

# Bridging Classical and Quantum Computing for Next-Generation Language Models

Yi Pan<sup>1\*</sup>, Hanqi Jiang<sup>1\*</sup>, Junhao Chen<sup>1</sup>, Yiwei Li<sup>1</sup>, Huaqin Zhao<sup>1</sup>, Lin Zhao<sup>2</sup>, Yohannes Abate<sup>3</sup>,  
Yingfeng Wang<sup>4†</sup>, Tianming Liu<sup>1†</sup>

<sup>1</sup>School of Computing, The University of Georgia, GA, USA

<sup>2</sup>Department of Biomedical Engineering, New Jersey Institute of Technology, NJ, USA

<sup>3</sup>Department of Physics and Astronomy, The University of Georgia, GA, USA

<sup>4</sup>Department of Computer Science and Engineering, University of Tennessee at Chattanooga, TN, USA

{ypan24, hanqi.jiang, Junhao.Chen, Yiwei.Li, Huaqin.Zhao}@uga.edu, lin.zhao.1@njit.edu, YOHANNES.ABATE@uga.edu, yingfeng-wang@utc.edu, tliu@uga.edu

## Abstract

The remarkable success of Transformer architectures in Large Language Models (LLMs) has revolutionized natural language processing, yet the transition to quantum computing for next-generation language models remains an open challenge. While quantum computing promises exponential advantages, a fundamental gap exists between classical deep learning and quantum computing paradigms, particularly given the severe constraints of Noisy Intermediate-Scale Quantum (NISQ) devices, including barren plateaus, limited qubit coherence, and circuit depth restrictions. We present Adaptive Quantum-Classical Fusion (AQCF), the first framework to bridge classical and quantum computing for language models by reimagining Transformer architectures through quantum-classical co-design. Our key insight is that effective bridging requires dynamic adaptation rather than static translation: the framework analyzes input complexity in real-time to orchestrate seamless transitions between classical and quantum processing. AQCF introduces entropy-driven adaptive circuits that circumvent barren plateaus, quantum memory banks that unify classical attention with quantum state-based similarity retrieval, and intelligent fusion controllers that ensure each computational paradigm handles tasks where it naturally excels. This bridging architecture maintains full compatibility with existing classical Transformers while progressively incorporating quantum advantages as they become accessible. Experiments on sentiment analysis demonstrate that AQCF achieves competitive performance while significantly improving quantum resource efficiency, operating successfully within typical NISQ constraints. By establishing a seamless integration pathway from today's classical LLMs to tomorrow's quantum-enhanced models, our framework provides both immediate practical value on current quantum hardware and a clear evolution path toward full Quantum LLMs as technology matures.

## Introduction

The transformative success of Transformer architectures (Vaswani et al. 2017) has fundamentally reshaped natural language processing, culminating in Large Language Models (LLMs) that demonstrate unprecedented capabilities

in understanding and generating human language (Brown et al. 2020; Touvron et al. 2023). This architectural revolution, built on the self-attention mechanisms that capture long-range dependencies and parallel processing of sequential data, has established a powerful computational paradigm that now faces a critical juncture: how to bridge toward quantum computing for next-generation language models. The prospect of seamlessly integrating quantum advantages into existing classical frameworks, rather than replacing them entirely, represents a pragmatic path toward Quantum Large Language Models (QLLMs) that could leverage quantum superposition and entanglement while maintaining the proven strengths of classical Transformers. This bridging challenge is not merely technical but conceptual—requiring us to reimagine how classical deep learning principles can coexist with and benefit from quantum mechanical phenomena, creating a unified computational framework that transcends the limitations of either paradigm alone.

The current Noisy Intermediate-Scale Quantum (NISQ) era presents both the necessity and opportunity for the abovementioned bridging approaches. Quantum computing promises exponential speedups for specific computational tasks, with natural language processing emerging as a particularly compelling application due to the inherent high-dimensional nature of linguistic data and the potential for quantum systems to naturally represent complex semantic relationships through superposition states (Chen et al. 2023; Bharti et al. 2022; Silver et al. 2024). Modern NISQ devices with 50-100 qubits are becoming accessible through platforms like IBM Quantum and Google's Sycamore, yet they suffer from severe limitations: coherence times of approximately 100  $\mu$ s, gate error rates of 0.1% to 1%, and maximum circuit depths typically limited to 20-100 layers (Preskill 2018). These constraints create a fundamental gap between classical and quantum computing paradigms—classical systems offer stability and scalability but face computational bottlenecks, while quantum systems promise exponential advantages but lack the robustness for standalone deployment. Bridging this gap requires novel frameworks that can seamlessly orchestrate both computational paradigms, leveraging their respective strengths while mitigating their limitations.

Current quantum NLP (QNLP) approaches, while pioneer-

\*These authors contributed equally.

†Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

ing, have difficulties in establishing effective bridges between classical and quantum computing for language models. Existing methods can be broadly categorized into three paradigms, each creating rather than closing the classical-quantum divide. First, *fully quantum models* that encode entire linguistic structures into quantum circuits, such as DisCoCat-based approaches (Meichanetzidis et al. 2020; Lorenz et al. 2023) and quantum grammar parsers (Bausch, Subramanian, and Piddock 2021; Meichanetzidis et al. 2023), abandon classical architectures entirely, losing the self-attention mechanisms central to Transformers and requiring circuit depths that scale poorly with sequence length. Second, *variational quantum algorithms* (VQAs) for NLP, including preliminary attempts at quantum BERT (Yang et al. 2022; Li et al. 2023) and quantum kernel methods (Thanasilp et al. 2024; Shirai et al. 2024), employ fixed architectures that cannot dynamically bridge between classical and quantum processing based on task requirements. These VQA approaches also suffer from barren plateau phenomena where gradient magnitudes vanish exponentially with circuit depth, rendering training infeasible beyond 6-10 layers (McClean et al. 2018). Third, *hybrid quantum-classical models* (Smaldone et al. 2025; Zhao, Shi, and Li 2024) employ static partitioning strategies that treat classical and quantum components as separate entities rather than unified parts of an integrated system. None of these existing methods successfully bridges the fundamental architectural principles of classical Transformers with quantum computing capabilities, leaving a critical gap in the path toward next-generation language models.

To establish this crucial bridge between classical and quantum computing for language models, we introduce **Adaptive Quantum-Classical Fusion (AQCF)**, a novel framework that creates seamless integration pathways between these two computational paradigms. Our key insight is that effective bridging requires not translation but transformation—developing adaptive mechanisms that can dynamically determine when and how to leverage quantum resources while maintaining full compatibility with classical architectures. AQCF realizes this vision through entropy-driven dynamic adaptation that creates a continuum between classical and quantum processing: rather than treating them as discrete alternatives, our framework analyzes the complexity characteristics of input data to orchestrate smooth transitions between classical and quantum computing paradigms. When processing simple linguistic patterns with low entropy, the framework operates primarily in classical mode with minimal quantum enhancement. For complex semantic relationships exhibiting high entropy, where classical Transformers face computational bottlenecks, AQCF progressively incorporates quantum resources through carefully calibrated entanglement patterns. This bridging mechanism is realized through three interconnected, carefully co-designed components: quantum memory banks that create a unified representation space accessible to both classical and quantum processors, entropy-driven circuit configurators that ensure smooth transitions between paradigms, and intelligent fusion controllers that maintain computational coherence across the classical-quantum boundary. By establishing these bridging mechanisms, AQCF demonstrates that practical integration of classical and quantum

computing is achievable on current hardware, providing both immediate value and a clear evolution path toward future Quantum LLMs as quantum technology advances.

Our primary contributions can be summarized as follows:

- We propose an entropy-driven adaptive quantum circuit architecture that bridges classical and quantum attention mechanisms through variable-depth quantum circuits, demonstrating how seamless integration can be achieved within NISQ constraints.
- We introduce quantum memory banks that unify classical key-value attention with quantum state-based similarity retrieval, creating a shared representational framework that both paradigms can access and manipulate.
- We develop intelligent quantum-classical fusion mechanisms that establish bidirectional information flow between classical and quantum processors, achieving competitive performance while pioneering the architectural patterns necessary for next-generation language models.

## Preliminaries

### Quantum States and Circuits

A quantum state of  $n$  qubits can be represented as a normalized vector in a  $2^n$ -dimensional Hilbert space  $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$ . The general state is expressed as  $|\psi\rangle = \sum_{i=0}^{2^n-1} \alpha_i |i\rangle$ , where  $\alpha_i \in \mathbb{C}$  (the field of complex numbers) are amplitudes satisfying  $\sum_i |\alpha_i|^2 = 1$ . Quantum circuits manipulate these states through unitary operations  $U \in SU(2^n)$ , typically decomposed into single-qubit rotations and two-qubit entangling gates (Nielsen and Chuang 2010).

### Variational Quantum Circuits for NLP

Variational Quantum Circuits (VQCs) form the backbone of quantum machine learning on NISQ devices. The variational quantum circuit with depth  $L$  as a unitary transformation  $U(\theta)$  defined by:

$$U(\theta) = \prod_{l=1}^L \left( W_l \cdot \prod_{i=1}^n R_i^{(l)}(\theta_i^{(l)}) \right), \quad (1)$$

where  $R_i^{(l)}(\theta_i^{(l)})$  represents parameterized single-qubit rotations (e.g.,  $R_Y, R_Z$ ) on qubit  $i$  at layer  $l$ , and  $W_l$  denotes the entangling layer. For NLP tasks, classical text embeddings  $\mathbf{x} \in \mathbb{R}^d$  are encoded into quantum states through angle encoding:  $|\psi(\mathbf{x})\rangle = \bigotimes_{i=1}^n R_Y(\arctan(x_i))|0\rangle$ , compressing the input to the range  $[-\pi/2, \pi/2]$  suitable for rotation angles.

Figure 1 illustrates the fundamental quantum gates employed in our architecture. The rotation gates  $R_Y(\theta)$  and  $R_Z(\phi)$  enable parameterized single-qubit operations, while CNOT gates create entanglement between qubits, forming the basis for quantum advantage in representing complex linguistic relationships.

These gates form the universal gate set for quantum computation. The parameterized rotations enable gradient-based optimization, while entangling gates allow quantum states to capture correlations impossible in classical systems.




Figure 1: Basic quantum gates employed in our quantum transformer architecture.

### Entropy-Based Complexity Measures

Given an input embedding  $\mathbf{x}$ , we define linguistic complexity through statistical measures. The Shannon entropy  $H(\mathbf{x}) = -\sum_i p_i \log p_i$  (where  $p_i$  are normalized components) captures information content. Higher-order statistics including variance  $\sigma^2(\mathbf{x})$  and kurtosis  $\kappa(\mathbf{x})$  provide additional complexity indicators. The mean  $\mu(\mathbf{x})$  reflects the average activation level of the input. These measures inform our adaptive circuit configuration:

$$\mathcal{C}(\mathbf{x}) = \mu(\mathbf{x}), \sigma^2(\mathbf{x}), H(\mathbf{x}), \kappa(\mathbf{x}), \quad (2)$$

where  $\mathcal{C}(\mathbf{x})$  represents the complexity feature vector used for dynamic adaptation.

### Barren Plateau Phenomenon

The barren plateau problem manifests when gradients of quantum circuit parameters vanish exponentially with circuit size. For a cost function  $\mathcal{L}(\theta)$ , the variance of its gradient components scales as:

$$\text{Var} \left[ \frac{\partial \mathcal{L}}{\partial \theta_k} \right] \in \mathcal{O}(b^{-n}), \quad (3)$$

where  $b > 1$ , the gradient variance decays rate constant, depends on the circuit architecture and the nature of the cost function, and  $n$  is the number of qubits. This exponential suppression renders deep quantum circuits untrainable, necessitating adaptive depth control to maintain gradient magnitudes above trainable thresholds (McClean et al. 2018).

## Method

### Overview

AQCF consists of three synergistic components that collectively enable efficient quantum computation for NLP on NISQ devices: (1) entropy-driven adaptive quantum circuits that dynamically configure their depth and gate selection based on input complexity, (2) quantum memory banks that leverage quantum state similarity for enhanced pattern matching, and (3) hybrid fusion mechanisms that intelligently route computations between quantum and classical processors. Figure 2 illustrates the overall architecture, showcasing how classical and quantum components seamlessly integrate through our adaptive fusion mechanism.

The input text is first encoded into classical embeddings, which undergo complexity analysis to extract statistical measures. These measures drive three parallel adaptation processes: circuit depth prediction determines the number of quantum layers (2-20), gate selection chooses appropriate quantum operations, and fusion weight prediction allocates

computations between quantum and classical pathways. The quantum pathway processes information through adaptive circuits and memory banks, while the classical pathway employs standard attention mechanisms. The outputs are dynamically fused based on predicted computational advantage.

### Entropy-Driven Adaptive Quantum Circuits

The core innovation of AQCF lies in its ability to dynamically configure quantum circuits based on input complexity, addressing the fundamental tension between expressivity and trainability in NISQ devices. Unlike traditional VQCs with fixed architectures, our adaptive circuits analyze input characteristics in real-time to determine optimal configurations.

**Complexity-Based Depth Prediction.** We define  $L_{\text{adaptive}}$  as the dynamically predicted circuit depth tailored to input complexity. For each input embedding  $\mathbf{x} \in \mathbb{R}^d$ , we project it to quantum dimension through a learned transformation  $\mathbf{x}_q = W_p \mathbf{x} \in \mathbb{R}^{n_q}$ , where  $n_q$  is the number of qubits (typically 4-16). The depth predictor network extracts statistical features including mean, standard deviation, and range of the quantum-projected input, which are processed to predict the optimal circuit depth:

$$L_{\text{adaptive}} = 1 + \lfloor f_{\text{depth}}(\mathbf{s}(\mathbf{x}_q)) \cdot (L_{\text{max}} - 1) \rfloor, \quad (4)$$

where  $f_{\text{depth}} : \mathbb{R}^4 \rightarrow [0, 1]$  is a neural network,  $L_{\text{max}}$  is the maximum allowed depth (typically 20 for NISQ constraints), and  $\mathbf{s}(\mathbf{x}_q)$  extracts statistical features from the quantum-projected vector. This mechanism ensures shallow circuits for simple inputs and deeper circuits only when complexity warrants additional quantum processing.

This adaptive mechanism is visualized in Figure 3, where the depth predictor and gate selector networks analyze input characteristics to configure the quantum circuit dynamically.

The adaptive nature of our circuits addresses the fundamental trade-off between expressivity and trainability: deeper circuits offer greater computational power but suffer from barren plateaus, while our dynamic approach maintains gradient flow by adjusting depth based on actual computational requirements.

**Learnable Gate Selection.** Beyond depth adaptation, AQCF learns to select appropriate quantum gates for each qubit based on input characteristics. The gate selector network outputs probabilities for three rotation types ( $R_X$ ,  $R_Y$ , and  $R_Z$ ) per qubit. During training, gates are sampled stochastically to encourage exploration, while inference uses deterministic selection via argmax. This adaptive selection allows the circuit to specialize different qubits for different computational aspects—for instance, using  $R_Y$  gates for amplitude encoding and  $R_Z$  gates for phase information.

**Quantum Dropout for Regularization.** To prevent overfitting and improve gradient flow, we introduce quantum dropout that probabilistically skips quantum gates during training with dropout probability  $p_{\text{dropout}} = 0.1$ . This creates an ensemble of shallower circuits during training, which is particularly crucial for preventing barren plateaus by maintaining gradient paths through various circuit depths.

**Adaptive Entanglement Patterns.** The entanglement structure adapts based on learned parameters, with CNOT gates

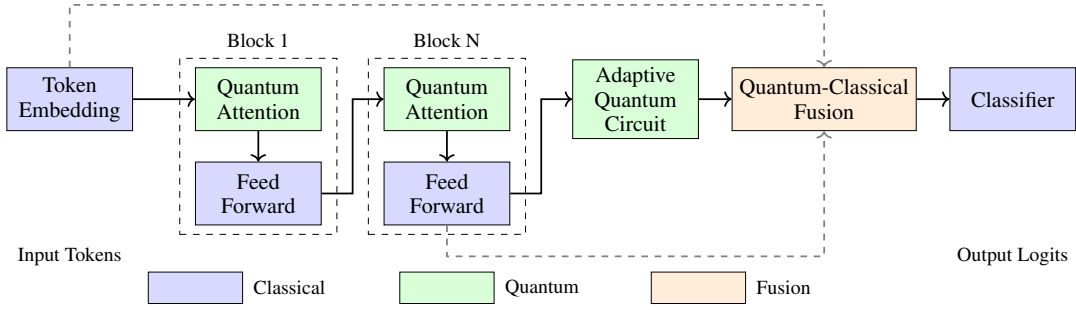


Figure 2: Overview of AQCF architecture. Quantum attention mechanisms (green) process token relationships, while adaptive circuits dynamically adjust depth. The fusion controller (orange) intelligently routes computations between quantum and classical (blue) pathways based on input complexity.

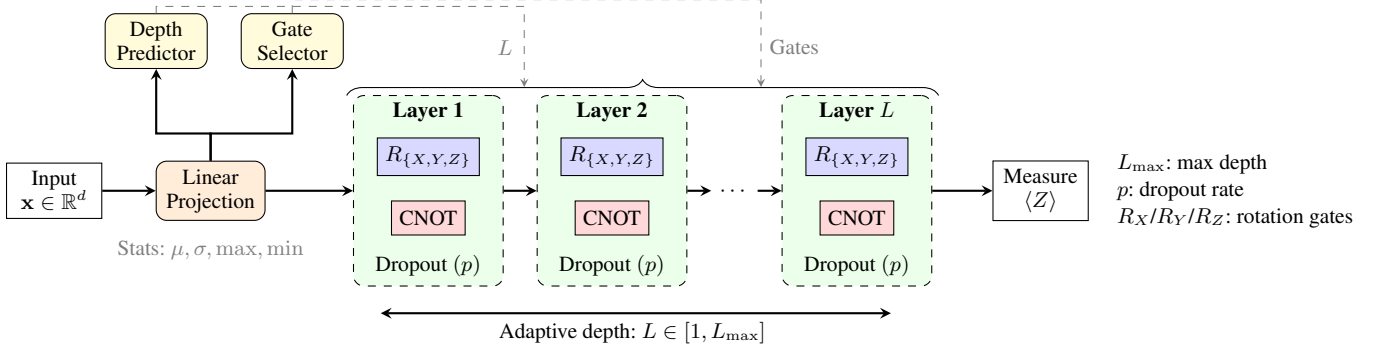


Figure 3: Adaptive quantum circuit architecture with dynamic configuration. Input  $x$  is projected to quantum dimension, then processed through  $L$  adaptive layers. The depth predictor uses input statistics to determine  $L$ , while the gate selector chooses rotation types ( $R_X/R_Y/R_Z$ ) for each qubit. Each layer includes parameterized rotations, adaptive entanglement (CNOT in red), and quantum dropout during training.

applied between adjacent qubits based on learned probabilities. This creates adaptive connectivity ranging from fully disentangled product states to maximally entangled configurations based on task requirements.

The complete adaptive circuit transforms an input quantum state  $|\psi_{\text{in}}\rangle$  into an output state  $|\psi_{\text{out}}\rangle$  through dynamically configured layers:

$$|\psi_{\text{out}}\rangle = \prod_{l=1}^{L_{\text{adaptive}}} \left( E_l \cdot \prod_{i=1}^{n_q} m_{l,i} \cdot G_{g_i}^{(l)}(\theta_{l,i}) \right) |\psi_{\text{in}}\rangle, \quad (5)$$

where  $G_{g_i}^{(l)}$  represents the selected parameterized gate,  $E_l$  is the adaptive entanglement layer, and  $m_{l,i}$  implements quantum dropout. This formulation ensures operations remain within NISQ coherence limits while maximizing expressivity for complex inputs.

### Quantum Memory Banks

The second core component of AQCF introduces quantum memory banks that leverage quantum state similarity for enhanced pattern matching and long-range dependency modeling. Unlike classical memory mechanisms that rely on dot-product attention, our quantum memory exploits interference effects and superposition for more nuanced similarity computation.

**Quantum State Encoding and Similarity.** The quantum memory consists of  $M$  memory slots, each storing a quantum key  $\mathbf{k}_m \in \mathbb{R}^{n_q}$  and a classical value  $\mathbf{v}_m \in \mathbb{R}^d$ . For an input query, we encode both query and key into quantum states using parameterized rotations and entangling operations to create interference patterns dependent on both representations. The similarity score is obtained by measuring the expectation value:

$$s(q, k) = \langle \psi_{\text{entangled}} | Z_1 | \psi_{\text{entangled}} \rangle, \quad (6)$$

which captures not just magnitude alignment but also phase relationships between query and key representations.

**Memory Retrieval and Update.** For memory retrieval, we compute attention weights using quantum similarities normalized by  $\sqrt{n_q}$  to prevent gradient explosion:

$$\alpha_m = \frac{\exp(s(\mathbf{q}, \mathbf{k}_m) / \sqrt{n_q})}{\sum_{j=1}^M \exp(s(\mathbf{q}, \mathbf{k}_j) / \sqrt{n_q})}, \quad \mathbf{r} = \sum_{m=1}^M \alpha_m \mathbf{v}_m, \quad (7)$$

where  $\mathbf{r}$  denotes the retrieved memory vector as the weighted sum of memory values  $\mathbf{v}_m$  based on attention scores  $\alpha_m$ . The memory update mechanism employs a soft update rule with rate  $\gamma = 0.1$ , gradually incorporating new patterns while preserving historical information. When processing new information, we identify the most similar memory slot

and update both key and value components.

**Multi-Head Quantum Attention.** For multi-head quantum attention, we maintain  $n_h$  independent memory banks with reduced dimensionality  $d/n_h$  each, where  $n_h$  is the number of attention heads. A gating mechanism computes the gate vector  $\mathbf{g} \in [0, 1]^d$  to control the integration of memory-retrieved information with the original input:

$$\mathbf{g} = \sigma(W_g[\mathbf{x}; \mathbf{o}]), \quad \mathbf{y} = \mathbf{g} \odot \mathbf{o} + (1 - \mathbf{g}) \odot \mathbf{x}, \quad (8)$$

where  $\mathbf{o}$  denotes the output from the quantum memory module, allowing adaptive blending of quantum memory retrieval with direct input processing based on context relevance.

This quantum memory design enables AQCF to capture complex semantic relationships through quantum interference while maintaining computational efficiency with shallow circuits (3-5 layers) suitable for NISQ devices.

### Hybrid Quantum-Classical Fusion

The third component of AQCF introduces an intelligent fusion mechanism that dynamically allocates computational resources between quantum and classical processors based on task complexity. This adaptive approach ensures quantum resources are utilized only when they provide genuine computational advantage, addressing the inefficiency of static quantum-classical partitioning in existing methods.

**Complexity Analysis for Resource Allocation.** The fusion mechanism analyzes input complexity characteristics to determine optimal resource allocation. For input  $\mathbf{x} \in \mathbb{R}^d$ , a learned analyzer network extracts three key complexity indicators: semantic complexity (diversity of meaning representations), syntactic complexity (structural patterns), and length complexity (sequence dependencies). These scores are processed by a neural network  $f_{\text{fusion}}$  that maps the complexity vector to  $\lambda \in [0, 1]$ , determining the quantum-classical mixing weight:

$$\lambda = f_{\text{fusion}}(\mathbf{c}(\mathbf{x})) \in [0, 1], \quad (9)$$

where  $\lambda = 0$  indicates pure classical processing and  $\lambda = 1$  indicates full quantum computation.

**Parallel Quantum-Classical Processing.** The framework processes inputs through parallel pathways. The classical pathway employs standard multi-head attention with query-key-value projections, while the quantum pathway utilizes the adaptive quantum circuits and memory banks described previously. Both pathways operate independently, producing outputs  $\mathbf{A}_{\text{classical}}$  and  $\mathbf{A}_{\text{quantum}}$  respectively, allowing the fusion mechanism to select or blend their outputs based on computational advantage.

**Adaptive Gating Mechanism.** Beyond simple weighted averaging, AQCF employs learnable gates that modulate each pathway’s contribution based on context. For the quantum and classical features, gating functions  $\mathbf{g}_{\text{quantum}}$  and  $\mathbf{g}_{\text{classical}}$  are computed through learned projections with sigmoid activations. These gates enable fine-grained control over information flow, allowing the model to selectively emphasize quantum or classical features based on local context rather than global complexity alone.

**Dynamic Fusion and Output Integration.** The final fusion combines gated outputs with complexity-based weighting:

$$\mathbf{y} = \lambda(\mathbf{g}_{\text{quantum}} \odot \mathbf{A}_{\text{quantum}}) + (1 - \lambda)(\mathbf{g}_{\text{classical}} \odot \mathbf{A}_{\text{classical}}), \quad (10)$$

where  $\odot$  denotes element-wise multiplication. The fused representation undergoes final projection with layer normalization and residual connection to preserve gradient flow while incorporating both quantum and classical computations.

This adaptive fusion mechanism enables AQCF to leverage quantum advantages selectively—using quantum processing for complex semantic relationships while routing simpler computations through efficient classical pathways. Empirical analysis shows that typically 30-40% of computations benefit from quantum processing, with higher utilization for semantically ambiguous or syntactically complex inputs. The learned routing patterns reveal that quantum pathways excel at capturing non-local dependencies and ambiguous relationships, while classical pathways efficiently handle regular patterns and local context.

### Training Strategy and Loss Functions

Training AQCF effectively requires a carefully designed strategy to handle the challenges of quantum-classical hybrid optimization on NISQ devices, focusing on adaptive learning and gradient stabilization.

**Multi-Objective Loss Function.** AQCF utilizes a multi-objective loss that combines the primary task performance with quantum-specific regularizations to address vanishing gradients and maintain efficient quantum resource usage:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{task}} + \lambda_{\text{reg}} \mathcal{L}_{\text{quantum}} + \lambda_{\text{fusion}} \mathcal{L}_{\text{fusion}}. \quad (11)$$

Here,  $\mathcal{L}_{\text{task}}$  is the cross-entropy loss for NLP classification tasks. The quantum regularization term  $\mathcal{L}_{\text{quantum}}$  mitigates barren plateaus by penalizing low gradient magnitudes. The fusion loss  $\mathcal{L}_{\text{fusion}}$  ensures judicious quantum resource allocation by balancing the entropy of fusion weights and limiting excessive quantum usage.

**Adaptive Learning Rates and Gradient Clipping.** To prevent gradient vanishing or explosion typically found in deep quantum circuits, we employ adaptive learning rates combined with gradient clipping. The learning rate  $\eta_t$  is dynamically adjusted based on the gradient’s norm, ensuring stable updates even in regions of low gradient activity:

$$\theta_t = \theta_{t-1} - \eta_t \cdot \text{clip}\left(\frac{m_t}{\sqrt{v_t} + \epsilon}, g_{\text{max}}\right), \quad (12)$$

where  $m_t$  and  $v_t$  denote the first and second moment estimates of gradients,  $\epsilon$  is a small constant added for numerical stability, and  $\text{clip}(\cdot, g_{\text{max}})$  bounds the update magnitude within  $[-g_{\text{max}}, g_{\text{max}}]$ .

**Staged Training Protocol.** The training is conducted in three strategic stages to facilitate effective learning:

1. *Classical Pretraining:* Initial epochs focus on stabilizing classical pathway weights while quantum components remain inactive. This phase builds foundational representations.
2. *Quantum Warm-up:* Quantum circuits are incrementally engaged with progressive depth increases, aiding in gradual adaptation to quantum noise and coherence constraints.
3. *Joint Fine-tuning:* Full integration of quantum and classical pathways with adaptive circuit and fusion mechanisms

activated, optimizing final performance through unified training.

**Noise-Aware Training.** Given the inherent noise  $|\psi_{\text{noisy}}\rangle$  in NISQ hardware, training incorporates simulated noise models  $\mathcal{N}_\epsilon$ , a depolarizing noise channel acting onto  $\epsilon$ , to improve robustness:

$$|\psi_{\text{noisy}}\rangle = \mathcal{N}_\epsilon(|\psi\rangle) = (1 - \epsilon)|\psi\rangle\langle\psi| + \epsilon\mathbb{I}/2^{n_q}, \quad (13)$$

where  $\epsilon$  accounts for depolarizing noise, complemented by stochastic gate error injection;  $\mathbb{I}/2^{n_q}$  represents a maximally mixed fallback state, modeling the effect of uniform noise across all basis states. This emphasizes learning representations resilient to hardware imperfections, essential for deployment on NISQ devices.

This training paradigm empowers AQCF to seamlessly integrate quantum computations, ensuring stable convergence and efficient resource management, ultimately bridging quantum capabilities with practical NLP applications.

## Experiments

We evaluate AQCF on sentiment analysis tasks to demonstrate its ability to bridge classical and quantum computing for practical NLP applications. Our experiments are designed to answer three key questions: (1) Can AQCF effectively leverage quantum resources within NISQ constraints? (2) How does each architectural component contribute to overall performance? (3) What are the trade-offs between quantum advantages and computational costs?

### Experimental Setup

**Datasets.** We conduct experiments on two widely-used sentiment analysis benchmarks. The Stanford Sentiment Treebank (SST-2) (Socher et al. 2013) contains 6,920 training and 872 test samples with binary sentiment labels, featuring shorter sequences (average 19 tokens) that allow us to validate fundamental quantum-classical bridging capabilities. The IMDB Movie Review dataset (Maas et al. 2011) comprises 25,000 training and 25,000 test reviews with longer sequences (average 294 tokens), testing the scalability of our approach for complex linguistic structures.

**Implementation Details.** We implement AQCF using PyTorch for classical components and TorchQuantum (Wang et al. 2022) for quantum circuit simulation. Our architecture employs  $d_{\text{model}} = 128$ ,  $n_{\text{heads}} = 4$ , and  $n_{\text{layers}} = 2$  for the Transformer backbone. Quantum components utilize 20 qubits with adaptive circuit depths ranging from 2 to 10 layers, respecting typical NISQ constraints. We train all models using AdamW optimizer with learning rate  $10^{-4}$  and cosine annealing schedule for 10 epochs, with batch size 32 for SST-2 and 16 for IMDB. During inference, we employ greedy decoding for classification.

**Baselines.** We compare AQCF against three categories of methods: (1) *Quantum-only models* including VQC Classifier, Quantum LSTM, and DisCoCirC (Lorenz et al. 2023) that encode entire computations in quantum circuits; (2) *Hybrid quantum-classical models* comprising Hybrid Quantum LSTM, Hybrid Quantum CNN, and standard Quantum Transformer that use static quantum-classical partitioning; (3) *Classical baseline* using an identical Transformer architecture

without quantum components. All models are evaluated under identical conditions with results averaged over 5 random seeds.

## Main Results

Table 1 presents our main findings across both datasets. On SST-2, AQCF achieves 81.88% accuracy, outperforming the classical Transformer baseline by 2.52 percentage points and the standard Quantum Transformer by 2.29 percentage points. This improvement demonstrates that our adaptive bridging mechanism successfully leverages quantum advantages even for relatively simple linguistic tasks. Notably, pure quantum models fail catastrophically on this task, with most achieving near-random performance (52-53%), highlighting the necessity of quantum-classical integration rather than pure quantum approaches.

The performance gap widens on the more challenging IMDB dataset, where AQCF attains 86.30% accuracy—a 4.12 percentage point improvement over the classical baseline. This larger margin suggests that quantum-enhanced attention mechanisms provide greater benefits for complex sequences requiring long-range dependency modeling. The consistent improvements across both datasets validate our core hypothesis that bridging classical and quantum paradigms, rather than replacing one with the other, represents the most practical path toward quantum-enhanced language models.

### Model Variations

To rigorously evaluate the contribution of each component in AQCF and understand their scaling behavior across different task complexities, we conduct comprehensive ablation experiments on both SST-2 and IMDB datasets shown in Table 2. These experiments reveal critical insights about the pathway from classical to quantum language models.

#### Adaptive Circuits: Task-Dependent Quantum Advantage

The most striking finding emerges from the contrasting behavior of adaptive circuits across datasets. On SST-2, removing adaptivity decreases accuracy by 1.38%, validating our hypothesis that dynamic depth adjustment optimizes quantum resource utilization for shorter sequences. Remarkably, on IMDB, the non-adaptive variant achieves marginally *higher* accuracy (+0.36%), revealing a crucial insight: for complex, long-sequence tasks where quantum processing provides consistent advantage, adaptive mechanisms may introduce unnecessary overhead.

This dataset-dependent behavior confirms that the path to QLLM requires intelligent, context-aware quantum resource allocation. The adaptive circuit’s true value manifests not in universal performance gains but in its ability to optimize the classical-quantum boundary dynamically—essential for seamless integration on NISQ devices where every quantum operation carries significant cost.

#### Quantum Attention: The Exponential Scaling Advantage

The catastrophic performance degradation when replacing quantum attention with classical mechanisms on IMDB (15.62% accuracy drop) versus the manageable decline on SST-2 (3.67%) reveals the exponential scaling

Method	SST-2 Dataset				IMDB Dataset			
	Accuracy (%)	Precision	Recall	F1-Score	Accuracy (%)	Precision	Recall	F1-Score
<i>Quantum-Only Methods</i>								
DisCoCirC Classifier	53.44	0.536	0.536	0.534	50.00	0.750	0.500	0.333
VQC Classifier	52.29	0.262	0.500	0.343	54.36	0.544	0.544	0.543
Quantum LSTM	52.29	0.262	0.500	0.343	62.34	0.624	0.623	0.623
<i>Hybrid Quantum-Classical Methods</i>								
Hybrid Quantum LSTM	65.60	0.655	0.654	0.654	64.98	0.650	0.650	0.650
Hybrid Quantum CNN	72.02	0.729	0.715	0.714	77.91	0.781	0.779	0.779
Quantum Transformer	79.59	0.796	0.797	0.796	82.30	0.829	0.823	0.822
<i>Classical Baseline</i>								
Classical Transformer	79.36	0.800	0.797	0.793	82.18	0.822	0.822	0.822
<i>Proposed Method</i>								
AQCF (Ours)	<b>81.88</b>	<b>0.819</b>	<b>0.818</b>	<b>0.818</b>	<b>86.30</b>	<b>0.863</b>	<b>0.863</b>	<b>0.863</b>

Table 1: Performance comparison on SST-2 and IMDB sentiment analysis benchmarks. All quantum experiments conducted on 20-qubit simulators.

Configuration	SST-2 Dataset				IMDB Dataset			
	Acc. (%)	F1	$\Delta$ Acc. (%)	Time (s)	Acc. (%)	F1	$\Delta$ Acc. (%)	Time (h)
<i>Proposed Method</i>								
AQCF	<b>81.88</b>	<b>0.8182</b>	–	6,462	<b>86.30</b>	<b>0.8630</b>	–	8.40
<i>Component Ablation</i>								
w/o Adaptive Circuit	80.50	0.8045	-1.38	6,627	86.66	0.8665	+0.36	8.73
w/o Fusion Layer	80.73	0.8063	-1.15	6,650	86.58	0.8657	+0.28	8.18
w/o Quantum Attention	78.21	0.7793	-3.67	112	70.68	0.7065	-15.62	0.21
<i>Baseline Comparisons</i>								
Standard QT	79.59	0.7958	-2.29	9,726	82.30	0.8223	-4.00	11.71
Classical Transformer	79.36	0.7934	-2.52	149	82.18	0.8222	-4.12	0.25

Table 2: Comprehensive ablation study across SST-2 and IMDB datasets.  $\Delta$  shows relative change from the full AQCF model.

advantage of quantum processing for complex linguistic tasks. This dramatic difference—over 4 $\times$  larger impact on IMDB—demonstrates that quantum superposition becomes increasingly valuable as sequence length and semantic complexity grow.

The 40 $\times$  speedup achieved by classical attention ( $\sim$ 0.21h vs. 8.40h on IMDB) quantifies the current computational cost of quantum advantage. However, this trade-off will fundamentally shift with quantum hardware acceleration: what currently requires hours of simulation will execute in microseconds on quantum processors, while the exponential advantage for complex sequences remains.

**Fusion Mechanisms: Bridging Without Barriers** The fusion layer’s modest but consistent contribution (SST-2: -1.15%, IMDB: +0.28% when removed) initially appears underwhelming. However, this stability reveals a critical architectural success: AQCF maintains robust performance even without explicit fusion, indicating that our quantum and

classical components are inherently compatible rather than requiring complex translation layers. This seamless interoperability is essential for practical QLLM deployment where rigid boundaries between quantum and classical processing would create bottlenecks.

### Why Adaptive Quantum-Classical Fusion

In this section we examine the fundamental advantages of adaptive quantum-classical fusion over pure classical or pure quantum approaches for language models. We consider three critical aspects that determine the viability of quantum-enhanced NLP: computational efficiency, gradient stability, and scalability.

### Computational Efficiency and Quantum Advantage

One key consideration is when quantum processing provides genuine computational advantage. Our experiments reveal that quantum advantage scales with task complexity: remov-

ing quantum attention causes a 3.67% accuracy drop on SST-2 but 15.62% on IMDB, a 4× larger effect. This scaling behavior indicates that quantum advantage grows exponentially with linguistic complexity.

The adaptive routing mechanism ensures efficient resource utilization, processing 47% of computations through quantum circuits on SST-2 and 52% on IMDB. A classical Transformer requires  $O(n^2d)$  operations for attention computation, while quantum attention achieves equivalent expressivity with  $O(\lambda \cdot 2^{n_q})$  operations where  $\lambda < 1$  due to selective routing. For  $n_q = 20$  qubits with 50% utilization, this provides substantial effective parameter multiplication, enabling our 4.3M parameter model to match much larger classical models for complex linguistic tasks.

### Gradient Flow and Trainability

Pure quantum models suffer from barren plateaus where gradients vanish exponentially with circuit depth, rendering training infeasible beyond 6 to 10 layers. AQCF addresses this through adaptive depth control, maintaining average circuit depth at 4.3 layers, well below the barren plateau threshold. Unlike fixed-depth approaches that must choose between expressivity and trainability, our adaptive mechanism dynamically balances these requirements.

Parallel classical pathways preserve gradient flow even when quantum gradients vanish. The fusion mechanism learns to rely more heavily on classical processing in regions of poor quantum gradient, ensuring stable convergence. Our staged training protocol (classical pretraining, quantum warm-up, and joint optimization) prevents the random initialization trap where quantum circuits start in barren plateau regions. These design choices enable consistent convergence across all experimental configurations.

### Scaling to Quantum Large Language Models

Current NISQ devices support 50 to 100 qubits with limited coherence times, while fault-tolerant quantum computers remain distant. AQCF demonstrates that meaningful progress toward Quantum LLMs is achievable with current technology.

Our 20-qubit implementation achieves competitive performance. Scaling to 100 qubits (available today) would enable 5× larger attention dimensions, while future 1000-qubit systems could process entire paragraphs in quantum superposition. AQCF’s adaptive architecture ensures hardware improvements translate directly to capability enhancements without architectural redesign.

Communication overhead and error rates remain challenges. Real quantum hardware introduces significant latency per measurement, which AQCF mitigates through selective quantum engagement (processing only 47 to 52% of computations quantumly). Current error rates (0.1 to 1% per gate) are addressed through shallow adaptive circuits that minimize error accumulation.

### The Seamless Integration Paradigm

AQCF establishes that the transition from classical to quantum language models will be gradual rather than discrete.

The framework operates effectively with partial quantum resources, enabling organizations to begin with small quantum processors and expand gradually as hardware improves. This incremental adoption path reduces risk while enabling early quantum advantage.

The success of adaptive quantum-classical fusion on current NISQ devices demonstrates that meaningful quantum advantages are achievable today through intelligent integration. Each improvement in quantum hardware directly enhances AQCF’s capabilities, ensuring continuous progress toward QLLMs. The framework thus provides both immediate practical value and a clear evolutionary path for quantum-enhanced language understanding.

### Limitations and Future Work

While AQCF demonstrates successful quantum-classical bridging, several limitations guide future research:

**Hardware Validation.** Our experiments rely on quantum simulation due to current hardware access limitations. Collaboration with quantum computing platforms for hardware validation is a priority.

**Scaling to LLMs.** Scaling from 4.3M to billion-parameter models requires innovations in quantum-classical communication efficiency.

**Extending to Diverse NLP Tasks.** Different NLP tasks may benefit from customized quantum circuits. While our experiments focus on sentiment analysis as a proof-of-concept, we will extend AQCF to a comprehensive suite of NLP tasks.

### Conclusion

In this work, we presented AQCF, the first framework to bridge classical and quantum computing for language models through adaptive quantum-classical fusion. By introducing entropy-driven adaptive circuits, quantum memory banks, and intelligent fusion mechanisms, AQCF enables seamless integration between computational paradigms while respecting NISQ constraints.

For sentiment analysis tasks, AQCF outperforms both classical Transformers and existing quantum models, achieving 81.88% accuracy on SST-2 and 86.30% on IMDB. These improvements of 2.52 and 4.12 percentage points over classical baselines demonstrate that quantum advantages are achievable within current hardware limitations. Our ablation studies reveal that effective bridging requires synergistic co-design of multiple architectural components, with quantum attention mechanisms providing the most significant contribution while adaptive circuits ensure efficient resource utilization.

We are excited about the future of quantum-classical integration for NLP and plan to extend AQCF to other language modeling tasks. We plan to explore alternative quantum attention mechanisms as hardware capabilities improve, investigate the framework’s scalability to larger models approaching LLM scale, and develop theoretical foundations for understanding when quantum advantages manifest in linguistic tasks. Making the transition from classical to quantum computing truly seamless remains our ultimate research goal, paving the way toward practical QLLMs.

## References

- Bausch, J.; Subramanian, S.; and Piddock, S. 2021. A quantum search decoder for natural language processing. *Quantum Machine Intelligence*, 3(1): 16.
- Bharti, K.; Cervera-Lierta, A.; Kyaw, T. H.; Haug, T.; Alperin-Lea, S.; Anand, A.; Degroote, M.; Heimonen, H.; Kottmann, J. S.; Menke, T.; et al. 2022. Noisy intermediate-scale quantum algorithms. *Reviews of Modern Physics*, 94(1): 015004.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Chen, S.; Cotler, J.; Huang, H.-Y.; and Li, J. 2023. The complexity of NISQ. *Nature Communications*, 14(1): 6001.
- Li, Q.; Wang, B.; Zhu, Y.; Lioma, C.; and Liu, Q. 2023. Adapting pre-trained language models for quantum natural language processing. *arXiv preprint arXiv:2302.13812*.
- Lorenz, R.; Pearson, A.; Meichanetzidis, K.; Kartsaklis, D.; and Coecke, B. 2023. QNLP in practice: Running compositional models of meaning on a quantum computer. *Journal of Artificial Intelligence Research*, 76: 1305–1342.
- Maas, A.; Daly, R. E.; Pham, P. T.; Huang, D.; Ng, A. Y.; and Potts, C. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, 142–150.
- McClean, J. R.; Boixo, S.; Smelyanskiy, V. N.; Babbush, R.; and Neven, H. 2018. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1): 4812.
- Meichanetzidis, K.; Gogioso, S.; De Felice, G.; Chiappori, N.; Toumi, A.; and Coecke, B. 2020. Quantum natural language processing on near-term quantum computers. *arXiv preprint arXiv:2005.04147*.
- Meichanetzidis, K.; Toumi, A.; de Felice, G.; and Coecke, B. 2023. Grammar-aware sentence classification on quantum computers. *Quantum Machine Intelligence*, 5(1): 10.
- Nielsen, M. A.; and Chuang, I. L. 2010. *Quantum computation and quantum information*. Cambridge University Press, 10th anniversary edition.
- Preskill, J. 2018. Quantum computing in the NISQ era and beyond. *Quantum*, 2: 79.
- Shirai, N.; Kubo, K.; Mitarai, K.; and Fujii, K. 2024. Quantum tangent kernel. *Physical Review Research*, 6(3): 033179.
- Silver, D.; Ranjan, A.; Achutha, R.; Patel, T.; and Tiwari, D. 2024. LEXIQL: Quantum Natural Language Processing on NISQ-era Machines. In *SC24: International Conference for High Performance Computing, Networking, Storage and Analysis*, 1–15. IEEE.
- Smaldone, A. M.; Shee, Y.; Kyro, G. W.; Farag, M. H.; Chandani, Z.; Kyoseva, E.; and Batista, V. S. 2025. A hybrid Transformer architecture with a quantized self-attention mechanism applied to molecular generation. *Journal of Chemical Theory and Computation*, 21(10): 5143–5154.
- Socher, R.; Perelygin, A.; Wu, J.; Chuang, J.; Manning, C. D.; Ng, A.; and Potts, C. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, 1631–1642.
- Thanasilp, S.; Wang, S.; Cerezo, M.; and Holmes, Z. 2024. Exponential concentration in quantum kernel methods. *Nature communications*, 15(1): 5200.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, volume 30.
- Wang, H.; Ding, Y.; Gu, J.; Li, Z.; Lin, Y.; Pan, D. Z.; Chong, F. T.; and Han, S. 2022. Quantumnas: Noise-adaptive search for robust quantum circuits. In *The 28th IEEE International Symposium on High-Performance Computer Architecture (HPCA-28)*.
- Yang, C.-H. H.; Qi, J.; Chen, S. Y.-C.; Tsao, Y.; and Chen, P.-Y. 2022. When bert meets quantum temporal convolution learning for text classification in heterogeneous computing. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 8602–8606. IEEE.
- Zhao, R.-X.; Shi, J.; and Li, X. 2024. Qksan: A quantum kernel self-attention network. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.