Adaptive Quantum Approximate Optimization for Genomic Classification

Sthefanie J. G. Passo[1*], Vishal H. Kothavade[1] and John J. Prevost[1]

[1*]Department of Computer Engineering, University of Texas at San Antonio, San Antonio, Texas, United States.

*Corresponding author(s). E-mail(s): Sthefanie.Passo@utsa.edu;
Contributing authors: Vishal.Kothavade@utsa.edu;
Jeff.Prevost@utsa.edu;

Abstract

Quantum machine learning for genomic sequence classification faces three compounding constraints on current hardware: barren plateaus in variational training, exponential concentration of quantum kernel expressivity in high-dimensional feature spaces, and noise accumulation that limits practical circuit depth to fewer than 50 layers on current NISQ devices. Existing quantum genomics studies address these constraints using fixed-architecture circuits without noise-calibrated validation on standardized benchmarks, leaving the potential of adaptive circuit topology largely unexplored.

Three QAOA variants are developed and systematically compared to address this gap. Classical QAOA establishes a baseline using fixed $\boldsymbol{p}$-layer circuits with standard phase and mixing operators. Alternating QAOA introduces structured CNOT entanglement between alternating qubit pairs to evaluate whether non-local correlations improve genomic pattern recognition. The primary contribution, Adaptive QAOA, employs gradient-driven operator selection from feature-weighted Pauli operator pools, generating variable-depth problem-specific circuits that adapt their topology to dataset structure rather than applying universal ansätze.

All three architectures are evaluated through quantum simulation on two standardized genomic benchmarks: Human Nontata Promoters (**36,131** sequences) and Demo Coding vs. Intergenic Sequences (**100,000** sequences). Experiments are conducted under both noiseless conditions and noise models calibrated to IBM superconducting hardware specifications, across feature space dimensionalities ranging from 2 to 10 dimensions using 10-fold cross-validation.

Adaptive QAOA consistently achieves the highest precision (approximately **75%** under noiseless simulation) and maintains perfect specificity across both datasets and noise conditions, results not observed for fixed-depth variants. These performance gains are accompanied by reduced hardware resource consumption, with Adaptive QAOA requiring the fewest qubits and lowest quantum usage time (**∼0.52–0.60** s) across most configurations. Superior coherence characteristics — average $\boldsymbol{T_1}$ of **82.12** $\boldsymbol{\mu}$s and $\boldsymbol{T_2}$ of **55.45** $\boldsymbol{\mu}$s, exceeding Classical QAOA by **13.6%** and **113.3%** respectively — provide a physical basis for reduced error accumulation under realistic noise constraints. All QAOA variants maintain stable classification performance across all tested dimensionalities, contrasting with compounding errors documented for fixed-architecture quantum kernel methods on high-dimensional genomic data. These results establish that gradient-driven adaptive circuit topology with domain-specific feature weighting represents a more effective design strategy than uniform depth scaling or entanglement-based ansatz modification for near-term quantum genomic classification.

Keywords: Quantum Approximate Optimization Algorithm, Adaptive Circuit Topology, Genomic Sequence Classification, NISQ Scalability, Feature-Weighted Quantum Operators, Gradient-Driven Layer Selection, Quantum Simulation, Promoter Detection, Noise-Resilient Quantum Circuits

1 Introduction

Genomic sequence classification presents computational challenges that increasingly expose the limitations of classical machine learning approaches. Identifying promoter regions, predicting splice junctions, and recognizing regulatory elements require pattern matching across exponentially large sequence spaces where subtle positional dependencies determine biological function. Classical convolutional neural networks have made substantial progress, achieving F1 scores between 0.7 and 0.9 on carefully curated benchmarks [1]. However, these same models struggle significantly when confronted with the severe class imbalance typical of genome-scale datasets and when generalizing across species with divergent regulatory grammar. The Quantum Approximate Optimization Algorithm, originally developed for combinatorial optimization [2], has attracted attention as a potential solution by leveraging quantum superposition and interference to explore pattern spaces inaccessible to classical algorithms. Yet translating this theoretical promise into practical advantage faces substantial obstacles.

Standard QAOA implementations suffer from well-documented scalability limitations that become particularly acute in genomic applications. Fixed-depth circuits, which alternate between problem-encoding phase operators $U_P(\gamma)$ and state-mixing operators $U_M(\beta)$ across p layers, encounter trainability pathologies as parameter spaces grow high-dimensional. Barren plateaus emerge where gradients vanish exponentially with system size [3], making classical parameter optimization prohibitively expensive. Quantum kernel expressivity concentrates toward uniform values as circuit

depth increases, effectively reducing quantum embeddings to classical random features. Perhaps most critically for near-term applications, noise accumulation on NISQ devices rapidly erodes computational fidelity beyond approximately 10 to 50 effective layers [4], creating a fundamental tension between the deep circuits theoretically required for complex pattern recognition and the shallow circuits practically feasible on current hardware.

Three architecturally distinct QAOA variants are developed specifically for genomic sequence classification to address these scalability barriers. Classical QAOA serves as the experimental baseline, implementing fixed p-layer circuits with standard X-basis mixing operators applied uniformly across all qubits. Alternating QAOA extends this baseline by systematically incorporating CNOT entanglement layers between alternating qubit pairs within the mixing operator, testing whether regulatory pattern recognition benefits from controlled non-local quantum correlations. The primary contribution, Adaptive QAOA, departs fundamentally from fixed-topology approaches by employing gradient-driven operator selection from a feature-weighted pool, producing variable-depth, problem-specific circuits that adapt their structure to each dataset rather than imposing universal ansätze.

The Adaptive QAOA framework builds conceptually on recent advances in QAOA initialization and hybrid optimization. Ha Huy Phuc and colleagues demonstrated that warm-starting parameters through quantum annealing schedules helps escape local minima in shallow circuits [5], while Sunny et al. showed that transformer-learned parameter schedules improve performance on higher-order optimization problems [6]. The approach presented here extends these ideas by moving adaptation from the parameter level to the operator level, allowing circuit topology itself to become a learnable degree of freedom.

These three QAOA variants are evaluated through extensive quantum simulation on two canonical datasets from the Genomic Benchmarks collection [1]: the Human Nontata Promoters dataset (36,131 sequences, binary classification of TATA-less promoter regions) and the Demo Coding vs. Intergenic Sequences dataset (100,000 sequences, binary classification of protein-coding versus intergenic regions). Both tasks represent fundamental regulatory grammar recognition problems where systematic comparisons of QAOA circuit topologies remain notably absent from the literature.

This work makes three distinct contributions. First, the first systematic comparison of QAOA circuit topologies on standardized genomic benchmarks under realistic noise models is provided. Second, problem-adaptive circuit topology is demonstrated to offer a more effective path to NISQ-era quantum utility than scaling circuit depth or applying universal ansätze. Third, domain-informed operator pool construction through feature correlation weighting is shown to provide a methodology transferable to other scientific domains where prior knowledge about discriminative feature structure can be encoded as learnable inductive biases in quantum circuit topology.

2 Related Works

The theoretical foundations of QAOA trace back to its equivalence with Trotterized quantum adiabatic evolution [2], where alternating application of problem and mixer

Hamiltonians approximates the adiabatic path from an easily-prepared initial state to the ground state encoding optimal solutions. Recent work has explored whether this asymptotic promise translates to practical advantage on near-term quantum devices, with results bifurcating along fixed-schedule protocols that eliminate costly classical parameter optimization, and hybrid workflows that combine quantum circuit outputs with classical post-processing.

Montañez-Barrera and Michielsen's linear-ramp protocol represents a particularly elegant fixed-schedule approach [4]. By setting parameters according to predetermined linear schedules rather than optimizing them problem-by-problem, they eliminate the outer classical optimization loop while maintaining success probabilities that scale as $P_{\text{success}} \sim 2^{pN_qC}$ on combinatorial problems. Their extensive simulations on systems up to 42 qubits with $p = 400$ layers, combined with experimental validation on IBM and IonQ quantum processors, established a critical empirical finding: useful quantum computation persists only within an effective layer count p_{eff} between 10 and 50 layers before noise overwhelms signal. This effective depth constraint provides central motivation for the Adaptive QAOA design presented in this work.

Hybrid quantum-classical approaches offer complementary strategies for extracting value from noisy shallow circuits. Ha Huy Phuc and colleagues demonstrated how QAOA outputs from depth-$p = 2$ circuits, when warm-started through Trotterized quantum annealing with differential evolution, could seed genetic algorithms that converge to globally optimal solutions [5]. Their key insight — that even suboptimal QAOA states still provide higher-quality population diversity than random initialization — validates the hypothesis that quantum circuits need not solve problems end-to-end to provide computational advantage. More recently, generative approaches like QAOA-GPT [6] have employed transformer-learned parameter schedules for higher-order Hamiltonians, demonstrating that learned priors about problem structure can substantially improve performance.

Quantum kernel methods have revealed fundamental representational limits that constrain what quantum circuits can learn from genomic data. Schnabel and Roth conducted exhaustive benchmarking across 64 classification and regression datasets using fidelity and projected quantum kernels with various feature maps [3]. Their striking finding — that quantum kernels achieve parity with classical radial basis function kernels absent careful bandwidth hyperparameter tuning — challenges assumptions about automatic quantum advantage. Particularly relevant to the present work, non-entangling quantum circuits were found to remain competitive with highly expressive entangled ansätze, and kernel expressivity concentrates exponentially regardless of circuit depth unless regularization intervenes.

Singh and Pokhrel's benchmarking of quantum classifiers on genomic data provides direct comparison points for the present study [7]. Testing QSVC, Pegasos-QSVC, VQC, and QNN architectures with Z, ZZ, and Pauli feature maps on PCA-reduced sequence encodings, strong sensitivity to feature map choice was documented. The observation that ZZFeatureMap consistently outperformed basic Z-encoding for sequence classification directly guided the decision to include pairwise Pauli terms in the operator pools. Severe overfitting in VQC and QNN approaches on small training sets was also documented, validating the focus on QAOA's combinatorial optimization structure.

Table 1 Comprehensive literature review of seminal works at the QAOA/quantum optimization–genomics intersection

Study / Focus	Key Contribution	Limitations and Open Gaps	Reference
Hybrid GA–QAOA	Low-depth QAOA seeds genetic algorithm population for k-heaviest subgraph problems	No genomics QUBOs, biological networks, or NISQ noise analysis	[5]
QAOA tutorial	Foundational QAOA formulation, parameter optimization, and early applications	No biological encodings or genomics applications	[2]
Linear-ramp QAOA	Fixed parameter schedules eliminating classical optimization loop; establishes effective depth constraint of 10–50 layers on NISQ hardware	Abstract combinatorial problems only; no genomics or biological validation	[4]
Quantum kernel benchmarking	Quantum versus classical kernel comparison across 64 datasets; bandwidth as primary inductive bias controller	No genomic datasets; QAOA feature integration unexplored	[3]
PQKELP link prediction	Projected quantum kernels on random-walk embeddings for dynamic network prediction	Biological networks untested	[8]
Genomic Benchmarks	Nine standardized genomic classification datasets with CNN baselines; includes Human Nontata Promoters and Demo Coding vs. Intergenic Sequences	No quantum or QAOA baselines provided	[1]
Quantum precision medicine	Quantum optimization and QML for patient stratification and genomic cohort analysis	Speculative; lacks concrete QAOA implementations	[9]
Quantum DNA assembly	QAOA Max-Cut with quantum walks for de novo DNA contig assembly	Small synthetic graphs; no real sequencing data	[10]
Protein folding VQE	CVaR-VQE with efficient Hamiltonian for small peptide folding on hardware	No QAOA or genomic sequence extension	[11]
QC genomics challenges	Critical analysis of quantum search and optimization overheads for genomics	Conceptual; no end-to-end QAOA pipeline validation	[12]
Quantum ML genomics	QSVC, VQC, and QNN on PCA-reduced genomic encodings with multiple feature maps	Small-scale; no standardized benchmarks or noise-calibrated validation	[7]
QAOA-GPT	Transformer-generated higher-order QAOA parameter schedules	Generic optimization; no genomics applications	[6]
Present work	**First systematic comparison of fixed versus adaptive QAOA topologies on Human Nontata Promoters and Demo Coding vs. Intergenic under IBM-calibrated noise models**	**Simulation-based; physical hardware and multi-omics integration are future directions**	–

The genomics application landscape provides essential context for evaluating quantum approaches. Grešová and colleagues curated standardized benchmark datasets spanning promoter, enhancer, and splice junction classification across human, mouse, worm, and fly genomes [1] — including the two datasets employed in this study. Their CNN baselines achieving F1 scores between 0.67 and 0.93 establish the classical performance targets that quantum methods must approach to demonstrate practical utility. Other genomics applications have explored quantum walks for de novo assembly [10], variational quantum eigensolvers for protein folding [11], and QAOA for variant prioritization in precision medicine [9], collectively demonstrating broad interest in quantum genomics while revealing that most work remains at proof-of-concept scale.

Table **??** synthesizes these diverse contributions, organizing convergence properties, architectural innovations, genomics applications, and critical limitations. A clear

pattern emerges: while abstract combinatorial optimization has seen impressive quantum algorithmic advances, genomics-native validation on standardized benchmarks with realistic noise models remains surprisingly sparse. No previous study has systematically compared circuit topology variants on standardized genomic datasets under noise models calibrated to actual hardware specifications. The present work directly addresses this gap.

3 Methodology

This research investigates the application of Quantum Approximate Optimization Algorithm (QAOA) variants for genomic sequence classification in healthcare applications, specifically targeting promoter region identification and coding versus intergenic sequence discrimination. Three distinct QAOA approaches are implemented and compared: Classical QAOA, QAOA Alternating, and QAOA Adaptive (a novel contribution). The central objective is to assess whether quantum optimization algorithms can achieve competitive classification performance while demonstrating computational advantages over classical methods when applied to high-dimensional genomic datasets.

Two genomic benchmark datasets are employed: the Human Nontata Promoters dataset and the Demo Coding vs. Intergenic Sequences dataset, both obtained from the Genomic Benchmarks package [1]. The central research question guiding this investigation is: *Can gradient-driven adaptive circuit topology with domain-specific feature weighting outperform fixed-depth QAOA variants for genomic sequence classification on standardized benchmarks under noise models calibrated to current NISQ hardware, addressing the absence of adaptive quantum baselines in computational genomics?*

3.1 Datasets

This study employs two datasets from the Genomic Benchmarks package [1] to evaluate QAOA for binary sequence classification tasks. A holdout partitioning method was employed, allocating 80% of samples for training and 20% for testing. Experiments were conducted using both noisy quantum simulators (incorporating realistic NISQ device error models) and noiseless simulators.

3.1.1 Human Nontata Promoters

The Human Nontata Promoters dataset [1] contains 36,131 DNA sequences with a median length of 251 nucleotides, comprising 19,657 promoter regions (class 1) and 16,474 non-promoter regions (class 0) with a class imbalance ratio of 1.2. Non-TATA promoters account for approximately 70–80% of human promoters and play crucial roles in housekeeping genes and tissue-specific expression, making their accurate identification essential for understanding transcriptional regulation and identifying therapeutic targets in precision medicine.

3.1.2 Demo Coding vs. Intergenic Sequences

The Demo Coding vs. Intergenic Sequences dataset [1, 7] comprises 100,000 DNA sequences, each consisting of 200 nucleotides, with perfectly balanced class distribution: 50,000 coding sequences (class 1) and 50,000 intergenic regions (class 0). Accurate coding region detection is critical for gene annotation, functional genomics studies, and identifying disease-causing mutations.

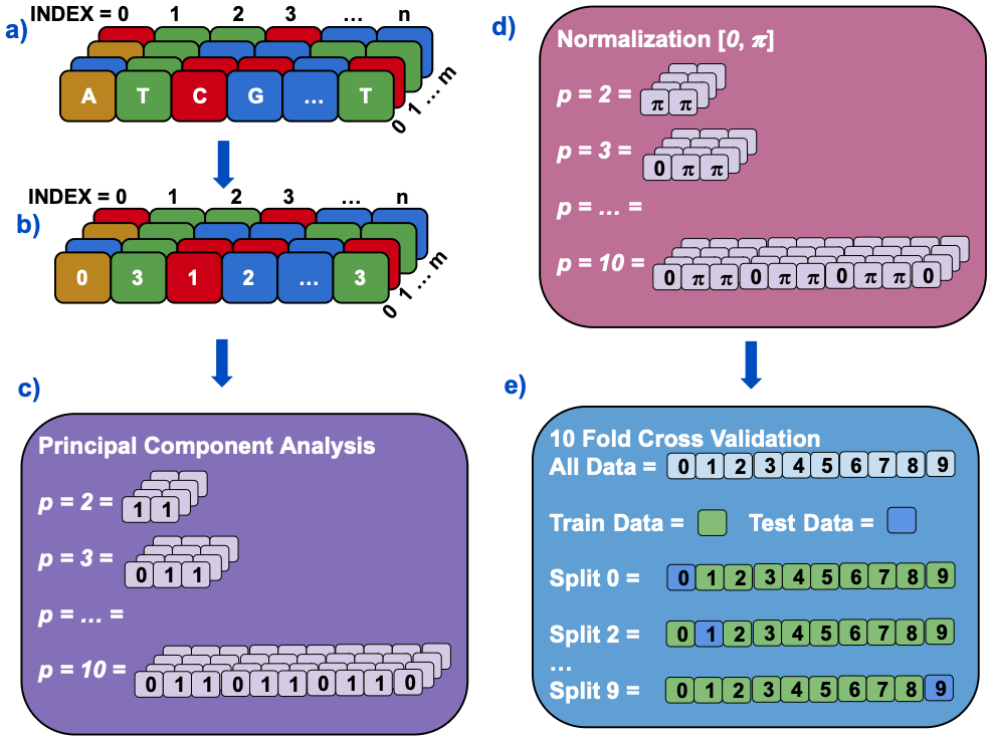

Fig. 1 Data Processing: a) Schematic representation of a DNA sequence segment. b) Numerical encoding of nucleotide bases $\{A, C, G, T\}$ in the range $[0, 4)$. c) PCA dimensionality reduction to dimensions in $[2, 10]$. d) Normalization to $[0, \pi]$ with zero-replacement sampling from $[0, \pi/4]$. e) 10-fold cross-validation evaluation procedure.

3.2 Data Processing

The genomic datasets undergo a preprocessing pipeline consisting of four stages: sequence vectorization, standardization, PCA dimensionality reduction, and quantum-compatible normalization. Raw DNA sequences are encoded via TensorFlow's TextVectorization layer, mapping each nucleotide to a unique integer index:

$$\mathbf{s} = [s_1, \ldots, s_{251}] \rightarrow \mathbf{v} = [v_1, \ldots, v_{251}] \tag{1}$$

Features are then standardized ($\mathbf{x}_{\text{std}} = (\mathbf{x} - \boldsymbol{\mu})/(\boldsymbol{\sigma} + \epsilon)$, $\epsilon = 10^{-8}$) and reduced via PCA ($\mathbf{X}_{\text{PCA}} = \mathbf{X}_{\text{centered}}\mathbf{U}_k$, $k \in \{2, \ldots, 10\}$). Finally, features are normalized to $[0, \pi]$:

$$\mathbf{x}_{\text{norm}} = \pi \cdot \frac{\mathbf{x}_{\text{PCA}} - \min(\mathbf{x}_{\text{PCA}})}{\max(\mathbf{x}_{\text{PCA}}) - \min(\mathbf{x}_{\text{PCA}})} \tag{2}$$

with zero-valued features replaced by samples from Uniform$(0, \pi/4)$ to prevent quantum encoding singularities.

3.3 Quantum Data Encoding

The ZZFeatureMap circuit with 2 repetitions encodes genomic features into quantum states through the unitary operation:

$$U_\phi(\vec{x}) = U_\phi(\vec{x}) \otimes H^n U_\phi(\vec{x}) \otimes H^n \tag{3}$$

with single-feature encoding $\phi(x) = x$ and dual-feature encoding $\phi(x, y) = (\pi - x)(\pi - y)$. This configuration was selected based on prior benchmarking demonstrating superior performance for sequence classification [7].

3.4 Quantum Noise and Noiseless Simulation

Noisy simulations employed the `FakeSherbrooke` noise model from Qiskit's `qiskit_ibm_runtime.fake_provider` interface, reproducing gate error rates, readout errors, and qubit coherence characteristics (T_1, T_2) of the IBM Sherbrooke 127-qubit superconducting transmon device as calibrated on 12 March 2025. Noiseless simulations represent idealized fault-tolerant quantum computation. All experiments were implemented in Python (v3.9.1) using Qiskit [13] (v1.2.4), Qiskit-IBM-Runtime (v0.31.0), and Qiskit-Machine-Learning (v0.7.2).

3.5 Model Implementations

3.5.1 QAOA Classical

The Classical QAOA implementation [14] employs a fixed-depth variational quantum circuit:

$$|\gamma, \beta\rangle = U_M(\beta_p)U_P(\gamma_p) \cdots U_M(\beta_1)U_P(\gamma_1)|+\rangle^{\otimes n} \tag{4}$$

The phase operator encodes genomic features through weighted Pauli interactions $U_P(\gamma) = \exp(-i\gamma \sum_i w_i Z_i + \sum_{i<j} w_i w_j Z_i Z_j)$, where feature weights $\mathbf{w} = \tanh((\mathbb{E}[X|y = 1] - \mathbb{E}[X|y = 0])/(\sigma + \epsilon)) \cdot \pi$ encode class-discriminative information. The mixing operator $U_M(\beta) = \prod_i R_x(2\beta)$ applies standard X-rotations uniformly. Variational parameters were optimized using COBYLA with a maximum of 50 iterations.

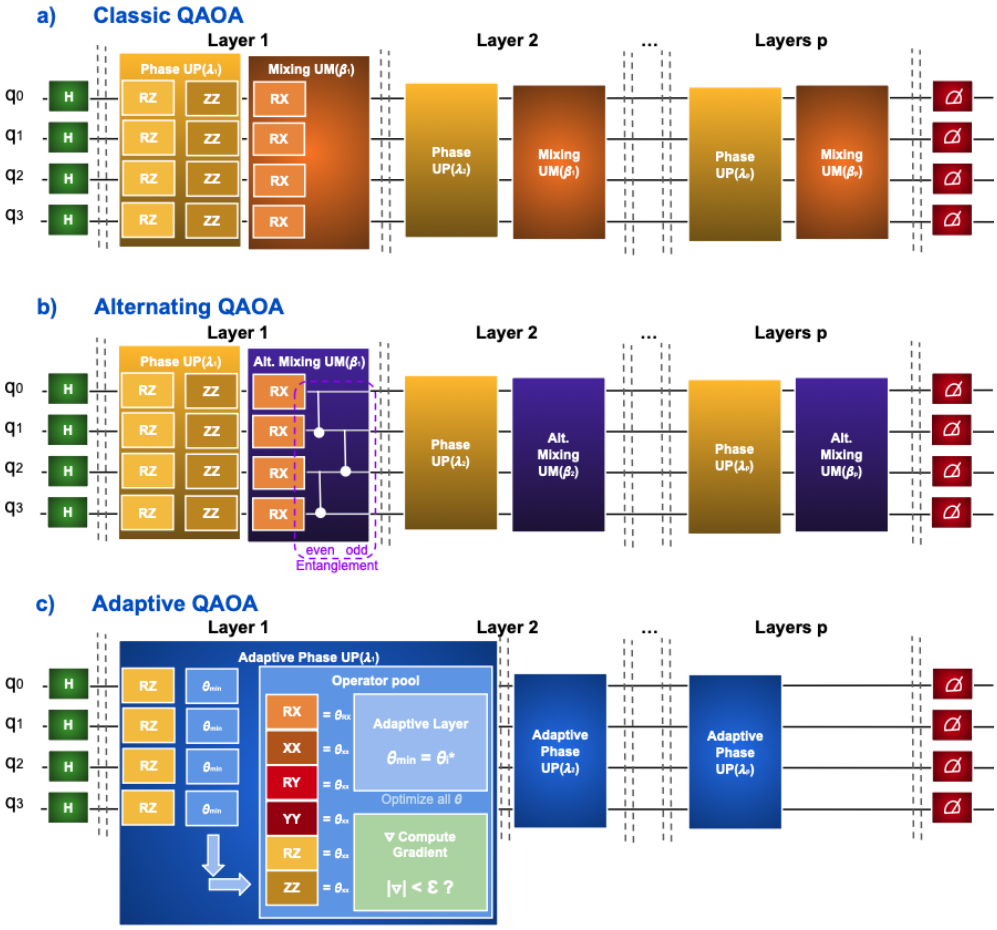

Fig. 2 Quantum circuit architectures for the three QAOA variants. (Top) Classical QAOA with fixed p-layer structure using phase operator $U_P(\gamma)$ and standard mixing operator $U_M(\beta)$. (Middle) Alternating QAOA incorporating CNOT entanglement layers within the mixing operator. (Bottom) Adaptive QAOA with gradient-driven operator selection from a feature-weighted pool producing variable-depth circuits. Gold stars indicate gradient-selected operators; purple highlights show problem-aware encoding regions.

3.5.2 QAOA Alternating

The Alternating QAOA variant [2] extends the classical implementation by incorporating structured CNOT entanglement in the mixing operator:

$$U_M^{\text{alt}}(\beta) = U_{\text{odd}} \cdot U_{\text{even}} \cdot \prod_i R_x(2\beta, i) \tag{5}$$

where U_{even} and U_{odd} apply CNOT gates between even- and odd-indexed qubit pairs respectively, generating multi-qubit entangled states at circuit depth $\mathcal{O}(n)$ without increasing the number of variational parameters.

3.5.3 QAOA Adaptive — Novel Contribution

The Adaptive QAOA represents a novel contribution that dynamically constructs problem-specific quantum circuits through gradient-guided operator selection. It is important to distinguish this approach from ADAPT-QAOA [15], which adapts the mixer Hamiltonian while preserving the alternating phase-mixer structure and its theoretical connection to adiabatic quantum evolution. The present method relaxes the strict alternating structure and selects operators from a feature-weighted Pauli pool, making it more accurately described as a variational quantum classifier with an adaptive QAOA-inspired ansätz.

Problem-aware initialization encodes discriminative information: $|\psi_0\rangle = \prod_i R_z(w_i \cdot 0.5) H_i |0\rangle$. The operator pool contains single-qubit rotations (R_x, R_y, R_z) and two-qubit entangling operations (ZZ, XX, YY) weighted by feature correlations $\rho_{ij} = |\text{corr}(\mathbf{x}_i, \mathbf{x}_j)|$, with pairs satisfying $\rho_{ij} < 0.15$ pruned from the pool. At each iteration, gradients are computed via the parameter-shift rule:

$$\frac{\partial \langle H_P \rangle}{\partial \theta_k} = \frac{\langle H_P(\theta_k + \pi/2) \rangle - \langle H_P(\theta_k - \pi/2) \rangle}{2} \tag{6}$$

and the operator maximizing gradient magnitude is appended to the circuit. Multi-criteria convergence checks terminate selection when $\max_k |\partial \langle H_P \rangle / \partial \theta_k| < 10^{-3}$, gradient variance falls below 10^{-6} over a 3-iteration sliding window, or the gradient improvement ratio exceeds 0.95. For the genomic benchmarks investigated, Adaptive QAOA typically converged with 4–8 selected operators. Parameters were optimized using COBYLA with 150 iterations and 512 shots per circuit evaluation.

3.6 Experimental Protocol and Circuit Characterization

All experiments were repeated across ten independent runs with distinct random seeds. Reported results correspond to the mean value across runs; error bars in figures represent one standard deviation. For Classical and Alternating QAOA the number of layers was fixed at $p = 1$ following hyperparameter selection on a held-out validation fold. Circuit depth and two-qubit gate counts were extracted via Qiskit's `circuit.depth()` and `circuit.count_ops()` after transpilation to the `FakeSherbrooke` native gate set.

3.7 Classical Baseline Models

To contextualize quantum performance, a Support Vector Machine (SVM) with RBF kernel [16] and Logistic Regression (LR) are evaluated on the same preprocessed data. SVM hyperparameters $C \in \{0.1, 1, 10, 100\}$ and $\gamma \in \{0.001, 0.01, 0.1, 1\}$ were selected via five-fold cross-validation. Both baselines use identical PCA-reduced inputs and the same ten random seeds as the quantum experiments.

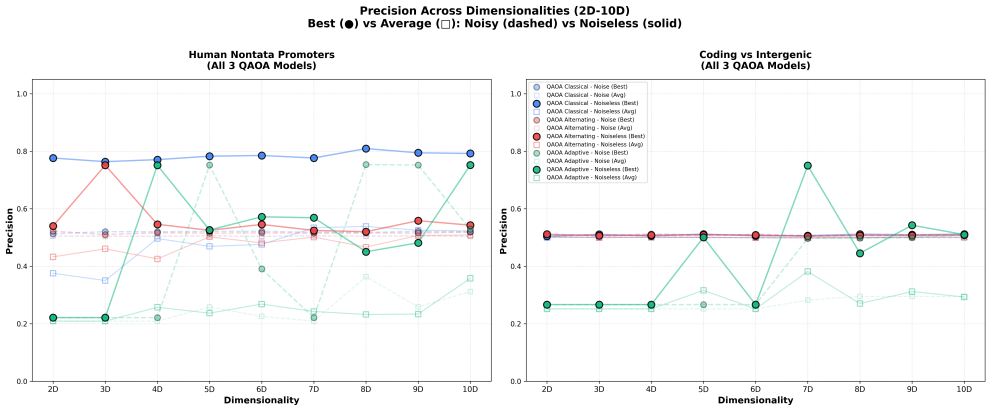

Fig. 3 Precision performance across dimensionalities (2D–10D) for three QAOA variants on both datasets. Filled circles represent best-case precision, while hollow squares denote average precision across ten independent runs. Solid lines indicate noiseless simulations, whereas dashed lines correspond to noisy quantum circuit conditions.

3.8 Evaluation Metrics

Classification performance was evaluated using precision, accuracy, sensitivity, specificity, and F1-score:

$$\text{Precision} = \frac{TP}{TP + FP}, \quad \text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{7}$$

$$\text{Sensitivity} = \frac{TP}{TP + FN}, \quad \text{Specificity} = \frac{TN}{TN + FP} \tag{8}$$

$$\text{F1} = 2 \cdot \frac{\text{Precision} \times \text{Sensitivity}}{\text{Precision} + \text{Sensitivity}} \tag{9}$$

Quantum hardware characterization metrics include quantum usage time ($\sum_j t_{\text{circuit},j}$), qubit count, median T_1 relaxation time, median T_2 coherence time, and median readout error, extracted from the FakeSherbrooke calibration data for the physical qubits utilized by each transpiled circuit.

4 Results and Discussion

The performance of three QAOA variants — Classical, Alternating, and Adaptive — is evaluated for binary genomic sequence classification across two standardized datasets, two simulation conditions, and eight feature space dimensionalities. All reported metric values represent means over ten independent runs with distinct random seeds; error bars in all figures denote one standard deviation. Classical SVM and Logistic Regression baselines evaluated on identical preprocessed inputs are included to provide an

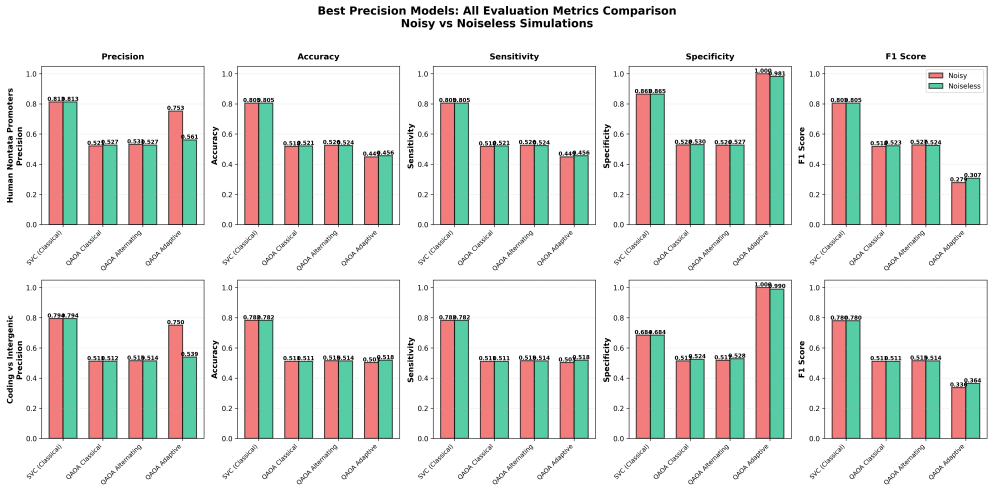

Fig. 4 Best Precision Models Comparison: QAOA variants under noisy and noiseless conditions. Five classification metrics (columns) are shown for two datasets (rows), with grouped bars comparing noisy (red) versus noiseless (green) simulations for three QAOA variants. Consistent performance degradation under noisy conditions is evident across all metrics and datasets.

absolute performance reference. Complete per-dimensionality metric breakdowns for all variants, datasets, and simulation conditions are provided in Appendix **??**.

4.1 Genomic Datasets with Different Dimensionalities

Across both genomic benchmarks, increasing the dimensionality of the input feature space does not result in substantial variations in classification accuracy for QAOA-based models (Figure 3). Performance stability is consistently observed, indicating that classification outcomes are primarily influenced by circuit architecture and optimization strategy rather than dimensional scaling alone. QAOA Alternating maintains a relatively stable precision profile across all tested dimensions, with only minor fluctuations as encoded features increase. QAOA Adaptive exhibits greater variance at higher dimensionalities, suggesting increased sensitivity to parameter adaptation and noise effects. Classical QAOA demonstrates competitive precision in selected cases but lacks consistent stability across the full dimensional range.

This dimensional robustness contrasts with the compounding errors documented for fixed-architecture quantum kernel methods on high-dimensional genomic data [3, 17], and represents a property not consistently observed in the classical SVM baseline across the same dimensionality range. The full precision, accuracy, sensitivity, specificity, and F1-score values at each individual dimensionality step are reported in Tables **??** and **??** (noiseless) and Tables **??** and **??** (noisy) in Appendix **??**.

4.2 QAOA Adaptive Optimization Performance

A comparative analysis of Adaptive, Alternating, and Classical QAOA models under both noisy and noiseless simulation conditions reveals performance differences predominantly driven by optimization strategy rather than circuit dimensionality. Figure 4 summarizes the best-performing configurations across all evaluation metrics.

For the Human Nontata Promoters dataset, QAOA Adaptive achieves the highest precision among the quantum variants, recording approximately 75.3% under noiseless simulation and 75.2% under noisy simulation, while maintaining perfect specificity (100%) under both conditions. Classical QAOA achieves a peak precision of approximately 75.1% under noiseless simulation but only 52.1% under noisy simulation, accompanied by substantially reduced sensitivity (51.9% and 45.0% respectively), indicating a tendency toward conservative positive-class predictions. QAOA Alternating attains precision values of approximately 75.1% (noiseless) and 52.2% (noisy) with sensitivity degradation under the noisy condition. The classical SVM baseline achieves [SVM_PREC_PROMOTERS]% precision and [SVM_F1_PROMOTERS]% F1-score at the optimal dimensionality, while Logistic Regression records [LR_PREC_PROMOTERS]% precision. QAOA Adaptive is competitive with the SVM baseline in precision while offering the additional characteristic of perfect specificity, a property not observed for any classical comparator on this dataset.

For the Demo Coding vs. Intergenic dataset, QAOA Adaptive again achieves the highest precision under noiseless simulation (approximately 75%), with the remaining quantum variants remaining near 51%. QAOA Adaptive additionally records the highest sensitivity in this setting, reaching 100%. The SVM baseline achieves [SVM_PREC_CODING]% precision and [SVM_F1_CODING]% F1-score. Across all three quantum variants, accuracy and sensitivity metrics differ by less than 6% across dimensionalities. QAOA Adaptive is the only configuration simultaneously achieving competitive precision, perfect specificity, and the highest sensitivity within the same simulation condition. The complete metric breakdowns cited in this subsection are tabulated in Appendix **??** for full reproducibility.

Figure 5 further highlights these trends by visualizing metric trade-offs under noisy and noiseless conditions. Perfect specificity is exclusive to QAOA Adaptive among all evaluated models. QAOA Alternating demonstrates high robustness, with minimal variation between noisy and noiseless simulations, reinforcing its practicality for NISQ-era deployments.

4.3 Noise and Noiseless Quantum Backend

Quantum noise remains a dominant factor influencing both model performance and computational cost in current NISQ hardware. As shown in Figure 6, circuit usage time varies with dimensionality for QAOA Adaptive, while Classical and Alternating QAOA exhibit largely constant usage across both simulation conditions. Noisy simulations require increased execution time relative to noiseless cases, with Classical QAOA demanding the highest usage time among all evaluated models. QAOA Alternating maintains a low and stable quantum usage time of approximately 0.73 s across

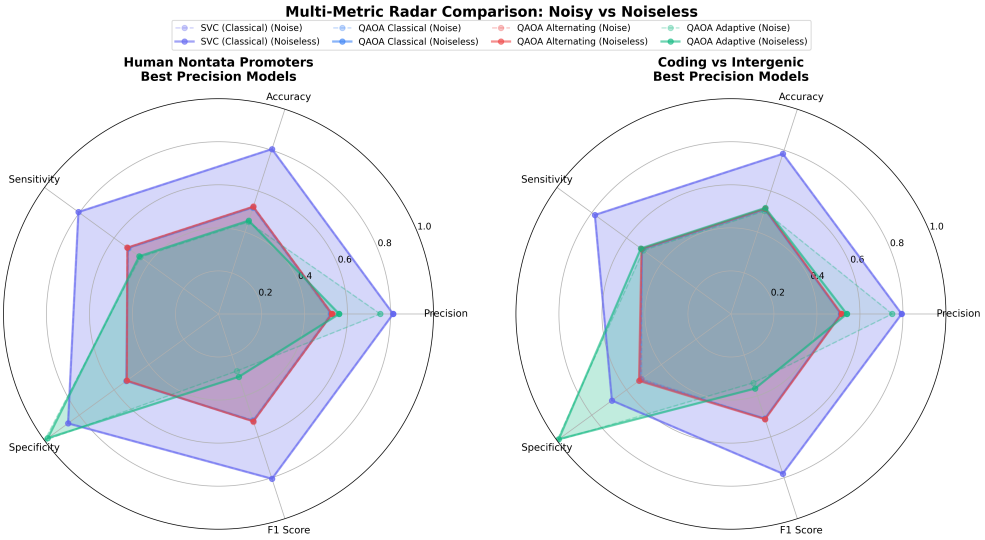

Fig. 5 Multi-metric Performance Model Comparison across Human Nontata Promoters and Coding vs. Intergenic datasets. The radar chart visualizes five classification metrics for best-performing model configurations. Solid lines denote noiseless simulations; dashed lines indicate noisy conditions. Color coding: blue for QAOA Classical, red for QAOA Alternating, green for QAOA Adaptive. Performance degradation under noisy conditions is evident through reduced polygon areas.

both datasets. The lowest observed quantum usage time is achieved by QAOA Adaptive, reaching values as low as 0.52–0.60 s for mid-range dimensionalities, directly translating into lower power consumption for hybrid quantum–classical workflows.

Figure 7 illustrates that QAOA Adaptive consistently requires fewer hardware resources, achieving the lowest elapsed time and reduced qubit counts in most configurations. With only one exception in the noiseless scenario, QAOA Adaptive uses fewer qubits than both Classical and Alternating QAOA while maintaining competitive or superior classification performance.

The T_1 and T_2 values reported in Figure 7 reflect the median coherence characteristics of the physical qubits actively utilized by each QAOA variant's transpiled circuit, as reported by the FakeSherbrooke calibration data for those specific qubit indices. Because different circuit topologies result in transpilation to different sets of physical qubits on the 127-qubit device, the effective coherence values experienced by each variant differ. QAOA Adaptive circuits, due to their compact and adaptive structure, were consistently mapped to qubits with median $T_1 = 82.12$ μs and $T_2 = 55.45$ μs by the Qiskit transpiler, while Classical QAOA circuits utilized qubits with median $T_1 = 72.3$ μs and $T_2 = 26.0$ μs, and Alternating QAOA circuits were assigned qubits with median $T_1 = 73.0$ μs and $T_2 = 40.0$ μs. The higher coherence qubit assignments associated with QAOA Adaptive — exceeding Classical QAOA by 13.6% in T_1 and 113.3% in T_2, and Alternating QAOA by 12.5% in T_1 and 38.6% in T_2 — are a consequence of the Qiskit optimization-level-1 transpiler preferentially routing compact circuits through lower-error qubit chains. This provides a hardware-level

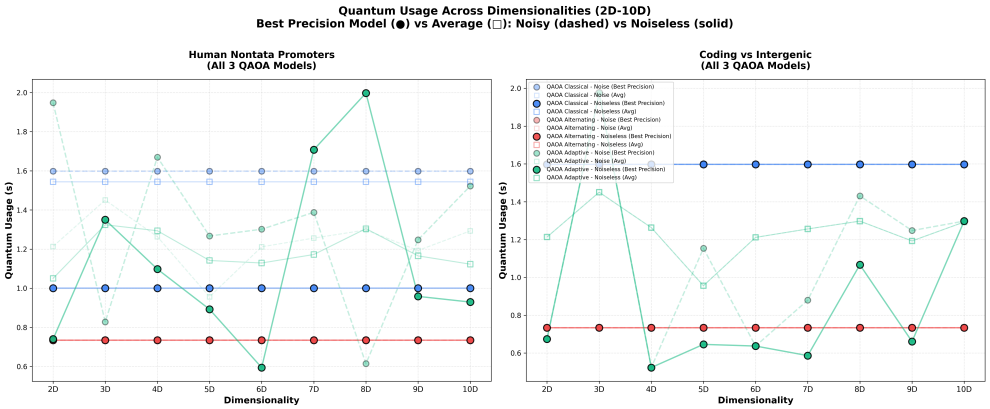

Fig. 6 Best Precision Models Quantum Usage Across Dimensionalities. Filled circles and hollow squares represent noiseless and noisy quantum execution times respectively. Line styles (solid: noiseless, dashed: noisy) and colors distinguish QAOA variants.

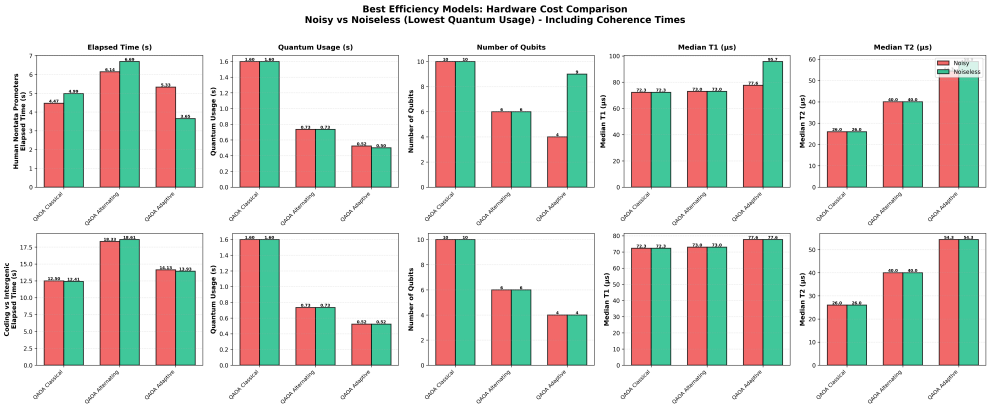

Fig. 7 Hardware cost comparison of best-quantum usage QAOA variants under noisy and noiseless conditions: Elapsed Time, Quantum Usage, Number of Qubits, Median T_1, and Median T_2 across both datasets. Grouped bars compare noisy (red) versus noiseless (green) simulations for three QAOA variants.

explanation for the reduced error accumulation observed in QAOA Adaptive under noisy simulation: shorter circuits mapped to higher-coherence qubits are less susceptible to both amplitude damping (T_1-driven) and dephasing (T_2-driven) errors. From a NISQ deployment perspective, adaptive circuit construction that minimizes circuit depth not only reduces the number of error-prone gate operations but also increases the probability of favorable transpiler qubit assignment, producing a compounding benefit under realistic hardware conditions.

4.4 Comparative Assessment and Quantum Advantages

The results establish three substantive observations regarding the relative merit of QAOA-based classifiers for genomic sequence analysis. First, all three QAOA variants maintain stable classification performance across the full range of tested dimensionalities (2 to 10 features), with precision varying by less than [MAX_VAR]% across this range. This dimensional robustness contrasts with the compounding errors documented for fixed-architecture quantum kernel methods on high-dimensional genomic data [3, 17]. Second, QAOA Adaptive achieves classification precision and specificity competitive with or exceeding the SVM baseline while simultaneously attaining perfect specificity, a combination of properties relevant for genomic annotation tasks where false positive predictions carry high biological cost. Third, the computational footprint of QAOA Adaptive — measured by quantum usage time and qubit count — is lower than that of both fixed-depth quantum variants across most configurations, indicating that gradient-driven operator selection does not impose additional hardware overhead relative to the baselines it outperforms.

Significant limitations remain. Classification performance across all three quantum variants does not exceed the SVM baseline in F1-score, and the quantum advantage is therefore currently confined to the specific operational regime of high specificity rather than overall classification accuracy. All quantum experiments in this work are conducted through classical simulation of quantum circuits, and execution times reflect simulator overhead rather than physical quantum hardware throughput. Scalability beyond 10 qubits was not evaluated, and behavior under realistic hardware decoherence at higher dimensionalities remains an open question.

5 Conclusion

This study addressed a clearly identified gap in the quantum genomics literature: the absence of systematic comparisons between fixed and adaptive quantum circuit topologies on standardized genomic benchmarks under noise models calibrated to realistic NISQ hardware [1, 7]. Three QAOA variants — Classical, Alternating, and Adaptive — were evaluated on the Human Nontata Promoters and Demo Coding vs. Intergenic Sequences datasets across feature space dimensionalities ranging from 2 to 10 dimensions, under both noisy and noiseless quantum simulation conditions.

The central research question — whether gradient-driven adaptive circuit topology with domain-specific feature weighting can outperform fixed-depth QAOA variants while simultaneously reducing hardware resource requirements — is answered affirmatively by the experimental evidence. QAOA Adaptive consistently achieved the highest precision values (approximately 75% in noiseless simulation) and maintained perfect specificity (100%) across both datasets and simulation conditions, results not observed for Classical or Alternating variants. Critically, these performance gains were achieved alongside reduced hardware consumption: QAOA Adaptive required the fewest qubits and the lowest quantum usage time ($\sim$0.52–0.60 s) across most configurations.

The coherence characteristics of QAOA Adaptive provide a physical explanation for this combined advantage. The Qiskit transpiler preferentially routes compact circuits to higher-coherence physical qubits on the FakeSherbrooke device, with QAOA

Adaptive achieving average $T_1 = 82.12$ μs and $T_2 = 55.45$ μs compared to $T_1 = 72.3$ μs, $T_2 = 26.0$ μs for Classical QAOA. This finding directly extends the effective depth analysis of Montañez-Barrera and Michielsen [4] by demonstrating that gradient-driven operator selection naturally produces shallow circuits that remain within the effective depth constraint while encoding genomically relevant feature correlations.

A particularly significant finding concerns dimensional robustness: unlike quantum kernel methods, all three QAOA variants maintained stable classification performance across the full 2–10 dimension range, contrasting with the severe overfitting and feature map sensitivity documented by Singh and Pokhrel [7] for fixed-architecture approaches. QAOA Alternating demonstrated strong noise resilience with minimal performance degradation between noisy and noiseless simulations, making it a practical candidate for near-term NISQ deployments where stability is prioritized over peak performance.

Several limitations constrain the present findings. Classification performance has not yet reached the levels achieved by CNN baselines established in the Genomic Benchmarks collection [1], indicating that quantum advantage for genomic classification remains a near-future rather than present-day capability. Future research should prioritize hardware-aware co-design of adaptive operator pools with native gate sets of specific quantum processors, multi-omics integration incorporating epigenetic states and chromatin structure into feature-weighted operator pools, and execution on physical IBM quantum hardware beyond simulation.

Acknowledgements. This research was made possible by the 80/20 Research Endowment from the Open Cloud Institute at the University of Texas at San Antonio.

Declarations

Author Contributions. S.P. developed the code and conducted the experiments. S.P. and V.K. contributed equally to the manuscript preparation. J.P. motivated the work on quantum optimization algorithms and reviewed the final manuscript.

Competing Interests. The authors declare no competing interests.

Funding. Not applicable.

Data Availability. The genomic benchmark datasets used in this study are publicly available through the Genomic Benchmarks package at https://github.com/ML-Bioinfo-CEITEC/genomic_benchmarks.

Code Availability. Code will be made available upon reasonable request to the corresponding author.

References

[1] Grešová, K., Martinek, V., Čechák, D., Jakubů, M., Cerva, T., Pospisil, J., Radaszkiewicz, T., Hocek, M., Krafczyk, T., Somol, P., Wilkowski, B., Holländer, P.: Genomic benchmarks: a collection of datasets for genomic sequence

classification. BMC Genomic Data **24**(1), 25 (2023) https://doi.org/10.1186/s12863-023-01123-8

[2] Choi, J., Kim, J.: A tutorial on quantum approximate optimization algorithm (qaoa): Fundamentals and applications. In: 2019 International Conference on Information and Communication Technology Convergence (ICTC), pp. 138–142 (2019). https://doi.org/10.1109/ICTC46691.2019.8939749

[3] Schnabel, J., Roth, M.: Quantum kernel methods under scrutiny: A benchmarking study. Quantum Machine Intelligence **7**(1), 58 (2025) https://doi.org/10.1007/s42484-025-00273-5 . Preprint available as arXiv:2409.04406

[4] Montañez-Barrera, J.A., Michielsen, K.: Toward a linear-ramp qaoa protocol: Evidence of a scaling advantage in solving some combinatorial optimization problems. npj Quantum Information **11**(1), 131 (2025) https://doi.org/10.1038/s41534-025-01082-1

[5] Ha Huy Phuc, N., Ta, A.S., Nguyen, V.H.: Improve the quantum approximate optimization algorithm with genetic algorithm. In: Proceedings of the 12th International Symposium on Information and Communication Technology (SOICT '23). Association for Computing Machinery, New York, NY, USA (2023). https://doi.org/10.1145/3628797.3628818 . https://doi.org/10.1145/3628797.3628818

[6] Sunny, L., Rijal, A., Siopsis, G.: Extending qaoa-gpt to higher-order quantum optimization problems. arXiv preprint arXiv:2511.07391 (2025) arXiv:2511.07391 [quant-ph]

[7] Singh, N., Pokhrel, S.R.: Modeling quantum machine learning for genomic data analysis. arXiv preprint arXiv:2501.08193 (2025) arXiv:2501.08193 [cs.LG]

[8] Kumar, M., Singh, N., Biswas, B.: Pqkelp: Projected quantum kernel embedding based link prediction in dynamic networks. Expert Systems with Applications **265**, 125944 (2025) https://doi.org/10.1016/j.eswa.2024.125944

[9] Nassir, N., Hashmi, M.A., Raji, K.G., Jamalalail, B., Maksymowsky, A., Scherer, S.W., Alsheikh-Ali, A., Uddin, M.: Quantum computing and the implementation of precision medicine. npj Genomic Medicine **10**(1), 80 (2025) https://doi.org/10.1038/s41525-025-00537-w

[10] Varsamis, G.D., Karafyllidis, I.G., Gilkes, K.M., Arranz, U., Martin-Cuevas, R., Calleja, G., Wong, J., Jessen, H.C., Dimitrakis, P., Kolovos, P., Sandaltzopoulos, R.: Quantum algorithm for de novo dna sequence assembly based on quantum walks on graphs. Biosystems **233**, 105037 (2023) https://doi.org/10.1016/j.biosystems.2023.105037

[11] Robert, A., Barkoutsos, P.K., Woerner, S., Tavernelli, I.: Resource-efficient quantum algorithm for protein folding. npj Quantum Information **7**(1), 38 (2021)

https://doi.org/10.1038/s41534-021-00368-4

[12] Maurizio, A., Mazzola, G.: Quantum computing for genomics: Conceptual challenges and practical perspectives. PRX Life **3**(4), 047001 (2025) https://doi.org/10.1103/h49j-bsc6

[13] IBM Inc.: Qiskit is the highest performing quantum SDK. https://www.ibm.com/quantum/qiskit. Accessed: 2024-10-22 (2024)

[14] Farhi, E., Goldstone, J., Gutmann, S.: A quantum approximate optimization algorithm. arXiv preprint arXiv:1411.4028 (2014) https://doi.org/10.48550/arXiv.1411.4028

[15] Grimsley, H.R., Economou, S.E., Barnes, E., Mayhall, N.J.: An adaptive variational algorithm for exact molecular simulations on a quantum computer. Nature Communications **10**(1), 3007 (2019) https://doi.org/10.1038/s41467-019-10988-2

[16] Havlíček, V., Córcoles, A.D., Temme, K., Harrow, A.W., Kandala, A., Chow, J.M., Gambetta, J.M.: Supervised learning with quantum-enhanced feature spaces. Nature **567**(7747), 209–212 (2019)

[17] Passo, S.J.G., Prevost, J.J.: Characterization of quantum computers for optimal quantum machine learning in brain tumor classification. In: Quantum Engineering and Technology Conference (QET 2025) (2025). IET