Physics*, 48(9-11), 771-783.
11. Shor, P. W. (1994, November). Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th annual symposium on foundations of computer science* (pp. 124-134). Ieee.
12. Grover, L. K. (1997). Quantum mechanics helps in searching for a needle in a haystack. *Physical review letters*, 79(2), 325.
13. Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J. C., Barends, R., ... & Martinis, J. M. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574 (7779), 505-510.
14. Bharti, K., Cervera-Lierta, A., Kyaw, T. H., Haug, T., Alperin-Lea, S., Anand, A., ... & Aspuru-Guzik, A. (2022). Noisy intermediate-scale quantum algorithms. *Reviews of Modern Physics*, 94(1), 015004.
15. Bruza, P., Busemeyer, J., & Gabora, L. (2009). Introduction to the special issue on quantum cognition. *Journal of Mathematical Psychology*, 53, 303-305.
16. Busemeyer, J. R., & Bruza, P. D. (2012). Quantum models of cognition and decision. *Cambridge University Press*.
17. Atmanspacher, H., & Römer, H. (2012). Order effects in sequential measurements of non-commuting psychological observables. *Journal of Mathematical Psychology*, 56(4), 274-280.
18. A.Y. Khrennikov. Ubiquitous quantum structure: From psychology to finance. Springer, 2010.
19. Trouillon, T., Welbl, J., Riedel, S., Gaussier, É., & Bouchard, G. (2016, June). Complex embeddings for simple link prediction. In *International conference on machine learning* (pp. 2071-2080). PMLR.
20. Sun, Z., Deng, Z. H., Nie, J. Y., & Tang, J. (2019). Rotate: Knowledge graph embedding by relational rotation in complex space. *arXiv preprint arXiv:1902.10197*.
21. Zhang, S., Tay, Y., Yao, L., & Liu, Q. (2019). Quaternion knowledge graph embeddings. *Advances in neural information processing systems*, 32.
22. Havlíček, V., Córcoles, A. D., Temme, K., Harrow, A. W., Kandala, A., Chow, J. M., & Gambetta, J. M. (2019). Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747), 209-212.
23. Beer, K., Bondarenko, D., Farrelly, T., Osborne, T. J., Salzmann, R., Scheiermann, D., & Wolf, R. (2020). Training deep quantum neural networks. *Nature communications*, 11(1), 808.

-
24. Arrasmith, A., Cerezo, M., Czarnik, P., Cincio, L., & Coles, P. J. (2021). Effect of barren plateaus on gradient-free optimization. *Quantum*, 5, 558.
 25. Coecke, B., Sadrzadeh, M., & Clark, S. (2010). Mathematical foundations for a compositional distributional model of meaning. *arXiv preprint arXiv:1003.4394*.
 26. Laine, T. A. (2025). Semantic Wave Functions: Exploring Meaning in Large Language Models Through Quantum Formalism. *OA J Applied Sci Technol*, 3(1), 01-22, 2025.
 27. Laine, T. A. (2025). The Quantum LLM: Modeling Semantic Spaces with Quantum Principles. *OA J Applied Sci Technol*, 3(2), 01-13, 2025.
 28. Merzbacher, E. (1998). Quantum mechanics john wiley & sons. Inc., New York, 50.

Copyright: ©2026 Timo Aukusti Laine, et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.